

administrator  
guide

# HP StorageWorks Continuous Access EVA V2.0

**Product Version:** 2.0

First Edition (December 2004)

**Part Number:** AA-RW1EA-TE

This document describes concepts, high-level operations, back-up and recovery procedures, and other ways to use HP StorageWorks Continuous Access EVA.



© Copyright 2003–2004 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided “as is” without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Java™ is a U.S. trademark of Sun Microsystems, Inc.

Linux is a U.S. registered trademark of Linus Torvalds

Microsoft, MS-DOS, MS Windows, Windows, and Windows NT are U.S. registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group.

Printed in the U.S.A.

HP Continuous Access EVA administrator guide  
First Edition (December 2004)  
Part Number: AA–RW1EA–TE

# Contents

<b>About this guide</b> .....	<b>9</b>
Intended audience .....	9
Prerequisites .....	9
Related documentation .....	10
Document conventions and symbols .....	11
HP technical support .....	12
HP authorized reseller .....	12
HP storage web site .....	12
<b>1 About HP Continuous Access EVA</b> .....	<b>13</b>
Overview of HP Continuous Access EVA .....	13
Features .....	15
What's new .....	15
Setup and configuration assumptions .....	16
Licensing .....	16
Hardware and software components .....	16
Enterprise Virtual Array .....	16
Hosts and host operating systems .....	16
Virtual Controller Software .....	16
Zoning .....	17
Management component .....	18
Interface options .....	19
HP Command View EVA .....	20
Previous products .....	22
<b>2 Concepts</b> .....	<b>23</b>
Remote data replication .....	23
Bidirectional replication .....	23
Local replication .....	23
Snapclone .....	24

Snapshot . . . . .	24
DR groups . . . . .	25
Source and destination DR groups . . . . .	25
Replication direction . . . . .	25
Home designation . . . . .	27
DR group presentation . . . . .	27
Presenting DR group members to the same FCA . . . . .	27
DR group properties . . . . .	27
DR group log . . . . .	28
DR group log states . . . . .	28
DR group log size . . . . .	28
Managed sets . . . . .	29
Members . . . . .	29
Managed set properties . . . . .	29
Managed set actions . . . . .	29
Considerations . . . . .	30
General considerations . . . . .	30
Storage resource considerations . . . . .	30
DR group considerations . . . . .	30
Failover . . . . .	31
Considerations . . . . .	31
Replication manager logs . . . . .	33
Failsafe mode . . . . .	34
<b>3 Replication tasks . . . . .</b>	<b>35</b>
About HP Replication Solutions Manager . . . . .	36
Starting the replication manager . . . . .	36
From the replication manager icon (application mode) . . . . .	37
From an authorized client (applet mode) . . . . .	38
Adding remote access IP addresses . . . . .	39
Monitoring events with the replication manager . . . . .	40
Data replication using the replication manager . . . . .	40
Remote replication tasks . . . . .	41
Local replication tasks . . . . .	41
Managing storage with multiple management servers . . . . .	42
Considerations when using multiple management servers to manage storage . . . . .	42
Creating the same configuration on each standby server . . . . .	43
Exporting and importing the database . . . . .	43
Changing the password for the management server . . . . .	45

---

Importing user accounts and group data . . . . .	46
Migrating data from HP Continuous Access user interface to the replication manager . . . .	47
Upgrading array firmware . . . . .	49
Moving storage management to another server . . . . .	50
Synchronizing time on the servers . . . . .	52
Setting array time . . . . .	53
Jobs . . . . .	55
Command line user interface . . . . .	55
<b>4 Recovery . . . . .</b>	<b>57</b>
Planning for a disaster . . . . .	57
Failover . . . . .	58
Backing up configuration information . . . . .	62
Using SSSU to capture configuration information . . . . .	62
Manually capturing configuration information . . . . .	63
Possible event scenarios . . . . .	65
Performing recovery actions . . . . .	65
Planned failover . . . . .	65
Unplanned failover . . . . .	66
Resume operations if unable to access destination while source in failsafe-locked state (extended period of time) . . . . .	66
Return operations to Home array . . . . .	66
Return operations to replaced new storage hardware . . . . .	67
Disk group hardware failure on the source array . . . . .	67
Disk group hardware failure on the destination array . . . . .	67
Failover and recovery procedures . . . . .	68
Planned failover . . . . .	68
Unplanned failover . . . . .	76
Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time) . . . . .	81
Revert to Home (failback) . . . . .	83
Return operations to replaced new storage hardware . . . . .	89
Recovering from a disk group hardware failure . . . . .	98
Disk group hardware failure on the source array . . . . .	99
Disk group hardware failure on the destination array . . . . .	103
<b>5 Troubleshooting . . . . .</b>	<b>107</b>
LUN inaccessible to host . . . . .	107
DR groups in unknown state . . . . .	107

Tunnel thrash .....	108
Remote server cannot detect a destination LUN .....	108
Long delays or time-outs on HP-UX .....	109
IP address on UNIX systems .....	109
Troubleshooting storage problems .....	110
<b>6 Best practices .....</b>	<b>117</b>
Creating a destination snapclone before making a full copy .....	117
Data movement using a snapclone .....	119
Manually specifying disk group membership for a log .....	121
Three-site cascaded data replication using snapclones .....	123
Before you begin .....	125
Procedure .....	126
Post procedure .....	126
Bootless DR group planned failover with Linux using LVM in standalone mode or with SuSE SLES 8 running LifeKeeper 4.4.3 .....	127
Source host procedure .....	127
Destination host procedure .....	128
Red Hat and SuSE Linux Lifekeeper clusters .....	129
Throttling of merge I/O after logging .....	130
Backing up replication jobs and configurations .....	130
Optimizing discovery refresh intervals .....	130
Optimizing discovery performance .....	130
Optimizing browser-based GUI performance .....	131
Coordinating enabled-host downtime .....	131
Minimizing simultaneous jobs .....	131
Avoiding configuration changes while jobs are running .....	131
Optimizing the number of active enabled hosts .....	131
Coordinating enabled host shutdowns .....	132
Coordinating replication server shutdowns .....	132
Avoiding network identification changes .....	132
Maintaining network connections .....	132
Using log files for troubleshooting jobs .....	133
Making CD-ROMs of replication product Web download files .....	133
Managing replication events .....	133
Minimizing simultaneous replication events on an array .....	133
Avoiding simultaneous replication events for the same virtual disk .....	133
Job scheduling .....	133
Complying with EVA snapshot rules .....	134

---

Complying with EVA snapclone rules. . . . .	134
Caching in Microsoft Windows . . . . .	135
Asynchronous replication with failsafe enabled. . . . .	135
<b>A Event message descriptions . . . . .</b>	<b>137</b>
Array messages. . . . .	137
DR group messages . . . . .	138
Job messages. . . . .	139
Error messages . . . . .	148
Dialog box error messages . . . . .	155
 <b>Glossary. . . . .</b>	 <b>157</b>
 <b>Index . . . . .</b>	 <b>163</b>
 <b>Figures</b>	
1 Basic HP Continuous Access EVA configuration with redundant servers . . . . .	14
2 Functional relationship between controllers and server. . . . .	17
3 HP Continuous Access EVA DR group replication. . . . .	26
4 Replicating relationships among DR groups . . . . .	32
5 Set System Time page . . . . .	54
6 Planned and unplanned transfer of operations . . . . .	69
7 Resumption of operations if unable to access destination in failsafe mode. . . . .	82
8 Return operations to new hardware . . . . .	90
9 Creating a DR group from a snapclone . . . . .	120
10 Data movement using snapclones example . . . . .	124

**Tables**

1	Document conventions . . . . .	11
2	Other interfaces and products that perform replication . . . . .	19
3	Previous products . . . . .	22
4	When to and when not to fail over a DR group, managed set, or array . . . . .	60
5	Array configuration record . . . . .	63
6	Failed disk group hardware indicators . . . . .	98
7	Identifying your troubleshooting situation . . . . .	110
8	Array messages . . . . .	137
9	DR group messages . . . . .	138
10	Job messages . . . . .	139
11	Error messages . . . . .	148
12	Dialog box error messages . . . . .	155



## About this guide

This guide provides information about:

- Understanding HP StorageWorks Continuous Access EVA product features.
- Understanding network configuration and replication concepts.
- Monitoring events.
- Planning for and performing failover and recovery procedures.
- Contacting technical support for additional assistance.

## Intended audience

This document is intended for system and network administrators who are experienced with the following:

- Storage Area Network (SAN) fabric configurations
- Host operating system (OS) environments
- Enterprise Virtual Arrays using HSV110 or HSV100 controllers, and running virtual controller software

## Prerequisites

Before using this document, read and follow the information in “[Setup and configuration assumptions](#)” on page 16.

## Related documentation

In addition to this document, please see other documents for this product.

For the following documentation, go to the following web site:

<http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html>

- *HP StorageWorks Continuous Access EVA overview*
- *HP StorageWorks Continuous Access EVA planning guide*
- *HP StorageWorks Continuous Access EVA Performance Estimator user guide*
- *HP StorageWorks EVA replication compatibility reference*
- *HP StorageWorks Replication Solutions Manager installation guide*
- *HP StorageWorks Replication Solutions Manager online Help and user guide*
- *HP StorageWorks Continuous Access EVA license key installation instructions*
- *HP StorageWorks Replication Solutions Manager Command Line User Interface reference guide*

The following document is located on the JRE server CD:

- *HP StorageWorks JRE Server installation instructions*

For the EVA3000, go to:

<http://h18006.www1.hp.com/products/storageworks/eva3000/index.html>

For the EVA5000, go to:

<http://h18006.www1.hp.com/products/storageworks/enterprise/index.html>

For Storage Management Appliances, go to:

<http://h18006.www1.hp.com/products/sanworks/managementappliance/index.html>

For SAN design or SAN extensions, go to:

<http://h18006.www1.hp.com/products/storageworks/san/documentation.html>

## Document conventions and symbols

Table 1: Document conventions

Convention	Element
Blue text: <a href="#">Figure 1</a>	Cross-reference links and email addresses
Blue, underlined text: <a href="http://www.hp.com">http://www.hp.com</a>	Web site addresses
<b>Bold font</b>	GUI elements that are clicked or selected, such as: <ul style="list-style-type: none"> <li>■ Menu and list items</li> <li>■ Buttons</li> <li>■ Check boxes.</li> </ul>
<i>Italics font</i>	Text emphasis
Monospace font	<ul style="list-style-type: none"> <li>■ File and directory names</li> <li>■ System output</li> <li>■ Code</li> <li>■ Text typed at the command-line</li> </ul>
<i>Monospace, italic font</i>	<ul style="list-style-type: none"> <li>■ Code variables</li> <li>■ Command-line variables</li> </ul>
<b>Monospace, bold font</b>	<ul style="list-style-type: none"> <li>■ Emphasis of file and directory names</li> <li>■ System output</li> <li>■ Code</li> <li>■ Text typed at the command-line</li> </ul>



**WARNING:** Indicates that failure to follow directions in the warning could result in bodily harm or death.



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

**Note:** Provides additional information.

## HP technical support

Telephone numbers for worldwide technical support are listed on the HP web site: <http://www.hp.com/support/>.

Collect the following information before calling:

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

For continuous quality improvement, calls may be recorded or monitored.

HP strongly recommends that customers sign up online using the Subscriber's choice web site: <http://www.hp.com/go/e-updates>.

- Subscribing to this service provides you with email updates on the latest product enhancements, newest versions of drivers, and firmware documentation updates as well as instant access to numerous other product resources.
- After signing up, you can quickly locate your products by selecting **Business support**, and then **Storage** under Product Category.

## HP authorized reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518.
- Elsewhere, visit <http://www.hp.com> and click **Contact HP** to find locations and telephone numbers.

## HP storage web site

For third-party product information, see the following vendor web sites:

- <http://www.hp.com>
- <http://www.hp.com/go/storage>
- <http://www.hp.com/support>
- <http://www.docs.hp.com>

# About HP Continuous Access EVA




This chapter provides an overview of replication features, and gives a brief description of required hardware components and software applications.

## Overview of HP Continuous Access EVA

HP StorageWorks Continuous Access EVA is the remote replication component of HP StorageWorks Enterprise Virtual Array (EVA) controller software. When this component is licensed, the controller copies data online, in real time, to a remote array over a local or extended storage area network (SAN). Properly configured, HP Continuous Access EVA is a disaster-tolerant storage solution that guarantees data integrity if an array or site fails.

[Figure 1](#) shows a basic configuration with redundant arrays and fabrics. One array is located at a local (or active) site and the other at a remote (or standby) site. In the figure, one fabric is called the black fabric and the other is called the gray fabric. Each array can perform primary data processing functions as a source, with data replication occurring on the destination array. The replication process can also be bidirectional, with some I/O streams moving to the array and other I/O streams moving simultaneously from the array. This feature allows the array to be the source for some data groups and the destination for others.



CXO8165c

- |                                   |                                      |
|-----------------------------------|--------------------------------------|
| ❶ Local active management server  | ❹ Remote standby management server 1 |
| ❷ Local host                      | ❺ Remote host                        |
| ❸ Local controller 1              | ❻ Remote controller 1                |
| ❹ Local black fabric switch       | ❼ Remote black fabric switch         |
| ❺ Local controller 2              | ❽ Remote controller 2                |
| ❻ Local gray fabric switch        | ❿ Remote gray fabric switch          |
| ❼ Local standby management server | ⓫ Remote standby management server 2 |
| ❽ Interswitch link—black fabric   | ⓬ Interswitch link—gray fabric       |

**Figure 1: Basic HP Continuous Access EVA configuration with redundant servers**

In [Figure 1](#), the management server represents the server where EVA management software is installed, including HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager. HP recommends at least two management servers at each site to avoid a single point of failure. If a significant failure occurs at the source array location, redundancy allows data processing to quickly resume at the destination array. This process is called failover. When the cause of the array failure has been resolved, data processing can be moved back to the original source array by performing another failover.

## Features

The following list describes HP Continuous Access EVA's more prominent features:

- Synchronous remote replication—ensures both source and destination copies are always identical and concurrent.
- Asynchronous remote replication—allows write completion back to the host for a faster host to storage I/O response time.
- Automated failover support
- Normal and failsafe data protection modes of operation
- Intersite link suspend-and-resume operations
- Bidirectional replication enables two sites in a remote replication connection to use each other to maintain synchronized copies of online data.
- Source and destination pair size of 1 GB to 2 TB in 1 GB increments
- Up to 128 DR groups per array
- Up to 128 remote source and destination pairs per array
- Up to 8 source and destination pairs per DR group

## What's new

This release introduces the following features:

- HP StorageWorks Replication Solutions Manager graphical interface, including multiple tree views and sortable lists of replication resources
- HP StorageWorks Replication Solutions Manager host agents, with the ability to mount volumes and run jobs directly on storage hosts
- Job language and job templates to automate replication tasks
- HP StorageWorks Replication Solutions Manager Command Line User Interface

## Setup and configuration assumptions

Before you can perform replication, the hardware and software components must be installed and a license activated. This guide assumes that you have the following components installed and configured.

### Licensing

To perform remote replication, an HP StorageWorks array must have a valid Continuous Access replication license (also known as a license-to-use, or LTU). To perform local replication—snapshots and snapclones—the array must have a Business Copy replication license. See the QuickSpecs for information on ordering a license. To activate licenses, see the documentation that comes with the license agreements.

### Hardware and software components

The following sections describe the hardware and software components that must be installed before you can perform replication. See the *HP StorageWorks Continuous Access EVA planning guide* for more information.

### Enterprise Virtual Array

You must have at least two configured EVA arrays: one at a primary site and one at an alternate site a safe distance away.

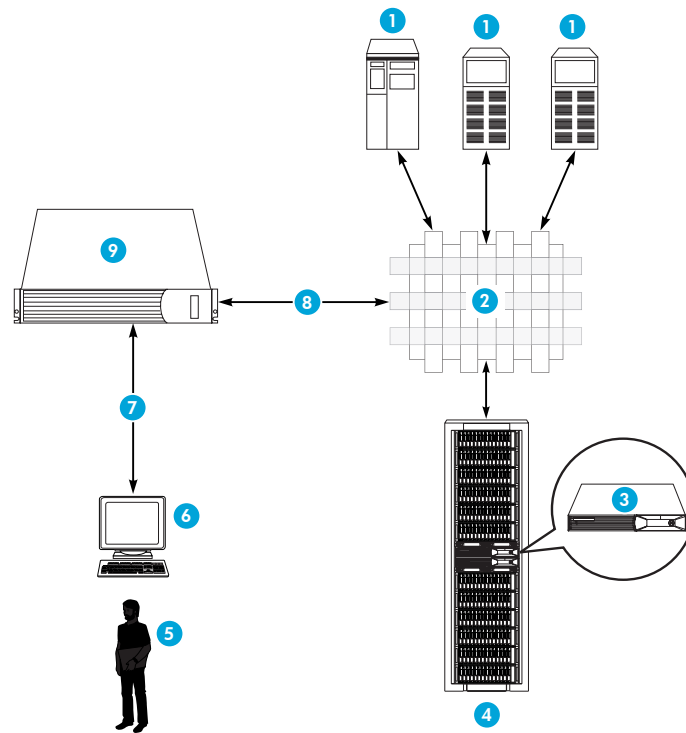
### Hosts and host operating systems

Use a supported operating system. See the *HP StorageWorks EVA replication compatibility reference* for a list of supported host operating systems and versions.

### Virtual Controller Software

HP Continuous Access EVA is enabled in the Virtual Controller Software on redundant EVA arrays. The controller software provides the functionality for the array controller. [Figure 2](#) on page 17 illustrates the role of the controller software in a single array. For additional information, see the *HP StorageWorks Enterprise Virtual Array user guide*.





CXO8166b

- |                               |   |
|-------------------------------|---|
| ① Host                        | ⑥ Browser   |
| ② Fabric                      | ⑦ Control input monitoring output                 |
| ③ Virtual Controller Software | ⑧ Control and monitor commands                    |
| ④ Array                       | ⑨ HP Command View EVA and the replication manager |
| ⑤ Administrator               |   |

**Figure 2: Functional relationship between controllers and server**

## Zoning

Proper zoning ensures that hosts in an HP Continuous Access EVA environment do not conflict with each other. See the *HP StorageWorks Continuous Access EVA planning guide* for more information.

## Management component

Two management applications are required for HP Continuous Access EVA: HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager. You must install and run HP Command View EVA on one of the following server types:

- Dedicated server—Runs only storage EVA software.
- Storage Management Appliance (SMA)—Runs software applications at a centralized location for managing and monitoring your storage. The SMA arrives preloaded with SMA software upon which other management applications can be loaded. HP StorageWorks Command View EVA and HP StorageWorks Element Manager for HSG software (preinstalled on the SMA), provide physical and logical views of your array through a graphical user interface.
- General purpose server—Runs other applications in addition to storage EVA software.

---

**Note:** In this guide, the term management server applies to any of the server types above.

---

## Interface options

Performing replication can involve one or more of the interfaces and products described in [Table 2](#).

**Table 2: Other interfaces and products that perform replication**

HP product	Interface	Remarks
HP StorageWorks Command View EVA	Browser-based graphical user interface (GUI)	<ul style="list-style-type: none"> <li>■ Required.</li> <li>■ Creates snapshots and snapclones from a GUI.</li> <li>■ Replicates by specifying an array and virtual disk; cannot replicate by specifying a host and host volume.</li> <li>■ Cannot perform dynamic mounting or interactions with hosts.</li> <li>■ Does not provide jobs, job templates, or scripting capabilities.</li> <li>■ See “<a href="#">HP Command View EVA</a>” on page 20 for more information.</li> </ul>
HP StorageWorks Replication Solutions Manager	Browser-based GUI and command line user interface (CLUI)	<ul style="list-style-type: none"> <li>■ Optional.</li> <li>■ Specialized storage replication software that supports both local replication and remote replication.</li> <li>■ Creates snapshots and snapclones using a GUI, jobs, and a CLUI.</li> <li>■ Replicates by specifying an array and virtual disk (LUN) or by specifying a host and host volume.</li> <li>■ Performs dynamic mounting and interactions with hosts.</li> <li>■ Includes integrated job editor, job templates and job scripting capabilities.</li> <li>■ Includes integrated job management.</li> </ul>

**Table 2: Other interfaces and products that perform replication**

HP product	Interface	Remarks
HP StorageWorks Storage System Scripting Utility (SSSU)	Host command line	<ul style="list-style-type: none"> <li>■ Optional.</li> <li>■ EVA host platform software that creates snapshots and snapclones from a command line or custom script.</li> <li>■ See snapshot and copy commands in the EVA SSSU reference guide.</li> <li>■ Replicates by specifying an array and virtual disk; cannot replicate by specifying host and host volume.</li> <li>■ Dynamic mounting and host interactions is accomplished by writing custom scripts.</li> </ul>
HP StorageWorks SMI-S Interface for Command View EVA	WEBM client-server using XML	<ul style="list-style-type: none"> <li>■ Optional.</li> <li>■ Provides an SMI-S compliant interface for HP StorageWorks Command View EVA.</li> </ul>

## HP Command View EVA

HP Command View EVA is a user interface that communicates with the EVA controllers to control and monitor the storage. HP Command View EVA maintains a database for each managed array and the database resides on that array. Instructions for its use can be found in the *HP StorageWorks Command View EVA Online Help*.

To use HP Continuous Access EVA you must first use HP Command View EVA to:

- Add licenses
- Initialize controllers—This process binds the controllers together as an operational pair and establishes preliminary data structures on the disk array.
- Create disk groups—A disk group is a set of physical disks from which storage pools are created. When your array is initialized, one default disk group is created. For the highest performance and availability, a disk group should only contain one disk drive model. If you are performing bidirectional replication, the disk groups should be symmetric with respect to the capacity on both arrays.
- Create hosts—The host connects to a fabric through a Fibre Channel adapter (FCA) and accesses storage through the controllers. Hosts contain a pair of FCA ports to connect with each fabric. Before a host can access any storage, it must be discovered by the array.

- Create virtual disks—Variable disk capacity that is defined and managed by the array controller and presentable to hosts as a disk. You can assign a combination of characteristics to a virtual disk, such as a name, redundancy level, size, and other performance characteristics.
- Present virtual disks to hosts—Assigning a virtual disk to a host results in a host presentation. You may present a virtual disk to a host during the virtual disk creation process, or wait until a later time. However, the virtual disk must be presented to a host in order to use it for replication.

You can also use HP Command View EVA to create snapshots and snapclones.

## Previous products

Table 3 describes previous products.

**Table 3: Previous products**

HP product	Interface	Remarks
HP StorageWorks Business Copy EVA/MA/EMA	Browser-based GUI and host command line	<ul style="list-style-type: none"> <li>■ Optional.</li> <li>■ Specialized storage replication software that creates snapshots and snapclones using Business Copy jobs.</li> <li>■ Replicates by specifying an array and virtual disk (LUN) or by specifying a host and host volume.</li> <li>■ Performs dynamic mounting and interactions with hosts.</li> <li>■ Includes integrated job editor, job templates, and job scripting capabilities.</li> <li>■ Includes integrated job management and scheduling.</li> <li>■ Also provides similar local replication features on supported MA/EMA arrays.</li> </ul>
Continuous Access User Interface	Browser-based GUI	<ul style="list-style-type: none"> <li>■ Optional.</li> <li>■ Specialized storage replication software that supports remote replication.</li> <li>■ Replicates by specifying an array and virtual disk (LUN) or by specifying a host and host volume.</li> <li>■ Performs dynamic mounting and interactions with hosts.</li> </ul>

# Concepts

## 2

This section describes some basic terminology and concepts you need to understand before replicating data.

### Remote data replication

The array at the active location is connected to a partner array at the standby location. To replicate data remotely, a source virtual disk is configured at the active array. When data replication is selected, the destination virtual disk (called a remote copy) is automatically created by software at the standby array. Any data written to the source virtual disk is then copied to the destination virtual disk. Applications continue to run while data replicates in the background over a separate connection.

### Bidirectional replication

When an array contains both source virtual disks and destination virtual disks, it is said to be *bidirectional*. An array can have a bidirectional data replication relationship with up to two other arrays; and an individual virtual disk can have a unidirectional replicating relationship with only one other virtual disk.

Use remote copies for disaster recovery, data migration, and data distribution.

### Local replication

Local replication is a licensed feature of HP StorageWorks arrays that allows you to quickly create local, point-in-time copies of your data. These copies are known as snapshots and snapclones.

In a typical environment, a server is running HP Command View EVA and a local replication automation product, for example Replication Solutions Manager (hereafter called the replication manager).

The replication server is connected by LAN to multiple hosts running the replication manager host agents (hereafter called enabled hosts). These enabled hosts are running production database and backup applications which perform I/O with multiple storage arrays in the SAN. An operator or administrator automates operations by running the replication manager jobs from a browsing computer, a scheduler, or from a host using the replication manger CLUI.

In this environment, the replication manager jobs can perform simultaneous, non-disruptive tape backups of the databases. To run daily backups, for example, you could create two jobs, each performing backups with their respective storage arrays and hosts.

You can also perform local replication using HP Command View EVA or any of the other supported HP Business Copy EVA interfaces.

## Snapclone

Snapclone replication instantly creates a copy of a virtual disk that begins as a fully allocated snapshot and then becomes an independent virtual disk.

## Snapshot

Snapshot replication instantly creates a demand allocated or fully allocated point-in-time copy of a source virtual disk. A demand allocated snapshot is a virtual copy in which the allocated disk space can change on demand from an initial minimum amount, up to the full capacity of the source. A fully allocated snapshot is a virtual copy in which the allocated disk space is initially set to, and remains fixed at, the full capacity of the source.



---

## DR groups

A data replication (DR) group is a named group of virtual disks selected from one or more disk groups so that they remotely replicate to the same destination, fail over together, share a log (DR group log disk), and preserve write order within the group.

---

**Note:** To use this feature with specific source and destination arrays, each array must have its own Continuous Access replication license.

---

---

### Notes:

- Only virtual disks eligible for remote replication can belong to a DR group.
  - A DR group can contain one or more virtual disks. For optimal performance, however, limit each DR group to one virtual disk.
  - You can create up to 128 DR groups per storage array.
- 


## Source and destination DR groups

DR groups are always created in pairs, consisting of a source DR group and a destination DR group. The source DR group contains the virtual disk you want to replicate remotely. This is the source virtual disk. The destination virtual disk (called the remote copy) resides in the destination DR group.

## Replication direction

The replication direction of a DR group is always from a source to a destination. When replicating from a source to a destination, the DR Group is in an original state. When replication occurs from an array that was created as a destination (for example, after a failover), the DR group is failed over.

[Figure 3](#) depicts the replication of one DR group between separate sites.



CX07989b

- |                      |  |
|----------------------|--|
| ❶ Host               | ❹ Array B  |
| ❷ Switch             | ❺ Source virtual disk  |
| ❸ Host I/O           | ❻ Destination virtual disk   |
| ❹ Replication writes | ❼ DR group containing a pair of source and destination virtual disks |
| ❺ Array A            |  |

**Figure 3: HP Continuous Access EVA DR group replication**

## Home designation

Home is the preferred source DR group in a remote replication relationship. Having a preferred source allows you to designate a point of reference for remote replication. Home designation is an HP Continuous Access EVA concept only.

## DR group presentation

All members of a DR group are presented to a host through the same controller and move together. For optimal performance, limit each DR group to one virtual disk.

---

**Note:** Use the same preferred path for all DR group members, with presentation to same Fibre Channel adapters (FCAs).

---

## Presenting DR group members to the same FCA

All members of a DR group must be presented to the same FCA on hosts with more than one FCA per fabric (for example, multiple FCA pairs, multiple dual-channel FCAs, or a combination of single and dual-channel FCAs). All DR group members must also be preferred to the same controller with the same failover characteristics.

This restriction is required to keep the DR group members using the same host FCA to EVA path. In the event of a path or controller failure, the members collectively fail over to the other path, thus preserving write order across the members of the DR group.

---

**Note:** Additional members added to a DR group will have the parameters of the original member. This can affect multipathing of OS applications.

---

## DR group properties

Every DR group has defined properties such as name, operational state, and so on. See the online help for an explanation of the properties.

## DR group log

The DR group log is a designated virtual disk that stores a source DR group's host writes while remote replication to the destination DR group is stopped or if the writes to the source side are faster than transfer rate to the destination side of the DR group. When replication is re-established, the contents of the log are written to the destination virtual disks within the destination DR group to synchronize the destinations with their sources. This process of writing the log disk contents to the destination in the order that the writes occurred is called *merging*.

Once a log reaches full capacity, it is often more practical to copy the changed blocks on the source virtual disk to the destination virtual disk. This copy operation is called a fast synchronization—all 8-MB block increments written on a source virtual disk since it lost connection with the destination are copied to the destination virtual disk.

At other times, a full copy copies the complete source virtual disk to the destination virtual disk. You cannot manually force fast synchronization or a full copy, rather it is an automatic process that occurs when a log has reached full capacity.

## DR group log states

A DR group log can be in one of the following states:

- Unused (Normal)—No source virtual disk is logging or merging.
- Logging—At least one source virtual disk is logging but none are merging.
- Merging—At least one source virtual disk is merging and logging.

## DR group log size

When a DR group is created, a log disk consists of 136 MB of Vraid1 space; however, the array can automatically allocate more space on demand. When a DR group is logging, the log disk grows in proportion to the amount of writes to the source virtual disk.

When creating disk groups and the virtual disks within them, ensure that sufficient space remains for DR group log disks to expand to their maximum size. HP recommends creating DR group log disks in near-online disk groups, if available. Otherwise, create log disks in the online disk groups with the most free space.

Array software considers the log disk full when any of the following conditions occurs:

- The log disk is twice the size of the DR group's virtual disks.
- No free space remains in the disk group.
- The log disk reaches 2 TB of Vraid1 (4 TB total).

When the log disk is declared full, DR group members are marked for full copy and the log disk is deleted.

## Managed sets

A managed set is a named collection of resources banded together for the purpose of management. For example, the managed set `Sales_Disks` might include two virtual disks, `West_Sales` and `East_Sales`.

Performing an action on a managed set, performs the action on all the members in the set. For example, if you perform the `New Snapshot` action on the managed set `Sales_Disks`, the interface creates a snapshot of `West_Sales` and a snapshot of `East_Sales`.

---

**Note:** The order in which members are added to the group has no effect on the order in which actions are performed on the members.

---

## Members

A managed set can comprise the following types of resource: DR groups, enabled hosts, host volumes, storage systems, or virtual disks.

## Managed set properties

A managed set is defined by its properties, such as name and type. Use the replication manager to view the properties.

## Managed set actions

Each type of managed set supports common actions, plus resource-specific actions (see `List` tab actions and `Members` tab actions). Use the `Managed Sets List` and `Tree` tabs to view the managed sets and members managed by the replication manager, and to perform actions.

## Considerations

### General considerations

- All resources, or members, in a single managed set must be of the same type (for example, all virtual disks).
- You can add a specific resource to more than one managed set.

### Storage resource considerations

You can add resources on more than one storage system to a managed set.

### DR group considerations

- Source and destination sides of different DR groups can be part of the same managed set. However, both the source and destination sides of the same DR Group cannot be part of the same managed set.
- Some DR group actions are permitted only on source DR groups; other actions are only permitted on destination DR groups. See the DR group actions table for more information.
- Create separate managed sets for source and destination DR groups so that if a failover occurs, you can perform the actions that correspond to the new identity of the managed set.
- If you plan to use DR group managed sets for failover operations, ensure the managed sets are controlled by the same server at the time of failover.

---

## Failover

Failover is an operation to reverse replication direction. When you initially set up data replication, you replicate from a source to a destination. Failover simply changes the direction of replication; the destination array assumes the role of the source and the source assumes the role of the destination. For example, if a DR group is replicating from array Alpha to array Bravo, a failover operation would change the direction of the replication so that data from array Bravo would be replicated to array Alpha.

You can fail over:

- DR groups—Reverses the replication direction of a DR group.
- Managed sets containing DR Groups—Reverses the replication direction of all DR groups within the managed set.

[Figure 4](#) shows a data replication relationship among DR groups at three locations. Arrays A and D are located at the primary site and arrays B and C are located at two alternate sites. If contact with the primary site is broken, failover can manually occur with the destination DR groups that were replicating to array B. Array B then acts as the primary site for these DR groups after failover. At array C, the source DR groups begin logging until new remote sites are re-established.

## Considerations

- Failing over is only permitted on a destination DR group.
- If only one component fails, repairing that single component may be preferable to performing a complete site failover.
- Do not fail over a DR group more frequently than once every 15 minutes.

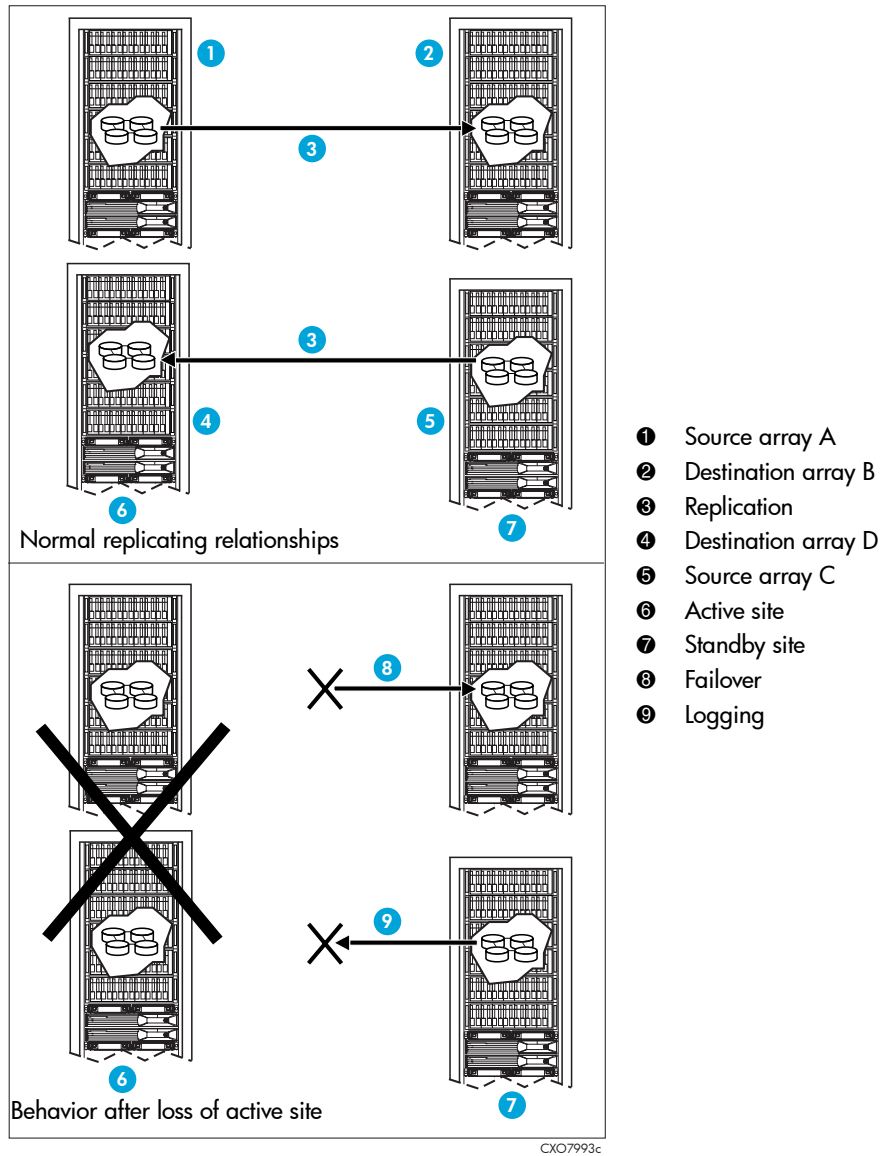


Figure 4: Replicating relationships among DR groups



## Replication manager logs

The following logs track and record activity in the replication manager:

- Event log—Contains system-generated messages resulting from:
  - User-initiated actions (for example, “suspend DR group”).
  - Replication-related storage events (for example, “DR group is constructing”).
  - Jobs (for example, “job complete”).

Messages are written to a set of five rotating files. The default size limit for each of these files is 1 MB. As the current file reaches its size limit, it is closed, rotated out, and a new file opened. You can view the messages in the Event pane of the replication manager. See the online help for more information. See the Appendix on page 137 for event descriptions.

- Trace log—Contains all events. Intended for HP personnel. Events that are useful to the user are transferred to the Event log and displayed in the Event pane. The Trace log is allotted 60 MB of space. As the log fills the space, old messages are discarded. You can view the Trace log in the Configuration window in the replication manager.
- Security log—Contains the following security activities:
  - Successful attempts to access the replication manager
  - Failed attempts to access the replication manager
  - Failed authorization credentials
  - Changes to security user accounts and group membership

Thirty days of history are saved by default. On the thirty-first day, the oldest entries are discarded. You can change the default to any number of days. HP recommends you save no more than 30 days. The file to change is called `security.cfg` and is located on the replication manager server in `C:\Program Files\Hewlett-Packard\sanmgr\security\config\`.

Security logs are `.txt` files located on the replication manager server in `C:\Program Files\Hewlett-Packard\sanmgr\security\logs`.

- Transaction log—Contains the replication manager database contents. Service personnel can access this log to recover data if a database is corrupted. No user action is needed or allowed on this log. Space needs vary for the log depending on database activity. However, space needs are generally less than 20 MB.

## Failsafe mode

The failsafe mode specifies how host writes and remote replication behave when a group member fails. The failsafe mode can be either:

- Failsafe enabled—If any virtual disk within the DR group fails or becomes unreachable, all host writes and remote replication automatically stop. This preserves the order of the replicated data. A failsafe-enabled DR group can be in one of two states:
  - Unlocked (failsafe-unlocked)—Host writes and remote replication occur.
  - Locked (failsafe-locked)—Host writes and remote replication automatically stop.
- Failsafe disabled—If any destination virtual disk (remote copy) within the DR group fails or becomes unreachable, all host writes to the source DR group continue, but all remote replication to the destination DR group automatically stops; the source DR group logs its host writes to the DR group log until remote replication is re-established. If a source virtual disk fails, host writes to the failed disk stop, as well as remote replication to its remote copy; host writes and remote replication to the other members of the DR group continue normally.

# Replication tasks

## 3

This chapter assumes that you have set up and configured your arrays' hardware and software, and that you are ready to perform replication tasks. This chapter discusses the HP Replication Solutions Manager—how to access it, what tasks you can perform with it, and how to use the alternative command line user interface instead of the replication manager, if desired. It also discusses the replication manager jobs feature.

## About HP Replication Solutions Manager

After you create hosts, virtual disks and their presentations, and your hosts can access your virtual disks, you are ready to replicate data using HP Replication Solutions Manager. HP Replication Solutions Manager (hereafter called the replication manager) is an interactive, visual environment for managing data replication. For installation instructions, see the *HP StorageWorks Replication Solutions Manager installation guide*. Operational information is in the online help system.

HP recommends that you use the replication manager to:


- Create and delete DR groups.
- Change properties for DR groups.
- Add and remove members from DR groups.
- Failover, suspend, resume, and change failsafe mode by DR group.
- Create and delete managed sets.
- Failover, suspend, resume, enable and change failsafe mode, and revert to Home by managed set.
- Direct commands to the appropriate array regardless of selected array.
- Actively monitor arrays (copy, failsafe, logging, merging).
- Create and restore configuration databases for managed sets and jobs.

## Starting the replication manager

Methods of starting the application include:

- From an icon placed on the desktop during installation.  
Use this method if you are logged on the management server directly or through Terminal Services. Do not browse to the replication manager from the management server.
- From an authorized client specified during installation or added later (see [Adding remote access IP addresses](#), page 39).

## From the replication manager icon (application mode)

1. Log on to the management server. Use Terminal Services to log on remotely.
2. Double-click the replication manager icon  on the server desktop.

A DOS window is displayed while the replication manager starts.

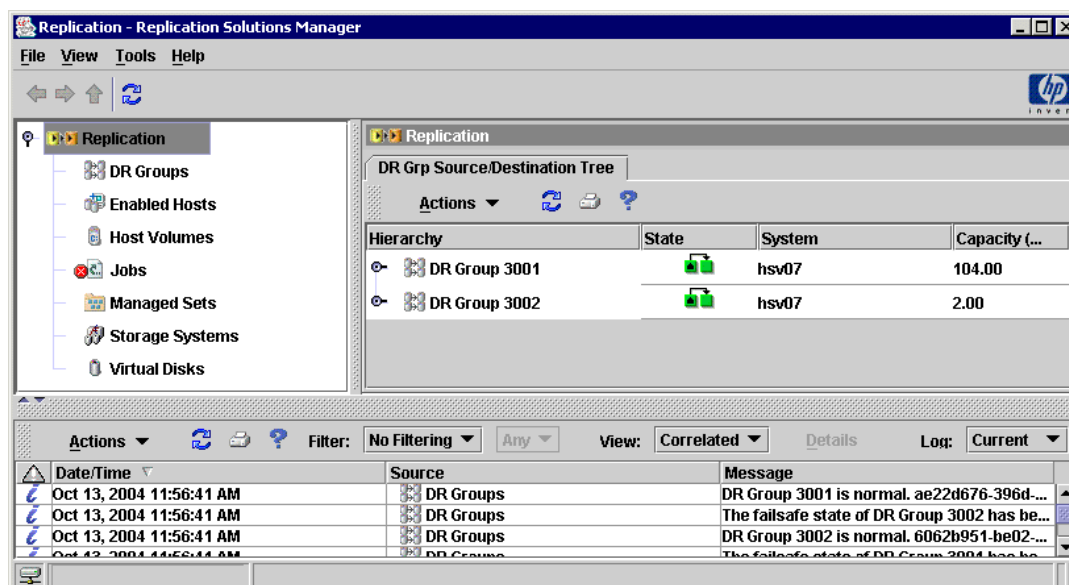
---

**Tip:** An alternative to double-clicking the replication manager icon is to double-click the file: `C:\Program Files\Hewlett-Packard\sanmgr\commandview\server\browser\start_gui.bat`.

---

3. Log on to the replication manager.
4. Enter the user name (originally **admin**) and password (originally **nimda**).

The Replication Solutions Manager window opens.



## From an authorized client (applet mode)

An authorized client is any computer whose IP address is entered in the management server's access list. If you did not enter the IP address during installation, you can add it now. See “[Adding remote access IP addresses](#)” on page 39.

Using a browser on an authorized client, you can start the replication manager:

- By direct access.
- Through the SMA software interface.

Before using either method, verify that supported JRE and browser versions are installed (see *HP StorageWorks EVA replication compatibility reference*). Supported JREs and setup instructions are provided on the HP StorageWorks JREserver CD (in the product kit).

---

**Note:** If you browse away from the replication manager, you must close and reopen the application or press SHIFT+F5 to re-establish the connection.

---

### Browsing directly to the replication manager

1. If you installed a new JRE, close all open browser sessions on the client. Closing and restarting the browser allows a new Java version to be recognized.
2. Open a Web browser and browse to the replication manager on the management server:

```
http://<MyServerName_or_IP_Address>:4096
```

3. Log on to the replication manager.  
Enter the user name (originally **admin**) and password (originally **nimda**).

### Browsing to the replication manager through the SMA

1. If you installed a new JRE, close all open browser sessions on the client. Closing and restarting the browser allows a new Java version to be recognized.
2. Open a Web browser and browse to the SMA:  

```
//http:<MyServerName_or_IP_Address>
```
3. Log on to the SMA (see *HP OpenView Storage Management Appliance User Guide* for login information).

4. Select **Settings > Manage Tools**.  
Verify that the replication manager is installed. (The version is “unknown”).
5. Select **Tools > replication solutions manager**.
6. Log on to the replication manager.  
Enter the user name (originally **admin**) and password (originally **nimda**).

## Adding remote access IP addresses

Use the following procedure to add remote access IP addresses. Host names *cannot* be added to this list in place of IP addresses. This procedure requires administrator-level access.



**Caution:** If your SAN security is high, HP recommends adding an entry that allows all clients to access the server, such as \*.\*.\*.\*. *Do not* add a general entry if you think your SAN is vulnerable to insecure access.

---

1. Access the replication manager server and log on as an administrator.
2. Navigate to the following file:
  - On management servers without HP Storage Area Manager:

```
C:\Program Files\Hewlett-Packard\sanmgr\hostagent\config\access.dat
```
  - On management servers with HP Storage Area Manager installed:

```
C:\Program Files\Hewlett-Packard\sanmgr\managementserver\config\authorizedClients.dat
```
3. Open the file using a text editor.
4. Review all IP addresses listed.



**Caution:** Both server and client IP address are listed. Modification or removal of a server IP address prevents communication between the associated replication manager server and host agents.

---

5. Modify or remove existing IP addresses or add new addresses.

**Note:** Host names *cannot* be added to this list in place of IP addresses.

---

## Monitoring events with the replication manager

An event is a system-generated status message, resulting from a:

- User-initiated action (for example, “suspend DR group”).
- Replication-related or array incident (for example, “retrieved data for array”).
- Job (for example, “job complete”).

Event message categories include:

- DR group
- Storage system
- Job
- Error event
- Dialog box

Event messages can be helpful when troubleshooting. Messages are collected by monitoring the messages written to the Trace log files. Whenever a relevant message is written, it is posted in the replication manager Event pane. The Event pane provides a description of the event codes, including the severity, the date and time when the event occurred, the source component of the event, and an explanation.

You can perform several event handling actions, such as:

- Sort and filter columns.
- Display events in a chronological order (standard view) or display the latest event for each resource (correlated view).
- View historical logs.
- Connect to a resource in the navigation pane by clicking the corresponding source entry in the Event pane.

The Event pane periodically polls the server for messages to display. You can also manually refresh content in the pane. See the online help for more information on using the event pane.

See page 137 for an explanation of the event codes.

## Data replication using the replication manager

After the replication manager is installed and initialized, you can replicate data if you have the appropriate replication license.



## Remote replication tasks

You need an HP Continuous Access EVA license to create DR groups.

## Local replication tasks

You need a Business Copy EVA license for the following functions:

- Create a snapshot—Creates a demand-allocated (space efficient) or fully allocated (space inefficient) point-in-time copy of a source virtual disk.
- Create a snapclone—Creates a copy of a virtual disk that begins as a fully allocated snapshot and then becomes an independent virtual disk

See the online help for instructions.

## Managing storage with multiple management servers

An array can have only one instance of HP Command View EVA managing it at a time. If there are multiple management servers with HP Command View EVA installed within the same SAN, one management server has the ability to acquire the control of any array in the management zone of the SAN. After HP Command View EVA has control of the array, Replication Solutions Manager can manage it.

To ensure disaster tolerance, you must have at least two management servers in your configuration. HP recommends that you have at least one management server at each site. The active management server actively manages the array, while all other standby management servers take control of the array if there is a failure on the active server. You can also use multiple management servers to divide the management responsibilities so that each server controls different arrays. However, the server must control both arrays when you create DR groups.

## Considerations when using multiple management servers to manage storage

Consider the following when using multiple management servers:

- Synchronize the clocks of the servers, taking time zones into account. This is important for event and log time stamps (see “Synchronizing time on the servers” on page 52).
- Decide which server will manage each array and which server can be used to take control of that array, if needed (see “Moving storage management to another server” on page 50). A single server must manage both arrays when creating DR groups. Later, you can use separate servers to monitor the state of the arrays.
- Know which applications are installed and running on each server. If another server needs to control an array, that server needs the same applications running to allow the same functionality. Applications can also be stopped and started on a server (see the *HP StorageWorks Replication Solutions Manager installation guide* for information on stopping and restarting servers).
- Practice server management role changes regularly. This allows you to verify that zones, clock times, installed applications, users, groups, enabled hosts, and storage configurations are correct.

## Creating the same configuration on each standby server

This procedure allows you to create the same configuration on each standby server by exporting the database from the active server and importing it to each standby server. This procedure has two parts:

- Exporting and importing the database—Saves jobs, managed sets, and enabled hosts data from an active to a standby server.
- Importing user names and groups—Saves user accounts and groups data from an active server to a standby server.

---

**Note:** Importing a database also imports the password from the active management server. If that password is different than the password on the standby management server, the DR groups are displayed in an unknown state and no action can be taken on them.

---

---

**Note:** Before you import a database to a standby server, an HP Command View EVA instance must have control of the arrays presented to the replication manager on the standby server.


---

## Exporting and importing the database

To create the same configuration on each standby server:

1. On the active server, access the replication manager.
2. Select **Tools > Export RSM V1.x Database**.


The Export RSM V1.x Database window opens.



3. Enter a local directory for the location of the exported database.

4. Copy the exported database from the active management server to a local drive on the standby management server.
5. Using Command View EVA on the standby server, take control of the arrays.
6. Access the replication manager.
7. Select **Tools > Import RSM V1.x Database**.

The Import RSM V1.x Database window opens.



8. Enter the location and name of the exported database that you copied from the active management server.
9. Click **OK**.

---

**Note:** If a resource already exists with the same name, the imported name of the resource is named *resourcename\_n*, where *n* is a number starting with 1.

---

---

**Note:** Any time the configuration changes regarding user-defined jobs, managed sets, and enabled hosts, you must perform this procedure again.

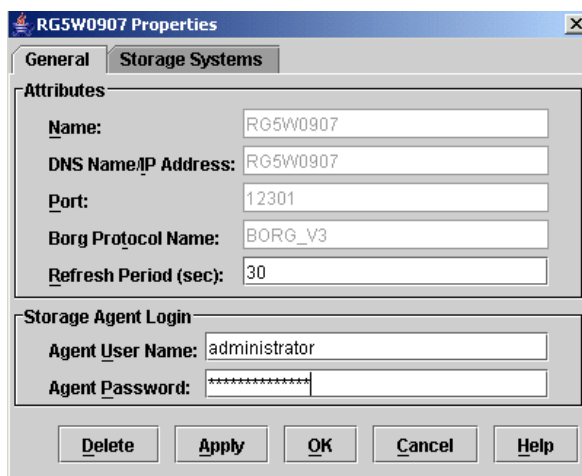
---

## Changing the password for the management server

Passwords for the active and standby management servers must match in order for the database import to succeed. To change the password for a management server:

1. In the replication manager, select **Tools > Configure > Storage Access > Management Server**.
2. Select the management server whose password you want to change.
3. Click **Properties**.

The Properties window opens, displaying the General tab.



The screenshot shows a dialog box titled "RG5W0907 Properties" with a close button (X) in the top right corner. The dialog has two tabs: "General" (selected) and "Storage Systems". Under the "General" tab, there are two sections: "Attributes" and "Storage Agent Login".

**Attributes:**

- Name: RG5W0907
- DNS Name/IP Address: RG5W0907
- Port: 12301
- Borg Protocol Name: BORG\_V3
- Refresh Period (sec): 30

**Storage Agent Login:**

- Agent User Name: administrator
- Agent Password: [masked with asterisks]

At the bottom of the dialog, there are five buttons: "Delete", "Apply", "OK", "Cancel", and "Help".

4. In the Agent Password box, under the Storage Agent Login area, enter the correct password.
5. Click **OK**.


## Importing user accounts and group data

The same data is imported, whether you import user accounts or groups.

To import user accounts and group data:

1. On the standby server, select  
**Tools > Configure > Security > Users > Import**  
or  
**Tools > Configure > Security > Groups > Import.**

The Import User Accounts window opens.



2. Enter either the server name or IP address of the active server (where the user accounts and groups are located), and click **OK**.
3. Enter the user name and password for the active server.
4. Click **OK**.

---

**Note:** Any time the configuration changes regarding user names or groups, you must perform this procedure again.

---

---

## Migrating data from HP Continuous Access user interface to the replication manager

Migrating an HP Continuous Access user interface (hereafter called HP CA UI) database to the HP Replication Solutions Manager allows you to keep the managed sets you created in HP CA UI. Use the back up database feature in HP CA UI and the import feature in the replication manager.

In HP CA UI, you could add the source and destination side of a DR group into a managed set. However, in the replication manager, you can add only one side (the source or destination). When migrating, the replication manager updates the managed sets to include whichever side you added first to the managed set in HP CA UI.

---

**Note:** Before you import a database into the replication manager, an HP Command View EVA instance must control the arrays presented to the replication manager.

---

1. In HP CA UI, back up the database (see the HP CA UI online help).

---

**Note:** If possible, save the database on the same server where the replication manager will be installed; otherwise, you must move it later.

---

2. Uninstall HP CA UI and install the replication manager.

---


**Note:** For installation instructions, see *HP StorageWorks Replication Solutions Manager installation guide*.

---

3. Using HP Command View EVA on the server where the replication manager resides, take control of the arrays.
4. Access the replication manager.
5. Select **Tools > Import CAUI V1.x Database**.

The Import legacy CAUI Database window opens.

6. Enter the path and file name of the database you want to import.



7. Click **OK**.



## Upgrading array firmware

Before installing a new firmware version on arrays that are licensed for remote replication, you must:

- Be running required prior versions of Virtual Controller Software (VCS) and HP Command View EVA (see *HP StorageWorks EVA replication compatibility reference*).
- Not be running more than two VCS versions (one of which must be the new version) in multiple relationships between the source and destination arrays.
- Suspend DR group replication for the duration of the code load, even if the intersite links are unavailable.
- Perform the code load under conditions of no leveling and no logging.

A code load of one controller forces both controllers on a storage array to reboot. If the controller being upgraded is used as a destination for replication I/O, the source controller logs new I/O while the destination controller reboots. Resume the replication to allow the controller to perform a merge. Wait for the merge to conclude after a code load, or wait three minutes—whichever is longer—before performing the code load on the other array in the relationship.

All arrays involved in a data replication relationship must be running with fully functional controllers during a code load. Only two VCS versions are supported at any one time when upgrading a multiple array relationship. For example, three arrays in a relationship cannot have three different VCS versions. If one is running VCS v3.010 and two are running v3.014, upgrade the v3.010 to v3.014, and then start upgrading all to v3.020. HP does not support different VCS versions among controllers in a replicating relationship for a duration longer than a week (168 hours).

---

**Note:** Some high-performance applications with low time-out thresholds may time out during the code load process. HP recommends that you perform VCS upgrades during periods of minimal activity.

---

## Moving storage management to another server

The following procedure describes how to move management of arrays from one server to another server.

1. Log on to HP Command View EVA on the server that you want to use to control your array.
2. Click **Discover**, and then **OK**.

The storage icons displayed in the navigation pane are gray to indicate that another server is managing the storage.


3. In the navigation pane, select a storage icon.

A page similar to the following is displayed to inform you that another instance of HP Command View EVA is managing the arrays.



4. Click **OK**.

A message similar to the following indicates that you are about to assume control of the storage with another server.



5. Click **OK**.

The storage icon in the navigation pane changes to green to indicate that the server you are logged on to has control of this array. The color change may take a few minutes.

6. Repeat steps 3 through 5 for the remaining gray storage icons for which you want to take control.
7. Start the replication manager (see “[About HP Replication Solutions Manager](#)” on page 36).

You are now able to manage your storage in the replication manager.

## Synchronizing time on the servers

It is important to synchronize the time clocks of your servers as closely as possible, keeping time zone differences in mind. This allows you to view event times in log files from multiple servers with a consistent time reference. This is especially important when comparing event logs of both arrays in a replicating relationship.

- To set the time on a management server, see the operating system-specific instructions.
- To synchronize SMAs to a Network Time Protocol (NTP) server, see the *HP OpenView Storage Management Appliance Software User Guide* at <http://h18000.www1.hp.com/products/sanworks/managementappliance/documentation.html#softv21>
- To set an array's time, see "Setting array time" on page 53.

## Setting array time

To set the time of an array to an individual server:

1. Log on to the server managing the desired array.
2. Log on to HP Command View EVA.
3. Select the desired array.

The Initialized Storage System Properties page is displayed.

The screenshot shows the HP Command View EVA interface. The top navigation bar includes 'Root View', 'Agent Options', and 'Help'. The main content area is titled 'Initialized Storage System Properties' and contains several sections:

- Buttons:** Save changes, Set options, View events, Uninitialize, Code load, Shut down, and a help icon (?).
- Identification:**
  - Name: HSV07
  - Node World Wide Name: 5000-1FE1-0015-31E0
  - UUID: 6005-08b4-0001-446f-0003-9000-0382-0000
- Condition/State:**
  - Operational state:  Good (Initialized)
- System:**
  - Type: HSV110
  - Version: 3010
  - Console LUN ID: 0

4. Click **Set Options**.

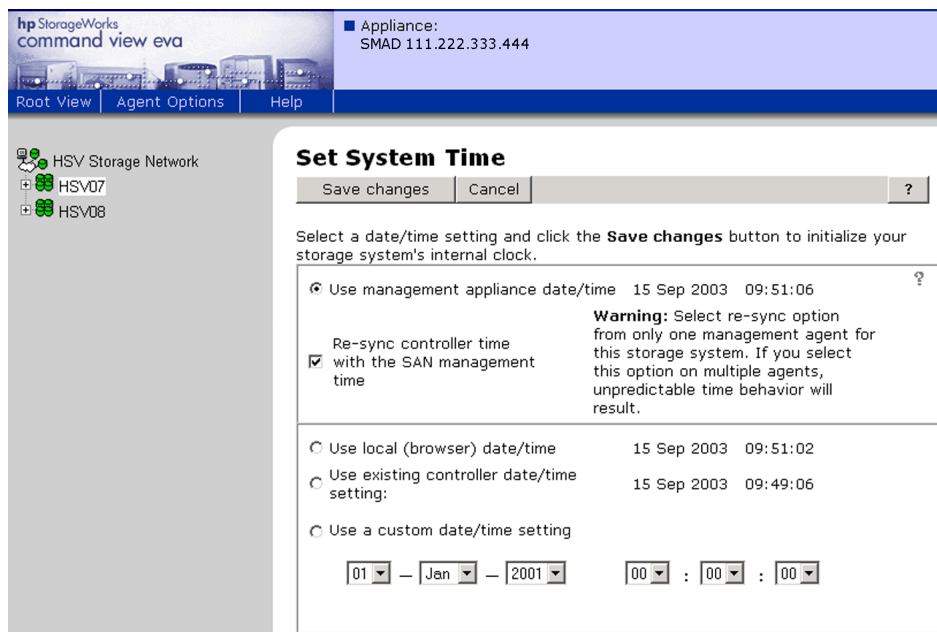
The System Options page is displayed.

The screenshot shows the 'System Options' page in the HP Command View EVA interface. It features an 'OK' button and four unchecked checkboxes:

- Configure event notification
- Configure host notification
- Set system operational policies
- Set time options

5. Click **Set time options**.

Various options for setting time are displayed on the Set System Time page.



**Figure 5: Set System Time page**

6. Select **Use management appliance date/time**.
7. Select **Re-sync controller time with the SAN management time**.

---

**Note:** Select this option from one management server only.

---

8. Click **Save changes**.

## Jobs

A job is a repeatable, replication-specific task or set of tasks that automates replication tasks. A job can be simple (for example, create a DR group) or complex (for example, perform cascaded replication). Jobs can be run from the replication manager, command line, or batch files, or by a scheduler.

You can automate replication and other tasks by using the jobs function in the replication manager. With the server's integrated job editor, you can generate a job template or create your own job by dragging and dropping tasks (script actions) into the Job Content pane. Complete task parameters by choosing resources from drop-down lists or by entering resource names.

For more information, see the online help.

## Command line user interface

The command line user interface (CLUI) allows you to perform various replication tasks, using individual commands and command scripts.

Using the CLUI, you can:

- Perform remote replication.
- Mount and unmount storage volumes.
- Run jobs from a host's command line user interface.
- Write scripts that run jobs.
- Use job return codes for conditional interactions between jobs and host scripts.
- Display resource information.

The CLUI is installed when you install the replication manager.





# Recovery

## 4

This chapter provides recovery information for performing failovers, and resuming operations after encountering a failsafe-locked condition or a disk group failure. Procedures use the HP StorageWorks Replication Solutions Manager, when applicable. Several scenarios are provided that cover most situations you could encounter, with procedures for handling each scenario.

## Planning for a disaster

When a disaster occurs at one of your storage sites, your first priority is to get your data back online in the shortest amount of time. Planning helps to minimize the downtime brought on by a disaster:

- Operating with a supported disaster-tolerant HP Continuous Access EVA configuration is of primary importance. Ensure that there are two fabrics with at least one intersite link per fabric.
- Ensure that your controllers are cabled in the supported “cross-cabling” configuration to your fabrics. See the *HP StorageWorks Continuous Access EVA planning guide* for more information.
- Have at least one management server available at every site, in case of a hardware or communication failure.
- Verify that each destination virtual disk within a DR group has been presented to a host. This allows the host access to the virtual disk immediately after a failover.
- Ensure that local and remote hosts are installed with the latest patches, virus protection, EVA platform kits, and Secure Path versions for that operating system.
- Keep your configuration current and documented at all sites. Install the latest versions of the virtual controller software firmware, management server software, HP Command View EVA, and the replication manager.

- Keep a record of your virtual disks and DR groups. Capture the configuration information after significant changes or at scheduled intervals (see “[Backing up configuration information](#)” on page 62).
- Keep the replication manager on every management server synchronized with any configuration changes (see “Managing storage with multiple management servers” on page 42).
- Back up the replication manager database for any configuration that the management server may use. These databases contain managed set and system folder information that you can quickly restore when a management server changes its role. For example, two management servers can control different arrays at one time, but each management server can have a database created with all the arrays under its control in case the other one becomes inoperative.
- Practice the recovery plan. Ensure that everyone involved in your storage administration practices for disaster recovery. Practice different failure scenarios and make decisions ahead of time about them. For example, if a controller fails, is it more important not to disrupt processing by doing a planned failover, or to not be at risk for a second controller failure that will result in an unplanned failover? In the case of multiple sites, which site has precedence for troubleshooting?

Scheduling practice disaster recoveries is a good way to verify that your records are up-to-date and that all required patches are installed.

## Failover

Failover can take several forms with HP Continuous Access EVA:

- Controller failover is the process that takes place when one controller in a pair assumes the workload of a failed or redirected companion controller in the same cabinet.
- DR group or managed set failover is an operation to reverse the replicating direction of the DR group or managed set to its partner. This reversal is possible because all data generated at a source array has been replicated to a destination array, in readiness for such a situation.
- Fabric or path failover is the act of transferring I/O operations from one fabric or path to another.

This chapter discusses the failover of DR groups and managed sets. It does not discuss controller failover within a cabinet, or path or fabric failover, because redundancy is assumed.

The failover method used with DR groups, managed sets, or arrays is determined by the severity of the failure or the reason for the failover. A planned failover can be used for situations such as an anticipated power disruption, scheduled equipment maintenance at the source array, or the need to transfer operations to another array. An unplanned failover is used for events such as multiple controller failures, multiple host failures, or an unplanned power outage at the source array.

If the source array fails, or if you are planning downtime with the source array, you must decide whether to perform a failover to the destination array. Always verify that all components at the destination array are operational before you begin a failover.

When you perform a failover, the destination array assumes the role of the source and becomes the active array. It remains the source array until you fail over to another system. By transferring control of system operation to the destination array, you can ensure minimal interruption of data access after a failure.

---

**Note:** When you perform a failover for a DR group or a managed set, you must fail over *all* components of the group or set. Therefore, if only one component fails, repairing that single component may be preferable to performing a complete failover.

HP recommends that you not perform a planned or unplanned failover of one or more DR groups more frequently than once every 15 minutes. The planned or unplanned failover of a controller should also not be performed more frequently than once every 15 minutes.

---

Table 4 outlines example situations that require a failover and those that do not. For each type of failover, an action is recommended, which may require action at the source or destination array. Since replication can be bidirectional, one array can be the source and destination for different DR groups. You can use this table to customize contingency plans within your specific environment.

**Table 4: When to and when not to fail over a DR group, managed set, or array**

Type of failure	Recommended action	
	DR group in normal mode	DR group in failsafe mode
Total loss of source array	Manually intervene to fail over data to destination array and then restart processing at the destination array (see <a href="#">“Unplanned failover” on page 66</a> ).	
Loss of both source controllers		
Loss of single source controller	Failover not necessary.	
Total destination array loss	Failover not necessary.	Manually intervene to continue processing at source array. See <a href="#">“Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 81</a> .
Loss of both destination controllers		
Loss of all intersite links		
Loss of both source intersite switches	Manually intervene to fail over data to destination array and then restart processing at the destination array (see <a href="#">“Unplanned failover” on page 66</a> ).	
Loss of single source intersite switch	Failover not necessary.	
Extended power outage at primary site	Manually intervene to fail over data to destination array and then restart processing at the destination array (see <a href="#">“Unplanned failover” on page 66</a> ).	
Loss of a managing server	Failover not necessary. Browse to standby managing server (see <a href="#">“Moving storage management to another server” on page 50</a> ).	
Loss of single disk in redundant storage	Failover not necessary.	

**Table 4: When to and when not to fail over a DR group, managed set, or array (continued)**

Type of failure	Recommended action	
	DR group in normal mode	DR group in failsafe mode
Loss of single host of cluster	Failover not necessary.	
Disk group hardware failure (loss of redundancy) on the source array	Fail over to destination, and repair array (see <a href="#">"Disk group hardware failure on the source array"</a> on page 99).	
Disk group hardware failure (loss of redundancy) on the destination array	Failover not necessary (see <a href="#">"Disk group hardware failure on the destination array"</a> on page 103).	

## Backing up configuration information

It is important to have a record of your storage configuration in case of a hardware failure (see “[Return operations to replaced new storage hardware](#)” on page 67).

## Using SSSU to capture configuration information

One way to capture your configuration information is with the Storage System Scripting Utility (SSSU).

---

**Note:** The SSSU is not available on Novell NetWare hosts or on the management server.

---

Using the SSSU `Capture Configuration` command, five scripts are run that append a user-defined configuration name to the file with a name in the form *UserName\_StepX.txt*, where StepX is one of these five configuration text files:

- Step1A—Captures the data needed to re-create the array itself, disk groups, hosts, virtual disks that are not used for data replication (either source or destination), and LUNS for the disks created.
- Step1B—Captures the data needed to re-create all source virtual disks used in DR groups on this array. However, the data captured by this step is not currently used in any recovery procedures.
- Step1C—Captures the data needed to present all source virtual disks (creates LUNs) used for DR groups to their hosts. However, the data captured by this step is not currently used in any recovery procedures.
- Step2—Captures the data needed to re-create all data replication specific configuration information only, and only DR-specific information for which the array is the source. This consists of source DR groups and their members only.
- Step3—Captures the data needed to create an SSSU script that will again present all remote virtual disks to their hosts.

### Example:

You have a management server with an IP address of 111.222.333.444. You want to back up your configuration for an array named HSV01 with the SSSU. Follow these steps:

1. Run the SSSU executable (*SSSU.exe*) to get a command prompt.

- At the command prompt, log onto the management server using your user name and password. For example:

```
select manager 111.222.333.444 username=user1 password=admin
```

- Select the array whose configuration you want to save. Use the following command with your array name (HSV01 is the array name used in this example):

```
select system HSV01
```

The command line prompt changes to reflect your array is selected.

- At the SSSU command prompt, enter the `Capture Configuration` command along with a path and file name where you want the configuration text files to reside. For example, to copy these files to a folder called `storage_systems\hsv01`, use the command:

```
capture configuration c:\storage_systems\hsv01.txt
```

Messages on the screen confirm that each step was successfully saved.

A current copy of the data that defines the array configuration is saved. If the array configuration changes later, rerun this procedure. Procedures to recover a configuration using the SSSU are in “Return operations to replaced new storage hardware” on page 67.

## Manually capturing configuration information

You can capture configuration information by writing it down manually. Use the following form as a guideline for capturing the information.

**Table 5: Array configuration record**

Array name:	
Array WWID:	
Console LUN ID:	(default = 0)
<b>Disk group information</b>	
Disk group name:	(default = default disk group)
Device count:	
Spare policy:	(none, single, or double)
Disk type:	(online or nearline)

**Table 5: Array configuration record (continued)**

Occupancy alarm:	(default = 95%)
<b>Host information</b>	
Folder name:	(default = \Hosts\)
Host name:	
Operating system:	
For each FCA port:	WWID:
<b>Virtual disk information</b>	
Folder name:	(default = \Virtual Disks\)
Virtual disk name:	
Disk group:	
Size:	
Redundancy level:	(Vraid0, Vraid1, Vraid5)
Write cache policy:	(Mirrored write-back, unmirrored write-back)
Read cache policy:	(On, off)
Read/write permissions:	(Read/write, read-only)
OS unit ID:	(default = 0)
Preferred path:	(None, Path A FO, Path A FO/FB, Path B FO, Path B FO/FB)
Presentation:	Host name:
	LUN:
<b>DR group information</b>	
Source:	Array name:
	Virtual disk members:
Destination:	Array name:
	Virtual disk members:
Parameters:	Failsafe mode: (disabled, enabled)
	Write mode: (synchronous, asynchronous)
	Destination mode: (none, read-only)



---

## Possible event scenarios

This section discusses the following scenarios that require manual intervention:

- [Planned failover](#)
- [Unplanned failover](#)
- [Resume operations if unable to access destination while source in failsafe-locked state \(extended period of time\)](#)
- [Return operations to Home array](#)
- [Return operations to replaced new storage hardware](#)
- [Disk group hardware failure on the source array](#)
- [Disk group hardware failure on the destination array](#)

## Performing recovery actions

Recovery procedures require such actions as failover, suspend, resume, disable failsafe, mounting, and unmounting. You can perform these actions using various interfaces and tools:

- The replication manager
- Command line user interface
- Job scripting
- Storage System Scripting Utility
- HP Command View EVA

This chapter describes generic recovery procedures. For specific procedures, see the documentation for your preferred method.

## Planned failover

**Situation:** Due to scheduled maintenance at the primary site, you need to perform a planned move of operations from the source array to the destination array.

**Action:** Prepare the source array for the failover, and then perform a failover to the destination array. After the failover is complete, you can continue to operate from this array and revert back to failsafe mode, if desired. When the maintenance is complete you can failover to the original source array (see the procedure “Planned failover” on page 68).

**Note:** You can set the Home designation in the replication manager interface to identify the preferred array. By default, an array created as a source is designated as the Home array. Because the role of the source array can change during failover, this Home designation allows you to identify your preferred array.

---

## Unplanned failover

**Situation:** You have experienced an unplanned loss of the primary site or the source array. The duration of the outage at the source is unknown. The HP Continuous Access EVA hardware components (hosts, controllers, and switches, for example) at the primary site may or may not remain intact.

**Action:** Perform an immediate failover to the remote site or passive array. When the primary site is back online, you can choose to return to the Home array or to one with new hardware (see the procedure “Unplanned failover” on page 76).

## Resume operations if unable to access destination while source in failsafe-locked state (extended period of time)

**Situation:** You have experienced an unplanned loss of the destination array, or a loss of the connection to the destination array, due to failure of the intersite links, loss of power at the alternate site, loss of both destination switches, and so on. The duration of the outage is unknown. The DR groups are in failsafe-enabled mode and host I/O is paused because the DR groups are failsafe-locked.

**Action:** Change from failsafe-enabled to normal mode, and then resume host I/O until the connection to the destination array is re-established. When the connection to the destination site is stable, change back to the failsafe-enabled mode (see the procedure “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 81).

## Return operations to Home array

**Situation:** You are operating from an array that was not originally designated as Home within the replication manager. You need to perform a planned move of operations from this alternate source array to the Home array.

**Action:** Prepare the Home array for the failover, and then perform a failover to the Home array (see the procedure “Revert to Home (failback)” on page 83).

## Return operations to replaced new storage hardware

**Situation:** Some type of disaster (lightning, flood, fire, severe equipment failure, or so on) has damaged equipment at the Home site and forced a failover to an alternate site. You are operating from an array that was not originally designated as Home.

**Action:** When the damaged components at the Home site (hosts, controllers, or switches, for example) have been repaired, and the site is operational and back online, perform a failover to new hardware at the Home site (see the procedure “Return operations to replaced new storage hardware” on page 89).

## Disk group hardware failure on the source array

**Situation:** A hardware failure on your source array causes a disk group to become inoperative. This can be caused by the loss of enough disks to create a loss of redundancy within the disk group and affects all Vraid types present on the disk group.

**Action:** If you plan to recover using data on the destination array, then failover to the destination array. Delete DR groups and virtual disks on the failed array. Repair the failed disk group. Re-create DR groups, virtual disks, and host presentations.

If the failed source array was logging at the time of the hardware failure, you must recover with data at the destination site or from a backup.

## Disk group hardware failure on the destination array

**Situation:** A hardware failure on your destination array causes a disk group to become inoperative. This can be caused by the loss of enough disks to create a loss of redundancy within the disk group and affects all Vraid types present on the disk group.

**Action:** Delete the DR groups on the source array that replicated to the failed disk group. Repair the failed disk group on the destination array. Re-create your DR groups on the source array and make host presentations at the destination array.

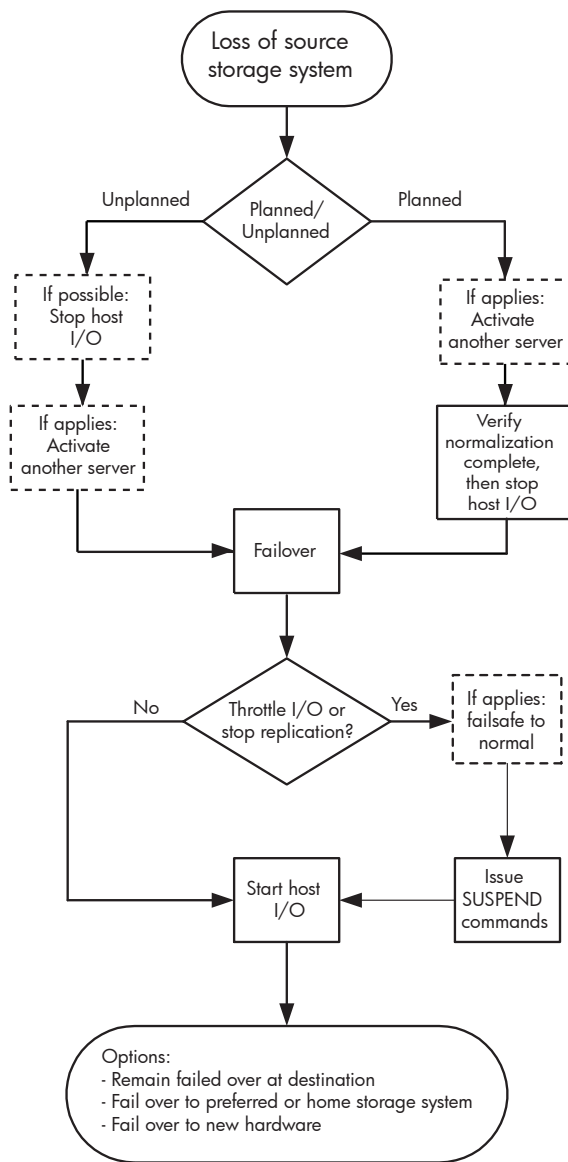
## Failover and recovery procedures

The following procedures explain how to resolve the scenarios mentioned previously. HP recommends that you practice these procedures so that you are prepared in a crisis. Customize these procedures for your own use, if needed.

### Planned failover

See [Figure 6](#) for the flow of steps required for a planned transfer of operations to a remote site. Complete the following steps:

1. If desired, move storage management to another management server.
2. Check to ensure that full normalization has occurred. If a merge and full copy are occurring, wait for them to complete.
3. Stop all host I/O on the source array. Follow the steps listed below for each operating system in your heterogeneous configuration:
  - a. **HP OpenVMS**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and unmount the volumes associated with these virtual disks.
  - b. **HP Tru64 UNIX**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and unmount the volumes associated with these virtual disks.
  - c. **HP-UX**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and then unmount the file systems associated with the virtual disks.
  - d. **IBM AIX**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, then unmount the file systems associated with the virtual disks.



CX08065b

**Figure 6: Planned and unplanned transfer of operations**

- e. **Linux**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and then unmount the file systems associated with the virtual disks.  
  
If you are running Logical Volume Manager (LVM) with or without clustering, see “Bootless DR group planned failover with Linux using LVM in standalone mode or with SuSE SLES 8 running LifeKeeper 4.4.3” on page 127.
  - f. **Microsoft® Windows® NT-X86**—If the operating system is up and running, shut it down.
  - g. **Microsoft Windows 2000/2003**—If the operating system is up and running, shut it down.
  - h. **Novell NetWare**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and then dismount the volumes associated with these virtual disks.
  - i. **Sun Solaris**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and unmount the volumes associated with these virtual disks.
4. Perform the failover operation.
  5. If you plan to throttle I/O to specific arrays, suspend your less important DR groups at your new source. This forces the controllers to replicate the most important data first when the links to the previous source controller are re-established.
  6. If you plan to operate for an extended time at the alternate site (Home array and Fibre Channel links must be functioning properly) and you have a DR group that needs failsafe mode enabled, perform these steps:
    - a. If DR groups were suspended, resume copying on affected destination DR groups. Wait for the log disk to finish merging.
    - b. Change affected DR groups to failsafe mode.

---

**Note:** You can enable failsafe mode at the destination array while a merge or full copy is being performed.

---

7. Issue operating system–dependent commands for presentation of units to remote hosts to start host I/O:
  - a. **HP OpenVMS**—Allow the hosts to recognize new units.
    - 1) If the remote hosts are shut down, boot them now. Booting the hosts enables OpenVMS to recognize the drives.
    - 2) If the remote hosts are not shut down, use the following command from a privileged account to enable OpenVMS to recognize the drives:

```
MCR SYSMAN IO AUTOCONFIGURE/LOG
```

- b. **HP Tru64 UNIX**—Allow the hosts to recognize new units.
      - 1) If the remote hosts are shut down, boot them now. Booting the hosts enables Tru64 UNIX to recognize the drives.
      - 2) If the remote hosts are not shut down, use the following command to recognize the drives:

```
hwmgr -scan scsi
```

This may take a while for large configurations. If this is the case, scan only the SCSI buses that have new units added. Scan only one bus at a time. Use the following command:

```
hwmgr -scan scsi -bus x
```

(where *x* is the SCSI bus number)

- c. **HP-UX**—Allow the remote hosts to recognize the new units.
        - 1) If the remote hosts are shut down, boot them now. Booting the hosts enables HP-UX to recognize the drives.
        - 2) If the remote hosts are not shut down, use the following command to enable HP-UX to recognize the drives and verify that they are present. This command displays only the previous configured failed-over LUNs:

```
ioscan -fnCdisk
```

If the device special files are not displayed, run `insf -e`, and then run `ioscan -fnCdisk` again.

Run the command:

```
vgimport VolumeGroupName DeviceSpecialFile
```

Repeat the previous command for each new failed-over LUN.

Use the following command to mount the LUNs:

```
mount -a
```

---

**Note:** *VolumeGroupName* is the name of the volume group you originally created at the local site. The *DeviceSpecialFiles* are from the `ioscan` in the form of `/dev/dsk/c_t_d/`.

For consistency, configure the same *DeviceSpecialFiles* with the same volume groups, logical volumes, and file systems for the failed-over LUNs at the remote site with the same LUNs at the local site.

---

d. **IBM AIX**—Allow the hosts to recognize new units.

- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables AIX to recognize the drives.
- 2) If the remote hosts are not shut down, use the following commands to recognize the drives and verify that they are present:

```
cfgmgr -v  
lsdev -Cc disk
```

Use the following commands to access file systems on the failed-over virtual disks:

```
importvg -y VolumeGroupName hdiskx  
mount all
```

---

**Note:** *VolumeGroupName* is the name of the volume group you originally created at the local site, and *x* is the number of the `hdisk` assigned to the failed-over virtual disk. If the `-y VolumeGroupName` parameter is omitted, AIX creates a default volume group name for you (for example, `vg00`).

---



- 
- e. **Linux**—Reboot the servers at the remote site, and then remount the file system.
  - f. **Microsoft Windows NT**—Allow the remote hosts to recognize new units.

Reboot the servers at the remote site and log on using an account that has administrative privileges. You should be able to see all of the units in My Computer.

- g. **Microsoft Windows 2000/2003**—Allow the remote hosts to recognize new units.
  - 1) On each host, log on using an account that has administrative privileges.
  - 2) Open Computer Management and click **Disk Management**.
  - 3) After Disk Management has initialized, select **Actions > Rescan Disks**. If the units fail to appear, click **F5** (Refresh). All of the failed-over units are displayed in the right pane.
- h. **Novell NetWare**—Allow the hosts to recognize new units.
  - 1) If the remote hosts are shut down, boot them now. If you are using traditional NetWare volumes, booting the hosts allows Novell NetWare to recognize the drives and automatically mount the volumes. If you are using NSS logical volumes, booting the hosts will recognize the NSS pools and activate them. However, you must manually mount each individual NSS volume by entering `MOUNT VolumeName` at the NetWare console.
  - 2) If the remote hosts are already up and running, or if they do not recognize the drives, issue the following command from the console before mounting the volumes:

```
SCAN FOR NEW DEVICES
```

Alternatively, you can use the *NWCONFIG* utility to issue this same command.

Next, mount the volumes with the following commands:

```
MOUNT ALL (for traditional NetWare volumes)
```

```
MOUNT VolumeName (for NSS logical volumes)
```

- i. **Sun Solaris**—Allow the remote hosts to recognize new units.
  - 1) Reboot the remote hosts using the `reboot -- -r` command, or use the following version-dependent commands to update the Secure Path Manager:

**Solaris 6, 7, and 8**—Run the following commands:

```
drvconfig -v
disks
/opt/CPQswsp/bin/spmgr display
```

**Solaris 9**—

- a) Present new units with LUN numbers sequentially following the old LUNs.
- b) Run the following commands:

```
drvconfig -v
disks
/opt/CPQswsp/bin/spmgr display
```

All of the units with two paths are displayed in the Secure Path Manager. View the units by using the `format` command.

- 2) If Secure Path was not configured for these units, use the following version-dependent commands to add them to the Secure Path Manager.

**Solaris 6, 7, and 8**—Run the following commands:

```
/opt/CPQswsp/bin/spconfig
/opt/CPQswsp/bin/spmgr/display -u
/opt/CPQswsp/bin/spmgr add WWLUNID
drvconfig -v
disks
/opt/CPQswsp/bin/spmgr display
```

---

**Solaris 9—**

- a) Add the units with `spmgr add WWLUNID` or `spmgr add-r WWNN all`.
- b) Run `update_drv -f sd` to inform the system about attribute changes to the sd driver.
- c) Run `disks` to create `/dev` entries for the new units.

You can now view the drives using the `format` command. See the current version of Secure Path documentation for additional assistance.

After the transfer of operation is complete, you have three options after the cause of the failover is resolved:

- Remain failed over at the alternate (or destination) site.
- Return operations to the Home array (see [“Revert to Home \(failback\)”](#) on page 83).
- Return operations to new hardware (see [“Return operations to replaced new storage hardware”](#) on page 89).

## Unplanned failover

See [Figure 6](#) for the flow of steps required for an unplanned transfer of operations to a remote site. Complete the following steps:

1. If your hosts are running on the source array, and you are able to access these hosts, then stop all host I/O.
2. If you cannot access the management server managing the arrays, establish management control with another management server (see “[Managing storage with multiple management servers](#)” on page 42).
3. Perform a failover to the destination site.
4. If you plan to throttle I/O to specific arrays, suspend your less important DR groups at your new source. This forces the controllers to replicate the most important data first when the links to the previous source controller are re-established.
5. Issue operating system-dependent commands for presentation of units to remote hosts to start host I/O:

a. **HP OpenVMS**—Allow the hosts to recognize new units.

1) If the remote hosts are shut down, boot them now. Booting the hosts enables OpenVMS to recognize the drives.

2) If the remote hosts are not shut down, use the following command from a privileged account to enable OpenVMS to recognize the drives:

```
MCR SYSMAN IO AUTOCONFIGURE/LOG
```

b. **HP Tru64 UNIX**—Allow the hosts to recognize new units.

1) If the remote hosts are shut down, boot them now. Booting the hosts enables Tru64 UNIX to recognize the drives.

2) If the remote hosts are not shut down, use the following command to recognize the drives:

```
hwmgr -scan scsi
```

This may take awhile for large configurations. If this is the case, scan only the SCSI buses that have new units added. Scan only one bus at a time. Use the following command:

```
hwmgr -scan scsi -bus x
```

(where *x* is the SCSI bus number)

- 
- c. **HP-UX**—Allow the remote hosts to recognize the new units.
- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables HP-UX to recognize the drives.
  - 2) If the remote hosts are not shut down, use the following command to enable HP-UX to recognize the drives and verify that they are present. This command will display only the previous configured failed-over LUNs:

```
ioscan -fnCdisk
```

If the device special files were not displayed, run `insf -e`, and then run `ioscan -fnCdisk` again.

Run the command:

```
vgimport VolumeGroupName DeviceSpecialFile
```

Repeat the previous command for each new failed-over LUN.

Use the following command to mount the LUNs:

```
mount -a
```

---

**Note:** *VolumeGroupName* is the name of the volume group you originally created at the local site. The *DeviceSpecialFiles* are from the `ioscan` in the form of `/dev/dsk/c_t_d/`.

For consistency, configure the same *DeviceSpecialFiles* with the same volume groups, logical volumes, and file systems for the failed-over LUNs at the remote site with the same LUNs at the local site.

---

- d. **IBM AIX**—Allow the hosts to recognize new units.
- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables AIX to recognize the drives.
  - 2) If the remote hosts are not shut down, use the following commands to recognize the drives and verify that they are present:

```
cfgmgr -v
```

```
lsdev -Cc disk
```

Use the following commands to access file systems on the failed-over virtual disks:

```
importvg -y VolumeGroupName hdiskx  
mount all
```

---

**Note:** *VolumeGroupName* is the name of the volume group you originally created at the local site, and *x* is the number of the hdisk assigned to the failed-over virtual disk. If the *-y VolumeGroupName* parameter is omitted, AIX will create a default volume group name for you, for example, *vg00*.

---

- e. **Linux**—Reboot the servers at the remote site and then remount the file system.
- f. **Microsoft Windows NT**—Allow the remote hosts to recognize new units.

Reboot the servers at the remote site and log in using an account that has administrative privileges. You should be able to see all of the units in My Computer.

- g. **Microsoft Windows 2000/2003**—Allow the remote hosts to recognize new units.
  - 1) On each host, log on using an account that has administrative privileges.
  - 2) Open Computer Management, and then click **Disk Management**.
  - 3) After Disk Management has initialized, select **Action > Rescan Disks**. If the units fail to appear, click **F5** (Refresh). All of the failed-over units are displayed in the right pane.
- h. **Novell NetWare**—Allow the hosts to recognize new units.
  - 1) If the remote hosts are shut down, boot them now. If you are using traditional NetWare volumes, booting the hosts allows Novell NetWare to recognize the drives and automatically mount the volumes. If you are using NSS logical volumes, booting the hosts allows NetWare to recognize and activate the NSS pools. However, you must manually mount each individual NSS volume by entering `MOUNT VolumeName` at the NetWare console.

- 2) If the remote hosts are already up and running, or if they do not recognize the drives, issue the following command from the console before mounting the volumes:

```
SCAN FOR NEW DEVICES
```

Alternatively, you can use the *NWCONFIG* utility to issue this same command.

Next, mount the volumes with these commands:

```
MOUNT ALL (for traditional NetWare volumes)
```

```
MOUNT VolumeName (for NSS logical volumes)
```

- i. **Sun Solaris**—Allow the remote hosts to recognize new units.

- 1) Reboot the remote hosts using the `reboot -- -r` command, or use the following version-dependent commands to update the Secure Path Manager:

**Solaris 6, 7, and 8**—Run the following commands:

```
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

**Solaris 9**—

- a) Present new units with LUN numbers sequentially following the old LUNs.
- b) Run `drvconfig:disks` to be able to see the devices in the `spmgr display -u list`.

All of the units with two paths are displayed in the Secure Path Manager. View the units by using the `format` command.

- 2) If Secure Path was not configured for these units, use the following version-dependent commands to add them to the Secure Path Manager:

**Solaris 6, 7, and 8**—Run the following commands:

```
/opt/CPQswsp/bin/spconfig  
/opt/CPQswsp/bin/spmgr/display -u  
/opt/CPQswsp/bin/spmgr add WWLUNID  
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

**Solaris 9**—

- a) Add the units with `spmgr add WWLUNID` or `spmgr add-r WWNN all`.
- b) Run `update_drv -f sd` to inform the system about attribute changes to the sd driver.
- c) Run `disks` to create `/dev` entries for the new units.

You can now view the drives using the `format` command. See the current version of Secure Path documentation for additional assistance.

With Solaris, you may need to execute the `fsck` command before the operating system can mount the LUN. HP recommends that you run the `fsck` command with the `-m` option before running `fsck` to repair the file system.

When the transfer of operation is complete and the cause of the failover is resolved, you have three options:

- Remain failed over at the alternate (or destination) site.
- Return operations to the Home array (see “[Revert to Home \(failback\)](#)” on page 83).
- Return operations to new hardware (see “[Return operations to replaced new storage hardware](#)” on page 89).



---

## Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)

See [Figure 7](#) for the flow of steps required to resume operations if you are unable to access the destination while in a failsafe-locked state. Complete the following steps:

1. Change affected source DR groups from failsafe-enabled mode to normal mode.
2. If necessary, issue operating system-dependent commands to the local hosts to start I/O again on the units that were failsafe-locked.
3. If you plan to throttle I/O to specific arrays, suspend your less important DR groups. This forces the controllers to replicate the most important data first when the links are re-established. When ready to merge to destination from source, issue the `Resume` command.

---

**Note:** If you stay in a suspended state for an extended length of time, you can overrun the log, which initiates a full copy. During a full copy, the data is not usable until the full copy completes.

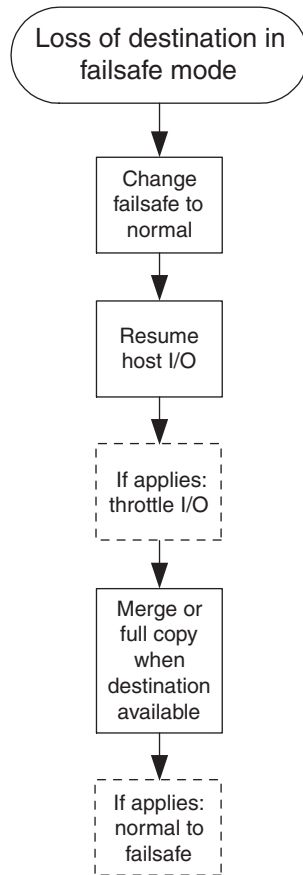
---

4. When connections to the destination site are re-established and merging is complete, change DR groups from normal mode to failsafe-enabled mode, if desired.

---

**Note:** If source DR groups go into full copy mode, you can also enable failsafe mode.

---



CXO8066A

**Figure 7: Resumption of operations if unable to access destination in failsafe mode**

## Revert to Home (failback)

Failback (also known as reverting to Home) is similar to a planned failover. After performing a failover from a source to a destination array, a failback can be performed in the other direction after waiting a minimum of 15 minutes.

Complete the following steps:

1. If desired, move storage management to another management server.
2. Ensure that full normalization occurred. If a merge and full copy are occurring, wait for them to complete.
3. Stop all host I/O on the source array. Follow the steps below for each operating system in your heterogeneous configuration:
  - a. **HP OpenVMS**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and then unmount the volumes associated with these virtual disks.
  - b. **HP Tru64 UNIX**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and unmount the volumes associated with these virtual disks.
  - c. **HP-UX**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and then unmount the file systems associated with the virtual disks.
  - d. **IBM AIX**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and then unmount the file systems associated with the virtual disks.
  - e. **Linux**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and then unmount the file systems associated with the virtual disks.

If you are running Logical Volume Manager (LVM) with or without clustering, see “Bootless DR group planned failover with Linux using LVM in standalone mode or with SuSE SLES 8 running LifeKeeper 4.4.3” on page 127.

- f. **Microsoft Windows NT-X86**—If the operating system is up and running, shut it down.
- g. **Microsoft Windows 2000/2003**—If the operating system is up and running, shut it down.
- h. **Novell NetWare**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and then dismount the volumes associated with these virtual disks.

- i. **Sun Solaris**—If the operating system is up and running, remove all I/O to the virtual disks that will be failed over, and then unmount the volumes associated with these virtual disks.
4. Perform the “revert to Home” operation. Or, perform the failover operation on affected DR groups.
5. If you plan to throttle I/O to specific arrays, suspend your less important DR groups at your new source. This forces the controllers to replicate the most important data first when the links to the previous source controller are re-established.
6. If you plan to operate for an extended time at the alternate site (Home array and Fibre Channel links must be functioning properly) and you have a DR group that needs failsafe mode enabled, perform the following steps:
  - a. If DR groups were suspended, resume copying on affected destination DR groups. Wait for the log disk to finish merging.
  - b. Enable failsafe mode of the affected DR groups.

---

**Note:** You can enable failsafe mode at the destination array while a merge or full copy is being performed.

---

7. Issue operating system–dependent commands for presentation of units to remote hosts to start host I/O.
  - a. **HP OpenVMS**—Allow the hosts to recognize new units.
    - 1) If the remote hosts are shut down, boot them now. Booting the hosts enables OpenVMS to recognize the drives.
    - 2) If the remote hosts are not shut down, use the following command from a privileged account to enable OpenVMS to recognize the drives:
  - b. **HP Tru64 UNIX**—Allow the hosts to recognize new units.
    - 1) If the remote hosts are shut down, boot them now. Booting the hosts enables Tru64 UNIX to recognize the drives.
    - 2) If the remote hosts are not shut down, use the following command to recognize the drives:

```
MCR SYSMAN IO AUTOCONFIGURE/LOG
```

```
hwmgr -scan scsi
```

---

This may take awhile for large configurations. If this is the case, scan only the SCSI buses that have new units added. Scan only one bus at a time. Use the following command:

```
hwmgr -scan scsi -bus x
```

(where *x* is the SCSI bus number)

- c. **HP-UX**—Allow the remote hosts to recognize the new units.
- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables HP-UX to recognize the drives.
  - 2) If the remote hosts are not shut down, use the following command to enable HP-UX to recognize the drives and verify that they are present. This command displays only the previously-configured failed-over LUNs:

```
ioscan -fnCdisk
```

If the device special files are not displayed, run `insf -e`, and then run `ioscan -fnCdisk` again.

Run the command:

```
vgimport VolumeGroupName DeviceSpecialFile
```

Repeat the previous command for each new failed-over LUN.

Use the following command to mount the LUNs:

```
mount -a
```

---

**Note:** *VolumeGroupName* is the name of the volume group you originally created at the local site. The *DeviceSpecialFiles* are from the `ioscan` in the form of `/dev/dsk/c_t_d/`.

For consistency, configure the same *DeviceSpecialFiles* with the same volume groups, logical volumes, and file systems for the failed-over LUNs at the remote site with the same LUNs at the local site.

---

- d. **IBM AIX**—Allow the hosts to recognize new units.
- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables AIX to recognize the drives.
  - 2) If the remote hosts are not shut down, use the following commands to recognize the drives and verify that they are present:

```
cfgmgr -v  
lsdev -Cc disk
```

Use the following commands to access file systems on the failed-over virtual disks:

```
importvg -y VolumeGroupName hdiskx  
mount all
```

---

**Note:** *VolumeGroupName* is the name of the volume group you originally created at the local site, and *x* is the number of the hdisk assigned to the failed-over virtual disk. If the *-y VolumeGroupName* parameter is omitted, AIX will create a default volume group name for you (for example, vg00).

---

- e. **Linux**—Reboot the servers at the remote site, and then remount the file system.
- f. **Microsoft Windows NT**—Allow the remote hosts to recognize new units.

Reboot the servers at the remote site and log on using an account that has administrative privileges. You can now view all of the units in My Computer.

- g. **Microsoft Windows 2000/2003**—Allow the remote hosts to recognize new units.
- 1) On each host, log in using an account that has administrative privileges.
  - 2) Open Computer Management and click **Disk Management**.
  - 3) After Disk Management has initialized, select **Action > Rescan Disks**. If the units fail to appear, click **F5** (Refresh). All of the failed-over units are displayed in the right pane.

- 
- h. **Novell NetWare**—Allow the hosts to recognize new units.
- 1) If the remote hosts are shut down, boot them now. If you are using traditional NetWare volumes, booting the hosts allows Novell NetWare to recognize the drives and automatically mount the volumes. If you are using NSS logical volumes, booting the hosts allows NetWare to recognize and activate the NSS pools. However, you must manually mount each individual NSS volume by entering `MOUNT VolumeName` at the NetWare console.
  - 2) If the remote hosts are already up and running, or if they do not recognize the drives, issue the following command from the console before mounting the volumes:

```
SCAN FOR NEW DEVICES
```

Alternatively, you can use the *NWCONFIG* utility to issue this same command.

Next, mount the volumes with the following commands:

```
MOUNT ALL (for traditional NetWare volumes)
```

```
MOUNT VolumeName (for NSS logical volumes)
```

- i. **Sun Solaris**—Allow the remote hosts to recognize new units.
- 1) Reboot the remote hosts using the `reboot -- -r` command, or use the following version-dependent commands to update the Secure Path Manager:

**Solaris 6, 7, and 8**—Run the following commands:

```
drvconfig -v
```

```
disks
```

```
/opt/CPQswsp/bin/spmgr display
```

**Solaris 9—**

a) Present new units with LUN numbers sequentially following the old LUNs.

b) Run the following commands:

```
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

You can view all the units with two paths in the Secure Path Manager, using the `format` command.

2) If Secure Path was not configured for these units, use the following version-dependent commands to add them to the Secure Path Manager.

**Solaris 6, 7, and 8—**Run the following commands:

```
/opt/CPQswsp/bin/spconfig  
/opt/CPQswsp/bin/spmgr/display -u  
/opt/CPQswsp/bin/spmgr add WWLUNID  
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

**Solaris 9—**

a) Add the units with `spmgr add WWLUNID` or `spmgr add-r WWNN all`.

b) Run `update_drv -f sd` to inform the system about attribute changes to the `sd` driver.

c) Run `disks` to create `/dev` entries for the new units.

You can now view the drives using the `format` command. See the current version of Secure Path documentation for additional assistance.



After the transfer of operation is complete, you have three options after the cause of the failover is resolved:

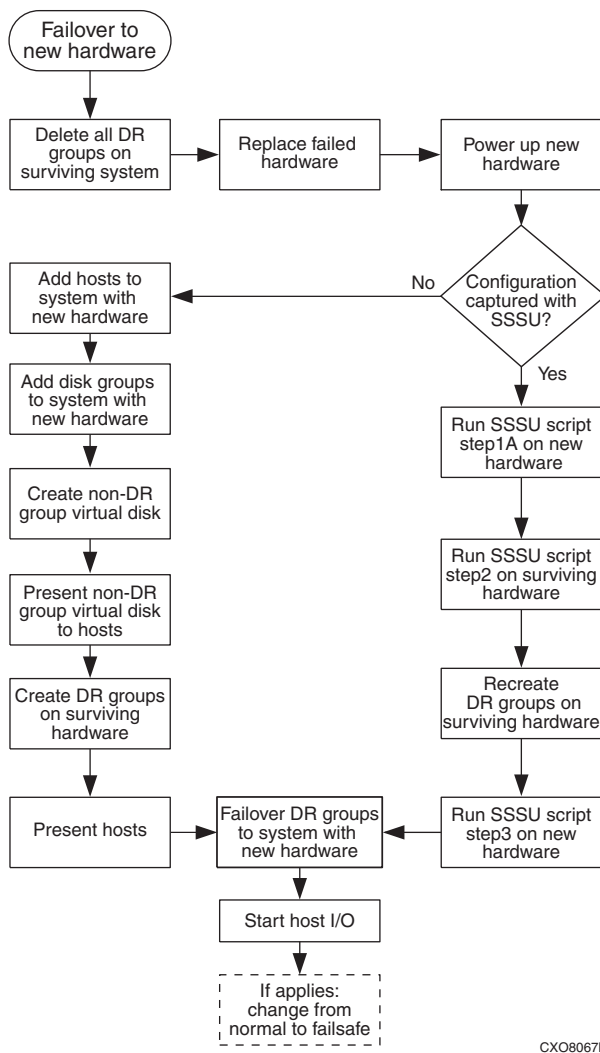
- Remain failed over at the alternate (or destination) site.
- Return operations to the Home array (see [“Revert to Home \(failback\)”](#) on page 83).
- Return operations to new hardware (see [“Return operations to replaced new storage hardware”](#) on page 89).

## Return operations to replaced new storage hardware

This procedure is used after a failure that results in the replacement of array hardware at what was the source array. The procedure does not include steps to rebuild servers using the storage (this should be part of your overall disaster plan). The new hardware now acts as the destination array after a failover and is referred to in this procedure as the system with new hardware, or the array with failed hardware. The surviving system is now your source array after the failover. The steps below explain the process to return operations to a system having replaced new hardware:

1. Denote your array names having failed or new hardware (destination) and your surviving array (source) in the table provided below. For example, your array with new hardware may be named HSV01 and your surviving array may be named HSV02. See the table during the procedure as needed.

	Array with failed or new hardware	Surviving array
Array Name		
Array Name		
Array Name		
Array Name		
Array Name		
Array Name		
Array Name		
Array Name		



CXO8067B

**Figure 8: Return operations to new hardware**

2. Delete all DR groups on the surviving system that ever had a relationship with the failed hardware.
3. Replace the failed hardware. Depending on the failure, this means replacing hard drives or controllers, deleting disk groups, and so on.

4. Remove the connection between the source and destination arrays. This can be accomplished by removing it from the SAN, disabling the intersite links, or by placing the arrays into separate zones.

To place the array into separate zones, you need two zones. One zone contains the source array, source hosts, and the management server. The second zone contains the destination arrays, destination hosts, and the management server.

5. (Optional. Not needed if entire array was replaced.) Delete any destination DR groups on the previously failed array. If this is not successful, the source-destination connection still exists, so go to the previous step.
6. (Optional. Not needed if entire array was replaced.) Delete all virtual disks that were members of DR groups on the destination array.
7. Re-establish communication between the source and destination arrays. Either add the array back into the SAN, enable the intersite links, or place the arrays into the same zone.
8. Perform one of the following:
  - a. If the replaced array configuration was captured with the Storage System Scripting Utility (SSSU), execute the script *ConfigName\_step1A* on the new hardware, and then proceed to step 13. See the SSSU documentation for instructions. *ConfigName* is a user-assigned name given to the SSSU script at the time of creation. See the procedure titled “[Backing up configuration information](#)” on page 62.
  - b. If you are not using an SSSU script for recovery, initialize the newly replaced array using the information you recorded in the “Array configuration record” on page 63. See the HP Command View EVA documentation for initialization instructions.

---

**Note:** To preserve your existing zoning, give the new hardware the same World Wide Names as they existed with the failed hardware.

---

9. Add the disk groups on the new hardware.
10. Add the hosts for the system with new hardware.
11. Create the non-DR group virtual disks.
12. Present all non-DR group virtual disks to their hosts.

13. Perform one of the following:
  - a. If the surviving array configuration was captured with the SSSU, execute *ConfigName\_step2* on the surviving array. *ConfigName* is a user-assigned name given to the SSSU script at the time of creation. DR groups are re-created with the SSSU if they were performing as the source when the configuration was captured. This step may take some time to complete.
  - b. If you are not using an SSSU script for recovery, re-create all DR groups on the surviving system with destinations set to the new system using the information you recorded in the “Array configuration record” on page 63.
14. If, in the previous step, you used the SSSU to re-create DR groups on the surviving array, you must manually re-create any additional DR groups that had their source on the failed hardware on the surviving array. This is necessary because the SSSU will not re-create those DR groups on the surviving array if they performed as the destination when the configuration was captured. After you perform this step, all DR groups reside on the surviving array.
15. At this point, you have the option of setting all affected DR groups from normal mode to failsafe-enabled mode.
16. Perform one of the following:
  - a. If the original array configuration was captured with the SSSU, then execute *ConfigName\_step3* on the new hardware. *ConfigName* is a user-assigned name given to the SSSU script at the time of creation.
  - b. If you are not using an SSSU script for recovery, present the destination virtual disks on the system with new hardware to the appropriate hosts using the information you recorded in the “Array configuration record” on page 63.
17. If, in the previous step, you used the SSSU to present destination virtual disks to their hosts, manually present any additional virtual disks that originally had their sources on the failed hardware to their hosts on the array with new hardware. This is necessary because the SSSU will not present virtual disks whose destinations were the surviving array when the configuration was captured. After performing this step, all destination virtual disks are presented to hosts.
18. If the system is to be the source for the DR groups, fail over any DR groups to the new array using the procedure “Planned failover” on page 65.

- 
19. Issue operating system–dependent commands for presentation of units to hosts to start I/O:
- a. **HP OpenVMS**—Allow the hosts to recognize new units:
    - 1) If the remote hosts are shut down, boot them now. Booting the hosts enables OpenVMS to recognize the drives.
    - 2) If the remote hosts are not shut down, use the following command from a privileged account to enable OpenVMS to recognize the drives:

```
MCR SYSMAN IO AUTOCONFIGURE/LOG
```
  - b. **HP Tru64 UNIX**—Allow the hosts to recognize new units.
    - 1) If the remote hosts are shut down, boot them now. Booting the hosts enables Tru64 UNIX to recognize the drives.
    - 2) If the remote hosts are not shut down, use the following command to recognize the drives:

```
hwmgr -scan scsi
```

This may take awhile for large configurations. If this is the case, scan only the SCSI buses that have new units added. Scan only one bus at a time. Use the following command:

```
hwmgr -scan scsi -bus x
```

(where x is the SCSI bus number)

- c. **HP-UX**—Allow the remote hosts to recognize the new units.
- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables HP-UX to recognize the drives.
  - 2) If the remote hosts are not shut down, use the following command to enable HP-UX to recognize the drives and verify that they are present. This command will display only the previous configured failed-over LUNs:

```
ioscan -fnCdisk
```

If the device special files were not displayed, run `insf -e`, then run `ioscan -fnCdisk` again.

Run the command:

```
vgimport VolumeGroupName DeviceSpecialFile
```

Repeat the previous command for each new failed-over LUN.

Use the following command to mount the LUNs:

```
mount -a
```

---

**Note:** *VolumeGroupName* is the name of the volume group you originally created at the local site. The *DeviceSpecialFiles* are from the `ioscan` in the form of `/dev/dsk/c_t_d/`.

For consistency, configure the same *DeviceSpecialFiles* with the same volume groups, logical volumes, and file systems for the failed-over LUNs at the remote site with the same LUNs at the local site.

---

- d. **IBM AIX**—Allow the hosts to recognize new units.
- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables AIX to recognize the drives.
  - 2) If the remote hosts are not shut down, use the following commands to recognize the drives and verify that they are present:

```
cfgmgr -v
```

```
lsdev -Cc disk
```

Use the following commands to access file systems on the failed-over virtual disks:

```
importvg -y VolumeGroupName hdiskx  
mount all
```

---

**Note:** *VolumeGroupName* is the name of the volume group you originally created at the local site, and *x* is the number of the hdisk assigned to the failed-over virtual disk. If the *-y VolumeGroupName* parameter is omitted, AIX will create a default volume group name for you, for example, vg00.

---

- e. **Linux**—Reboot the servers at the remote site and then remount the file system.
- f. **Microsoft Windows NT**—Allow the remote hosts to recognize new units.

Reboot the servers at the remote site and log in using an account that has administrative privileges. You should be able to see all of the units in My Computer.

- g. **Microsoft Windows 2000/2003**—Allow the remote hosts to recognize new units.
  - 1) On each host, log on using an account that has administrative privileges.
  - 2) Open Computer Management and click **Disk Management**.
  - 3) After Disk Management has initialized, select **Action > Rescan Disks**. If the units fail to appear, click **F5** (Refresh). All of the failed-over units are now displayed in the right pane.
- h. **Novell NetWare**: Allow the hosts to recognize new units.
  - 1) If the remote hosts are shut down, boot them now. If you are using traditional NetWare volumes, booting the hosts allows Novell NetWare to recognize the drives and automatically mount the volumes. If you are using NSS logical volumes, booting the hosts allows NetWare to recognize and activate the NSS pools. However, you must manually mount each individual NSS volume by entering `MOUNT VolumeName` at the NetWare console.

- 2) If the remote hosts are already up and running, or if they do not recognize the drives, issue the following command from the console before mounting the volumes:

```
SCAN FOR NEW DEVICES
```

Alternatively, you can use the *NWCONFIG* utility to issue this same command.

Next, mount the volumes with these commands:

```
MOUNT ALL (for traditional NetWare volumes)
```

```
MOUNT VolumeName (for NSS logical volumes)
```

- i. **Sun Solaris**—Allow the remote hosts to recognize new units.

- 1) Reboot the remote hosts using the `reboot -- -r` command, or use the following version-dependent commands to update the Secure Path Manager:

**Solaris 6, 7, and 8**—Run the following commands:

```
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

**Solaris 9**—

- a) Present new units with LUN numbers sequentially following the old LUNs.
- b) Run `drvconfig:disks` to be able to see the devices in the `spmgr display -u list`.

You can view all the units with two paths in the Secure Path Manager, using the `format` command.



- 2) If Secure Path was not configured for these units, use the following version-dependent commands to add them to the Secure Path Manager.

**Solaris 6, 7, and 8**—Run the following commands:

```
/opt/CPQswsp/bin/spconfig
/opt/CPQswsp/bin/spmgr/display -u
/opt/CPQswsp/bin/spmgr add WWLUNID
drvconfig -v
disks
/opt/CPQswsp/bin/spmgr display
```

**Solaris 9**—

- a) Add the units with `spmgr add WWLUNID` or `spmgr add-r WWNN all`.
- b) Run `update_drv -f sd` to inform the system about attribute changes to the sd driver.
- c) Run `disks` to create `/dev` entries for the new units.

You can view the drives using the `format` command. See the current version of Secure Path documentation for additional assistance.

20. (Optional) Set the DR groups to the desired Home setting.

## Recovering from a disk group hardware failure

Disk group hardware failure occurs when a disk group loses a quantity of disks beyond the capability from which a given Vraid type can recover. It is a loss of redundancy that results in an inoperative disk group. This condition can occur from the loss of one disk for Vraid0 to as few as two disks for Vraid1 and Vraid5. In each case, the hardware failure needs to be fixed, and the disk group data has to be structurally rebuilt. This section describes the symptoms and recovery methods for an inoperative disk group at either the source or destination array.





If an array only has one disk group, and that disk group fails, the array becomes inoperative. Re-initialize the array to manage it (see “Return operations to replaced new storage hardware” on page 67).

If you have multiple disk groups and one fails, follow the procedures on page 99 and page 103.

### Failed disk group hardware indicators

If disk group hardware fails, the replication manager displays the following:

**Table 6: Failed disk group hardware indicators**

Resource	Symbol	Description
Array		Indicates the array is in an abnormal state and requires attention.
Virtual disks		Indicates a catastrophic failure and requires immediate action.
DR groups	 or 	Red indicates a failure; yellow indicates the DR group is in a degraded state. Either condition requires immediate attention.

---

## Disk group hardware failure on the source array

There are two ways to recover from a disk group hardware failure on the source array:

- If data replication was occurring normally when the source disk group became inoperative, the data at the destination array is current. A failover is performed to the destination array, DR groups are deleted, the inoperative disk group is repaired, and the DR groups are re-created. Data is then copied back.
- If your disk group becomes inoperative when your DR groups are logging (for example, your DR groups were suspended, or the intersite links are down), your data is stale on the destination array. Stale data is older data that is not as current as what exists on its partnered array. If you prefer to use stale data for recovery, the steps are the same as if replication was occurring normally. However, if you prefer to continue from a point-in-time, the inoperative disk group is repaired, and data is restored from a backup or full copy.

---

**Note:** When you delete DR groups to recover from a disk group hardware failure, you lose the redundancy of the other site or disaster tolerance of your data.

---

Perform this procedure using HP Command View EVA when a disk group hardware failure occurs on the source array and the data on the destination array is current.

1. Navigate to each DR group on the surviving array and perform a failover (see “Unplanned failover” on page 66).
2. In HP Command View EVA, begin troubleshooting the disk group problem.
3. Navigate to the failed disk group.

A list of failed virtual disks and DR groups is displayed.

hp StorageWorks  
command view eva

Appliance:  
SMAD123AB45C678 111.222.333.444

Root View | Agent Options | Help

HSV Storage Network  
 HSV06  
 HSV05  
 Virtual Disks  
 Hosts  
 Disk Groups  
 Default Disk Group  
 Disk 001  
 Disk 001  
 Disk 001  
 Disk 002  
 Disk 003  
 Disk 004  
 Disk 005  
 Disk 006  
 Disk 007  
 Disk 008  
 Disk 009  
 Disk 010  
 Disk 011  
 Disk 012  
 Disk 013  
 Disk 014

**Disk Group Hardware Failure**

The following Vdisks and data replication groups have failed as a result of a hardware failure in this disk group.

**Failed Vdisks**  
 Virtual Disks\standalone vdisk\SA1\ACTIVE  
 Virtual Disks\win05\_vdisk\win05\_vdisk1\ACTIVE

**Failed Data Replication Groups**  
 DR\_Group 001  
 win05\_dr1

- Click the **Continue with no changes** button to continue operation without changing anything in your system at this time. In some cases, repairing the hardware will allow your system's virtualization features to automatically repair the Vdisk failure. Your disk group will display as degraded until you repair your hardware or return to this page to delete the failed Vdisks and DR groups shown above.
- Click the **Start deletion process** button to delete the failed Vdisks and data replication groups shown above. If the hardware failure is small, your disk group will be restored to normal operation, but without the failed objects above. If the hardware failure is severe, and your disk group cannot be restored, it will be deleted along with the objects listed above.

Continue with no changes | Start deletion process

4. Click the **Start deletion process** tab.

A message displays requesting confirmation.

5. Click **OK**.

A list of affected DR groups requiring deletion is displayed.

hp StorageWorks  
command view eva

Appliance:  
SMAD123AB45C678 111.222.333.444

Root View | Agent Options | Help

HSV Storage Network  
 HSV06  
 HSV05  
 Virtual Disks  
 Hosts  
 Disk Groups  
 Default Disk Group  
 Ungrouped Disks  
 Data Replication  
 Hardware

**Data Replication Group Manual Deletion**

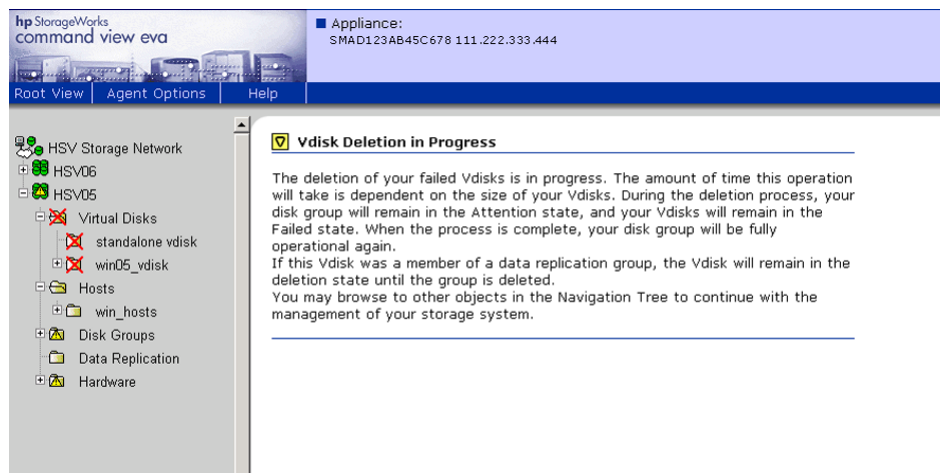
To resolve your disk group's failure condition, you must first navigate to each affected data replication group shown below and delete it.

You can only delete a DR group at its source system, so you may have to navigate to other systems to delete some or all of the DR groups listed below. The details of your data replication configuration will determine exactly what steps you must follow to delete your DR groups.

**Affected Data Replication Groups**  
 DR\_Group 001  
 win05\_dr1

You must monitor the state of your data replication groups to determine when the deletion process is complete. When it is finished, navigate to this disk group to finish restoring its normal operation by deleting any failed Vdisks it contains.

6. Select an affected DR group and click **Delete**.  
A message is displayed to inform you that a DR group is being deleted.
7. Click **OK**.  
The affected DR groups are deleted.
8. Select the failed virtual disks that were members of the affected DR groups.  
A message is displayed while virtual disks are being deleted.



When the deletion completes, an HP Command View EVA virtual disk Folder Properties screen is displayed, showing the virtual disk was deleted.

9. Navigate to the disk group and click **Finish**.
10. On the surviving array, delete the source DR group associated with the failed DR groups deleted in [step 6](#).
11. (Optional) Repair your hard drives and re-create your disk group (see the HP Command View EVA documentation).
12. In the replication manager on the surviving array, re-create the DR groups and set up host presentation on the repaired array.
13. After normalization occurs between the source and destination arrays, fail over to the repaired array using the procedure “Planned failover” on page 68.

### **Recovery when data replication was logging before failure**

If data is logging when a source disk group hardware failure occurs, the data on the destination array is stale (not current). You have the following options:

- Recover using the stale data on the destination array (see [“Disk group hardware failure on the source array”](#) on page 67).
- Recover from a known, good point using a backup.
  - If you want to perform a failover to quickly activate the destination array before repairing the inoperative disk group, use the procedure on page 67, and then restore from a backup.
  - If you want to repair the inoperative disk group first, perform the repair, delete the inoperative DR groups and virtual disks on the failed system, re-create your virtual disks and DR groups, and then restore your data from an external backup.

## Disk group hardware failure on the destination array

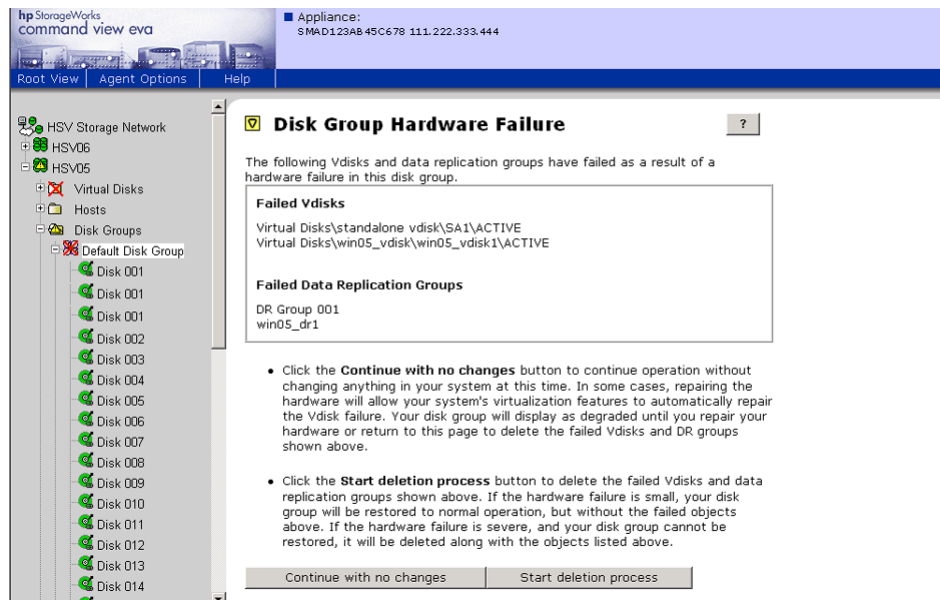
This section describes how to recover from an inoperative disk group on your destination array. Your first indications that a disk group has become inoperative may be icons like those shown in “[Failed disk group hardware indicators](#)” on page 98, except that your destination disk group status is Unknown.

**Note:** When you delete DR groups to recover from a disk group hardware failure, you lose the redundancy of the other site or disaster tolerance of your data.

Perform this procedure using HP Command View EVA when a disk group hardware failure occurs on the destination array and the data on the source array is current.

1. In HP Command View EVA, begin troubleshooting the disk group problem.
2. Navigate to the failed disk group.

A list of failed virtual disks and DR groups is displayed.

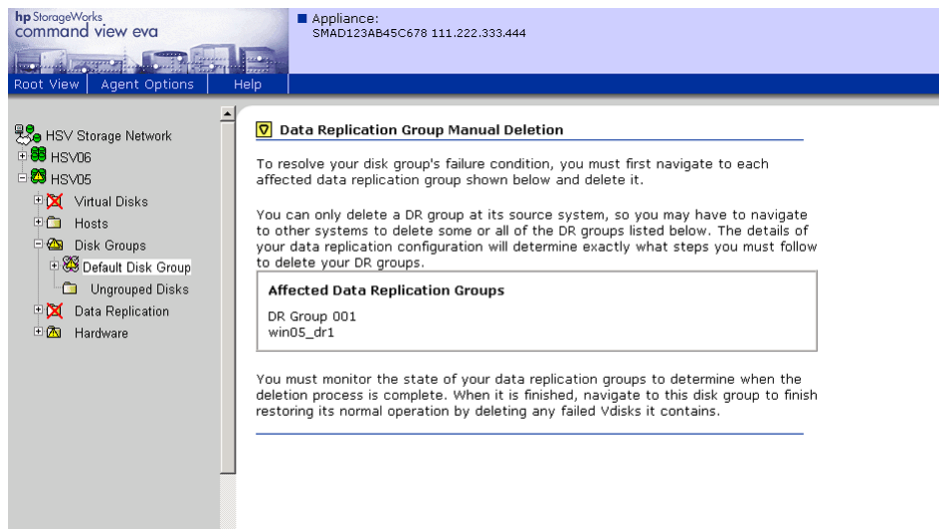


3. Click the **Start deletion process** tab.

A confirmation message is displayed.

4. Click **OK**.

A list of affected DR groups requiring deletion is displayed.

5. Select an affected DR group and click **Delete**.

A confirmation message is displayed.

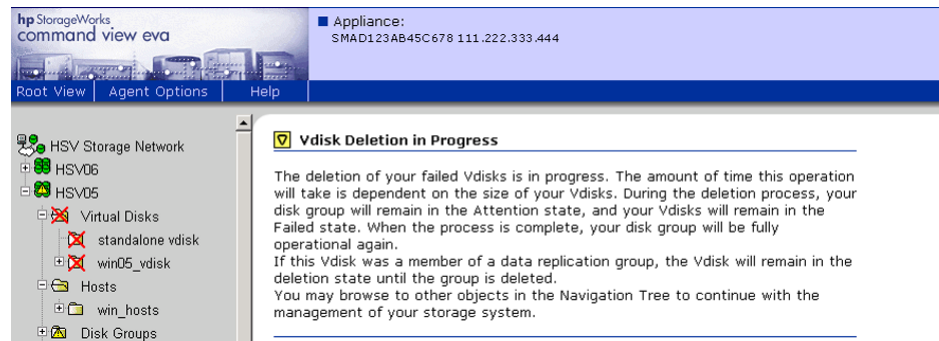
6. Click **OK**.

The affected DR group is deleted.

7. Repeat [step 5](#) and [step 6](#) for each affected DR group.

## 8. Select failed virtual disks that were members of the affected DR groups.

A message is displayed while virtual disks are being deleted.





When the deletion completes, an HP Command View EVA virtual disk Folder Properties screen is displayed showing the virtual disk was deleted.

9. Navigate to the disk group and click **Finish** to resolve the disk group hardware failure.
10. On the surviving array, delete the source DR group associated with the failed DR groups deleted in [step 6](#).
11. (Optional) Repair your hard drives and re-create your disk group (see the HP Command View EVA documentation).
12. In the replication manager on the surviving array, re-create the DR groups and set up host presentation on the repaired array.



# Troubleshooting

## 5

This chapter provides troubleshooting guidance for arrays and links between multiple sites.

### LUN inaccessible to host

A “stalled LUN” event (4206001b) in HP Command View EVA indicates that a LUN has been inaccessible to the host for at least four minutes, causing the LUN to be in a quiesced state. Take the following actions to troubleshoot this situation and prevent possible data loss:

1. Verify that the host still cannot access the LUN.
2. Try to resynchronize the controller from the HP Command View EVA field service page.
3. If the situation still exists, unpresent and re-present the LUN to the host.
4. If the situation still exists, restart the controller and its partner controller, if necessary.

### DR groups in unknown state

If your DR groups are in an unknown state, check to see if you have recently imported the replication manager database from an active management server to the management server where the DR groups are in an unknown state. If so, the problem is probably that the passwords do not match on the management servers. See “[Changing the password for the management server](#)” on page 45 for more information.

## Tunnel thrash

Tunnel thrash is the frequent closing and re-opening of a tunnel while holding host I/O in the transition. This occurs when peer controllers can see each other, but cannot sustain replication data with any path, even when throttled to the minimum. Some possible causes of tunnel thrash are:

- High volumes of packet loss
- Incorrectly configured routers
- Re-routed IP circuits
- Oversubscribed circuits

Although tunnel thrash is rare, a critical event (c23670c) is generated and displayed in HP Command View EVA for each DR group that shares the affected tunnel. You must intervene to prevent possible data loss.

Take the following actions to resolve this situation:

- Check all routers and look for high volumes of packet loss.
- Ensure that all router are configured correctly.
- Contact your service provider to check if the circuit has been alternate routed.
- Check to see if thrashing occurs during peak times and not during low volume times. If so, the circuit may be over subscribed and you may need to increase bandwidth.

---

**Note:** An informational event (c22000c) is generated for an open tunnel. No action is required.

---

## Remote server cannot detect a destination LUN

If you have a remote server that cannot detect a destination LUN, it could be that the DR group access mode is set to “disabled.” A remote server can detect a LUN with a “read-only” access mode, but cannot detect it if the mode is set to “disabled.” The replication manager allows you to change the DR group’s access mode from “disabled” to “read-only,” thereby allowing the remote server to detect the destination LUN. See the online help for information on editing a DR group’s properties.

## Long delays or time-outs on HP-UX

If an HP-UX host has multiple disk devices with failed or no longer presented LUNS behind them, it can take an increasingly long time to gather host information as the number of disk devices increases. If an HP-UX host exhibits time-outs on host discovery or failed jobs while waiting for host operations to complete, take the following actions:

- Check the disk devices showing long time-outs. Secure Path can display the status of the disk devices it is managing. For disk devices not managed by Secure Path, check for I/O time-outs by running an OS tool such as `diskinfo` on each disk device.

Remove any disk devices that show long time-outs, if they are no longer needed.

- If the disk devices are intentionally in this state, improve performance by modifying the I/O time-out setting for those disks with the `pvchange -t` command. HP-UX has a default I/O time-out of 30 seconds for SCSI disks. The `pvchange -t` command allows you to reduce the amount of time before a time-out on a given disk occurs. Reducing the time-out decreases the amount of time a host discovery takes.

## IP address on UNIX systems

During a normal installation of HP Continuous Access EVA on a UNIX system, you are prompted to enter the IP address of the local host. The IP address is saved in `/opt/sanmgr/hostagent/config/commIpAddr.txt`. Un-installing HP Continuous Access EVA does not remove the IP address, and you are not prompted for the IP address if you re-install. If you entered the IP address incorrectly or changed the IP address of the local host, you must manually edit the `commIpAddr.txt` file.

To change the IP address of the local host:

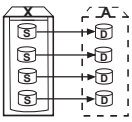
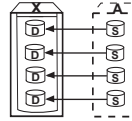
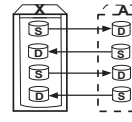
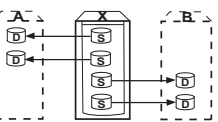
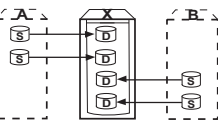
1. Open the `/opt/sanmgr/hostagent/config/commIpAddr.txt` file.
2. Change the IP address in the `commIpAddr.txt` file.
3. Restart the OVSAM hostagent.

## Troubleshooting storage problems

Deciding when to do a failover and which procedure to perform is complicated when you have multiple sites. With remote replication, each array is allowed two replication relationships. This means that an array may have source or destination virtual disks in DR groups that replicate to as many as two other arrays.

**Table 7** lists several replication situations. In each situation, “X” marks the array that receives the troubleshooting actions. Arrays marked “A” or “B” are remote arrays in a replication relationship with array X. Dotted lines depict remote arrays. Virtual disks in each array are marked either “S” (for source) or “D” (for destination). Match the situation closest to the environment you want to troubleshoot and go to the page listed for troubleshooting steps.

**Table 7: Identifying your troubleshooting situation**


Graphical depiction of your environment	Situation description	Location
	Situation 1—One source array with one relationship.	page 112
	Situation 2—One destination array with one relationship.	page 112
	Situation 3—An array has DR groups with one bidirectional relationship to one other array.	page 112
	Situation 4—A source array has DR groups with replication to two destination arrays.	page 113
	Situation 5—A destination array has DR groups with replication from two other arrays.	page 113

**Table 7: Identifying your troubleshooting situation**

Graphical depiction of your environment	Situation description	Location
	<p>Situation 6—An array has DR groups with replication from a source array and to a destination array.</p>	<p>page 114</p>
	<p>Situation 7—An array has DR groups with source and destination virtual disks on another array (site A) and only source virtual disks to a second array (site B).</p>	<p>page 114</p>
	<p>Situation 8—An array has DR groups with source and destination virtual disks on another array (site A) and only destination virtual disks from a second array (site B).</p>	<p>page 115</p>
	<p>Situation 9—An array has DR groups with source and destination virtual disks to two other arrays (sites A and B).</p>	<p>page 116</p>

### Situation 1—One source array with one relationship


If site X is not operational, perform a failover to continue operations at site A. See “[Unplanned failover](#)” on page 76.



If site X has had hardware failures, but it is still operating through redundant components, you do not need to move operations to site A. See your EVA documentation for further troubleshooting information.

### Situation 2—One destination array with one relationship

If site X is not operational, some or all of the DR groups are logging at site A or, if failsafe-enabled, are failsafe-locked.




If the DR groups are logging, logging continues at site A until the problem with site X is resolved.

If the DR groups are failsafe-locked at site A, and you want to continue operations, see “[Resumption of operations if unable to access destination while source in failsafe-locked state \(extended period of time\)](#)” on page 81.

To resolve the problem with site X, see your EVA documentation for further troubleshooting information. To replace the array with an uninitialized array, see “[Return operations to replaced new storage hardware](#)” on page 89. To resolve a disk group problem, see “[Disk group hardware failure on the destination array](#)” on page 67.

### Situation 3—An array has DR groups with one bidirectional relationship to one other array

If site X is not operational, some or all of the DR groups with sources at site A are logging at site A or, if failsafe-enabled, are failsafe-locked. The DR groups with sources at site X need to execute a failover to continue operations at site A.



#### *For the source DR groups located at site A:*

If the DR groups are failsafe-locked at site A, and you want to continue operations, see “[Resumption of operations if unable to access destination while source in failsafe-locked state \(extended period of time\)](#)” on page 81.

If the DR groups are logging, logging continues at site A until the problem with site X is resolved.




**For the destination DR groups located at site A:**

If site X is not operational, perform a failover to continue operations at site A (see “[Unplanned failover](#)” on page 76).

If site X has had hardware failures, but it is still operating through redundant components, you do not need to move operations to site A. See your EVA documentation for further troubleshooting information.


**Situation 4—A source array has DR groups with replication to two destination arrays**



If site X is not operational, perform a failover to continue operations at site A and site B (see “[Unplanned failover](#)” on page 76 to failover to site A). Repeat the procedure for site B.

If site X has had hardware failures, but it is still operating through redundant components, you do not need to move operations. See your EVA documentation for further troubleshooting information.

**Situation 5—A destination array has DR groups with replication from two other arrays**




If the site X is not operational, some or all of the DR groups are logging at site A or B or, if failsafe-enabled, are failsafe-locked.

If the DR groups are logging, logging continues at site A and B until the problem with site X is resolved.

If the DR groups are failsafe-locked at site A and B, and you want to continue operations, see “[Resumption of operations if unable to access destination while source in failsafe-locked state \(extended period of time\)](#)” on page 81 for site A. Repeat the procedure for site B.

Resolve the problem with site X (see your EVA documentation for further troubleshooting information). To replace the array with an uninitialized array, see “[Return operations to replaced new storage hardware](#)” on page 89. To resolve a disk group problem, see “[Disk group hardware failure on the destination array](#)” on page 67.

**Situation 6—An array has DR groups with replication from a source array and to a destination array**



If site X is not operational, some or all of the DR groups are logging at site A or, if failsafe-enabled, are failsafe-locked.

If the DR groups are logging, logging will continue at site A until the problem with site X is resolved.


If the DR groups are failsafe-locked at site A, and you want to continue operations, see [“Resumption of operations if unable to access destination while source in failsafe-locked state \(extended period of time\)”](#) on page 81.

Resolve the problem with site X (see your EVA documentation for further troubleshooting information). To replace the array with an uninitialized array, see [“Return operations to replaced new storage hardware”](#) on page 89. To resolve a disk group problem, see [“Disk group hardware failure on the destination array”](#) on page 67.

If site X is not operational, perform a failover to continue operations at site B (see [“Unplanned failover”](#) on page 76 for failover to site B).

If site X has had hardware failures, but is operating through redundant components, you do not need to move operations to site B (see your EVA documentation for further troubleshooting information).

**Situation 7—An array has DR groups with source and destination virtual disks on another array (site A) and only source virtual disks to a second array (site B)**



If site X is not operational, some or all of the DR groups with sources at site A are logging at site A or, if failsafe-enabled, are failsafe-locked. The DR groups with sources at site X need to execute a failover to continue operations at site A.

***For the source DR groups located at site A:***

If the DR groups are failsafe-locked at site A, and you want to continue operations, see [“Resumption of operations if unable to access destination while source in failsafe-locked state \(extended period of time\)”](#) on page 81.

If the DR groups are logging, logging continues at site A until the problem with site X is resolved.

***For the destination DR groups located at site A:***


If site X is not operational, perform a failover to continue operations at site A (see [“Unplanned failover”](#) on page 76 to failover to site A).

***For the destination DR groups located at site B:***

If site X is not operational, perform a failover to continue operations at site B (see [“Unplanned failover”](#) on page 76 to failover to site B).

If site X has had hardware failures, but is operating through redundant components, you do not need to move operations to site A or B (see your EVA documentation for further troubleshooting information).

**Situation 8—An array has DR groups with source and destination virtual disks on another array (site A) and only destination virtual disks from a second array (site B)**



If site X is not operational, some or all of the DR groups with sources at site A and B are logging at site A and B, or, if failsafe-enabled, are failsafe-locked. The DR groups with sources at site X need to execute a failover to continue operations at site A.

***For the source DR groups located at site A:***

If the DR groups are failsafe-locked at site A, and you want to continue operations, see [“Resumption of operations if unable to access destination while source in failsafe-locked state \(extended period of time\)”](#) on page 81.


If the DR groups are logging, logging continues at site A until the problem with site X is resolved.

***For the destination DR groups located at site A:***

If site X is not operational, perform a failover to continue operations at site A (see [“Unplanned failover”](#) on page 76).

If site X has had hardware failures, but is operating through redundant components, you do not need to move operations to site A (see your EVA documentation for further troubleshooting information).

**Situation 9—An array has DR groups with source and destination virtual disks to two other arrays (sites A and B)**



If site X is not operational, some or all of the DR groups with sources at site A or B are logging at site A or B, or, if failsafe-enabled, are failsafe-locked. The DR groups with sources at site X need to execute a failover to continue operations at site A.

***For the source DR groups located at site A and site B:***

If the DR groups are failsafe-locked at site A or B, and you want to continue operations, see [“Resumption of operations if unable to access destination while source in failsafe-locked state \(extended period of time\)”](#) on page 79 for site A. Repeat the procedure for site B.

If the DR groups are logging, logging continues at site A and B until the problem with site X is resolved.

***For the destination DR groups located at site A and site B:***

If site X is not operational, perform a failover to continue operations at the alternate sites. For failover to site A, see [“Unplanned failover”](#) on page 74. Repeat the procedure at site B.

If site X has had hardware failures, but is operating through redundant components, you do not need to move operations (see your EVA documentation for further troubleshooting information).

# Best practices

## 6

This chapter describes replication best practice procedures. It covers creating and using snapclones, bootless DR group planned failovers with Linux or SuSE Linux, throttling merge I/O after logging, and other miscellaneous best practices.

### Creating a destination snapclone before making a full copy

---

**Note:** A Business Copy EVA license is required for the following procedure.

---

When logging occurs on a source array, a temporary disparity occurs between data being processed at the local site and the data that exists at the remote location. A merge or full copy corrects this disparity later when the reason for the interruption is remedied. A merge sends data from the log disk in write order so it remains crash consistent. However, this is not the case with a full copy.

When a log fills to the point where it is marked for a full copy, there is a risk to the destination copy of the data once the process begins. This risk is due to the nature of a full copy, which copies data in 1-MB chunks starting at the beginning of the source virtual disk. This full-copy process does not maintain write ordering, and if it is not completed due to a second error, such as a loss of the source array, it leaves the destination array in an indeterminate state. Therefore, to prevent loss of data, best practice suggests creating a snapclone of destination virtual disks containing critical or important data prior to starting a full copy. If a major failure occurs at the local site during a full copy, the snapclone provides a clean copy of data as it existed before full copy writes were started to the destination array. However, any new writes that occurred on the source between the time the snapclone was created and the major failure would result in the loss of the new writes.

The following procedure describes the steps to take in a situation where you lose the connection between a source and destination array, and want to protect against a second failure when performing a full copy. Best practice suggests the creation of a destination snapclone whenever the link outage is expected to last more than several minutes.

---

**Note:** You cannot use this procedure if a full copy has been started.

---

1. Using the replication manager, navigate to each affected DR group and suspend replication.
2. When able, use the managing server to make a snapclone of the destination virtual disks, using the procedures described in HP StorageWorks Business Copy documentation.
3. Using the replication manager, navigate to each affected DR group and resume replication. This will only enable replication if the links are still down.

---

## Data movement using a snapclone

---

**Note:** A Business Copy EVA license is required for the following procedure.

---

Use this procedure to move a copy of your data residing on a virtual disk to a remote location by the use of a snapclone. A snapclone is an exact copy of your virtual disk at the particular point-in-time it was created. The virtual disk being copied to the remote site becomes part of a DR group that can then be used as a new source virtual disk. You can use this procedure with data movement services such as:

- **Data distribution**—Pushing copies of data to other geographic locations to make it locally accessible.
- **Data migration**—Moving data to a new location or to one with a larger storage capacity.

To move data using a snapclone:

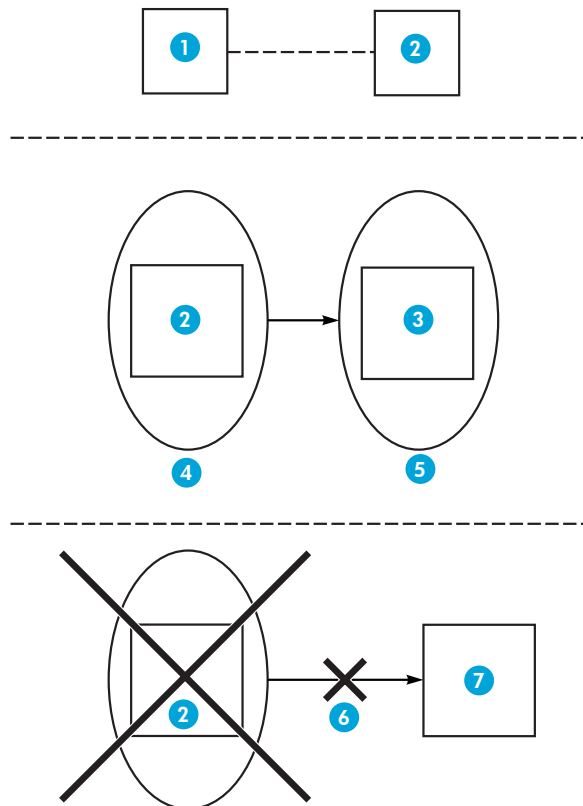
1. Make a snapclone of the virtual disk containing the data to be moved. See the online help for procedures on creating snapclones.

After the snapclone is created, the link from the snapclone to its original virtual disk dissolves, and the snapclone becomes a separate virtual disk.

2. Create a DR group with the new snapclone-created virtual disk linked to the remote array where you want the data to reside. The creation of the DR group replicates the virtual disk to your desired destination. For procedures on creating a DR group, see “Creating DR Groups” in Chapter 4.
3. Delete the source-cloned virtual disk. You have the option of keeping the remote virtual disk or deleting the remote virtual disk.
4. Choose to keep remote virtual disk.

The data now resides as a new virtual disk on the remote array. It can be used as a source for another DR group, subject to the restriction that an array can be involved in a replicating relationship with only one other array.

Figure 9 provides a high-level summary of the following steps that perform data movement using a snapclone.



CXO8068b

- ❶ HSV05 array
- ❷ HSV06 array
- ❸ HSV18 array
- ❹ DR group
- ❺ Virtual disk 1
- ❻ Virtual disk 2
- ❼ Replication
- ❽ Virtual disk snapclone

**Figure 9: Creating a DR group from a snapclone**



---

## Manually specifying disk group membership for a log

During the creation of a DR group, a 139-MB log is automatically created for the source and destination arrays. Placement of the log into a disk group on each system is based on the amount of free space in the disk group, the number of other logs in each disk group, the potential size of each existing log, and the potential size of the new log. It is possible that the disk group automatically selected for the log may not be the same disk group where the DR group resides. For example, disk groups containing near-online FATA drives are automatically selected for the location of the log. If near-online drives are not present, disk groups with the largest amount of average free log space are chosen. The location of the log is rarely an issue, but you can control the placement of a log into a specific disk group to separate the log from the data.

By default, a log is created in a near-online disk group. You cannot force a log into an online disk group if a near-online disk group exists. This restriction applies to both the source and destination arrays.

If possible, creating DR groups and specifying log disk location should be done before other non-DR virtual disks are created. Doing this ensures sufficient available space in the disk groups for a log.

If there is enough available space in the destination array to create the destination virtual disks and log disk, then use the following procedure to manually specify the disk group where a log will be created:

1. Create Vraid0 virtual disks on the destination array in all disk groups except the one that will be used for the DR group destination virtual disk and log, so that all available space is filled. These virtual disks are temporary and are not used for data storage.
2. Select each disk group to check the available space on the destination array. Each disk group should report zero space available except the disk group where the DR group destination virtual disk and log will be created.
3. Wait until the virtual disks have finished allocating space and are in a normal state.
4. Create Vraid0 virtual disks on the source array in all disk groups except the one that will be used for the DR group source virtual disk and log, so that all available space is filled. These virtual disks are temporary and will not be used for data storage.
5. Select each disk group to check the available space on the source array. Each disk group should report zero space available except the disk group where the DR group source virtual disk and log will be created

6. Wait until the virtual disks have finished allocating space and are in a normal state.
7. Create the DR group. Specify the disk group on the destination array where you want the destination virtual disk and log created. Do not allow the option to automatically select.
8. Delete the Vraid0 virtual disks that were created in step 1, and then delete those created in step 4.

The log disk is now located in the disk group you selected. The disk group membership of the log is not visible to HP Command View EVA or to the replication manager interface.

## Three-site cascaded data replication using snapclones

---

**Note:** A Business Copy license is required for the following procedure.

---

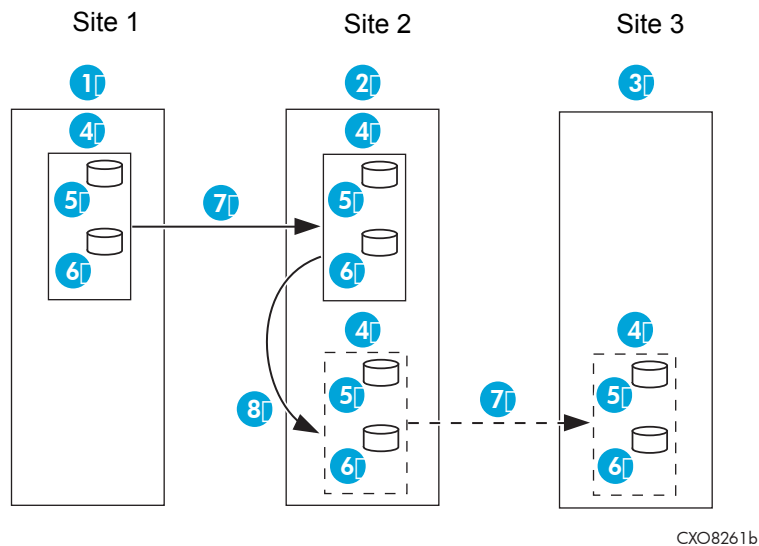
This procedure allows you to move copies of your data to a second remote location using HP Command View EVA and snapclones. The remote location can be an array without a replicating relationship to the array where the data was created. Exact copies of the virtual disks containing the data are created by using snapclones, and these are placed into a DR group for movement to the remote system.

For example, in [Figure 10](#) a production environment contains a DR group that replicates between arrays at Site 1 and Site 2. The DR group contains two virtual disks (05-06vdisk1 and 05-06vdisk2) that are to be archived on another array (Site 3). A snapclone of each virtual disk is created on the Site 2 array. After presentation to a host (set up only for presentation purposes, but required for the creation of a DR group), these members are added to a DR group called DR snapclone1. This DR group now resides on a source array that replicates to the desired destination array (Site 3). At the remote location, you can remove the virtual disk members from the DR group, renamed, and archived.

---

**Note:** For this procedure, Site 1 is called the source array, Site 2 (the destination for the DR group from Site 1) is called the intermediate array, and Site 3 is referred to as the remote array.

---



- |               |                          |
|---------------|--------------------------|
| ❶ HSV05 array | ❺ Virtual disk 1         |
| ❷ HSV06 array | ❻ Virtual disk 2         |
| ❸ HSV18 array | ❼ Replication            |
| ❹ DR group    | ❽ Virtual disk snapclone |

**Figure 10: Data movement using snapclones example**

You can perform cascaded replication using any of the following methods:

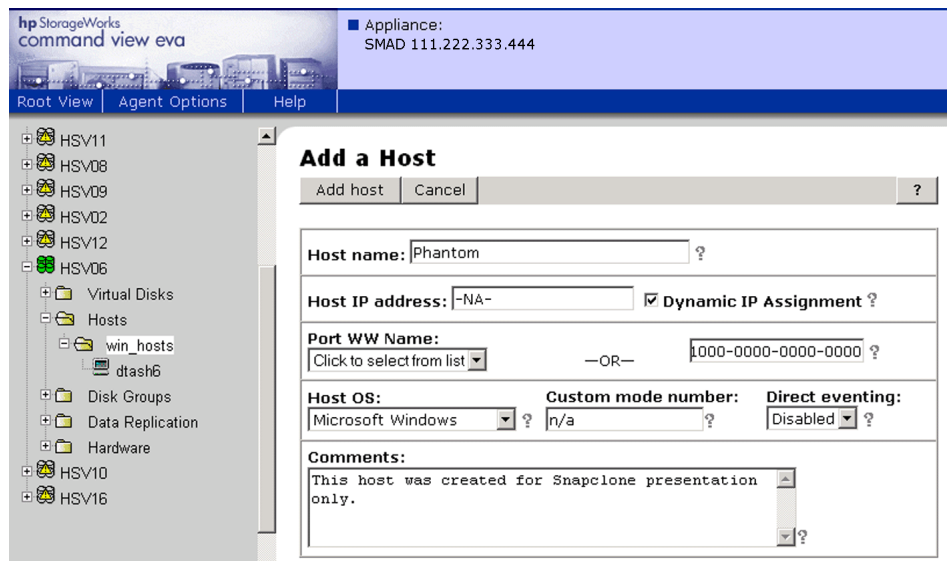
- Using the job template in the replication manager
- Manually in the replication manager
- Manually with the command line user interface (CLUI)

The following procedure describes the steps you must perform, regardless of the method you use. See the replication manager online help and the *HP StorageWorks Replication Solutions Manager Command Line User Interface reference guide* for additional instructions.


## Before you begin

Before you begin the procedure, ensure that you have done the following tasks:

1. Set up DR groups at the source (site 1 in [Figure 10](#)) and the destination (site 3 in [Figure 10](#)).
2. Set up the host on the intermediate site (site 2 in [Figure 10](#)) using HP Command View EVA:
  - a. In the Add a Host window, enter a host name in the Host name box and a WWN in the Port WW Name box. Click the **Add host** tab.



An Operation completed page is displayed.



## Procedure

1. Enable failsafe mode for any DR group containing more than one replication pair.
2. Set synchronous write mode for any DR groups in this procedure.
3. If normalization is occurring to members of the DR group to be moved, wait for the members to normalize.
4. If an application requires that I/O be suspended before creation of a snapclone, suspend I/O at this time.
5. Create a snapclone of each virtual disk on the intermediate array (Site 2 in [Figure 10](#)).
6. Set the DR group back to asynchronous, if applicable.
7. Set failsafe mode back to Disabled for the DR group, if applicable.
8. If the application was suspended in step 4, restart the host application.
9. Present the snapclone(s) to the host on the intermediate site (site 2 in [Figure 10](#)).
10. Place the snapcloned virtual disks into a new DR group.
11. Wait for normalization to complete.
12. Unpresent the host from the snapcloned virtual disks in the DR group on the intermediate array (Site 2 in [Figure 10](#)).
13. Remove the remaining snapcloned virtual disks from the DR group by deleting the DR group. Leave the remote virtual disks intact by not discarding them during the deletion.
14. Delete the remaining snapcloned virtual disks from the intermediate array (Site 2 in [Figure 10](#)).

## Post procedure

1. In HP Command View EVA, on the remote array (Site 3 in [Figure 10](#)), change the write protection of the mirrored snapcloned virtual disks to that of no write protection.
2. Rename the virtual disk to a useful name in the HP Command View EVA interface.

The virtual disks are now available on the remote array for any purpose.

---

## Bootless DR group planned failover with Linux using LVM in standalone mode or with SuSE SLES 8 running LifeKeeper 4.4.3

The following procedures describe how to perform a bootless DR group failover when running the Logical Volume Manager (LVM) with Linux. Separate procedures for running in standalone host mode or with clusters (LifeKeeper) are listed. Perform the procedures for the source host, followed by the procedures for the destination host.

---

**Note:** This procedure is not supported for unplanned failovers. The term “bootless” means that after the LUNs are first presented to a destination host, which requires an initial reboot, no further reboot of that host should be required.

---

### Source host procedure

Perform one of the following steps on the source host, depending on whether or not you are running LifeKeeper 4.4.3.

1. If you are running LifeKeeper 4.4.3, proceed to step 2. If you are not running LifeKeeper, perform the following steps:
  - a. From your source host, stop I/O to your LUNs. Allow enough time for the I/O to complete before proceeding to the next step.
  - b. Unmount the volumes contained in the DR group.  
Example: `umount /mounts/lv011`
  - c. Change the status of the LUNs to inactive with the following command:  
Example: `vgchange VolumeGroupName -a n`
  - d. Make the group unknown to the system with the `vgexport` command.  
Example: `vgexport vg01`
  - e. Perform a failover of the DR group using the Continuous Access user interface.

- f. Depending on the number of LUNs, do one of the following to prevent Secure Path from detecting a failed disk:
  - For individual LUNs, run `spmgr quiesce -p path` (for each path visible to the LUNs).
  - For all LUNs at once, run `spmgr set -p off`. This method will turn off path verification for all LUNs still visible to the system.
2. If you are running LifeKeeper 4.4.3 clusters:
  - a. Bring your resources “out of service” with the LifeKeeper GUI.
  - b. From the system console:
    - 1) Enter the `mount` command to verify the volume is unmounted.
    - 2) Enter the `vgscan` command to verify that the volume group was exported.
  - c. Perform a failover of the DR group with the Continuous Access user interface.
  - d. Depending on the number of LUNs, do one of the following to prevent Secure Path from detecting a failed disk:
    - For individual LUNs, run `spmgr quiesce -p path` (for each path visible to the LUNs).
    - For all LUNs at once, run `spmgr set -p off`. This method will turn off path verification for all LUNs still visible to the system.

## Destination host procedure

If this is the first time that LUNs are being presented to the destination host, reboot the host to pick up the new LUNs. If a reboot is not required (LUNs have been previously presented), and the paths are quiesced, issue the `spmgr restart all` command to unquiesce the paths.

Perform one of the following steps on the destination host, depending on whether or not you are running LifeKeeper 4.4.3.

1. If you are running LifeKeeper 4.4.3, proceed to step 2. If you are not running LifeKeeper, perform the following steps:
  - a. Issue the following command to make the volume known to the system:

```
vgimport VolumeGroupName PhysicalVolumePath
```

Example: `vgimport vg01 /dev/sda1`



- b. Mount the file systems.

Example: `mount -t reiserfs /dev/vg01/lvol1 /mounts/lvol1`

- c. Start host I/O.
- d. If the verification path is turned off, issue the following command:

```
spmgr set -p on
```

2. If you are running LifeKeeper 4.4.3 clusters:

- a. If this is the first time LUNs are being presented to the destination host, you must build the resource hierarchies for each new LUN presented (see the documentation on the LifeKeeper CD).
- b. Bring your resources “out of service” with the LifeKeeper GUI.
- c. Start host I/O.
- d. If the verification path is turned off, issue the following command:

```
spmgr set -p on
```

## Red Hat and SuSE Linux Lifekeeper clusters

Lifekeeper clusters must be zoned so that clustered hosts can see only one controller port per fabric. The operating system host mode of the controller must also be set to *custom*.

## Throttling of merge I/O after logging

When I/O has been halted, DR groups not in failsafe-enabled mode automatically resume replication when links to the remote arrays are restored. If there are dozens of DR groups with large logs, they compete for bandwidth as they try to synchronize simultaneously.

By suspending the merging or copying of non-critical DR groups, the controllers merge only the most critical data first, allowing this data to be synchronized and become accessible before the less important data. As the more important groups finish merging, resume the I/O of the groups that were suspended. This concentration or channeling of I/O to specific groups by the use of suspend and resume commands is called throttling I/O.

## Backing up replication jobs and configurations

HP recommends that you perform regular backups of jobs and configurations using the save and restore features in your replication products. This ensures that job and configuration data can be easily restored during planned or unplanned maintenance of the replication server. See

## Optimizing discovery refresh intervals

Use configuration settings in your replication products to optimize discovery refresh intervals. Replication products require up-to-date information on SAN resources. Find a balance between a discovery refresh interval that is too short (slowing overall performance with frequent discovery) and an interval that is too long (producing job failures due to out-of-date SAN information).

## Optimizing discovery performance

Use configuration settings in your replication products to eliminate discovery of resources that are infrequently used. Discovery of all resources in a large SAN can slow performance. For example, if a SAN includes 15 arrays, but you are only using local replication on 10 arrays, deselecting 5 will improve performance during discovery.

## Optimizing browser-based GUI performance

Keep simultaneous browser sessions to the same replication manager to a minimum. A large number of sessions decreases responsiveness of the replication manager.

## Coordinating enabled-host downtime

Ensure that planned downtime for enabled hosts is coordinated with replication jobs. A job will fail if any of a referenced host is not available when the job is run.

## Minimizing simultaneous jobs

If you are using jobs, minimize simultaneous job execution, even if the jobs involve different arrays. For example, running too many replication manager jobs at the same time can reduce the overall responsiveness of the replication manager and of other applications on the replication management server.

## Avoiding configuration changes while jobs are running

Avoid changing storage and host configurations while jobs are running. For example, don't change an array configuration with one interface (for example, HP Command View EVA) while replication jobs are running in another interface. Changing resources can lead to job failures and require manual intervention to restore resources to operational readiness.

## Optimizing the number of active enabled hosts

Keep the number of active enabled hosts to the minimum needed to perform required operations. Consider stopping host agent processes on hosts when they are not needed for jobs. If operational changes result in a host no longer being used in jobs, consider removing the host agent. Reducing the number of host agents that communicate with the local replication server will result in better overall performance of the replication management server.

## Coordinating enabled host shutdowns

Ensure that planned shutdowns are coordinated. Stopping an enabled host causes running jobs to fail if the job involves that host.

## Coordinating replication server shutdowns

Ensure that planned shutdowns are coordinated. Stopping a replication server:

- Stops any local replication applications on the server.
- Causes running jobs to fail.
- Prevents scheduled jobs from starting.

## Avoiding network identification changes

If possible, avoid changing the network identification (computer network name or IP address) of a replication management server or the computers on which replication host agents are running.

- **Server identification change**—If identification of a replication server is changed, use documented procedures to update all associated replication host agents so they can communicate with the replication management server. If the replication host agents are not updated to reflect the new replication server identification, jobs that involve the enabled host will fail.
- **Host agent identification change**—If identification of a replication host agent is changed, update impacted replication jobs so the jobs provide the correct references to enabled hosts. If the impacted jobs are not updated to reflect the new host agent identification, the jobs will fail.

## Maintaining network connections

Ensure that network connections between a replication management server and enabled hosts are maintained, especially while jobs are running. Jobs that interact with enabled hosts fail if the network connection is not maintained throughout the job.

## Using log files for troubleshooting jobs

The replication management server and host agents generate log files for job events. These detailed event log files can be helpful to HP support personnel when troubleshooting replication jobs.

## Making CD-ROMs of replication product Web download files

HP recommends that you make installation and archive CD-ROMs of local replication files that you download from the HP Web site. Making copies of Web download files ensures that you can quickly restore a replication management server and host agents to a given version without repeating download procedures.

## Managing replication events

Develop operation guidelines and best practices to address the following situations and concerns.

### Minimizing simultaneous replication events on an array

Minimize the number of replication requests to the same array at the same time. Consider limiting access to the various management and command line interfaces. Too many simultaneous replication events can reduce array performance.

### Avoiding simultaneous replication events for the same virtual disk

Avoid making multiple replication requests to the same virtual disk at the same time. Multiple replication events to the same virtual disk not only slow performance, but in the case of automated jobs, can lead to job failures. For example, if the maximum number of snapshots per virtual disk is exceeded when the job is running, the job will fail.

## Job scheduling

When using an external scheduler to schedule jobs, consider the timing of each job relative to other jobs that involve the array and host resources. Tune and load-balance demands to maximize performance.

## Complying with EVA snapshot rules

The following general EVA snapshot rules apply:

- The array must have a local replication license.
- Each snapshot is created in the same disk family as the source virtual disk.
- All snapshots of a given virtual disk must have the same allocation policy.
- No more than seven snapshots of a given virtual disk can exist at one time.
- When managing array resources, snapshots are counted as a virtual disks.

Snapshots cannot be made when a storage volume:

- Is itself a snapshot.
- Is in the process of unsharing or being deleted.

When a local replication interface indicates that a EVA storage volume (virtual disk) does not support snapshot replication, or if a snapclone script action fails in a job, it is probable that some of these rules have been violated.

## Complying with EVA snapclone rules

The following general EVA snapclone rules apply:

- The array must have a local replication license.
- Each snapclone is created in the same disk group as its source, but in a different disk family.

Snapclones cannot be made when a storage volume:

- Is itself a snapshot or has a snapshot.
- Is in the process of unsharing or being deleted.

When a local replication interface indicates that an EVA storage volume (virtual disk) does not support snapclone replication, or if a snapclone script action fails in a job, it is probable that some of these rules have been violated.

## Caching in Microsoft Windows

Small files in Microsoft Windows can be held in cache, disrupting replication to the remote controller. Flush all cache files, if possible, before performing a failover. One source of information for flushing data caches on CPU and kernel architecture can be obtained from:

<http://msdn.microsoft.com/library/en-us/wcedsn40/html/cgconimplementingcacheflushroutines.asp>

Another option is to use the HP StorageWorks Business Copy EVA application to flush the cache. For more information, go to:

<http://h18006.www1.hp.com/products/storage/software/bizcopyeva/index.html>

---

**Note:** Rebooting of the source host(s) is the only qualified procedure at this time.

---

## Asynchronous replication with failsafe enabled

Running in asynchronous replication mode with failsafe enabled is not supported. There is no benefit when using these two modes together in normal operation, and doing so may induce LUN instability after the loss of intersite links.





# Event message descriptions



This appendix describes event messages that you may encounter in the event pane.

## Array messages

The following messages are generated by the background server discovery process.

{1} is the name of the array.

**Table 8: Array messages**

Description
Invalid license on {1}.
Licensed capacity exceeded for {1}.
License will expire within 15 days for {1}.
Licenses on {1} are valid.
Discovered storage array {1} WWN.
Ignoring unmanaged storage array {1} WWN.
Retrieved data for storage array {1} WWN.
Refresh of data failing for storage array {1} WWN.

## DR group messages

The following messages are generated by the background server discovery process.

{1} is the name of the DR group.

{2}, {3} is additional information to describe the event.

**Table 9: DR group messages**

Description
{1} was created.
{1} was deleted.
The failsafe state of {2} has been {1}.
The replication mode of {1} has changed from {2} to {3}.
{1} has been suspended.
{1} has been resumed.
{1} is constructing.
{1} is copying.
{1} is degraded.
{1} is disabled.
{1} is failed.
{1} is normal.
{1} state is unknown.
{2} was added to {1}.
{2} was removed from {1}.

## Job messages

The following messages are generated by the server job engine.

{1} is the name of the job.

{2}, {3}, {4} is additional information to describe the event.

**Table 10: Job messages**

Description
Operation started.
Operation complete.
Operation cancelled.
Operation failed.
Job created {1}.
Job started {1}.
Job complete {1}.
Job cancelled {1}.
Job failed {1}.
Job deleted {1}.
Job paused {1}.
Cancelling Job.
Internal error occurred starting job {1}.
Starting job {1}.
Startup failed for job {1}.
Job failed runtime validation checks {1}.
Failed to find description for job {1}.
Failed to find specification for job {1}.
Invalid branch target in job {1}.
Job validation failed in job {1}.
Job preprocessing failed in job {1}.
Failed to process task {2} in job {1}.
Argument {2} was expected to be type {3} in job {1}.
Task {2} has incorrect number of arguments in job {1}.

**Table 10: Job messages (continued)**

Description
Argument {2} may not be a variable in job {1}.
Failed to find connection object for consistency set {2} in job {1}.
Failed to find storage system {2} in job {1}.
Failed to validate argument {2} for operation {3} in job {1}.
Incorrect number of arguments for operation {2} in job {1}.
Database error trying to validate arguments for operation {2} in job {1}.
Operation {2} is not yet implemented in job {1}.
Failed to retrieve variable {2} for operation {3} in job {1}.
Launching host command {2} on host {3} for job {1}.
Host command succeeded returning {2} for job {1}.
Host command failed with code {2} ({3}) for job {1}.
Delaying job until time/date {2} for job {1}.
Resuming job {1} after wait.
Wait operation failed with code {2} ({3}) for job {1}.
Argument value {2} has invalid format in job {1}.
Starting discovery refresh for job {1}.
Completed discovery refresh for job {1}.
Discovery refresh failed with code {2} ({3}) for job {1}.
Starting host agent discovery refresh for job {1}.
Completed host agent discovery refresh for job {1}.
Host agent discovery refresh failed with code {2} ({3}) for job {1}.
Pausing job {1}.
Exiting job {1} with condition {2}.
Delaying until time/date {2} in job {1}.
Resuming job {1} after delay.
Error parsing date/time ({2}) in Latently operation for job {1}.
Starting operation for job {1}.
Operation completed successfully for job {1}.
Unable to locate DR group {2} for job {1}.

**Table 10: Job messages (continued)**

Description
Creating presentation target for job {1}.
Successfully created presentation target for job {1}.
Creation of presentation target failed—already exists for job {1}.
Invalid RAID level specified for job {1}.
Object not found failure for job {1}.
Invalid copy type specified for job {1}.
Deleting DR group for job {1}.
Successfully deleted DR group for job {1}.
Delete DR group failed for job {1}.
Failed to delete DR group because storage system was not available for job {1}.
Failed to delete DR group because neither storage system was available for job {1}.
Could not delete remote volume for DR group because one of the storage systems is unavailable for job {1}.
Failed to delete remote volume for DR group for job {1}.
Deleting DR group for job {1}.
Successfully deleted DR group for job {1}.
Delete DR group failed for job {1}.
Invalid delete mode for job {1}.
Deleting host volume started for job {1}.
Successfully deleted host volume for job {1}.
Delete host volume failed for job {1}.
Deleting storage volume for job {1}.
Successfully deleted storage volume for job {1}.
Deleting storage volumes for job {1}.
Successfully deleted storage volumes for job {1}.
Delete volume operation started for job {1}.
Delete volume operation successful for job {1}.
Unable to delete volume for job {1}.
Failed to write XML data to file for job {1}.

**Table 10: Job messages (continued)**

Description
Exporting XML data to file {2} for job {1}.
Exported {2} jobs for job {1}.
Failed to export jobs—job not found for job {1}.
Exported {2} host agents for job {1}.
Failed to export host agents for job {1}.
Exported {2} managed sets for job {1}.
Failed to get members of managed set {2} for job {1}.
Failed to get managed set list for job {1}.
Failed to export managed set member {2} for job {1}.
Exported storage management servers for job {1}.
Failed to export storage management servers for job {1}.
Exported storage licenses to {2} for job {1}.
Skipped license export—unable to find the license file ({2}) for job {1}.
Starting a fast failover operation on storage system for job {1}.
Completed fast failover operation on storage system for job {1}.
Starting get source UNC DR group name operation on storage system for job {1}.
Completed get source UNC DR group name operation on storage system for job {1}.
Starting get virtual disk for host storage volume operation for job {1}.
Completed get virtual disk for host storage volume operation for job {1}.
Get virtual disk for host storage volume failed for direct attach volume for job {1}.
Import failed for job {1}.
Failed to read XML data from file for job {1}.
Importing XML data from file {2} for job {1}.
Imported {2} jobs for job {1}.
Failed to import jobs—database error for job {1}.
Imported {2} host agents for job {1}.
Failed to import host agents for job {1}.
Failed to refresh host agent data for job {1}.

**Table 10: Job messages (continued)**

Description
Imported {2} managed sets for job {1}.
Failed to import managed sets for job {1}.
Invalid managed set type detected for job {1}.
Imported storage management servers for job {1}.
Failed to import storage management servers for job {1}.
Imported licenses from {2}.
License import (Autopass) of file ({2}) failed for job {1}.
License import failed—unable to initialize Autopass module for job {1}.
License import failed—file ({2}) does not have XML extension for job {1}.
Launch started for job {1}.
Launch successful for job {1}.
Launch failed for job {1}.
Object not found for job {1}.
Starting mount operation for job {1}.
Mount operation completed successfully for job {1}.
Mount failed for job {1}.
Presenting storage volume {2} for job {1}.
Presentation of storage volume {2} succeeded for job {1}.
Presentation of storage volume {2} failed with code {3} ({4}) for job {1}.
Presentation host {2} not found for job {1}.
Virtual disk {2} already presented to {3}; representing for job {1}.
Unpresenting storage volume {2} for job {1}.
Unpresentation of storage volume {2} succeeded for job {1}.
Unpresentation of storage volume {2} failed with code {3} ({4}) for job {1}.
Starting post present storage volume operation for job {1}.
Post present storage volume successful for job {1}.
Post present storage volumes failed for job {1}.
Starting post unpresent host volumes operation for job {1}.
Post unpresent host volumes successful for job {1}.

**Table 10: Job messages (continued)**

Description
Post unrepresent host volumes failed for job {1}.
Starting pre unrepresent host volumes operation for job {1}.
Pre unrepresent host volumes successful for job {1}.
Pre unrepresent host volumes failed for job {1}.
Presenting storage volumes to host {2} for job
Presentation of storage volumes to host {2} succeeded for job {1}.
Presentation of storage volumes to host {2} failed with code {3} ({4}) for job {1}.
Unrepresenting storage volumes to host {2} for job.
Unrepresentation of storage volumes to host {2} succeeded for job {1}.
Unrepresentation of storage volumes to host {2} failed with code {3} ({4}) for job {1}.
Removing member from DR group for job {1}.
Successfully removed member from DR group for job {1}.
Resuming DR group for job {1}.
Successfully resumed DR group for job {1}.
Failed to resume DR group for job {1}.
Setting destination modes for DR group for job {1}.
Successfully set destination modes for job {1}.
Setting source modes for DR group for job {1}.
Successfully set source modes for job {1}.
Starting set DR group comments operation for job {1}.
DR group set comments completed successfully for job {1}.
Unable to locate DR group for job {1}.
Starting set DR group destination access operation for job {1}.
DR group set destination access completed successfully for job {1}.
Invalid destination access mode specified ({2}) for job {1}.
Unable to locate DR group {2} for job {1}.
Starting set DR group failsafe operation for job {1}.
DR group set failsafe completed successfully for job {1}.
Invalid failsafe mode specified ({2}) for job {1}.



**Table 10: Job messages (continued)**

Description
Unable to locate DR group {2} for job {1}.
Starting set DR group IO mode operation for job {1}.
DR group set IO mode completed successfully for job {1}.
Invalid IO mode specified ({2}) for job {1}.
Unable to locate DR group {2} for job {1}.
Starting set DR group suspend operation for job {1}.
DR group set suspend completed successfully for job {1}.
Invalid suspend mode specified ({2}) for job {1}.
Unable to locate DR group connection for job {1}.
Unable to locate DR group {2} for job {1}.
Starting local replication for job {1}.
Successfully started local replication for job {1}.
Starting remote replication for job {1}.
Successfully started remote replication for job {1}.
Start replica operation started for job {1}.
Start replica operation successful for job {1}.
Unable to start replica for job {1}.
Starting replication for job {1}.
Successfully started replication for job {1}.
Start of replication failed for job {1}.
Suspending DR group started for job {1}.
Successfully suspended DR group for job {1}.
Suspend/resume DR Group started for job {1}.
Suspend/resume successful for DR group for job {1}.
Starting unmount operation for job {1}.
Unmount operation completed successfully for job {1}.
Unmount failed for job {1}.
Unpresent storage volume started for job {1}.
Unpresent storage volume successful for job {1}.

**Table 10: Job messages (continued)**

Description
Validate snapclone host volume started for job {1}.
Validate snapclone host volume successful for job {1}.
Validation failed for snapclone host volume for job {1}.
Validate snapclone storage volume started for job {1}.
Validate snapclone storage volume successful for job {1}.
Validation failed for snapclone storage volume for job {1}.
Validate snapshot host volume started for job {1}.
Validate snapshot host volume successful for job {1}.
Validation failed for snapshot host volume for job {1}.
Validate snapshot storage volume started for job {1}.
Validate snapshot storage volume successful for job {1}.
Validation failed for snapshot storage volume for job {1}.
Validate storage system started for job {1}.
Validate storage system successful for job {1}.
Validation failed for storage system for job {1}.
Validate storage volume started for job {1}.
Validate storage volume successful for job {1}.
Validation failed for storage volume for job {1}.
Validate host started for job {1}.
Validate host successful for job {1}.
Validation failed for host for job {1}.
Validate host volume started for job {1}.
Validate host volume successful for job {1}.
Validation failed for host volume for job {1}.
Validate not host volume started for job {1}.
Validate not host volume successful for job {1}.
Validation failed for not host volume for job {1}.
Wait for DR group normalization started for job {1}.
Wait for DR group normalization completed successfully for job {1}.

**Table 10: Job messages (continued)**

Description
Wait for DR group normalization failed for job {1}.
Wait for remote replica to complete started for job {1}.
Wait for remote replica to complete has completed successfully for job {1}.
Wait for remote replica to complete has failed for job {1}.
Wait for storage volume to become available started for job {1}.
Wait for storage volume to become available completed successfully for job {1}.
Wait for storage volume to become available failed for job {1}.
Waiting for connection for DR group: {2} to become available {1}.
Wait for storage volume to normalize started for job {1}.
Wait for storage volume to normalize completed successfully for job {1}.
Wait for storage volume to normalize failed for job {1}.
Failed to set task return value {2} for job {1}.
No task return variable defined for job {1}.
Create DR group started for job {1}.
Create DR group succeeded for job {1}.
Create DR group failed—invalid IO mode {2} for job {1}.
Remove DR group member started for job {1}.
Remove DR group member succeeded for job {1}.
Remove DR group member failed for job {1}.
Remove DR group member failed—invalid delete mode {2} for job {1}.

## Error messages

The following messages are generated by server exceptions. These are general errors and several different actions could cause these messages.

{0} is the name of the storage resource that is the source of the error.

{1}, {2} is additional information to describe the event.

**Table 11: Error messages**

Description
EVA simulator failure: {2} {1} {0}.
EVA simulator storage volume not found failure: {2} {1} {0}.
Remote exception occurred: {2} {1} {0}.
Method not supported: {2} {1} {0}.
Not implemented in presentation set data: {2} {1} {0}.
Not supported in base DR group: {2} {1} {0}.
Refresh not implemented in base storage system: {2} {1} {0}.
Not available: {2} {1} {0}.
No information yet: {2} {1} {0}.
Method not implemented in device DR group: {2} {1} {0}.
Refresh not implemented in device management path: {2} {1} {0}.
Performance monitor not available: {2} {1} {0}.
Get updating refresh count not implemented in device storage system: {2} {1} {0}.
Method not implemented in device presentation set: {2} {1} {0}.
Device log destination enum list required: {2} {1} {0}.
Storage system not located: {2} {1} {0}.
Failed to save device configuration to file: {2} {1} {0}.
Remote method invocation error: {2} {1} {0}.
Proxy call failed: {2} {1} {0}.
Get performance monitor exception: {2} {1} {0}.
Unexpected exception: {2} {1} {0}.
Failed to load device configuration from file: {2} {1} {0}.
No targets specified: {2} {1} {0}.

**Table 11: Error messages (continued)**

Description
Removal of existing presentations failed: {2} {1} {0}.
Target not found: {2} {1} {0}.
Controller configuration service not located: {2} {1} {0}.
Presentation failed. Cannot set LUN access: {2} {1} {0}.
Presentation not confirmed on array: {2} {1} {0}.
Cannot locate presentation info for host: {2} {1} {0}.
Cannot locate presented unit interface for host: {2} {1} {0}.
Unsupported copy type: {2} {1} {0}.
Storage volume does not belong to this storage system: {2} {1} {0}.
Disk group not from target storage system: {2} {1} {0}.
Unable to locate target storage system: {2} {1} {0}.
Unable to locate connection: {2} {1} {0}.
System element not found: {2} {1} {0}.
Target storage system not specified: {2} {1} {0}.
Source storage system not specified: {2} {1} {0}.
Disk group not found: {2} {1} {0}.
CA Basic exception: {2} {1} {0}.
Cannot create replica: {2} {1} {0}.
Cannot create DR group: {2} {1} {0}.
DR groups have no members: {2} {1} {0}.
DR groups must have equal members: {2} {1} {0}.
Replication failed: {2} {1} {0}.
Cannot make copy of snapshot: {2} {1} {0}.
Maximum number of snapshots exceeded: {2} {1} {0}.
All snapshots must have same copy type: {2} {1} {0}.
Delete DR group failed: {2} {1} {0}.
Method not implemented: {2} {1} {0}.
Bad management path ID: {2} {1} {0}.
No management path info: {2} {1} {0}.

**Table 11: Error messages (continued)**

Description
DR group not found: {2} {1} {0}.
Presentation not found: {2} {1} {0}.
DR group not found: {2} {1} {0}.
Host Fibre Channel port not defined: {2} {1} {0}.
Storage modification exception: {2} {1} {0}.
Method only valid for remote replication: {2} {1} {0}.
Detach not available: {2} {1} {0}.
Mode not supported: {2} {1} {0}.
Failed to set Home: {2} {1} {0}.
Snaps found—cannot delete volume: {2} {1} {0}.
Member of a remote copy—cannot delete volume: {2} {1} {0}.
DR EVA exception: {2} {1} {0}.
Host name already in use: {2} {1} {0}.
Port WWN already in use: {2} {1} {0}.
Resume operation failed: {2} {1} {0}.
Unable to set comments for DR group: {2} {1} {0}.
SMI-S virtual disk already a member of DR group: {2} {1} {0}.
Storage pool path not found: {2} {1} {0}.
Failed to locate storage remote replication service: {2} {1} {0}.
Transaction failed due to element manager error: {2} {1} {0}.
Exception reported: {2} {1} {0}.
Remove member delete replica not valid for DR group: {2} {1} {0}.
Set comments failed for DR group: {2} {1} {0}.
Metadata not available: {2} {1} {0}.
Set metadata failed for DR group: {2} {1} {0}.
SMI-S failed to locate source DR group: {2} {1} {0}.
SMI-S failed to locate destination DR group: {2} {1} {0}.
SMI-S bad intermediate state two destination DR group: {2} {1} {0}.
SMI-S bad intermediate state two source DR group: {2} {1} {0}.

**Table 11: Error messages (continued)**

Description
SMI-S destination virtual disk presented: {2} {1} {0}.
SMI-S virtual disk not member of a DR group: {2} {1} {0}.
Modify DR group failed—cannot locate remote replication service: {2} {1} {0}.
SMI-S storage array does not support this LUN access enum: {2} {1} {0}.
SMI-S no information yet: {2} {1} {0}.
Modify DR group set failed with element manager error: {2} {1} {0}.
SMI-S remove virtual disk failed to get storage volume: {2} {1} {0}.
Device exception: {2} {1} {0}.
SMI-S error—creating storage hardware—ID failed: {2} {1} {0}.
Create presentation target failed: {2} {1} {0}.
SMI-S error—create protocol controller failed: {2} {1} {0}.
View protocol controller ref not found: {2} {1} {0}.
SMI-S error—assign access failed: {2} {1} {0}.
SMI-S error—invoke attach device: {2} {1} {0}.
SMI-S error—invoke detach device: {2} {1} {0}.
Storage configuration service not located: {2} {1} {0}.
Storage capabilities not located: {2} {1} {0}.
Replication failed—element manager error: {2} {1} {0}.
Element manager interface disabled: {2} {1} {0}.
Set sync write mode failed—element manager error failure: {2} {1} {0}.
Virtual disk delete failed—element manager error failure: {2} {1} {0}.
Storage hardware ID—management service not located: {2} {1} {0}.
Delete presentation target failed.
Privilege management service not located: {2} {1} {0}.
Suspend failed—element manager error failure: {2} {1} {0}.
Resume failed—element manager error failure: {2} {1} {0}.
Unable to set comments for DR group: {2} {1} {0}.
No information yet: {2} {1} {0}.
MSA manager disabled: {2} {1} {0}.

**Table 11: Error messages (continued)**

Description
MSA manager connection failed: {2} {1} {0}.
MSA virtual disk not found: {2} {1} {0}.
Unable to set comments for DR group: {2} {1} {0}.
Set name failed for storage pool—element manager interface error: {2} {1} {0}.
Set name failed for storage pool—appliance connection failed: {2} {1} {0}.
Virtual disk not found: {2} {1} {0}.
Virtual disk is already member of a DR group: {2} {1} {0}.
Data replication failed due to appliance connection failure: {2} {1} {0}.
Data replication failed while adding virtual disk: {2} {1} {0}.
Transaction failed due to bad appliance: {2} {1} {0}.
Set name failed for DR group—device error: {2} {1} {0}.
Set name failed for DR group—appliance connection failed: {2} {1} {0}.
Set comments failed for DR group—device error: {2} {1} {0}.
Set comments failed for DR group—bad appliance: {2} {1} {0}.
Device metadata not available: {2} {1} {0}.
Set metadata failed for DR group—device error: {2} {1} {0}.
Set metadata failed for DR group—bad appliance: {2} {1} {0}.
Device bad intermediate state two source DR group: {2} {1} {0}.
Device failed to locate source DR group: {2} {1} {0}.
Device failed to locate destination DR group: {2} {1} {0}.
Exception initializing Home: {2} {1} {0}.
Device bad intermediate state two destination DR group: {2} {1} {0}.
Device destination virtual disk presented: {2} {1} {0}.
Device virtual disk is not a member of a DR group: {2} {1} {0}.
Data replication failed while removing virtual disk: {2} {1} {0}.
Set name failed for virtual disk—device error: {2} {1} {0}.
Set name failed for virtual disk—appliance connection failed: {2} {1} {0}.
Set comments failed for virtual disk—device error: {2} {1} {0}.
Set comments failed for virtual disk—appliance connection failed: {2} {1} {0}.



**Table 11: Error messages (continued)**

Description
Virtual disk member of destination DR group—LUN access cannot be set: {2} {1} {0}.
Appliance connection failed—LUN access cannot be set: {2} {1} {0}.
Set LUN access failed—device error: {2} {1} {0}.
Set LUN access transaction failed due to bad appliance: {2} {1} {0}.
Device handler is disabled: {2} {1} {0}.
Unable to locate storage system: {2} {1} {0}.
Delete presentation target—device error: {2} {1} {0}.
Appliance is disabled: {2} {1} {0}.
Set name failed for storage system—device error: {2} {1} {0}.
Set name failed for storage system—appliance connection failed: {2} {1} {0}.
Set comments failed for storage system—device error: {2} {1} {0}.
Set comments failed for storage system—appliance connection failed: {2} {1} {0}.
Data replication failed—device error: {2} {1} {0}.
Set comments failed for storage pool—appliance connection failed: {2} {1} {0}.
Set comments failed for storage pool—device error: {2} {1} {0}.
Device removal of virtual disk failed to get storage volume: {2} {1} {0}.
Appliance connection failed: {2} {1} {0}.
Present failed—element manager interface error: {2} {1} {0}.
Unpresent failed—element manager interface error: {2} {1} {0}.
Element manager interface error—no such element exception: {2} {1} {0}.
Replication failed—element manager interface error: {2} {1} {0}.
Delete virtual disk failed—element manager interface error: {2} {1} {0}.
Failover failed—element manager interface error: {2} {1} {0}.
Delete DR group failed—element manager interface error: {2} {1} {0}.
Suspend failed—element manager interface error: {2} {1} {0}.
Resume failed—element manager interface error: {2} {1} {0}.
Set failsafe mode failed—element manager interface error: {2} {1} {0}.
Set synchronous write mode failed—element manager interface error: {2} {1} {0}.

**Table 11: Error messages (continued)**

Description
Set access mode failed—element manager interface error: {2} {1} {0}.
Unable to set comments for DR group: {2} {1} {0}.
No information yet: {2} {1} {0}.
Unknown exception in HAL layer: {2} {1} {0}.
Remote exception occurred: {2} {1} {0}.
Validation failed for operation: {2} {1} {0}.
Unknown exception in transport layer: {2} {1} {0}.
I/O error during transport: {2} {1} {0}.
XML response empty when not expected to be: {2} {1} {0}.
XML in a bad format: {2} {1} {0}.
Unknown exception in java layer on host: {2} {1} {0}.
Cannot generate the XML response on host: {2} {1} {0}.
Error loading java library on host: {2} {1} {0}.
Invalid input to request on host: {2} {1} {0}.
Unknown exception on host: {2} {1} {0}.
Mount operation failed on host: {2} {1} {0}.
Unmount operation failed on host: {2} {1} {0}.
Storage system {1} required for this action was not found in the database. This is probably due to a corrupt database. Please contact customer support.
The presentation target object {0} was not found.
The DR group object {0} was not found.
The connection object {0} was not found.
The presentation object {0} was not found.
The virtual disk {0} was not found.
The storage pool {0} was not found.
The performance monitor, index {0}, was not found.
The host agent {0} was not found.
The host bus adaptor {0} was not found.
The host volume {0} was not found.

**Table 11: Error messages (continued)**

Description
The mount point {0} was not found.
Unknown type {0}.
Invalid parameter(s) for {0}.
Command-line UI error: {0}.
Parsing error encountered: {0}.
SQL error: {0}.
Unknown host {0}.

## Dialog box error messages

The following general error messages are displayed in a GUI dialog box. Several different actions could cause one of these messages.

{0} is the name of the storage resource that is the source of the error.

{1}, {2} is additional information to describe the event.

**Table 12: Dialog box error messages**

Description
Error code not localized for exception: {2} {1} {0}.
The {0} method is not implemented.
Method {0} not found.
Class {0} not found.
Could not instantiate {0}.
Illegal access to {0}.
Error invoking {0}.
Invalid parameter {0}.
Remote error: {0}
Security error: {0}.
Not all of the storage systems could be refreshed. See the storage systems resource for the current state.
Operation failed with exception: {2} for job {1} {0}.



This glossary defines terms used in this guide or related to this product and is not a comprehensive glossary of computer terms.

**active management server**

*See* management server.

**array**

*See* virtual array and storage system.

**asynchronous**

A descriptive term for computing models that eliminate timing dependencies between sequential processes. In asynchronous replication, the array controller acknowledges that data has been written at the source before the data is copied at the destination. Asynchronous replication is an optional DR group property. *See also* synchronous.

**bidirectional**

A descriptive term for an array that contains both source and destination virtual disks. This configuration allows multidirectional I/O flow among several arrays

**B-series switches**

Fibre Channel core and SAN switches made by Brocade and sold by HP.

**C-series switches**

Fibre Channel switches made by Cisco and sold by HP.

**client**

An intelligent device that requests services from other intelligent devices. In the context of HP StorageWorks Replication Manager, a client is a computer that is used to access the replication manager remotely using a supported browser.

**HP Continuous Access EVA**

HP Continuous Access EVA is a storage-based HP StorageWorks product consisting of two or more storage systems performing disk-to-disk replication, along with the management user interfaces that facilitates configuring, monitoring, and maintaining the replicating capabilities of the storage systems.

**default disk group**

The disk group that is created when an array is initialized. The minimum number of disks the group can contain is eight. The maximum is the number of installed disks.

**destination**

The targeted recipient (for example, a DR group, array, virtual disk) of replicated data. See also source.

**disk group**

A named group of disks selected from all the available disks in an array. One or more virtual disks can be created from a disk group.

**DR group**

Data replication group. A named group of virtual disks selected from one or more disk groups so that they replicate to the same destination, fail over together, and preserve write order within the group.

**dual fabric**

Two independent fabrics providing multipath connections between Fibre Channel end devices.

**EVA**

Enterprise Virtual Array, an HP StorageWorks product that consists of one or more virtual arrays. See also virtual arrays.

**event**

A system-generated status message, resulting from a:

- User-initiated action (for example, “suspend DR group”)
- Replication or system transaction (for example, “retrieved data for storage system”)
- Job operation (for example, “job complete”)

**fabric**

A network of Fibre Channel switches or hubs and other devices.

**failover**

An operation that reverses replication direction so that the destination becomes the source and the source becomes the destination. Failovers can be planned or unplanned and can occur between DR groups, managed sets, fabrics or paths, and array controllers.

**failsafe**

A descriptive term for devices that automatically assume a safe condition after a malfunction. Failsafe DR groups stop accepting host input and stop logging write history if a member of the group becomes unreachable.

**general purpose server**

A server that runs customer applications such as file and print services. HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager can be used on a general purpose server in limited configurations.

**Home**

The DR group that is the preferred source in a replication relationship. By default, Home is the original source, but it can be set to the destination DR group.

**host**

A computer that runs user applications and uses the information stored on an array.

**host volume**

Storage capacity that is defined and mountable by a host operating system. In HP Replication Solutions Manager, host volumes are disks or volumes that are reported by an enabled host.

**initialization**

A configuration step that binds the controllers together and establishes preliminary data structures on the array. Initialization also sets up the first disk group, called the default disk group, and makes the array ready for use.

**LUN**

Logical unit number. Logical units are the components within SCSI targets that execute I/O commands. Virtual disks that are presented to hosts correspond to logical units and are identified by LUN IDs. See also present.

**M-series switches**

Fibre Channel Director and Edge switches made by McDATA and sold by HP.

**managed set**

Selected resources grouped together for convenient management. For example, you can create a managed set to manage all DR groups whose sources reside in the same rack.

**management server**

A server where HP StorageWorks Enterprise Virtual Array (EVA) management software is installed, including HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager, if used. A dedicated management server runs EVA management software exclusively. Other management servers are general purpose servers, HP ProLiant Storage Server (NAS) models, and the HP OpenView Storage Management Appliance.

When there are multiple management servers in a SAN, one is active and all others are standby. The active management server actively manages the array, while the standby management server takes control of the array if there is a failure on the active management server. There is only one active management server at a time for any given management zone in a SAN.

**merge**

The act of transferring log contents to the destination virtual disk to synchronize the source and destination.

**mount point**

The file system path and directory where a host volume is accessed.

**normalization**

The initial full copy that occurs between source and destination virtual disks.

**(to) present**

The array controller act of making a virtual disk accessible to a host computer.

**remote copy**

A replica virtual disk on the destination array.

**resource**

An object in the Replication Solutions Manager navigation pane; namely, DR groups, enabled hosts, host volumes, managed sets, storage systems, and virtual disks. Replication is performed using these resources.

**SAN**

Storage area network, a network of storage devices and the initiators that store and retrieve information on those devices, including the communication infrastructure.

**snapclone**

A copy that begins as a fully allocated snapshot and becomes an independent virtual disk. Applies only to the HP StorageWorks EVA.

**snapshot**

A nearly instantaneous copy of the contents of a virtual disk created without interruption of operations on the source virtual disk. Snapshots are typically used for short-term tasks such as backups.

**source (Home)**

A descriptive term for the virtual disk, DR group, or virtual array where an original I/O is stored before replication. See also destination.

**standby management server**

*See* management server.

**Storage Management Appliance**

HP OpenView Storage Management Appliance, an HP hardware-software product designed to run SAN management applications such as HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager.



**storage system**

Synonymous with virtual array. The HP StorageWorks Enterprise Virtual Array consists of one or more storage systems. See also virtual array.

**synchronous**

A descriptive term for computing models that perform tasks in chronological order without interruption. In synchronous replication, the source waits for data to be copied at the destination before acknowledging that it has been written at the source. See also asynchronous.

**UUID**

Unique Universal Identifier, a unique 128-bit identifier for each component of an array. UUIDs are internal system values that users cannot modify.

**VCS**

Virtual Controller Software. The software in the HP StorageWorks Enterprise Virtual Array controller. Controller software manages all aspects of array operation, including communication with HP StorageWorks Command View EVA.

**virtual array**

Synonymous with disk array and storage system, a group of disks in one or more disk enclosures combined with control software that presents disk storage capacity as one or more virtual disks. See also virtual disks.

**Virtual Controller Software**

See VCS.

**virtual disk**

Variable disk capacity that is defined and managed by the array controller and presentable to hosts as a disk.

**Vraid**

Techniques for configuring virtual disks to provide fault tolerance and increase performance. Vraid techniques are identified by level numbers.

<u>Level</u>	<u>Redundancy</u>	<u>Technique</u>
Vraid0	None	Striping
Vraid1	High	Mirroring
Vraid5	Medium	Striping and parity

**Vraid0**

A virtualization technique that provides no data protection. Data chunks are distributed across the disk group from which the virtual disk is created. Reading and writing to a Vraid0 virtual disk is very fast and uses available storage to the fullest, but provides no data protection (redundancy) unless there is parity.

**Vraid1**

A virtualization technique that provides the highest level of data protection. All data blocks are mirrored, or written twice, on separate disks. For read requests, the block can be read from either disk, which can increase performance. Mirroring requires the most storage space because twice the storage capacity must be allocated for a given amount of data.

**Vraid5**

A virtualization technique that uses parity striping to provide moderate data protection. For a striped virtual disk, data is broken into chunks and distributed across the disk group. If the striped virtual disk has parity, another chunk (a parity chunk) is calculated from the data chunks and written to the disks. If a data chunk becomes corrupted, the data can be reconstructed from the parity chunk and the remaining data chunks.

# Index

## A

- active site [13](#)
- adding remote access IP addresses [39](#)
- alternate site [16](#)
- array messages [137](#)
- assumptions
  - EVA [16](#)
  - host operating systems [16](#)
  - hosts [16](#)
- audience [9](#)
- authorized reseller, HP [12](#)
- avoiding configuration changes while jobs are running [131](#)
- avoiding network identification changes [132](#)
- avoiding simultaneous replication events for the same virtual disk [133](#)

## B

- backing up the configuration [62](#)
- best practices [117](#)
  - avoiding configuration changes while jobs are running [131](#)
  - avoiding network identification changes [132](#)
  - avoiding simultaneous replication events for the same virtual disk [133](#)
  - backing up replication jobs and configurations [130](#)
  - bootless DR group failover (Linux) [127](#)
  - coordinating enabled host shutdowns [132](#)
  - coordinating enabled-host downtime [131](#)

- coordinating replication server shutdowns [132](#)
- creating a destination snapclone before making a full copy [117](#)
- maintaining network connections [132](#)
- making CD-ROMs of replication product [133](#)
- managing replication events [133](#)
- minimizing simultaneous replication events on an array [133](#)
- optimizing browser-based GUI performance [131](#)
- optimizing discovery performance [130](#)
- optimizing discovery refresh intervals [130](#)
- optimizing the number of active enabled hosts [131](#)
- scheduling jobs [133](#)
- snapclone rules [134](#)
- snapshot rules [134](#)
- specifying disk group membership for a log [121](#)
- support procedures [117](#)
- three-site cascaded replication [123](#)
- throttling of merge I/O after logging [130](#)
- using a snapclone to move data [119](#)
- using log files for troubleshooting [133](#)
- Windows caching [135](#)
- bidirectional replication [13](#), [23](#)
- bootless DR group failover (Linux) [127](#)
- Business Copy EVA
  - license [117](#), [119](#)

**C**

- caching, Windows 135
- capturing configuration information
  - manually 63
  - using SSSU 62
- cascaded replication 123
- codes
  - event 40
- component repair vs. failover 59
- concepts 23
  - bidirectional replication 23
  - DR group 25
  - failover 31
  - failsafe mode 34
  - Home designation 27
  - local replication 23
  - remote replication 23
  - replication direction 25
  - snapclone 24
  - snapshot 24
  - zoning 17
- configuration
  - backing up 62
  - text files 62
- configuration information
  - capturing
    - manually 63
    - using SSSU 62
- configuring standby servers 43
- considerations for managing storage 42
- conventions
  - document 11
- coordinating enabled host shutdowns 132
- coordinating enabled-host downtime 131
- coordinating replication server shutdowns 132
- creating
  - destination snapclone before full copy 117

**D**

- data movement using a snapclone 119
- data replication 23

- data replication group 25
- database
  - exporting 43
  - importing 43
  - migrating 47
- destination array 13
- destination virtual disk 23
- dialog box error messages 155
- disaster planning 57
- disk group
  - hardware failure 98
    - definition 98
    - on source array 99
    - on the destination array 103
- document conventions 11
- DR group
  - bootless failover (Linux) 127
  - creating from snapclone 120
  - depiction 26
  - description 25
  - failover 58
  - log disks 28
  - log size 28
  - logging state 28
  - messages 138
  - presentation 27
  - presenting DR group members to same FCA 27
  - properties 27
  - unknown state 107

**E**

- Element Manager for HSG 18
- Enterprise Virtual Array
  - description 16
- environment 17
- error messages 148
  - dialog box 155
- EVA 16
  - assumptions 16
- event
  - codes 40

logs 33  
messages 137  
monitoring 40  
scenarios 65  
exporting database 43

## F

failover 14  
  concept 31  
  controller 58  
  defined 58  
  DR group 58  
  fabric or path 58  
  managed set 58  
  planned procedure 68  
  planned scenario 59, 65  
  unplanned procedure 76  
  unplanned scenario 59, 66  
failover vs. component repair 59  
failsafe mode 34  
fast synchronization 28  
FATA drives 121  
FCA  
  presenting DR group members 27  
Fibre Channel adapter  
  in host 20  
full copy 28, 117

## G

general purpose server  
  starting replication manager 38  
getting help 12  
group data  
  importing 46

## H

Home designation  
  DR group 27  
host operating systems 16  
  assumptions 16

hosts  
  assumptions 16  
HP Command View EVA  
  description 18, 20  
  interface options 19  
HP Continuous Access EVA  
  configuration 14  
  failover 57  
  features 15  
  overview 13  
  prerequisites 9  
  related documentation 10  
HP OpenVMS  
  privileges 71, 76, 84, 93  
HP Replication Solutions Manager  
  interface options 19  
HP Replication Solutions Manager interface  
  monitoring 40  
HP resources  
  authorized reseller 12  
  storage website 12  
  technical support 12  
HP Storage Area Manager  
  starting replication manager 38  
HP StorageWorks SMI-S  
  interface options 20  
HP-UX time-outs and delays 109

## I

importing  
  database 43  
  group data 46  
  user accounts 46  
inoperative disk group 98  
interface options 19  
  HP Command View EVA 19  
  HP Replication Solutions Manager 19  
  HP StorageWorks SMI-S 20  
  SSSU 20  
IP address  
  changing on UNIX systems 109

**J**

- jobs [55](#)
  - messages [139](#)
  - minimizing simultaneous [131](#)
  - scheduling [133](#)

**L**

- license keys [16](#)
- licensing [16](#)
- Lifekeeper clusters
  - Red Hat and SuSE Linux [129](#)
- local replication
  - concept [23](#)
- local site [13](#)
- log disks [28](#)
  - description [28](#)
  - fast synchronization [28](#)
  - full copy [28](#)
  - group membership [121](#)
  - marked for full copy [117](#)
- log size [28](#)
- logging [99](#)
- logging states [28](#)
- logs
  - event [33](#)
  - security [33](#)
  - trace [33](#)
  - transaction [34](#)
- long delays on HP-UX [109](#)
- loss of redundancy [98](#)
- LUN
  - inaccessible to host
    - troubleshooting [107](#)
  - read-only access for destination [108](#)

**M**

- maintaining network connections [132](#)
- making CD-ROMs of replication product [133](#)
- managed set [29](#)
  - failover [58](#)

- management server
  - changing password [45](#)
- managing replication events [133](#)
- merging [117](#)
- messages
  - array [137](#)
  - dialog box errors [155](#)
  - DR group [138](#)
  - error [148](#)
  - event [137](#)
  - jobs [139](#)
- migrating database [47](#)
- minimizing simultaneous replication events on an array [133](#)
- mode
  - failsafe [34](#)
- multiple servers
  - managing storage [42](#)

**O**

- optimizing browser-based GUI performance [131](#)
- optimizing discovery performance [130](#)
- optimizing discovery refresh intervals [130](#)
- optimizing the number of active enabled hosts [131](#)
- original state [25](#)
- overview
  - HP Continuous Access EVA [13](#)

**P**

- password
  - changing
    - management server [45](#)
- path failover [58](#)
- placement of log into disk group [121](#)
- planned failover [59](#), [68](#)
  - scenario [65](#)
- planned transfer of operations [69](#)
- planning for disaster [57](#)
- prerequisites [9](#)

presentation  
  DR group [27](#)  
primary site [16](#)  
procedures  
  post replication manager installation tasks  
    adding remote access IP addresses [39](#)  
  starting replication manager  
    applet mode via browsing [38](#)

## R

read-only access  
  destination LUN [108](#)  
Red Hat Linux  
  Lifekeeper clusters [129](#)  
related documentation [10](#)  
remote access IP addresses, adding [39](#)  
remote copy [23](#)  
remote replication  
  bidirectional replication [23](#)  
  concept [23](#)  
  tasks [41](#)  
remote site [13](#)  
replication direction  
  concept [25](#)  
replication manager  
  adding remote access IP addresses [39](#)  
  browsing  
    directly [38](#)  
    SMA [38](#)  
  data replication [40](#)  
  monitoring events [40](#)  
  starting [36](#)  
    applet mode [38](#)  
    applet mode via browsing [38](#)  
    application mode [37](#)  
Replication Solutions Manager interface  
  managed sets [29](#)  
  recommended uses [36](#)  
replication tasks [35](#)  
Resume command [118](#)  
resumption of operations procedure [81](#)  
resumption of operations scenario [66](#)

return operations to Home array [83](#)  
return operations to Home array scenario [66](#)  
return operations to replaced new hardware  
  [67](#), [89](#)

## S

scheduling jobs [133](#)  
security logs [33](#)  
servers  
  multiple  
    managing storage [42](#)  
    synchronizing time [52](#)  
setting read-only access for a destination LUN  
  [108](#)  
single component failure [59](#)  
site failover, description [59](#)  
SMA  
  starting replication manager [38](#)  
snapclone  
  concept [24](#)  
  creating a DR group [120](#)  
  creating before full copy [117](#)  
  data movement [119](#)  
  rules [134](#)  
  three-site cascaded replication [123](#)  
snapshot  
  concept [24](#)  
  rules [134](#)  
source array [13](#)  
source virtual disk [23](#)  
specifying disk group membership for a log [121](#)  
SSSU [91](#)  
  interface options [20](#)  
stale data [99](#)  
standby server  
  configuring [43](#)  
standby site [13](#)  
starting replication manager  
  applet mode via browsing [38](#)  
starting the replication manager [36](#)

- storage management
  - moving storage management to another server [50](#)
- storage management considerations [42](#)
- storage problems
  - troubleshooting [110](#)
- Storage System Scripting Utility [91](#)
- support procedures [117](#)
- supported operating systems [16](#)
- SuSE Linux
  - Lifekeeper clusters [129](#)
- Suspend command [118](#)
- synchronizing time on servers [52](#)

## T

- technical support, HP [12](#)
- three-site cascaded replication [123](#)
- throttling I/O [130](#)
- throttling of merge I/O after logging [130](#)
- time-outs on HP-UX [109](#)
- trace logs [33](#)
- transaction logs [34](#)
- troubleshooting
  - DR groups in unknown state [107](#)
  - IP address on UNIX [109](#)
  - long delays or time-outs on HP-UX [109](#)
  - LUN detection [108](#)
  - LUN inaccessible to host [107](#)
  - storage problems [110](#)

- tunnel thrash [108](#)
- tunnel thrash
  - troubleshooting [108](#)

## U

- UNIX systems
  - changing IP address of local host [109](#)
- unknown state
  - DR groups [107](#)
- unplanned failover [59, 66](#)
  - procedure [76](#)
- unplanned transfer of operations [69](#)
- user accounts
  - importing [46](#)
- using log files for troubleshooting [133](#)

## V

- VCS [16](#)
- Virtual Controller Software
  - description [16](#)

## W

- websites
  - HP storage [12](#)
  - technical support [12](#)
- when or when not to failover [60](#)
- Windows caching [135](#)