



EMC SnapView
For EMC Navisphere

ADMINISTRATOR'S GUIDE

P/N 069001180
REV A06

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748 -9103
1-508 -435 -1000
www.EMC.com

Copyright © 2002 - 2005 EMC Corporation. All rights reserved.

Published July, 2005

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

Trademark Information

EMC², EMC, CLARiON, Navisphere, and PowerPath are registered trademarks and Access Logix, FLARE, MirrorView, Powerlink, SAN Copy, and SnapView are trademarks of EMC Corporation.

All other trademarks used herein are the property of their respective owners.

Preface	vii
----------------------	-----

Chapter 1 About EMC SnapView Software

Introduction to SnapView	1-2
Clones Overview	1-3
Snapshots Overview.....	1-4
Clone and Snapshot Tradeoffs.....	1-6
SnapView Components.....	1-7
Navisphere Manager	1-8
Navisphere CLI	1-8
admsnap	1-9
SnapView Servers	1-10
Configuration Guidelines	1-12
SnapView Limits.....	1-12
Event Monitor	1-13
MirrorView	1-14
SAN Copy	1-14
SnapView Behavior with AX-Series Storage Systems	1-15
AX-Series to CX-Series Terminology Differences for SnapView	1-15
Navisphere Manager Right-Click Menu Options for SnapView	1-16
Using Online Help	1-20

Chapter 2 Setting Up SnapView

Setting Up SnapView to Use Clones	2-2
Prerequisites for Setting Up Clones	2-2
Overview of Setting Up SnapView to Use Clones.....	2-3
Allocating Clone Private LUNs	2-4
Deallocating/Reallocating Clone Private LUNs.....	2-6

Creating a Clone Group.....	2-7
Adding a Clone to a Clone Group	2-9
Setting Up SnapView to Use Snapshots	2-12
Prerequisites for Setting Up Snapshots	2-12
Overview of Setting Up SnapView to Use Snapshots.....	2-14
Reserved LUN Pool with SnapView	2-15
Starting a SnapView Session.....	2-17
Creating a Snapshot	2-24
Adding a Snapshot to a Storage Group	2-27
Chapter 3	Using SnapView
Using Clones.....	3-2
Clone States	3-2
Fracturing a Clone.....	3-5
Synchronizing a Fractured Clone.....	3-11
Reverse Synchronizing a Fractured Clone.....	3-13
Removing a Clone from a Clone Group.....	3-18
Destroying a Clone Group	3-19
Clone and Source LUN Trespasses.....	3-20
Bad Blocks on Clones	3-21
Using Snapshots.....	3-23
Activating a Snapshot	3-23
Deactivating a Snapshot	3-25
Rolling Back a SnapView Session	3-27
Stopping a SnapView Session.....	3-33
Destroying a Snapshot.....	3-34
Snapshot/SnapView Session and Source LUN Trespasses.....	3-35
Bad Blocks on Rollback.....	3-36
Chapter 4	Displaying SnapView Properties
Displaying and Modifying Clone Properties.....	4-2
Clone Properties	4-2
Clone Feature Properties	4-3
Clone Group Properties.....	4-3
Source LUN Properties	4-4
Displaying and Modifying Snapshot Properties.....	4-5
Snapshot Name Properties.....	4-5
SnapView Session Properties.....	4-6
Displaying Status of All Snapshots and SnapView Sessions	4-7

Chapter 5	SnapView Examples	
	Clones Example	5-2
	Summary	5-2
	Hardware and Software Configuration.....	5-2
	Operations Overview	5-3
	Illustrated Overview.....	5-5
	Snapshots Example	5-10
	Summary	5-10
	Hardware and Software Configuration.....	5-10
	Operations Overview	5-11
	Illustrated Overview.....	5-13
	Snapshots Example with Rollback.....	5-16
	Summary	5-16
	Hardware and Software Configuration.....	5-16
	Operations Overview	5-17
	Illustrated Overview.....	5-19
Appendix A	Using SnapView on a Tru64 Server	
	Determining a Tru64 Source LUN	A-2
	Setting Up Clones	A-8
	Creating a Clone Group and Adding a Clone	A-8
	Using Clones	A-9
	Verifying the Clone	A-9
	Accessing the Clone.....	A-9
	Setting Up Snapshots	A-13
	Creating a Snapshot and Starting a SnapView Session...	A-13
	Using Snapshots	A-14
	Activating a Snapshot	A-14
	Accessing the Snapshot.....	A-14
Appendix B	Customer Support	
	Overview of Detecting and Resolving Problems	B-2
	Troubleshooting the Problem	B-3
	Before Calling the Customer Support Center	B-4
	Documenting the Problem	B-5
	Reporting a New Problem	B-6
	Sending Problem Documentation	B-7
Glossary		g-1
Index		i-1

As part of an effort to improve and enhance the performance and capabilities of its product line, EMC from time to time releases revisions of its hardware and software. Therefore, some functions described in this guide may not be supported by all revisions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.

If a product does not function properly or does not function as described in this guide, please contact your EMC representative.

About This Manual

This manual describes the tasks for setting up, configuring, and managing a storage system using EMC® SnapView™ software. Each major section includes introductory information and a general procedure for completing a task. This manual is not intended for use during the actual setup, configuration, and management of storage systems so the steps in the procedures purposely do not include screen captures of the dialog boxes.

The introductory information and *detailed* steps for each procedure appear in the SnapView online help so you have the complete information available when you actually set up, configure, and manage storage systems, should you require help.

Audience

This guide is part of the SnapView product documentation set, and is intended for use by customers and service providers who use EMC Navisphere® Manager software to set up and manage SnapView software.

Readers of this guide should be familiar with Navisphere Manager.

Organization

This manual is organized as follows:

- | | |
|------------|---|
| Chapter 1 | Introduces the EMC SnapView software application, including clones and snapshots. This chapter also lists the configuration guidelines and the right-click menu options available for SnapView in Navisphere Manager. |
| Chapter 2 | Describes the steps required for setting up clones and snapshots. |
| Chapter 3 | Describes the options available for using clones and snapshots. |
| Chapter 4 | Describes how to display and/or modify the properties dialog boxes for each SnapView component. |
| Chapter 5 | Contains examples, from setting up clones and snapshots to using them. Each example also contains an illustrated overview that shows the main steps outlined in the examples. |
| Appendix A | Describes how to use SnapView with a Tru64 server. |
| Appendix B | Reviews the EMC process for detecting and resolving software problems, and provides essential questions that you should answer before contacting the EMC Customer Support Center. |
| Glossary | Defines SnapView and other terms used in this guide. |

Related Documentation

Related documents include:

- ◆ *EMC Navisphere Manager Administrator's Guide* (P/N 069001125)
- ◆ *EMC SnapView Command Line Interfaces Reference* (P/N 069001181)
- ◆ Release notes for SnapView and admsnap

Conventions Used in This Guide

EMC uses the following conventions for notes and cautions.

A note presents information that is important, but not hazard-related.



CAUTION

A caution contains information essential to avoid data loss or damage to the system or equipment. The caution may apply to hardware or software.

Typographical Conventions

This manual uses the following format conventions:

This typeface	Indicates text (including punctuation) that you type verbatim, all commands, pathnames, filenames, and directory names. It indicates the name of a dialog box, field in a dialog box, menu, menu option, or button.
<i>This typeface</i>	Represents variables for which you supply the values; for example, the name of a directory or file, your username or password, and explicit arguments to commands.
This typeface	Represents a system response (such as a message or prompt), a file or program listing.
x > y	Represents a menu path. For example, Operations > Poll All Storage Systems tells you to select Poll All Storage Systems on the Operations menu.
[]	Encloses optional entries.
	Separates alternative parameter values; for example: <i>LUN-name LUN-number</i> means you can use either the LUN-name or the LUN-number.

Finding Current Information

The most up-to-date information about the SnapView software is posted on the EMC Powerlink™ website. We recommend that you download the latest information before you start the SnapView software.

To access EMC Powerlink, use the following link:

<http://powerlink.emc.com>

After you log in, select **Support > Document Library** and find the following:

- ◆ Release notes for SnapView and admsnap
- ◆ The latest version of this guide that is applicable to your software revision.
- ◆ *EMC Installation Roadmap for CX-Series, AX-Series and FC-Series Storage Systems*, which provides a checklist of the tasks that you must complete to install your storage system in a storage area network (SAN) or direct attach configuration.

Where to Get Help

For questions about technical support, call your local sales office or service provider.

If you have a valid EMC service contract, contact EMC Customer Service at:

United States: (800) 782-4362 (SVC-4EMC)

Canada: (800) 543-4782 (543-4SVC)

Worldwide: (508) 497-7901

Follow the voice menu prompts to open a service call and select the applicable product support.

Sales and Customer Service Contacts

For the list of EMC sales locations, please access the EMC home page at:

<http://www.EMC.com/contact/>

For additional information on the EMC products and services available to customers and partners, refer to the EMC Powerlink Web site at:

<http://powerlink.EMC.com>

Your Comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send a message to techpub_comments@EMC.com with your opinions of this guide.

About EMC SnapView Software

This chapter introduces the EMC® SnapView™ software and its user interface, as well as the two command line interfaces for it. The command line interfaces include the server-based admsnap utility and the EMC Navisphere® CLI interface.

If you are already familiar with SnapView from reading Chapter 1 in the *EMC SnapView Command Line Interfaces Reference* document, you can skip to Chapter 2, unless you want to learn about the SnapView right-click menu options; if so refer to page 1-16.

Major topics are

- ◆ Introduction to SnapView 1-2
- ◆ SnapView Components..... 1-7
- ◆ SnapView Servers..... 1-10
- ◆ Configuration Guidelines 1-12
- ◆ SnapView Behavior with AX-Series Storage Systems 1-15
- ◆ Navisphere Manager Right-Click Menu Options for SnapView..... 1-16
- ◆ Using Online Help 1-20

Introduction to SnapView

SnapView is a storage-system-based software application that allows you to create a copy of a LUN by using either clones or snapshots.

A clone, also referred to as a Business Continuance Volume (BCV), is an actual copy of a LUN and takes time to create, depending on the size of the source LUN. A snapshot is a virtual point-in-time copy of a LUN and takes only seconds to create.

SnapView has the following important benefits:

- ◆ It allows full access to production data with modest impact on performance and without the risk of damaging the original data.
- ◆ For decision support or revision testing, it provides a coherent, readable and writable copy of real production data.
- ◆ For backup, it practically eliminates the time that production data spends offline or in hot backup mode, and it offloads the backup overhead from the production server to another server.

Depending on your application needs, you can create clones, snapshots, or snapshots of clones. For a detailed overview of clones or snapshots, refer to the *Clones Overview* (see page 1-3) or to the *Snapshots Overview* (see page 1-4). For a comparison of using clones, snapshots, and snapshots of clones, refer to Table 1-1, *A Comparison of Clones and Snapshots*, on page 1-6.

Clones Overview

A clone is a complete copy of a source LUN. You specify a source LUN when you create a Clone Group. The copy of the source LUN begins when you add a clone LUN to the Clone Group. The software assigns each clone a clone ID. This ID remains with the clone until you remove the clone from its group.

While the clone is part of the Clone Group and unfractured (not accessible to a secondary server), any server write requests made to the source LUN are simultaneously copied to the clone. Once the clone contains the desired data, you can fracture the clone. Fracturing the clone breaks it from its source LUN and, once activated, makes it available to a secondary server.

Clone private LUNs record information that identifies data chunks on the source LUN and clone LUN that have been modified after you fractured the clone. A modified data chunk is a chunk of data that a server changes by writing to the clone or source LUN. A log in the clone private LUN records this information, but no actual data is written to the clone private LUN. This log reduces the time it takes to synchronize or reverse synchronize a clone and its source LUN since the software copies only modified (changed) chunks.

Writes made to the source LUN from the production server are copied to the fractured clone only when you manually perform a synchronization, which unfractures the clone and updates the contents on the clone with its source.

Figure 1-1 shows an example of how a fractured clone works. Note, as the production server writes to the source LUN, and the secondary server writes to the clone, the clone private LUN tracks areas on the source and clone that have changed since the clone was fractured.

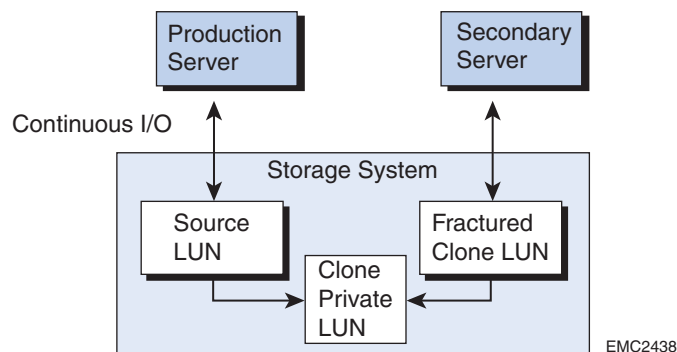


Figure 1-1 Clones Example

Snapshots Overview

A snapshot is a virtual LUN that allows a secondary server to view a point-in-time copy of a source LUN. You determine the point in time when you start a SnapView session. The session keeps track of the source LUN's data at a particular point in time.

During a session, the production server is still able to write to the source LUN and modify data. When this happens, the software stores a copy of the original point-in-time data on a reserved LUN in the SP's LUN pool. This operation is referred to as copy-on-first-write because it occurs only when a data chunk is first modified on the source LUN.

As the session continues and additional I/O modifies other data chunks on the source LUN, the amount of data stored in the reserved LUN pool grows. If needed, you can increase the size of the reserved LUN pool by adding more LUNs to the LUN pool.

Important An adequate number of reserved LUNs is essential since SnapView terminates sessions if the reserved LUN runs out of space and no additional LUNs are in the SP's LUN pool.

From a secondary server, you can view a session's point-in-time data by activating (mapping) a snapshot to the session. You can activate only one snapshot at a time to a session. If another point-in-time view is desired, you can deactivate (unmap) a snapshot from a session and activate it to another session of the same source LUN.

Though a snapshot appears as a conventional LUN to other servers, its data does not reside on a disk like a conventional LUN. A snapshot is a composite of the unchanged data chunks on the source LUN and data chunks that have changed on the reserved LUN. The data chunks on the source LUN are those that have not been modified since you started the session. The data chunks in the reserved LUN pool are copies of the original source LUN data chunks that have been modified since you started the session.

Figure 1-2 shows an example of how snapshots work. The snapshot is a composite of the source LUN and the reserved LUN in the SP's LUN pool.




Figure 1-2 Snapshot Example

SnapView also allows you to instantly restore a session's point-in-time data back to the source LUN, if the source LUN were to become corrupt or if a session's point-in-time data is desired for the source. You can do this by using SnapView's rollback feature.

Clone and Snapshot Tradeoffs

The following table describes the benefits and tradeoffs of using clones, snapshots, and snapshots of clones.

Table 1-1 A Comparison of Clones and Snapshots

	Clones	Snapshots	Snapshots of Clones
Benefits	<ul style="list-style-type: none"> Provides immediacy in replacing the contents of the source LUN with the contents of the clone LUN, should the source become corrupted. Makes backup operation nondisruptive. Provides enhanced protection against critical data loss because it is an actual LUN. Provides immediacy in redirecting servers from the source LUN to the clone LUN, should the source become corrupted. 	<ul style="list-style-type: none"> Provides immediacy in replacing the contents of the source LUN with the contents of the session, should the source LUN become corrupted. Makes backup operation nondisruptive. Provides a quick and instant copy because it is a virtual LUN. 	<ul style="list-style-type: none"> Provides immediacy in replacing the contents of the source LUN with the contents of the session, should the source LUN become corrupted. Makes backup operation nondisruptive. Provides an extra level of protection against critical data loss if both the source LUN and clone LUN become corrupted.
Creation Time	Minutes to hours. The creation time depends on the size of the source LUN. Subsequent synchronizations are incremental.	Instantaneous	Instantaneous
Disk Space Used	Uses the same amount of disk space as the source LUN.	Uses reserved LUN pool space, which is usually 10% to 20% of the source LUN size per session, but will vary depending on how much data has changed on the source LUN.	Uses reserved LUN pool space (for the snapshot) and full disk space (for the clone), which usually totals 100% of the source LUN size for clones and 10% to 20% of the source LUN size per session, but will vary depending on how much data has changed on the source LUN.
Data Recovery Time After Source LUN Failure/Corruption	Instantaneous after initializing a reverse synchronization.	Instantaneous after initializing a rollback operation.	Restore time from network or backup tape if fractured clone LUN received server I/O.
Performance Impact on the Storage System	<ul style="list-style-type: none"> There is no performance impact when a clone LUN is in a fractured state. For the initial synchronization of the clone LUN, performance decreases. Subsequent synchronizations or reverse synchronizations have minimal impact on performance. Impact is also determined by the synchronization rate, which is set when the clone LUN is added to the Clone Group. 	A slight performance decrease due to the copy-on-first-write.	Combination of both clone LUNs and snapshot LUNs.

SnapView Components

SnapView consists of the following software components:

- ◆ A set of drivers that provides the SnapView functionality, and resides on the storage system with the LUNs you want to copy.

All CX-Series storage systems ship from the factory with SnapView software installed, but *not* enabled. To use the SnapView software functionality, the SnapView Enabler must be installed on the storage system.

- ◆ The admsnap utility that provides a command line executable to let you manage clone and snapshot devices on the server. The admsnap utility ships with the SnapView enabler and resides on any servers connected to storage systems that have the SnapView software installed and enabled.
- ◆ User Interface (UI) - Navisphere Manager, which must be installed on at least one storage system on the same network as the SnapView storage system.

CX300, CX500, and CX700 storage systems ship from the factory with Navisphere Manager installed and enabled. CX200, CX400, and CX600 storage systems ship from the factory with Navisphere Manager installed, but *not* enabled. To use the Manager functionality on a CX200, CX400, or CX600 storage system, the Manager enabler must be installed on the storage system.

- ◆ Command Line Interface (CLI) - Navisphere CLI, which ships as part of the Navisphere Host Agent packages.

You must use Navisphere Manager or Navisphere CLI (not admsnap) to set up SnapView; then you can use admsnap and either Navisphere Manager or Navisphere CLI to manage ongoing SnapView operations.

Navisphere Manager

Navisphere Manager is a centralized storage-system management tool for configuring and managing CLARiiON® storage systems. It provides the following basic functionality:

- ◆ Discovery of CLARiiON storage systems
- ◆ Status and configuration information display
- ◆ Event management
- ◆ Storage configuration and allocation

Navisphere Manager 6.X is a web-based user interface that lets you securely manage CLARiiON storage systems locally on the same LAN or remotely over the Internet, using a common browser. Navisphere Manager 6.X resides on a CX-Series or FC-Series storage system or a Microsoft Windows Server 2003, Windows 2000, or Windows NT server that is running the Storage Management Server software, and is downloaded to the browser when the Storage Management Server software is accessed.

Navisphere CLI

The Navisphere CLI provides another management interface (along with Navisphere Manager and admsnap) to clones and snapshots. You can use Navisphere CLI commands and admsnap commands together to manage clones and snapshots. You need both admsnap and Navisphere CLI because admsnap interacts with the server operating system and CLI interacts with the storage system.

For additional information on how to use Navisphere CLI for SnapView and admsnap, refer to the *EMC SnapView Command Line Interface Reference*.

admsnap

The admsnap utility is an executable program that you can run interactively or via a script to manage clones and snapshots. The admsnap utility resides on the servers connected to the storage system with the SnapView driver.

The admsnap utility runs on the following server platforms:

- ◆ Hewlett Packard HP-UX
- ◆ IBM AIX (RS/6000 and RS/6000 SP servers)
- ◆ Linux (32-bit Intel platform, 64-bit AMD processor Linux, 64-bit Intel Xeon processor, and 64-bit Intel Itanium processor)

Important Separate admsnap installation packages are available for the 32-bit Intel platform, 64-bit AMD processor Linux/64-bit Intel Xeon processor, and the 64-bit Intel Itanium processor. The admsnap packages for the 64-bit AMD processor Linux and the 64-bit Intel Xeon processor are the same. For minimum supported Linux kernel revisions for each platform, refer to the *Admsnap Release Notes*.

- ◆ Microsoft Windows (Windows Server 2003, Windows 2000, and Windows NT)

Important Separate admsnap installation packages are available for Windows Server 2003, Windows 2000, and Windows NT. Unless indicated otherwise, all references to a Windows server in this document refer to Windows Server 2003, Windows 2000, and Windows NT servers *only*.

- ◆ Novell NetWare
- ◆ SGI IRIX
- ◆ Sun Solaris
- ◆ VMware® ESX Server™

For the supported versions of these servers/operating systems, see the release notes for SnapView and admsnap.

SnapView Servers

SnapView requires at least two servers: one server (called the production server) contains the LUN you want to copy, and another server (called the secondary server) lets you view the clone or snapshot. You can have multiple secondary servers.

The production server:

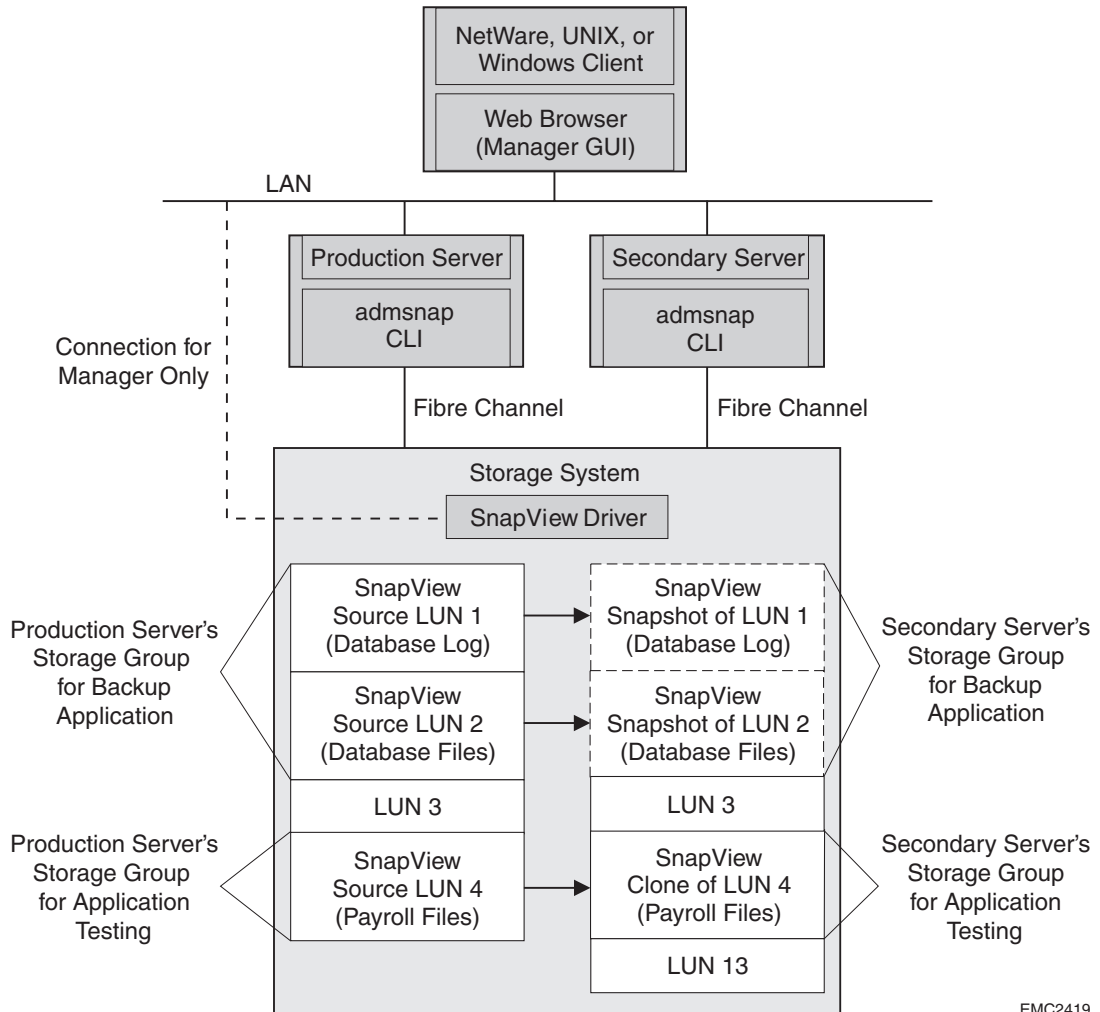
- ◆ Runs the customer applications
- ◆ Owns the source LUN

The secondary server (or any other server):

- ◆ Owns the clone or snapshot
- ◆ Reads from or writes to the fractured clone or activated snapshot
- ◆ Performs backup tasks using the clone or snapshot or an independent analysis (such as, decision support or revision testing).

You can configure a clustered server to access a source LUN, but not both the source LUN and its clone or snapshot. Only a server outside the cluster can access the clone or snapshot.

Figure 1-3 shows a sample SnapView environment with two servers creating clones and snapshots of two database files and their logs.



EMC2419

Figure 1-3 Sample SnapView Environment with Clones and Snapshots

Configuration Guidelines

This section lists the configuration guidelines for SnapView.

SnapView Limits

Table 1-2 lists the maximum SnapView limits for your storage system.

A metaLUN is a single entity; therefore it counts as one of your clone or snapshot limits. For example, if a metaLUN is composed of five LUNs and you create a clone of this LUN, it counts as one of your clone limits, not five. SnapView also supports the new, larger LUNs that FLARE™ supports (refer to the FLARE release notes).

Table 1-2 SnapView Configuration Guidelines

Parameter	CX700	CX600 or FC4700	CX500 or CX500i	CX400	CX300 or CX300i	AX100
Clones						
Per Storage System	200 ^a	100 ^b	100 ^a	50 ^b	100 ^a	Not supported
Per Source LUN	8	8	8	8	8	Not supported
Clone Groups						
Per Storage System	100	50	50	25	50	Not supported
Clone Private LUNs						
Per Storage System (required)	2	2	2	2	2	Not supported
Snapshots^c						
Per Storage System	300	300	150	150	100	4
Per Source LUN	8	8	8	8	8	1
SnapView Sessions^c						
Per Source LUN	8	8	8	8	8	1
Reserved LUNs						
Per Storage System	100	100	50	50	25	20

- Includes MirrorView™ primary and secondary images. The source is no longer counted towards the image limit.
- Includes source LUN and MirrorView primary and secondary images.
- The limits for snapshots and sessions include SnapView snapshots or SnapView sessions as well as reserved snapshots or reserved sessions used in other applications, such as SAN Copy™ (incremental sessions) and MirrorView™ / Asynchronous applications.

Event Monitor

Use the Event Monitor to monitor events specific to SnapView. The Event Monitor is an enterprise tool that supports Centralized or Distributed monitoring of storage systems in a heterogeneous environment. The Event Monitor software consists of two distinct parts: the Event Monitor User Interface (UI) and the Event Monitor.

The Event Monitor user interface (UI) is part of Navisphere Manager and runs on the web browser. The user interface provides you with an intuitive tool to set up responses for events and to choose which storage systems to observe. The user interface lets you customize a configuration to use any of the supported notification methods. You can easily configure it to e-mail, page, or send an SNMP trap to an industry standard event-management tool. The user interface need only be used when setting up configurations or viewing the Event History log.

Event Monitor resides on both Navisphere SP Agent and Host Agent and is available on many operating systems. Once configured, the Event Monitor runs continuously as a service or daemon, observing the state of all specified storage systems and notifying you when selected events occur.

To configure Event Monitor for SnapView, refer to the online help Table of Contents entry, *Monitoring and responding to events*, or Chapter 5 in the *EMC Navisphere Manager Administrator's Guide*, P/N 069001125.

You can configure Navisphere Manager to send an event notification when the reserved LUN pool becomes 50, 75, 80, 85, 90, or 95 percent full. The application will generate the following warning event code and description when each of the reserved LUN pool capacities is reached:

0x71004000 The reserved LUNs assigned to LUN X have become Y% percent full, where X is the source LUN and Y is either 50, 75, 80, 85, 90, or 95.

MirrorView

If a LUN is a MirrorView primary or secondary image, you cannot create a Clone Group for that image. Similarly, if a LUN is a member of a Clone Group, as the source or clone, it cannot serve as a MirrorView primary or secondary image.

If the MirrorView/Synchronous option is installed, you can create a snapshot of the primary or secondary image. However, we recommend that you only take a snapshot of a mirror's secondary image if the image's state is either Synchronized or Consistent. If the image is Synchronizing or Out-of-Sync, the snapshot's data will not be useful.

If the MirrorView/Asynchronous option is installed, you can create a snapshot of the primary or secondary image. However, we recommend that you only take a snapshot of a mirror's secondary image if the last update started has completed successfully. If the update did not complete successfully, for example, the image fractured or the update is still in progress, the snapshot's data will not be useful.

SAN Copy

You can use SnapView with SAN Copy to create a snapshot or a clone of the destination LUN, so that the SnapView replica can be put in the secondary server Storage Group, rather than the SAN Copy destination. This allows the SAN Copy destination to maintain consistency with its source, and be available on an ongoing basis for incremental updates. Keep in mind that SAN Copy tracks server writes to the SAN Copy source LUN (from the production server); but SAN Copy does not track server writes to the SAN Copy destination LUN (from the secondary server).

SnapView Behavior with AX-Series Storage Systems

As with Navisphere Manager on CX-Series storage systems, Navisphere Express on AX-Series storage systems allows users to manage SnapView. No new functionality is added. Since Manager's user interface must include functionality for all storage system types it supports, some SnapView commands and features are limited or unavailable for AX-Series storage systems. For more information, refer to the *Configuration Guidelines* section on page 1-12.

AX-Series to CX-Series Terminology Differences for SnapView

The following table lists and defines the SnapView terminology differences between AX-Series and CX-Series storage systems. For terminology differences between Navisphere Manager and Navisphere Express, refer to the *EMC Navisphere Manager Administrator's Guide*.

AX-Series Term	CX-Series Term	Navisphere Manager Definition
snapshot	SnapView session and snapshot	<p>A SnapView session is a point-in-time copy of a source LUN. The session keeps track of how the source LUN looks at a particular point in time.</p> <p>A snapshot is a virtual LUN and when activated, it allows a secondary server to view a SnapView session.</p> <p>You can create a snapshot before or after you start a SnapView session; however, the snapshot has no use until a secondary server activates it to a session.</p>
disk resources	reserved LUN pool	<p>The reserved LUN pool works with replication software, such as SnapView, SAN Copy, and MirrorView, to store data or information required to complete a replication task. For example, with SnapView, after you start a SnapView session and as the production server writes to the source LUN, the software stores a copy of the original data in the reserved LUN pool in chunks. When a secondary server activates the snapshot to the SnapView session, the snapshot views the original source LUN data chunks that have been modified since you started the session from the reserved LUN pool and unmodified data chunks from the source LUN.</p>

Navisphere Manager Right-Click Menu Options for SnapView

Table 1-3 describes the right-click options available for SnapView in Navisphere Manager's **Storage** tree. Depending on the status of the component or the version of SnapView that is running on the storage system you are managing, some menu options are unavailable. Some options may also be dimmed and unavailable to AX-Series storage systems.

Table 1-3 SnapView Basic Storage Tree Icons: Images and Descriptions





Image	Description	Menu Option	Used to
	Container for the storage system and all of its components. When the storage system is working normally, the software displays this icon.	SnapView > Start SnapView Session	Start a point-in-time copy of a source LUN(s).
		SnapView > SnapView Summary	Display the status of any snapshot and session for the selected storage system.
		SnapView > Create Clone Group	Designate a source LUN that you want to clone at some time.
		SnapView > Clone Feature Properties	Allocate and deallocate Clone Private LUNs and globally enable the Protected Restore feature.
	Container for SP A's or SP B's reserved LUN pool. SP A's and SP B's reserved LUN pool consists of any reserved LUNs owned by the selected SP.	Configure	Add or remove reserved LUNs to or from an SP's LUN pool.
		Properties	Display the properties of SP A's or SP B's reserved LUN pool.
	Container for the Clones and Snapshots icon.	None	NA
	Container for all Clone Groups. Note: This icon is expandable only if a Clone Group is created.	None	NA

Table 1-3 SnapView Basic Storage Tree Icons: Images and Descriptions




Image	Description	Menu Option	Used to
	Clone Group icon and container for clone IDs.	Add Clone	Create a relationship with the source and LUN you are adding, which will become the clone LUN.
		Destroy Clone Group	Destroy the relationship between the source LUN and the Clone Group.
		Properties	Display Clone Group, source LUN, and clone properties.
	Clone ID icon and container for clone LUNs.	Synchronize	Update the clone LUN with the data on its source LUN.
		Reverse Synchronize	Replace the data on the source LUN with the data on the clone.
		Fracture	Break off the clone from its source LUN so a secondary server can access the clone. The clone is still part of the Clone Group but any changes made to the source LUN are not copied to the clone.
		Remove	Break off the clone from its source LUN and remove it from the Clone Group. The clone LUN becomes a conventional LUN.
		Properties	Display the Clone Group, source LUN, and clone properties.
	Container for snapshots, reserved snapshots, SnapView sessions, and reserved sessions.	None	NA

Table 1-3 SnapView Basic Storage Tree Icons: Images and Descriptions







Image	Description	Menu Option	Use to
	Container for snapshot LUNs and the reserved snapshot container. <i>Snapshot Names</i> - individual snapshots.	None	NA
		Add to Storage Group	Add a snapshot to a Storage Group.
		Destroy Snapshot	Destroy a snapshot and any server writes made to the snapshot (if the snapshot is active).
		Activate Snapshot	Map a snapshot to a SnapView session and make it visible to a secondary server. Note: This option is available only if a snapshot is inactive.
		Deactivate Snapshot	Unmap a snapshot from a SnapView session and destroy any secondary server writes made to the snapshot. Note: This option is available only if a snapshot is active.
		Properties	Display the properties of a snapshot.
		SAN Copy > Create Session from LUN	Create a full SAN Copy session from the selected snapshot as the source LUN. This option is available only if the SAN Copy software is installed.
	Container for all the reserved snapshots running on the storage system. <i>Reserved Snapshots</i> - individual snapshot reserved for another application.	None	List the snapshots reserved for other applications, such as for the SAN Copy or MirrorView/Asynchronous software. This option is available only if the SAN Copy or MirrorView/Asynchronous software is installed. Note: You cannot perform any operations from this icon.
	Container for all SnapView sessions and the reserved sessions container. This icon appears even when no sessions are active in the storage system. <i>Sessions</i> - individual SnapView sessions.	None	List the sessions reserved for other applications, such as for the SAN Copy or MirrorView/Asynchronous software. This option is available only if the SAN Copy or MirrorView/Asynchronous software is installed. Note: You cannot perform any operations from this icon.
		Start Rollback	Replace the data on the source LUN with the data on the SnapView session.
		Stop Session	Destroy the point-in-time session and free the reserved LUN space used by the session.
		Properties	Display the properties of a SnapView session.

Table 1-3 SnapView Basic Storage Tree Icons: Images and Descriptions

Image	Description	Menu Option	Use to
	Container for all the reserved sessions running on the storage system. <i>Reserved Sessions</i> - individual sessions reserved for another application.	None	List the sessions reserved for other applications, such as for the incremental SAN Copy or MirrorView/Asynchronous software. This option is available only if the incremental SAN Copy or MirrorView/Asynchronous software is installed. Note: You cannot perform any operations from this icon.
	LUN icon - LUNs in a Storage Group.	SnapView > Create Snapshot	Create a virtual LUN that allows a secondary server to view a SnapView session.
		SnapView > Start SnapView Session	Start a point-in-time copy of the selected source LUN.
		SnapView > Create Clone Group	Create a Clone Group with the selected LUN as the source.
	MetaLUN icon -Type of LUN whose capacity is the combined capacities of all the LUNs that compose it.	SnapView > Create Snapshot	Create a virtual LUN that allows a secondary server to view a SnapView session.
		SnapView > Start SnapView Session	Start a point-in-time copy of the selected source LUN.
		SnapView > Create Clone Group	Create a Clone Group with the selected LUN as the source.

Using Online Help

The following online help is available from the Navisphere Manager user interface:

- ◆ **A set of organized, linked help topics**

To access the online help table of contents, click **Help > Help Topics** on the Menu bar in the application's Main window, or click the help icon on the Toolbar.

- ◆ **Context-sensitive help topics**

To display context-sensitive help, click the **Help** button displayed in each dialog box.

This chapter gives a general overview of how to set up SnapView in order to use clones or snapshots. For detailed information on these topics, please refer to the online help.

Major sections in this chapter are

Setting Up Clones

- ◆ Prerequisites for Setting Up Clones.....2-2
- ◆ Overview of Setting Up SnapView to Use Clones2-3
- ◆ Allocating Clone Private LUNs2-4
- ◆ Deallocating/Reallocating Clone Private LUNs2-6
- ◆ Creating a Clone Group2-7
- ◆ Adding a Clone to a Clone Group.....2-9

Setting Up Snapshots

- ◆ Prerequisites for Setting Up Snapshots.....2-12
- ◆ Overview of Setting Up SnapView to Use Snapshots2-14
- ◆ Reserved LUN Pool with SnapView2-15
- ◆ Starting a SnapView Session2-17
- ◆ Creating a Snapshot.....2-24
- ◆ Adding a Snapshot to a Storage Group2-27

Setting Up SnapView to Use Clones

This section describes how to set up SnapView to use clones.

Prerequisites for Setting Up Clones

Before you can set up and use clones, the following must be true:

- ◆ **Source LUNs must be bound.** SnapView requires Access Logix™ software and Access Logix requires LUNs to be bound on a RAID Group. For a client or production server to access a source LUN, you must assign the source LUN to a Storage Group and connect the Storage Group to the production server. To do this, Access Logix must be enabled.
- ◆ **LUNs that you will plan on using as clone LUNs must be bound.** These LUNs must be the same size as the source LUNs that you will clone. EMC strongly recommends that you bind your clone LUNs on RAID Groups that are different than their source LUNs. The clone's RAID Group does *not* have to be the same RAID type as the source LUN.
- ◆ **If the source or clone LUNs are on a VMware® ESX Server, the LUNs must be configured as raw device mapping volumes.** If you will be using source or clone LUNs that are not already configured as a raw device mapping volume and set to physical compatibility mode, use the VMware **vmkfstools** utility to reconfigure them. For information on using this utility, refer to the *Managing Raw Device Mappings - Utilities* section found at http://www.vmware.com/pdf/esx25_rawdevicemapping.pdf.
- ◆ **For VMware ESX Servers, verify that the source LUN is presented to the Virtual Machine** (guest operating system running on the Virtual Machine). For information on how to present a LUN to the Virtual Machine, refer to the VMware documentation that shipped with your ESX Server.

- ◆ **For a secondary server to access the clone LUN, the clone must be assigned to a Storage Group** (but you cannot read the clone until you fracture it). The Storage Group must be connected to the secondary server that will access the clone. You must assign the clone LUN to a Storage Group other than the Storage Group that holds the source LUN. EMC supports placing a clone in the same Storage Group as its source LUN *only if* you use RM/Local, RM/SE, or Powerpath Volume Manager to put the clone in the Storage Group. This software provides *same host access* to the clone and the source LUN. For information on using these software products, refer to the documentation for the product.

If you have a VMware ESX Server, the clone and source LUNs must be accessed by different VMs, unless the VM is running one of the software programs that support same host access.

- ◆ **Configure Event Monitor, if you want to be notified of SnapView events.** Event Monitor is part of the Navisphere Agent and is available on many operating systems. Once configured, the Event Monitor runs continuously as a service or daemon, observing the state of all specified storage systems and notifying you when selected events occur. To configure Event Monitor for SnapView, refer to the online help Table of Contents entry, *Monitoring and responding to events*, or to the *EMC Navisphere Manager Administrator's Guide*, P/N 069001125.

What Next? To set up clones, continue to the next section.

Overview of Setting Up SnapView to Use Clones

The following is a checklist for setting up SnapView clones.

- Allocate clone private LUNs, refer to page 2-4.

This step is required only in the initial setup of clones.

- Create a Clone Group, refer to page 2-7.
- Add a clone to a Clone Group, refer to page 2-9.

To learn about the possible clone states after you add a clone to a Clone Group, refer to *Clone States* on page 3-2.

What Next? Use the clone as described in Chapter 3.

Allocating Clone Private LUNs

Clone private LUNs record information that identifies data chunks on the source LUN and clone LUN that have been modified after you fractured the clone. A modified data chunk is a chunk of data that a server changes by writing to the clone or source LUN. A log in the clone private LUN records this information, but no actual data is written to the clone private LUN. This log reduces the time it takes to synchronize or reverse synchronize a clone and its source LUN since the software copies only modified (changed) chunks.

You must allocate one clone private LUN for each SP before you can create a Clone Group.

Eligible Clone Private LUNs

Each clone private LUN must be a minimum of 250000 blocks. You must bind these LUNs before you allocate them as a clone private LUN. You can use any LUN that is at least 250000 blocks in size as a clone private LUN, *except* for the following:

- ◆ Hot spare LUNs
- ◆ MirrorView remote mirror LUNs (LUNs used as either a primary or secondary image)
- ◆ SnapView clone, snapshot, or source LUNs
- ◆ SAN Copy source or destination logical units.
- ◆ Private LUNs (LUNs reserved as clone private LUNs or for use by the reserved LUN pool).

Clone private LUNs larger than 250000 blocks provide no performance benefit.

To bind these LUNs refer to the section, *Creating LUNs on RAID Groups*, in the *EMC Navisphere Manager Administrator's Guide*.

You should bind clone private LUNs in a RAID Group that normally does not see heavy I/O.

To Allocate Clone Private LUNs

1. From any client that is managing the storage system, in the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Right-click the storage system icon for which you want to allocate clone private LUNs and select **SnapView > Clone Feature Properties**. The **Clone Feature Properties** dialog box opens.
3. Select an eligible LUN from the **Available LUNs** list.
4. Click the right-arrow button to move the selected LUN to the **Clone Private LUNs** list.
5. Repeat steps 3 and 4, so that two LUNs (one for each SP) are in the **Clone Private LUNs** list.

You cannot specify which SP the clone private LUN is assigned to.

6. If you plan on using the protected restore feature, select the **Allow Protected Restore** option to globally enable it. For information about the protected restore feature, refer to page 3-14.

Important When you select the **Allow Protected Restore** option, the SnapView driver automatically allocates 8 MB in additional memory per SP. The additional memory is fixed and is used to copy the data from the clone LUN to the source LUN in order to satisfy server write requests to the source LUN. This additional memory counts against the total memory budget for storage-system-based drivers.

7. Click **OK**, and then **Yes** to confirm the allocation of the selected clone private LUNs.

What Next?

To continue setting up clones, go to the section, *Creating a Clone Group* on page 2-7, to create a Clone Group.

To reallocate the clone private LUNs, continue to the next section.

Deallocating/Reallocating Clone Private LUNs

If a clone private LUN fails or if you decide that you want to use a different LUN as your clone private LUN, you can reallocate both clone private LUNs with existing Clone Groups and clones. However, if you reallocate one clone private LUN, you must reallocate the other clone private LUN.

If you want to deallocate the clone private LUNs, you must remove all clones and destroy any Clone Groups. You must also globally disable the **Allow Protected Restore** option in the **Clone Features Properties** dialog box.

To Deallocate/ Reallocate Clone Private LUNs

1. From any client that is managing the storage system, in the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Right-click the icon for the storage system for which you want to reallocate the clone private LUNs and select **SnapView > Clone Feature Properties**. The **Clone Feature Properties** dialog box opens.
3. Select the clone private LUN you want to remove from the **Clone Private LUNs** list, and then use the left-arrow button to move it to the **Available LUNs** list.
4. Select an eligible LUN (see page 2-4) from the **Available LUNs** list, and then click the right-arrow button to move the selected LUN to the **Clone Private LUNs** list.

You do not specify which SP the clone private LUN is assigned to; Navisphere does this for you.

5. Deallocate the remaining clone private LUN by repeating steps 3 and 4 and select another eligible LUN.
6. Click **OK**, and then **Yes** to confirm the new allocation of the selected clone private LUNs.

The software transfers the information stored in the previous clone private LUN to the new clone private LUN.

What Next?

To continue setting up clones, continue to the next section to create a Clone Group.

Creating a Clone Group

A Clone Group contains a source LUN and all of its clones. When you create a Clone Group you specify a LUN to be cloned. This LUN is referred to as the source LUN. Once you create the Clone Group, the SnapView software assigns a unique ID to the group. No clones of the specified source LUN exist until you add a clone to the Clone Group. The purpose of creating a Clone Group is to designate a source LUN that you want to clone at some time.

If you have not allocated two clone private LUNs, you must allocate them before you create a Clone Group.

- Eligible LUNs** Any source LUN is eligible to be cloned, except for the following:
- ◆ Hot spare LUNs
 - ◆ Remote mirror LUNs (LUNs participating as either a primary or secondary image)
 - ◆ Clone LUNs (LUNs participating in any Clone Group as either a source LUN or a clone LUN)
 - ◆ Snapshot LUNs
 - ◆ Private LUNs (LUNs reserved as clone private LUNs, reserved LUN pool, or write intent log)

To Create a Clone Group

1. From any client that is managing the storage system, in the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Right-click the icon for the storage system for which you want to create a Clone Group and select **SnapView > Create Clone Group**. The **Create Clone Group** dialog box opens.
3. In **Name**, enter a valid name for the Clone Group.

A valid Clone Group name must be at least one character. This character must not exceed 64 ASCII characters and must be unique per storage system. The software removes any preceding or following whitespaces from the name prior to verifying the validity of, and setting, the name.

4. In **Description**, enter an optional description about the Clone Group.

The information is limited to 256 ASCII characters. The software removes any preceding or following whitespaces from the description prior to verifying the validity of, and setting, the description.

5. In **Quiesce Threshold**, enter the quiesce threshold for the Clone Group.

The quiesce threshold is the time period after which the source LUN is not receiving any server writes. After this time period, the system transitions the unfractured clone, which is in a Consistent state, to a Synchronized state.

You set the quiesce threshold on a per Clone Group basis. Any clone you add to this Clone Group will retain this quiesce value. Valid values are 10 – 3600 seconds. The default is 60 seconds.

6. From the **LUN to be Cloned** list, select an eligible source LUN.

If you created the Clone Group by right-clicking a storage-system icon, the list includes all source LUNs available on the storage system and you can select multiple source LUNs.

If you created the Clone Group by right-clicking a source LUN icon, the list includes only the selected source LUN.

7. Click **OK** to create the Clone Group, and then **Yes** to confirm the creation of the Clone Group.

If the action is successful, Navisphere closes the dialog box and places an icon for the Clone Group in the Storage tree.

If the action is not successful, an error message is displayed.

You can expand the capacity of a LUN or metaLUN that is participating in a Clone Group, but neither the production or secondary server can access this added capacity until the expansion is complete and you perform some additional operations. The expansion process involves removing the clone from the Clone Group and destroying the Clone Group. For detailed information on expanding a LUN or metaLUN, see the online help or the *EMC Navisphere Manager Administrator's Guide* (P/N 069001125).

What Next? Continue to the next section to add a clone to the Clone Group.

Adding a Clone to a Clone Group

When you add a clone to a Clone Group, you create an actual copy of the source LUN you specified when creating the group. The copy of the source LUN begins when you add a clone LUN to the Clone Group and the **Initial Sync Required** option is selected.

The software assigns a clone ID to each clone. This ID remains with the clone until you remove the clone from its group.

Important for a Windows server - You must delete all file entries in the recycling bin of the source LUN before adding the clone to the Clone Group. If you do not delete these entries, the clone you are adding will copy them byte for byte.

Source LUN Writes and Unfractured Clone States

While the clone is part of the Clone Group and unfractured (not accessible to a secondary server), server writes to the source LUN can continue, with the following results:

Server writes to the source — When a source LUN receives a server write request, the clone transitions into a Consistent state because the clone is no longer a byte-for-byte copy of its source. Writes made to the source LUN are simultaneously copied to the clone, which results in the following possible states:

- ◆ Synchronized — the clone transitions into a Synchronized state if no other server writes are made to the source LUN for the duration of the quiesce threshold. The quiesce threshold is the amount of time that must pass after a server write request before the software transitions any unfractured clones to a Synchronized state. You specify the quiesce threshold when you create a Clone Group.
- ◆ Consistent - the clone remains in a consistent state if server writes are made to the source LUN during the duration of the quiesce threshold.

When you add a clone to the Clone Group, with the **Initial Sync Required** property selected, the clone state is Synchronizing. The software transitions the clone to Synchronized or Consistent state only after the initial synchronization is complete.

No server writes to the source — If a source LUN does not receive any server writes for the duration of the quiesce threshold, any unfractured clone in a Consistent state transitions into a Synchronized state.

To Add a Clone to a Clone Group

For additional information on the `admsnap` commands described below, refer to the *EMC SnapView Command Line Interfaces Reference*.

Using Navisphere from any client that is managing the storage system, do the following:

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the **Clones** icon, and then to the **Clone Group Name** icon to which you want to add a clone.
2. Right-click the **Clone Group Name**, and then select **Add Clone**. The **Add Clone** dialog box opens.
3. In the **Select Clone LUN** list, select the LUN you want to become a clone. The **Select Clone LUN** list only contains LUNs of the same size as the source LUN.

Important If the selected LUN has data and you selected the **Initial Sync Required** option, SnapView will destroy the current contents of the LUN.

If the clone LUN you are adding does not belong to the same SP as its source LUN, the clone LUN will trespass over to the source LUN's SP. The **Trespassed LUNs** dialog box reports any LUNs that have different default and current owners. To open this dialog select **Tools** from the main menu and click **Trespassed LUNs Status**.

4. To specify an initial synchronization on the clone you are adding, select **Initial Sync Required**. This option is not required but Navisphere selects it by default. Depending on the size of the source LUN, the initial synchronization process can take a few minutes to a few hours.

The **Initial Sync Required** is necessary unless your source LUN does not contain any data, for example, if you bind a source LUN and have not added it to a Storage Group. If you select **Initial Sync Required** with this empty source LUN, resources are used to synchronize the empty source LUN to the clone.

5. To specify the **Protected Restore** feature on the clone you are adding, select **Use Protected Restore**.

To use this feature, you must globally enable it by selecting the **Allow Protected Restore** option in the **Clone Features Properties** dialog box. For information on the Protected Restore feature, refer to *Protected Restore Feature* on page 3-14.

6. In **Recovery Policy**, select the method for recovering the clone after a failure. Options are **Automatic** or **Manual**.
7. In **Synchronization Rate**, specify a relative value (low, medium, or high) for the priority of completing updates. **High** completes updates faster, but may significantly affect storage system performance for host I/O requests. **Low** completes updates slower, but also minimizes the impact on other storage system operations.

EMC recommends that you do not use a **High** synchronization rate on a storage system with a single SP.

8. Click **Apply** to add a clone to the Clone Group, and then click **Yes** to confirm the addition of the clone. A **Success: Add Clone** dialog box opens.
9. Click **OK**. The application places a **Clone** icon under the associated **Clone Group Name** icon.
10. Wait for the synchronization to complete and verify that the clone is in a Consistent or Synchronized state by right-clicking the **Clone Group Name** icon that contains the clone you added, select **Properties**, and then click the appropriate **Clone** tab.

To create multiple clones of this source LUN, repeat steps 2 through 10.

You can expand the capacity of a LUN or metaLUN that is participating in a Clone Group, but neither the production or secondary server can access this added capacity until the expansion is complete and you perform some additional operations. The expansion process involves removing the clone from the Clone Group and destroying the Clone Group. For detailed information on expanding a LUN or metaLUN, see the online help or the EMC *Navisphere Manager Administrator's Guide* (P/N 069001125).

What Next?

To start using the clone or for information on possible clone states, refer to Chapter 3, *Using SnapView*.

Setting Up SnapView to Use Snapshots

This section describes how to set up SnapView to use snapshots.

Prerequisites for Setting Up Snapshots

Before you can set up and use snapshots, the following must be true:

- ◆ **Source LUNs must be bound.** SnapView requires Access Logix and Access Logix requires LUNs to be bound on a RAID Group. For a client or production server to access a source LUN, you must assign the source LUN(s) to a Storage Group and connect the Storage Group to the production server. To do this, you must enable Access Logix.
- ◆ **If the source LUN is on a VMware ESX Server, the LUN must be configured as a raw device mapping volume.** If you will be using source LUNs that are not already configured as a raw device mapping volume and set to physical compatibility mode, use the VMware **vmkfstools** utility to reconfigure them. For information on using this utility, refer to the *Managing Raw Device Mappings - Utilities* section found at http://www.vmware.com/pdf/esx25_rawdevicemapping.pdf.
- ◆ **For VMware ESX Servers, verify that the source LUN is presented to the Virtual Machine** (guest operating system running on the Virtual Machine). For information on how to present a LUN to the Virtual Machine, refer to the VMware documentation that shipped with your ESX Server.
- ◆ **For a secondary server to access the snapshot, you must assign the snapshot to a Storage Group.** The Storage Group must be connected to the secondary server that will activate the snapshot. You must assign the snapshot to a Storage Group other than the Storage Group that holds the source LUN. EMC supports placing a snapshot in the same Storage Group as its source LUN only if you use RM/Local, RM/SE, or Powerpath Volume Manager to put the snapshot in the Storage Group. This software provides same host access to the snapshot and the source LUN. For information on using these software products, refer to the documentation for the product.

If you have a VMware ESX Server, the snapshot and source LUNs must be accessed by different VMs, unless the VM is running one of the software programs that support same host access.

VMware ESX Servers *must* activate the snapshot before adding it to a Storage Group.

- ◆ **Reserved LUNs must be bound and added to the reserved LUN pool before starting a SnapView session.** Since each SP has its own reserved LUN pool, you must add a reserved LUN to the LUN pool of the source LUN's SP. To configure the reserved LUN pool for SnapView, refer to the online help Table of Contents entry, **Managing Storage Systems > Configuring and Monitoring the Reserved LUN Pool** or the chapter on the reserved LUN pool in the latest revision of the *EMC Navisphere Manager Administrator's Guide*.
- ◆ **Configure Event Monitor, if you want to be notified of SnapView events.** Event Monitor is part of the Navisphere Agent and is available on many operating systems. Once configured, the Event Monitor runs continuously as a service or daemon, observing the state of all specified storage systems and notifying you when selected events occur. To configure Event Monitor for SnapView, refer to the online help Table of Contents entry, *Monitoring and responding to events*, or to the *EMC Navisphere Manager Administrator's Guide*.

What Next? To set up snapshots, continue to the next section.

Overview of Setting Up SnapView to Use Snapshots

The following is a checklist for setting up SnapView snapshots.

Important You must complete the prerequisites for setting up snapshots, as listed on page 2-12, before you can perform any of the following procedures.

- Start a SnapView session, refer to page 2-17.

You can create a snapshot before starting a session but the snapshot has no use until you start a session on it. A secondary server can then activate the snapshot to the session.

- Create a snapshot, refer to page 2-24.
- If you do *not* have a VMware ESX Server** - Add the snapshot to a Storage Group connected to the server that will access the snapshot, refer to page 2-27.

If you have a VMware ESX Server - Activate the snapshot, refer to page 3-23.

What Next? Use the snapshot as described in Chapter 3.

Reserved LUN Pool with SnapView

Important With SnapView Version 02.03.xxx (or higher), the snapshot cache is referred to as the reserved LUN pool. The reserved LUN pool works with SnapView in the same way as the snapshot cache. However, unlike the snapshot cache, which was used solely for SnapView, the reserved LUN pool shares its LUN resources with other applications such as SAN Copy and MirrorView/Asynchronous. The only visible change in the Navisphere User Interface (UI) is in the tree structure. The reserved LUN pool is now structured directly under the **Storage System** icon instead of the **SnapView** icon.

The reserved LUN pool consists of one or more private LUNs and works with SnapView sessions and snapshots. The reserved LUN pool stores the original source LUN data chunks that have been modified since the start of the session. For any one session, the contents of a reserved LUN(s) and any unchanged source LUN(s) blocks compose the snapshot.

Server writes made to an activated snapshot are also stored on a reserved LUN in the SP's LUN pool. When you deactivate the snapshot, the reserved LUN space is freed and all server writes are destroyed.

Each SP has its own reserved LUN pool, and before starting a session, the reserved LUN pool must contain at least one LUN for each source LUN that will be starting a session. You can add any LUNs that are available to either SP's reserved LUN pool. Each SP manages its own LUN pool and assigns a separate reserved LUN (or multiple LUNs) to each SnapView source LUN. Multiple sessions of a single source LUN will share the same reserved LUN or LUNs.

If the reserved LUN fills up and the SP's LUN pool has no additional LUNs, the software automatically terminates the session that is trying to allocate reserved LUN space, logs an error, releases the reserved LUN(s) used by this session, and returns them to the SP's LUN pool. The software also destroys all copy-on-first-write data stored in the reserved LUN pool for that session. At this point, the snapshot becomes inactive and any server that has mounted volumes on the snapshot will receive I/O errors and lose access.

If you have multiple sessions of a single source LUN and the reserved LUN fills up, when the production server modifies a chunk on the source LUN, resulting in a copy-on-first-write, every session that has the same chunk will be terminated if no additional LUNs are

available in the SP's LUN pool. Other sessions that did not have this chunk will continue to run and use the reserved LUN space that the terminated sessions were using.

Important SnapView, SAN Copy, and MirrorView / Asynchronous share the LUN resources of the reserved LUN pool. For example, if you are running an incremental SAN Copy session on one LUN and a SnapView session on another LUN, the reserved LUN pool must contain at least two LUNs - one for each source LUN. If both sessions are running on the same source LUN, the sessions will share a reserved LUN.

Configuring the Reserved LUN Pool

You must configure the reserved LUN pool before you start a SnapView session. Use Navisphere Manager to configure the reserved LUN pool (refer to the online help topic **Managing Storage Systems > Configuring and Monitoring the Reserved LUN Pool** or refer to the chapter on the reserved LUN pool in the latest revision of the *EMC Navisphere Manager Administrator's Guide*).

To manually estimate a suitable LUN pool size, refer to **Managing Storage Systems > Configuring and Monitoring the Reserved LUN Pool** in the Table of Contents for the Navisphere Manager online help and select the **Estimating the Reserved LUN Pool Size** topic or refer to the chapter on the reserved LUN pool in the latest revision of the *EMC Navisphere Manager Administrator's Guide*.

Starting a SnapView Session

A SnapView session is a point-in-time copy of a source LUN. The session keeps track of how the source LUN looks at a particular point in time.

After you start a SnapView session and as the production server writes to the source LUN(s), the software stores a copy of the original data in the reserved LUN pool in chunks. This copy is referred to as copy-on-first-write and occurs only once, which is when the server first modifies a data chunk on the source LUN(s).

A secondary server can then activate (map) a snapshot to the SnapView session. The snapshot views the original source LUN data chunks that have been modified since you started the session from the reserved LUN pool and unmodified data chunks from the source LUN(s).

Session Name

You must give each session a name when you start the session. The name persists throughout the session and is viewable through Navisphere Manager. Use the name to determine session status and to stop the session.

Before starting a SnapView session, the SP of the source LUN(s) must contain at least one free (unallocated) LUN in its reserved LUN pool.

Optional Modes

When you start a SnapView session, you can specify that the session run in persistent and/or consistent mode.

A SnapView session can run in both persistent and consistent mode. For AX-Series storage systems, persistent mode is always enabled and consistent mode is not supported.

Consistent mode — Preserves the point-in-time restartable copy across a set of source LUNs. The SnapView driver will delay any I/O requests to the set of source LUNs until the session has started on all LUNs (thus preserving the point-in-time restartable copy on the entire set of LUNs). In the event of a failure, the software will not start the session on any source LUN and will display an error message. Consistent mode also prevents you from adding other LUNs to the session.

A restartable copy is a data state having dependent write consistency and where all internal database/application control information is consistent with a Database Management System/application image.

You would use consistent mode if

- ◆ you want a consistent replica across a set of LUNs.
- ◆ you do *not* want to stop I/O to the source LUN(s) before starting a session.
- ◆ you want to prevent other LUNs from being added to the session. If desired, you can start a consistent session on a single source LUN to prevent other LUNs from being added to the session.
- ◆ the set of source LUNs that you will use to start the session spans both SP's; however, the source LUNs cannot span storage systems.

The consistent feature is available on a per session basis (not per snapshot or source LUN) and counts as one of the eight sessions per source LUN limit.

Persistent mode — Creates a session that can withstand the following failures and trespasses:

- ◆ SP reboot or failure
- ◆ Storage system reboot or power failure
- ◆ Server I/O trespassing to the peer SP

The persistence feature is available on a per session basis (not per snapshot or source LUN). In the event of a failure, reserved LUNs along with the source LUNs will trespass to the other SP. Depending on your failover software, once the failed SP is running, you may need to issue a restore command in order to restore the proper source LUNs and reserved LUNs back to their original SP. For the appropriate restore command, refer to the documentation that shipped with your failover software.

For information on supported failover software for the storage system you are managing, refer to the Release Notes for SnapView and admsnap.

Multiple Sessions

If you have an AX-Series storage system, you can start only one SnapView sessions per source LUN(s).

If you have a CX-Series or FC4700-Series storage system, you can start up to eight concurrent sessions per source LUN(s). However, each snapshot must be activated to a different SnapView session and accessed by different servers. For example, if you create eight snapshots and start three sessions for a single source LUN, three different servers can activate three of the snapshots to three different sessions. The remaining five snapshots will be inactive.

You can also use the deactivate and activate functions to change the focus of a snapshot from one session to another. You must deactivate (unmap) a snapshot before you can activate (map) it to another session. For example, if you start eight sessions and create one snapshot for a single source LUN, a secondary server can not activate to another session until they deactivate from the current session. Once the secondary server deactivates the snapshot from the session, you can activate and deactivate between the remaining sessions.

Refer to Chapter 3 to activate or deactivate the snapshot.

The eight-session limit includes SnapView sessions and any reserved sessions used in another application such as SAN Copy and MirrorView/Asynchronous.

Before Starting a SnapView Session

Before starting a SnapView session, verify the following:

- You have configured the reserved LUN pool.
- The source LUNs are currently not being used to start another consistent session. Once the consistent session starts on all the source LUNs, you can start another consistent session using the same source LUNs.
- If MirrorView/Asynchronous is installed, the source LUNs that you will use to start the consistent session cannot be in the process of starting an update to a group. If they are, you may receive an error when you attempt to start the consistent session. Once the mirror or group is updating or has been updated and before the next update occurs, you can start the consistent session using the same source LUNs.

You can expand the capacity of a LUN or metaLUN that is participating in a SnapView session, but neither the production or secondary server can access this added capacity until the expansion is complete and you perform some additional operations. The expansion process involves stopping all SnapView sessions and destroying all snapshots. For detailed information on expanding a LUN or metaLUN, see the online help or the *EMC Navisphere Manager Administrator's Guide* (P/N 069001125).

To Start a SnapView Session

For additional information on the `admsnap` commands described below, refer to the *EMC SnapView Command Line Interfaces Reference*.

1. If you will start a session in consistent mode, skip to step 2. If you will *not* start a session in consistent mode, you must do the following:
 - a. From the production server, stop I/O to the source LUN(s).
 - b. From the production server, flush all cached data to the source LUN(s), by issuing the appropriate command for your operating system.
 - For a Windows server, use the **admsnap flush** command.
 - For Solaris, HP-UX, AIX, and Linux servers, unmount the file system by issuing the **umount** command. If you are unable to unmount the file system, issue the **admsnap flush** command.

Neither the **flush** command nor the **sync** command is a substitute for unmounting the file system. Both commands only complement unmounting the file system.

- For an IRIX server, the **admsnap flush** command is not supported. Unmount the file system by issuing the **umount** command. If you cannot unmount the file system, use the **sync** command to flush buffers. The **sync** command reduces the number of times you need to issue the **fsck** command on the secondary server's file system. Refer to your system's man pages for **sync** command usage.
- For a Novell NetWare® server, use the **dismount** command on the volume to dismount the file system.

With some operating systems, you may need to shut down the application to flush the data. Specific operating systems have different requirements.

2. From any client that is managing the storage system, in the **Enterprise Storage** dialog box, click the **Storage** tab.
 - a. Navigate to the icon of the source LUN(s) on which you want to start a session and select **SnapView > Start Session**.

If you want to group logically related LUNs together, especially for consistency operations, you can start a single session on multiple source LUNs. To do this, navigate to the storage system icon and select **SnapView > Start Session**.

If the **Start Session** option is unavailable, make sure that at least one free (unallocated) LUN exists in the SP's reserved LUN pool or that you have not exceeded the eight session limit per source LUN. These limits include any reserved sessions. Reserved sessions are sessions used with another application such as SAN Copy and MirrorView/Asynchronous.

- b. Enter a unique name for the session. If the name you specify is already assigned to another session, an error message appears.

To agree with an admsnap requirement for a valid session name, include only characters, numbers, spaces, and underscores in **Session Name**. Symbols are not supported.

- c. In **Optional Mode(s)**, you may select the following modes:
 - Select **persistent mode**, if you want the session to withstand failures and trespasses. Persistent mode is always enabled on AX-Series storage systems.
 - Select **consistent mode** if you want a consistent replica across a set of LUNs, you do *not* want to stop I/O to the source LUNs before starting a session, you want to prevent other LUNs from being added to the session, or the set of source LUNs that you will use to start the session span both SP's. Consistent mode is not supported on AX-Series storage systems.

A SnapView session can run in both persistent and consistent mode. If you do not select persistent or consistent mode, the session will run in normal mode.

- d. In **Source LUNs Available**, select the source LUN(s) on which you want to start the session.

If you start the session by right-clicking a storage-system icon, the list includes all source LUNs available on the storage system and you can select multiple source LUNs.

If you start a session by right-clicking a source LUN icon, the list includes only the selected source LUN.

- e. Click **OK** to start the session.

If the start of the session is successful, Navisphere places an icon along with the name you assigned to the session below the session's source LUN(s) and the storage system in the **Storage** tree.

3. From the production server, resume I/O to the source LUN(s).

If you started a consistent session, while the session is in the process of starting on all LUNs, you cannot

- stop the session. However, the software may stop the session immediately after the start completes if I/O is queued and there is not enough space in the SP's reserved LUN pool.
- perform a rollback operation. You must wait until the session starts on the entire set of source LUNs.
- activate a snapshot to the session. You must wait until the session starts on the entire set of source LUNs.

While the session is running, use the **SnapView Session Properties** dialog box to monitor the reserved LUN pool usage for the SP, see *SnapView Session Properties* on page 4-6.

SnapView Session States

After you start a SnapView session, each session is in a particular state. The state of the session is either normal, rolling back, or stopping, as described in Table 2-1.

Table 2-1 Session States

Session State	Description
Normal	A session that has started.
Rolling Back	A session is in the process of restoring its point-in-time copy data back to the source LUN(s). This option is available with SnapView Version 2.2 (or higher).
Stopping	A session that has been stopped and is in the process of cleaning up some driver data structures.

What Next?

What you do next depends on whether you have created a snapshot to map (activate) to this session.

- ◆ If you have not created a snapshot, continue to the next section.
- ◆ If you have created a snapshot but have not added the snapshot to a Storage Group, go to *Adding a Snapshot to a Storage Group* on page 2-27.

Creating a Snapshot

A snapshot is a virtual LUN and when activated, it allows a secondary server to view a SnapView session. An active snapshot is a composite of a source LUN and reserved LUN data that lasts until you destroy the snapshot. You can create a snapshot before or after you start a session; however, the snapshot has no use until a secondary server activates it to a session.

If the storage system loses power while the session is running, and the session is not running in persistent mode, the session is lost and the snapshot becomes inactive. If the session is running in persistent mode, both the session and snapshot would survive the storage system power failure.

Unless you have additional software that supports same host access, you must assign the snapshot to a Storage Group other than the Storage Group that holds the source LUN(s). You also must assign multiple snapshots, of the same source LUN(s), to different Storage Groups. For information on software that supports same host access, refer to the *Prerequisites for Setting Up Snapshots* on page 2-12.

Multiple Snapshots

If you have an AX-Series storage system, you can create only one snapshot per source LUN(s). If you have a CX-Series or FC4700-Series storage system, you can create up to eight snapshots per source LUN(s). However, each snapshot must be activated to a different SnapView session and accessed by different servers. For example, if you create three snapshots and start eight sessions for a single source LUN, three servers can each activate a snapshot to a different sessions. Once the servers activate a snapshot to a session, the session is not available to another snapshot until it is deactivated.

You can also use the deactivate and activate functions to change the focus of a snapshot from one session to another. You must deactivate (unmap) a snapshot before you can activate (map) it to another session. For example, if you start eight sessions and create one snapshot for a single source LUN, a secondary server can not activate to another session until they deactivate from the current session. Once the secondary server deactivates the snapshot from the session, you can activate and deactivate between the remaining sessions. Refer to Chapter 3 to activate or deactivate the snapshot.

The eight-snapshot limit includes SnapView snapshots and any reserved snapshots used in another application such as SAN Copy and MirrorView / Asynchronous.

To Create a Snapshot

1. From any client that is managing the storage system, in the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Navigate to the icon for the source LUN(s) on which you want to create a snapshot.
3. Right-click on the source LUN icon and select **Create Snapshot**.

If the **Create Snapshot** option is unavailable, be sure that you have not exceeded the eight-snapshot limit per source LUN or the snapshot limit per storage system. For a complete list of snapshots the storage-system, right-click on the storage system icon and select **SnapView > SnapView Summary**. These limits include any reserved snapshots used for another application such as SAN Copy and MirrorView / Asynchronous.

4. Next to the **Storage System** and **Snapshot Source LUN** names, verify that you are creating the snapshot for the correct source LUN(s) on the correct storage system.
5. In **Snapshot Name**, enter a unique name for the snapshot. If the **Snapshot Name** is left blank, the default name *SCLUN 0xn* is assigned to the snapshot, where *n* is an incremented internal number that is assigned to the snapshot.
6. In **Storage Group**, select a Storage Group for the snapshot. **Storage Group** lists all Storage Groups that are attached to the storage system. The server(s) assigned to that Storage Group can access the snapshot.

Important If you have a VMware ESX Server, **do not select a Storage Group**. You *must* activate the snapshot before you add it to a Storage Group.

7. Click **OK** to create the snapshot.

The application creates the snapshot and, if the action is successful, does the following:

- Places an icon and description for the new snapshot in the **Snapshots** container associated with the snapshot's source LUN(s). The format for the snapshot description is *snapshot name[LUN X; servername; Inactive]* where **Inactive** means the snapshot is not active, that is not mapped to a session.

The *servername* does not appear in the snapshot description when viewed from the **Snapshot Names** node.

- Adds the text, **Snapshot Inactive**, to the description for the snapshot's source LUN(s). For example:
LUN 18 [0x12; RAID 1; servername; Snapshot Inactive]
where **Snapshot Inactive** means none of the snapshots on this source LUN(s) have been activated.

If an active snapshot exists on the source LUN(s), the text **Snapshot Active** appears next to the *servername*.

- For shared storage systems, assigns the snapshot to a Storage Group and adds the snapshot to the specified Storage Group.

What Next?

If you have not started a session, go to *Starting a SnapView Session* on page 2-17.

To start using your snapshot, go to *Activating a Snapshot* on page 3-23.

Adding a Snapshot to a Storage Group

Before accessing your snapshot data, the snapshot you map to a session must belong to a Storage Group and the Storage Group must connect to a server.

Unless you have additional software that supports same host access, you must assign the snapshot to a Storage Group other than the Storage Group that holds the source LUN(s). You also must assign multiple snapshots, of the same source LUN(s), to different Storage Groups. For information on software that supports same host access, refer to the *Prerequisites for Setting Up Snapshots* on page 2-12.

If the server that will have access to the snapshot already connects to a Storage Group, add the snapshot to that Storage Group. If you create a new Storage Group for the snapshot and then connect the server to the new Storage Group, the software removes the server from the original Storage Group and it will no longer have access to the LUNs in that Storage Group.

To Add a Snapshot to a Storage Group

Important If you have a VMware ESX Server, you *must* activate the snapshot before you add it to a Storage Group.

1. From any client that is managing the storage system, in the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Navigate to the snapshot you want to add to a Storage Group.
3. Right-click on the snapshot icon and select **Add to Storage Group**.
4. In the **Available Storage Groups** list, select the Storage Group for which you want to add the snapshot.

The Storage Group moves to **Selected Storage Groups**.

5. Click **OK** to add the snapshot to the Storage Group.

To connect the server to the Storage Group, refer to *Setting Up Access Logix*, in *EMC's Navisphere Manager Administrator's Guide*.

What Next? To start using the snapshot, see Chapter 3, *Using SnapView*.

To set up clones, see *Setting Up SnapView to Use Clones* on page 2-2.

This chapter gives a general overview of how to use clones and snapshots. For detailed information on these topics, please refer to the online help.

Major sections in this chapter are

Using Clones

- ◆ Clone States 3-2
- ◆ Fracturing a Clone 3-5
- ◆ Synchronizing a Fractured Clone 3-11
- ◆ Reverse Synchronizing a Fractured Clone 3-13
- ◆ Removing a Clone from a Clone Group 3-18
- ◆ Destroying a Clone Group 3-19
- ◆ Clone and Source LUN Trespasses 3-20
- ◆ Bad Blocks on Clones 3-21

Using Snapshots

- ◆ Activating a Snapshot 3-23
- ◆ Deactivating a Snapshot 3-25
- ◆ Rolling Back a SnapView Session 3-27
- ◆ Stopping a SnapView Session 3-33
- ◆ Destroying a Snapshot 3-34
- ◆ Snapshot/SnapView Session and Source LUN Trespasses 3-35
- ◆ Bad Blocks on Rollback 3-36

Using Clones

This section describes how to use clones. This section also describes how SnapView handles clone and source LUN trespasses and bad blocks of data.

After you have added a clone to a Clone Group, you can fracture it (see page 3-5) to make it available to another server, then you can do any of the following:

- ◆ Synchronize the fractured clone (see page 3-5)
- ◆ Reverse synchronize the fractured clone (see page 3-13)
- ◆ Remove the fractured clone from the Clone Group (see page 3-18)

What Next?

To learn about the possible states of a clone, continue to the next section. To start using the clone refer to one of the referenced pages listed above.

Clone States

Each clone you add to a Clone Group has its own state that indicates if it contains usable data. The possible clone states are: Reverse Synchronizing, Reverse Out-of-Sync, Synchronized, Synchronizing Out-of-Sync, or Consistent. Depending on the state of the clone, some operations may be unavailable (refer to Table 3-1 on page 3-3).

When you remove a clone from the Clone Group, it is no longer associated with its source LUN or Clone Group. It retains the copied data and becomes a conventional (regular) LUN.

Table 3-1, on page 3-3, lists when the clone is available for server I/O. The source LUN you specify when creating a Clone Group is available for server I/O during any clone state except for a Reverse Out-of-Sync state. Any server writes made to the source LUN during a reverse synchronization are copied to the clone. If you do not want incoming source writes copied to the clone during a reverse synchronization, you must select the **Protected Restore** feature in the **Add Clone** or **Clone Properties - Clone LUN tab** dialog box before issuing a reverse synchronization. However, before you can select the **Protected Restore** feature, it must be globally enabled by selecting the **Allow Protected Restore** option in the **Clone Features Properties** dialog box.

Table 3-1 Clone States

Clone State	Description	Cause of State	Permitted Operations	Clone Available for I/O
Consistent	A clone was in a Synchronized state and received incoming server writes via the source (if the clone is unfractured) or to the clone (if the clone is fractured). A clone in a Consistent state is usable but may not contain the most up-to-date information since writes made to the source may have not been copied to the clone.	<ul style="list-style-type: none"> • A clone is fractured while in a Consistent or Synchronized state. • A clone is unfractured and has yet to transition to a Synchronized state. 	<ul style="list-style-type: none"> • Fracture (only if clone is not already fractured) • Remove (only if the clone is fractured) • Synchronize (only if clone is fractured) • Reverse Synchronize (only if clone is fractured) 	Yes, if clone is fractured
Out-of-Sync	A clone was in the process of synchronizing but failed, so the clone is not a byte-for-byte copy of its source LUN and, therefore, is unusable.	A synchronization operation failed to complete successfully.	<ul style="list-style-type: none"> • Synchronize • Remove • Fracture (only if the clone was fractured by the system due to an error in the software or storage system; refer to the Event Log for cause of system fracture) 	No
Reverse Out-of-Sync	A clone was in the process of reverse synchronizing but failed, and therefore, the source is unusable and another reverse sync operation is recommended.	A reverse synchronization operation failed to complete successfully.	<ul style="list-style-type: none"> • Reverse synchronize • Remove • Fracture (only if the clone was fractured by the system due to an error in the software or storage system; refer to the Event Log for the cause of the system fracture) 	Yes

Table 3-1 Clone States (*continued*)

Clone State	Description	Cause of State	Permitted Operations	Clone Available for I/O
Reverse Synchronizing	A clone is in the process of copying its data to its source LUN.	<ul style="list-style-type: none"> Manually administered with Navisphere or the CLI from the production server. Automatically restarted following an SP failure <i>only</i> if a reverse synchronization was in progress during the SP failure and the recovery policy was set to Auto. <p>Note Before you can reverse synchronize a clone, you must fracture it. Also, the clone cannot be in an Out-of-Sync state, and no other clone in the Clone Group can be synchronizing or reverse synchronizing.</p>	Fracture	No
Synchronized	A clone is a byte-for-byte copy of its source.	<ul style="list-style-type: none"> After a synchronization, reverse synchronization, or a server write request, an unfractured clone automatically transitions to this state when the source LUN does not receive any server write requests during a specified period of time, which is defined by the Clone Group's quiesce threshold. A clone is added to a Clone Group without setting "Initial Sync Required". 	<ul style="list-style-type: none"> Fracture Remove 	No
Synchronizing	A clone is in the process of copying data from its source LUN.	<ul style="list-style-type: none"> Manually administered through Navisphere by the production server. Automatically started when you add a clone to a Clone Group with the Initial Sync Required property selected. Automatically restarted following an SP failure (if a synchronization was in progress during an SP failure and the recovery policy was set to Auto). <p>Note Before you can manually synchronize a clone, you must fracture the clone, and the clone cannot be in a Reverse Out-of-Sync state.</p>	Fracture	No

Fracturing a Clone

When you fracture a clone or a set of clones (consistent fracture), you split the clone(s) from its source LUN to make it available to a secondary server. A secondary server can access the fractured clone(s) if the clone belongs to a Storage Group that is connected to the secondary server. The secondary server can then use the clone for operations such as system backups, data modeling, or software application testing.

Unless you have additional software that supports same host access, you must assign the clone LUN to a Storage Group other than the Storage Group that holds the source LUN(s). You also must assign multiple fractured clones, of the same source LUN(s), to different Storage Groups. For information on software that supports same host access, refer to the *Prerequisites for Setting Up Clones* on page 2-2.

Consistent Fracture

A consistent fracture is when you fracture more than one clone at the same time in order to preserve the point-in-time restartable copy across the set of clones. The SnapView driver will delay any I/O requests to the source LUNs of the selected clones until the fracture has completed on all the clones (thus preserving the point-in-time restartable copy on the entire set of clones).

A restartable copy is a data state having dependent write consistency and where all internal database/application control information is consistent with a Database Management System/application image.

The clones you want to fracture must be within different Clone Groups. You cannot perform a consistent fracture on clones belonging to different storage systems.

After the consistent fracture completes, there is no group association between the clones. If there is a failure on any of the clones, the consistent fracture will fail on all of the clones. If any clones within the group were fractured prior to the failure, the software will re-synchronize those clones.

Consistent fracture is supported on CX-Series storage systems only. If you have a CX600 or CX700 storage system, you can fracture up to 16 clones at the same time. If you have another supported CX-Series storage system, you can only fracture up to 8 clones at the same time.

A maximum of 32 consistent fracture operations can be in progress simultaneously per storage system.

Types of Fractures

Two types of fractures can occur - administrative or system.

An *administrative fracture* will occur after any of the following events:

- ◆ The production server performed a fracture operation on the clone or set of clones.
- ◆ The secondary server performed a reverse synchronization operation on another clone in the Clone Group.
- ◆ The clone driver encountered a media failure and fractured the clone as a result of the failure.

A media failure can occur during a synchronization or reverse synchronization. Two types of media failures can occur: a read error or a write error. Once you correct the failure, you must re-issue the synchronization or reverse synchronization operation.

A *system fracture* will occur after any of the following events:

- ◆ The clone driver encountered an internal driver error and fractured the clone as a result.
- ◆ The source LUN trespasses to the peer SP, either manually or due to an SP failure.

Important If a clone is system fractured, you must manually fracture the clone because the software does not recognize the clone as being fractured.

When to Fracture a Clone LUN

You can fracture a clone in any state. However, to make a fractured clone available for server write requests, you must fracture the clone while it is in a *Synchronized* or *Consistent* state. Once you fracture the clone, it transitions into a *Consistent* state.

If you fracture a clone in a *Consistent* state, the source may be receiving writes from the server that have not been copied to the clone. If you want the clone to be a byte-for-byte copy of the source, you must ensure that no server writes are occurring when you fracture the clone.

By waiting until the clone is in a *Synchronized* state, you can have some confidence that server writes are no longer occurring. The *synchronized* state indicates that some number of seconds, defined by the *quiesce threshold* property you set when creating the Clone Group, have elapsed since the last server write. However, it does not

guarantee that another server write will not occur around the same time you fracture the clone.

If the clone never enters the Synchronized state, server writes are still continuing. If there is not active I/O, it may be difficult to determine the reason for those writes. However, it may be difficult to determine the reason for those writes. For example, they can be due to heartbeats in some clustered systems.

In some cases it is useful to fracture the clone in a Consistent state. For instance, you would fracture a clone in a Consistent state when

- ◆ you can determine that there are no writes being issued to the source LUN and you prefer not to wait for the quiesce threshold to elapse.
- ◆ you only need a crash consistent image (the image you would have if the server failed or shut down improperly).
- ◆ your application has a hot backup mode in which the application continues to write to the source, but internally understands how to disregard writes generated after your specified point-in-time.

Writes to the Fractured Clone LUN

When a fractured clone receives a server write request, the software marks the clone as dirty, which indicates that the clone has been modified and is no longer a byte-for-byte copy of its source LUN.

Writes to the Source LUN

After you fracture the clone, the software does not copy any server write requests made to the source LUN unless you manually perform a synchronization. Synchronizing a fractured clone unfractures the clone and updates the contents on the clone with its source LUN (see *Synchronizing a Fractured Clone* on page 3-11).

To Fracture a Clone

For additional information on the `admsnap` commands described below, refer to the *EMC SnapView Command Line Interfaces Reference*.

1. If you are fracturing a set of clones (consistent fracture), skip to step 2. If you are fracturing a single clone, you must do the following:
 - a. From the production server, stop I/O to the source LUN or LUNs.
 - b. From the production server, flush all cached data to the source LUN or LUNs by issuing the appropriate command for your operating system.
 - For a Windows server, use the **admsnap flush** command.
 - For Solaris, HP-UX, AIX, and Linux servers, unmount the file system by issuing the **umount** command. If you are unable to unmount the file system, issue the **admsnap flush** command. The **flush** command flushes all cached data.
 - For an IRIX server, the **admsnap flush** command is not supported. Unmount the file system by issuing the **umount** command. If you cannot unmount the file system, use the **sync** command to flush all cached data. The **sync** command reduces the number of times you need to issue the **fsck** command on the secondary server's file system. Refer to your system's man pages for **sync** command usage.
 - For a Novell NetWare server, use the **dismount** command on the volume to dismount the file system.

Neither the **flush** command nor the **sync** command is a substitute for unmounting the file system. Both commands only complement unmounting the file system.

With some operating systems, you may need to shut down the application to flush the data. Specific operating systems have different requirements.

2. Using Navisphere from any client that is managing the storage system, do the following:
 - a. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the **Clones** icon, and then to the **Clone Group Name** icon(s).
 - b. Right-click the **Clone Group Name** icon that contains the clone you want to fracture and select **Properties**.
 - c. Select the **Clone** tab for the clone you want to fracture and verify that its state is *Synchronized*. If its state is not *Synchronized*, you must wait until it is *Synchronized* before closing the dialog box and continuing to the next step.

The state of the clone may also be *Consistent*. Refer to *When to Fracture a Clone LUN* on page 3-6.

- d. If you are fracturing only one clone, continue to the next step. If you are performing a consistent fracture, repeat steps a through c to verify the state of the other clones.
- e. Double-click the **Clone Group Name** icon that contains the clone you want to fracture. If you are performing a consistent fracture, press and hold the **Ctrl** key and repeat this step for each clone you want to fracture.

If you are performing a consistent fracture, the clones you want to fracture must be within different Clone Groups.

- f. Right-click the clone you want to fracture, and select **Fracture**. If you are fracturing more than one clone, you can right-click on any of the clones you have selected.

A **Confirm: Fracture Clone** dialog box opens if you are fracturing one clone or a **Confirm: Consistent Fracture Clones** dialog box opens if you are fracturing more than one clone.

- g. Click **Yes** to confirm the fracture of the clone or clones, and then click **OK** upon succession.

Note for consistent fractures - If there is a failure on any of the clones, the consistent fracture will fail on all of the clones. If any clones within the group were fractured prior to the failure, the software will re-synchronize those clones.

The clone is now fractured from its source LUN. For information about the state of a clone see , *Clone States*, on page 3-2.

3. From the production server, resume I/O to the source LUN.
4. From the secondary server, activate the clone(s) by issuing the appropriate command for your operating system. In order for the secondary server to access the fractured clone(s), the clone(s) must belong to a Storage Group that is connected to the secondary server.

- On a Windows server, use the following admsnap command:
admsnap clone_activate

After a delay, the **admsnap clone_activate** command finishes rescanning the system and assigns drive letters to newly discovered clone devices.

Important: If the secondary server is running Windows NT and the clone was already mounted on a secondary server, a reboot is required after you activate the fractured clone. If the secondary server is running Windows 2000, a reboot is recommended but not required.

- On a UNIX server, use the same commands you used to discover new FLARE™ LUNs.
- On a NetWare server, run the **list devices** command or use the **scan all LUNs** command on the console.

Depending on your operating system, additional steps may be required from the secondary server in order to make the clone visible to the operating system.

5. If you have a VMware ESX Server, do the following:
 - a. Rescan the bus at the ESX Server level.
 - b. If a Virtual Machine (VM) is already running, power off the VM and use the Service Console of the ESX Server to assign the clone to the VM.

If a VM is *not* running, create a VM on the ESX Server and assign the clone to the VM.
 - c. Power on the VM and scan the bus at the VM level. For VMs running Windows, you can use the **admsnap activate** command to rescan the bus.

Synchronizing a Fractured Clone

The purpose of synchronizing a fractured clone is to update the contents on the clone with its source LUN. Once you synchronize a fractured clone, the clone becomes unfractured and behaves the same way as a newly added clone in a Clone Group.

When you initiate a synchronization, the software copies the data chunks that are different between the source LUN and the clone LUN to the clone. The software also copies any server writes from the source LUN to the clone LUN. The software does not copy any unmodified chunks to the clone.

If the same chunk of data is modified on the source LUN more than once, only the last modification is copied to the clone.

Before Starting a Synchronization

To synchronize a fractured clone, the following must be true:

- ◆ the clone *cannot* be in a Reverse-Out-of-Sync state.
- ◆ another clone in the Clone Group *cannot* be in a Reverse Synchronizing state.
- ◆ If a clone is system fractured, you must manually fracture the clone before you can synchronize it.

During a Synchronization

While the clone is synchronizing, you

- ◆ *cannot* remove the clone in a Synchronizing state.
- ◆ *cannot* perform a reverse synchronization with any other clone in the Clone Group.

A maximum of 20 concurrent synchronizations and/or reverse synchronizations can exist per SP. The software queues all other synchronization and reverse synchronization requests.

To Synchronize a Fractured Clone

You must explicitly follow the procedure for synchronizing a clone to avoid data loss. For additional information on the `admsnap` commands described below, refer to the *EMC SnapView Command Line Interfaces Reference*.

1. If the clone is mounted on a secondary server, deactivate the clone from the server it is mounted on by issuing the appropriate command for your operating system.
 - On a Windows server, use the following `admsnap` command:
`admsnap clone_deactivate -o clone drive letter`

Note For Windows 2000 Terminal Services Edition - if a terminal server is running on the secondary server, you must remove and disable it before using the `clone_deactivate` command. If you have the terminal server enabled, it prevents the `clone_deactivate` command from working properly.

- On a UNIX server, unmount the file system by issuing the **`umount`** command.
 - On a NetWare server, use the **`dismount`** command on the volume to dismount the file system.
2. Using Navisphere from any client that is managing the storage system, do the following:
 - a. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the **Clones** icon, and then double-click the **Clone Group Name** icon that contains the clone you want to synchronize.
 - b. Right-click the clone you want to synchronize, and then click **Synchronize. A Confirm: Synchronize Clone** dialog box opens.
 - c. Click **Yes** to confirm the synchronization of the clone, and then **OK**. The application does not wait for the synchronization to complete in order to display the synchronization successful message. The clone begins synchronization and is unfractured.

- d. Wait for the synchronization to complete. To verify this right-click the clone's **Clone Group Name** icon and select **Properties**. Then select the **Clone** tab for the clone you synchronized and verify that the clone condition is *Normal*. If it is not *Normal*, you must wait until it is before closing the dialog box and continuing to the next step.
3. To make the clone LUN available to the secondary server, you must fracture the clone as described in *To Fracture a Clone* on page 3-8.

Reverse Synchronizing a Fractured Clone

Reverse synchronizing a fractured clone replaces the data on the source LUN with the data on the clone. This allows you to revert to an earlier copy of the source LUN, for instance if the source became corrupted.

To ensure that there is no data corruption on the source LUN, you have to take the source LUN offline before you initiate the reverse synchronization. Once the operation begins, you can bring the source LUN back online.

When you initiate a reverse synchronization from Navisphere, the software immediately fractures all the clones in the Clone Group and then the software unfractures the clone from which you initiated the reverse synchronization. The software then copies the data chunks that differ between the source and the clone to the source LUN. The source LUN can instantly access the data, while the actual copying continues in the background.

Writes to the Source LUN

Any server writes made to the source LUN after the clone was originally fractured are overwritten with the data on the clone. The software does not copy any unmodified clone chunks to the source. After the reverse synchronization has completed, the clone that initiated the reverse synchronization remains unfractured.

If you modify the same data chunk on the clone more than once, the software copies only the last modification to the source LUN.

During a reverse synchronization, the software automatically copies any incoming server writes to the source LUN to the clone. If you do not want source writes copied to the clone during a reverse synchronization, you must check the **Use Protected Restore** feature in the **Add Clone** dialog box *before* initiating a reverse synchronization.

If you check the **Use Protected Restore** feature, after the reverse synchronization has completed, SnapView *fractures* the clone that initiated the reverse synchronization.

Protected Restore Feature

The **Protected Restore** feature protects the data on a clone during a reverse synchronization. When you select this feature, during a reverse synchronization the software will not copy any server writes made to the source LUN to the clone. Instead, the software records information in the clone private LUN to identify the source LUN writes for subsequent synchronizations.

Once you initiate a reverse synchronization, the software immediately unfractures the clone that initiated the reverse synchronization. Then the software fractures any other clone in the Clone Group in order to protect them from corruption should the reverse synchronization operation fail. The software then begins to copy its data to its source LUN. After the reverse synchronization has completed, the software fractures the clone that initiated the reverse synchronization.

You enable the Protected Restore feature on a per clone basis and not on a per Clone Group basis. You can select this feature when you first add a clone to a Clone Group (from the **Add Clone** dialog box) or at any time *before* you perform a reverse synchronization (from the **Clone Properties - Clone LUN** tab).

Before you can select the **Protected Restore** feature, you must globally enable it by selecting the **Allow Protected Restore** option in the **Clone Features Properties** dialog box. When you select this option, the SnapView driver automatically allocates 8 MB in additional memory per SP. The additional memory is fixed and is used to monitor modified blocks on the source LUN, in order to prevent these blocks from being overwritten by the clone during a reverse synchronization. This additional memory counts against the total memory budget for storage-system-based drivers.

Before Starting a Reverse Synchronization

To reverse synchronize a fractured clone, the following must be true:

- ◆ The clone is *not* in an Out-of-Sync state.
- ◆ Another clone in the Clone Group is *not* in a Synchronizing or Reverse Synchronizing state.

During a Reverse Synchronization

While the clone is reverse synchronizing, you *cannot*

- ◆ add a clone to the Clone Group
- ◆ remove the clone that is reverse synchronizing from the Clone Group
- ◆ synchronize any clone in the Clone Group
- ◆ reverse synchronize any clone in the Clone Group

To Reverse Synchronize a Fractured Clone

You must explicitly follow the procedure for reverse synchronizing a clone to avoid data loss. For additional information on the `admsnap` commands described below, refer to the *EMC SnapView Command Line Interfaces Reference*.

1. From the production server, stop I/O to the source LUN.
2. From the production server, deactivate the source LUN by issuing the appropriate command for your operating system.
 - On a Windows server, use the following `admsnap` command:
`admsnap clone_deactivate -o source-drive-letter`
 - On a UNIX server, unmount the file system by issuing the **`umount`** command. You must unmount the file system to avoid data corruption.
 - On a NetWare server, use the **`dismount`** command on the volume to dismount the file system.
3. From the secondary server, deactivate the clone LUN by issuing the appropriate command for your operating system.
 - On a Windows server, use the following `admsnap` command:
`admsnap clone_deactivate -o clone-drive-letter`
 - On a UNIX server, unmount the file system by issuing the **`umount`** command. You must unmount the file system to avoid data corruption.
 - On a NetWare server, use the **`dismount`** command on the volume to dismount the file system.

4. Using Navisphere from any client that is managing the storage system, do the following:

To use the Protected Restore feature, you must select it from the **Clone Properties - Clone LUN** tab before initiating a reverse synchronization.

- a. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the **Clones** icon, and then to the **Clone Group Name** icon that contains the clone you want to reverse synchronization.
- b. Double-click the **Clone Group Name** icon that contains the clone you want to reverse synchronize.
- c. Right-click the clone you want to reverse synchronize, and then click **Reverse Synchronize**. A **Confirm: Reverse Synchronize Clone** dialog box opens.

Important When the reverse synchronization begins, the software automatically fractures all the clones in the Clone Group.

- d. Click **Yes** to confirm the reverse synchronization of the clone, and then **OK** upon succession. The application does not wait for the reverse synchronization to complete in order to display the reverse synchronization successful message.

Depending on whether or not you enabled the **Protected Restore** feature, the following occurs to the clone that initiated the reverse synchronization:

- **With the Protected Restore feature** - the software fractures the clone after the reverse synchronization completes.
- **Without the Protected Restore feature** - the software leaves the clone unfractured.

If only minor differences exist between the clone and its source, the software may not have time to transition the state of the clone to reverse synchronization. This means that the state of the clone is still displayed as Consistent in Navisphere, even though the reverse synchronization was successful.

5. From the production server, activate the source LUN by issuing the appropriate command for your operating system.

- On a Windows server, use the following admsnap command:
admsnap clone_activate

After a delay, the **admsnap clone_activate** command finishes rescanning the system and assigns drive letters to newly discovered source LUN.

Important: If the production server is running Windows NT and the source LUN was already mounted on a production server, a reboot is required after you activate the LUN. If the secondary server is running Windows 2000, a reboot is recommended but not required.

- On a UNIX server, use the same commands you used to discover new FLARE™ LUNs.
- On a NetWare server, run the **list devices** command or use the **scan all LUNs** command on the console.

Depending on your operating system, additional steps may be required from the secondary server in order to make the LUN visible to the operating system.

6. From the production server, resume I/O to the source LUN.
7. Wait for the reverse synchronization to complete. To verify this right-click the clone's **Clone Group Name** icon and select **Properties**. Then select the **Clone** tab for the clone you reversed synchronized and verify that the clone condition is *Normal*. If it is not *Normal*, you must wait until it is before closing the dialog box and continuing to the next step.
8. If you did not enable the **Protected Restore** feature, you must fracture the clone LUN to make it available to the secondary server (see *To Fracture a Clone* on page 3-8).

Removing a Clone from a Clone Group

When you remove a clone from a Clone Group, the clone becomes a conventional (regular) LUN and is no longer associated with its source or Clone Group.

Before Removing a Clone from a Clone Group

To remove a clone from its Clone Group, the clone

- ◆ must be fractured or in a Synchronized or Consistent state.
- ◆ *cannot* be in a Synchronizing or Reverse Synchronizing state.
- ◆ *cannot* be in a queue to be synchronized.

In an effort to control resource consumption, the SnapView driver limits a maximum of four concurrent synchronizations. Once this limit is met, the software queues all other synchronization requests.

- ◆ *cannot* be removed if it is in a Reverse-Out-of-Sync state and it is the only clone in the Clone Group. If you remove a clone under these conditions, you cannot restore it. It is recommended that you perform a reverse synchronization operation from the clone and, when the reverse synchronization is successful, you should remove the last clone from the Clone Group.

To Remove a Clone from a Clone Group

1. From any client that is managing the storage system, in the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Navigate to the **Clones** icon, and then double-click the **Clone Group Name** icon that contains the clone you want to remove.
3. Right-click the clone you want to remove, and click **Remove**. A **Confirm: Remove Clone** dialog box opens.
4. Click **Yes** to confirm the removal of the clone and then **OK** when is successful.

The application removes the clone from its Clone Group. This clone is now a conventional LUN and it no longer counts against the clone and mirror limits.

Destroying a Clone Group

Destroying a Clone Group removes it permanently from the storage system and releases the source LUN specified when you created the Clone Group. Before you can destroy a Clone Group, you must remove all clones in the Clone Group.

To Destroy a Clone Group

1. From any client that is managing the storage system, in the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Navigate to the **Clones** icon, right-click the **Clone Group Name** icon that you want to destroy and, then click **Destroy Clone Group**.
3. Click **OK** to destroy the Clone Group, and then **Yes** to confirm.
A **Success: Destroy Clone Group** dialog box opens.
4. Click **OK**.

The application removes the Clone Group from the **Storage** tree and the source LUN specified when you created the Clone Group is now available for use in other operations.

Clone and Source LUN Trespasses

Trespasses can occur when you perform a manual trespass or when the software performs an automatic trespass.

Manual Trespass — You can manually trespass a fractured clone or a source LUN that is participating in a Clone Group to its peer SP.

Automatic Trespass — The software performs an automatic trespass due to one of the following reasons:

- ◆ If a clone belongs to the peer SP of its source LUN when you add the clone to the Clone Group.
- ◆ If the clone is fractured and you initiate a synchronization or reverse synchronization, the software trespasses the clone to the SP of its source LUN (if the clone and source LUN belong to different SPs).
- ◆ If an SP fails, the software trespasses the fractured clone and source LUNs to the peer (working) SP.

If a clone or source LUN is trespassed to the peer SP as a result of an SP failure, you must restore the failed SP before the clone or source LUN can be trespassed back to its original SP.

Effects of a Manual or Automatic Trespass — Whether the trespass is manual or automatic, the following effects occur after a clone or source LUN trespasses to the peer SP:

- ◆ When a fractured clone is trespassed, nothing in addition to the trespass occurs.
- ◆ When a source LUN that is participating in a Clone Group is trespassed to the peer SP, this SP:
 - Acquires the Clone Group and any unfractured clones,
 - Acquires the clone private LUN data for this source LUN, and
 - Assumes any SnapView operations, such as fractures, synchronizations, or reverse synchronizations.

Bad Blocks on Clones

This section describes what bad blocks are, how SnapView handles them, and what you can do to correct them.

Bad Blocks Overview

A bad block is an unreadable block on the LUN. The unreadable block is due to an incomplete write to the disk. Since there is an incomplete write to the disk, you cannot read the bad block on the LUN.

Although bad blocks are rare, they can occur anywhere on a LUN. If they occur in data or metadata, most operating systems will detect them and log errors, which causes applications to fail. If a bad block occurs in a file system's free space or in a database's free space, the server does not detect the bad block and it is essentially harmless.

Bad Blocks and Clones

If a bad block is encountered on the source LUN during a synchronization, SnapView generates a bad block at the same location on the clone LUN. If a bad block is encountered on the clone LUN during a reverse synchronization, SnapView generates a bad block at the same location on the source LUN. SnapView then proceeds with the synchronization or reverse synchronization operation. The LUNs are then a full copy of the source (if a synchronization was issued) or the clone (if a reverse synchronization was issued), which includes the original bad blocks.

SnapView generates a message in the event log to inform the user of the bad blocks on the LUN. After SnapView generates twenty messages, it stops logging bad blocks in the event log, but continues generating bad blocks on the clone LUN if one is encountered on the source LUN during a synchronization, or it continues generating bad blocks on the source LUN if one is encountered on the clone LUN during a reverse synchronization. If SnapView encounters more than 32,708 bad blocks, it aborts the synchronization or reverse synchronization operation.

If the bad blocks occur in a file system's free space or in a database's free space, SnapView detects them during a full synchronization (the initial synchronization) or reverse synchronization. Subsequent partial synchronizations encounter bad blocks only if they occur in a chunk that the server has written to.

Correcting Bad Blocks

Although bad blocks in a database's free space may be harmless, they can cause a synchronization or reverse synchronization operation to take longer than usual, in addition to generating excessive log messages.

You can correct a bad block by successfully writing to it. However, writing to it may be impossible if it is in free space. Instead, you can use a server-based utility to back up the data from the LUN with bad blocks. Then reformat or unbind/bind the LUN and restore the data from backup.

Using Snapshots

This section describes how to use and destroy snapshots. This section also describes how SnapView handles snapshot and source LUN(s) trespasses.

After you have started a SnapView session, you can do any of the following:

- ◆ Activate a snapshot (see the next section on this page)
- ◆ Deactivate a snapshot (see page 3-25)
- ◆ Recover source LUN data with the rollback feature (see page 3-27)
- ◆ Stop a session (see page 3-33)

Activating a Snapshot

The Navisphere Manager activate option maps the snapshot to a SnapView session. When you administer the activate option from Navisphere, you must reboot the secondary server, or use some other means, so that this server recognizes the new device created when you started the session. When you execute the **activate** command from the admsnap server utility, the command scans the secondary server's system buses for storage-system devices and determines if any device is part of a SnapView session. To use admsnap to activate a snapshot, refer to the *EMC SnapView Command Line Interfaces Reference*.

A secondary server can activate a snapshot to any SnapView session on the same source LUN as the snapshot. Once a secondary server activates a snapshot to a session, this server can write to the activated snapshot. The software stores all writes made to the snapshot in the reserved LUN pool. If the secondary server deactivates the snapshot from the session, the software destroys all writes made to the session.

You can create up to eight snapshots and activate (map) each snapshot to a single session provided that there is a different server for each snapshot. Only one snapshot at a time can activate a session.

The production and secondary servers must be running the same operating system (not a requirement for raw data access).

If the secondary server is running Windows NT, you must execute Disk Administrator the first time you activate a snapshot. Subsequent activations of this snapshot do not require this extra step.

To Activate a Snapshot

1. From the secondary server in the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Navigate to the snapshot you want to activate and select **Activate Snapshot**.
3. In **Available Sessions**, select the session name to which you want to map (activate) the snapshot.
4. Click **OK** to activate the snapshot to the session you selected.
If the action is successful, Navisphere closes the dialog box. Otherwise, Navisphere displays an error message and the dialog box remains open.
5. **If you do not have a VMware ESX Server** - Reboot the secondary server, or use some other means, such as the **admsnap activate** command, so that it recognizes the new device created when you started the session.

Depending on your operating system, additional steps may be required from the secondary server in order to make the snapshot visible to the operating system. For more information, see the product release notes.

If you have a VMware ESX Server - do the following:

- a. Add the snapshot to a Storage Group connected to the ESX Server that will access the snapshot, refer to page 2-27.
- b. Rescan the bus at the ESX Server level.
- c. If a Virtual Machine (VM) is already running, power off the VM and use the Service Console of the ESX Server to assign the snapshot to the VM.
If a VM is *not* running, create a VM on the ESX Server and assign the snapshot to the VM.
- d. Power on the VM. If the snapshot is already presented to the VM, rescan the bus at the VM level. For VMs running Windows, you can use the **admsnap activate** command to rescan the bus.

Deactivating a Snapshot

The deactivate function unmaps a snapshot from a SnapView session and destroys any secondary server writes made to the snapshot. The snapshot and session still exist but are not visible from the secondary server.

The secondary server must deactivate a snapshot before mapping it to another SnapView session. For example, if you start eight SnapView sessions on a single source LUN and create one snapshot for this same source LUN, a secondary server can activate (map) only one of the sessions at a time with the snapshot. If this secondary server wants to activate its snapshot to one of the other seven sessions, it must deactivate the snapshot and then activate it to another session.

To Deactivate a Snapshot

1. From the production server, flush all cached data to the source LUN(s) of the SnapView session by issuing the appropriate command for your operating system.
 - For a Windows server, use the **admsnap flush** command.

For a Windows 2000 server, after issuing the **admsnap flush** command, delete the drive letter.

 - For Solaris, HP-UX, AIX, and Linux servers, unmount the file system by issuing the **umount** command. If you are unable to unmount the file system, issue the **admsnap flush** command. The **flush** command flushes all cached data.
 - For an IRIX server, the **admsnap flush** command is not supported. Unmount the file system by issuing the **umount** command. If you cannot unmount the file system, use the **sync** command to flush cached data. The **sync** command reduces the number of times you need to issue the **fsck** command on the secondary server's file system. Refer to your system's man pages for **sync** command usage.
 - On a Novell NetWare server, use the **dismount** command on the volume to dismount the file system.

Neither the **flush** command nor the **sync** command is a substitute for unmounting the file system. Both commands only complement unmounting the file system.

With some operating systems, you may need to shut down the application to flush the data. Specific operating systems have different requirements. For more information, see the product release notes.

2. From the secondary server in the **Enterprise Storage** dialog box, click the **Storage** tab.
3. Navigate to the snapshot you want to deactivate and select **Deactivate Snapshot**.

The Navisphere Manager deactivate function does not flush all data on the secondary server. To flush I/O from this server, use the **admsnap deactivate** command on this server or any operating-system-specific commands to accomplish this task.

4. A message appears stating that you may need to flush I/O on the server operating system that is viewing the snapshot. If you click **Cancel**, no action is performed. If you click **Yes**, this deactivates the snapshot you selected from the SnapView session and destroys any writes made to the snapshot.

Depending on your operating system, additional steps may be required from the secondary server in order to flush all cached data on the secondary server. For more information, see the product release notes.

Rolling Back a SnapView Session

A rollback operation lets you restore the point-in-time copy of a SnapView session to the source LUN(s). You can roll back any persistent SnapView session. You select persistent mode only when you first start a session.

When you start a rollback operation, you have the option of starting a recovery session. A recovery session is just like any other SnapView session you start, except persistent mode is required. It is the default and you cannot be change the mode setting.

A recovery session allows you to undo the rollback operation because you can roll back the recovery session. The recovery session contains the point-in-time data of your source LUN(s) before you started your original rollback.

When you confirm the start of a rollback operation, the source LUN can instantly access the session's point-in-time data, while background copying continues until all the session's point-in-time blocks are copied to the source LUN. You cannot start a new session on the source LUN that is participating in the rollback.

The rollback operation itself does not count against the eight-session limit per source LUN. Starting a rollback recovery session will count as a single session against this limit.

Allocating Reserved LUN Pool Space

The rollback operation itself does not use reserved LUN pool space. However, the rollback may cause copy-on-first-write operations for sessions you started after the session that is rolling back. This in turn will consume more reserved LUN pool space. These copy-on-first-write operations occur because the original point-in-time data of those sessions is being modified by the rollback operation. Additionally, since you can resume I/O to the source LUN(s) or start another session while background copying continues, you may need additional reserved LUN space.

If you run out of reserved LUNs in the SP's LUN pool while the rollback operation is in progress (during the background copying), the software will terminate the session that is rolling back after it completes all background copying. The software will also terminate any other session that tries to allocate additional reserved LUN pool space.

Rolling Back with Snapshots

Each source LUN belonging to a session that is rolling back can have an activated or deactivated snapshot. If the session has an activated snapshot, SnapView copies any server writes made to this snapshot to the source LUN. However, you must unmount the snapshot *before* you start the rollback operation. When the rollback operation completes, including background copying, you can remount the snapshot.

If you deactivate the snapshot prior to starting the rollback operation, any server writes made to the snapshot are lost.

Before Starting a Rollback

Before starting a rollback on a SnapView session, verify the following:

- ❑ The session is running in persistent mode.
- ❑ The session is not in the process of being rolled back.
- ❑ The source LUN(s) is not participating in another rollback.
- ❑ If an activated snapshot is mapped to this session, dismount the snapshot.

You *cannot* roll back a SnapView session if the session's source LUN(s) is

- ◆ an unfractured clone. You must fracture the clone before rolling back the session.
- ◆ a fractured clone that is unavailable for I/O. For instance, if the clone was fractured while in an Out-of-Sync state.
- ◆ a Clone Group's source LUN that has any unfractured clone LUNs or clones that are in the process of reverse synchronizing. All clone LUNs for this source LUN must be fractured in a Synchronized or Consistent state before you can roll back a session to its source LUN.
- ◆ a primary image that has one or more unfractured secondary images. You must fracture the secondary image(s) before rolling back the session.
- ◆ a secondary mirror image. You must promote the secondary image and then fracture it before rolling back the session.

To Start a Rollback

Important for Windows operating systems only - To prevent data corruption during the rollback operation, you should disable the indexing service and recycle bin on the source LUN(s) of the session you will roll back.

1. From the production server, stop I/O to the source LUN(s).
2. From the production server, flush all cached data to the source LUN(s), by issuing the appropriate command for your operating system.

- For a Windows server, use the **admsnap flush** command.

For a Windows 2000 server, after issuing the **admsnap flush** command, delete the drive letter.

- For Solaris, HP-UX, AIX, and Linux servers, unmount the file system by issuing the **umount** command. If you are unable to unmount the file system, issue the **admsnap flush** command. The **flush** command flushes all cached data.
- For an IRIX server, the **admsnap flush** command is not supported. Unmount the file system by issuing the **umount** command. If you cannot unmount the file system, use the **sync** command to flush cached data. The **sync** command reduces the number of times you need to issue the **fsck** command on the secondary server's file system. Refer to your system's man pages for **sync** command usage.
- On a Novell NetWare server, use the **dismount** command on the volume to dismount the file system.

Neither the **flush** command nor the **sync** command is a substitute for unmounting the file system. Both commands only complement unmounting the file system.

With some operating systems, you may need to shut down the application to flush the data. Specific operating systems have different requirements.

3. If the session you want to roll back has an activated snapshot, and you want to keep any server writes made to this snapshot, you must unmount the snapshot by doing one of the following from the secondary server:

- For a Windows Server 2003, Windows 2000, or Windows NT operating system, use the **clone_deactivate** command.

Although the **clone_deactivate** command is prefixed with the term clone, you may issue this command on any Windows device that requires a volume to be taken offline, making it inaccessible to the operating system, and the removal of the assigned drive letter.

- On a UNIX system, use the **umount** command.
- On a Novell NetWare server, use the **dismount** command on the volume to dismount the file system.

If you do not want to keep any servers writes made to the snapshot, you must deactivate the snapshot now before continuing to the next step (see *Deactivating a Snapshot* on page 3-25).

4. Using Navisphere from any client that is managing the storage system, do the following:
 - a. In the **Enterprise Storage** dialog box, click the **Storage** tab.
 - b. Navigate to the **Sessions** icon, right-click the session you want to roll back and click **Start Rollback**.

A source LUN(s) can have only one session rolling back at a time. There is no limit to the number of concurrent rollback operations you can have on a storage system.

The **Start Rollback** dialog box opens.

If the selected session or another session on any of the same source LUNs are already rolling back, the software displays an error message. If you select multiple sessions, the **Start Rollback** option will be grayed out.

- c. In **Start Recovery Session**, select **Start Session**, if you want to start a recovery session before the rollback operation begins. A recovery session allows you to undo the rollback operation because you can roll back the recovery session. If selected, enter the name of the new session in **Session Name**.

A recovery session will always run in persistent mode. If you roll back a consistent session, the recovery session will run in consistent and persistent mode.

Starting a recovery session will count against the eight-session limit per source LUN. If you exceed the eight-session limit, which includes any reserved sessions, you will receive an error message when you click **OK** to start the rollback. If this occurs and you want to start a recovery session, you must stop another session on the source LUN. Be sure you do not stop the session you want to roll back.

- d. In **Rollback Rate**, select the rate at which the data will be copied back to the source LUN(s). Options are **High**, **Medium**, or **Low**. The default is **Medium**.
- e. Click **OK** to start the rollback.

Important: Once you start a rollback operation, you cannot stop it or the session that is being rolled back.

The software validates the session and, if no errors are reported, a **Confirm: Start Rollback** dialog box opens.

- f. To confirm the start of the rollback, click **OK**.

The rollback operation begins. The session's point-in-time data is instantly available on the source LUN(s), while background copying continues. To view the progress of the background copying, right-click the session that is rolling back and select **Session Properties** and then select the **Rollback** tab.

If you started a rollback recovery session, the software will start the session before the rollback operation (as long as you have not exceeded the eight-session limit per source LUN).

Important - If the source LUN is a clone, do not start a synchronization or reverse-synchronization on any clone in the Clone Group while the background copying is in progress. If the source LUN is a mirror, do not promote a secondary mirror to this source while the background copying is in progress.

5. From the production server, resume I/O to the source LUN(s).

Important: Be sure to verify that you have enough reserved LUNs in the SP's LUN pool before resuming I/O to the source LUN(s). Refer to *Allocating Reserved LUN Pool Space* on page 3-27.

The production server can resume I/O to the source LUN(s) while background copying continues. However, you cannot start another rollback operation for this source LUN(s) until the SnapView driver completes all background copying.

Important Server writes made to the source LUN(s) while the rollback is in progress will overwrite the data being rolled back.

6. If you unmounted an activated snapshot in step 3, do one of the following from the secondary server:
 - For a Windows Server 2003, Windows 2000, or Windows NT operating system, use the **clone_activate** command.

Although the **clone_activate** command is prefixed with the term clone, you may issue this command on any device that requires a scan for new LUNs and, for Windows only, for the assignment of a drive letter.

- On a UNIX and Novell NetWare system, use the **mount** command.

Stopping a SnapView Session

Stopping the last SnapView session of a source LUN frees the reserved LUN(s) used by the session and any SP memory used to maintain the session image. The newly freed reserved LUN(s) becomes available for other sessions.

If the snapshots participating in the session belong to one or more Storage Groups and you stop the session, the servers connected to the Storage Groups will no longer have access to the snapshots in those Storage Groups. EMC recommends that you do one of the following:

- ◆ Flush the I/O on the server activated to the snapshot before stopping the session.
- ◆ Deactivate the snapshot before stopping the session.

Stopping a session with an active snapshot makes the snapshot appear inactive.

To Stop a SnapView Session

1. From any client that is managing the storage system, in the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Navigate to the **Sessions** icon to select the SnapView session you want to stop.
3. Right-click the session and select **Stop Session**.
4. Click **Yes** to stop the session.

The application removes the snapshot icon from the **Session** container in the **Storage** tree.

Destroying a Snapshot

When you destroy a snapshot, the following is true:

- ◆ If the snapshot is inactive, the software destroys only the selected snapshot.
- ◆ If the snapshot is active, a warning message appears indicating that you should deactivate the snapshot before destroying it. If you accept the warning message, the software deactivates the snapshot, and destroys it (the snapshot) and any server writes made to the snapshot.

The Navisphere Manager deactivate function does not flush all cached data on the secondary server. To flush I/O from this server, do not accept the warning message; use the **admsnap deactivate** utility command on this server or any operating-system-specific commands to accomplish this task.

- ◆ If the snapshot belongs to a Storage Group(s), an error message appears indicating that you cannot destroy a snapshot that is in a Storage Group. Remove the snapshot from the Storage Group(s), and then destroy the snapshot.

To Destroy a Snapshot

1. From any client that is managing the storage system, in the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Navigate to the snapshot you want to destroy.
3. Right-click on the snapshot and select **Destroy Snapshot**.
4. In the confirmation dialog box, click **Yes** to destroy the snapshot.

The application removes the snapshot icon from the **Snapshots** container in the **Storage** tree.

Snapshot/SnapView Session and Source LUN Trespasses

Trespasses can occur manually or automatically on any SnapView session that is running in persistent mode.

Important If your session is not running in persistent mode, it will not trespass to the peer SP. The software destroys your session and deactivates any activated snapshots.

Manual Trespass — You can manually trespass a source LUN that has a session to its peer SP.

Automatic Trespass — The software performs an automatic trespass due to one of the following reasons:

- ◆ SP reboot or failure
- ◆ Array reboot or power failure
- ◆ Server I/O trespassing to the peer SP

Once the failed SP is running, the production server must issue a restore command in order to restore the proper source LUNs, sessions, snapshots, and reserved LUNs back to their original SP (for the appropriate restore command, refer to the documentation that shipped with your failover software).

Effects of a Manual or Automatic Trespass — When a source LUN(s) is trespassed to the peer SP, this SP:

- ◆ Resumes rolling back any LUNs that were in the middle of a rollback at the time of the trespass.
- ◆ Acquires the reserved LUNs that are assigned to that source LUN.
- ◆ Acquires any persistent sessions and any snapshots associated with these sessions. These snapshots can be activated or deactivated. If the snapshot is activated, the production server or secondary server can issue a restore command in order to restore the snapshot (once the original SP is up and running). If the snapshot is deactivated, only the production server can issue the restore command.
- ◆ Assumes any SnapView operations, such as starting a session and allocating additional reserved LUNs.

Important For information on how SPs manage the reserved LUN pool, refer to the latest revision of the *EMC Navisphere Manager Administrator's Guide*.

Bad Blocks on Rollback

This section describes what bad blocks are, how SnapView handles them, and what you can do to correct them.

Bad Blocks Overview

A bad block is an unreadable block on the LUN. The unreadable block is due to an incomplete write to the disk. Since there is an incomplete write to the disk, you cannot read the bad block on the LUN.

Although bad blocks are rare, they can occur anywhere on a LUN. If they occur in data or metadata, most operating systems will detect them and log errors, which causes applications to fail. If a bad block occurs in a file system's free space or in a database's free space, the server does not detect the bad block and it is essentially harmless.

Bad Blocks and Rollback

If the software encounters a bad block on a reserved LUN during a rollback operation, SnapView generates a bad block on the source LUN. SnapView generates the bad block at the same location the block was supposed to be rolled back.

SnapView generates a message in the event log to inform you of the bad blocks on the source LUN. After twenty messages, SnapView stops logging bad blocks in the event log, but it continues generating bad blocks on the source LUN until the rollback operation is complete.

Correcting Bad Blocks

Although bad blocks in a database's free space may be harmless, they can cause a rollback operation to take longer than usual, in addition to generating excessive log messages.

You can correct a bad block by successfully writing to it. However, writing to it may be impossible if it is in free space. Instead, you can use a server-based utility to back up the data from the LUN with bad blocks. Then reformat or unbind/bind the LUN and restore the data from backup.

This chapter gives a general overview of how to display the **Properties** dialog box for each SnapView component. For detailed information on these topics, please refer to the on-line help.

Major sections are

Clone Properties

- ◆ Displaying and Modifying Clone Properties 4-2
- ◆ Clone Properties 4-2
- ◆ Clone Feature Properties 4-3
- ◆ Clone Group Properties 4-3
- ◆ Source LUN Properties 4-4

Snapshot Properties

- ◆ Displaying and Modifying Snapshot Properties 4-5
- ◆ Snapshot Name Properties 4-5
- ◆ SnapView Session Properties 4-6
- ◆ Displaying Status of All Snapshots and SnapView Sessions 4-7

Displaying and Modifying Clone Properties

Each Clone Group, source LUN, and clone has a **Properties** dialog box associated with it that provides a variety of information about the component. This section describes how to open the **Properties** dialog box for each component and view the component's properties.

Clone Properties

To display the properties of the clone, do the following:

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the Clone Group icon for which you want to display the clone properties.
2. Right-click the **Clone Group** icon and select **Properties**.
3. In the **Clone Properties - Clone** tab dialog box, the following information is available:

- Storage system (readable *only*)
- Clone ID (readable *only*)
- State (readable *only*)
- Available for I/O (readable *only*)
- Is fractured (readable *only*)
- Condition (readable *only*)

If you fractured a clone or a group of clones, the clone condition will display as Administratively Fractured.

- Is Dirty (readable *only*)

If this option displays **yes**, the clone received a server write request, which means it is not a byte-for-byte duplicate of its source LUN.

- Synchronized (readable *only*)
- Clone LUN name, ID, capacity, drive type (readable *only*)
- Recovery policy (readable and writable)
- Synchronization rate (readable and writable)
- Protected restore (readable and writable)

You can modify the Protected Restore feature only if the **Allow Protected Restore** option is globally enabled from the **Clone Features Properties** dialog box. If the **Allow Protected Restore** option is not enabled, the **Protected Restore** feature is grayed out and unavailable.

4. Click **Help** for information on the properties in the dialog box, click **OK** to accept the changes and close the dialog box or click **Close** to close the dialog box.

Clone Feature Properties

To display the Clone Feature Properties, which contain clone private LUN information, do the following:

1. From the **Storage** tab of the **Enterprise Storage** dialog box, right-click the storage system for which you want to display the Clone Feature Properties and select **SnapView > Clone Feature Properties**.
2. In the **Clone Feature Properties** dialog box, the following information is available:
 - Clone private LUNs (readable and writable)
 - Available LUNs, including source LUN, ID, capacity, and drive type.
 - Clone private LUNs, including clone private LUN, capacity, and drive type.
 - Allow Protected Restore (readable and writable)
3. Click **Help** for information on the properties in the dialog box. Click **OK** to accept the changes and close the dialog box or click **Cancel** to close the dialog box without applying any changes.

Clone Group Properties

To display the properties of the Clone Group, do the following:

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the Clone Group icon for which you want to display the properties.
2. Right-click the **Clone Group** icon and select **Properties**.
3. In the **Clone Properties - General** tab dialog box, the following information is available:
 - Storage system (readable *only*)
 - Clone Group name (readable and writable)
 - Unique ID (readable *only*)
 - Description (readable and writable)
 - Quiesce threshold (readable and writable)
4. Click **Help** for information on the properties in the dialog box. Click **OK** to accept the changes and close the dialog box or click **Cancel** to close the dialog box without applying any changes.

Source LUN Properties

To display the properties of the source LUN of a clone, do the following:

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the Clone Group icon for which you want to display the properties.
2. Right-click the **Clone Group** icon and select **Properties**.
3. In the **Clone Properties - Source** tab dialog box, the following information is available:
 - Storage system (readable *only*)
 - Clone ID (readable *only*)
 - Available for I/O (readable *only*)
 - Source LUN name, ID, capacity, and drive type (readable *only*)
4. Click **Help** for information on the properties in the dialog box, or click **Close** to close the dialog box.

Displaying and Modifying Snapshot Properties

Each snapshot and SnapView session has a **Properties** dialog box associated with it that provides a variety of information about the component. This section describes how to open the **Properties** dialog box for each component and view or modify the component's properties.

To view the properties of the reserved LUN pool, refer to the latest revision of the *EMC Navisphere Manager Administrator's Guide*.

Snapshot Name Properties

To display or modify the properties of the snapshot, do the following:

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the **snapshot name** icon for which you want to display the properties.
2. Right-click the **snapshot name** icon and select **Properties**.
3. In the **Snapshot Properties** dialog box, you can view the following information:
 - Snapshot name (readable for AX and CX storage systems and writable for CX storage systems only)
 - Storage system in which the snapshot resides (readable *only*)
 - Snapshot source LUN (readable *only*)
 - WorldWide Unique ID (readable *only*)
 - Session name (readable *only*)
 - Snapshot state (readable *only*)
 - Storage Group (readable *only*)
 - Server (host) name (readable *only*)
4. Click **Help** for information on the properties in the dialog box, or click **Close** to close the dialog box.

SnapView Session Properties

To display the properties of a SnapView session, do the following:

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the **Session** icon for which you want to display the properties.
2. Right-click the session for which you want to display properties and, then select **Properties**.
3. In the **SnapView Session Properties** dialog box, you can view the following information:
 - Creation time (readable *only*)
 - Optional modes selected (readable *only*)
 - Snapshot source LUN used (readable *only*)
 - Snapshot LUN name assigned to the session (readable *only*)
 - Snapshot unique ID (readable *only*)
 - Statistics, such as total reads in a session, reads from the reserved LUN pool, reads from the snapshot source LUN, total writes in a session, writes from the reserved LUN pool, writes from the snapshot source LUN, and writes larger than the reserved LUN pool entry size. (readable *only*)
 - Rollback information, such as the rollback rate and progress. This information is available only when there is a rollback operation in progress. (readable and writable)
4. Click **Help** for information on the properties in the dialog box, or click **Close** to close the dialog box.

Displaying Status of All Snapshots and SnapView Sessions

You can view the status of all snapshots on all managed storage systems as well as the status of all SnapView sessions. To display the **SnapView Summary** dialog box, do the following:

All SnapView Summary properties are read *only*.

1. On the **Tools** menu in the **Main** Navisphere window, click **SnapView Summary**.
2. In the **SnapView Summary** dialog box, you can view the following:
 - From the **Source LUNs** tab
 - Snapshot source LUN
 - Number of snapshots
 - Number of sessions
 - Storage system
 - From the **Snapshots** tab
 - Snapshot
 - Status
 - Snapshot source LUN
 - Storage system
 - From the **Sessions** tab
 - Sessions
 - Snapshot Source LUN
 - Snapshot
 - Modes
 - Storage system
3. Click **Help** for information on the properties in the dialog box, or click **Close** to close the dialog box.

SnapView Examples

This chapter contains examples, from setting up clones and snapshots to using them. Each example also contains an illustrated overview that shows the main steps outlined in the examples.

Major sections are

- ◆ Clones Example 5-2
- ◆ Snapshots Example 5-10
- ◆ Snapshots Example with Rollback 5-16

Clones Example

This section provides an example of how to set up and use clones.

The server names, files, and applications used in this section are intended for example purposes only.

Summary

In this example, you are creating two clones to perform software testing on a database file and its log file. Once you have completed testing, you decide that you want to keep the modified data and replace the data on the source LUNs with this modified data. To do this, you start a reverse synchronization on the clone LUNs. The reverse synchronization will replace the contents of the source LUN with the contents of the clone LUN.

Hardware and Software Configuration

The server in this example is a typical Microsoft SQL Server 2000. Your environment has the following configurations.

Servers ^a	Storage Groups	LUN Names
Production ph12345	Source NWDataLog	Source NWData storing Northwind.mdf and NWLog storing Northwind.ldf
Secondary sh12345	Clone NWDataLogTest	Clone NWDataClone storing Northwind.mdf and NWLogClone storing Northwind.ldf

- a. Servers are used for example purposes only. Any client that is managing a connected storage system can perform most of these functions but not all. For information on which server can perform specific functions, refer to the reference sections listed in the operations overview.

Clone Groups	LUNs in Clone Group
NWDataCG	NWData and NWDataClone
NWLogCG	NWLog and NWLogClone

Operations Overview

Important The following table does not provide detailed steps for each task. It is important that you refer to the *Prerequisites for Setting Up Clones* on page 2-2 and to the reference sections listed below before completing any tasks.

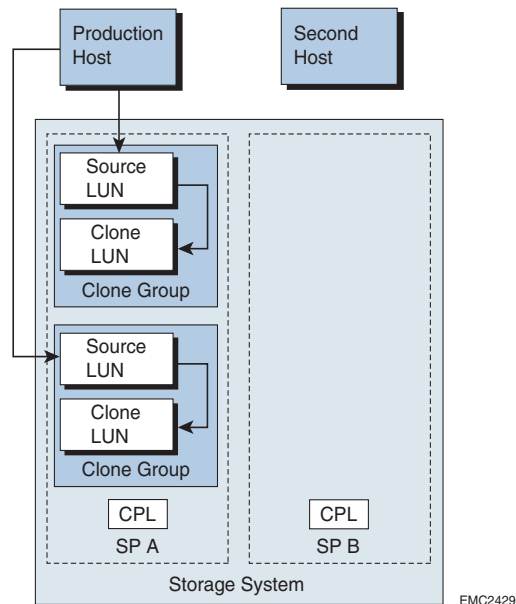
Task	Task Description	Reference Section
1. Setting up LUNs to be used as clones	<ul style="list-style-type: none"> <input type="checkbox"/> From server ph12345 (production server), create LUNs. You will need two LUNs that will become your clone LUNs (NWDataClone and NWLogClone). These LUNs must be the same sizes as source LUN NWData (Northwind.mdf) and source LUN NWLog (Northwind.ldf), but they can be different RAID types. <input type="checkbox"/> From server ph12345 (production server), assign the newly created LUNs (NWDataClone and NWLogClone) to Storage Group NW_DataLogTest and connect this Storage Group to server sh12345 (secondary server). 	<i>Prerequisites for Setting Up Clones</i> on page 2-2
2. Allocating Clone Private LUNs	<ul style="list-style-type: none"> <input type="checkbox"/> Create two LUNs that are at least 250000 blocks. These LUN will be used as Clone Private LUNs. <input type="checkbox"/> From server ph12345 (production server), allocate the two LUNs you just created as Clone Private LUNs. 	<i>Allocating Clone Private LUNs</i> on page 2-4
3. Creating a Clone Group	<ul style="list-style-type: none"> <input type="checkbox"/> From server ph12345 (production server), select source LUN NWData and create a Clone Group called NWDataCG. <input type="checkbox"/> From server ph12345 (production server), select source LUN NWLog and create a Clone Group called NWLogCG. 	<i>Creating a Clone Group</i> on page 2-7
4. Adding a clone to the Clone Group	<ul style="list-style-type: none"> <input type="checkbox"/> From server ph12345 (production server), add LUN NWDataClone to the NWDataCG Clone Group and LUN NWLogClone to the NWLogCG Clone Group. <p>Note Select Initial Sync Required for both clone LUNs.</p>	<i>Adding a Clone to a Clone Group</i> on page 2-9

Task	Task Description	Reference Section
5. Fracturing the clone	<input type="checkbox"/> From server ph12345 (production server), verify that LUN NWDataClone and LUN NWLogClone are in a Synchronized or Consistent state. <input type="checkbox"/> From server ph12345 (production server), fracture the LUNs NWDataClone and NWLogClone (clone LUNs).	<i>Clone Properties</i> on page 4-2 <i>Fracturing a Clone</i> on page 3-5
6. Activating clones	<input type="checkbox"/> From the sh12345 server (secondary server), activate the clone LUNs (NWDataClone and NWLogClone).	<i>Fracturing a Clone</i> on page 3-5 (step 4)
7. Trespassing clones	<input type="checkbox"/> From server ph12345 (production server), trespass LUNs NWDataClone and NWLogClone (clone LUNs) to the peer SP.	Clone and Source LUN Trespasses3-20
8. Starting software testing	<input type="checkbox"/> From server sh12345 (secondary server), start software testing on LUNs NWDataClone and NWLogClone (clone LUNs). Note If you want to go back to an earlier copy of either source LUN NWData or NWLog , you can synchronize clone LUNs NWDataClone or NWLogClone , but you must fracture them after the synchronization is completed and before you start testing again.	None <i>Synchronizing a Fractured Clone</i> on page 3-11
9. Stopping software testing	<input type="checkbox"/> From server sh12345 (secondary server), stop software testing.	None
10. Reverse Synchronizing Fractured clones	<input type="checkbox"/> From server ph12345 (production server), reverse synchronize LUNs NWDataClone and NWLogClone (clone LUNs).	<i>Reverse Synchronizing a Fractured Clone</i> on page 3-13
11. Removing clones from Clone Group	<input type="checkbox"/> From server ph12345 (production server), remove LUN NWDataClone from the NWDataCG Clone Group and remove LUN NWLogClone from the NWLogCG Clone Group.	<i>Removing a Clone from a Clone Group</i> on page 3-18
12. Destroying a Clone Group	<input type="checkbox"/> From server ph12345 (production server), destroy the NWDataCG and NWLogCG Clone Groups.	<i>Destroying a Clone Group</i> on page 3-19

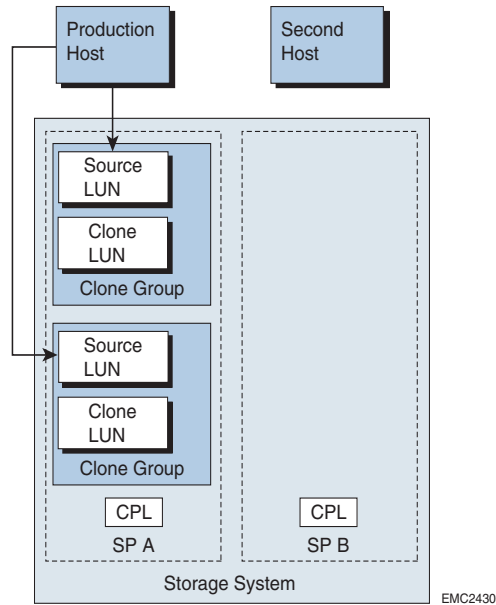
Illustrated Overview

The following section provides an illustrated description of the *main* operations described in the table in the previous section.

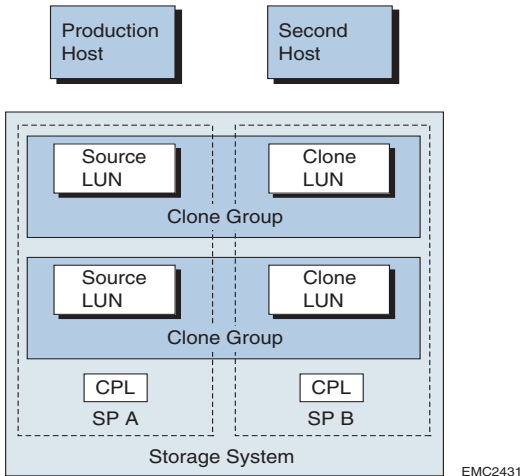
1. Production server adds clone LUNs to Clone Groups and initial synchronization begins (the contents of the source LUN are copied to the clone LUN). I/O to the source LUNs from the production server continues.



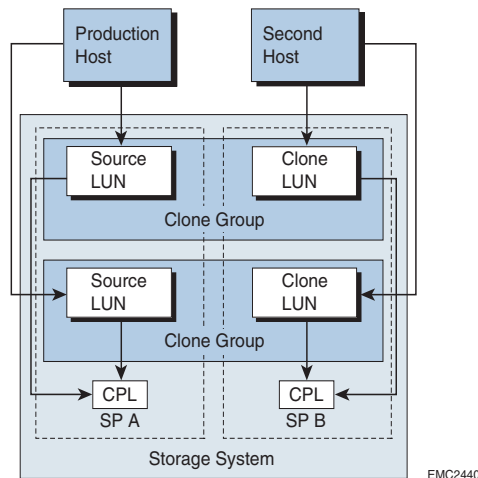
- Initial synchronization completes. I/O to the source LUNs from the production server continues.



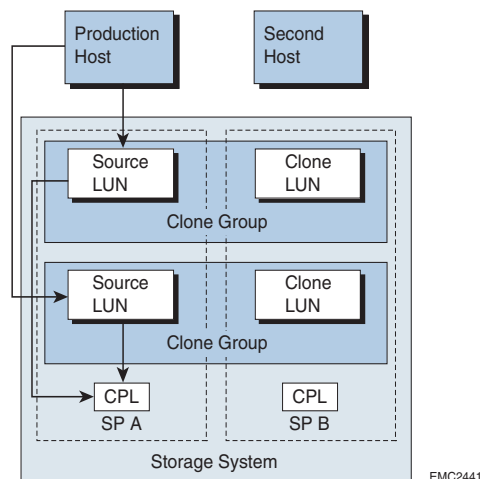
- I/O stops to the source LUNs from the production server. The production server then fractures and respresents the clones to SP B.



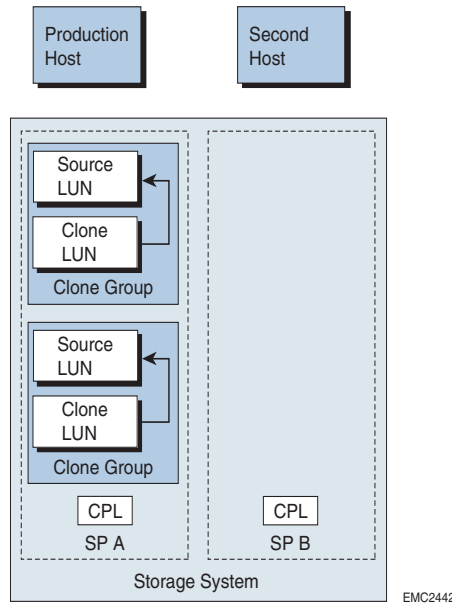
4. I/O resumes to the source LUNs from the production server. The secondary server activates the clone LUNs. I/O to the clone LUNs from the secondary server begins and software testing starts. As I/O modifies the fractured clones and source LUNs, the clone private LUNs record information that identifies these modified data chunks but no actual data is written to the clone private LUNs.



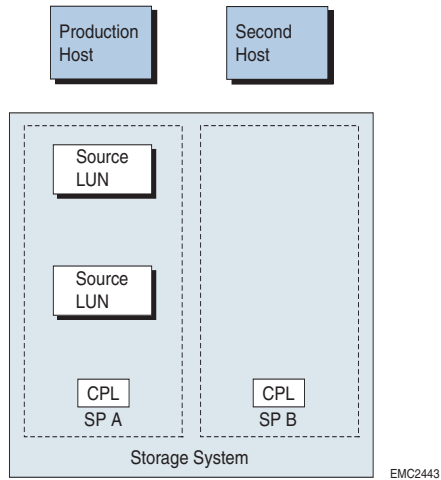
5. Software testing stops and I/O to the clone LUNs from the secondary server also stops.



- I/O stops to the source LUNs from the production server. The production server then initiates a reverse synchronization (without the Protected Restore feature enabled) to replace the contents of the source LUNs with the contents of the clone LUNs. The reverse synchronization causes the clone LUNs to trespass back to SP A.



- The production server removes the synchronized clones from the Clone Groups and destroys the Clone Groups.



Snapshots Example

This section provides an example of how to set up and use snapshots.

The server names, files, and applications used in this section are intended for example purposes only.

Summary

In this example, you are starting two SnapView sessions and creating two snapshots of a database file and its log file. You will then back up the two snapshots onto tape.

Hardware and Software Configuration

Your environment has the following configurations:

Hardware	Servers ^a	Storage Groups	LUN Names	SnapView Sessions
Storage System CX600	Production ph12345	Source NWDataLog	Source NWData storing Northwind.mdf and NWLog storing Northwind.ldf	NWDataSession (started on LUN NWData and activated to NWDataSnap)
Server Typical MS SQL Server 2000	Secondary sh12345	Snapshot NWDataLog_backup	Snapshot NWDataSnap storing Northwind.mdf and NWLogSnap storing Northwind.ldf	NWLogSession (started on LUN NWLog and activated to NWLogSnap)

- a. These servers are used for example purposes only. Any client that is managing a connected storage system can perform most of these functions but not all. For information on which server can perform specific functions, refer to the reference sections listed in the operations overview.

Operations Overview

Important The following table does not provide detailed steps for each task. It is important that you refer to the *Prerequisites for Setting Up Snapshots* on page 2-12 and to the reference sections listed below before completing any tasks.

Task	Task Description	Reference Section/Document
1. Configure the reserved LUN pool	<ul style="list-style-type: none"> <input type="checkbox"/> On each SP, determine the size of the reserved LUN pool. <input type="checkbox"/> On the storage system, bind one or more LUNs on each SP to the size you determined for the reserved LUN pool. <input type="checkbox"/> From the ph13245 server (production server), allocate the reserved LUNs to the SP's LUN pool. 	The latest revision of the <i>EMC Navisphere Manager Administrator's Guide</i> .
2. Start a SnapView session	<ul style="list-style-type: none"> <input type="checkbox"/> From the ph13245 server (production server), start two SnapView sessions (NWDataSession and NWLogSession). 	<i>Starting a SnapView Session</i> on page 2-17
3. Create a snapshot	<ul style="list-style-type: none"> <input type="checkbox"/> From the ph13245 server (production server), create two snapshots (NWDataSnap and NWLogSnap). 	<i>Creating a Snapshot</i> on page 2-24
4. Add the snapshot to a Storage Group	<ul style="list-style-type: none"> <input type="checkbox"/> From the ph13245 server (production server), add snapshot NWDataSnap and snapshot NWLogSnap to Storage Group NWDataLog_backup. 	<i>Adding a Snapshot to a Storage Group</i> on page 2-27
5. Activate the snapshot	<ul style="list-style-type: none"> <input type="checkbox"/> From the sh12345 server (secondary server), activate the NWDataSnap snapshot to the NWDataSession sessions. <input type="checkbox"/> From the sh12345 server (secondary server), activate the NWLogSnap snapshot to the NWLogSession sessions. 	<i>Activating a Snapshot</i> on page 3-23

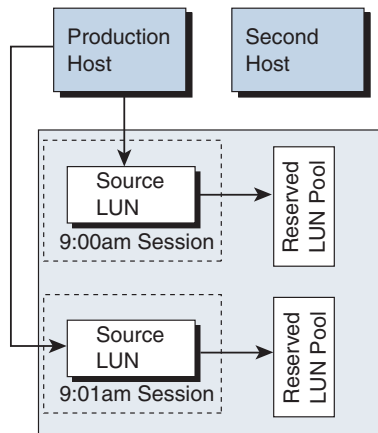
Task	Task Description	Reference Section/Document
6. Identify the snapshot	<input type="checkbox"/> From the sh12345 server (secondary server), identify the snapshots (NWDataSnap and NWLogSnap) to the operating system. This procedure depends on the operating system; for example, on a Windows server you need to run Disk Admin . Note This step is needed only once, as part of the SnapView initial setup.	None
7. Start backup application	<input type="checkbox"/> From the sh12345 server (secondary server), start the backup application for the snapshots (NWDataSnap and NWLogSnap). Note While the snapshot is backing up to tape, there is a performance impact on the source LUNs (NWData and NWLog). <input type="checkbox"/> Verify that the backup has completed before continuing to the next step.	None
8. Deactivate the snapshot	<input type="checkbox"/> From the sh12345 server (secondary server), deactivate the snapshots (NWDataSnap and NWLogSnap).	<i>Deactivating a Snapshot</i> on page 3-25
9. Stop the SnapView session	<input type="checkbox"/> From the ph12345 server (production server), stop the SnapView sessions (NWDataSession and NWLogSession).	<i>Stopping a SnapView Session</i> on page 3-33

Illustrated Overview

The following section provides an illustrated description of the *main* operations described in the table in the previous section.

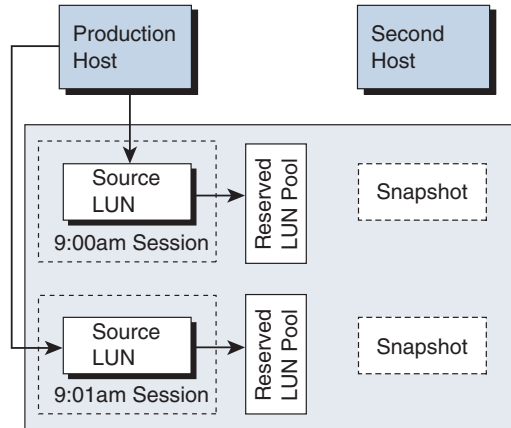
I/O to the source LUNs from the production server continues while backing up the snapshots.

1. Production server starts the SnapView sessions. Any modifications made to the source LUNs from the production server are written to the source LUNs. The software copies the original data that was just modified and stores it in the reserved LUN pool.



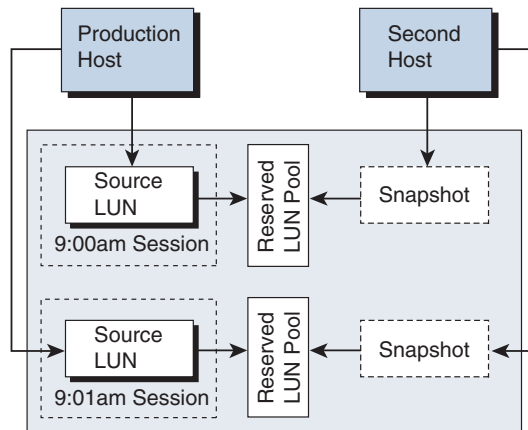
EMC2756

- Production server creates the snapshots.



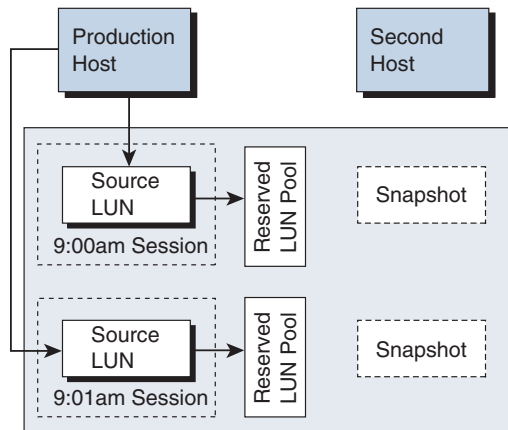
EMC2757

- The secondary server identifies and activates the snapshots, and then begins the backup application. Since the secondary server is reading the snapshot to perform the backup application, modifications can occur to the snapshot and would be stored in the reserved LUN pool.



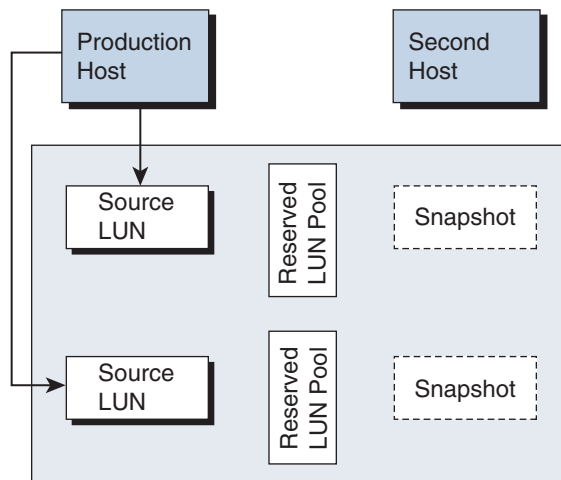
EMC2758

4. Backup application completes and secondary server deactivates the snapshots.



EMC2757

5. Production server stops the SnapView sessions. Any additional modifications made to the source LUNs from the production server (after the server stopped the sessions) are written to the source LUNs. The software no longer copies the original data to the reserved LUN pool.



EMC2759

Snapshots Example with Rollback

This section provides an example of how to set up and use snapshots and how to roll back a SnapView session.

The server names, files, and applications used in this section are intended for example purposes only.

Summary

In this example, you have an application that runs on an Oracle database. Every Monday through Friday you start a single session of both databases. On Friday, you realize that the database is corrupted or contains changes that you do not want, so you go back and test your daily sessions. You discover that Thursday and Wednesday's sessions are bad but Tuesday's session contains the data you want. You then roll back Tuesday's session to the source LUN, which will bring you back to Tuesday's data.

Hardware and Software Configuration

Your environment has the following configurations:

Hardware	Servers ^a	Storage Groups	LUN Names	SnapView Sessions
Storage System CX400	Production ph12345	Source OracleDB	Source OracleData storing employeefiles.mdf and OracleLog storing employeefiles.ldb	OracleSession1 OracleSession2 OracleSession3 OracleSession4 OracleSession5
Server Typical Solaris Oracle Server	Secondary sh12345	Snapshot OracleDB_backup	Snapshot OracleDataSnap storing employeefiles.mdf and OracleLogSnap storing employeefiles.ldb	

- a. These servers are used for example purposes only. Any client that is managing a connected storage system can perform most of these functions but not all. For information on which server can perform specific functions, refer to the reference sections listed in the operations overview.

Operations Overview

Important The following table does not provide detailed steps for each task. It is important that you refer to the *Prerequisites for Setting Up Snapshots* on page 2-12 and to the reference sections listed below before completing any tasks.

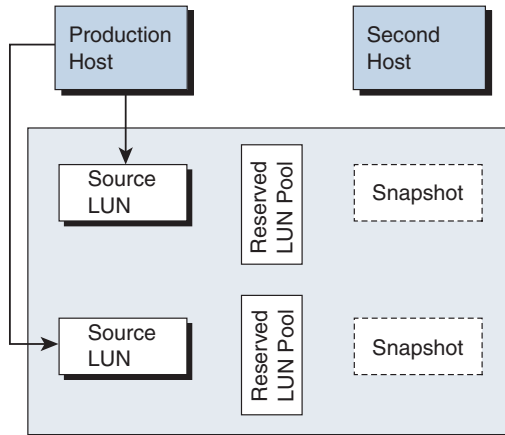
Task	Task Description	Reference Section/Document
1. Configure the reserved LUN pool	<ul style="list-style-type: none"> <input type="checkbox"/> On each SP, determine the size of the reserved LUN pool. <input type="checkbox"/> On the storage system, bind one or more LUNs on each SP to the size you determined for the reserved LUN pool. <input type="checkbox"/> From the ph13245 server (production server), allocate the reserved LUNs to the SP's LUN pool. 	The latest revision of the <i>EMC Navisphere Manager Administrator's Guide</i> .
2. Create a snapshot	<ul style="list-style-type: none"> <input type="checkbox"/> From the ph13245 server (production server), create two snapshots (OracleDataSnap and OracleLogSnap). 	<i>Creating a Snapshot</i> on page 2-24
3. Start a SnapView session	<ul style="list-style-type: none"> <input type="checkbox"/> Monday - From the ph13245 server (production server), start a SnapView session (OracleSession1). <input type="checkbox"/> Tuesday - From the ph13245 server (production server), start a SnapView session (OracleSession2). <input type="checkbox"/> Wednesday - From the ph13245 server (production server), start a SnapView session (OracleSession3). <input type="checkbox"/> Thursday - From the ph13245 server (production server), start a SnapView session (OracleSession4). <input type="checkbox"/> Friday - From the ph13245 server (production server), start a SnapView session (OracleSession5). 	<i>Starting a SnapView Session</i> on page 2-17
4. Add the snapshot to a Storage Group	<ul style="list-style-type: none"> <input type="checkbox"/> From the ph13245 server (production server), add snapshot OracleDataSnap and snapshot OracleLogSnap to Storage Group OracleDB_backup. 	<i>Adding a Snapshot to a Storage Group</i> on page 2-27

Task	Task Description	Reference Section/Document
5. Activate the snapshot	<input type="checkbox"/> From the sh12345 server (secondary server), activate the snapshots (OracleDataSnap and OracleLogSnap) to Friday's SnapView session (OracleSession5).	<i>Activating a Snapshot</i> on page 3-23
6. Identify the snapshot	<input type="checkbox"/> From the sh12345 server (secondary server), identify the snapshots (OracleDataSnap and OracleLogSnap) to the operating system. This procedure depends on the operating system; for example, on a Windows server you need to run Disk Admin . Note This step is needed only once, as part of the SnapView initial setup.	None
7. Deactivate the snapshot	<input type="checkbox"/> While viewing Friday's session, you realize that the database and its log file are corrupted or contains changes that you do not want. So from the sh12345 server (secondary server), deactivate the snapshots (OracleDataSnap and OracleLogSnap) from Friday's SnapView session (OracleSession5), so you can view the sessions that were started earlier in the week.	<i>Deactivating a Snapshot</i> on page 3-25
8. Verify other SnapView sessions	<input type="checkbox"/> Repeat steps 5 to 7 until you get to Tuesday's session (OracleSession2), which contains the data you want.	
9. Start rollback operation	<input type="checkbox"/> From the ph12345 server (production server), start the rollback operation on Tuesday's SnapView session (OracleSession2). When you confirm the start of a rollback operation, the source LUN can instantly access the session's point-in-time data, while data copying continues in the background.	<i>Rolling Back a SnapView Session</i> on page 3-27
10. Continue daily sessions	<input type="checkbox"/> Once the rollback completes, which includes all background copying, from the ph12345 server (production server), resume starting your daily sessions.	<i>Starting a SnapView Session</i> on page 2-17

Illustrated Overview

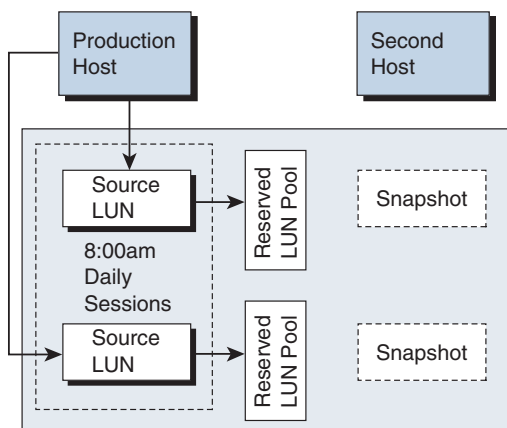
The following section provides an illustrated description of the *main* operations described in the table in the previous section.

1. Production server creates the snapshots.



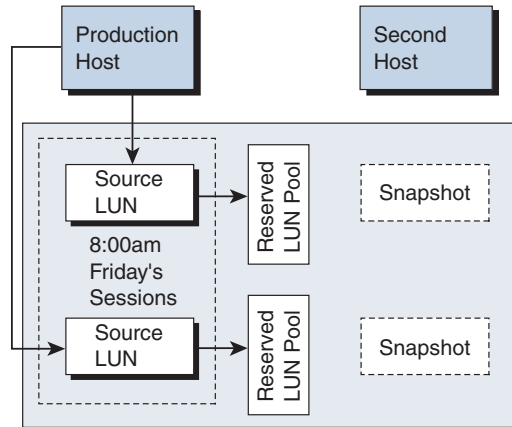
EMC2759

2. Production server starts a daily weekday SnapView session at 8:00am. Any modifications made to the source LUNs from the production server are written to the source LUNs. The software copies the original data that was just modified and stores it in the reserved LUN pool.



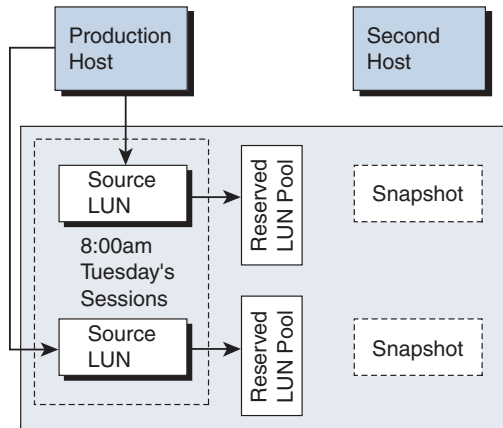
EMC2760

- The secondary server identifies and activates the snapshots to Friday's session. While viewing Friday's session, you realize that the database and its log file are corrupted or contains changes that you do not want.



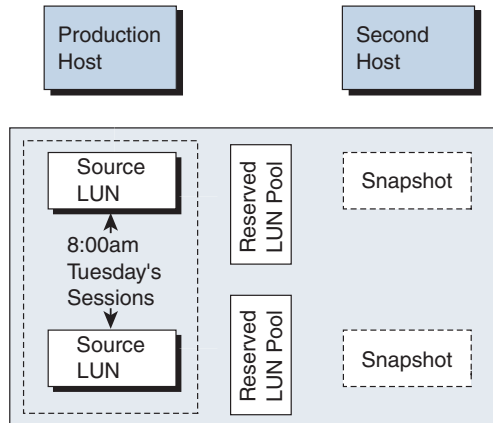
EMC2761

- The secondary server deactivates the snapshot from Friday's session and activates it to Tuesday's session, which contains the data you want.



EMC2762

5. Start rollback operation. Tuesday's point-in-time data is copied to the source LUN.



EMC2763

Using SnapView on a Tru64 Server

This appendix describes how to use SnapView clones and snapshots with a Tru64 server.

Major sections in this appendix are

Prerequisites

- ◆ Determining a Tru64 Source LUN A-2

Clones

- ◆ Setting Up Clones..... A-8
- ◆ Using Clones..... A-9

Snapshots

- ◆ Setting Up Snapshots..... A-13
- ◆ Using Snapshots..... A-14

Determining a Tru64 Source LUN

It is recommended that you determine a Tru64 source LUN when you first create the source LUN. The following steps are necessary only for those cases where the relationship between the file system or raw device and the storage-system LUN is not known with certainty. It is important that you follow these steps to ensure that the correct source LUN is being used.

The following procedure starts with an example of a mounted file system named **/source**. In the case of a raw disk partition, start at step 3.

1. Select the file system you want to copy and show where it is mounted with the **mount** command. The following is an example output for the server with the **/source** file system:

```
root_domain#root on / type advfs (rw)
/proc on /proc type procfs (rw)
usr_domain#usr on /usr type advfs (rw)
usr_domain#var on /var type advfs (rw)
source_domain#source_fset on /source type advfs (rw)
```

The **/source** file system is listed as a mount point of the **source_fset** AdvFS fileset, which is part of the **source_domain** AdvFS domain.

2. Determine the disk device that the mount point is using.

For **ufs** file systems the disk device is part of the **mount** command output so you can skip this step.

The **/source** file system in this example is an **advfs** file system, so you can use the **showfdmn** command to list the device(s) that are part of the domain.

The following is the output from the **showfdmn source_domain** command:

```

Id   Date Created      LogPgs  Version  Domain Name
3c88d6c8.0002d460 Fri Mar8 10:20:40  512  4
source_domain
Vol   512-Blks  Free   % Used Cmode Rblks Wblks Vol
Name
1L    85729280 83176608 3%   on  256  256
/dev/disk/dsk347c

```

Source_domain as shown comprises the **/dev/disk/dsk347c** disk device.

3. Determine the LUN and SCSI bus number associated with the disk device.
 - a. Use the **hwmgr -view devices -dsf** command to determine the Hardware Identifier (HWID) of the particular device. The output is similar to the following:

```

hwmgr -view devices -dsf /dev/disk/dsk347c
HWID: Device Name      Mfg   Model  Location
417:/dev/disk/dsk347c DGC   RAID 5 IDENTIFIER=264

```

In the above output the HWID is 417.

The IDENTIFIER entry matches the User Defined Identifier (UDID) that the SRM console WWIDMGR program uses.

- b. Use the HWID to get detailed information for the particular device using the **hwmgr -show scsi -id** command. The following is output from the **hwmgr -show scsi -id 417 -full** command:

```

          SCSI                DEVICE DEVICE DRIVER NUM DEVICE
FIRST
HWID:DEVICEID HOSTNAME TYPE  SUBTYPE OWNER  PATH
FILE VALID PATH
417: 56          182ba209 disk   none   2      2
dsk347 [5/1/8]

WWID:01000010:6006-0173-1460-0000-9518-a222-272d-d6
11

```

```

          BUS   TARGET  LUN   PATH STATE
          -----
          5     1       8     valid
          5     0       8     valid

```

In this example, the devices reside on SCSI bus 5, target 0 or 1 and LUN 8. Both SPs are also connected to a single HBA; if the SPs are connected to two HBAs, then you will see two SCSI bus numbers.

- c. If you know the CX600, CX400, or FC4700-Series storage system on which the above LUN resides, go to step 8. Otherwise, continue to step 4 to identify the storage system.
4. Identify the server bus adapter that makes up this SCSI bus by using the following **grep** command in the **/var/adm/messages** file:

```
grep scsi5 /var/adm/messages
```

where **scsi5** is SCSI bus number 5. This command produces the following output:

```
Mar 12 10:47:48 182ba209 vmunix: scsi5 at emx4
slot 0 rad 0
```

The adapter number is emx4.

- Use the **emxmgr** utility to display the topology information for this adapter. Issue the **emxmgr -t emx4** command to get the following:

emx4 state information:

```
Link : connection is UP
      Point to Point
      Fabric attached
      FC DID 0x011a00
```

```
Link is SCSI bus 5 (e.g. scsi5)
```

```
SCSI target id -1
portname is 1000-0000-C924-0CD5
nodename is 2000-0000-C924-0CD5
```

```
N_Port at FC DID 0x011400 - SCSI tgt id 0 :
```

```
portname 5006-0160-4004-D3A5
nodename 5006-0160-2004-D3A5
```

```
Present, Logged in, FCP Target, FCP Logged in,
```

```
N_Port at FC DID 0x011500 - SCSI tgt id 1 :
```

```
portname 5006-0168-4004-D3A5
nodename 5006-0160-2004-D3A5
```

```
Present, Logged in, FCP Target, FCP Logged in,
```

```
N_Port at FC DID 0xfffffc - SCSI tgt id -1 :
```

```
portname 20FC-0060-6950-08EA
nodename 1000-0060-6950-08EA
```

```
Present, Logged in, Directory Server,
```

```
N_Port at FC DID 0xfffffe - SCSI tgt id -1 :
```

```
portname 200A-0060-6950-08EA
nodename 1000-0060-6950-08EA
```

```
Present, Logged in, F_PORT,
```

- From the above output, verify that this adapter is SCSI bus number 5 and that target IDs 0 and 1 are portnames 5006-0160-4004-D3A5 and 5006-0168-4004-D3A5 respectively. The complete World Wide Name is constructed by prefacing the port name with the node name.
- In Navisphere Manager, display the properties of each Storage Group attached to the server by right-clicking the Storage Group name and selecting **Properties**. From the **Storage Group Properties** dialog box, click the **Advanced** tab.

Verify that the SP Port World Wide Name information for both SPs matches the target information from step 5.

- Determine which storage-system LUN is represented by the Server LUN 8 we identified in step 3. For this step you must have Navisphere CLI (navicli) installed.

Important Navisphere CLI is *not* supported on a Tru64 server. You must install CLI on another server that is connected to the Tru64 server. To install the CLI, refer to the EMC Host Agent and CLI Installation Guide for the operating system of the other server.

Enter the following CLI command on the other server:

```
navicli -h SP-servername storagegroup -list -gname  
storage-groupname
```

where

SP-servername is the server name of the FC4700 storage system.

storage-groupname is the name of the Storage Group.

This command provides a mapping of storage-system (array) LUNs (ALUs) to server LUNs (HLUs), as shown in the following output:

ALU is the ID that is specified when you bind a LUN.

```
Storage Group Name:      Group-RB-1
Storage Group UID:
2A:CF:C7:8B:5E:22:D6:11:80:3D:08:00:1B:41:30:3F
HBA/SP Pairs:
  HBA UID
SP Name      SP Port

20:00:00:00:C9:24:0C:D5:10:00:00:00:C9:24:0C:D5
SP A          0

20:00:00:00:C9:24:0C:D5:10:00:00:00:C9:24:0C:D5
SP B          0
HLU/ALU Pairs:
  HLU Number      ALU Number
    0              0
    1              1
    2              2
    3              3
    4              4
    5              5
    6              6
    7              7
    8              113
Shareable:          NO
```

HLU Number 8 maps to ALU Number 113. Thus array LUN 113 is the logical unit that we want to use as the source for the SnapView operation.

What Next?

To use SnapView clones, continue to the next section, *Setting Up Clones*.

To use SnapView snapshots, go to *Using Snapshots* on page A-14.

Setting Up Clones

To set up clones with a Tru64 server, you must determine a source LUN, as described on page A-2. Once you have determined a source LUN, you must allocate Clone Private LUNs. Then you can create a Clone Group and add a clone to that group.

Creating a Clone Group and Adding a Clone

The following steps outline the order in which you should perform tasks when using SnapView clones with a Tru64 server.

The order of the following steps may vary between operating systems, depending on the utilities that are available for a particular environment.

For a detailed description of how to perform each task, refer to the reference section listed.

1. Allocate Clone Private LUNs (see *Allocating Clone Private LUNs* on page 2-4).
2. Create a Clone Group using the appropriate source LUN (see *Creating a Clone Group* on page 2-7).
3. Add a clone to the Clone Group to create the copy of the source LUN (see *Adding a Clone to a Clone Group* on page 2-9).

Allow the clone to synchronize with the source LUN before fracturing the clone. Another server cannot use the clone until it is fractured. You should fracture the clone only while it is in a Synchronized or Consistent state.

4. Fracture the clone from the Clone Group using Navisphere (see *Fracturing a Clone* on page 3-5).

Using Clones

This section describes how a secondary server can activate and access a clone.

Verifying the Clone

To verify the clone LUN, do the following:

1. Verify that the clone LUN is in a Storage Group and issue the following command from the Tru64 UNIX server in this Storage Group:

```
hwmgr -scan scsi
```

This command updates the SCSI device database and will make the clone LUN visible. Wait for this command to complete its actions, which are performed asynchronously.

2. Verify the completion by identifying the newly found clone LUN using the following command:

```
hwmgr -show scsi
```

What Next?

To access the clone, continue to the next section.

Accessing the Clone

The specific method for accessing the data on the clone LUN depends on the particular format of the source LUN. This section shows how to access clone LUNs for the following types of file systems:

- ◆ ufs – UNIX[®] file system
- ◆ AdvFS – Tru64 UNIX Advanced file system

ufs File System

In the case of a ufs file system, all file system information will reside on the disk and you do not need to take any additional steps to identify the file system to the server. When you create a clone of a ufs file system, the file system will be dirty unless it was not mounted at the time the clone was added to the Clone Group. When attempting to mount a dirty file system a message such as the following will be displayed:

```
/dev/disk/dsk352c on /ufscopy: Dirty file system
```

When such a message appears, it is necessary to run **fsck** on this file system. For the file system in the above example issue the following **fsck** command:

```
fsck -y /dev/disk/dsk352c
```

This command produces the following output, indicating the progress through the check and fix of the affected file system:

```
/sbin/ufs_fsck -y /dev/disk/dsk352c
** /dev/rdisk/dsk352c
** Last Mounted on /ufssource
** Phase 1 - Check Blocks and Sizes
** Phase 2 - Check Pathnames
** Phase 3 - Check Connectivity
** Phase 4 - Check Reference Counts
** Phase 5 - Check Cyl groups
23819 files, 545252 used, 11930300 free (3956
frags, 1490793 blocks, 0.0% fragmentation)
Filesystem '/dev/rdisk/dsk352c' Tru64 UNIX UFS
v.3 UFS
```

Once the command has completed, if necessary, create the mount point. You can then mount the file system using the following **mount** command:

```
mkdir /ufscopy
mount /dev/disk/dsk352c /ufscopy
```

You can now use the clone. Remember to unmount the clone prior to resynchronizing it on the source LUN.

AdvFS File System

To access an AdvFS file system that has been snapped, you must take several steps to properly inform the backup server of the file system. An AdvFS file system exists on a file domain that contains links to the device special files that are part of the domain. This link has to be created manually, since the standard **mkfdmn** command will write new information to the device thus destroying the clone.

The file domain information is kept in a set of subdirectories of **/etc/fdmns**. To create the correct information on the secondary server, complete the following steps:

1. Create the following subdirectory in **/etc/fdmns**:

```
mkdir /etc/fdmns/new_domain_name
```

where

new_domain_name is the name you want to give this domain on the backup server.

The domain name must be unique to the server; for example, **mkdir /etc/fdmns/CopyDomain**.

2. Create a symbolic link between the device special file and the new domain as follows:

```
ln -s dev_special_file /etc/fdmns/new_domain_name/dev_special_file
```

This link points AdvFS to the correct device and the file-set information on this device. For example:

```
ln -s /dev/disk/dsk351c /etc/fdmns/CopyDomain/dsk351c
```

Use the device name within the domain to identify the particular device.

At this point the server has the required information to access the AdvFS file domain and file set. The file-set information will stay the same regardless of the domain name; that is if the file set was named **source** on the source server, it will still be named **source** on the backup server.

3. Create a mount-point directory, if necessary, as follows:

```
mkdir mountpoint
```

4. To ensure the integrity of the AdvFS domain, verify it using the following command:

```
/sbin/advfs/verify new_domain_name
```

This **verify** command returns an error if the clone and the source LUN are on the same server.

If this is the case, modify the **mount** command for the file set with the **-o dual** option (if the clone and source are on separate servers, this is not necessary):

```
mount [-o dual] new_domain_name#fileset mountpoint
```

An example of the command for a case where both source and clone are on the same server is:

```
mount -o dual CopyDomain#SourceSet /advfsCopy
```

The following message will be generated for this case (but not if the servers are distinct):

```
Dual mounting a split mirror AdvFS filesystem.  
This takes a short while to update the domain's ID.
```

The clone is now ready for use. Again, remember to unmount the clone before resynchronizing it.

What Next?

Depending on your applications needs, do one of the following:

- ◆ To resynchronize the clone, refer to *Synchronizing a Fractured Clone* on page 3-11.
- ◆ To reverse synchronize the clone, refer to *Reverse Synchronizing a Fractured Clone* on page 3-13.
- ◆ To remove the clone from the Clone Group, refer to *Removing a Clone from a Clone Group* on page 3-18.

Setting Up Snapshots

To set up snapshots with a Tru64 server, you must determine a source LUN, as described on page A-2. Once you have determined a source LUN, you must configure the reserved LUN pool. Then you can create a snapshot and start a SnapView session.

Creating a Snapshot and Starting a SnapView Session

The following steps outline the order in which tasks should be performed when using SnapView snapshots with a Tru64 server.

The order of the following steps may vary between operating systems, depending on the utilities that are available for a particular environment.

For a detailed description of how to perform each task, refer to the reference section listed.

1. Determine a suitable reserved LUN pool size (refer to *EMC Navisphere Manager Administrator's Guide*).
2. Bind one or more LUNs on an SP to the size you determined for the reserved LUN and add them to the SP's reserved LUN pool (refer to the *EMC's Navisphere Manager Administrator's Guide*).
3. Create a snapshot (see *Creating a Snapshot* on page 2-24).

You start a SnapView session first or create a snapshot first. However, a secondary server cannot view the session data unless a snapshot is activated to the session and that snapshot is accessible to the secondary server.

4. Start a SnapView session (see, *Starting a SnapView Session* on page 2-17).

What Next? To activate and access the snapshot, continue to the next section.

Using Snapshots

This section describes how a secondary server can activate and access a snapshot.

Activating a Snapshot

To activate the snapshot from a secondary server, do the following:

1. Activate the snapshot (see, *Activating a Snapshot* on page 3-23).
2. Enter the following command to initiate the update of the SCSI device database:

```
hwmgr -scan scsi
```

Once the system completes this command, the snapshot LUN is accessible to the secondary server.

You can use the Navisphere CLI command from a non-Tru64 server to identify the server LUN number of the new snapshot LUN, as described in the *Determining a Tru64 Source LUN* section.

What Next? To access the snapshot, continue to the next section.

Accessing the Snapshot

The specific method for accessing the data on the snapshot LUN depends on the particular format of the source LUN. This section shows how to access snapshot LUNs for the following types of file systems:

- ◆ ufs – UNIX® file system
- ◆ AdvFS – Tru64 UNIX Advanced file system

ufs File System

In the case of a ufs file system, all file system information will reside on the disk and you do not need to take any additional steps to identify the file system to the server. When a snapshot is taken of a ufs file system, the file system will be dirty unless it was not mounted at the time of the snapshot. When attempting to mount a dirty file system a message such as the following will be displayed:

```
/dev/disk/dsk352c on /ufscopy: Dirty file system
```

When such a message appears, it is necessary to run **fsck** on this file system. For the file system in the above example issue the following **fsck** command:

```
fsck -y /dev/disk/dsk352c
```

This command produces the following output, indicating the progress through the check and fix of the affected file system:

```
/sbin/ufs_fsck -y /dev/disk/dsk352c
** /dev/rdisk/dsk352c
** Last Mounted on /ufssource
** Phase 1 - Check Blocks and Sizes
** Phase 2 - Check Pathnames
** Phase 3 - Check Connectivity
** Phase 4 - Check Reference Counts
** Phase 5 - Check Cyl groups
23819 files, 545252 used, 11930300 free (3956
frags, 1490793 blocks, 0.0% fragmentation)
Filesystem '/dev/rdisk/dsk352c' Tru64 UNIX UFS
v.3 UFS
```

Once the command has completed, if necessary, create the mount point. You can then mount the file system using the following **mount** command:

```
mkdir /ufscopy
mount /dev/disk/dsk352c /ufscopy
```

You can now use the snapshot. Remember to unmount the snapshot prior to destroying it on the storage system.

AdvFS File System

To access an AdvFS file system that has been snapped, you must take several steps to properly inform the backup server of the file system. An AdvFS file system exists on a file domain that contains links to the device special files that are part of the domain. This link has to be created manually, since the standard **mkfdmn** command will write new information to the device thus destroying the snapshot.

The file domain information is kept in a set of subdirectories of **/etc/fdmns**. To create the correct information on the snapshot server, complete the following steps:

1. Create the following subdirectory in **/etc/fdmns**:

```
mkdir /etc/fdmns/new_domain_name
```

where

new_domain_name is the name you want to give this domain on the backup server.

The domain name must be unique to the server; for example, **mkdir /etc/fdmns/CopyDomain**.

2. Create a symbolic link between the device special file and the new domain as follows:

```
ln -s dev_special_file /etc/fdmns/new_domain_name/dev_special_file
```

This link points AdvFS to the correct device and the file-set information on this device. For example:

```
ln -s /dev/disk/dsk351c /etc/fdmns/CopyDomain/dsk351c
```

Use the device name within the domain to identify the particular device.

At this point the server has the required information to access the AdvFS file domain and file set. The file-set information will stay the same regardless of the domain name; thus if the file set was named **source** on the source server, it will still be named **source** on the backup server.

3. Create a mount-point directory, if necessary, as follows:

```
mkdir mountpoint
```

4. To ensure the integrity of the AdvFS domain, verify it using the following command:

```
/sbin/advfs/verify new_domain_name
```

This **verify** command returns an error if the snapshot and the source LUN are on the same server.

If this is the case, modify the **mount** command for the file set with the **-o dual** option (if the snapshot and source are on separate servers, this is not necessary):

```
mount [-o dual] new_domain_name#fileset mountpoint
```

An example of the command for a case where both source and snapshot are on the same server is:

```
mount -o dual CopyDomain#SourceSet /advfsCopy
```

The following message will be generated for this case (but not if the servers are distinct):

```
Dual mounting a split mirror AdvFS filesystem.  
This takes a short while to update the domain's ID.
```

The snapshot is now ready for use. Again, remember to unmount the snapshot before destroying it.

What Next?

Depending on your application's needs, do one of the following:

- ◆ To stop a session, refer to *Stopping a SnapView Session* on page 3-33.
- ◆ To destroy a snapshot, refer to *Destroying a Snapshot* on page 3-34.

Customer Support

This appendix reviews the EMC process for detecting and resolving software problems, and provides essential questions that you should answer before contacting the EMC Customer Support Center.

This appendix covers the following topics:

- ◆ Overview of Detecting and Resolving Problems B-2
- ◆ Troubleshooting the Problem B-3
- ◆ Before Calling the Customer Support Center B-4
- ◆ Documenting the Problem..... B-5
- ◆ Reporting a New Problem B-6
- ◆ Sending Problem Documentation..... B-7

Overview of Detecting and Resolving Problems

EMC software products are supported directly by the EMC Customer Support Center in the United States.

EMC uses the following process to resolve customer problems with its software products (Figure B-1).

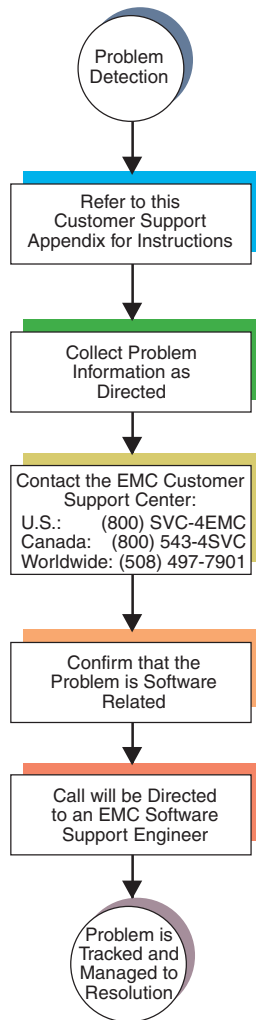


Figure B-1 Problem Detection and Resolution Process

Troubleshooting the Problem

Please perform the relevant diagnostic steps before you contact the EMC Customer Support Center:

1. Read the documentation carefully.
2. Reconstruct the events leading up to the problem and describe them in writing.
3. Run some test cases to reproduce the problem.

If you encounter a problem that requires technical programming or analysis, call the nearest EMC office or contact the EMC Customer Support Center at one of the following numbers:

United States: **(800) 782-4362 (SVC-4EMC)**

Canada: **(800) 543-4782 (543-4SVC)**

Worldwide: **(508) 497-7901**

Please do not request a specific support representative unless one has already been assigned to your particular system problem.

For additional information on EMC products and services available to customers and partners, refer to the EMC Powerlink website at:

<http://powerlink.EMC.com>

Before Calling the Customer Support Center

Have the following information available before calling the Customer Support Center or your support representative (if one has been assigned to you):

- Your company name
- Your name
- Your phone number
- Your site ID, if known
- For an existing problem, the problem tracking system ID, if one was previously assigned to the problem by a support representative

Documenting the Problem

If the EMC Customer Support Center requests information regarding the problem, please document it completely, making sure to include the following information:

- Your company name and address
- Your name
- Your telephone number
- Your site ID
- The importance of the problem, so that it can be assigned a priority level

To expedite the processing of your support request, you can photocopy this list and include it with the package.

Reporting a New Problem

For a new problem, please provide the following information:

- Release level of the software that you are running
- Software installation parameters
- Host type on which you are running
- Operating system you are running and its release number
- Functions of the software that you are running
- Whether you can reproduce the problem
- Previous occurrences of the problem
- Whether the software has ever worked correctly
- Time period that the software did work properly
- Conditions under which the software worked properly
- Changes to your system between the time the software worked properly and the problem began
- Exact sequence of events that led to the system error
- Message numbers and complete text of any messages that the system produced
- Log file dated near the time the error occurred
- Results from tests that you have run
- Other related system output
- Other information that may help solve the problem

Sending Problem Documentation

Use one of the following methods to send documentation of the problem to the EMC Customer Support Center:

- ◆ E-mail
- ◆ FTP
- ◆ U.S. mail to the following address:

EMC Customer Support Center
171 South Street
Hopkinton, MA 01748-9103

If the problem was assigned a number or a specific support representative, please include that information in the address as well.

This glossary contains terms related to SnapView. Many of these terms are used in this manual.

A

Active A snapshot is currently participating in a SnapView session and is accessible to secondary servers.

Activate An operation on a snapshot that maps it to a SnapView session. This feature is available in Navisphere Manager and admsnap and CLI.

admsnap Server-based software that provides a command line interface to SnapView software running in a storage-system SP. With admsnap, you can start and stop sessions and activate and deactivate snapshots by typing commands on a secondary server system.

B

Business Continuance Volumes (BCVs) Another term used for clones. See *Clone*.

C

Chunk An aggregate of multiple disk blocks that SnapView uses to perform copy-on-first-write operations. The selectable chunk sizes are 16 Kbytes, 32 Kbytes, 64 Kbytes, 128 Kbytes, 256 Kbytes, and 512 Kbytes. The default size is 64 Kbytes (128 blocks in Navisphere). For SnapView Version 2.1 or higher, the chunk size is set to 64K (128 blocks). You cannot change this value.

CLI	Navisphere Command Line Interface. SnapView uses two CLIs: a clone CLI and a snapshot CLI. The clone CLI is implemented in Java and the snapshot CLI is not.
Client	A server (computer or laptop) that has an Internet browser and connects to a storage application server via a network. You use it to manage connected storage systems.
Clone	A LUN that is an actual copy of a specified source LUN. The state of the clone determines if it is a byte-for-byte copy of its source. You create a clone when you add a clone to the Clone Group.
Clone Group	A collection of a source LUN and all of its clones. The purpose of creating a Clone Group is to establish a source LUN that you may want to clone at some time.
Clone Private LUNs	LUNs that record information that identifies areas on the source and clone that have changed since the clone was fractured. A log in the clone private LUN records this information but no actual data is written to the clone private LUN. This log is a bitmap and reduces the time it takes to synchronize and reverse synchronize a clone and its source.
Clone States	Each clone in a Clone Group has its own state. The state of the clone determines whether or not the clone is usable. The possible clone states are Consistent, Out-of-Sync, Reverse Out-of-Sync, Reverse Synchronizing, Synchronized, or Synchronizing.
Consistent Fracture	Fracturing more than one clone at the same time. The clones can be within the same Clone Group or within different Clone Groups. After the consistent fracture completes, there is no group association between the clones.
Consistent Mode	Preserves the point-in-time copy across a set of source LUNs. The SnapView driver will any delay any I/O requests to the set of source LUNs until the session has started on all LUNs (thus preserving the point-in-time on the entire set of LUNs).
Consistent State	A clone in a Synchronized state that receives server I/O to the source (if the clone is unfractured) or to the clone (if the clone is fractured). A consistent clone is usable but may not contain the most up-to-date information since writes made to the source have not been copied to the clone.

Copy-On-First-Write An algorithm that copies current contents of a source LUN before it is modified (written to). The copy-on-first-write operation is on a chunk: before the first modification of any disk blocks within a chunk, the software reads and stores the original data of the chunk in the reserved LUN pool. This policy applies only to the first modification of the data. Overwrite of any data that has already had a copy-on-first-write does not require any extra processing since the software saved the original data in the reserved LUN pool.

D

Deactivate An operation on a snapshot that unmaps it from a SnapView session to make it invisible to any secondary servers. The software destroys any writes made to the snapshot but the snapshot and SnapView session still exist. This feature is available in Navisphere Manager and admsnap, however, the Manager deactivate function does not flush all data and clear all buffers on the secondary server.

F

Fracture The process of breaking off a clone from its source. Once a clone is fractured, it can receive server I/O requests.

H

Host Agent Navisphere Agent that runs on a server system.

I

Inactive A snapshot that is not currently participating in a SnapView session and is invisible to any secondary servers.

M

Modified Data Chunk A chunk of data that a server changes by writing to the clone, snapshot, or source LUN.

N

Navisphere Manager The EMC Navisphere Manager application.

O

Out-of-Sync State A clone that was in the process of synchronizing but failed. An Out-of-Sync clone is not a byte-for-byte copy of its source LUN and therefore, is unusable.

P

Persistent Mode Creates a session that can withstand an SP reboot or failure, a storage system reboot or power failure, or server I/O trespassing to the peer SP.

Point-in-Time The moment a SnapView session starts.

Private LUN A LUN that cannot be assigned to a Storage Group. Once you add a LUN to the reserved LUN pool or allocate a LUN as a clone private LUN, it becomes a private LUN.

Protected Restore When selected, a process that prevents source writes from being copied to the clone during a reverse synchronization.

Q

Quiesce Threshold The time period after which, without I/O from the server, any clone in the Consistent state and not fractured is marked as being in the Synchronized state. Valid values are 10 – 3600 seconds. The default is 60 seconds.

R

Recovery Policy The policy used to determine how a clone is recovered after a failure. Options are **auto** or **manual**.

Reserved LUN A private LUN (a LUN to which a server cannot perform I/O) assigned to an SP's reserved LUN pool.

Reserved LUN Pool The disk storage used to store blocks of original data chunks when you first modify that chunk on the source LUN(s) after the start of a session. Each SP manages its own LUN pool space and assigns a separate reserved LUN (or multiple LUNs) to each source LUN.

Reserved Sessions Sessions used for another application such as SAN Copy and MirrorView / Asynchronous.

Reserved Snapshots	Snapshots used for another application such as SAN Copy and MirrorView/Asynchronous.
Restartable Copy	A data state having dependent write consistency and where all internal database/application control information is consistent with a Database Management System/application image.
Reverse Out-of-Sync State	A clone that was in the process of reverse synchronizing but failed. Therefore, the source LUN is unusable and another reverse synchronization is recommended.
Reverse Synchronizing State	A clone that is unfractured and in the process of copying its data to its source LUN.
Rollback	Restores the point-in-time copy of a SnapView session to the source LUN(s).
S	
Snapshot	Views a point-in-time image of a source LUN(s). A snapshot occupies no disk space, but appears like a normal LUN to secondary servers and can serve for backup or another use. Other, older terms for snapshot, which are no longer used, include SnapshotCopy LUN (SCLUN) and SnapCopy LUN (SLU) .
SnapView	Allows you to obtain a copy of a LUN by creating a clone or snapshot. The clone or snapshot can serve for backup, decision support scenarios, or as a base for temporary operations on the production data without damaging the original data on the source LUN.
SnapView Session	The period of time that SnapView is managing a reserved LUN pool region. The SnapView session begins when you start a session using Navisphere Manager, Navisphere CLI, or admsnap and ends when you stop the session. You can give each session a name (the session name) when you start the session. The name persists throughout the session and is viewable through Navisphere. You use the name to check session status and to end the session.
Source LUN	The original LUN from which a clone or snapshot is generated. An older term for source LUN, which is no longer used, is Target LUN (TLU) .
SP Agent	The Navisphere Agent that runs in an SP (CX-Series and FC-Series storage systems).

Synchronization Rate	Specifies a relative value (low, medium, or high) for the priority of completing updates. High completes updates faster, but may significantly affect storage system performance for host I/O requests. Low completes updates slower, but also minimizes the impact on other storage system operations.
Synchronized State	A clone that is a byte-for-byte copy of its source and, therefore, is usable.
Synchronizing State	An unfractured clone that is in the process of copying data from its source LUN.

A

access, snapshot 2-27
 activate, snapshot 3-23
 adding to a Storage Group 2-27
 admsnap 1-9
 definition of g-1
 introduction to 1-7
 AdvFS file system, *see* Tru64 server

B

backup
 benefit with SnapView 1-2
 snapshot example 5-10
 see also secondary server
 bad blocks
 on clones 3-21
 on rollback 3-36
 basic storage component icons 1-16
 business continuance volumes (BCVs), *see* clone

C

chunk g-1
 chunk, modified data g-3
 CLI 1-8
 clone
 compared to snapshot 1-6
 creating 2-9, 2-10
 definition of 1-3, g-2
 dirty 4-2
 example 5-2
 fracturing 3-5
 ID 1-3, 2-9, 4-2
 overview 1-3

 properties of 4-2
 removing from Clone Group 3-18
 setting up 2-2
 states 3-3, g-2
 Clone Group
 adding a clone to 2-9, 2-10
 creating 2-7
 definition of 2-7, g-2
 destroying 3-19
 properties 4-3
 removing clone from 3-18
 clone private LUNs
 allocating 2-4, 2-5
 deallocating 2-6
 definition of g-2
 reallocating 2-6
 clone state
 consistent 3-3, g-2
 out-of-sync 3-3, g-4
 reverse out-of-sync 3-3, g-5
 reverse synchronizing 3-4, g-5
 synchronized 3-4, g-6
 synchronizing 3-4, g-6
 configuration guidelines, SnapView 1-12
 consistent clone state, definition of 3-3, g-2
 copy-on-first-write, definition of 1-4, g-3
 Customer support B-3
 Customer Support Center B-7

D

decision support
 benefit with SnapView 1-2
 see also secondary server
 dirty clone 4-2

E

emxmgr utility, *see* Tru64 server
 expansion, LUN 2-8, 2-11, 2-19

F

fractured clone
 remove 3-18
 fracturing a clone 3-5
 fracturing, definition of g-3

G

grep command, *see* Tru64 server

I

icons, basic storage components 1-16
 inactive snapshot 2-26, g-3

L

LUN expansion 2-8, 2-11, 2-19
 LUNs 2-4
 bind 2-4

M

metaLUN 2-8, 2-11, 2-19
 modified data chunk g-3
 multiple snapshots 2-24
 multiple SnapView sessions 2-24

N

Navisphere Manager 1-8

O

out-of-sync clone state, definition of 3-3, g-4

P

private LUN g-4
 properties
 clone 4-2
 Clone Feature Properties 4-3
 Clone Group 4-3
 snapshot 4-5

 snapshots and SnapView sessions,
 displaying all 4-7
 SnapView session 4-6
 source LUN 4-4
 protected restore feature 3-14

Q

Quiesce Threshold
 definition of g-4
 setting the 2-8

R

reserved
 LUN, definition of g-4
 sessions, definition of g-4
 snapshots, definition of g-5
 reserved LUN pool
 definition of g-4
 icon for 1-16
 with SnapView 2-15
 reverse out-of-sync clone state, definition of 3-3,
 g-5
 reverse synchronizing
 clone state, definition of g-5
 fractured clone 3-13
 protected restore feature 3-14
 revision testing
 see secondary server
 revision testing, benefit with SnapView 1-2
 rollback 3-27
 background copying 3-27, 3-31
 definition of 3-27
 progress of 3-31
 rate 3-31
 recovery session 3-27
 session 3-27
 starting 3-28, 3-29

S

secondary server 1-10
 server
 production 1-10
 secondary 1-10
 setting up clones 2-2

- snapshot
 - activating 3-23
 - compared to clone 1-6
 - creating a 2-24
 - deactivating 3-25
 - deactivating a g-3
 - definition of g-5
 - destroy 3-34
 - example 5-10
 - inactive 2-26
 - multiple 2-24
 - overview 1-4
 - properties of 4-5
 - SnapView
 - components 1-7
 - configuration guidelines 1-12
 - definition of g-5
 - introduction to 1-2
 - servers 1-10
 - SnapView session
 - definition of g-5
 - multiple 2-24
 - properties of 4-6
 - rolling back 3-27
 - starting a 2-17, 2-20
 - stopping a 3-33
 - software components
 - command line interface 1-7
 - graphical user interface 1-7
 - SnapView 1-7
 - source LUN
 - definition of g-5
 - properties of 4-4
 - SP (storage processor), reserved LUN pool icon 1-16
 - stopping a SnapView session 3-33
 - Storage Group 2-2, 2-27
 - sync command
 - flush function in UNIX 2-20, 3-8, 3-25, 3-29
 - synchronized clone state, definition of 3-4, g-6
- T**
- Technical support B-3
 - tradeoffs between snapshots and clones 1-6
 - Tru64 server A-1
 - AdvFS file system A-11, A-16
 - determining a source LUN A-2
 - emxmgr utility A-5
 - grep command A-4
 - setting up clones A-8
 - adding a clone A-8
 - creating a Clone Group A-8
 - setting up snapshots A-13
 - creating a snapshot A-13
 - starting a SnapView session A-13
 - ufs file system A-9, A-14
 - using clones A-9
 - accessing a clone A-9
 - using snapshots A-14
 - accessing the snapshot A-14
 - activating a snapshot A-14
- U**
- ufs file system, see Tru64 server

