



EMC® VNX™ Series

Release 7.1

Using RecoverPoint/SE with VNX™ for File for Disaster Recovery

P/N 300-013-440 Rev 01

EMC Corporation

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

Copyright © 2011 - 2012 EMC Corporation. All rights reserved.

Published July 2012

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

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

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

For the most up-to-date regulatory document for your product line, go to the Technical Documentation and Advisories section on EMC Powerlink.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

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

Corporate Headquarters: Hopkinton, MA 01748-9103

Preface	7
Chapter 1: Introduction	9
System requirements.....	10
Restrictions.....	11
User interface choices.....	19
Related information.....	20
Chapter 2: Concepts	23
RecoverPoint product family.....	24
RecoverPoint.....	24
RecoverPoint/SE.....	24
RecoverPoint replication configurations.....	24
Communication between VNX systems.....	26
VNX primary and standby system compatibility.....	27
RecoverPoint configurations.....	27
RecoverPoint/SE active/passive.....	28
RecoverPoint/SE active/active'.....	28
RecoverPoint replication modes.....	29
RecoverPoint hardware and software.....	31
RecoverPoint Appliance.....	32
Splitters.....	34
RecoverPoint interfaces.....	35
RecoverPoint logical entities.....	35
Consistency groups.....	35
Copies.....	39
Replication sets.....	40

Replication policies.....	40
Journals.....	40
Volumes.....	41
Snapshots.....	43
Planning considerations.....	45
File system recommendations.....	45
Remote administration account recommendations.....	45
VNX for Block configuration.....	46
VNX volume and Data Mover decisions/flexibility.....	47
Data Mover configuration checklist.....	49
RecoverPoint/SE configuration sheet.....	50
Task overview.....	52
Sample configuration.....	55
Chapter 3: Configuring VNX backends.....	57
Implement RecoverPoint CRR by using Deployment Manager.....	58
Configure connection between RPAs and SPs	58
Configure the source VNX.....	59
Identify VNX for File LUNs.....	59
Verify LUNs in the RPA storage group at the source site.....	61
Identify source servers.....	61
Configure the destination VNX.....	61
Create LUNs to match the LUNs on the source VNX	62
Add LUNs to the VNX for File storage group at the destination site.....	64
Verify LUNs in the VNX for File storage group at the destination site.....	64
Verify LUNs in the RPA storage group at the destination site.....	65
Identify destination servers.....	65
Configure the consistency group with the production and remote site.....	66
Chapter 4: Configuring RecoverPoint (active/passive).....	69
Preinitialize the configuration.....	70
Preinitialize from the first (source) system.....	70
Preinitialize from the second (destination) system.....	71
Verify the preinitialization.....	73
Initialize the configuration (active/passive).....	73
Initialize from the destination system (active/passive).....	75

Verify configuration (active/passive).....	81
Failover the source system (active/passive).....	88
Activate a failover from the destination VNX (active/passive).....	89
Verify active/passive operations after failover.....	94
Ensure access after failover.....	99
Failback the source system (active/passive).....	100
Prepare for the failback (active/passive).....	100
Failback the source system from the destination system(active/passive).....	101
Verify operations after a failback (active/passive).....	103
Chapter 5: Configuring RecoverPoint (active/active').....	107
Initialize the configuration (active/active').....	108
Verify remote administration account (Optional).....	110
Initialize one system (active/active').....	113
Initialize the second system (active/active').....	117
Verify configuration (active/active').....	121
Failover the source system (active/active').....	128
Failover from the destination system.....	128
Verify active/active' operations after failover.....	132
Failback the source system (active/active').....	134
Prepare for the failback (active/active').....	135
Failback the source system from its destination (active/active').....	135
Chapter 6: Managing.....	139
Manage RPA information.....	140
List RPA information.....	140
Add RPA entry to the NAS database.....	140
Get detailed RPA information.....	141
Update RPA.....	141
Check cabinet-level RecoverPoint/SE information.....	142
List consistency group information.....	145
Get detailed consistency group information.....	146
Display information for each consistency group.....	147
Suspend consistency group operation.....	149
Resume consistency group operations	150
Ensure Data Mover eligibility.....	151
Ensure Data Mover network device compatibility.....	153
Get additional information during initialization.....	154

Verify configuration during initialization.....	155
Change the Data Mover configuration.....	159
Modify VNX for Block security information after a failover.....	162
Change file system configuration.....	163
Change a local file system to a replicated file system.....	164
Change a replicated file system to a local file system.....	165
Chapter 7: Troubleshooting.....	167
EMC E-Lab Interoperability Navigator.....	168
Known problems and limitations.....	168
Retrieve information from log files.....	170
Resolve initialization failures.....	172
Resolve failover failures.....	174
Resolve failback failures.....	177
Additional situations involving Data Movers.....	180
Error messages.....	181
EMC Training and Professional Services.....	182
Glossary.....	183
Index.....	185

Preface

EMC RecoverPoint is a data replicating facility, which provides remote disaster recovery with point-in-time recovery that utilizes continuous remote replication technology. It provides NAS file system and VDM remote replication to a secondary site by implementing Data Mover level replication and cabinet-level failover. RecoverPoint maintains a synchronized remote replica of production data between source and destination VNX for File/Block pairs at separate locations.

[Chapter 2](#) provides more details.

This document describes how to use RecoverPoint with VNX systems for a cabinet-level (file system data only) disaster recovery.

This document is intended for those who are responsible for configuring RecoverPoint with VNX for File. The document also identifies administrative tasks, such as:

- ◆ Initial backend setup performed by your local EMC Service Provider.
- ◆ Initial RecoverPoint setup and verification, including preinitialization, initialization, test failover, and test failback performed by your local EMC Service Provider.

Note: The source VNX shuts down even during test failover and test failback.

- ◆ Troubleshooting procedures, including resolution or potentially complex change to the VNX performed by your local EMC Service Provider or EMC Customer Service.

Topics include:

- ◆ [System requirements on page 10](#)
- ◆ [Restrictions on page 11](#)
- ◆ [User interface choices on page 19](#)
- ◆ [Related information on page 20](#)

System requirements

Table 1 on page 10 describes the EMC® VNX™ Series software, hardware, network, and storage configurations.

Table 1. System requirements

Software	<ul style="list-style-type: none"> ◆ VNX operating environment (OE) on the source and destination VNX for Block systems, which must also have EMC Unisphere™, RecoverPoint/SE, and EMC Access Logix™ installed and enabled. VNX OE on the source and destination, supports the consistency groups for VNX for File and VNX for Block systems. VNX with RecoverPoint/SE replicates at the VNX for File cabinet level. This includes system volumes, and user file system volumes. ◆ The VNX for Block remote replication sets and the consistency group must be set up by your local EMC Service Provider, as SAN administrator, through EMC Unisphere software, NaviCLI, or RecoverPoint CLI. ◆ The same VNX major version on both the source and destination Control Stations.
Hardware	<ul style="list-style-type: none"> ◆ RecoverPoint appliance (RPA), which manages all aspects of reliable data replication at all sites. ◆ Array-based write splitters installed on the VNX system. ◆ Two VNX systems, one at each site. VNX systems come in various configurations (VNX5300, VNX5500, VNX5700, and VNX7500), but there is no restrictions on the configuration of source and destination site VNX systems for the disaster recovery setup. ◆ Two physical FC links from each VNX storage processor zoned to Recoverpoint appliances on corresponding sites. There are no restrictions on the ports to be used. <p style="margin-left: 20px;">Note: Port 0 is reserved for MirrorView™ /S connections.</p> <ul style="list-style-type: none"> ◆ For EMC VNX VG2 and EMC VNX VG8 setup you can use VNX for Block as the backend. Only LUNs from boot (storage system that provisions the VNX for File control LUNs) VNX for Block storage are supported for the VNX for File cabinet DR setup. ◆ If you have multiple VNX for Block systems, only the one consistency group on the boot VNX for Block system is available for failover to the remote site. <p>The EMC E-Lab™ Interoperability Navigator provides the full set of VNX and gateway configuration guidelines.</p>
Network	<ul style="list-style-type: none"> ◆ IP network for the Control Station of the source VNX and the Control Station of the destination VNX. ◆ LAN or WAN links for communication between the Control Stations. ◆ Dedicated FC links for connecting the storage arrays. ◆ Both VNX systems must be in the same management domain and therefore must be able to communicate over the IP network.

Restrictions

RecoverPoint/SE supports VNX for File by offering protection for a full-NAS Cabinet. While VNX Replicator should typically be used for replicating file systems, the RecoverPoint/SE solution is suitable for use cases where synchronous replication is required, and where you would have otherwise used file system replication with MirrorView/S. These restrictions must be taken into consideration and accepted when designing such a solution, and before committing to it.

- ◆ In a RecoverPoint/SE environment, you cannot cascade VNX for File platforms. A RecoverPoint/SE configuration consists of a VNX for File with an attached VNX for Block at one site and another VNX for File with an attached VNX for Block at the other site. Dedicated RecoverPoint/SE FC/IP links are established between the storage processors (SPs) of the source and destination Block storage arrays (SPA-SPA and SPB-SPB). If the source SPB fails, then replication continues between the source and destination SPA.
- ◆ A RecoverPoint/SE active/passive configuration can have only one NAS consistency group per VNX for File platform. The total consistency groups on an RPA is limited to 128, which includes both NAS and SAN consistency groups. A RecoverPoint/SE active/active' configuration can have only one NAS consistency group for each direction.
- ◆ RecoverPoint/SE is geared for new installations. The consistency group is limited to 2048 replication sets.
- ◆ Using RecoverPoint/SE with VNX requires three control LUNs in the NAS consistency group. The remaining 2048 LUNs can be data LUNs.
- ◆ The maximum storage capacity of a RecoverPoint/SE consistency group is 600 TB.
- ◆ Disaster Recovery using RecoverPoint/SE is different from that using EMC Symmetrix[®] Remote Data Facility (EMC SRDF[®]). With RecoverPoint/SE, VNX for File does not see the replicated LUNs at the destination site until a failover is activated. However, with SRDF the remote replication sets (R2 volumes) of the Control Station LUNs are visible to the Data Movers. With RecoverPoint/SE, there is no access to the destination LUNs until a failover is activated. Therefore, the destination side of a RecoverPoint/SE CRR configuration cannot be used for tape backup or loadsharing prior to a failover activation.

Note: RecoverPoint/SE configurations are not suitable for backups from the destination.

- ◆ During the failback operation, the storage restoration phase of the process waits until both sides are 100-percent synchronized to ensure that the source system gets the latest data. The length of the synchronization is based on the amount of data that has been updated since the destination system experienced a failover. After that, the destination VNX for File fails back the Data Movers, and the failback process stops client access, fails back the storage system, and then restores the source VNX for File. For RecoverPoint/SE configuration, the local file systems and RecoverPoint-protected file systems have their own, dedicated Data Movers. To avoid failover problems, Data Movers that are RecoverPoint-protected must contain only file systems that are intended to be protected by RecoverPoint/SE, not local file systems. Also, make sure that the file systems do not span multiple storage systems.

- ◆ While configuring Primary and DR site VNX/Gateway systems for the disaster recovery setup:
 - The SLIC IO modules on both Primary and DR site Data Movers must be identical in type and quantity.
 - The order of SLIC IO modules on both Primary and DR site Data Movers must be identical.
- ◆ All Data Movers must use their default names, for example, server_2, server_3, and so on, when you configure them for RecoverPoint/SE. If you change the default Data Mover names, it may cause RecoverPoint/SE activation to stop responding when you add the NBS devices to the Operating Environment.
- ◆ Performance Constraints - The current implementation imposes significant performance restrictions for synchronous replication of file systems:
 - All file systems in the pool (cabinet) will be replicated and this will “waste” RecoverPoint resources by replicating less important information. To avoid this, you can use multiple pools in the cabinet. But this makes the solution a lot more complex.
 - Because the whole NAS cabinet must reside in a single RecoverPoint Consistency Group, you cannot balance the load of the NAS across several Consistency Groups, and thus cannot benefit from the RP scale out architecture utilizing up to 8 RPAs.
 - Using synchronous replication further amplifies these restrictions because the IOPS and throughput limits for synchronous replication are lower compared to asynchronous replication, and they are greatly affected by the (FC) WAN latency.
 - Considering all these restrictions, typical performance numbers that can be achieved by using this solution are listed in the following table. This table assumes synchronous replication, a single pool in the cabinet, 8 KB write size, and roundtrip latency of 0.5ms (about 50km distance). Another assumption is that you want to protect only 30 percent of the writes, but have to replicate all of them:

Number of RPAs per side	2	4	8
Average sustained IOPS for the whole cabinet	6,400	6,400	6,400
Average sustained IOPS “available” for the applications that need to be protected (30%)	1,920	1,920	1,920
RecoverPoint cluster throughput (MB/s)	50	50	50

For more details, please see performance related information in the RecoverPoint and RecoverPoint/SE Release Notes.

- ◆ Some of the RecoverPoint features and differentiators are not available with the current NAS Cabinet replication solution. Specifically, the following key features are not supported:

- DR testing – You cannot access the replica image for DR testing (a key operation for disaster recovery readiness). Also, you cannot back up the replica without performing a failover to the DR site.
- "Any point-in-time" recovery - The solution does not support any point-in-time recovery. You can only fail over to the latest point in time.
- GUI management for RecoverPoint - The solution does not support GUI management for RecoverPoint, including Unisphere integration on the production frame after failover. You can manage RecoverPoint only through CLI.
- Consistency Group replication - You can replicate only one Consistency Group in one direction.

Initial setup

- ◆ RecoverPoint/SE requires a license.
- ◆ RecoverPoint/SE supports VNX for File, VNX for Block, and gateway configurations.
- ◆ With RecoverPoint/SE, VNX for File does not support Continuous Data Protection (CDP) and Concurrent Local and Remote data protection (CLR) replication configurations.
- ◆ RecoverPoint/SE does not support CNS and CFS storage platforms.
- ◆ For gateway configurations, only boot (control LUNs) back-end VNX for Block configurations are allowed.
- ◆ The VNX for Block source LUNs and equivalent destination LUNs must have the same size.

Note: Create LUNs to match the LUNs on the source VNX.

Note: [System requirements on page 10](#) identifies the RecoverPoint/SE with VNX basic hardware and software requirements. The RecoverPoint family documentation available on the [EMC Online Support](#) website provides information about the RecoverPoint products.

Disk types and storage pools

- ◆ RecoverPoint/SE integration with VNX requires:
 - CMSTD — RAID group FC MirrorView/RecoverPoint/SE mirror
 - CMATA — RAID group SATA MirrorView/RecoverPoint/SE mirror
 - CMSAS — RAID group SAS MirrorView/RecoverPoint/SE mirror
 - CMEFD — RAID group EFD MirrorView/RecoverPoint/SE mirror

- MMIXD — Unified mixed MirrorView/RecoverPoint/SE mirror
 - CMNLS — RAID group NL-SAS MirrorView/RecoverPoint/SE mirror
 - MPERF — Unified FC MirrorView/RecoverPoint/SE mirror - Performance
 - MCAPA — Unified SATA MirrorView/RecoverPoint/SE mirror - Capacity
 - MPERF — Unified SAS MirrorView/RecoverPoint/SE mirror - Performance
 - MCAPA — Unified NL-SAS MirrorView/RecoverPoint/SE mirror - Capacity
 - MFLSH — Unified EFD MirrorView/RecoverPoint/SE mirror
- ◆ The integration of RecoverPoint/SE with VNX uses the following system-defined Automatic Volume Management (AVM) storage pools:
- cm_r1 — Designed for high performance and availability at low cost. This storage pool uses VNX CMSTD disk volumes created from RAID 1 replicated-pair disk groups.
 - cm_r5_performance — Designed for medium performance and availability at low cost. This storage pool uses VNX CMSTD disk volumes created from 4+1 RAID 5 disk groups.
 - cm_r5_economy — Designed for medium performance and availability at lowest cost. This storage pool uses VNX CMSTD disk volumes created from 8+1 RAID 5 disk groups.
 - cmata_archive — Designed for use with infrequently accessed data, such as archive retrieval. This storage pool uses VNX ATA CMATA disk drives in a RAID 5 configuration.
 - cmata_r3 — Designed for archival performance and availability at lowest cost. This AVM storage pool uses VNX ATA CMATA disk drives in a RAID 3 configuration.
 - cm_r6 — Designed for high performance and availability at low cost. This storage pool uses VNX CMSTD disk volumes created from RAID 6 replicated-pair disk groups.
 - cmata_r6 — Designed for archival performance and availability at lowest cost. This AVM storage pool uses VNX ATA CMATA disk drives in a RAID 6 configuration.
 - cmata_r10 — Designed for archival performance and availability at lowest cost. This AVM storage pool uses VNX ATA CMATA disk drives in a RAID 10 configuration.
 - cmsas_r10 — Designed for archival performance and availability at lowest cost. This AVM storage pool uses VNX CLSAS disk drives in a RAID 10 configuration.
 - cmsas_archive — Designed for use with infrequently accessed data, such as archive retrieval. This storage pool uses VNX CLSAS disk drives in a RAID 5 configuration.
 - cmsas_r6 — Designed for archival performance and availability at lowest cost. This AVM storage pool uses VNX CLSAS disk drives in a RAID 6 configuration.

- `cm_r10` — Designed for high performance and availability at low cost. This storage pool uses VNX CMSTD disk volumes created from RAID 10 replicated-pair disk groups.
 - `cmefd_r5` — Designed for archival performance and availability at lowest cost. This AVM storage pool uses VNX CLEFD disk drives in a RAID 5 configuration.
 - `cmefd_r10` — Designed for archival performance and availability at lowest cost. This AVM storage pool uses VNX CLEFD disk drives in a RAID 10 configuration.
- ◆ You can view the RecoverPoint/SE disk types and view or manipulate AVM storage pools by using the Unisphere software. Also, note that replicated devices are converted automatically from CLSTD to CMSTD during the RecoverPoint/SE initialization process.

Note: Disks in a storage pool must be all replicated or all unreplicated. The `/nas/sbin/nas_rp-cabinetdr -init` command fails if a mix of disk types is created by changing the storage system configuration.

Note: *Managing VNX Volumes and File Systems with AVM* describes how to manage volumes and file systems. *Managing VNX Volumes and File Systems Manually* describes how to manage volumes and file systems.

RecoverPoint/SE volume restrictions

- ◆ There must be at least one journal volume on each site per consistency group. These journal volumes must be visible only to the RPA.
- ◆ The minimum supported journal volume size is 5 GB.
- ◆ LUNs used for a journal volume should not be of the Thin type.
- ◆ The repository volume must be visible only to the RPA that are on the same site as the repository volume.
- ◆ The maximum supported repository volume size is 3 GB.

Management interface and Control Station restrictions

- ◆ After the initial setup, which involves steps performed by your local EMC Service Provider to establish the VNX storage system configuration, you should manage RecoverPoint/SE disaster recovery only through the VNX for File CLI by using commands, such as `/nas/sbin/nas_rp` in the VNX environment from the primary Control Station, Control Station in slot 0 (CS0). Unisphere and the Navisphere CLI support a full suite of RecoverPoint functions, but VNX for File users should limit the use of these VNX for Block management interfaces to monitoring operations. VNX for File users should not routinely manage RecoverPoint/SE by using the VNX for Block management interfaces because it might lead to data inconsistency or the

inability to boot the VNX for File system. After RecoverPoint/SE initialization, changes affecting the configuration of the VNX for Block require the use of Unisphere or Navisphere CLI commands. For example, to add or delete replicated LUNs. These changes are typically made by your local EMC Service Provider.

- ◆ After a failover is activated at the destination VNX, you cannot use Unisphere to manage the components of the original source site configuration from the destination. Use the CLI to manage the original source site configuration from the destination.
- ◆ For sites with redundant Control Stations, run all RecoverPoint/SE management commands from CS0, never Control Station in slot 1 (CS1). For example, `/nas/sbin/nas_rp -cabinetdr -init, -failover, and -failback`, and `nas_rp -cg`. When running `/nas/sbin/nas_rp -cabinetdr -init, -failover` or `-failback` commands, make sure that CS1 is halted. However, for steady state non-RecoverPoint operations, CS1 can be operational.
- ◆ Initialization, failover activation, and initial failback operations are all performed from the designated destination VNX in the configuration.
- ◆ All procedures involving root user (su) privileges should be followed very carefully. Do not use su unless explicitly mentioned. See the man page for su for more information and usage of su.
- ◆ The distance between source and destination Control Station systems is limited to 200 km.
- ◆ Do not activate a failover unless you have a valid RecoverPoint/SE configuration in place. Ensure that the source Data Mover is configured with a remote standby Data Mover at the destination.
- ◆ Although RecoverPoint/SE administrators can access and modify any consistency group including NAS consistency groups, they must ensure that they synchronize all operations on NAS consistency groups with the NAS administrators.
- ◆ During the NAS consistency group failover, the services on the source Control Station are shut down. Consequently, you cannot manage VNX for block through the Unisphere software. However, you can manage VNX for block through NaviCLI.

Equivalent destination LUN restrictions

Observe the following restrictions when selecting the RAID group or HLU IDs for an equivalent destination LUN:

- ◆ A source LUN and its equivalent destination LUN must be the same size as the source. If you use the Unisphere software, the destination LUN created is of the same size as its source LUN.
- ◆ For control LUNs: Ensure that the equivalent destination control LUNs (with HLU IDs 6, 7, and 9 to represent the secondary images of the source control LUNs 0, 1, and 4), reside in a RAID group other than the one used for the destination system's own control LUNs. For example, if the destination system's own control LUNs are already using RAID group 0, you can use a RAID group such as 3 for the equivalent destination control LUNs.

- ◆ For user data LUNs: There are no specific restrictions for RAID group IDs, but use an HLU ID of 16 or higher (0–15 are reserved) based on what is available. In a new install, you might want to make the ALU match the HLU ID for simplicity.

Account and password restrictions

- ◆ The global account information, such as username and password, that you supply during initialization must match the global account information established for VNX. The account must be established for the VNX for Block administrator role. If you rerun initialization and the global account information is already known, you are not prompted for it again. If you change this account information, you must follow the guidelines summarized in [Guidelines for changing your VNX configuration on page 47](#).
- ◆ The remote administration account you establish can have the same information for both the VNX systems in an active/active' configuration. Again, during a rerun of the initialization process, you are not prompted for the information, but a message indicates that the established remote administration account is being used.
- ◆ You must log in to the remote administration account, for example, dradmin, to activate a failover or perform a failback of the source VNX.
- ◆ Output differs for informational commands, such as `nas_server -info -all` depending on how you run the command as `nasadmin`, `root`, or the established remote administration account. For example, after initialization of remote standby Data Movers, running `nas_server -info -all` as `nasadmin` does not display the Data Movers owned by the remote administration account, instead you view them when you run the command as `root`. Also, the information in the output's `acl` field for `owner=` is the remote administration account when run as `root`. For example, `owner=dradmin`.
- ◆ To avoid potential configuration errors, use `root` access only when required, not for routine administration tasks, such as creating a file system, creating a checkpoint, or performing a manual local standby failover.

Failover restrictions

- ◆ RecoverPoint/SE does not support partial failovers. When a failover occurs, all file systems associated with RecoverPoint-protected Data Movers fail over. To avoid failover problems, it is critical that replicated file systems reside only on RecoverPoint-protected Data Movers. [Chapter 7](#) provides description of the potential failover and failback problems for RecoverPoint.
- ◆ Only one system in an active/active' configuration can have a failover activated. You cannot have failovers activated on both systems at the same time.
- ◆ When replicating NAS with RecoverPoint/SE, this feature will only support failover to the latest image. You will not be able to select a specific point in time.

RecoverPoint/SE with VNX Usermapper

- ◆ If you want to continue the use of the Usermapper service after a RecoverPoint/SE failover activation, make sure that the Internal Usermapper service is running on the source RecoverPoint-protected Data Mover.
- ◆ With External Usermapper, access to the Usermapper service is lost after a RecoverPoint/SE failover.

RecoverPoint/SE with SnapSure checkpoints

- ◆ EMC SnapSure™ SavVol cannot be created on local storage if the Production File System (PFS) is mounted on a Data Mover configured with a remote standby. If you plan to create checkpoints of a PFS that resides on a RecoverPoint/SE LUN, ensure that the entire SnapSure SavVol with the checkpoints resides in the same pool of RecoverPoint/SE LUNs used to create the PFS. If any part of the SavVol is stored on a local volume rather than completely on the pool of RecoverPoint/SE LUNs, the checkpoints are not failed over and therefore are not recoverable in the event of a failover. Evaluate the use of checkpoints carefully.
- ◆ After a failover is activated, checkpoint scheduling is not supported until a failback is performed.
- ◆ Checkpoint autoextend of the SavVol is not supported. If the SavVol fills to capacity, writes to the PFS continue and the oldest checkpoint gets deactivated.

SNMP or e-mail event notification

After a failover is activated, SNMP or e-mail for event notification is not supported.

Other VNX feature-specific restrictions

- ◆ EMC VNX Replicator™ works with disaster recovery replication products such as SRDF/Synchronous (SRDF/S) and SRDF/Asynchronous (SRDF/A), MirrorView/Synchronous (MirrorView/S) or RecoverPoint/SE. You can run SRDF, MirrorView/S or RecoverPoint products and VNX Replicator on the same data. However, if there is an SRDF, MirrorView/S, or RecoverPoint/SE site failover, you cannot manage VNX Replicator sessions on the SRDF, MirrorView/S, or RecoverPoint/SE failover site. Existing VNX Replicator sessions will continue to run on the failed over Data Mover and data will still be replicated. On the primary site, you can continue to manage your SRDF, MirrorView/S, or RecoverPoint/SE replication sessions after the failback.
- ◆ Automatic File System Extension is not supported from the disaster recovery site after failover, but manual file system extension is possible. After failback, normal Automatic File System Extension operations can be performed.
- ◆ RecoverPoint/SE does not support Multi-Path File System (MPFS).

- ◆ In a RecoverPoint/SE environment involving iSCSI, features that are not supported by RecoverPoint/SE are also not supported by iSCSI. For example, Automatic File System Extension and EMC TimeFinder[®]/FS, NearCopy, and FarCopy.
- ◆ Products that require Symmetrix storage, such as TimeFinder/FS, NearCopy, and FarCopy, do not work with RecoverPoint/SE.

RecoverPoint/SE error and informational messages help prevent configuration of unsupported features, such as Automatic File System Extension. [Error messages on page 181](#) provides more information.

[EMC E-Lab Interoperability Navigator on page 168](#) provides information about product interoperability. For example, products such as EMC SnapView[™] for Block do not work in the VNX for File environment. Your local EMC sales organization can provide information about using products with RecoverPoint/SE.

Note: The *VNX Release Notes*, available at <http://Support.EMC.com>, contain the latest information about changes to documented restrictions.

Note: The *RecoverPoint and RecoverPoint/SE Release Notes*, available on the [EMC Online Support](#) website provides detailed information on other RecoverPoint/SE specific restrictions.

User interface choices

This document describes how to configure RecoverPoint/SE and integrate it with VNX by using the command line interface (CLI). You cannot use other VNX management applications to configure RecoverPoint.

You can use the Unisphere software to view the storage pools and disk types used in the RecoverPoint/SE CRR configuration. For example, a disk associated with RecoverPoint/SE appears with the disk type CMSTD, CMATA, CMSAS, CMEFD, MMIXD, MPERF, MCAPA, CMNLS when viewing disk or volume properties after RecoverPoint/SE initialization. The AVM storage pools for RecoverPoint/SE are cm_r1, cm_r5_performance, cm_r5_economy, cmata_archive, cmata_r3, cm_r6, cmata_r6, cmata_r10, cmsas_r10, cmsas_archive, cmsas_r6, cm_r10, cmefd_r5, or cmefd_r10, depending on the disk configuration. While you cannot use Unisphere to configure RecoverPoint/SE, you can use it to manage storage objects, such as file systems that reside on the source RecoverPoint-protected Data Mover.

Note: When creating new storage objects on this RecoverPoint-protected Data Mover, you are restricted to space in the RecoverPoint/SE storage pools. Objects that require local storage must be created on local Data Movers.

In addition, when you use Unisphere to view information:

- ◆ You can view an alert that addresses RecoverPoint/SE consistency group conditions. [Retrieve information from log files on page 170](#) provides more information about the events that correspond to the Unisphere alerts.

- ◆ When you view Data Mover properties, note that the Standby For Movers field identifies the primary Data Movers for which the current Data Mover is a local standby and the Standby Movers field identifies the local standby Data Movers for the current primary Data Mover. The Standby Movers field is blank if no local standby Data Movers are configured.

Related information

Specific information related to the features and functionality described in this document are included in:

- ◆ *RecoverPoint Deployment Manager Product Guide*
- ◆ *EMC RecoverPoint Administrator's Guide*
- ◆ *EMC RecoverPoint Installation Guide*
- ◆ *EMC RecoverPoint CLI Reference Guide*
- ◆ *EMC RecoverPoint Security Configuration Guide*
- ◆ *EMC RecoverPoint Deploying RecoverPoint with SANTap and SAN-OS Technical Notes*
- ◆ *EMC RecoverPoint Deploying RecoverPoint with SANTap and NX-OS Technical Notes*
- ◆ *EMC RecoverPoint Deploying RecoverPoint with Connectrix AP-7600B and PB-48K-AP4-18 Technical Notes*
- ◆ *RecoverPoint and RecoverPoint/SE Release Notes*
- ◆ *EMC VNX Command Line Interface Reference for File*
- ◆ *Celerra Network Server Error Messages Guide*
- ◆ *Problem Resolution Roadmap for VNX*
- ◆ Online VNX man pages
- ◆ *Using SRDF/S with VNX for Disaster Recovery*
- ◆ *Using MirrorView/Synchronous with VNX for File for Disaster Recovery*
- ◆ *Using SRDF/A with VNX*
- ◆ *Using VNX Replicator*

EMC VNX documentation on the EMC Online Support website

The complete set of EMC VNX series customer publications is available on the EMC Online Support website. To search for technical documentation, go to <http://Support.EMC.com>. After logging in to the website, click the VNX Support by Product page to locate information for the specific feature required.

VNX wizards

Unisphere software provides wizards for performing setup and configuration tasks. The Unisphere online help provides more details on the wizards.

Topics include:

- ◆ RecoverPoint product family on page 24
- ◆ RecoverPoint replication configurations on page 24
- ◆ Communication between VNX systems on page 26
- ◆ VNX primary and standby system compatibility on page 27
- ◆ RecoverPoint configurations on page 27
- ◆ RecoverPoint replication modes on page 29
- ◆ RecoverPoint hardware and software on page 31
- ◆ RecoverPoint logical entities on page 35
- ◆ Planning considerations on page 45
- ◆ Task overview on page 52

RecoverPoint product family

EMC RecoverPoint is a high-performance, cost-effective solution for local and remote data protection, replication, and disaster recovery. The RecoverPoint family has two core products:

- ♦ [RecoverPoint on page 24](#)
- ♦ [RecoverPoint/SE on page 24](#)

RecoverPoint

RecoverPoint brings you continuous data protection and continuous remote replication for on-demand protection and recovery at any point in time. Advanced capabilities include policy-based management, application integration, and WAN acceleration.

With RecoverPoint, you implement a single, unified solution to protect and replicate data across heterogeneous storage. You simplify management and reduce costs, recover data at a local or remote site at any point in time, and ensure continuous replication to a remote site without impacting performance.

RecoverPoint/SE

RecoverPoint/SE brings continuous data protection and continuous remote replication to your VNX storage. RecoverPoint/SE gives you on-demand protection and recovery at any point in time and advanced capabilities such as policy-based management and bandwidth optimization.

With RecoverPoint/SE, you implement a single unified solution for data protection, simplify management, reduce costs, and avoid data loss due to server failures or data corruption.

RecoverPoint replication configurations

RecoverPoint supports three replication configurations:

- ♦ [Continuous data protection \(CDP\) on page 25](#)
- ♦ [Continuous remote replication \(CRR\) on page 25](#)
- ♦ [Concurrent local and remote data protection \(CLR\) on page 25](#)

RecoverPoint Continuous Data Protection (CDP) and Continuous Remote Replication (CRR) provide bi-directional replication and a point-in-time recovery mechanism. These features enable replica volumes to be rolled back to a previous point-in-time and used for read/write operations without affecting ongoing replication or data protection.

Continuous data protection (CDP)

RecoverPoint CDP can instantly recover data to any point-in-time by leveraging bookmarks from the replica journal.

In a CDP configuration, RecoverPoint replicates data within the same site or to a local bunker site some distance away over any distance. Data can be replicated locally at a distance that does not exceed the limitation specified in the *Release Notes for EMC RecoverPoint Release and Service Pack Releases*, and the data is transferred by Fibre Channel. Writes from the splitter to the RPA are written synchronously, and snapshot granularity is set to *per second*, so the exact data size and contents are determined by the number of writes made by the host application per second.

If necessary, the snapshot granularity can be set to *per write*. The replication mode can also be set to synchronous, when an RPO time of zero is required.

CDP configurations include:

- ◆ *(Standard) CDP*, in which all components (splitters, storage, RPAs, and hosts) are located at the same site.
- ◆ *Stretch CDP*, in which the production host is located at the local site, splitters and storage are located at both the bunker site and the local site, and the RPAs are located at the bunker site. The repository volume and both the production and local journals are located at the bunker site.

Continuous remote replication (CRR)

In CRR configurations, data is transferred between a local and a remote site over Fibre Channel or a WAN. The RPAs, storage, and splitters are located at both the local and the remote site.

By default, the replication mode is set to asynchronous, and the snapshot granularity is set to **dynamic**, so the exact data size and contents are determined by the policies set by the user and system performance. This provides protection to application consistent and other specific points in time.

Note: Synchronous replication is only supported when the local and remote sites are connected using Fibre Channel. The *EMC RecoverPoint Installation Guide* provides information on limitations.

Concurrent local and remote data protection (CLR)

In a CLR configuration, RecoverPoint replicates data to both a local and a remote site simultaneously.

Communication between VNX systems

In a RecoverPoint/SE CRR replication configuration, the source and destination site VNX systems communicate through:

- ◆ IP network
- ◆ RecoverPoint/SE CRR replication link
- ◆ Fibre Channel or IP connection

Figure 1 on page 26 shows a RecoverPoint/SE CRR configuration with two sites, the source (primary) site and the destination (disaster recovery) site. Each site has an attached VNX for File and VNX for Block pair. Dedicated RecoverPoint/SE CRR FC/IP links exist between the storage processors of the source and destination VNX for Block systems. Each VNX for File system has a path to each storage processor through the switch fabric. All the components of the configuration have IP network connectivity.

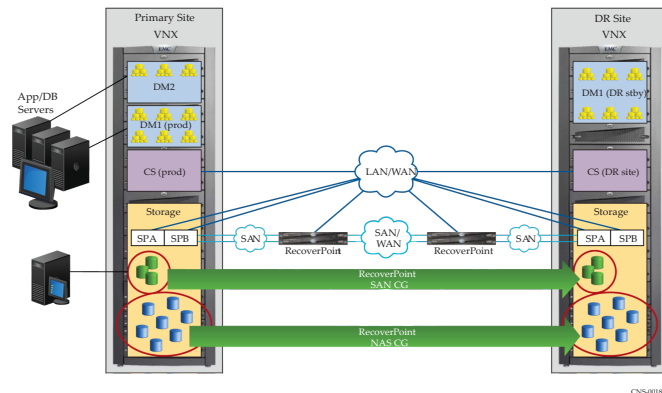


Figure 1. Sample RecoverPoint/SE with VNX configuration

Each storage processor on a VNX for Block system has a unique communication link. Each link has the possibility of failing independently. RecoverPoint/SE uses the dedicated links to send data to the mirrored LUNs and pass state information between the storage arrays. The storage processor (SPA-SPA and SPB-SPB) connections usually use the highest-numbered port.

The two Control Stations and both storage arrays must be connected through IP. The Control Stations can communicate over LAN/WAN links. Each Control Station must also use the IP network to communicate with both storage processors on the Site A and Site B VNX for Block systems. The same Control Station (CS0), which can promote the destination-site image should the source image become unavailable, must be able to manage the VNX for File/Block pairs.

VNX primary and standby system compatibility

All gateway and unified systems that support VNX version 7.0 and later are supported by RecoverPoint/SE. It is possible to replicate VNX for File LUNs with RecoverPoint/SE from different VNX systems and gateway models as per the failover rules outlined in [Table 2 on page 27](#).

Table 2. VNX primary and standby system compatibility matrix

Primary/Standby	VNX5300	VNX5500	VNX5700	VNX7500	Gateway VG2	Gateway VG8
VNX5300	Y	Y ¹	Y ¹	Y ¹	Y ¹	Y ¹
VNX5500	Y	Y	Y ¹	Y ¹	Y ¹	Y ¹
VNX5700	Y	Y	Y	Y ¹	Y ¹	Y ¹
VNX7500	Y	Y	Y	Y	Y ¹	Y ¹
Gateway VG2	Y ¹	Y	Y ¹	Y ¹	Y	Y ¹
Gateway VG8	Y ¹	Y ¹	Y ¹	Y ¹	Y	Y

Note: The primary systems are listed horizontally and the standby systems are listed vertically.

The gateway models can only be connected to CX4 or VNX series backends.

It is recommended that the standby Data Mover be at least as performance and capacity capable as its primary Data Mover. These restrictions are not enforced by the code.

The `nas_rp -cabinetdr -init` command validates and enforces standby compatibility. The compatibility rule is simple: the standby Data Mover must have the same or more network devices as the primary Data Mover.

RecoverPoint configurations

RecoverPoint supports the following configurations:

- ◆ [RecoverPoint/SE active/passive on page 28](#)
- ◆ [RecoverPoint/SE active/active' on page 28](#)

¹ To make the primary Data Mover compatible with the standby Data Mover (from six cge ports to four cge ports), two ports out of the six must be masked. Masking will be done by using the `hidden_interfaces` parameter. [Ensure Data Mover network device compatibility on page 153](#) provides more information about the `hidden_interfaces` parameter.

RecoverPoint/SE active/passive

RecoverPoint/SE uses consistency groups, which have LUNs and their copies as replication sets, to replicate data sets. In the active/passive configuration, as shown in [Figure 2 on page 28](#) the consistency groups are configured for different LUN sets. SAN CGs are configured for SAN LUNs that are used for applications running on hosts other than the VNX for File system. NAS CGs are configured for user defined VNX for File LUNs and the VNX for File system's Control LUNs. Some unprotected LUNs are local LUNs that will not be accessible on the DR site after failover. The active/passive configuration has Data Mover DM1 on the Primary site, which is hosting file systems that are protected by RecoverPoint/SE. The Data Mover DM1 on the DR site is a standby for the Primary site DM1. The Data Mover DM2 on the Primary site hosts file systems that are not protected and the LUNs are not replicated. When a Primary site failure occurs, recovery of this unprotected data set is not possible until the Primary site is recovered and restored.




Figure 2. Active/passive configuration

RecoverPoint/SE active/active'

The active/active' configuration is more complex than the active/passive configuration because some LUNs are replicated in the reverse direction. As shown in [Figure 3 on page 29](#) one NAS consistency group exists for VNX for File LUNs on Primary A (A -> B) and another NAS consistency group exists for VNX for File LUNs on Primary B in the opposite direction (A <- B). Additionally, Data Mover DM3 on Primary B hosts file systems that are protected, and DM3 on Primary A is the remote standby for that DM3 on Primary B. If either Primary site fails, the other Primary site takes over 100 percent of the functionality, covering all protected file systems.

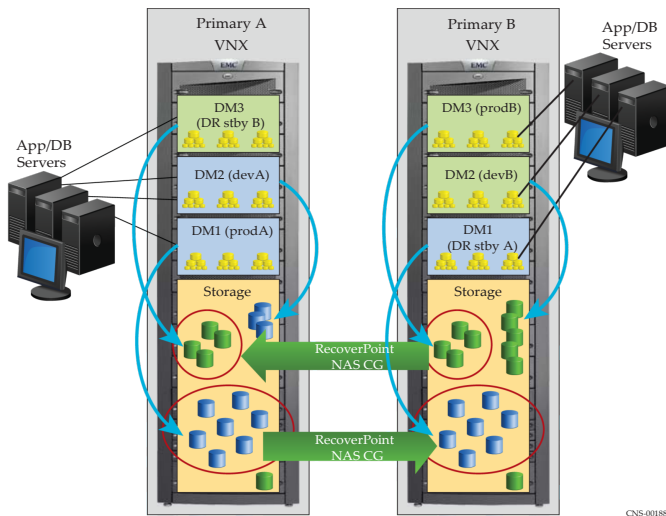


Figure 3. Active/active' configuration

RecoverPoint replication modes

RecoverPoint supports two data replication modes:

- ◆ [Asynchronous replication mode on page 29](#)
- ◆ [Synchronous replication mode on page 30](#)

Asynchronous replication mode

In *asynchronous replication mode*, the host application initiates a write and does not wait for an acknowledgment from the remote RPA before initiating the next write. The data of each write is stored in the local RPA, and acknowledged at the local site. The RPA decides based on the lag policy, system loads, and available resources when to transfer the writes in the RPA to the replica storage. This is the default replication mode.

The primary advantage of asynchronous replication is its ability to provide synchronous-like replication without degrading the performance of host applications.

Asynchronous replication, however, is not the preferred mode of replication in all situations. Asynchronous replication does not conserve bandwidth. Furthermore, and particularly as volumes increase, more data can be lost, as larger chunks of data that have been acknowledged at the local site may not be delivered to the target side in the case of a disaster.

RecoverPoint replicates asynchronously only in situations in which doing so enables superior host performance without resulting in an unacceptable level of potential data loss.

Asynchronous replication mode provides a **Snapshot Granularity** policy that regulates data transfer. The following granularities can be defined:

- ◆ **Fixed (per write)**: To create a snapshot from every write operation.
- ◆ **Fixed (per second)**: To create one snapshot per second.
- ◆ **Dynamic**: To have the system determine the snapshot granularity according to available resources.

Note: Fixed (per write) and fixed (per second) are not applicable to the NAS consistency group.

Synchronous replication mode

In *synchronous replication mode*, the host application initiates a write and then waits for an acknowledgment from the remote RPA before initiating the next write.

By default, new consistency groups are created with asynchronous mode enabled, and must be set to replicate synchronously.

Note: New consistency groups are created with the Measure lag when writes reach the target RPA (as opposed to the journal) setting enabled. When replicating synchronously, performance is substantially higher when this setting is enabled.

To ensure that no subsequent writes are made until an acknowledgment is received from the remote RPA, host applications are regulated by RecoverPoint. If your applications cannot be regulated for any reason, choose asynchronous replication mode.

Replication in synchronous mode produces a replica that is always one hundred percent up-to-date with its production source. The trade-off is that to ensure that no subsequent writes are made until an acknowledgment is received from the remote RPA, host applications can be regulated by RecoverPoint, and this could impact application performance. Alternatively, you can configure RecoverPoint to dynamically alternate between synchronous and asynchronous replication modes, according to predefined lag or throughput conditions. [Dynamic sync mode on page 31](#) provides more information about configuring RecoverPoint to dynamically alternate between the two modes.

Synchronous replication mode is only supported for replication to a remote replica over Fibre Channel.

Note: Synchronous replication mode is not supported for replication over the WAN.

Synchronous replication mode is efficient for replication both within the local SAN environment (as in CDP configurations, referred to as local replication), as well as for replication over Fibre Channel (as in CRR configurations, referred to as remote replication). However, when replicating synchronously, the longer the distance between the production source and the replica copy, the greater the latency.

To replicate data synchronously, your current RecoverPoint license must support synchronous replication. To verify that your current RecoverPoint license supports synchronous replication:

1. On the RecoverPoint view, click **Configure RecoverPoint Settings** in the RecoverPoint task block.
2. Expand **Account Settings**.
3. Verify that the word *Supported* is displayed next to **Synchronous Replication** in the **License Usage** section of this dialog box.

If you want to replicate synchronously and this feature is not supported in your current version of RecoverPoint, contact *EMC Customer Service*.

Dynamic sync mode

When replicating synchronously over longer distances, you can set RecoverPoint to replicate in *dynamic sync mode*, a submode of synchronous replication mode. In this mode, you can define group protection policies that enable the group to automatically begin replicating asynchronously whenever the group's data throughput or latency reaches a maximum threshold, and then automatically revert to synchronous replication mode when the throughput or latency falls below a minimum threshold.

You can also switch manually between replication modes, using the RecoverPoint CLI. This is useful, for example, if you generally require synchronous replication, but wish to use CLI scripts and your system scheduler to manually switch between replication modes during different times in the day, like during your nightly backups.

When the replication policy is controlled dynamically by both throughput and latency (both **Dynamic by latency** and **Dynamic by throughput** are enabled), it is enough that one of the two values of **Start async replication above** are met for RecoverPoint to automatically start replicating asynchronously to a replica. However, both **Resume sync replication below** settings must be met before RecoverPoint will automatically revert to synchronous replication mode.

To prevent jittering, the values specified for **Resume sync replication below** must be lower than the values specified for **Start async replication above**, or the system will issue an error.

Note: Groups undergo a short initialization every time the replication mode changes (for example, from synchronous to asynchronous). During initialization, data is not transferred synchronously.

RecoverPoint hardware and software

The following replication hardware and software are used in the RecoverPoint solution:

- ♦ [RecoverPoint Appliance on page 32](#)
- ♦ [Splitters on page 34](#)

- ◆ [RecoverPoint interfaces on page 35](#)

RecoverPoint Appliance

RecoverPoint appliances (RPAs) manage data replication at all sites. During replication, an RPA at the production site decides when and what data to transfer to the replica site. The RPA analyzes application load and resource availability so that it prevents degradation of host application performance while adhering to the specified replication policy. The RPAs at the replica site distribute the data to the replica storage.

In a failover, the RPA roles can be reversed. Moreover, RecoverPoint supports simultaneous bi-directional replication, where the same RPA can serve as the production RPA for one consistency group and the replica RPA for another consistency group.

The RPAs at each site form an RPA cluster. Each cluster can include between two and eight RPAs, as set during RecoverPoint system installation. The cluster size must be the same at all sites in an installation.

Physically, the RPA cluster can be located in the same facility as the host and storage subsystems. Alternatively, because RPAs have their own independent computing and storage resources, they can be located at a separate facility some distance away from the host or storage subsystems. This provides greater data protection in the event of a localized disaster. The *Release Notes for EMC RecoverPoint and Service Pack Releases* provide information on limitations.

During normal operation, all RPAs in a cluster are active all of the time. Consequently, if one of the RPAs in a cluster goes down, the RecoverPoint system supports immediate switchover of the functions of that RPA to another RPA in the cluster.

Each RPA has the following dedicated interfaces:

- ◆ *Fibre Channel*—Used for data exchange with local host applications and storage subsystems. The RPA supports a dual-port configuration to the Fibre Channel, thereby providing redundant connectivity to the SAN-attached storage and the hosts.
- ◆ *WAN*—Used to transfer data to other sites (Ethernet).
- ◆ *Management*—Used to manage the RecoverPoint system (Ethernet).

You can access an RPA directly through an SSH connection to the RPA's dedicated box-management IP address. You can also access all RPAs in the RecoverPoint configuration through the virtual site-management IP address of each site in the RecoverPoint configuration. This allows you to manage the entire installation from a single location.

Terminology

Preferred RPA

Each consistency group must be assigned one or more preferred RPAs. A non-distributed consistency group has one preferred RPA, called the Primary RPA. A distributed consistency group has multiple preferred RPAs: a minimum of one primary RPA and one secondary RPA, and a maximum of three secondary RPAs.

RecoverPoint cluster

A group of inter-linked RPAs, working together to provide replication services. The RPAs in a RecoverPoint cluster are called nodes. The nodes of the RecoverPoint cluster are connected to each other through the local area network, a wide area network, or by Fibre Channel. A RecoverPoint cluster can be deployed within a single site for CDP (or stretched CDP) configurations, or deployed in two sites for CRR and CLR configuration.

To scale-up and support a higher throughput rate, more RPA nodes can be added to the RecoverPoint cluster. RecoverPoint can be deployed with two to eight nodes per site. The cluster size must be the same at both sites in a RecoverPoint installation.

The RPA nodes provide high availability. If a node fails, the consistency groups using that node (that is, the consistency groups whose Primary or Secondary RPA are set to the failed RPA) flip-over to a different node in the RecoverPoint cluster.

The RecoverPoint cluster at each site is managed by a process called the site control. The RPA node that hosts the site control is selected using cluster leader arbitration (LEP) and can only be RPA1 or RPA2.

Physically, the RPA cluster can be located in the same facility as the host and storage subsystems. Alternatively, because RPAs have their own independent computing and storage resources, they can be located at a separate facility some distance away from the host or storage subsystems. This provides greater data protection in the event of a localized disaster. The *Release Notes for EMC RecoverPoint* and *Service Pack Releases* provide information on limitations.

During normal operation, all RPAs in a cluster are active all of the time. Consequently, if one of the RPAs in a cluster goes down, the RecoverPoint system supports immediate switchover of the functions of that box to another RPA in the cluster.

Primary RPA

The RPA that, whenever possible, handles replication for a consistency group. If an error occurs in the primary RPA, replication can in most cases be switched over to another RPA at the same side. The Primary RPA is set through the consistency group Policy tab.

RPA1

The RPA node that was designated as RPA1 of a RecoverPoint cluster, during a RecoverPoint installation.

RPA2

The RPA node that was designated as RPA 2 of a RecoverPoint cluster, during a RecoverPoint installation.

Note:

Only the first two RPAs (RPA 1 and RPA 2) in a RecoverPoint cluster can host the Site control services.

Site management

Also known as site control. The process that manages the RecoverPoint cluster at each site. In CDP configurations, there is only one site control. In CRR and CLR configurations, there are separate site controls for each site in the configuration. The active instance of the site control is run only by RPA 1 or RPA 2.

The user accesses the site control to manage and monitor RecoverPoint, using a Site Management IP address. To run the EMC RecoverPoint Management Application GUI, the user connects to a RecoverPoint cluster by opening a browser window and typing the site management IP of the RecoverPoint cluster into the browser address bar.

To run the EMC RecoverPoint Command Line Interface, the user would connect to a RecoverPoint cluster by opening an SSH connection with the site management IP of the RecoverPoint cluster.

To identify the site control, log in to the RecoverPoint Management Application as a user with SE privileges, and click on the RPAs section of the Navigation Pane. The ID column of the RPAs table, displays an asterisk for each RPA acting as the Site Control.

Note:

The RPA node that can be used to host the site control is selected using cluster leader arbitration (LEP) and can only be RPA1 or RPA2.

Site management IP

A virtual, floating IP address assigned to the RPA that is currently active (runs the site control). In this RPA fails, the floating IP address dynamically switches to the RPA that assumes operation (either RPA1 or RPA2). Although using the site management IP is best practice, all management activities can also be performed on a specific RPA, by entering its dedicated IP address.

Splitters

A *splitter* is proprietary software that is installed on either host operating systems, storage subsystems, or intelligent fibre switches. Splitters access replica volumes; that is, volumes that contain data to be replicated. The primary function of a splitter is to "split" application writes so that they are sent both to the storage volumes and the RPA simultaneously. The splitter has no perceptible impact on host performance, since all CPU-intensive processing necessary for replication is performed by the RPA.

An array-based splitter is used by EMC VNX as well as EMC CLARiiON storage arrays to split data. This feature is only available for EMC VNX and EMC CLARiiON storage. The RecoverPoint array-based splitter, therefore, is referred to as the *CLARiiON splitter*. The array-based splitter can either be bundled with the VNX operating environment for Block or the FLARE[®], or a splitter enabling package can be non-disruptively upgraded.

Note: The EMC Support Matrix on the EMC Online Support website provides exact support statements including OS versions, and other caveats in supported and unsupported configurations.

RecoverPoint interfaces

RecoverPoint provides the following interfaces that you can use to manage a RecoverPoint environment:

- ◆ The *RecoverPoint management console* provides a graphical user interface (GUI) that allows you to manage RecoverPoint through a web browser. The *RecoverPoint Administrator's Guide* provides more information.
- ◆ The *RecoverPoint/SE for Unisphere management console* integrates RecoverPoint/SE functionality into the VNX Unisphere management interface. Using RecoverPoint/SE for Unisphere, you can manage a RecoverPoint/SE environment from the same central Unisphere management console that you use to manage your VNX or CLARiiON storage environment. You do not need to invoke the classic RecoverPoint interface to manage your RecoverPoint environment. This help system describes the RecoverPoint/SE for Unisphere management console.

Note:

This help system describes how to use the RecoverPoint/SE for Unisphere management console.

-
- ◆ The *RecoverPoint command line interface (CLI)* allows you to manage RecoverPoint interactively from the command line or through scripts. The *EMC RecoverPoint CLI Reference Guide* provides more information.

RecoverPoint logical entities

The following logical entities constitute your replication environment:

- ◆ [Consistency groups on page 35](#)
- ◆ [Copies on page 39](#)
- ◆ [Replication sets on page 40](#)
- ◆ [Replication policies on page 40](#)
- ◆ [Journals on page 40](#)
- ◆ [Volumes on page 41](#)
- ◆ [Snapshots on page 43](#)

Consistency groups

A consistency group consists of one or more replication sets. Each *replication set* consists of a production volume and the replica volumes to which it is replicating. The consistency

group ensures that updates to the replicas are always consistent and in correct write order. The replicas can be used to continue working or to restore the production source, in case it is damaged.

The consistency group monitors all the volumes added to it to ensure consistency and write-order fidelity. If two data sets are dependent on one another (for instance, a database and a database log), they must be in the same consistency group. Imagine a motion picture film. The video frames are saved on one volume, the audio on another. But neither will make sense without the other. The saves must be coordinated so that they will always be consistent with one another. In other words, the volumes must be replicated together in one consistency group. That will guarantee that at any point, the saved data will represent a true state of the film.

A consistency group consists of:

- ◆ *Settings*, such as consistency group name, preferred RPA, reservation support.
- ◆ *Policies*, such as compression, bandwidth limits, and maximum lag, that govern the replication process.
- ◆ *Replication set*: A production source volume and the replica volumes to which it replicates.
- ◆ *Journals*: Receive changes to data. Each copy has a single journal. Changes are distributed from the replica journal to storage. The replica journals also retain rollback data for their replica.

The production journal does not contain rollback information. The system marking information is contained in the production journal.

When a consistency group, copy, or volume is defined in RecoverPoint for the first time, its volumes are initialized.

In RecoverPoint, a consistency group can be:

- ◆ Non-distributed (regular)
- ◆ Distributed

Note:

Throughout the RecoverPoint documentation, the term *consistency group* is used to refer to groups when no differentiation is required between distributed and non-distributed groups.

For the complete set of limitations associated with consistency groups, see the *EMC RecoverPoint Release Notes*.

Distributed consistency groups

The distributed consistency group feature allows you to create consistency groups that require a total throughput and IOPS rate that exceeds the supported throughput and IOPS rate of a single RecoverPoint appliance. This prevents you from having to split data that requires strict write-order fidelity into multiple consistency groups. Distributed groups can

handle a much higher throughput and IOPS rate regardless of the amount of data being replicated.

How do distributed consistency groups work

Distributed consistency groups are divided into four segments and these segments are transferred through one primary RPA and up to three secondary RPAs.

RecoverPoint data recovery processes are affected in the following way:

- ◆ The primary RPAs at both sites (if two sites exist) are responsible for the receipt and handling of all system process requests.
- ◆ All of the marking information is handled by the primary RPA at the source-side.

When to use a distributed consistency group

You should consider setting a consistency group as *distributed* when:

- ◆ The maximum throughput rate of a single RPA is not sufficiently sustaining the write-rate or peaks of the consistency group.
- ◆ A consistency group often experiences high loads.
- ◆ You expect that a consistency group will require a higher throughput rate than that of a single RPA. In this case, it is preferable to initially create the consistency group as distributed, rather than modifying an existing consistency group after creation.

Limitations

Up to eight distributed consistency groups can be defined in RecoverPoint, and the total number of distributed and non-distributed consistency groups is 128.

When setting a group as distributed, the following limitations apply:

- ◆ The snapshot granularity of all links in the consistency group can be no finer than 1 second.
- ◆ Journal loss occurs when modifying a group's topology (setting a non-distributed group as distributed, or setting a distributed group as non-distributed).
- ◆ When configuring journals for a distributed consistency group, keep the following in mind:
 - All copies of distributed consistency groups must have a journal that is at least 20 GB in size.
 - The recommended journal size for distributed groups with snapshot consolidation enabled is at least 40 GB.
 - If the capacity of an existing copy journal is less than the minimum journal size required for distributed consistency groups, the consistency group will need to be disabled and then enabled again after adding journal volumes, and this will cause a full sweep.

- ◆ Distributed consistency groups are only supported if there is a Fibre Channel connection between all RPAs in a RecoverPoint cluster (per site). Therefore:
 - In Fibre Channel environments, make sure all of the RPAs at each site are connected to the SAN through a Fibre Channel switch, and zoned together so that they see each other in the SAN.
 - In iSCSI environments, make sure all RPAs are physically connected to each other through their HBA Fibre Channel ports.

Note:

In iSCSI configurations, there can be a maximum of two RPAs per RecoverPoint cluster (that is, per site) because two of the four existing Fibre Channel ports in the RPA's HBA are already connected directly to the storage. If more than two RPAs per RecoverPoint cluster are required, connect all of the RPAs in the cluster through Fibre Channel switches (two should be used for high availability) and zone them together.

- ◆ If any of the primary or secondary RPAs associated with a consistency group becomes unavailable, there will be a brief pause in transfer on all of the group's primary and secondary RPAs, and all of the group segments will undergo a short initialization.
- ◆ Under certain circumstances (for example, if one of the primary or secondary RPAs becomes unavailable) two consistency group segments could be handled by the same RPA.
- ◆ In general, distributed consistency groups offer better performance than non-distributed (regular) consistency groups, as distributed groups run on a minimum of two RPAs (one primary RPA and one secondary RPA). There is only a small improvement in performance when a group is run on three RPAs. However, there is a steep improvement in performance when a group is run on four RPAs.

For the complete set of limitations associated with distributed consistency groups, see the *EMC RecoverPoint Release Notes*.

How to verify distributed consistency group support

To verify that your RecoverPoint environment supports distributed consistency groups, select **Configure RecoverPoint Settings** and click on the **Account Settings** tab. In the **License Usage** section, you should see the text *Distributed Groups: Supported*. If this text does not appear, contact EMC Customer Support.

How to set a consistency group as distributed

To set a non-distributed group as distributed, or set a distributed group as non-distributed:

1. On the Consistency Group tab, select the consistency group and click **Properties**. Select **Policy** and expand **Advanced**.
2. Select or clear the **Distribute group** checkbox.

3. If you are enabling this feature, select the secondary RPAs for the distributed consistency group.

Note: When modifying a consistency group's topology, journal loss occurs.

Standard (non-distributed) consistency groups

New consistency groups are by default defined as non-distributed. *Non-distributed consistency groups* transfer data through one primary RPA that is designated when the consistency group was created. You can select the primary RPA through the group policy settings.

A maximum of 128 consistency groups can be defined in RecoverPoint, and a single RPA cannot be configured to have more than 64 consistency groups. In the event of RPA failure, groups that transfer data through one RPA will move to other RPAs in the cluster.

In such a case, an RPA can temporarily hold up to 128 groups, and the data of all groups will continue to be transferred between sites. This state is temporary, however, as an RPA with more than 64 groups may run into high loads, and if this state is prolonged, group policies could be affected.

Each RPA has a maximum throughput rate (the *EMC RecoverPoint Release Notes* provides information on this limit), which together with the host write-rate and available network resources, limits the maximum size of the consistency group. For a higher throughput rate, and to balance the load of extra-large consistency groups, define the consistency group as distributed.

Copies

A copy constitutes all of the volumes defined for replication at a given location (production, local, or remote). A copy includes a journal size limit setting that defines RTO, journal compression policies, and protection policies that define snapshot consolidation and the required protection window.

In CDP and CRR configurations, there is one production copy and one replica. In CLR configurations, there is one production copy and two replicas (one local copy at the production site and one remote copy at the disaster recovery site).

Note:

The term "replica" is used to differentiate between production and non-production copies, whenever necessary. In CLR configurations there are two replicas, or non-production copies (also known as targets).

The production copy consists of production volumes and the production journal, which may consist of one or more journal volumes. The non-production copies (that is, replica copies) each consist of replica volumes and a replica journal, which may consist of one or more journal volumes.

Replication sets

Every SAN-attached storage volume in the production storage must have a corresponding volume at each copy. A production volume and its associated replica volumes are called a *replication set*. Each consistency group contains as many replication sets as there are volumes in the production storage to replicate. Data consistency and write-order fidelity are maintained across all volumes assigned to a consistency group, including volumes on different storage systems.

At least one volume must be added to the journal of each copy in a replication set.

Replication policies

A replication policy is a set of parameters driven by business objectives that control system operation during replication. The set of parameters that compose a policy that matches the customer's replication business objective, and enforce it automatically through preset RPO, RTO, and resource allocation and compression policies.

Replication with the RecoverPoint system is policy-driven. A replication policy, based on the particular business needs of your company, is uniquely specified for each consistency group, and each copy. The policy comprises a set of parameters that collectively governs the way in which replication is carried out. Replication behavior changes dynamically during system operation in light of the policy, the level of system activity, and the availability of network resources.

Journals

A *journal* is one or more volumes dedicated on the storage at each copy in a RecoverPoint configuration. Journals are defined per copy, and can consist of multiple journal volumes. Each journal is maintained independently.

There are two types of journals:

- ◆ Each consistency group has a single *production journal*, which contains the system delta marking information. The production journal does not contain snapshots used for PIT recovery.
- ◆ Each replica (non-production copy) has a dedicated *replica journal*. Replica journals are used to hold snapshots that are either waiting to be distributed, or that have already been distributed to the replica storage, metadata for each image, and bookmarks. The replica journals receive changes to data. Changes are distributed from the replica journal to the replica storage. Then, the data of the replica storage is stored in the UNDO stream of the replica journal, so that the replica storage can be rolled back to a previous PIT.

Each replica journal holds as many snapshots as its capacity allows. The oldest snapshot—after being distributed to the copy storage—is removed to make room for the

newest snapshot, in a cyclic manner. The actual number of snapshots in the journal varies, depending on the size of the snapshot and the capacity of the storage dedicated to this purpose.

You can address individual snapshots in a replica journal. Therefore, if required due to a disaster, you can roll back the stored data image to an earlier PIT that was unaffected by the disaster. Frequent small-aperture snapshots provide high granularity for achieving maximum data recovery in the event of a disaster.

The replica journal snapshot data can be compressed, or deduped, using snapshot consolidation.

There is a minimal size for each journal. The *EMC RecoverPoint Release Notes* provides more information on for this limitation. By default, 20 percent of the journal's capacity is dedicated to the image access log used during Logged access (although this default value can be modified), another 5 percent of the journal's capacity is dedicated to the calculation of indexes that are used for Virtual access, and an additional ~1 GB is dedicated to handling bursts during distribution. This means that only ~75 percent of the journal is available to store snapshots.

For efficiency purposes, ensure that all volumes contained in a journal are of the same size for disk striping purposes. Otherwise, RecoverPoint defines the capacity of all of the volumes by the smallest volume in the journal, and any available space beyond the capacity of the smallest volume is not used.

Volumes

The RecoverPoint interface represents LUNs as *volumes*. Therefore, this help file refers to *LUNs* when referencing the storage entity and to *volumes* when referencing the RecoverPoint entity.

The following types of volumes exist in all RecoverPoint configurations:

- ◆ [Production volumes on page 42](#)
- ◆ [Replica volumes on page 42](#)
- ◆ [Production journal volume on page 42](#)
- ◆ [Replica journal volume on page 42](#)

- ◆ [Repository volume on page 43](#)

Production volumes

The production volumes are the volumes that are written to by the host applications at the production site.

Replica volumes

The replica volumes are the volumes to which the production volumes are replicated.

Production journal volume

The production journal volume stores information about the replication process—called *marking information*—that is used to make synchronization of the replication volumes at the two sites, when required, much more efficient. Each production journal can consist of one or more LUNs.

Note: In RecoverPoint/SE, the production and local replica journals and repository volume must all reside on the same VNX for Block or CLARiiON system.



CAUTION Since the production journal contains the system marking information, the removal of a journal volume from the production site will cause a full-sweep synchronization.

Replica journal volume

Each replica journal can consist of one or more LUNs.

Note: In RecoverPoint/SE only: The production and local replica journals and repository volume must all reside on the same VNX for Block or CLARiiON system.

If more than one volume at a time is added to the journal, it is recommended that all added volumes have the same capacity for best performance and efficiency. If the added volumes have the same or nearly the same capacity (at least 85 percent of the largest volume), data is striped across those journal volumes, improving performance. When striped, the capacity used in each journal volume is equal to the capacity of the smallest journal volume in the group of added volumes; remaining capacity in those volumes is not used.

Volumes of very different capacities will be concatenated and not striped. In most cases, this will affect performance, but all capacity will be used.

If two groups of volumes of two different capacities are added, they are striped in two groups. If additional volumes are added afterwards, the new volumes will be considered as a group by themselves according to the criteria above. Existing volumes and newly added volumes will not be striped together.

In the case that the combined physical size of all journal volumes at a given copy is larger than the combined physical size of the journal volumes at the other copy, the protection window at the first copy will be larger than the protection window at the other copy.

Note: Journal volumes cannot reside on LUNs that are virtually provisioned (thin LUNs).

Repository volume

A special volume—the *repository volume*— must be dedicated to the SAN-attached storage at each site, for each RPA cluster. The repository volume serves all RPAs of the particular cluster and splitters associated with that cluster. It stores configuration information about the RPAs and consistency groups, which enables a properly functioning RPA to seamlessly assume the replication activities of a failing RPA from the same cluster.

Note: In RecoverPoint/SE only: The production and local replica journals and repository volume must all reside on the same VNX for Block or CLARiiON system.

Snapshots

A *snapshot* is a point in time marked by the system for recovery purposes. A snapshot includes only that data that has changed from the previous image. When distributed, it creates a new current image on the remote storage.

A *snapshot* is the difference between one consistent image of stored data and the next. Snapshots are taken seconds apart. The application writes to storage; at the same time, the splitter provides a second copy of the writes to the RecoverPoint appliance. In asynchronous replication, the appliance gathers several writes into a single snapshot. The exact time for closing the snapshot is determined dynamically depending on replication policies and the journal of the consistency group. In synchronous replication, each write is a snapshot. When the snapshot is distributed to a replica, it is stored in the journal volume, so that it is possible to revert to previous images by using the stored snapshots.

The snapshots at a copy are displayed in the copy **Journal Tab**.

For each consistency group, a *Snapshot Granularity* policy can be configured to regulate data transfer, and the following granularities can be defined:

- ◆ *Fixed (per write)*: To create a snapshot from every write operation.
- ◆ *Fixed (per second)*: To create one snapshot per second.
- ◆ *Dynamic*: To have the system determine the snapshot granularity according to available resources.

Note: Fixed (per write) and fixed (per second) are not applicable to the NAS consistency group.

Bookmarks

A *bookmark* is a named snapshot. The bookmark uniquely identifies an image. Bookmarks can be set and named manually. Bookmarks can also be created automatically by the system either at regular intervals or in response to a system event. Bookmarked images are listed by name.



Time	Bookmark	App
01/01 12:51:30		
01/01 12:51:14	 Server restored.	
01/01 12:50:33		
01/01 12:49:39	 Server crashed.	
01/01 12:49:36		
01/01 12:48:39		
01/01 12:47:42		

Figure 4. Examples of snapshots and bookmarks.

You can bookmark a snapshot at any time. Bookmarks are useful to mark particular points in time, such as an event in an application, or a point in time to which you want to fail over.

The bookmarked snapshots at a copy are displayed in the copy **Journal Tab**.

Group sets

A *group set* is a set of consistency groups to which the system applies parallel bookmarks at a user-defined frequency. Group sets are useful for consistency groups that are dependent on one another or that must work together as a single unit.

A group set allows you to automatically bookmark a set of consistency groups so that the bookmark represents the same recovery point in each consistency group in the group set. This allows you to define consistent recovery points for consistency groups that are distributed across different RPAs.

The automatic periodic bookmark consists of the name you specified for the group set and an automatically incremented number. Numbers start at zero, are incremented up to 65535, and then begin again at 0.

The same bookmark name is used across all the groups. To apply automatic bookmarks, the sources must be at the same site (replicating in the same direction) and transfer must be enabled for each consistency group included in the group set.

Note: It is recommended that the NAS consistency group not be added to any group sets.

Planning considerations

Before configuring RecoverPoint with VNX, consider:

- ◆ [File system recommendations on page 45](#)
- ◆ [Remote administration account recommendations on page 45](#)
- ◆ [VNX for Block configuration on page 46](#)
- ◆ [VNX volume and Data Mover decisions/flexibility on page 47](#)
- ◆ [Data Mover configuration checklist on page 49](#)
- ◆ [RecoverPoint/SE configuration sheet on page 50](#)

File system recommendations

- ◆ With a new installation of RecoverPoint/SE on VNX for File, perform a RecoverPoint/SE initialization by using the `/nas/sbin/nas_rp -cabinetdr -init` command before building the file systems. This avoids the extra configuration steps associated with the conversion of disk types from unreplicated to replicated during the initialization process.
- ◆ Ensure that you mount the local file systems on local Data Movers and the replicated file systems on RecoverPoint-protected Data Movers. [Ensure Data Mover eligibility on page 151](#) describes the Data Mover configurations that are checked during RecoverPoint/SE initialization. Also, [Chapter 7](#) describes RecoverPoint/SE DR failure scenarios and errors.

Note: If the configuration includes a replicated file system mounted on a local Data Mover, the replicated file system is not protected in the event of a failover.

- ◆ For high-availability configurations such as RecoverPoint/SE, do not span a file system across multiple storage systems.

Remote administration account recommendations

In an active/active' configuration, ensure that a different user ID (UID) for a remote administration account user is used for each RecoverPoint/SE CRR direction. For example, `dradmin`. Having different UIDs for the remote administration account user in each direction in the active/active' configuration ensures that the correct Data Mover (server) information is always displayed for the appropriate command when a failover is activated. [Verify remote administration account \(Optional\) on page 110](#) provides details on how to create a remote administration account user.

VNX for Block configuration

The VNX for Block configuration setup tasks for RecoverPoint/SE, performed by your local EMC Service Provider, ensure that the LUNs on the VNX storage arrays are properly configured in conformance with RecoverPoint/SE requirements and that the RecoverPoint/SE links are operational.

Summary of VNX for Block configuration steps

On each VNX for Block system, the configuration steps performed by your local EMC Service Provider by using Unisphere (or Navicli) include:

1. A single domain must be established to contain the source and destination arrays. All devices that are replicated must be in the same domain.
2. A RecoverPoint/SE connection between the source and destination arrays must be established before any data can be replicated. The RecoverPoint/SE links connecting the storage processors on the arrays usually use the highest-numbered port.
3. Each source LUN, which represents a primary image, must be prepared. If a source LUN does not exist, it must be bound on the source storage system. Then, it is assigned to the primary VNX for File system's storage group. As part of the basic VNX for Block setup, three source Control Station LUNs with host LUN IDs numbered 0,1, and 4 must exist for the dos, log, and nas LUNs, respectively.
4. For each source LUN, an equivalent destination LUN must be created on the destination storage system. A destination LUN represents a secondary image. The source and its destination LUN must be the same size. As part of this process, the destination LUN should finish binding, and the backup LUN should not be assigned to any storage group. You can assign replicated LUNs to user defined pools.

Note: [Create LUNs to match the LUNS on the source VNX on page 62](#) provides more information.

5. For each source data LUN on the source storage system, a remote replication set must be created.
6. The destination LUNs must be assigned to the destination VNX for File system's storage group. Although these LUNs are not visible to the destination VNX for File until after being promoted, they can still be in a storage group after each is added to a replication set as a secondary image. This step requires fixed LUN mappings, the dos, log, and nas control LUNs must have host LUN IDs 0, 1, and 4 on the source and are mapped to 6, 7, and 9, respectively on the destination.
7. From the source site, a RecoverPoint/SE consistency group must be created. You can see the assigned name of the consistency group during the RecoverPoint/SE

initialization procedure. Then, add the source and remote LUN pairs as replication sets.

Guidelines for changing your VNX configuration

- ◆ Storage system configuration changes can be performed by your local EMC Service Provider after initial configuration as long as a failover has not been activated. No storage system configuration changes should be made if a failover has been activated.
- ◆ Any VNX configuration change that affects any of the replication sets, storage groups, or consistency groups used by the system requires you to rerun the `/nas/sbin/nas_rp -cabinetdr -init` command from the destination system. If you want to add storage or remote replication sets to the consistency group, consult your local EMC Service Provider.
- ◆ If the global account password changes, the VNX for File storage security information on the Control Station must also be updated. After initialization, you can capture this change by rerunning the `/nas/sbin/nas_rp -cabinetdr -init` command. After a failover, you must use the `nas_storage` command with the `modify` option. [Modify VNX for Block security information after a failover on page 162](#) provides more information.

VNX volume and Data Mover decisions/flexibility

With either active/passive or active/active' RecoverPoint/SE configuration, you can choose which LUNs (disk volumes) and Data Movers to protect with remote replication. For example, a typical active/passive configuration provides a full backup of the source site, Data Mover for Data Mover. However, in general, you can remotely replicate some LUNs and Data Movers while others are only locally protected. You do not have to protect all volumes and Data Movers.

When planning the RecoverPoint/SE Data Mover configuration:

- ◆ For each source Data Mover you choose to protect with a remote RecoverPoint/SE standby Data Mover, you must provide a dedicated standby Data Mover at the destination site. There must be a one-to-one relationship between a given source Data Mover that you choose to protect and a dedicated remote standby Data Mover at the destination site. In addition, the source Data Mover and its remote RecoverPoint/SE standby must use the same slot number. For example, `server_2` in slot 2 on the source and `server_2` in slot 2 at the destination.
- ◆ If you want one or more source site RecoverPoint/SE-protected Data Movers to also have a local standby Data Mover, the requirements for the local standby are:
 - The local standby must serve only RecoverPoint/SE-protected source Data Movers. The local standby cannot serve a mix of RecoverPoint/SE-protected and non-RecoverPoint/SE Data Movers.
 - Any local standby that serves a RecoverPoint/SE-protected source Data Mover must also have a dedicated remote RecoverPoint/SE standby Data Mover. In addition, the local standby and its remote standby must use the same slot number. For example,

local standby server_3 in slot 3 must have a remote standby that also uses server_3 in slot 3. Having multiple local standby Data Movers serving a given RecoverPoint/SE-protected source Data Mover is also supported as long as each one of these local standby Data Movers has a remote standby Data Mover, and the remote standby uses the same slot number as the local standby.

Note: After you have established a RecoverPoint/SE configuration, you should use the `/nas/sbin/nas_rp -cabinetdr -init` command, not the `server_standby -create` command (run as root), to manage the remote standby Data Mover relationships and verify relationships after standby configuration. Using the `nas_rp -cabinetdr -init` command helps prevent misconfiguration because the initialization procedure checks Data Mover eligibility, and validates and enforces standby Data Mover compatibility with the source Data Mover. The possible Data Mover conditions that you might see during initialization are listed in [Ensure Data Mover eligibility on page 151](#). It might be helpful to review these conditions in advance. You can use the `server_standby` command to change a local-only standby configuration and then use the `nas_rp -cabinetdr -init` command to validate the configuration change.

- ◆ The network configuration of the RecoverPoint/SE standby Data Mover must be a superset of the network configuration of the source Data Mover. The network devices of the standby Data Mover must match those of the source Data Mover or be a superset. The RecoverPoint/SE initialization procedure validates and enforces standby compatibility. If you have questions about Data Mover compatibility between source and standby Data Movers, consult the E-Lab Interoperability Navigator, which is available on Powerlink.
- ◆ The network configuration of the RecoverPoint/SE standby Data Mover must be able to access the destination data LUNs corresponding to the source Data Mover.

[Table 3 on page 48](#) and [Table 4 on page 49](#) list the Data Mover configurations you can have at the local and remote sites in active/passive and active/active' RecoverPoint/SE environments.

Table 3. Permissible active/passive Data Mover configurations

Source site Data Mover	Direction of DR relationship	Destination site Data Mover
Local source Data Mover that will be RecoverPoint/SE-protected	----->	RecoverPoint/SE remote standby Data Mover
Local standby Data Mover for RecoverPoint/SE-protected source Data Mover	----->	RecoverPoint/SE remote standby Data Mover using same slot number as the source-site local standby
Local source Data Mover (non-RecoverPoint/SE)	No relationship	Local source Data Mover (non-RecoverPoint/SE)
Local standby Data Mover (non-RecoverPoint/SE)	No relationship	Local standby Data Mover (non-RecoverPoint/SE)

Table 4. Permissible active/active' Data Mover configuration

Site A Data Mover	Direction	Site B Data Mover
Local source Data Mover that will be RecoverPoint/SE-protected	----->	RecoverPoint/SE remote standby Data Mover
Local standby Data Mover for RecoverPoint/SE-protected source Data Mover	----->	RecoverPoint/SE remote standby Data Mover using same slot number as the other source-site local standby
Local source Data Mover (non-RecoverPoint/SE)	No relationship	Local production Data Mover (non-RecoverPoint/SE)
Local standby Data Mover (non-RecoverPoint/SE)	No relationship	Local standby Data Mover (non-RecoverPoint/SE)
RecoverPoint/SE remote standby Data Mover	----->	Local source Data Mover that will be RecoverPoint/SE-protected
RecoverPoint/SE remote standby Data Mover using same slot number as the other source-site local standby	----->	Local standby Data Mover for RecoverPoint/SE-protected source Data Mover

Data Mover configuration checklist

To ensure proper Data Mover configuration:

1. Decide and list which Data Movers to designate as RecoverPoint/SE source and RecoverPoint/SE remote standby Data Movers. This is a one-to-one failover relationship. A Data Mover can be a RecoverPoint/SE standby for only one source Data Mover, and a RecoverPoint/SE source Data Mover can only have one RecoverPoint/SE standby. In the initialization procedure, these failover relationships are designated and assigned.

By default, Data Movers are referred to as servers and are named `server_n`, starting with `server_2`.

2. Make sure each local standby Data Mover providing local standby coverage for a RecoverPoint/SE-protected Data Mover has a corresponding remote RecoverPoint/SE standby.
3. For source Data Mover and remote standby Data Mover compatibility, consult the EMC E-Lab Interoperability Navigator and make sure:
 - The source Data Mover and its remote standby are of a similar model type, that is, the local RecoverPoint/SE Data Mover and its corresponding remote standby Data Mover should be the same model, or the remote standby represents a superset of the source Data Mover's network configuration.

- The remote standby Data Mover has the same performance and capacity as its source Data Mover.
4. If you see a Data Mover condition of not compatible during initialization, the source Data Mover has one or more network devices not available on the remote standby Data Mover. The network devices in the two cabinets do not have the same configuration.
 5. Consider IP data network connectivity problems when planning the Data Mover assignments.
 6. Ensure that network interfaces for the RecoverPoint/SE source and RecoverPoint/SE standby Data Movers are identical, and the same set of network clients can access the RecoverPoint/SE source and corresponding remote standby Data Movers.
 7. Be aware that a failover operation moves the network configuration associated with the source Data Mover to the destination Data Mover. For example, if cge0 on the source Data Mover is connected to subnet A on the source side, cge0 on the destination Data Mover must be connected to subnet A on the destination side. If the destination side has a different subnet network configuration, you must change the IP addresses on the destination Data Mover manually after the activation completes. [Ensure access after failover on page 99](#) provides more information.
 8. Evaluate the infrastructure of the destination site, such as its subnet addresses, the availability of NIS/DNS servers in the right UNIX domain, the availability of WINS/PDC/BDC/DC in the right Windows domain, and the availability of NTmigrate or Usermapper hosts. The CIFS environment requires more preparation to set up a RecoverPoint/SE configuration because of the higher demands on its infrastructure than with the NFS environment. For example, authentication is handled by the infrastructure versus client OS. For the CIFS environment, you must perform mappings between usernames/groups and UIDs/GIDs.
 9. Evaluate whether to establish network firewalls between the pair of source and destination systems. If so, any firewalls between the sites must allow for Control Station-to-Control Station communication across the IP network by using ports 80 and 8000 for HTTP, port 6389 for NaviCLI, and port 443 for Navisphere Secure CLI (naviseccli). Port 8000 must be open for Control Station-to-Control Station communication over an IP link in general.

RecoverPoint/SE configuration sheet

[Table 5 on page 51](#) provides a tracking sheet for RecoverPoint/SE configuration information. For an active/passive configuration, use one sheet and for an active/active' configuration, use two sheets.

Note: For detailed information about the VNX for Block configuration for RecoverPoint/SE, such as a list of the VNX for Block LUN mappings, consult the expanded configuration sheet put together by your local EMC Service Provider (as VNX for Block administrator) during the source and destination VNX for Block setup. It is expected that the detailed configuration sheet is left at your site after initial setup.

Table 5. Basic RecoverPoint/SE configuration tracking sheet

What you specify	Source-site information	Destination-site information
Control Station name		
Control Station IP address		
Password specified with nas_cel (must be the same for both)		
Existing VNX for Block storage information Global VNX for Block account user-name: _____ Global VNX for Block password: _____	APM #:	APM #:
Data Mover DR pair	Source server: Type:	Standby server: Type:
Data Mover DR pair	Source server: Type:	Standby server: Type:
Data Mover DR pair	Source server: Type:	Standby server: Type:
Data Mover DR pair	Source server: Type:	Standby server: Type:
Data Mover DR pair	Source server: Type:	Standby server: Type:
Remote administration account and password		Created on destination during initialization: Username: Password:
RPA administration account and password		
Management IP address		

Task overview

Table 6 on page 52 provides an overview of the basic tasks to establish a RecoverPoint/SE active/passive or active/active' configuration and their associated commands.

Before performing any of these tasks, review the requirements summarized in [Planning considerations on page 45](#) to ensure that the VNX for Block and VNX for File Data Movers are configured correctly. Your local EMC Service Provider establishes and verifies the initial RecoverPoint/SE CRR replication configuration including preinitialization, initialization, test failover, and a test failback, and also manages the VNX for Block configuration changes.

Table 6. RecoverPoint/SE CRR replication task overview

Task summary	Command used	What it does
Preinitialize the relationship between the source and destination VNX systems.	From each VNX (source and destination), using nasadmin: <pre>nas_cel -create <cel_name> -ip <ip> -passphrase <passphrase></pre> and <pre>nas_rp -rpa -add <name> -ip <rpa_ip> -admin <rpa_admin_name> [-password <password>]</pre>	<ul style="list-style-type: none"> ◆ Establishes trusted communication between source and destination VNX systems as a prerequisite to RecoverPoint/SE CRR replication active/passive or active/active' initialization. Must be performed on both VNX systems by using the same passphrase (6–15 characters). ◆ Adds the RPA information to the NAS database.
Initialize the RecoverPoint/SE CRR replication relationship between the source and destination VNX systems.	From the destination VNX, using nasadmin, su to root: <pre>/nas/sbin/nas_rp -cabinetdr -init <cel_name></pre> Note: The /nas/sbin/nas_rp -cabinetdr -init command should be run after any storage configuration change that affects any of the RPA, storage groups, or RP consistency groups used by the VNX for File system.	<ul style="list-style-type: none"> ◆ Identifies the VNX system that is paired with the current Culham system in the configuration. ◆ Prompts for the VNX for Block global user account information. ◆ Establishes the remote administration account that is used to manage the destination. ◆ Checks for Data Mover compatibility and establishes the Data Mover relationships from the source Data Mover to the standby. ◆ Checks the VNX for Block storage system configuration to ensure proper LUN configuration and mapping, storage group information, and device group information.

Table 6. RecoverPoint/SE CRR replication task overview *(continued)*

Task summary	Command used	What it does
Failover data and services from a source VNX to a destination VNX.	From the destination VNX Control Station, using the remote administration account set up during initialization (for example, rpdadmin), su to root: <code>/nas/sbin/nas_rp -cabinetdr -failover</code>	<ul style="list-style-type: none"> ◆ Performs a manual failover scenario (for example, a source VNX has become unavailable and requires a failover to the destination VNX). ◆ Attempts to shut down source services and Data Movers gracefully if they have not already been shut down. ◆ Promotes the destination LUNs, which become read/write. The source RecoverPoint/SE-protected file systems are no longer available on the source. ◆ Enables each RecoverPoint/SE standby Data Mover on the remote system to become active and acquire the following characteristics of its source counterpart: <ul style="list-style-type: none"> ◆ Network identity: IP and MAC addresses of all network interface cards (NICs) in the failed Data Mover ◆ Service identity: Network File System/Common Internet File Service (NFS/CIFS) characteristics of the exported file system controlled by the failed Data Mover

Table 6. RecoverPoint/SE CRR replication task overview *(continued)*

Task summary	Command used	What it does
Failback a source VNX after a failover.	From the destination VNX Control Station, using the remote administration account and su to root: <code>/nas/sbin/nas_rp -cabinetdr -failback</code>	Typically scheduled and performed under the guidance of your local EMC Service Provider or EMC Customer Service to ensure continuity between the VNX for Block components. Failback of a source VNX for File system involves a complete check of the VNX for Block system and RecoverPoint/SE and verification of full connectivity to the restored file systems on the source VNX for File: Copies data from destination-site LUNs to the corresponding source-site LUNs on the source VNX for Block system. Reboots RecoverPoint/SE CRR standby Data Movers into standby mode. Write-disables destination-site LUNs from the Data Movers. Synchronizes destination to source. Resumes replicating the source devices. Reboots each Data Mover on the source VNX and reacquires the IP addresses and file system control from the RecoverPoint/SE CRR standby Data Movers.
Manage the RecoverPoint Appliance	From the destination VNX system, using the remote administration account: <code>nas_rp -rpa</code> <code>-list</code> <code> -add <rpaname> -ip <ip> -admin <rpaadminname> [-password <password>]</code> <code> -update {<rpaname> id=<id>}</code> <code> -info {<rpaname> id=<id>} [-version] [-verbose]</code>	<ul style="list-style-type: none"> ◆ Displays the properties of a specified RPA and the license settings. ◆ Manages the RPA (adds RPA entries to the NAS database, repairs all SSH RSA key issues, or lists all the RPAs configured in the system).

Table 6. RecoverPoint/SE CRR replication task overview *(continued)*

Task summary	Command used	What it does
Manage Consistency groups	From the VNX for File, using the remote administration account: nas_rp -cg -list -info {<cgname> id=<id>} -suspend {<cgname> id=<id>} -resume {<cgname> id=<id>} -modify {<cgname> id=<id>} -rpo <time_in_seconds>	<ul style="list-style-type: none"> ◆ Displays the properties of a specified NAS consistency group. ◆ Manages the NAS consistency group (lists all the consistency groups discovered in the RPA, suspends/resumes data transfer for a specified NAS consistency group, and modifies the RPO settings on a specified NAS consistency group to a user-specified value).

Sample configuration

The RecoverPoint/SE CRR replication active/passive and active/active' configuration tasks described in this document involve the following:

`new_york`, which serves as the source VNX in the active/passive configuration and one source/destination in the active/active' configuration. `new_york` is located at the production data center. The IP address of the source Control Station in slot 0 (CS0) is 192.168.96.85.

`new_jersey`, which serves as the destination VNX in the active/passive configuration and one source/destination VNX for File in the active/active' configuration. `new_jersey` is located at a remote disaster-recovery data center. The IP address of the destination Control Station in slot 0 (CS0) is 192.168.96.87.

The number of replication sets and file system output is unique to each configuration.

Configuring VNX backends

To configure your VNX backend systems for disaster recovery with RecoverPoint, you can perform the following tasks by using the Navisphere command line interface (CLI) and the Unisphere software:

- ◆ [Implement RecoverPoint CRR by using Deployment Manager on page 58](#)
- ◆ [Configure connection between RPAs and SPs on page 58](#)
- ◆ [Configure the source VNX on page 59](#)
- ◆ [Configure the destination VNX on page 61](#)
- ◆ [Configure the consistency group with the production and remote site on page 66](#)

Implement RecoverPoint CRR by using Deployment Manager

After identifying the source and destination VNX systems, you must implement the Continuous Remote Replication (CRR) option of RecoverPoint by using Deployment Manager. Using the CRR option, RecoverPoint replicates over a WAN (IP) or Fibre Channel (FC) to another storage array at a remote site. There is no distance limit for replication over IP. Deploying RecoverPoint requires installing and configuring a number of components. This includes installing RecoverPoint Appliances (RPAs), RecoverPoint software, and RecoverPoint splitter software.

Configure the VNX for Block SPs (both A and B) with RecoverPoint splitters on both sites. Using the Unisphere software, create a RecoverPoint storage group on each VNX for Block, which includes:

- ◆ All RPAs
- ◆ RecoverPoint replication volumes—LUNs that belong to the storage group for VNX for File

You can use the Deployment Manager to create repository volumes at both the source and destination sites. The repository volumes:

- ◆ Must reside in a storage environment that guarantees high availability.
- ◆ Must be provisioned on a non-thin LUN.
- ◆ Must be accessible to all RPAs at the site.
- ◆ Should be a maximum of 2.86 GB.

The repository volume contains configuration data that is created in the replication process. Any post-installation changes in the repository volume will result in a full sweep of all consistency groups and require a new activation license.

EMC RecoverPoint Deployment Manager provides detailed information on implementing the RecoverPoint CRR configuration.

Configure connection between RPAs and SPs

For each disk, RecoverPoint/SE requires two SCSI paths from each SP to the RPA. Two ports must be available on each SP other than port 0. Port 0 is reserved for MirrorView/S. Zone

these ports to the RPAs. EMC RecoverPoint Deployment Manager provides more information on zoning. [Figure 5 on page 59](#) illustrates the connections between RPAs and SPs.

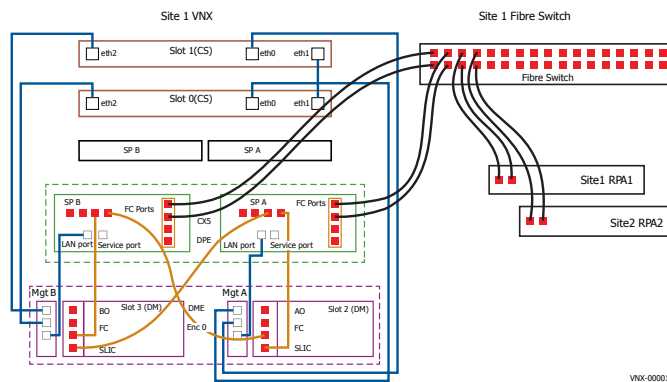


Figure 5. Connection between RPAs and SPs

Configure the source VNX

To prepare the source VNX for RecoverPoint configuration, the first step is to identify the VNX for File LUNs that you want to replicate. The ALU IDs associated with the control LUNs are retrieved as hexadecimal numbers. You must convert them to decimals to match them with the LUNs on the VNX for Block platform. After you identify the LUNs, you can add them to the RPA storage group at the source site and then identify the source servers.

To configure the source VNX, do the following:

- ◆ [Identify VNX for File LUNs on page 59](#)
- ◆ [Add LUNs to the RPA storage group at the source site](#)
- ◆ [Verify LUNs in the RPA storage group at the source site on page 61](#)
- ◆ [Identify source servers on page 61](#)

Identify VNX for File LUNs

To identify VNX for File LUNs, do the following:

- ◆ [Identify control LUNs on page 59](#)
- ◆ [Identify LUNs that are in use on page 60](#)
- ◆ [Identify LUNs that are not in use on page 60](#)

Identify control LUNs

For identifying control LUNs:

1. Ensure that you include VNX for File control LUNs in the RecoverPoint consistency group.
2. Run the `nas_storage -info` command.
3. From the VNX ~filesystem storage group details, select the LUNs that are mapped to the HLU 0,1, and 4.
4. Run "`nas_disk -info id={0|1|4}`" to get information about the ALUs associated with the control LUNs.
5. Convert these ALU IDs to decimal to match them with the LUNs on the VNX for Block platform.

Identify LUNs that are in use

For LUNs that are currently providing storage to existing file systems:

1. Select the file systems that you want to replicate.
2. For each file system, identify the disks that are provisioning the storage by typing the following command syntax:

```
$ nas_fs -query:Name==<FSNAME> -fields:Disks -format:"%q" -query:*
-fields:Id,Name,SymmDev -format:"Id=%s, Name=%s ALU=%s\n"
```

where:

<FSNAME> = name of the file system

Example:

```
$ nas_fs -query:Name==fs_test -fields:Disks -format:"%q" -query:*
-fields:Id,Name,SymmDev -format:"Id=%s, Name=%s ALU=%s\n"
```

Output:

```
Id=9, Name=d9 ALU=0039
```

3. Convert the ALU IDs to decimal to match them with the LUNs on the VNX for Block platform

Example:

An ALU ID of 0039 would translate to a decimal value of 57 ($16*3 + 9$)

Identify LUNs that are not in use

Identify the LUNs that are not in use by typing:

```
$ nas_disk -query:InUse=="n" -fields:Id,Name,SymmDev -format:"Id=%s,
Name=%s ALU=%s\n"
```

Output:

```

Id=7, Name=d7 ALU=0029
Id=8, Name=d8 ALU=002B
Id=10, Name=d10 ALU=0006
Id=11, Name=d11 ALU=0028
Id=12, Name=d12 ALU=002A

```

Verify LUNs in the RPA storage group at the source site

To list the LUNs that reside in the RPA storage group at the source site, type the following navisecli command:

```
$ /nas/sbin/navisecli -h A_FNM00103400378 -user nasadmin -password nasadmin -scope
0 storagegroup -messner -list
```

Output:

```
Storage Group Name:    site1_rpa_sg
Storage Group UID:    AA:19:A7:53:37:D3:DF:11:BF:26:00:60:16:36:E0:1E
```

HLU/ALU Pairs:

HLU Number	ALU Number
-----	-----
0	1018
...
4	4095
...
Shareable:	YES

By default, all the LUNs that exist on the VNX for Bock at the time of installing RecoverPoint/SE are added to the RPA storage group on each site. All other LUNs that are created after installing RecoverPoint/SE must be added exclusively to each storage group.

Identify source servers

Identify the source servers based on whether you are building an active/passive (unidirectional) or active/active' (bidirectional) configuration, as follows:

- ◆ [Table 7 on page 74](#) shows how to identify source servers for an active/passive configuration.
- ◆ [Table 8 on page 109](#) shows how to identify source servers for an active/active' configuration.

Configure the destination VNX

To prepare the destination VNX for RecoverPoint configuration, the first step is to create LUNs to match the LUNs that are selected for replication on the source VNX. After you create the LUNs, you can add them to the VNX for File storage group and the RPA storage

group at the destination site, and then identify the destination servers to match the servers selected at the source site.

To configure the destination VNX, do the following:

- ♦ [Create LUNs to match the LUNS on the source VNX on page 62](#)
- ♦ [Add LUNs to the VNX for File storage group at the destination site on page 64](#)
- ♦ [Verify LUNs in the VNX for File storage group at the destination site on page 64](#)
- ♦ [Add LUNs to the RPA storage group at the destination site](#)
- ♦ [Verify LUNs in the RPA storage group at the destination site on page 65](#)
- ♦ [Identify destination servers on page 65](#)

Create LUNs to match the LUNS on the source VNX

Before you begin

The recommendations for pairing production and disaster recovery LUNs are:

- ♦ Each pair of production and disaster site remote LUNs in the RecoverPoint consistency group must have the same storage properties:
 - Must be Raid-group based with the same raid-group on the disaster recover site
 - Must be Unified Pool-based LUNs with the same properties on the disaster recovery site
- ♦ Configure the following scheme when the Production and disaster recovery site remote LUNs cannot have the same storage properties:
 - Match all the LUNs that can be replicated from a Pool (for example, type A) on the source site to the LUNs that belong to same Pool on the disaster recovery site (for example, type B, but need not be of type A)

Example:

LUNs 10,11, and 12 are replicable and belong to clar_r5_performance on the Production site. Try to match these LUNs to the disaster recovery site remote LUNs, which are all part of the same pool (for example, clar_r5_economy, but may not be same as clar_r5_performance).

Source Site LUN ID (Pool)	DR Site LUN ID (Pool)
Id=10 (clar_r5_performance)	Id=20 (clar_r5_economy)
Id=11 (clar_r5_performance)	Id=21 (clar_r5_economy)
Id=12 (clar_r5_performance)	Id=22 (clar_r5_economy)

These can be TLUs or DLUs or MIXED types too.

Source Site LUN ID (Pool)	DR Site LUN ID (Pool)
Id=10 (clar_r5_performance)	Id=20 (Thin Pool1)
Id=11 (clar_r5_performance)	Id=21 (Thin Pool1)
Id=12 (clar_r5_performance)	Id=22 (Thin Pool1)

Note: Ensure that the LUNs on the DR site are from the same pool.

The following configuration is an example of a misconfiguration, where the LUNs of the DR site belong to different local pools.

Source Site LUN ID (Pool)	DR Site LUN ID (Pool)
Id=10 (clar_r5_performance)	Id=20 (clar_r5_economy)
Id=11 (clar_r5_performance)	Id=21 (Thin Pool1)
Id=12 (clar_r5_performance)	Id=22 (Thin Pool1)

The DR initialization operation will not detect such a misconfiguration. The DR failover process will detect the file systems that are spanning across different disk types and keep them offline (unmounted). You will have to manually mount these file systems by using the force option. File system operations such as Extensions/Checkpoints are not allowed for these types of file systems.

Procedure

You can create equivalent destination LUNs of source LUNs by using either the Unisphere software or the Navisphere CLI.

Using Unisphere

To create equivalent destination LUNs for each source LUN by using Unisphere:

1. In the systems list, select the storage system for which you want to create LUNs.
2. Select **Storage ► LUNs**.
3. From the task list, under LUNs, select **Create LUN**.

While creating matching LUNs, follow the recommendations in this section.

4. In the **General** tab, under **Storage Pool Properties**, select **Pool** or **Raid Group**.
5. Select the **RAID Type**. Pools support RAID 1/0, RAID 5, and RAID 6 only.
6. Select the RAID group or pool as the **Storage Pool** for the new LUN.
7. Under **LUN Properties**, set the **User Capacity** for the new LUN.

8. Select the **LUN ID** for the newly created LUN.
9. Select the number of LUNs to be created.
10. Under **LUN Name**, do one of the following:
 - a. Select **Name** and type a name for the newly created LUN.
 - b. Select **Automatically assign LUN IDs as LUN Names** to automatically use LUN IDs as LUN names.
11. Click **Apply**. The LUNs are created.

Using Navisphere CLI

To create the equivalent destination LUN of a source LUN by using the Navisphere CLI, type:

```
$ /nas/opt/Navisphere/bin/naviseccli -h <IP_address_or_hostname_of_SP A> -user
<username> -password <password> -scope 0 bind <raid-type [lun-number]> -rg <rgID> -rc
<0|1> -wc <0|1> -sp <a|b> -sq <mb|gb> -cap <same-size-as-primary>
```

Example (equivalent destination control dos LUN ALU ID 17 the same size as source dos control LUN 0, and in RAID group 3):

```
$ /nas/opt/Navisphere/bin/naviseccli -h element-spa -user nasadmin -password nasadmin
-scope 0 bind r5 17 -rg 3 -rc 1 -wc 1 -sp a -sq gb -cap 11
```

Add LUNs to the VNX for File storage group at the destination site

To add the LUNs to the VNX for File storage group at the destination site, type the following navisecli command for each identified destination LUN:

```
$ /nas/opt/Navisphere/bin/naviseccli -h 10.245.189.49 -user admin -password admin
-scope 0 storagegroup -messner -addhlu -gname site2_file_sg -hlu 2 -alu 1523
```

Note: The remote copy luns created on the destination VNX to be mapped to the source control luns 0, 1, and 4 must also be mapped to HLUs 6, 7 and 9 respectively in the destination VNX ~filesystem storage group.

Verify LUNs in the VNX for File storage group at the destination site

To list the LUNs that reside in the VNX for File storage group at the destination site, type the following navisecli command:

```
$ /nas/sbin/naviseccli -h A_FNM00103400378 -user nasadmin -password nasadmin -scope
0 storagegroup -messner -list
```

Output:


```

Storage Group Name:   site2_file_sg
Storage Group UID:   AA:19:A7:53:37:D3:DF:11:BF:26:00:60:16:36:E0:1E

HLU/ALU Pairs:

  HLU Number      ALU Number
  -----
    6              1523
    ...           ....
    7              2134
    ...           ....
Shareable:         YES

```

Verify LUNs in the RPA storage group at the destination site

To list the LUNs that reside in the RPA storage group at the destination site, type the following navisecli command:

```
$ /nas/sbin/navisecli -h A_FNM00103400378 -user nasadmin -password nasadmin -scope
0 storagegroup -messner -list
```

Output:

```

Storage Group Name:   site2_rpa_sg
Storage Group UID:   AA:19:A7:53:37:D3:DF:11:BF:26:00:60:16:36:E0:1E

HLU/ALU Pairs:

  HLU Number      ALU Number
  -----
    6              1523
    ...           ....
    7              2134
    ...           ....
Shareable:         YES

```

When you create new LUNs manually, you must add them exclusively into storage groups.

Identify destination servers

Identify the destination servers based on whether you are building an active/passive (unidirectional) or active/active' (bidirectional) configuration, as follows:

- ◆ [Table 7 on page 74](#) shows how to identify source servers for an active/passive configuration.
- ◆ [Table 8 on page 109](#) shows how to identify source servers for an active/active' configuration.

Configure the consistency group with the production and remote site

To set up a RecoverPoint consistency group:

1. On the main Unisphere page, from the **Replicas** ► **RecoverPoint** task list, click **Create CG Wizard**.
2. In the **Policy Configuration** dialog box, type the name of the consistency group, select the primary RPA to be associated with it, and click **Next**.
3. Select the **Production Site**.
4. Type the name of the Production copy.

Note: The *RecoverPoint Administrator's Guide* provides detailed information on configuring production policy options.

5. Enable the creation of a remote copy at the remote site.

Note:

- After you select the production site, the remote site is automatically selected.
 - Do not select the Create Local Copy option because it is not supported in this solution.
-

6. Type the name of the remote copy.

Note: The *RecoverPoint Administrator's Guide* provides detailed information on configuring remote copy policy options.

7. Select the LUNs at the production site that you want to add to a replication set from the list that appears.

Note: Ensure that each replication set consists of three control LUNs and the remaining user LUNs.

8. Select a LUN at the disaster recovery site to pair with each LUN selected at the production site.



CAUTION The control LUNs in the production RPA storage group must be mapped to the newly created remote control LUNs added to the remote RPA storage group. Do not map the control LUNs from the production RPA storage group to the local control LUNs of the remote VNX. This will result in data unavailability, data loss, or both.

9. Click **Next**. The **Journal Provisioning** dialog box appears.
10. To provision Journal Volumes:

- Accept the Journal configured automatically.
Or
- Select **Manually Select Journal Volumes**, click **Next** and manually configure Journal Volumes.

Note: Configure one Journal Volume for the production site and one for the remote site.

11. Click **Next** to review the Replication Set configuration. A summary of the consistency group appears.
12. To begin synchronization between the production and remote sites, select **Start data transfer immediately**.
13. Click **Finish** to complete the consistency group configuration.

Configuring RecoverPoint (active/passive)

The tasks to configure active/passive RecoverPoint/SE are:

- ◆ [Preinitialize the configuration on page 70](#)
- ◆ [Initialize the configuration \(active/passive\) on page 73](#)
- ◆ [Failover the source system \(active/passive\) on page 88](#)
- ◆ [Failback the source system \(active/passive\) on page 100](#)

Preinitialize the configuration

Preinitializing the configuration involves establishing a trusted communication between the source and destination systems and setting up the RPA. The preinitialization tasks are performed by your local EMC Service Provider.

Before you begin

- ◆ Preinitialization is a prerequisite to the disaster recovery initialization.
- ◆ Preinitialization must be performed on the source and destination systems.
- ◆ The system times of the source and destination Control Stations must be within 10 minutes of each other.
- ◆ Preinitialization must be performed by using the same 6–15 character passphrase for both the systems. For example, `nasadmin`.
- ◆ To preinitialize, the user must log in to the system as `nasadmin`.
- ◆ The `nas_cel` command must be run with the `-create` option on both the systems.
- ◆ The RPA object with its site management IP address and login information must be added to the NAS database.
- ◆ The preinitialization tasks are performed only once, after which the systems become ready for the VNX for File cabinet DR initialization procedures.
- ◆ The systems can be set up and made ready for production prior to initialization.

Procedure

The tasks to preinitialize the configuration are:

- ◆ [Preinitialize from the first \(source\) system on page 70](#)
- ◆ [Preinitialize from the second \(destination\) system on page 71](#)
- ◆ [Verify the preinitialization on page 73](#)

Before performing any of these tasks, review the requirements summarized in [Planning considerations on page 45](#) to ensure that the VNX for Block system and the VNX for File Data Movers are configured correctly. Your local EMC Service Provider establishes and verifies the initial RecoverPoint/SE CRR configuration, including preinitialization, initialization, test failover, and a test failback, and also manages VNX for Block configuration changes.

Preinitialize from the first (source) system

Step	Action
1.	Log in to the source system (<code>new_york</code>) as <code>nasadmin</code> .

Step	Action
2.	<p>Preinitialize the connection from the source system to the destination system by using this command syntax:</p> <pre>\$ nas_cel -create <cel_name> -ip <ip> -passphrase <passphrase></pre> <p>where:</p> <p><cel_name> = name of the destination system</p> <p><ip> = IP address of the destination Control Station in slot 0 (CS0)</p> <p><passphrase> = 6-15 character password</p> <p>Example:</p> <p>To preinitialize the connection from new_york to new_jersey with IP address 192.168.96.87 and passphrase nasadmin, type:</p> <pre>\$ nas_cel -create new_jersey -ip 192.168.96.87 -passphrase nasadmin</pre> <p>Output:</p> <pre>operation in progress (not interruptible)... id = 1 name = new_jersey owner = 0 device = channel = net_path = 192.168.96.87 celerra_id = APM000420008170000 passphrase = nasadmin</pre>

Preinitialize from the second (destination) system

Step	Action
1.	Log in to the destination system (new_jersey) as nasadmin.

Step	Action
2.	<p>Preinitialize the connection from the destination system to the source system by using this command syntax:</p> <pre>\$ nas_cel -create <cel_name> -ip <ip> -passphrase <passphrase></pre> <p>where:</p> <p><cel_name> = name of the source system</p> <p><ip> = IP address of the source Control Station in slot 0 (CS0)</p> <p><passphrase> = 6-15 character password</p> <p>Example:</p> <p>To preinitialize the connection from new_jersey to new_york with IP address 192.168.96.85 and passphrase nasadmin, type:</p> <pre>\$ nas_cel -create new_york -ip 192.168.96.85 -passphrase nasadmin</pre> <p>Output:</p> <pre>operation in progress (not interruptible)... id = 2 name = new_york owner = 0 device = channel = net_path = 192.168.96.85 celerra_id = APM000417005490000 passphrase = nasadmin</pre>
3.	<p>Add the RPA information to the NAS database by using this command syntax:</p> <pre>\$ nas_rp -rpa -add <name> -local_ip <local_rpa_ip> -remote_ip <remote_rpa_ip> -admin <rpa_admin_name> [-password <password>]</pre> <p>where:</p> <p>name = name of the RPA to be added</p> <p>local_rpa_ip = site management IP of the local RPA to be added</p> <p>remote_rpa_ip = site management IP of the remote RPA to be added</p> <p>rpa_admin_name = administrator username</p> <p>password = password corresponding to the administrator username</p> <p>Example:</p> <p>To add the RPA information to the NAS database, type:</p> <pre>\$/nas/sbin/nas_rp -rpa -add rpa1 -local_ip 10.245.64.16 -remote_ip 10.245.64.21 -admin admin -password admin</pre> <p>Output:</p> <pre>done</pre>

Verify the preinitialization

Step	Action
1.	<p>Log in to the source system (new_york) as nasadmin.</p> <hr/> <p>Note: Root user privilege is not required to run the <code>nas_cel -list</code> (or <code>nas_cel -info</code>) command. However, you must be logged in as root to run the <code>nas_cel</code> command for create, modify, update, or delete operations.</p> <hr/>
2.	<p>Verify the preinitialization of new_york and new_jersey by typing:</p> <pre>\$ nas_cel -list</pre> <p>Output:</p> <pre>id name owner mount_dev channel net_path CMU 0 new_york 0 192.168.96.85 APM000417005490000 2 new_jersey 0 192.168.96.87 APM000420008170000</pre> <hr/> <p>Note: The ID is 0 for the system from which you run the command.</p> <hr/>

Initialize the configuration (active/passive)

Initializing the active/passive configuration prepares the designated destination system to provide full file system access and functionality in the event of a source-site failure. The initialization tasks are performed by your local EMC Service Provider from the destination system.

Before you begin

- ◆ The VNX for File operating environment must be installed on the source and destination systems.
- ◆ The RecoverPoint/SE CRR link must be operational between the source and destination sites.
- ◆ The requirements summarized in [Planning considerations on page 45](#) must be met to ensure that the VNX for Block and VNX for File systems are set up correctly.
- ◆ If there is a default local standby, evaluate which standby relationships are needed for the RecoverPoint/SE CRR configuration. For example, in a four-Data Mover configuration, where server_5 is the default local standby, remove the server_5 standby relationships on both systems so that server_5 is not a local standby for server_2 and server_3. To remove the standby relationship, use the `server_standby` command, as described in [Data Mover configuration checklist on page 49](#). If you need to change the Data Mover type from standby to regular, use the `server_setup` command.

- ◆ Before running the `/nas/sbin/nas_rp -cabinetdr -init` command to select an active Data Mover as a remote standby Data Mover, clean up any part of the configuration, such as the network configuration, which no longer applies to that Data Mover.
- ◆ If you have a dual Control Station environment, halt CS1 before the initialization. You can perform RecoverPoint/SE operations by using CS0 only. After halting CS1, verify the halt by typing the `/nas/sbin/getreason` command, and checking for the line "0 - slot_1 powered off" in the output.
- ◆ Ensure that you have the global VNX for Block account password established. This account must be associated with the VNX for Block Manager or higher privileges to manage all storage system settings in the domain.

Note: If the existing global VNX for Block account password that you specify as part of the initialization procedure changes, but the VNX for File storage information has not been updated, you get an error. Step 4 of [Initialize from the destination system \(active/passive\) on page 75](#) shows an example of setting the password. After initialization is complete, you can capture the change by rerunning the `/nas/sbin/nas_rp -cabinetdr -init` command. However, after a failover, you must use the `nas_storage` command with the `modify` option to update the VNX for File storage security information on the Control Station. [Modify VNX for Block security information after a failover on page 162](#) provides more information.

- ◆ To avoid any path errors, always log in as `nasadmin`, switch (`su`) to root, and then run the `nas_rp -cabinetdr -init` command from `/nas/sbin`.

Note: When using `su`, ensure that you follow the steps in the procedure exactly. Do not use `su` unless explicitly instructed.

- ◆ Initialization fails if errors such as missing RecoverPoint/SE consistency group configuration, missing Control LUNS in the consistency group, and mismatched source and remote site LUN sizes occur. [Resolve initialization failures on page 172](#) describes failure scenarios associated with initialization.
- ◆ Verify the active/passive Data Mover configuration. [Table 7 on page 74](#) shows the Data Mover configuration used in this section. The Data Movers configured for RecoverPoint/SE are highlighted.

Table 7. Sample active/passive Data Mover configuration

Source site Data Mover (new_york)	Direction of DR relationship	Destination site Data Mover (new_jersey)
server_2 (source Data Mover)	----->	server_2 (remote standby)
server_3 (local standby for server_2)	----->	server_3 (remote standby)
server_4 (local Data Mover)	No relationship	server_4 (source Data Mover)

Table 7. Sample active/passive Data Mover configuration (continued)

Source site Data Mover (new_york)	Direction of DR relationship	Destination site Data Mover (new_jersey)
server_5 (local standby for server_4)	No relationship	server_5 (local standby for server_4)

In this configuration:

- new_jersey is the destination site system from which you run the commands to initialize the active/passive configuration. new_york is the source (active) system.
- Data Movers server_2 and server_3 on the source systemew_york are configured for remote disaster recovery with RecoverPoint/SE. Also, server_3 is a local standby for server_2 and server_5 is a local standby for server_4.
- Data Movers server_2 and server_3 on new_jersey are configured as remote standbys. server_2 for source server_2 and server_3 for source server_3. Also, server_4 is a local Data Mover and has a local standby of server_5.
- You can select the NAS LUNs for creating the RecoverPoint/SE consistency group between the new_york and new_jersey sites. [Configure the consistency group with the production and remote site on page 66](#) provides more details.

Procedure

The tasks to initialize the active/passive configuration are:

- ♦ [Initialize from the destination system \(active/passive\) on page 75](#)
- ♦ [Verify configuration \(active/passive\) on page 81](#)

Initialize from the destination system (active/passive)

Initializing the destination system establishes it to serve as the disaster recovery site.

Step	Action
1.	Log in to the destination system (new_jersey) as nasadmin and switch (su) to root.

Step	Action
2.	<p>Start the active/passive initialization process on the destination system by using this command syntax:</p> <pre># nas_rp -cabinetdr -init <cel_name></pre> <p>where:</p> <p><cel_name> = name of the source (remote) system as configured during the remote communication configuration</p> <p>Example:</p> <p>To initialize new_york as the source site for the destination site new_jersey, type:</p> <pre># /nas/sbin/nas_rp -cabinetdr -init new_york</pre> <p>Output:</p>

Step	Action
	<pre> Culham with RecoverPoint Disaster Recovery Initializing new_york --> new_jersey Contacting new_york for remote storage info Local storage system: FNM00093600019 Remote storage system: FNM00094700042 Discovering storage on new_york (may take several minutes) Setting security information for FNM00093600019 Discovering storage at 172.24.173.26 (may take several minutes) Contacting new_york for remote storage info Contacting new_york for server capabilities... Analyzing server information... Source servers available to be configured for remote DR ----- 1. server_2:new_york 2. server_3:new_york [local standby] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: 1 Destination servers available to act as remote standby ----- 1. server_2:new_jersey server_3:new_jersey [local standby] b. Back Select a new_jersey server: 1 Source servers available to be configured for remote DR ----- 1. server_2:new_york [remote standby is server_2:new_jersey] 2. server_3:new_york [local standby] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: c Standby configuration validated OK Enter user information for managing remote site new_york Username: dradmin Password: ***** Retype your response to validate Password: ***** Setting up server_2 on new_york Rebooting server_2 on new_jersey as standby ... done Setting acl for server_2 on new_jersey Updating the Culham domain information done </pre>

Step	Action
3.	<p>At the prompt, type the global account (nasadmin) information, such as username and password established during VNX for Block configuration.</p> <p>The account must be associated with the VNX for Block Manager or higher privileges.</p> <p>Example:</p> <pre> Enter the Global account information Username: nasadmin Password: ***** Retype your response to validate Password: ***** Discovering storage on new_york (may take several minutes) Setting security information for APM00042000817 Discovering storage APM00041700549 (may take several minutes) Discovering storage (may take several minutes) Contacting new_york for remote storage info Gathering server information... Contacting new_york for server capabilities... Analyzing server information... </pre>
4.	<p>At the prompt, configure each source Data Mover that is to be RecoverPoint/SE-protected with a remote standby Data Mover. To configure the Data Mover relationships, type the appropriate selection:</p> <ul style="list-style-type: none"> ◆ Type the selection number (not the server ID for a Data Mover) associated with the Data Mover to configure. For example, selection 1 for server_2. ◆ Type v to verify the current server configuration. You can verify the configuration after specifying each source Data Mover relationship. If there is an error, it is reported. ◆ Type q to quit the initialization process. ◆ Type c to continue with the initialization process after specifying the Data Mover relationships. ◆ Type d to display more information if the selection menu indicates that a Data Mover is ineligible for a remote disaster recovery configuration. In this case, the Data Mover is not selectable, and a not eligible for remote DR message appears within parentheses after the server name, as shown in Ensure Data Mover eligibility on page 151. ◆ If you are rerunning the initialization to change your Data Mover configuration, type r after specifying a selection number to remove the configuration of a destination Data Mover currently serving as a remote standby. Change the Data Mover configuration on page 159 contains more information and steps. ◆ Type b to return to the previous selection screen after specifying a destination Data Mover. <hr/> <p>Note: Review Data Mover configuration checklist on page 49 to determine your Data Mover configuration. A destination site local standby cannot be paired with a source site local standby for disaster recovery. The initialization procedure validates and enforces standby Data Mover compatibility and DR eligibility for the Data Movers. Ensure Data Mover eligibility on page 151 provides more information.</p> <hr/>

Step	Action
	<pre> Example: Source servers available to be configured for remote DR ----- 1. server_2:new_york 2. server_3:new_york [local standby] 3. server_4:new_york 4. server_5:new_york [local standby] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: 1 Destination servers available to act as remote standby ----- 1. server_2:new_jersey 2. server_3:new_jersey 3. server_4:new_jersey server_5:new_jersey [local standby] b. Back Select a new_jersey server: 1 Source servers available to be configured for remote DR ----- 1. server_2:new_york [remote standby is server_2:new_jersey] 2. server_3:new_york [local standby] 3. server_4:new_york 4. server_5:new_york [local standby] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: 2 Destination servers available to act as remote standby ----- server_2:new_jersey [is remote standby for server_2:new_york] 2. server_3:new_jersey 3. server_4:new_jersey server_5:new_jersey [local standby] b. Back Select a new_jersey server: 2 Source servers available to be configured for remote DR ----- 1. server_2:new_york [remote standby is server_2:new_jersey] 2. server_3:new_york [remote standby is server_3:new_jersey] 3. server_4:new_jersey 4. server_5:new_jersey [local standby] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: c Standby configuration validated OK </pre>

Step	Action
5.	<p>At the prompt, create the remote administration account (dradmin) to manage the remote (new_york) site.</p> <p>Example:</p> <pre>Enter user information for managing remote site new_york Username: dradmin Password: ***** Retype your response to validate Password: *****</pre> <p>Note: Remember the username and password. You must log in using this information when you activate a failover. This is a Linux-governed password on the Control Station and is covered in the Linux man pages for account passwords, such as man passwd.</p>
6.	<p>At the prompt, continue with the initialization by typing yes.</p> <p>Note: If you type no, the initialization aborts with an informational message. Note that both systems are informed about the consistency group creation.</p> <p>Example:</p> <pre>Initializing Active-->Passive (new_york-->new_jersey) Do you wish to continue? [yes or no] yes Updating RecoverPoint configuration cache Setting up server_3 on new_york Rebooting server_3 on new_jersey as standby ... done Setting up server_2 on new_york Rebooting server_2 on new_jersey as standby ... done Creating user account dradmin Setting acl for server_3 on new_jersey Setting acl for server_2 on new_jersey Updating the Celerra domain information Creating consistency group cg_new_york on new_jersey Creating consistency group cg_new_york on new_york done</pre>
7.	<p>Exit root by typing:</p> <pre># exit</pre> <p>Output</p> <pre>exit</pre> <p>Note: The initialization of the active/passive RecoverPoint/SE CRR configuration is complete.</p>

If the IP address of the Control Station changes after the initialization process runs, rerun the `/nas/sbin/nas_rp -cabinetdr -init` command to accommodate the change. If you change any hostnames or IP addresses and want to accommodate the change before you run the initialization process, edit and update the `/etc/hosts` file and ensure that each host can resolve its node name. If the storage configuration changes, which affects any of the RPA, storage

groups, or RP consistency groups used by the system, rerun the `/nas/sbin/nas_rp -cabinetdr -init` command.

Verify configuration (active/passive)

After initialization, you can verify elements of the active/passive RecoverPoint/SE CRR configuration by using the `nas_server -list`, `nas_server -info -all`, `nas_rp_cg -list`, `nas_rp_cg -info <name>`, `nas/sbin/nas_rp -cabinetdr -info`, `server_df ALL`, `nas_fs -info`, and `nas_disk -list` commands.

Step	Action
1.	Log in to the source system (new_york) as nasadmin.
2.	<p>List all the Data Movers by typing:</p> <pre>\$ /nas/bin/nas_server -list</pre> <p>Output</p> <pre>id type acl slot groupID state name 1 1 1000 2 0 server_2 2 4 1000 3 0 server_3 3 1 1000 4 0 server_4 4 4 1000 5 0 server_5</pre> <p>Note: Data Movers server_3 and server_5 have type 4 to identify them as standbys. server_3 is a local standby for server_2 and both have remote RecoverPoint/SE CRR standbys. server_5 is a local standby for server_4. If you run the command from the destination side as nasadmin, only server_4 and server_5 would appear in the list because server_2 and server_3 are managed by the dradmin account.</p>
3.	<p>Switch (su) to root by typing:</p> <pre>\$ su</pre> <p>Password:</p>

Step	Action
4.	<p>List all information for all Data Movers by typing:</p> <pre>\$ nas_server -info -all</pre> <p>Output:</p> <pre>id = 1 name = server_2 acl = 1000, owner=nasadmin, ID=201 type = nas slot = 2 member_of = standby = server_3, policy>manual RDFstandby= slot=2 status : defined = enabled actual = online, active id = 2 name = server_3 acl = 1000, owner=nasadmin, ID=201 type = standby slot = 3 member_of = standbyfor= server_2 RDFstandby= slot=3 status : defined = enabled actual = online, ready id = 3 name = server_4 acl = 1000, owner=nasadmin, ID=201 type = nas slot = 4 member_of = standby = server_5, policy=auto status : defined = enabled actual = online, active id = 4 name = server_5 acl = 1000, owner=nasadmin, ID=201 type = standby slot = 5 member_of = standbyfor= server_4 status : defined = enabled actual = online, ready</pre>

Step	Action																		
	<p>Note: When run from the source side, the output identifies Data Movers that have remote standbys (in the RDFstandby= field) and local standbys (standby=), and also indicates a Data Mover serving as a local standby (standbyfor=). When run from the destination, the output indicates which Data Movers are owned by dradmin and serve as remote standbys (in the acl= field and the type= field), and also indicates if a Data Mover serves as a local standby (standbyfor=). The standbyfor= has a value only if the Data Mover serves as a local standby.</p>																		
5.	<p>Exit root by typing:</p> <pre># exit</pre> <p>Output:</p> <pre>exit</pre>																		
6.	<p>Verify the creation of the consistency group on the source system after initialization by typing:</p> <pre>\$ nas_rp -cg -list</pre> <p>Output:</p> <table border="1"> <thead> <tr> <th>ID</th> <th>Name</th> <th>RPA ID</th> <th>Prod Copy</th> <th>Remote Copy</th> <th>Control</th> </tr> <tr> <th>LUN</th> <th>CG</th> <th></th> <th></th> <th></th> <th></th> </tr> </thead> <tbody> <tr> <td>1</td> <td>cg_new_york</td> <td>4</td> <td>Src_dev10</td> <td>DR_dev13</td> <td>True</td> </tr> </tbody> </table> <p>Note: cg_new_york represents the RecoverPoint/SE-protected data LUNs in the consistency group. The group in your configuration might have a different name, which is assigned when the consistency group is created as part of the storage system configuration process. Activate a failover from the destination VNX (active/passive) on page 89 shows an example of detailed consistency group information prior to failover activation.</p>	ID	Name	RPA ID	Prod Copy	Remote Copy	Control	LUN	CG					1	cg_new_york	4	Src_dev10	DR_dev13	True
ID	Name	RPA ID	Prod Copy	Remote Copy	Control														
LUN	CG																		
1	cg_new_york	4	Src_dev10	DR_dev13	True														

Step	Action
7.	<p>Get detailed information about the consistency group on the source VNX by using this command syntax:</p> <pre># nas_rp -cg -info {<name> id=<id>}</pre> <p>where:</p> <p><name> = name of the consistency group</p> <p><id> = ID assigned to the consistency group</p> <p>Example</p> <p>To get detailed information about the consistency group on the source system, type:</p> <pre>\$ nas_rp -cg -info cg_new_york</pre> <p>Output:</p> <pre>id = 4 name = cg_new_york rpa = rpa1 source copy = Src_dev10 remote copy = DR_dev13 source clar id = APM00102102333 remote clar id = APM00102400657 contains control luns = True transfer state = ACTIVE replication direction = local -> remote role = Primary transfer mode = Async rpo = SYSTEM Replication sets Id Name Src LUN Dst LUN Size 70 RSet 1 4 25 2147483648 71 RSet 10 119 119 214748364800 72 RSet 11 120 120 214748364800 73 RSet 12 101 101 214748364800 74 RSet 13 102 102 214748364800 75 RSet 14 103 103 214748364800 76 RSet 15 104 104 214748364800 77 RSet 16 105 105 214748364800 78 RSet 17 106 106 214748364800 79 RSet 18 107 107 214748364800 80 RSet 19 108 108 214748364800 81 RSet 2 0 18 11811160064 82 RSet 20 109 109 214748364800 83 RSet 21 110 110 214748364800 84 RSet 22 111 111 214748364800 85 RSet 23 112 112 214748364800 86 RSet 3 1 19 11811160064 87 RSet 4 113 113 214748364800 88 RSet 5 114 114 214748364800 89 RSet 6 115 115 214748364800 90 RSet 7 116 116 214748364800 91 RSet 8 117 117 214748364800 92 RSet 9 118 118 214748364800</pre>

Step	Action
8.	Switch (su) to root by typing: <pre>\$ su</pre> Password:
9.	Get information about the RecoverPoint/SE CRR replication configuration from the source VNX by typing: <pre># /nas/sbin/nas_rp -cabinetdr -info</pre> Output: <pre>***** Consistency Group Configuration ***** id = 4 name = cg_new_york rpa = rpa1 source copy = Src_dev10 remote copy = DR_dev13 source clar id = APM00102102333 remote clar id = APM00102400657 contains control luns = True ***** Servers configured with RPstandby ***** id = 1 name = server_2 acl = 0 type = nas slot = 2 member_of = standby = RDFstandby= slot=2 status : defined = enabled actual = online, active id = 2 name = server_3 acl = 0 type = nas slot = 3 member_of = standby = RDFstandby= slot=3 status : defined = enabled actual = online, ready ***** Servers configured as standby ***** No servers configured as standby</pre>

Step	Action
10.	<p>Exit root by typing:</p> <pre># exit</pre> <p>Output</p> <pre>exit</pre>
11.	<p>Display information about the file systems associated with Data Movers on the source VNX by typing:</p> <pre>\$ /nas/bin/server_df ALL</pre> <p>Output:</p> <pre>server_2 : Filesystem kbytes used avail capacity Mounted on ufs4 413030384 576 413029808 0% /ufs4 ufs3 413030384 576 413029808 0% /ufs3 ufs2 413030384 576 413029808 0% /ufs2 ufs1 413030384 576 413029808 0% /ufs1 root_fs_common 13624 5256 8368 39% /.etc_common root_fs_2 114592 728 113864 1% / server_3 : Error 2: server_3 : No such file or directory failed to complete command server_4 : Filesystem kbytes used avail capacity Mounted on root_fs_common 13624 5256 8368 39% /.etc_common ufslocal3 206515184 576 206514608 0% /ufslocal3 ufslocal2 413030384 576 413029808 0% /ufslocal2 ufslocal1 413030384 576 413029808 0% /ufslocal1 root_fs_4 114592 712 113880 1% / server_5 : Error 2: server_5 : No such file or directory failed to complete command</pre> <p>Note: No information is reported for the standby Data Movers as nasadmin. Information is only reported for the two source Data Movers; server_2 is RecoverPoint/SE-protected.</p>

Step	Action
12.	<p>Display information for one of the file systems from the source VNX by using this command syntax:</p> <pre>\$ nas_fs -info <fs_name></pre> <p>where:</p> <p><fs_name> = name of a specific file system</p> <p>Example:</p> <p>To view the information for replicated file system ufs1 on server_2 on new_york, type:</p> <pre>\$ nas_fs -info ufs1</pre> <p>Output:</p> <pre>id = 27 name = ufs1 acl = 0 in_use = True type = udfs worm = off volume = v202 pool = cm_r5_performance member_of = root_avm_fs_group_14 rw_servers = server_2 ro_servers = rw_vdms = ro_vdms = auto_ext = no,virtual_provision=no stor_devs = APM00041700549-0017,APM00041700549-0014,APM00041700549-0013 disks = d14,d9,d12 disk=d14 stor_dev=APM00041700549-0017 addr=c16t117 server=server_2 disk=d14 stor_dev=APM00041700549-0017 addr=c32t117 server=server_2 disk=d14 stor_dev=APM00041700549-0017 addr=c0t117 server=server_2 disk=d14 stor_dev=APM00041700549-0017 addr=c48t117 server=server_2 disk=d9 stor_dev=APM00041700549-0014 addr=c0t114 server=server_2 disk=d9 stor_dev=APM00041700549-0014 addr=c48t114 server=server_2 disk=d9 stor_dev=APM00041700549-0014 addr=c16t114 server=server_2 disk=d9 stor_dev=APM00041700549-0014 addr=c32t114 server=server_2 disk=d12 stor_dev=APM00041700549-0013 addr=c16t113 server=server_2 disk=d12 stor_dev=APM00041700549-0013 addr=c32t113 server=server_2 disk=d12 stor_dev=APM00041700549-0013 addr=c0t113 server=server_2 disk=d12 stor_dev=APM00041700549-0013 addr=c48t113 server=server_2</pre> <hr/> <p>Note: In the output, the pool reflects a RecoverPoint/SE pool type. In this example, cm_r5_performance.</p> <hr/>

Step	Action																																																																																																																																																			
13.	<p>List the disks on the source VNX and view the replicated and local-only disks by typing:</p> <pre>\$ nas_disk -list</pre> <p>Output:</p> <table border="1"> <thead> <tr> <th>id</th> <th>inuse</th> <th>sizeMB</th> <th>storageID-devID</th> <th>type</th> <th>name</th> <th>servers</th> </tr> </thead> <tbody> <tr><td>1</td><td>y</td><td>11263</td><td>APM00041700549-0000</td><td>CMSTD</td><td>root_disk</td><td>1,2,3,4</td></tr> <tr><td>2</td><td>y</td><td>11263</td><td>APM00041700549-0001</td><td>CMSTD</td><td>root_ldisk</td><td>1,2,3,4</td></tr> <tr><td>3</td><td>y</td><td>2047</td><td>APM00041700549-0002</td><td>CLSTD</td><td>d3</td><td>1,2,3,4</td></tr> <tr><td>4</td><td>y</td><td>2047</td><td>APM00041700549-0003</td><td>CLSTD</td><td>d4</td><td>1,2,3,4</td></tr> <tr><td>5</td><td>y</td><td>2047</td><td>APM00041700549-0004</td><td>CMSTD</td><td>d5</td><td>1,2,3,4</td></tr> <tr><td>6</td><td>y</td><td>2047</td><td>APM00041700549-0005</td><td>CLSTD</td><td>d6</td><td>1,2,3,4</td></tr> <tr><td>7</td><td>y</td><td>549623</td><td>APM00041700549-0010</td><td>CLSTD</td><td>d7</td><td>1,3,4,2</td></tr> <tr><td>8</td><td>y</td><td>549623</td><td>APM00041700549-0012</td><td>CMSTD</td><td>d8</td><td>1,3,4,2</td></tr> <tr><td>9</td><td>y</td><td>549623</td><td>APM00041700549-0014</td><td>CMSTD</td><td>d9</td><td>1,3,4,2</td></tr> <tr><td>10</td><td>y</td><td>549623</td><td>APM00041700549-0016</td><td>CMSTD</td><td>d10</td><td>1,3,4,2</td></tr> <tr><td>11</td><td>y</td><td>549623</td><td>APM00041700549-0011</td><td>CLSTD</td><td>d11</td><td>1,3,4,2</td></tr> <tr><td>12</td><td>y</td><td>549623</td><td>APM00041700549-0013</td><td>CMSTD</td><td>d12</td><td>1,3,4,2</td></tr> <tr><td>13</td><td>y</td><td>549623</td><td>APM00041700549-0015</td><td>CMSTD</td><td>d13</td><td>1,3,4,2</td></tr> <tr><td>14</td><td>y</td><td>549623</td><td>APM00041700549-0017</td><td>CMSTD</td><td>d14</td><td>1,3,4,2</td></tr> <tr><td>15</td><td>n</td><td>608270</td><td>APM00041700549-0018</td><td>CLATA</td><td>d15</td><td>2,1,3,4</td></tr> <tr><td>16</td><td>n</td><td>608270</td><td>APM00041700549-0019</td><td>CLATA</td><td>d16</td><td>2,1,3,4</td></tr> <tr><td>17</td><td>n</td><td>608270</td><td>APM00041700549-001C</td><td>CLATA</td><td>d17</td><td>2,1,3,4</td></tr> <tr><td>18</td><td>n</td><td>608270</td><td>APM00041700549-001D</td><td>CLATA</td><td>d18</td><td>2,1,3,4</td></tr> <tr><td>19</td><td>n</td><td>608270</td><td>APM00041700549-001A</td><td>CLATA</td><td>d19</td><td>2,1,3,4</td></tr> <tr><td>20</td><td>n</td><td>608270</td><td>APM00041700549-001B</td><td>CLATA</td><td>d20</td><td>2,1,3,4</td></tr> </tbody> </table> <p>Note: The output highlights the replicated disks associated with file system ufs1 in the pool cm_r5_perfor- mance. These disks have the disk type CMSTD.</p>	id	inuse	sizeMB	storageID-devID	type	name	servers	1	y	11263	APM00041700549-0000	CMSTD	root_disk	1,2,3,4	2	y	11263	APM00041700549-0001	CMSTD	root_ldisk	1,2,3,4	3	y	2047	APM00041700549-0002	CLSTD	d3	1,2,3,4	4	y	2047	APM00041700549-0003	CLSTD	d4	1,2,3,4	5	y	2047	APM00041700549-0004	CMSTD	d5	1,2,3,4	6	y	2047	APM00041700549-0005	CLSTD	d6	1,2,3,4	7	y	549623	APM00041700549-0010	CLSTD	d7	1,3,4,2	8	y	549623	APM00041700549-0012	CMSTD	d8	1,3,4,2	9	y	549623	APM00041700549-0014	CMSTD	d9	1,3,4,2	10	y	549623	APM00041700549-0016	CMSTD	d10	1,3,4,2	11	y	549623	APM00041700549-0011	CLSTD	d11	1,3,4,2	12	y	549623	APM00041700549-0013	CMSTD	d12	1,3,4,2	13	y	549623	APM00041700549-0015	CMSTD	d13	1,3,4,2	14	y	549623	APM00041700549-0017	CMSTD	d14	1,3,4,2	15	n	608270	APM00041700549-0018	CLATA	d15	2,1,3,4	16	n	608270	APM00041700549-0019	CLATA	d16	2,1,3,4	17	n	608270	APM00041700549-001C	CLATA	d17	2,1,3,4	18	n	608270	APM00041700549-001D	CLATA	d18	2,1,3,4	19	n	608270	APM00041700549-001A	CLATA	d19	2,1,3,4	20	n	608270	APM00041700549-001B	CLATA	d20	2,1,3,4
id	inuse	sizeMB	storageID-devID	type	name	servers																																																																																																																																														
1	y	11263	APM00041700549-0000	CMSTD	root_disk	1,2,3,4																																																																																																																																														
2	y	11263	APM00041700549-0001	CMSTD	root_ldisk	1,2,3,4																																																																																																																																														
3	y	2047	APM00041700549-0002	CLSTD	d3	1,2,3,4																																																																																																																																														
4	y	2047	APM00041700549-0003	CLSTD	d4	1,2,3,4																																																																																																																																														
5	y	2047	APM00041700549-0004	CMSTD	d5	1,2,3,4																																																																																																																																														
6	y	2047	APM00041700549-0005	CLSTD	d6	1,2,3,4																																																																																																																																														
7	y	549623	APM00041700549-0010	CLSTD	d7	1,3,4,2																																																																																																																																														
8	y	549623	APM00041700549-0012	CMSTD	d8	1,3,4,2																																																																																																																																														
9	y	549623	APM00041700549-0014	CMSTD	d9	1,3,4,2																																																																																																																																														
10	y	549623	APM00041700549-0016	CMSTD	d10	1,3,4,2																																																																																																																																														
11	y	549623	APM00041700549-0011	CLSTD	d11	1,3,4,2																																																																																																																																														
12	y	549623	APM00041700549-0013	CMSTD	d12	1,3,4,2																																																																																																																																														
13	y	549623	APM00041700549-0015	CMSTD	d13	1,3,4,2																																																																																																																																														
14	y	549623	APM00041700549-0017	CMSTD	d14	1,3,4,2																																																																																																																																														
15	n	608270	APM00041700549-0018	CLATA	d15	2,1,3,4																																																																																																																																														
16	n	608270	APM00041700549-0019	CLATA	d16	2,1,3,4																																																																																																																																														
17	n	608270	APM00041700549-001C	CLATA	d17	2,1,3,4																																																																																																																																														
18	n	608270	APM00041700549-001D	CLATA	d18	2,1,3,4																																																																																																																																														
19	n	608270	APM00041700549-001A	CLATA	d19	2,1,3,4																																																																																																																																														
20	n	608270	APM00041700549-001B	CLATA	d20	2,1,3,4																																																																																																																																														

Failover the source system (active/passive)

Activating a failover enables each RecoverPoint/SE standby Data Mover on the remote system to become active. After a successful failover, users have access to the same file systems by using the same network addresses as on the source system.

Before you begin

- After a failover has been activated, do not make any storage system configuration changes.
- Prior to a failover, storage system configuration changes affecting the initialized configuration can be captured by rerunning the `/nas/sbin/nas_rp -cabinetdr -init` command. A change to the existing global VNX for Block password requires an update of the VNX for File storage security information on the Control Station. After initialization is complete, this change can be captured by rerunning the `/nas/bin/nas_rp -cabinetdr -init` command. However, after a failover, you must use the `nas_storage` command with the `-modify` option. [Modify VNX for Block security information after a failover on page 162](#) provides more information.

Procedure

It is strongly recommended that your local EMC Service Provider performs a test failover and failback as part of the initial disaster recovery with the RecoverPoint/SE setup and verification process.

Note: The source will be shut down even if it is a test failover.

The tasks to activate an active/passive RecoverPoint/SE failover are:

- ◆ [Activate a failover from the destination VNX \(active/passive\) on page 89](#)
- ◆ [Verify active/passive operations after failover on page 94](#)
- ◆ [Ensure access after failover on page 99](#)

Activate a failover from the destination VNX (active/passive)

Failover is activated from the destination system Control Station by using the remote administration account (dradmin) that was created during initialization. In the sample scenario, the source system (new_york) is assumed to have become unavailable and requires a failover to the destination system.

To activate a test failover when no real disaster scenario exists, use the following guidelines, based on whether the goal is to perform a graceful failover or to simulate a true disaster scenario:

- ◆ To simply perform a graceful test failover, do the following before activating the failover:
 - Make sure the RecoverPoint/SE CRR links and the IP data network connections between Control Stations are up. You can also keep the source VNX for File powered up during the failover.

Note: The source VNX for File will be shut down at the end of the failover process.

- Check the consistency group information to verify that the group condition is in the appropriate state for the failover activation. The condition should be Active and the state should be Synchronized or Consistent.

Note: If you activate a test failover with the links down, or the consistency group is not in a proper state to activate a failover, a full synchronization is performed automatically as part of the failback to reconstruct the consistency group and replication sets, which is time-consuming and undesirable for a test scenario. If you activate the failover when the consistency group is not in the proper state to fail over, a warning message appears. In this case, abort the failover activation, make sure the consistency group is resynchronized, and then retry the failover activation.

- If you are without IP connectivity between the Control Stations prior to the failover activation, you can manually shut down the source Data Movers by using the `/nas/sbin/nas_halt now` command.

- After a test failover, you should reboot the source-site Control Station.
- You must initialize the configuration before the failover process.
- If any VNX for Block configuration changes have been made after initial configuration, these changes should be reflected by the VNX for File configuration prior to a failover. If the VNX for Block configuration for the NAS Storage Group has changed, the `/nas/sbin/nas_rp -cabinetdr -init` command must be rerun to update the VNX for File system before a failover is performed.
- When using `su`, make sure that you follow the steps in the procedure exactly. Do not use `su` unless explicitly instructed.

Note: For sites with redundant Control Stations, ensure that all RecoverPoint/SE CRR management commands, including `/nas/sbin/nas_rp -cabinetdr -init`, `-failover`, and `-failback`, are run from CS0. Ensure that CS1 is halted at both sites before you run the failover or failback commands. After halting CS1, verify the halt by typing the `/nas/sbin/getreason` command, and checking for the line "0 - slot_1 powered off" in the output.

Step	Action
1.	Log in to the destination system (<code>new_jersey</code>) as <code>nasadmin</code> .

Step	Action																																																																																																																													
2.	<p>List information for the consistency group and check the state of the consistency group prior to the failover by typing:</p> <pre>\$ nas_rp_cg -info cg_new_york</pre> <p>Output:</p> <pre>id = 1 name = cg_new_york rpa = rpa1 source copy = Src_dev10 remote copy = DR_dev13 source clar id = APM00102102333 remote clar id = APM00102400657 contains control luns = True transfer state = ACTIVE replication direction = local <- remote role = Primary transfer mode = Async rpo = SYSTEM</pre> <table border="1"> <thead> <tr> <th colspan="5">Replication sets</th> </tr> <tr> <th>Id</th> <th>Name</th> <th>Src LUN</th> <th>Dst LUN</th> <th>Size</th> </tr> </thead> <tbody> <tr><td>277</td><td>RSet 1</td><td>4</td><td>25</td><td>2147483648</td></tr> <tr><td>278</td><td>RSet 10</td><td>119</td><td>119</td><td>214748364800</td></tr> <tr><td>279</td><td>RSet 11</td><td>120</td><td>120</td><td>214748364800</td></tr> <tr><td>280</td><td>RSet 12</td><td>101</td><td>101</td><td>214748364800</td></tr> <tr><td>281</td><td>RSet 13</td><td>102</td><td>102</td><td>214748364800</td></tr> <tr><td>282</td><td>RSet 14</td><td>103</td><td>103</td><td>214748364800</td></tr> <tr><td>283</td><td>RSet 15</td><td>104</td><td>104</td><td>214748364800</td></tr> <tr><td>284</td><td>RSet 16</td><td>105</td><td>105</td><td>214748364800</td></tr> <tr><td>285</td><td>RSet 17</td><td>106</td><td>106</td><td>214748364800</td></tr> <tr><td>286</td><td>RSet 18</td><td>107</td><td>107</td><td>214748364800</td></tr> <tr><td>287</td><td>RSet 19</td><td>108</td><td>108</td><td>214748364800</td></tr> <tr><td>288</td><td>RSet 2</td><td>0</td><td>18</td><td>11811160064</td></tr> <tr><td>289</td><td>RSet 20</td><td>109</td><td>109</td><td>214748364800</td></tr> <tr><td>290</td><td>RSet 21</td><td>110</td><td>110</td><td>214748364800</td></tr> <tr><td>291</td><td>RSet 22</td><td>111</td><td>111</td><td>214748364800</td></tr> <tr><td>292</td><td>RSet 23</td><td>112</td><td>112</td><td>214748364800</td></tr> <tr><td>293</td><td>RSet 3</td><td>1</td><td>19</td><td>11811160064</td></tr> <tr><td>294</td><td>RSet 4</td><td>113</td><td>113</td><td>214748364800</td></tr> <tr><td>295</td><td>RSet 5</td><td>114</td><td>114</td><td>214748364800</td></tr> <tr><td>296</td><td>RSet 6</td><td>115</td><td>115</td><td>214748364800</td></tr> <tr><td>297</td><td>RSet 7</td><td>116</td><td>116</td><td>214748364800</td></tr> <tr><td>298</td><td>RSet 8</td><td>117</td><td>117</td><td>214748364800</td></tr> <tr><td>299</td><td>RSet 9</td><td>118</td><td>118</td><td>214748364800</td></tr> </tbody> </table> <p>As run from new_jersey in the active/passive configuration, the consistency group information reflects that this system is the secondary or destination system and the group represents a collection of replication sets. The replicated disks are not visible to hosts from the destination prior to a failover.</p>	Replication sets					Id	Name	Src LUN	Dst LUN	Size	277	RSet 1	4	25	2147483648	278	RSet 10	119	119	214748364800	279	RSet 11	120	120	214748364800	280	RSet 12	101	101	214748364800	281	RSet 13	102	102	214748364800	282	RSet 14	103	103	214748364800	283	RSet 15	104	104	214748364800	284	RSet 16	105	105	214748364800	285	RSet 17	106	106	214748364800	286	RSet 18	107	107	214748364800	287	RSet 19	108	108	214748364800	288	RSet 2	0	18	11811160064	289	RSet 20	109	109	214748364800	290	RSet 21	110	110	214748364800	291	RSet 22	111	111	214748364800	292	RSet 23	112	112	214748364800	293	RSet 3	1	19	11811160064	294	RSet 4	113	113	214748364800	295	RSet 5	114	114	214748364800	296	RSet 6	115	115	214748364800	297	RSet 7	116	116	214748364800	298	RSet 8	117	117	214748364800	299	RSet 9	118	118	214748364800
Replication sets																																																																																																																														
Id	Name	Src LUN	Dst LUN	Size																																																																																																																										
277	RSet 1	4	25	2147483648																																																																																																																										
278	RSet 10	119	119	214748364800																																																																																																																										
279	RSet 11	120	120	214748364800																																																																																																																										
280	RSet 12	101	101	214748364800																																																																																																																										
281	RSet 13	102	102	214748364800																																																																																																																										
282	RSet 14	103	103	214748364800																																																																																																																										
283	RSet 15	104	104	214748364800																																																																																																																										
284	RSet 16	105	105	214748364800																																																																																																																										
285	RSet 17	106	106	214748364800																																																																																																																										
286	RSet 18	107	107	214748364800																																																																																																																										
287	RSet 19	108	108	214748364800																																																																																																																										
288	RSet 2	0	18	11811160064																																																																																																																										
289	RSet 20	109	109	214748364800																																																																																																																										
290	RSet 21	110	110	214748364800																																																																																																																										
291	RSet 22	111	111	214748364800																																																																																																																										
292	RSet 23	112	112	214748364800																																																																																																																										
293	RSet 3	1	19	11811160064																																																																																																																										
294	RSet 4	113	113	214748364800																																																																																																																										
295	RSet 5	114	114	214748364800																																																																																																																										
296	RSet 6	115	115	214748364800																																																																																																																										
297	RSet 7	116	116	214748364800																																																																																																																										
298	RSet 8	117	117	214748364800																																																																																																																										
299	RSet 9	118	118	214748364800																																																																																																																										

Step	Action
3.	<p>Switch (su) to dradmin by typing:</p> <pre>\$ su - dradmin</pre> <p>Password: _____</p> <p>Note: The password is the same as specified during the initialization procedure, as shown in step 5 of Initialize from the destination system (active/passive) on page 75.</p> <p>_____</p>
4.	<p>Switch (su) to root by typing:</p> <pre>\$ su</pre> <p>Password: _____</p>
5.	<p>Failover from the source to the destination by typing:</p> <pre># /nas/sbin/nas_rp -cabinetdr -failover</pre> <p>Output:</p> <pre>Sync with CLARiiON backend done Validating consistency group configuration done</pre> <p>_____</p> <p>Note: As part of this process, all Data Movers at the source are halted. At the destination, do not shut down or reboot any Data Movers during a failover or a failback. Resolve failover failures on page 174 provides more information about errors that can occur during a failover.</p> <p>_____</p>

Step	Action
6.	<p>At the prompt, verify the readiness of the source site for shutdown and continue with the failover process by typing yes.</p> <p>If you type no, the failover aborts with an informational message.</p> <hr/> <p>Note: This process can take 10–20 minutes, depending on your configuration.</p> <hr/> <p>Example:</p> <pre> Is source site new_york ready for complete shut down (power OFF)? [yes or no] yes Contacting source site new_york, please wait... done Shutting down remote site new_york done </pre> <hr/> <p>Note: Next, the failover begins. A failback of the destination VNX for File write-enables the destination LUNs and write-disables the source LUNs, provided that the RecoverPoint/SE CRR link and the source VNX for Block are operational.</p> <hr/> <pre> Failing over ... consistency group : cg_new_york Failing over Devices ... done Adding NBS access to local server server_2 done Adding NBS access to local server server_3 done Adding NBS access to local server server_4 done Adding NBS access to local server server_5 done Activating the target environment ... done </pre> <hr/> <p>Note: The RecoverPoint/SE CRR standby Data Movers become active.</p> <hr/> <pre> server_2 : going offline rdf : going active replace in progress ...done failover activity complete server_3 : going offline rdf : going active replace in progress ...done failover activity complete commit in progress (not interruptible)...done commit in progress (not interruptible)...done commit in progress (not interruptible)...done commit in progress (not interruptible)...done done </pre> <hr/> <p>Note: The failover is complete.</p> <hr/>

Step	Action
7.	<p>Exit root by typing.</p> <pre># exit</pre> <pre>exit</pre> <hr/> <p>Note: After a failover has been activated, do not make any storage system configuration changes. In the event of a true disaster where the source VNX for Block is unavailable, contact your local EMC Service Provider or EMC Customer Service to coordinate fallback activities.</p> <hr/>

Verify active/passive operations after failover

Step	Action
1.	Log in to the destination system (new_jersey) as dradmin.
2.	<p>Display information about the Data Movers after the failover activation by typing:</p> <pre>\$ /nas/bin/nas_server -list</pre> <p>Output:</p> <pre>id type acl slot groupID state name 1 1 0 2 0 0 server_2 2 4 0 3 0 0 server_3</pre>
3.	<p>Display information about the activated RecoverPoint/SE standby Data Movers on the destination system by typing:</p> <pre>\$ /nas/bin/server_df ALL</pre> <p>Output:</p> <pre>server_2 : Filesystem kbytes used avail capacity Mounted on ufs4 413030384 576 413029808 0% /ufs4 ufs3 413030384 576 413029808 0% /ufs3 ufs2 413030384 576 413029808 0% /ufs2 ufs1 413030384 576 413029808 0% /ufs1 root_fs_common 13624 5256 8368 39% /.etc_common root_fs_2 114592 728 113864 1% / server_3 : Error 2: server_3 : No such file or directory failed to complete command</pre> <hr/> <p>Note: From dradmin, no information is reported for server_3 because it is a local standby.</p> <hr/>

Step	Action
4.	<p>List information for the consistency group after the failover by using this command syntax:</p> <pre># nas_rp -cg -info {<name> id=<id>}</pre> <p>where:</p> <p><name> = name of the consistency group</p> <p><id> = ID assigned to the consistency group (shown with nas_rp_cg -list)</p> <p>Example:</p> <p>To list information for the consistency group after the failover, type:</p> <pre>\$ nas_rp_cg -info cg_new_york</pre> <p>Output:</p> <pre>id = 1 name = cg_new_york rpa = rpa1 source copy = Src_dev10 remote copy = DR_dev13 source clar id = APM00102102333 remote clar id = APM00102400657 contains control luns = True transfer state = ACTIVE replication direction = local -> remote role = Primary transfer mode = Async rpo = SYSTEM Replication sets Id Name Src LUN Dst LUN Size 277 RSet 1 4 25 21474836480 278 RSet 10 119 119 214748364800 279 RSet 11 120 120 214748364800 280 RSet 12 101 101 214748364800 281 RSet 13 102 102 214748364800 282 RSet 14 103 103 214748364800 283 RSet 15 104 104 214748364800 284 RSet 16 105 105 214748364800 285 RSet 17 106 106 214748364800 286 RSet 18 107 107 214748364800 287 RSet 19 108 108 214748364800 288 RSet 2 0 18 11811160064 289 RSet 20 109 109 214748364800 290 RSet 21 110 110 214748364800 291 RSet 22 111 111 214748364800 292 RSet 23 112 112 214748364800 293 RSet 3 1 19 11811160064 294 RSet 4 113 113 214748364800 295 RSet 5 114 114 214748364800 296 RSet 6 115 115 214748364800 297 RSet 7 116 116 214748364800 298 RSet 8 117 117 214748364800 299 RSet 9 118 118 214748364800</pre>

Step	Action
5.	Switch (su) to root by typing: \$ su Password:

Step	Action
6.	<p>List cabinet-level RecoverPoint/SE information after the failover by typing:</p> <pre># /nas/sbin/nas_rp -cabinetdr -info</pre> <p>Output:</p> <pre>***** Consistency Group Configuration ***** id = 3 name = cg_new_york rpa = rpa1 source copy = Prod_98 remote copy = Remote_46 source clar id = FNM00093600019 remote clar id = FNM00094700042 contains control luns = True ***** Servers configured with RPstandby ***** id = 1 name = server_2 acl = 1000, owner=nasadmin, ID=201 type = nas slot = 2 member_of = standby = server_3, policy=auto RDFstandby= slot=2 status : defined = enabled actual = online, active id = 2 name = server_3 acl = 1000, owner=nasadmin, ID=201 type = standby slot = 3 member_of = standbyfor= server_2 RDFstandby= slot=3 status : defined = enabled actual = online, ready ***** Servers configured as standby ***** id = 4 name = server_5 acl = 1000, owner=nasadmin, ID=201 type = standby slot = 5 member_of = standbyfor= server_4 status : defined = enabled actual = online, ready</pre>

Step	Action																																																																																																																																																																																																																																
7.	<p>Get a list of the source site file systems available at the destination after the failover activation by typing:</p> <pre># nas_fs -list</pre> <p>Output:</p> <table border="1"> <thead> <tr> <th>id</th> <th>inuse</th> <th>type</th> <th>acl</th> <th>volume</th> <th>name</th> <th>server</th> </tr> </thead> <tbody> <tr><td>1</td><td>n</td><td>1</td><td>0</td><td>10</td><td>root_fs_1</td><td></td></tr> <tr><td>2</td><td>y</td><td>1</td><td>0</td><td>12</td><td>root_fs_2</td><td>1</td></tr> <tr><td>3</td><td>y</td><td>1</td><td>0</td><td>14</td><td>root_fs_3</td><td>2</td></tr> <tr><td>4</td><td>y</td><td>1</td><td>0</td><td>16</td><td>root_fs_4</td><td>3</td></tr> <tr><td>5</td><td>y</td><td>1</td><td>0</td><td>18</td><td>root_fs_5</td><td>4</td></tr> <tr><td>6</td><td>n</td><td>1</td><td>0</td><td>20</td><td>root_fs_6</td><td></td></tr> <tr><td>7</td><td>n</td><td>1</td><td>0</td><td>22</td><td>root_fs_7</td><td></td></tr> <tr><td>8</td><td>n</td><td>1</td><td>0</td><td>24</td><td>root_fs_8</td><td></td></tr> <tr><td>9</td><td>n</td><td>1</td><td>0</td><td>26</td><td>root_fs_9</td><td></td></tr> <tr><td>10</td><td>n</td><td>1</td><td>0</td><td>28</td><td>root_fs_10</td><td></td></tr> <tr><td>11</td><td>n</td><td>1</td><td>0</td><td>30</td><td>root_fs_11</td><td></td></tr> <tr><td>12</td><td>n</td><td>1</td><td>0</td><td>32</td><td>root_fs_12</td><td></td></tr> <tr><td>13</td><td>n</td><td>1</td><td>0</td><td>34</td><td>root_fs_13</td><td></td></tr> <tr><td>14</td><td>n</td><td>1</td><td>0</td><td>36</td><td>root_fs_14</td><td></td></tr> <tr><td>15</td><td>n</td><td>1</td><td>0</td><td>38</td><td>root_fs_15</td><td></td></tr> <tr><td>16</td><td>y</td><td>1</td><td>0</td><td>40</td><td>root_fs_common</td><td>4,2,3,1</td></tr> <tr><td>17</td><td>n</td><td>5</td><td>0</td><td>73</td><td>root_fs_ufslog</td><td></td></tr> <tr><td>18</td><td>n</td><td>5</td><td>0</td><td>76</td><td>root_panic_reserve</td><td></td></tr> <tr><td>19</td><td>n</td><td>5</td><td>0</td><td>77</td><td>root_fs_d3</td><td></td></tr> <tr><td>20</td><td>n</td><td>5</td><td>0</td><td>78</td><td>root_fs_d4</td><td></td></tr> <tr><td>21</td><td>n</td><td>5</td><td>0</td><td>79</td><td>root_fs_d5</td><td></td></tr> <tr><td>22</td><td>n</td><td>5</td><td>0</td><td>80</td><td>root_fs_d6</td><td></td></tr> <tr><td>27</td><td>y</td><td>1</td><td>0</td><td>202</td><td>ufs1</td><td>1</td></tr> <tr><td>29</td><td>y</td><td>1</td><td>0</td><td>205</td><td>ufs2</td><td>1</td></tr> <tr><td>30</td><td>y</td><td>1</td><td>0</td><td>207</td><td>ufs3</td><td>1</td></tr> <tr><td>31</td><td>y</td><td>1</td><td>0</td><td>209</td><td>ufs4</td><td>1</td></tr> <tr><td>32</td><td>n</td><td>1</td><td>0</td><td>211</td><td>ufs5</td><td></td></tr> <tr><td>33</td><td>n</td><td>1</td><td>0</td><td>213</td><td>ufs6</td><td></td></tr> <tr><td>34</td><td>y</td><td>1</td><td>0</td><td>217</td><td>ufslocal1</td><td>3</td></tr> <tr><td>36</td><td>y</td><td>1</td><td>0</td><td>220</td><td>ufslocal2</td><td>3</td></tr> <tr><td>40</td><td>y</td><td>1</td><td>0</td><td>230</td><td>ufslocal3</td><td>3</td></tr> </tbody> </table>	id	inuse	type	acl	volume	name	server	1	n	1	0	10	root_fs_1		2	y	1	0	12	root_fs_2	1	3	y	1	0	14	root_fs_3	2	4	y	1	0	16	root_fs_4	3	5	y	1	0	18	root_fs_5	4	6	n	1	0	20	root_fs_6		7	n	1	0	22	root_fs_7		8	n	1	0	24	root_fs_8		9	n	1	0	26	root_fs_9		10	n	1	0	28	root_fs_10		11	n	1	0	30	root_fs_11		12	n	1	0	32	root_fs_12		13	n	1	0	34	root_fs_13		14	n	1	0	36	root_fs_14		15	n	1	0	38	root_fs_15		16	y	1	0	40	root_fs_common	4,2,3,1	17	n	5	0	73	root_fs_ufslog		18	n	5	0	76	root_panic_reserve		19	n	5	0	77	root_fs_d3		20	n	5	0	78	root_fs_d4		21	n	5	0	79	root_fs_d5		22	n	5	0	80	root_fs_d6		27	y	1	0	202	ufs1	1	29	y	1	0	205	ufs2	1	30	y	1	0	207	ufs3	1	31	y	1	0	209	ufs4	1	32	n	1	0	211	ufs5		33	n	1	0	213	ufs6		34	y	1	0	217	ufslocal1	3	36	y	1	0	220	ufslocal2	3	40	y	1	0	230	ufslocal3	3
id	inuse	type	acl	volume	name	server																																																																																																																																																																																																																											
1	n	1	0	10	root_fs_1																																																																																																																																																																																																																												
2	y	1	0	12	root_fs_2	1																																																																																																																																																																																																																											
3	y	1	0	14	root_fs_3	2																																																																																																																																																																																																																											
4	y	1	0	16	root_fs_4	3																																																																																																																																																																																																																											
5	y	1	0	18	root_fs_5	4																																																																																																																																																																																																																											
6	n	1	0	20	root_fs_6																																																																																																																																																																																																																												
7	n	1	0	22	root_fs_7																																																																																																																																																																																																																												
8	n	1	0	24	root_fs_8																																																																																																																																																																																																																												
9	n	1	0	26	root_fs_9																																																																																																																																																																																																																												
10	n	1	0	28	root_fs_10																																																																																																																																																																																																																												
11	n	1	0	30	root_fs_11																																																																																																																																																																																																																												
12	n	1	0	32	root_fs_12																																																																																																																																																																																																																												
13	n	1	0	34	root_fs_13																																																																																																																																																																																																																												
14	n	1	0	36	root_fs_14																																																																																																																																																																																																																												
15	n	1	0	38	root_fs_15																																																																																																																																																																																																																												
16	y	1	0	40	root_fs_common	4,2,3,1																																																																																																																																																																																																																											
17	n	5	0	73	root_fs_ufslog																																																																																																																																																																																																																												
18	n	5	0	76	root_panic_reserve																																																																																																																																																																																																																												
19	n	5	0	77	root_fs_d3																																																																																																																																																																																																																												
20	n	5	0	78	root_fs_d4																																																																																																																																																																																																																												
21	n	5	0	79	root_fs_d5																																																																																																																																																																																																																												
22	n	5	0	80	root_fs_d6																																																																																																																																																																																																																												
27	y	1	0	202	ufs1	1																																																																																																																																																																																																																											
29	y	1	0	205	ufs2	1																																																																																																																																																																																																																											
30	y	1	0	207	ufs3	1																																																																																																																																																																																																																											
31	y	1	0	209	ufs4	1																																																																																																																																																																																																																											
32	n	1	0	211	ufs5																																																																																																																																																																																																																												
33	n	1	0	213	ufs6																																																																																																																																																																																																																												
34	y	1	0	217	ufslocal1	3																																																																																																																																																																																																																											
36	y	1	0	220	ufslocal2	3																																																																																																																																																																																																																											
40	y	1	0	230	ufslocal3	3																																																																																																																																																																																																																											

Step	Action																																																																																																																																																			
8.	<p>Get a list of the disks available after the failover activation by typing:</p> <pre># nas_disk -list</pre> <p>Output:</p> <table border="1"> <thead> <tr> <th>id</th> <th>inuse</th> <th>sizeMB</th> <th>storageID-devID</th> <th>type</th> <th>name</th> <th>servers</th> </tr> </thead> <tbody> <tr><td>1</td><td>y</td><td>11263</td><td>APM00042000817-0006</td><td>CMSTD</td><td>root_disk</td><td>1,2,3,4</td></tr> <tr><td>2</td><td>y</td><td>11263</td><td>APM00042000817-0007</td><td>CMSTD</td><td>root_ldisk</td><td>1,2,3,4</td></tr> <tr><td>3</td><td>y</td><td>2047</td><td>APM00041700549-0002</td><td>CLSTD</td><td>d3</td><td>1,2,3,4</td></tr> <tr><td>4</td><td>y</td><td>2047</td><td>APM00041700549-0003</td><td>CLSTD</td><td>d4</td><td>1,2,3,4</td></tr> <tr><td>5</td><td>y</td><td>2047</td><td>APM00042000817-0009</td><td>CMSTD</td><td>d5</td><td>1,2,3,4</td></tr> <tr><td>6</td><td>y</td><td>2047</td><td>APM00041700549-0005</td><td>CLSTD</td><td>d6</td><td>1,2,3,4</td></tr> <tr><td>7</td><td>y</td><td>549623</td><td>APM00041700549-0010</td><td>CLSTD</td><td>d7</td><td>1,3,4,2</td></tr> <tr><td>8</td><td>y</td><td>549623</td><td>APM00042000817-0012</td><td>CMSTD</td><td>d8</td><td>1,3,4,2</td></tr> <tr><td>9</td><td>y</td><td>549623</td><td>APM00042000817-0014</td><td>CMSTD</td><td>d9</td><td>1,3,4,2</td></tr> <tr><td>10</td><td>y</td><td>549623</td><td>APM00042000817-0016</td><td>CMSTD</td><td>d10</td><td>1,3,4,2</td></tr> <tr><td>11</td><td>y</td><td>549623</td><td>APM00041700549-0011</td><td>CLSTD</td><td>d11</td><td>1,3,4,2</td></tr> <tr><td>12</td><td>y</td><td>549623</td><td>APM00042000817-0013</td><td>CMSTD</td><td>d12</td><td>1,3,4,2</td></tr> <tr><td>13</td><td>y</td><td>549623</td><td>APM00042000817-0015</td><td>CMSTD</td><td>d13</td><td>1,3,4,2</td></tr> <tr><td>14</td><td>y</td><td>549623</td><td>APM00042000817-0017</td><td>CMSTD</td><td>d14</td><td>1,3,4,2</td></tr> <tr><td>15</td><td>n</td><td>608270</td><td>APM00041700549-0018</td><td>CLATA</td><td>d15</td><td>2,1,3,4</td></tr> <tr><td>16</td><td>n</td><td>608270</td><td>APM00041700549-0019</td><td>CLATA</td><td>d16</td><td>2,1,3,4</td></tr> <tr><td>17</td><td>n</td><td>608270</td><td>APM00041700549-001C</td><td>CLATA</td><td>d17</td><td>2,1,3,4</td></tr> <tr><td>18</td><td>n</td><td>608270</td><td>APM00041700549-001D</td><td>CLATA</td><td>d18</td><td>2,1,3,4</td></tr> <tr><td>19</td><td>n</td><td>608270</td><td>APM00041700549-001A</td><td>CLATA</td><td>d19</td><td>2,1,3,4</td></tr> <tr><td>20</td><td>n</td><td>608270</td><td>APM00041700549-001B</td><td>CLATA</td><td>d20</td><td>2,1,3,4</td></tr> </tbody> </table> <p>Note: The replicated disks have the disk type CMSTD.</p>	id	inuse	sizeMB	storageID-devID	type	name	servers	1	y	11263	APM00042000817-0006	CMSTD	root_disk	1,2,3,4	2	y	11263	APM00042000817-0007	CMSTD	root_ldisk	1,2,3,4	3	y	2047	APM00041700549-0002	CLSTD	d3	1,2,3,4	4	y	2047	APM00041700549-0003	CLSTD	d4	1,2,3,4	5	y	2047	APM00042000817-0009	CMSTD	d5	1,2,3,4	6	y	2047	APM00041700549-0005	CLSTD	d6	1,2,3,4	7	y	549623	APM00041700549-0010	CLSTD	d7	1,3,4,2	8	y	549623	APM00042000817-0012	CMSTD	d8	1,3,4,2	9	y	549623	APM00042000817-0014	CMSTD	d9	1,3,4,2	10	y	549623	APM00042000817-0016	CMSTD	d10	1,3,4,2	11	y	549623	APM00041700549-0011	CLSTD	d11	1,3,4,2	12	y	549623	APM00042000817-0013	CMSTD	d12	1,3,4,2	13	y	549623	APM00042000817-0015	CMSTD	d13	1,3,4,2	14	y	549623	APM00042000817-0017	CMSTD	d14	1,3,4,2	15	n	608270	APM00041700549-0018	CLATA	d15	2,1,3,4	16	n	608270	APM00041700549-0019	CLATA	d16	2,1,3,4	17	n	608270	APM00041700549-001C	CLATA	d17	2,1,3,4	18	n	608270	APM00041700549-001D	CLATA	d18	2,1,3,4	19	n	608270	APM00041700549-001A	CLATA	d19	2,1,3,4	20	n	608270	APM00041700549-001B	CLATA	d20	2,1,3,4
id	inuse	sizeMB	storageID-devID	type	name	servers																																																																																																																																														
1	y	11263	APM00042000817-0006	CMSTD	root_disk	1,2,3,4																																																																																																																																														
2	y	11263	APM00042000817-0007	CMSTD	root_ldisk	1,2,3,4																																																																																																																																														
3	y	2047	APM00041700549-0002	CLSTD	d3	1,2,3,4																																																																																																																																														
4	y	2047	APM00041700549-0003	CLSTD	d4	1,2,3,4																																																																																																																																														
5	y	2047	APM00042000817-0009	CMSTD	d5	1,2,3,4																																																																																																																																														
6	y	2047	APM00041700549-0005	CLSTD	d6	1,2,3,4																																																																																																																																														
7	y	549623	APM00041700549-0010	CLSTD	d7	1,3,4,2																																																																																																																																														
8	y	549623	APM00042000817-0012	CMSTD	d8	1,3,4,2																																																																																																																																														
9	y	549623	APM00042000817-0014	CMSTD	d9	1,3,4,2																																																																																																																																														
10	y	549623	APM00042000817-0016	CMSTD	d10	1,3,4,2																																																																																																																																														
11	y	549623	APM00041700549-0011	CLSTD	d11	1,3,4,2																																																																																																																																														
12	y	549623	APM00042000817-0013	CMSTD	d12	1,3,4,2																																																																																																																																														
13	y	549623	APM00042000817-0015	CMSTD	d13	1,3,4,2																																																																																																																																														
14	y	549623	APM00042000817-0017	CMSTD	d14	1,3,4,2																																																																																																																																														
15	n	608270	APM00041700549-0018	CLATA	d15	2,1,3,4																																																																																																																																														
16	n	608270	APM00041700549-0019	CLATA	d16	2,1,3,4																																																																																																																																														
17	n	608270	APM00041700549-001C	CLATA	d17	2,1,3,4																																																																																																																																														
18	n	608270	APM00041700549-001D	CLATA	d18	2,1,3,4																																																																																																																																														
19	n	608270	APM00041700549-001A	CLATA	d19	2,1,3,4																																																																																																																																														
20	n	608270	APM00041700549-001B	CLATA	d20	2,1,3,4																																																																																																																																														

Ensure access after failover

If you have not accounted for different IP subnets at the source and destination sites, perform these steps, either manually, or by creating and running a script, after RecoverPoint/SE failover to ensure access to the same file systems by using the same network addresses as on the source site, provided network access to the destination site system exists.

Step	Action
1.	Halt CIFS.
2.	Set IP addresses and default routes.
3.	Adjust services such as WINS, DNS, NIS, and NTP.
4.	Restart CIFS.

Failback the source system (active/passive)

Source system failback involves restoring full connectivity to the file systems on the source VNX.

Before you begin

Failback of the source system includes two phases:

- ◆ The storage failback phase provides network access to the destination VNX for File while the destination and source VNX for Block systems synchronize. The length of the synchronization is based on the amount of data that has been updated since the destination system experienced failover.
- ◆ The network failback phase suspends network clients from file system access. Note that with active write I/O it could take some time to perform the synchronization.

Procedure

It is strongly recommended that your local EMC Service Provider performs a test failover and failback as part of the initial disaster recovery with RecoverPoint/SE setup and verification process.

The tasks to failback the source system are:

- ◆ [Prepare for the failback \(active/passive\) on page 100](#)
- ◆ [Failback the source system from the destination system\(active/passive\) on page 101](#)
- ◆ [Verify operations after a failback \(active/passive\) on page 103](#)

Prepare for the failback (active/passive)

Step	Action
1.	<p>Request a complete system check of the VNX for Block system and RecoverPoint/SE. This is required to verify proper RecoverPoint/SE operations.</p> <hr/> <p>Important: If this is a true disaster scenario, keep the source VNX for File (new_york) and its VNX for Block powered off until you are instructed to power them up by your local EMC Service Provider or EMC Customer Service.</p> <hr/>
2.	<p>Power up the VNX for Block system after you have been told that it is safe to proceed, and then continue with the failback procedure on the destination VNX for File.</p> <hr/> <p>Important: Do not attempt to restart the source VNX for File (new_york).</p> <hr/>

Failback the source system from the destination system(active/passive)

Use the remote administration account (dradmin) on the destination system to failback the source VNX for File.

Step	Action
1.	Log in to the destination system (new_jersey) as dradmin and switch (su) to root.
2.	<p>Start the failback of the source VNX for File by typing:</p> <pre># /nas/sbin/nas_rp -cabinetdr -failback</pre> <p>_____</p> <p>Note: Use the absolute path. If you do not use the absolute path, the command fails with a general command error. If the command fails, rerun the command with the absolute path. Also, ensure that no other command is running and using dradmin when the /nas/sbin/nas_rp -cabinetdr -failback command is running. During the failback, do not shut down or reboot any Data Movers.</p> <p>_____</p>
3.	<p>At the prompt, proceed with the failback after confirmation from your local EMC Service Provider or EMC Customer Service by typing yes.</p> <p>This step validates the consistency group configuration, and then requests a shutdown of the source.</p> <p>If you type no, the failback aborts with an informational message.</p> <p>_____</p> <p>Note: This process can take 15–30 minutes, depending on your configuration.</p> <p>_____</p> <pre>Sync with CLARiiON backend done Validating consistency group configuration done Contacting source site new_york, please wait... done Running restore requires shutting down source site new_york. Do you wish to continue? [yes or no] yes Shutting down remote site new_york done</pre>
4.	<p>At the prompt, continue with the failback after your local EMC Service Provider or EMC Customer Service has helped to verify that the source VNX for Block and RecoverPoint/SE LUNs are ready (operational) by typing yes.</p> <pre>Is source site new_york ready for storage restoration ? [yes or no] yes</pre> <p>_____</p> <p>Important: The source VNX for Block and the RecoverPoint/SE CRR link must be operational. At this point, do not restart the source VNX for File.</p> <p>_____</p> <p>_____</p> <p>Note: Under certain conditions, LUNs might not be in the proper state to fail back, in which case the storage system synchronization does not complete and the failback command exits. If necessary, contact your local EMC Service Provider or EMC Customer Service for assistance.</p> <p>_____</p>

Step	Action
5.	<p>At the prompt, proceed with the failback after you ensure that the source VNX for File is powered on and its Control Station (CS0) is operational and on the data network by typing yes. Ensure that CS1 is halted.</p> <pre>Is source site ready for network restoration ? [yes or no] yes _____</pre> <p>Note: At this point, the standby Data Movers are halted and the destination file systems and shares become unavailable to network clients for a certain amount of time, depending on the total amount of replicated data. The destination devices are set to read-only while the source and destination Blocks fully synchronize.</p> <pre>_____ Restoring local servers done Waiting for local servers to reboot done Removing NBS access from local server server_2 .. done Removing NBS access from local server server_3 .. done _____</pre> <p>Note: Next, the consistency group failback occurs. If it is successful, the source site again becomes the production site.</p> <pre>_____ Failing back ... consistency group : cg_new_york _____</pre> <p>Note: The remainder of the failback can take up to 15 minutes, depending on the Data Mover configuration. The source site system and Data Movers are rebooted and restored.</p> <pre>_____ Restoring remote site new_york, please wait... done done _____</pre> <p>Note: The network restoration phase is over and the failback is complete.</p>
6.	<p>Exit root by typing:</p> <pre># exit exit</pre>
7.	<p>Exit dradmin by typing:</p> <pre>\$ exit logout</pre>
<p>After the failback process completes, wait 5–10 minutes for CS0 to come back up before logging in to the source VNX for File (new_york) and managing it directly from the source nasadmin account. If you have a dual Control Station environment, keep CS1 powered off after the failback. Also, if a Data Mover was replaced after a failover at the destination, the setup procedure performed for the hardware at the destination must also be performed at the source site after the failback process completes by your local EMC Service Provider or EMC Customer Service. Resolve failback failures on page 177 describes error scenarios that might occur during the failback process. If you have problems with the restored system, contact your EMC Service Provider or EMC Customer Service.</p>	

Verify operations after a failback (active/passive)

Step	Action
1.	Log in to the source system (new_york) as nasadmin.

Step	Action
2.	<p>Verify that the consistency group on the source system has returned to the primary role after the failback by using this command syntax:</p> <pre># nas_rp_cg -info {<name> id=<id>}</pre> <p>where:</p> <p><name> = name of the consistency group</p> <p><id> = ID assigned to the consistency group</p> <p>Example:</p> <p>To verify that the consistency group on new_york has returned to the primary role after the failback, type:</p> <pre>\$ nas_rp_cg -info cg_new_york</pre> <pre>id = 4 name = cg_new_york rpa = rpa1 source copy = Src_dev10 remote copy = DR_dev13 source clar id = APM00102102333 remote clar id = APM00102400657 contains control luns = True transfer state = ACTIVE replication direction = local -> remote role = Primary transfer mode = Async rpo = SYSTEM</pre> <pre>Replication sets Id Name Src LUN Dst LUN Size 277 RSet 1 4 25 2147483648 278 RSet 10 119 119 214748364800 279 RSet 11 120 120 214748364800 280 RSet 12 101 101 214748364800 281 RSet 13 102 102 214748364800 282 RSet 14 103 103 214748364800 283 RSet 15 104 104 214748364800 284 RSet 16 105 105 214748364800 285 RSet 17 106 106 214748364800 286 RSet 18 107 107 214748364800 287 RSet 19 108 108 214748364800 288 RSet 2 0 18 11811160064 289 RSet 20 109 109 214748364800 290 RSet 21 110 110 214748364800 291 RSet 22 111 111 214748364800 292 RSet 23 112 112 214748364800 293 RSet 3 1 19 11811160064 294 RSet 4 113 113 214748364800 295 RSet 5 114 114 214748364800 296 RSet 6 115 115 214748364800 297 RSet 7 116 116 214748364800 298 RSet 8 117 117 214748364800 299 RSet 9 118 118 214748364800</pre>

Step	Action
3.	<p>Display information about the file systems associated with Data Movers on restored new_york by typing:</p> <pre>\$ /nas/bin/server_df ALL</pre> <p>Output:</p> <pre>server_2 : Filesystem kbytes used avail capacity Mounted on ufs4 413030384 576 413029808 0% /ufs4 ufs3 413030384 576 413029808 0% /ufs3 ufs2 413030384 576 413029808 0% /ufs2 ufs1 413030384 576 413029808 0% /ufs1 root_fs_common 13624 5256 8368 39% /.etc_common root_fs_2 114592 728 113864 1% / server_3 : Error 2: server_3 : No such file or directory failed to complete command server_4 : Filesystem kbytes used avail capacity Mounted on root_fs_common 13624 5256 8368 39% /.etc_common ufslocal3 206515184 576 206514608 0% /ufslocal3 ufslocal2 413030384 576 413029808 0% /ufslocal2 ufslocal1 413030384 576 413029808 0% /ufslocal1 root_fs_4 114592 712 113880 1% / server_5 : Error 2: server_5 : No such file or directory failed to complete command</pre> <p>Note: No information is reported for the standby Data Movers as nasadmin. Information is only reported for the two source Data Movers; server_2 is RecoverPoint/SE-protected.</p>

Configuring RecoverPoint (active/active')

To preinitialize the configuration before configuring active/active' RecoverPoint/SE, refer to [Preinitialize the configuration on page 70](#).

The tasks to configure active/active' RecoverPoint/SE are:

- ◆ [Initialize the configuration \(active/active'\)](#) on page 108
- ◆ [Failover the source system \(active/active'\)](#) on page 128
- ◆ [Failback the source system \(active/active'\)](#) on page 134

Initialize the configuration (active/active')

Initializing an active/active' configuration, which is bidirectional, involves initializing the source and destination systems. The initialization tasks are performed by your local EMC Service Provider.

Before you begin

- ◆ The operating environment for VNX for File must be installed on the source and destination systems.
- ◆ RecoverPoint/SE CRR link must be operational between the source and destination sites.
- ◆ The requirements summarized in [Planning considerations on page 45](#) must be met to ensure that the VNX for Block and the VNX for File are set up correctly.
- ◆ If there is a default local standby, evaluate which standby relationships are needed for the RecoverPoint/SE configuration. For example, in a four-Data Mover configuration, where server_5 is the default local standby, remove the local standby relationship on new_york between server_5 and server_4 and set the server_5 type to nas. To remove the standby relationship, use the server_standby command, as described in [Data Mover configuration checklist on page 49](#). If you need to change the Data Mover type from standby to regular, use the server_setup command.
- ◆ Before running `/nas/sbin/nas_rp -cabinetdr -init` to select an active Data Mover as a remote standby Data Mover, clean up any part of the configuration, such as the network configuration that no longer applies to that Data Mover.
- ◆ If you have a dual Control Station environment, halt CS1 before the initialization. You can perform RecoverPoint/SE operations by using CS0 only. After halting CS1, verify the halt by typing the `/nas/sbin/getreason` command, and checking for the line "0 - slot_1 powered off" in the output.
- ◆ Make sure that you have the global Block account password established for VNX for Block. This account must be associated with Block Manager or higher privileges to manage all storage system settings in the domain.

Note: If the existing global Block account password that you specify as part of the initialization procedure changes, but the VNX for File storage information has not been updated, you get an error. Step 4 of [Initialize from the destination system \(active/passive\) on page 75](#) shows an example of setting the password. After initialization is complete, you can capture the change by rerunning the `/nas/sbin/nas_rp -cabinetdr -init` command. However, after a failover, you must use the `nas_storage` command with the `modify` option to update the VNX for File storage security information on the Control Station. [Modify VNX for Block security information after a failover on page 162](#) provides more information.

- ◆ In an active/active' configuration, ensure that a different UID for a remote administration account user (for example, `dradmin`) is used for each RecoverPoint/SE direction. [Remote administration account recommendations on page 45](#) provides the steps needed to perform this task.

- ♦ To avoid any path errors, always log in as nasadmin, switch (su) to root, and then run the nas_rp -cabinet -init command from /nas/sbin.
When using su, make sure that you follow the steps in the procedure exactly. Do not use su unless explicitly instructed.
- ♦ Initialization fails if errors such as missing RecoverPoint/SE consistency group configuration, missing Control LUNS in the consistency group, and mismatched source and remote site LUN sizes occur. [Resolve initialization failures on page 172](#) describes failure scenarios associated with initialization.
- ♦ Verify the active/active' Data Mover configuration. [Table 8 on page 109](#) shows the four-Data Mover configuration used in this section.

Table 8. Sample active/active' Data Mover configuration

new_york Data Mover	Direction	new_jersey Data Mover
server_2 (source Data Mover)	----->	server_2 (remote standby for source Data Mover on new_york)
server_3 (local standby for server_2)	----->	server_3 (remote standby for new_york source Data Mover's local standby)
server_4 (remote standby for source Data Mover on new_jersey)	<-----	server_4 (source Data Mover)
server_5 (remote standby for new_jersey source Data Mover's local standby)	<-----	server_5 (local standby for server_4)

In this configuration:

- Initialization is performed on new_york, where new_york serves as the destination for source new_jersey. Data Movers server_4 and server_5 on source new_jersey have RecoverPoint/SE standbys server_4 and server_5 on new_york. Here, server_5 on new_jersey is also a local standby for server_4. Consistency group cg_new_jersey is created on both systems.
- Initialization is performed on new_jersey, where new_jersey serves as the destination for source new_york. Data Movers server_2 and server_3 on source system new_york have RecoverPoint/SE standbys server_2 and server_3 on new_jersey. Consistency group cg_new_york is created on both systems.

Note: For the purpose of example, assume that this active/active' configuration is newly configured, not changed from an active/passive configuration.

Procedure

The tasks to initialize the configuration are:

- ◆ [Verify remote administration account \(Optional\) on page 110](#)
- ◆ [Initialize one system \(active/active'\) on page 113](#)
- ◆ [Initialize the second system \(active/active'\) on page 117](#)
- ◆ [Verify configuration \(active/active'\) on page 121](#)

Verify remote administration account (Optional)

Ensure that a different user ID (UID) for a remote administration account user (dradmin) is used for each direction. Having different UIDs for the remote administration account user in each direction in the active/active' configuration ensures that the correct Data Mover (server) information is always displayed for the appropriate command when a failover is activated.

Before you begin

- ◆ If the user ID for a RecoverPoint/SE remote administration account (for example, dradmin) user in both directions is the same, delete the user ID and then the remote administration account on one Control Station (not both).
- ◆ To list the user ID, use the `nas_acl -list` command. You do not need root privileges to list information, but you do need root privileges to delete the ID.
- ◆ Use the Unisphere software to manage the remote administration accounts.
- ◆ For a new installation, perform steps 3 and 5 only. If the same UID is already used for a dradmin user in both RecoverPoint/SE CRR directions, perform steps 1–5.

Procedure

Step	Action
1.	<p>Delete the user ID associated with the remote administration account (for example, 500) from one Control Station by using this command syntax:</p> <pre># nas_acl -delete -user <numerical_id></pre> <p>where:</p> <p><numerical_id> = user ID of the remote administration account to delete</p> <p>Example:</p> <p>To delete a user with ID 500, type:</p> <pre># nas_acl -delete -user 500</pre>
2.	<p>Delete the remote administration account on one Control Station (not both) through the Unisphere software by doing the following:</p> <ol style="list-style-type: none"> 1. Log on to the remote system through Unisphere. 2. From the list of systems, select the VNX system from which you want to delete the remote administration account. 3. Select Settings ► Security ► User Management ► Local Users for File. 4. Select the remote administration account that you want to delete. 5. Click Delete.

Step	Action
3.	<p>Create a remote administration user account to one Control Station in each RecoverPoint/SE CRR direction through Unisphere by doing the following:</p> <ol style="list-style-type: none"> 1. Log on to the remote system through Unisphere, and from the list of systems, select the VNX system on which you want to create the remote administration account. 2. Select Settings ► Security ► User Management ► Local Users for File. 3. Click Create. The Create User dialog box appears. 4. In the User Name field, type the name of the user account. 5. In the UID field, manually specify the UID for the account. 6. Type and confirm the password: <ol style="list-style-type: none"> a. In the New Password field, type the new password. b. In the Confirm New Password field, retype the password. c. In the Password Expiration (Days) field, type the number of days after which the password will expire. 7. From the Primary Group list, select a primary role. 8. From the Group Role Membership list, select one or more group roles of which the user account will be a member. The user account does not need to have the same role level on both systems. <hr/>Note: All users are required to belong to the nasadmin group.<hr/> 9. In the Client Access field, select CLI access allowed to allow the user account to be accessed through CLI. This option is selected by default.
4.	<p>In the CLI, switch (su) to root by typing:</p> <pre>\$ su</pre> <p>Password:</p>
5.	<p>Delete the relationship with the remote VNX for File by using this command syntax:</p> <pre># nas_cel -delete <cel_name></pre> <p>where: <cel_name> = current name of the remote VNX for File in the configuration</p> <p>Example: To delete the relationship with the remote VNX for File, type:</p> <pre># nas_cel -delete new_jersey</pre>

Step	Action
6.	<p>Re-create the relationship with the remote VNX for File by using this command syntax:</p> <pre># nas_cel -create <cel_name> -ip <ip> -passphrase <passphrase></pre> <p>where:</p> <p><ip> = current IP address of the remote Control Station in slot 0 (CS0)</p> <p><passphrase> = current 6–15 character password</p> <p>Example:</p> <p>To re-create the relationship with the remote VNX for File, type:</p> <pre># nas_cel -create new_jersey -ip 192.168.96.87 -passphrase nasadmin</pre> <hr/> <p>Note: Ensure you use the current VNX for File name, IP address, and passphrase configuration when the relationship is reestablished.</p> <hr/>
<hr/> <p>Note: After the relationship between the systems is established by nas_cel as described in Preinitialize the configuration on page 70, run /nas/sbin/nas_rp -cabinetdr -init, following the instructions in Initialize the configuration (active/active) on page 108. This ensures that the unique user IDs can be selected for each remote administration account user during RecoverPoint/SE initialization.</p> <hr/>	

Initialize one system (active/active')

Initialize one of the systems (new_york) that serve as a destination in the active/active' configuration.

Step	Action
1.	Log in to the system (new_york) that is to serve as a destination as nasadmin and switch (su) to root.

Step	Action
2.	<p>Start the active/active' initialization process and identify the other system by using this command syntax:</p> <pre># /nas/sbin/nas_rp -cabinetdr -init <cel_name></pre> <p>where:</p> <p><cel_name> = name of a source system</p> <p>Example:</p> <p>To initialize new_jersey as the source site to communicate with new_york, type:</p> <pre># /nas/sbin/nas_rp -cabinetdr -init new_jersey</pre> <p>Output:</p> <pre>Culham with RecoverPoint Disaster Recovery Initializing new_jersey --> new_york Contacting new_jersey for remote storage info Local storage system: FNM00094700042 Remote storage system: FNM00093600019 Discovering storage on new_jersey (may take several minutes) Setting security information for FNM00094700042 Discovering storage on new_york (may take several minutes) Contacting new_jersey for remote storage info Gathering server information... Contacting new_jersey for server capabilities... Analyzing server information...</pre>

Step	Action
3.	<p>At the prompt, configure each source Data Mover that is to be RecoverPoint/SE-protected with a remote standby Data Mover. To specify the Data Mover relationships, type the appropriate selection:</p> <p>Review Data Mover configuration checklist on page 49 to determine your Data Mover configuration. A destination site local standby cannot be paired with a source site local standby for disaster recovery. The initialization procedure validates and enforces standby Data Mover compatibility and DR eligibility for the Data Movers. Ensure Data Mover eligibility on page 151 provides more information.</p> <ul style="list-style-type: none"> ◆ Type the selection number (not the server ID for a Data Mover) associated with the Data Mover to configure. For example, selection 3 for server_4. ◆ Type v to verify the current server configuration. You can verify the configuration after specifying each source Data Mover relationship. If there is an error, it is reported. ◆ Type q to quit the initialization process. ◆ Type c to continue with the initialization process after specifying the Data Mover relationships. ◆ Type d to display more information if the selection menu indicates that a Data Mover is ineligible for a remote disaster recovery configuration. In this case, the Data Mover is not selectable, and is not eligible for remote DR message appears within parentheses after the server name, as shown in Ensure Data Mover eligibility on page 151. ◆ If you are rerunning the initialization to change your Data Mover configuration, type r after specifying a selection number to remove the configuration of a destination Data Mover currently serving as a remote standby. Change the Data Mover configuration on page 159 contains more information and steps. ◆ Type b to return to the previous selection screen after specifying a destination Data Mover. <p>Example:</p>

Step	Action
	<pre> Source servers available to be configured for remote DR ----- 1. server_2:new_jersey 2. server_3:new_jersey 3. server_4:new_jersey 4. server_5:new_jersey [local standby] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_jersey server: 3 Destination servers available to act as remote standby ----- 1. server_2:new_york server_3:new_york [local standby] 3. server_4:new_york 4. server_5:new_york b. Back Select a new_york server: 3 Source servers available to be configured for remote DR ----- 1. server_2:new_jersey 2. server_3:new_jersey 3. server_4:new_jersey [remote standby is server_4:new_york] 4. server_5:new_jersey [local standby] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_jersey server: 4 Destination servers available to act as remote standby ----- 1. server_2:new_york server_3:new_york [local standby] server_4:new_york [is remote standby is server_4:new_jersey] 4. server_5:new_york b. Back Select a new_york server: 4 Source servers available to be configured for remote DR ----- 1. server_2:new_jersey 2. server_3:new_jersey 3. server_4:new_jersey [remote standby is server_4:new_york] 4. server_5:new_jersey [remote standby is server_5:new_york] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_jersey server: c Standby configuration validated OK Using administrative user "rdfadmin" </pre>

Step	Action
4.	<p>At the prompt, continue with the initialization when the initialization sequence begins by typing yes.</p> <p>_____</p> <p>Note: If you type no, the initialization aborts with an informational message.</p> <p>_____</p> <p>Example:</p> <pre> Initializing (new_jersey-->new_york) Do you wish to continue? [yes or no] yes Setting up server_3 on new_jersey Setting up server_2 on new_jersey Setting acl for server_3 on new_york Setting acl for server_2 on new_york Updating the VNX for File domain information Creating consistency group on new_york 45 cg_new_york rpal Prod_98 Remote_46 True Creating replication sets on eng56446 Replication Sets 314 RSet 1 114 18 11811160064 315 RSet 2 115 19 11811160064 316 RSet 3 118 16 2147483648 317 RSet 4 25 90 10737418240 318 RSet 5 26 115 10737418240 done </pre>
5.	<p>Exit root by typing:</p> <pre> # exit exit </pre>

Initialize the second system (active/active')

Initialize the system (new_jersey) that serves as a destination.

Step	Action
1.	Log in to the system (new_jersey) that is the destination as nasadmin and switch (su) to root.

Step	Action
2.	<p>Start the active/active' initialization process and identify the source system by using this command syntax:</p> <pre># /nas/sbin/nas_rp -cabinetdr -init <cel_name></pre> <p>where:</p> <p><cel_name> = name of the source system</p> <p>Example:</p> <p>To make the configuration bidirectional and initialize new_york as the source site for destination site new_jersey, type:</p> <pre># /nas/sbin/nas_rp -cabinetdr -init new_york</pre> <p>Culham with RecoverPoint Disaster Recovery</p> <p>Initializing new_york --> new_jersey</p> <p>Contacting new_jersey for remote storage info</p> <p>Local storage system: FNM00093600019 Remote storage system: FNM00094700042</p> <p>Discovering storage on new_york (may take several minutes) Setting security information for FNM00093600019</p> <p>Discovering storage on new_jersey (may take several minutes)</p> <p>Contacting new_york for remote storage info Gathering server information... Contacting new_york for server capabilities... Analyzing server information...</p>

Step	Action
3.	<p>At the prompt, configure each source Data Mover that is to be RecoverPoint/SE-protected with a remote standby Data Mover. To specify the Data Mover relationships, type the appropriate selection:</p> <ul style="list-style-type: none"> ◆ Type the selection number (not the server ID for a Data Mover) associated with the Data Mover to configure. For example, selection 1 for server_2. ◆ Type v to verify the current server configuration. You can verify the configuration after specifying each source Data Mover relationship. If there is an error, it is reported. ◆ Type q to quit the initialization process. ◆ Type c to continue with the initialization process after specifying the Data Mover relationships. ◆ Type d to display more information if the selection menu indicates that a Data Mover is ineligible for a remote disaster recovery configuration. In this case, the Data Mover is not selectable, and is not eligible for remote DR message appears within parentheses after the server name, as shown in Ensure Data Mover eligibility on page 151. ◆ If you are rerunning the initialization to change your Data Mover configuration, type r after specifying a selection number to remove the configuration of a destination Data Mover currently serving as a remote standby. Change the Data Mover configuration on page 159 contains more information and steps. ◆ Type b to return to the previous selection screen after specifying a destination Data Mover. <hr/> <p>Note: Review Data Mover configuration checklist on page 49 to determine your Data Mover configuration. A destination site local standby cannot be paired with a source site local standby for disaster recovery. The initialization procedure validates and enforces standby Data Mover compatibility and DR eligibility for the Data Movers. Ensure Data Mover eligibility on page 151 provides more information.</p> <hr/> <p>Example:</p>

Step	Action
	<pre> Source servers available to be configured for remote DR ----- 1. server_2:new_york 2. server_3:new_york [local standby] server_4:new_york [is remote standby for server_4:new_jersey] server_5:new_york [is remote standby for server_5:new_jersey] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: 1 Destination servers available to act as remote standby ----- 1. server_2:new_jersey 2. server_3:new_jersey server_4:new_jersey [remote standby is server_4:new_york] server_5:new_jersey [remote standby is server_5:new_york] b. Back Select a new_jersey server: 1 Source servers available to be configured for remote DR ----- 1. server_2:new_york [remote standby is server_2:new_jersey] 2. server_3:new_york [local standby] server_4:new_york [is remote standby for server_4:new_jersey] server_5:new_york [is remote standby for server_5:new_jersey] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: 2 Destination servers available to act as remote standby ----- server_2:new_jersey [is remote standby for server_2:new_york] 2. server_3:new_jersey server_4:new_jersey [remote standby is server_4:new_york] server_5:new_jersey [remote standby is server_5:new_york] b. Back Select a new_jersey server: 2 Source servers available to be configured for remote DR ----- 1. server_2:new_york [remote standby is server_2:new_jersey] 2. server_3:new_york [remote standby is server_3:new_jersey] server_4:new_jersey [remote standby is server_4:new_york] server_5:new_jersey [remote standby is server_5:new_york] v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: c Standby configuration validated OK </pre>

Step	Action
	Using administrative user "rdfadmin"
4.	<p>At the prompt, continue with the initialization when the initialization sequence begins by typing yes.</p> <hr/> <p>Note: If you type no, the initialization aborts with an informational message. Note that both systems are informed about the consistency group creation.</p> <hr/> <p>Example:</p> <pre> Initializing (new_york-->new_jersey) Do you wish to continue? [yes or no] yes Setting up server_3 on new_york Setting up server_2 on new_york Setting acl for server_3 on new_jersey Setting acl for server_2 on new_jersey Updating the Celerra domain information Creating consistency group on new_jersey 45 cg_new_jersey rpal Prod_96 Remote_48 True Creating replication sets on new_jersey Replication Sets 314 RSet 1 114 18 11811160064 315 RSet 2 115 19 11811160064 316 RSet 3 118 16 2147483648 317 RSet 4 25 90 10737418240 318 RSet 5 26 115 10737418240 done </pre> <hr/> <p>Note: The active/active' configuration process is complete.</p> <hr/> <p>Note: Prior to a failover, any storage system configuration changes affecting the initialized configuration can be captured by rerunning /nas/sbin/nas_rp -cabinetdr -init.</p> <hr/>

Verify configuration (active/active')

You can verify the RecoverPoint/SE active/active' operations after initialization on either system in the configuration.

Step	Action
1.	Log in to one of the systems (new_jersey) as nasadmin and switch (su) to root.

Step	Action
2.	<p>List all information for all Data Movers by typing:</p> <pre># nas_server -info -all</pre> <hr/> <p>Note: The output indicates which Data Movers are owned by dradmin and serve as remote standbys (in the acl= field and the type= field), as well as which Data Movers have remote standbys (in the RDFstandby= field). The output also identifies a Data Mover that has a local standby (standby=), and a Data Mover that serves as a local standby (standbyfor=). The standbyfor= has a value only if a Data Mover serves as a local standby.</p> <hr/> <pre>id = 1 name = server_2 acl = 2000, owner=dradmin, ID=500 type = standby slot = 2 member_of = standbyfor = status : defined = enabled actual = online, ready id = 2 name = server_3 acl = 2000, owner=dradmin, ID=500 type = standby slot = 3 member_of = standbyfor = status : defined = enabled actual = online, ready id = 3 name = server_4 acl = 1000, owner=nasadmin, ID=201 type = nas slot = 4 member_of = standby = server_5, policy=auto RDFstandby= slot=4 status : defined = enabled actual = online, ready id = 4 name = server_5 acl = 1000, owner=nasadmin, ID=201 type = standby slot = 5 member_of = standbyfor = server_4 RDFstandby= slot=5 status : defined = enabled actual = online, ready</pre>

Step	Action																																																
3.	<p>List all Data Movers, including the ones owned by dradmin as remote standbys by typing:</p> <pre># nas_server -list</pre> <p>Output:</p> <table border="1"> <thead> <tr> <th>id</th> <th>type</th> <th>acl</th> <th>slot</th> <th>groupID</th> <th>state</th> <th>name</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>4</td> <td>2000</td> <td>2</td> <td></td> <td>0</td> <td>server_2</td> </tr> <tr> <td>2</td> <td>4</td> <td>2000</td> <td>3</td> <td></td> <td>0</td> <td>server_3</td> </tr> <tr> <td>3</td> <td>1</td> <td>1000</td> <td>4</td> <td></td> <td>0</td> <td>server_4</td> </tr> <tr> <td>4</td> <td>4</td> <td>1000</td> <td>5</td> <td></td> <td>0</td> <td>server_5</td> </tr> </tbody> </table> <p>Note: If you run the command as nasadmin, only server_4 and server_5 would appear in the list. server_2 and server_3 are not available to the nasadmin account because they are now managed by the dradmin account.</p>	id	type	acl	slot	groupID	state	name	1	4	2000	2		0	server_2	2	4	2000	3		0	server_3	3	1	1000	4		0	server_4	4	4	1000	5		0	server_5													
id	type	acl	slot	groupID	state	name																																											
1	4	2000	2		0	server_2																																											
2	4	2000	3		0	server_3																																											
3	1	1000	4		0	server_4																																											
4	4	1000	5		0	server_5																																											
4.	<p>List the consistency groups on the VNX for File by typing:</p> <pre># nas_rp -cg -list</pre> <p>Output:</p> <table border="1"> <thead> <tr> <th>ID</th> <th>name</th> <th>owner</th> <th>storage ID</th> <th>acl</th> <th>type</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>cg_new_york</td> <td>500</td> <td>APM00042000817</td> <td>0</td> <td>RP</td> </tr> <tr> <td>2</td> <td>cg_new_jersey</td> <td>0</td> <td>APM00042000817</td> <td>0</td> <td>RP</td> </tr> </tbody> </table> <table border="1"> <thead> <tr> <th>ID</th> <th>Name</th> <th>RPA ID</th> <th>Prod Copy</th> <th>Remote Copy</th> </tr> </thead> <tbody> <tr> <td>Control</td> <td>LUN CG</td> <td></td> <td></td> <td></td> </tr> <tr> <td>1</td> <td>cg_new_york</td> <td>4</td> <td>Src_dev10</td> <td>DR_dev13</td> </tr> <tr> <td></td> <td>True</td> <td></td> <td></td> <td></td> </tr> <tr> <td>2</td> <td>cg_new_jersey</td> <td>5</td> <td>?</td> <td>?</td> </tr> <tr> <td></td> <td>?</td> <td></td> <td></td> <td></td> </tr> </tbody> </table>	ID	name	owner	storage ID	acl	type	1	cg_new_york	500	APM00042000817	0	RP	2	cg_new_jersey	0	APM00042000817	0	RP	ID	Name	RPA ID	Prod Copy	Remote Copy	Control	LUN CG				1	cg_new_york	4	Src_dev10	DR_dev13		True				2	cg_new_jersey	5	?	?		?			
ID	name	owner	storage ID	acl	type																																												
1	cg_new_york	500	APM00042000817	0	RP																																												
2	cg_new_jersey	0	APM00042000817	0	RP																																												
ID	Name	RPA ID	Prod Copy	Remote Copy																																													
Control	LUN CG																																																
1	cg_new_york	4	Src_dev10	DR_dev13																																													
	True																																																
2	cg_new_jersey	5	?	?																																													
	?																																																

Step	Action																																																																																																																								
5.	<p>List the information for the source consistency group by typing:</p> <pre># nas_rp -cg -info cg_new_york</pre> <p>Output:</p> <pre>id = 4 name = cg_new_york rpa = rpa1 source copy = Src_dev10 remote copy = DR_dev13 source clar id = APM00102102333 remote clar id = APM00102400657 contains control luns = True transfer state = ACTIVE replication direction = local <- remote role = Primary transfer mode = Async rpo = SYSTEM</pre> <p>Replication sets</p> <table border="1"> <thead> <tr> <th>Id</th> <th>Name</th> <th>Src LUN</th> <th>Dst LUN</th> <th>Size</th> </tr> </thead> <tbody> <tr><td>277</td><td>RSet 1</td><td>4</td><td>25</td><td>2147483648</td></tr> <tr><td>278</td><td>RSet 10</td><td>119</td><td>119</td><td>214748364800</td></tr> <tr><td>279</td><td>RSet 11</td><td>120</td><td>120</td><td>214748364800</td></tr> <tr><td>280</td><td>RSet 12</td><td>101</td><td>101</td><td>214748364800</td></tr> <tr><td>281</td><td>RSet 13</td><td>102</td><td>102</td><td>214748364800</td></tr> <tr><td>282</td><td>RSet 14</td><td>103</td><td>103</td><td>214748364800</td></tr> <tr><td>283</td><td>RSet 15</td><td>104</td><td>104</td><td>214748364800</td></tr> <tr><td>284</td><td>RSet 16</td><td>105</td><td>105</td><td>214748364800</td></tr> <tr><td>285</td><td>RSet 17</td><td>106</td><td>106</td><td>214748364800</td></tr> <tr><td>286</td><td>RSet 18</td><td>107</td><td>107</td><td>214748364800</td></tr> <tr><td>287</td><td>RSet 19</td><td>108</td><td>108</td><td>214748364800</td></tr> <tr><td>288</td><td>RSet 2</td><td>0</td><td>18</td><td>11811160064</td></tr> <tr><td>289</td><td>RSet 20</td><td>109</td><td>109</td><td>214748364800</td></tr> <tr><td>290</td><td>RSet 21</td><td>110</td><td>110</td><td>214748364800</td></tr> <tr><td>291</td><td>RSet 22</td><td>111</td><td>111</td><td>214748364800</td></tr> <tr><td>292</td><td>RSet 23</td><td>112</td><td>112</td><td>214748364800</td></tr> <tr><td>293</td><td>RSet 3</td><td>1</td><td>19</td><td>11811160064</td></tr> <tr><td>294</td><td>RSet 4</td><td>113</td><td>113</td><td>214748364800</td></tr> <tr><td>295</td><td>RSet 5</td><td>114</td><td>114</td><td>214748364800</td></tr> <tr><td>296</td><td>RSet 6</td><td>115</td><td>115</td><td>214748364800</td></tr> <tr><td>297</td><td>RSet 7</td><td>116</td><td>116</td><td>214748364800</td></tr> <tr><td>298</td><td>RSet 8</td><td>117</td><td>117</td><td>214748364800</td></tr> <tr><td>299</td><td>RSet 9</td><td>118</td><td>118</td><td>214748364800</td></tr> </tbody> </table> <p>_____</p> <p>Note: In this example, new_york has two consistency groups, one for which it is the source, and one for which it is the destination. The entry for cg_new_york identifies the role as Primary. The entry for cg_new_jersey identifies the role as Secondary.</p> <p>_____</p>	Id	Name	Src LUN	Dst LUN	Size	277	RSet 1	4	25	2147483648	278	RSet 10	119	119	214748364800	279	RSet 11	120	120	214748364800	280	RSet 12	101	101	214748364800	281	RSet 13	102	102	214748364800	282	RSet 14	103	103	214748364800	283	RSet 15	104	104	214748364800	284	RSet 16	105	105	214748364800	285	RSet 17	106	106	214748364800	286	RSet 18	107	107	214748364800	287	RSet 19	108	108	214748364800	288	RSet 2	0	18	11811160064	289	RSet 20	109	109	214748364800	290	RSet 21	110	110	214748364800	291	RSet 22	111	111	214748364800	292	RSet 23	112	112	214748364800	293	RSet 3	1	19	11811160064	294	RSet 4	113	113	214748364800	295	RSet 5	114	114	214748364800	296	RSet 6	115	115	214748364800	297	RSet 7	116	116	214748364800	298	RSet 8	117	117	214748364800	299	RSet 9	118	118	214748364800
Id	Name	Src LUN	Dst LUN	Size																																																																																																																					
277	RSet 1	4	25	2147483648																																																																																																																					
278	RSet 10	119	119	214748364800																																																																																																																					
279	RSet 11	120	120	214748364800																																																																																																																					
280	RSet 12	101	101	214748364800																																																																																																																					
281	RSet 13	102	102	214748364800																																																																																																																					
282	RSet 14	103	103	214748364800																																																																																																																					
283	RSet 15	104	104	214748364800																																																																																																																					
284	RSet 16	105	105	214748364800																																																																																																																					
285	RSet 17	106	106	214748364800																																																																																																																					
286	RSet 18	107	107	214748364800																																																																																																																					
287	RSet 19	108	108	214748364800																																																																																																																					
288	RSet 2	0	18	11811160064																																																																																																																					
289	RSet 20	109	109	214748364800																																																																																																																					
290	RSet 21	110	110	214748364800																																																																																																																					
291	RSet 22	111	111	214748364800																																																																																																																					
292	RSet 23	112	112	214748364800																																																																																																																					
293	RSet 3	1	19	11811160064																																																																																																																					
294	RSet 4	113	113	214748364800																																																																																																																					
295	RSet 5	114	114	214748364800																																																																																																																					
296	RSet 6	115	115	214748364800																																																																																																																					
297	RSet 7	116	116	214748364800																																																																																																																					
298	RSet 8	117	117	214748364800																																																																																																																					
299	RSet 9	118	118	214748364800																																																																																																																					

Step	Action																																																																																																																								
6.	<p>List the information for the destination consistency group by typing:</p> <pre>nas_rp -cg -info cg_new_jersey</pre> <p>Output:</p> <pre>id = 25 name = cg_new_jersey rpa = rpa1 source copy = Src_dev10 remote copy = DR_dev13 source clar id = APM00102102333 remote clar id = APM00102400657 contains control luns = True transfer state = ACTIVE replication direction = local -> remote role = Secondary transfer mode = Async rpo = SYSTEM</pre> <p>Replication sets</p> <table border="1"> <thead> <tr> <th>Id</th> <th>Name</th> <th>Src LUN</th> <th>Dst LUN</th> <th>Size</th> </tr> </thead> <tbody> <tr><td>277</td><td>RSet 1</td><td>4</td><td>25</td><td>2147483648</td></tr> <tr><td>278</td><td>RSet 10</td><td>119</td><td>119</td><td>214748364800</td></tr> <tr><td>279</td><td>RSet 11</td><td>120</td><td>120</td><td>214748364800</td></tr> <tr><td>280</td><td>RSet 12</td><td>101</td><td>101</td><td>214748364800</td></tr> <tr><td>281</td><td>RSet 13</td><td>102</td><td>102</td><td>214748364800</td></tr> <tr><td>282</td><td>RSet 14</td><td>103</td><td>103</td><td>214748364800</td></tr> <tr><td>283</td><td>RSet 15</td><td>104</td><td>104</td><td>214748364800</td></tr> <tr><td>284</td><td>RSet 16</td><td>105</td><td>105</td><td>214748364800</td></tr> <tr><td>285</td><td>RSet 17</td><td>106</td><td>106</td><td>214748364800</td></tr> <tr><td>286</td><td>RSet 18</td><td>107</td><td>107</td><td>214748364800</td></tr> <tr><td>287</td><td>RSet 19</td><td>108</td><td>108</td><td>214748364800</td></tr> <tr><td>288</td><td>RSet 2</td><td>0</td><td>18</td><td>11811160064</td></tr> <tr><td>289</td><td>RSet 20</td><td>109</td><td>109</td><td>214748364800</td></tr> <tr><td>290</td><td>RSet 21</td><td>110</td><td>110</td><td>214748364800</td></tr> <tr><td>291</td><td>RSet 22</td><td>111</td><td>111</td><td>214748364800</td></tr> <tr><td>292</td><td>RSet 23</td><td>112</td><td>112</td><td>214748364800</td></tr> <tr><td>293</td><td>RSet 3</td><td>1</td><td>19</td><td>11811160064</td></tr> <tr><td>294</td><td>RSet 4</td><td>113</td><td>113</td><td>214748364800</td></tr> <tr><td>295</td><td>RSet 5</td><td>114</td><td>114</td><td>214748364800</td></tr> <tr><td>296</td><td>RSet 6</td><td>115</td><td>115</td><td>214748364800</td></tr> <tr><td>297</td><td>RSet 7</td><td>116</td><td>116</td><td>214748364800</td></tr> <tr><td>298</td><td>RSet 8</td><td>117</td><td>117</td><td>214748364800</td></tr> <tr><td>299</td><td>RSet 9</td><td>118</td><td>118</td><td>214748364800</td></tr> </tbody> </table>	Id	Name	Src LUN	Dst LUN	Size	277	RSet 1	4	25	2147483648	278	RSet 10	119	119	214748364800	279	RSet 11	120	120	214748364800	280	RSet 12	101	101	214748364800	281	RSet 13	102	102	214748364800	282	RSet 14	103	103	214748364800	283	RSet 15	104	104	214748364800	284	RSet 16	105	105	214748364800	285	RSet 17	106	106	214748364800	286	RSet 18	107	107	214748364800	287	RSet 19	108	108	214748364800	288	RSet 2	0	18	11811160064	289	RSet 20	109	109	214748364800	290	RSet 21	110	110	214748364800	291	RSet 22	111	111	214748364800	292	RSet 23	112	112	214748364800	293	RSet 3	1	19	11811160064	294	RSet 4	113	113	214748364800	295	RSet 5	114	114	214748364800	296	RSet 6	115	115	214748364800	297	RSet 7	116	116	214748364800	298	RSet 8	117	117	214748364800	299	RSet 9	118	118	214748364800
Id	Name	Src LUN	Dst LUN	Size																																																																																																																					
277	RSet 1	4	25	2147483648																																																																																																																					
278	RSet 10	119	119	214748364800																																																																																																																					
279	RSet 11	120	120	214748364800																																																																																																																					
280	RSet 12	101	101	214748364800																																																																																																																					
281	RSet 13	102	102	214748364800																																																																																																																					
282	RSet 14	103	103	214748364800																																																																																																																					
283	RSet 15	104	104	214748364800																																																																																																																					
284	RSet 16	105	105	214748364800																																																																																																																					
285	RSet 17	106	106	214748364800																																																																																																																					
286	RSet 18	107	107	214748364800																																																																																																																					
287	RSet 19	108	108	214748364800																																																																																																																					
288	RSet 2	0	18	11811160064																																																																																																																					
289	RSet 20	109	109	214748364800																																																																																																																					
290	RSet 21	110	110	214748364800																																																																																																																					
291	RSet 22	111	111	214748364800																																																																																																																					
292	RSet 23	112	112	214748364800																																																																																																																					
293	RSet 3	1	19	11811160064																																																																																																																					
294	RSet 4	113	113	214748364800																																																																																																																					
295	RSet 5	114	114	214748364800																																																																																																																					
296	RSet 6	115	115	214748364800																																																																																																																					
297	RSet 7	116	116	214748364800																																																																																																																					
298	RSet 8	117	117	214748364800																																																																																																																					
299	RSet 9	118	118	214748364800																																																																																																																					

Step	Action
7.	Get information about the RecoverPoint/SE CRR configuration from new_jersey by typing: <code># /nas/sbin/nas_rp -cabinet -info</code> Output:

Step	Action
	<pre> ***** Consistency Group Configuration ***** id = 3 name = cg_new_jersey rpa = rpa1 source copy = Prod_98 remote copy = Remote_46 source clar id = FNM00093600019 remote clar id = FNM00094700042 contains control luns = True ***** Servers configured with RPstandby ***** id = 3 name = server_4 acl = 1000, owner=nasadmin, ID=201 type = nas slot = 4 member_of = standby = server_5, policy=auto RDFstandby= slot=4 status : defined = enabled actual = online, ready id = 4 name = server_5 acl = 1000, owner=nasadmin, ID=201 type = standby slot = 5 member_of = standbyfor= server_4,server_3 RDFstandby= slot=5 status : defined = enabled actual = online, ready ***** Servers configured as standby ***** id = 1 name = server_2 acl = 2000, owner=dradmin, ID=500 type = standby slot = 2 member_of = standbyfor= status : defined = enabled actual = online, ready id = 2 name = server_3 acl = 2000, owner=dradmin, ID=500 type = standby slot = 3 member_of = standbyfor= status : </pre>

Step	Action
	<pre>defined = enabled actual = online, ready</pre>

Failover the source system (active/active')

You activate a RecoverPoint/SE failover in an active/active' configuration from one of the systems. In this example, the failover is performed on `new_jersey`, which serves as the destination for the source system (`new_york`). Therefore, the failover places `new_jersey` in the active role.

Before you begin

- ◆ After a failover has been activated, do not make any storage system configuration changes.
- ◆ Prior to a failover, storage system configuration changes affecting the initialized configuration can be captured by rerunning the `/nas/sbin/nas_rp -cabinetdr -init` command. A change to the existing global Block password requires an update of the VNX for File storage security information on the Control Station. After initialization is complete, this change can be captured by rerunning the `/nas/sbin/nas_rp -cabinetdr -init` command. However, after a failover, you must use the `nas_storage` command with the `-modify` option. [Modify VNX for Block security information after a failover on page 162](#) provides more information.

Note: It is strongly recommended that your local EMC Service Provider performs a test failover and restore as part of the initial RecoverPoint/SE setup and verification process.

Procedure

The tasks to activate an active/active' RecoverPoint/SE failover are:

- ◆ [Failover from the destination system on page 128](#)
- ◆ [Verify active/active' operations after failover on page 132](#)

Failover from the destination system

Activate the failover from a destination system by using the remote administration account (`dradmin`). In an unplanned failover scenario, one of the source systems (`new_york`) attached to a remote VNX for Block is assumed to have become unavailable and requires a failover to its destination VNX for File (`new_jersey`).

To activate a test failover when no real disaster scenario exists, use the following guidelines, based on whether the goal is to perform a graceful failover or to simulate a true disaster scenario:

- ◆ To simply perform a graceful test failover, do the following before activating the failover:

- Ensure that the RecoverPoint/SE links and the IP data network connections between Control Stations are up. You can also keep the source VNX for File powered up during the activation.
- Check your consistency group information to verify that the group condition is in the appropriate state for the failover activation.

Note: If you activate a test failover with the links down, or the consistency group is not in a proper state to activate a failover, a full synchronization is performed automatically during the failback to reconstruct the consistency group and replication sets, which is time-consuming and undesirable for a test scenario. If you activate the failover when the consistency group is not in the proper state to fail over, a warning message appears. In this case, abort the failover activation, ensure that the consistency group is resynchronized, and then retry the failover activation.

- ◆ Verify that the data at the destination is in a proper state before activating the failover. If not, the data might be in an unknown condition, requiring one or more file system checks by using fsck as well as a full synchronization.
- ◆ If you are without IP connectivity between the Control Stations prior to the failover activation, you can manually shut down the source Data Movers by using the `/nas/sbin/nas_halt now` command.
- ◆ After a test failover, you should reboot the source-site Control Station.
- ◆ You must have a valid RecoverPoint/SE configuration with a remote standby configured before you failover.
- ◆ If any VNX for Block configuration changes have been made after initial configuration, these changes should be reflected by the VNX for File configuration prior to a failover activation. Although the RecoverPoint/SE consistency group name should not be changed on the VNX for Block, if it is, then your local EMC Service Provider or EMC Customer Service must perform the procedure to correct the VNX for File consistency group configuration before a failover can be successfully activated.

Important: For sites with redundant Control Stations, ensure that all RecoverPoint/SE cabinet DR management commands, including `/nas/sbin/nas_rp -cabinetdr -init, -failover, and -failback`, are run from CS0. Ensure that CS1 is halted at both sites before you run the failover or failback commands. After halting CS1, verify the halt by typing the `/nas/sbin/getreason` command, and checking for the line "0 - slot_1 powered off" in the output.

Step	Action
1.	<p>Log in to the destination system (new_jersey) from which you must perform a failover as dradmin.</p> <hr/> <p>Note: The password you supply is the same one specified during initialization, as shown in Initialize the second system (active/active') on page 117.</p> <hr/>
2.	<p>Switch (su) to root by typing:</p> <pre>\$ su</pre> <p>Password:</p>

Step	Action
3.	<p>Fail over from the source system to the destination system by typing:</p> <pre data-bbox="418 365 912 390"># /nas/sbin/nas_rp -cabinetdr -failover</pre> <p>Output:</p> <pre data-bbox="418 466 1198 516">Sync with CLARiiON backend done Validating consistency group configuration done</pre> <hr data-bbox="418 533 656 537"/> <p>Note: As part of this process, all Data Movers at the source are halted. At the destination, do not shut down or reboot any Data Movers during a failover or a failback. Resolve failover failures on page 174 provides more information about errors that can occur during a failover.</p> <hr data-bbox="418 638 656 642"/>

Step	Action
4.	<p>At the prompt, continue with the failover process after verifying that the source system is ready for shutdown by typing yes.</p> <hr/> <p>Note: If you type no, the failover aborts with an informational message.</p> <hr/> <p>Note: This process can take 10–20 minutes, depending on your configuration.</p> <hr/> <pre> Is source site new_york ready for complete shut down (power OFF)? [yes or no] yes Contacting source site new_york, please wait... done Shutting down remote site new_york done Failing over ... consistency group : cg_new_york Failing over Devices ... done Adding NBS access to local server server_2 done Adding NBS access to local server server_3 done Adding NBS access to local server server_4 done Adding NBS access to local server server_5 done Activating the target environment ... done </pre> <hr/> <p>Note: At this point, the RecoverPoint/SE remote standby Data Movers become active.</p> <hr/> <pre> server_2 : going offline rdf : going active replace in progress ...done failover activity complete server_3 : going offline rdf : going active replace in progress ...done failover activity complete commit in progress (not interruptible)...done commit in progress (not interruptible)...done commit in progress (not interruptible)...done commit in progress (not interruptible)...done done </pre> <hr/> <p>Note: The failover is now complete.</p> <hr/>
5.	<p>Exit root by typing.</p> <pre> # exit exit </pre>

Step	Action
	Note: After a failover, do not perform any storage system configuration changes. In the event of a true disaster where the source VNX for Block is unavailable, contact your local EMC Customer Support Representative to coordinate failback activities.

Verify active/active' operations after failover

Step	Action
1.	Log in to the destination system (new_jersey) as nasadmin and switch (su) to root.

Step	Action
2.	<p>Display detailed information for the Data Movers by typing:</p> <pre># nas_server -info -all</pre> <hr/> <p>Note: This step can be done as nasadmin or root.</p> <hr/> <p>Output:</p> <pre>id = 1 name = server_2 acl = 0 type = nas slot = 2 member_of = standby = server_3, policy=auto RDFstandby= slot=2 status : defined = enabled actual = online, ready id = 2 name = server_3 acl = 0 type = standby slot = 3 member_of = standbyfor= server_2 RDFstandby= slot=3 status : defined = enabled actual = online, ready id = 3 name = server_4 acl = 2000, owner=rdf, ID=500 type = standby slot = 4 member_of = standbyfor= status : defined = enabled actual = online, ready id = 4 name = server_5 acl = 2000, owner=rdf, ID=500 type = standby slot = 5 member_of = standbyfor= status : defined = enabled actual = online, ready</pre>

Step	Action
3.	<p>Display information about the failed over RecoverPoint/SE standby Data Movers on new_jersey by typing:</p> <pre># server_df ALL</pre> <p>Output:</p> <pre>server_2 : Filesystem kbytes used avail capacity Mounted on root_fs_common 13624 5256 8368 39% /.etc_common root_fs_2 114592 728 113864 1% / server_3 : Error 2: server_3 : No such file or directory failed to complete command server_4 : Filesystem kbytes used avail capacity Mounted on root_fs_common 13624 5256 8368 39% /.etc_common root_fs_4 114592 688 113904 1% / server_5 : Error 2: server_5 : No such file or directory failed to complete command</pre> <p>----- Note: From dradmin, no information is reported for local standbys. For example, server_3, which is a local standby. -----</p>

Failback the source system (active/active')

Performing a failback of the source system involves restoring full connectivity to the file systems on the source system.

Before you begin

Restoration of the source system includes two phases:

- ♦ The storage restoration phase provides network access to the destination VNX for File while the destination and source VNX for Block systems synchronize. The length of the synchronization is based on the amount of data that has been updated since the destination system was activated.
- ♦ The network restoration phase suspends network clients from file system access. Note that with active write I/O it could take some time to perform the synchronization.

Procedure

The tasks to restore the source system are:

- ♦ [Prepare for the failback \(active/active'\) on page 135](#)

- ◆ [Failback the source system from its destination \(active/active'\)](#) on page 135

Note: It is strongly recommended that your EMC Service Provider performs a test failover and failback as part of the initial RecoverPoint/SE setup and verification process.

Prepare for the failback (active/active')

Under the guidance of your EMC Customer Support Representative, prepare for the failback before proceeding with the failback procedures on the destination system (new_jersey).

Step	Action
1.	Request a complete system check for the VNX for Block and RecoverPoint/SE. This is required to verify proper RecoverPoint/SE operations. Important: If this is a true disaster scenario, keep the source VNX for File (new_york) and its VNX for Block powered off until you are instructed to power them up by your EMC Customer Support Representative.
2.	Power up the VNX for Block after you have been told that it is safe to proceed, and then continue with the restore procedure on the destination VNX for File. Note: Do not attempt to restart the source VNX for File (new_york).

Failback the source system from its destination (active/active')

Use the remote administration account (dradmin) on the destination system (new_jersey) to failback the source system.

Step	Action
1.	Log in to the destination system (new_jersey) as dradmin and switch (su) to root.
2.	Start the failback of the source system from the destination system by typing: <pre># /nas/sbin/nas_rp -cabinetdr -failback</pre> Note: Use the absolute path. If you do not use the absolute path, the command fails with a general command error. If the command fails, rerun the command with the absolute path. Also, ensure that no other NAS command is running and using dradmin when the /nas/sbin/nas_rp -cabinetdr -failback command is running. During the failback, do not shut down or reboot any Data Movers.

Step	Action
3.	<p>At the prompt, shut down the source after ensuring that the synchronization is successful and after your local EMC Customer Support Representative has verified that you can proceed with a failback. Shut down the source-site system by typing yes.</p> <p>_____</p> <p>Note: If you type no, the failback aborts with an informational message.</p> <p>_____</p> <p>This step validates the consistency group configuration, and then requests a shutdown of the source.</p> <p>_____</p> <p>Note: This process can take 15–30 minutes, depending on your configuration.</p> <p>_____</p> <pre> Sync with CLARiion backend done Validating consistency group configuration done Contacting source site new_york, please wait... done Running restore requires shutting down source site new_york. Do you wish to continue? [yes or no] yes Shutting down remote site new_york done </pre>
4.	<p>At the prompt, continue with the failback after your local EMC Customer Support Representative has helped verify that the source VNX for Block and RecoverPoint/SE LUNs are ready (operational) by typing yes.</p> <pre> Is source site new_york ready for storage restoration ? [yes or no] yes </pre> <p>_____</p> <p>Important: The source VNX for Block and the RecoverPoint/SE CRR link must be operational. At this point, do not restart the source VNX for File.</p> <p>_____</p> <p>_____</p> <p>Note: Under certain conditions, LUNs might not be in the proper state to fail back, in which case the storage system synchronization does not complete and the restore command exits. If this occurs, Resolve failback failures on page 177 provides more information. If necessary, contact your local EMC Service Provider or EMC Customer Support Representative for assistance.</p> <p>_____</p>

Step	Action
5.	<p>At the prompt, proceed with the failback after you ensure that the source system is powered on and its Control Station (CS0) is operational and on the data network by typing yes. Ensure that CS1 is halted.</p> <pre>Is source site ready for network restoration ? [yes or no] yes _____</pre> <p>Note: At this point, the standby Data Movers are halted and the destination file systems and shares become unavailable to network clients for a certain amount of time, depending on the total amount of replicated data. The destination consistency groups are set to read-only.</p> <pre>_____</pre> <pre>Restoring local servers done Waiting for local servers to reboot done Removing NBS access from local server server_2 .. done Removing NBS access from local server server_3 .. done Removing NBS access from local server server_4 .. done Removing NBS access from local server server_5 .. done _____</pre> <p>Note: Next, the consistency group failback occurs. If it is successful, the source site again becomes the production site.</p> <pre>_____</pre> <pre>Failing back ... consistency group : cg_new_york _____</pre> <p>Note: The remainder of the failback can take up to 15 minutes, depending on the Data Mover configuration. The source site system and Data Movers are rebooted and restored.</p> <pre>_____</pre> <pre>Restoring remote site new_york, please wait... done done _____</pre> <p>Note: The network failback phase is over and the failback is complete.</p> <pre>_____</pre>
6.	<p>Exit root by typing:</p> <pre># exit exit</pre>
7.	<p>Exit rdfadmin by typing:</p> <pre>\$ exit logout</pre>
<p>Note: After the failback process completes, wait 5–10 minutes for CS0 to come back up before logging in to the source system (new_york) and managing it directly from the source nasadmin account. If you have a dual Control Station environment, keep CS1 powered off after the failback. Also, if a Data Mover was replaced after a failover at the destination, the setup procedure performed for the hardware at the destination must also be performed at the source site after the failback process completes by your local EMC Service Provider or EMC Customer Service. Resolve failback failures on page 177 describes error scenarios that might occur during the failback process. If you have problems with the restored system, contact your local EMC Customer Support Representative.</p>	

If the IP address of a Control Station changes after initialization, rerun the `/nas/sbin/nas_rp -cabinetdr -init` command. To change a VNX for File hostname or IP address, follow the procedures described in *Configuring and Managing Networking on VNX*. Also, rerun the `/nas/sbin/nas_rp -cabinetdr -init` command after any storage configuration change that affects any of the replication sets, storage groups, or consistency groups used by VNX for File. After a failover, however, do not perform any storage system configuration changes.

Note: Perform RecoverPoint management tasks by using either the `/nas/sbin/nas_rp -cabinetdr -info` or `nas_rp -cg` command. In a dual Control Station environment, run these commands only from CS0, the primary Control Station in slot 0.

The tasks to manage RecoverPoint/SE with VNX are:

- ◆ [Manage RPA information on page 140](#)
- ◆ [Check cabinet-level RecoverPoint/SE information on page 142](#)
- ◆ [List consistency group information on page 145](#)
- ◆ [Get detailed consistency group information on page 146](#)
- ◆ [Suspend consistency group operation on page 149](#)
- ◆ [Resume consistency group operations on page 150](#)
- ◆ [Ensure Data Mover eligibility on page 151](#)
- ◆ [Change the Data Mover configuration on page 159](#)
- ◆ [Modify VNX for Block security information after a failover on page 162](#)
- ◆ [Change file system configuration on page 163](#)

Manage RPA information

To manage RPA information, you must:

- ◆ [List RPA information on page 140](#)
- ◆ [Add RPA entry to the NAS database on page 140](#)
- ◆ [Get detailed RPA information on page 141](#)
- ◆ [Update RPA on page 141](#)

List RPA information

Basic RPA information includes the RPA name and IP address.

Step	Action
1.	Log in to the destination system (new_jersey) as nasadmin and switch (su) to root.
2.	List the RPAs configured in the system by typing: <pre>\$ nas_rp -rpa -list</pre> <pre>Id name ipaddress 1 rpa1 172.24.173.9</pre>

Add RPA entry to the NAS database

RPA entries must be added to the NAS database along with their IP address and login information. This can be done in a non-interactive mode as well with the -password option specified in the CLI.

Step	Action
1.	Log in to the destination system (new_jersey) as nasadmin.
2.	Add the RPA information to the NAS database by typing: <pre>\$ /nas/sbin/nas_rp -rpa -add rpa1 -ip 10.245.64.16 -admin admin -password admin</pre> <p>Output:</p> <pre>Done</pre>

Get detailed RPA information

RPA information includes details such as the RPA name, IP address and administrator, installation ID, license, activation code, and version.

Step	Action
1.	Log in to the destination system (new_jersey) as nasadmin.
2.	<p>List RPA information from the NAS database and license settings from the backend by typing:</p> <pre>\$ nas_rp -rpa -info rpa1</pre> <p>Output:</p> <pre>Name = rpa1 ID = 1 IP = 172.24.173.9 Administrator = admin Installation ID = f442ea06c265ac910b153fcf855d98a8b86fe1f0 License = Rts6q8aDCHx1175rYkKxuQs/NqZtosXmSWwXd0ZIoJIVaXw6Ot9TLuMXk-TD7/rAnVnvEUL2wfgF+5eBMw0RS0wAA Activation Code = 13dfacBpSohMgUIzHwgRJPYhboBvK1L56RnfU6CELkDoWX-Na/PJB7whLptoiXfGsASdTgaqtHs57UFP7M7wECwAA Version = 3.2.SP2.P2 (h.26)</pre>

Update RPA

Step	Action
1.	Log in to the destination system (new_jersey) as nasadmin.
2.	<p>Repair any SSH RSA key issues and enable the CS-RPA communication based on SSH Key authentication by using this command syntax:</p> <pre>\$ nas_rp -rpa -update {<rpaname> id=<id>}</pre> <p>where:</p> <pre><rpaname> = name of the RPA <id> = ID assigned to the RPA</pre> <p>Example</p> <pre>\$ nas_rp -rpa -update rpa1</pre> <p>Output:</p> <pre>done</pre>

Check cabinet-level RecoverPoint/SE information

Cabinet-level RecoverPoint/SE information includes:

- ◆ Information about the active consistency group that is eligible for failover, including the group state, condition, and number of replication sets. This is the same information displayed with the `nas_rp -cg -info` command for the active consistency group.
- ◆ Information about the Data Movers (servers) that have been configured with standbys or as a local standby, and if active/active', the Data Movers that are configured as RDF standbys, as well as the Data Mover status.

Step	Action
1.	Log in to the system (<code>new_york</code>) as <code>nasadmin</code> and switch (<code>su</code>) to root.

Step	Action
2.	List general information about RecoverPoint/SE by typing: <code># /nas/sbin/nas_rp -cabinetdr -info</code> Output:

Step	Action
	<pre> ***** Consistency Group Configuration ***** name = cg_new_york description = new_york_as_source uid = 50:6:1:60:B0:60:14:27:2:0:0:0:0:0:0:0 state = Consistent role = Primary condition = Active recovery policy = Automatic number of mirrors = 3 mode = SYNC owner = 0 mirrored disks = root_disk,root_ldisk,d5, local clarid = APM00044700306 remote clarid = APM00050601161 mirror direction = local -> remote ***** Servers configured with RPstandby ***** id = 1 name = server_2 acl = 1000, owner=nasadmin, ID=201 type = nas slot = 2 member_of = standby = RDFstandby= slot=2 status : defined = enabled actual = online, active id = 2 name = server_3 acl = 1000, owner=nasadmin, ID=201 type = nas slot = 3 member_of = standby = RDFstandby= slot=3 status : defined = enabled actual = online, ready defined = enabled actual = online, ready ***** Servers configured as standby ***** id = 3 name = server_4 acl = 2000, owner=dradmin, ID=500 type = standby slot = 4 member_of = standbyfor= status : defined = enabled actual = online, ready id = 4 name = server_5 </pre>

Step	Action
	<pre>acl = 2000, owner=dradmin, ID=500 type = standby slot = 5 member_of = standbyfor = status : defined = enabled actual = online, ready</pre>

List consistency group information

Consistency group information for the RecoverPoint/SE configuration includes basic information, such as the group name, owner, storage ID, owner (ACL ID) of the group, and type.

Step	Action
1.	<p>Log in to the VNX for File (new_jersey) as nasadmin.</p> <hr/> <p>Note: The <code>-list</code> and <code>-info</code> options do not require root permission. However, you must switch (su) to root on the source VNX for File to manipulate consistency group operations with options such as <code>suspend</code> and <code>resume</code>.</p> <hr/>
2.	<p>List general information about the RecoverPoint/SE consistency group by typing:</p> <pre>\$ nas_rp -cg -list</pre> <p>Output:</p> <pre>ID name owner storage ID acl type 1 cg_new_york 500 APM00042000817 0 MVIEW 2 cg_new_jersey 0 APM00042000817 0 MVIEW</pre> <pre>ID Name RPA ID Prod Copy Remote Copy Control LUN CG 13 NASCG_dev10_dev13 4 Src_dev10 DR_dev13 True</pre> <hr/> <p>Note: In this active/active' example, <code>cg_new_york</code> is owned by <code>dradmin</code> (500) because <code>new_jersey</code> serves as the destination system for <code>cg_new_york</code>. For <code>cg_new_jersey</code>, <code>new_jersey</code> is the source, so the owner is 0.</p> <hr/>

Get detailed consistency group information

Detailed RecoverPoint/SE consistency group information includes information for one group by using the group name or group ID.

Before you begin

For any consistency group, you can control whether the command synchronizes the Control Station's view with that of the VNX for Block before displaying the information. By default, the command performs the synchronization.

The detailed information is from the perspective of the system on which it is run and includes the consistency group ID, name, consistency group role of primary (source) or secondary (destination), transfer mode, RPA name, RPO, and replication direction.

Procedure

Data transfer for a RecoverPoint/SE consistency group can be in one of eight states, and is displayed per remote copy. [Table 9 on page 146](#) summarizes the consistency group transfer states.

Table 9. Consistency group data transfer states

Consistency group state	Description
Paused	Data is not being transferred to a copy, because transfer has been paused by the user.
Active	Data is being transferred asynchronously to a copy.
Active (Synchronized)	Data is being transferred synchronously to a copy.
High-load	The system enters a permanent high-load state while data is being transferred to a copy.
High-load (n%)	The system enters a temporary high-load state while data is being transferred to a copy.
Init (n%)	A copy is being initialized, or undergoing a full sweep or volume sweep.
Paused by system	Data is not being transferred to a copy because the transfer has been paused by the system.
N/A	Data is not being transferred to a copy because the copy has been disabled by the user.

Display information for each consistency group

Step	Action
1.	Log in to the VNX for File (new_york) as nasadmin.

Step	Action
2.	<p>List general information about a specific RecoverPoint/SE consistency group by using this command syntax:</p> <pre># nas_rp -cg -info {<name> id=<id>}</pre> <p>where:</p> <p><name> = name of the consistency group</p> <p><id> = ID assigned to the consistency group (shown with nas_rp -cg -list)</p> <hr/> <p>Note: To get information without synchronization, use the sync no option. For example, nas_rp -cg -info cg_new_york -sync no.</p> <hr/> <p>Example:</p> <p>To get general information about a specific RecoverPoint/SE consistency group with synchronization on, type:</p> <pre>\$ nas_rp -cg -info cg_new_york</pre> <p>Output:</p>

Step	Action																																																																																																																								
	id = 13																																																																																																																								
	name = cg_new_york																																																																																																																								
	rpa = rpa1																																																																																																																								
	source copy = Src_dev10																																																																																																																								
	remote copy = DR_dev13																																																																																																																								
	source clar id = APM00102102333																																																																																																																								
	remote clar id = APM00102400657																																																																																																																								
	contains control luns = True																																																																																																																								
	transfer state = ACTIVE																																																																																																																								
	replication direction = local -> remote																																																																																																																								
	role = Primary																																																																																																																								
	transfer mode = Sync																																																																																																																								
	rpo = SYSTEM																																																																																																																								
	Replication sets																																																																																																																								
	<table border="1"> <thead> <tr> <th>Id</th> <th>Name</th> <th>Src LUN</th> <th>Dst LUN</th> <th>Size</th> </tr> </thead> <tbody> <tr><td>277</td><td>RSet 1</td><td>4</td><td>25</td><td>2147483648</td></tr> <tr><td>278</td><td>RSet 10</td><td>119</td><td>119</td><td>214748364800</td></tr> <tr><td>279</td><td>RSet 11</td><td>120</td><td>120</td><td>214748364800</td></tr> <tr><td>280</td><td>RSet 12</td><td>101</td><td>101</td><td>214748364800</td></tr> <tr><td>281</td><td>RSet 13</td><td>102</td><td>102</td><td>214748364800</td></tr> <tr><td>282</td><td>RSet 14</td><td>103</td><td>103</td><td>214748364800</td></tr> <tr><td>283</td><td>RSet 15</td><td>104</td><td>104</td><td>214748364800</td></tr> <tr><td>284</td><td>RSet 16</td><td>105</td><td>105</td><td>214748364800</td></tr> <tr><td>285</td><td>RSet 17</td><td>106</td><td>106</td><td>214748364800</td></tr> <tr><td>286</td><td>RSet 18</td><td>107</td><td>107</td><td>214748364800</td></tr> <tr><td>287</td><td>RSet 19</td><td>108</td><td>108</td><td>214748364800</td></tr> <tr><td>288</td><td>RSet 2</td><td>0</td><td>18</td><td>11811160064</td></tr> <tr><td>289</td><td>RSet 20</td><td>109</td><td>109</td><td>214748364800</td></tr> <tr><td>290</td><td>RSet 21</td><td>110</td><td>110</td><td>214748364800</td></tr> <tr><td>291</td><td>RSet 22</td><td>111</td><td>111</td><td>214748364800</td></tr> <tr><td>292</td><td>RSet 23</td><td>112</td><td>112</td><td>214748364800</td></tr> <tr><td>293</td><td>RSet 3</td><td>1</td><td>19</td><td>11811160064</td></tr> <tr><td>294</td><td>RSet 4</td><td>113</td><td>113</td><td>214748364800</td></tr> <tr><td>295</td><td>RSet 5</td><td>114</td><td>114</td><td>214748364800</td></tr> <tr><td>296</td><td>RSet 6</td><td>115</td><td>115</td><td>214748364800</td></tr> <tr><td>297</td><td>RSet 7</td><td>116</td><td>116</td><td>214748364800</td></tr> <tr><td>298</td><td>RSet 8</td><td>117</td><td>117</td><td>214748364800</td></tr> <tr><td>299</td><td>RSet 9</td><td>118</td><td>118</td><td>214748364800</td></tr> </tbody> </table>	Id	Name	Src LUN	Dst LUN	Size	277	RSet 1	4	25	2147483648	278	RSet 10	119	119	214748364800	279	RSet 11	120	120	214748364800	280	RSet 12	101	101	214748364800	281	RSet 13	102	102	214748364800	282	RSet 14	103	103	214748364800	283	RSet 15	104	104	214748364800	284	RSet 16	105	105	214748364800	285	RSet 17	106	106	214748364800	286	RSet 18	107	107	214748364800	287	RSet 19	108	108	214748364800	288	RSet 2	0	18	11811160064	289	RSet 20	109	109	214748364800	290	RSet 21	110	110	214748364800	291	RSet 22	111	111	214748364800	292	RSet 23	112	112	214748364800	293	RSet 3	1	19	11811160064	294	RSet 4	113	113	214748364800	295	RSet 5	114	114	214748364800	296	RSet 6	115	115	214748364800	297	RSet 7	116	116	214748364800	298	RSet 8	117	117	214748364800	299	RSet 9	118	118	214748364800
Id	Name	Src LUN	Dst LUN	Size																																																																																																																					
277	RSet 1	4	25	2147483648																																																																																																																					
278	RSet 10	119	119	214748364800																																																																																																																					
279	RSet 11	120	120	214748364800																																																																																																																					
280	RSet 12	101	101	214748364800																																																																																																																					
281	RSet 13	102	102	214748364800																																																																																																																					
282	RSet 14	103	103	214748364800																																																																																																																					
283	RSet 15	104	104	214748364800																																																																																																																					
284	RSet 16	105	105	214748364800																																																																																																																					
285	RSet 17	106	106	214748364800																																																																																																																					
286	RSet 18	107	107	214748364800																																																																																																																					
287	RSet 19	108	108	214748364800																																																																																																																					
288	RSet 2	0	18	11811160064																																																																																																																					
289	RSet 20	109	109	214748364800																																																																																																																					
290	RSet 21	110	110	214748364800																																																																																																																					
291	RSet 22	111	111	214748364800																																																																																																																					
292	RSet 23	112	112	214748364800																																																																																																																					
293	RSet 3	1	19	11811160064																																																																																																																					
294	RSet 4	113	113	214748364800																																																																																																																					
295	RSet 5	114	114	214748364800																																																																																																																					
296	RSet 6	115	115	214748364800																																																																																																																					
297	RSet 7	116	116	214748364800																																																																																																																					
298	RSet 8	117	117	214748364800																																																																																																																					
299	RSet 9	118	118	214748364800																																																																																																																					

Suspend consistency group operation

Temporarily halting the replication from the source to the destination suspends the link.

Before you begin

When suspending consistency group operations:

- ◆ Changes can still be made to the source LUNs, but they are not applied to the destination LUNs (that is, secondary images) until you resume operations.
- ◆ Suspending a consistency group causes a pause in the data transfer of writes from the source to the destination replica. The data transfer state will change to paused.
- ◆ A pause stops replicating operations from the primary (source) image to a secondary (destination) mirror image. A pause can occur either automatically because of a failure

in the path to the destination image's storage processors, or manually by an administrative action, or both.

Procedure

Step	Action
1.	Log in to the source (new_york) as nasadmin.
2.	<p>Suspend RecoverPoint/SE consistency group operations by using this command syntax:</p> <pre># nas_rp -cg -suspend <name></pre> <p>where:</p> <p><name> = name of the consistency group</p> <p>Example:</p> <p>To suspend RecoverPoint/SE consistency group operations, type:</p> <pre># nas_rp -cg -suspend cg_new_york</pre> <p>done</p>

Resume consistency group operations

On the source VNX for File, resume consistency group operations and restart replicating operations.

Before you begin

- ◆ When the consistency group resumes operations, the destination LUNs are synchronized with the source LUNs. You can use this option after you suspend data transfer, or when the consistency group is in the Paused state.
- ◆ When the nas_rp -cg -resume command is run in response to a RecoverPoint/SE consistency group suspend command, the -resume command resumes data transfer of the RecoverPoint/SE consistency group by performing a manual, incremental synchronization of the group's replication set pairs, the source LUNs, and their equivalent destination LUNs. The resume option can be used to recover from a destination site failure such as a RecoverPoint/SE link down or storage processor failure.

Procedure

Step	Action
1.	Log in to the VNX for File (new_york) as nasadmin.

Step	Action
2.	<p>Resume consistency group operations and perform a manual synchronization by using this command syntax:</p> <pre># nas_rp -cg -resume <name></pre> <p>where:</p> <p><name> = name of the consistency group</p> <p>Example:</p> <p>To resume consistency group operations and perform a manual synchronization, type:</p> <pre># nas_rp -cg -resume cg_new_york</pre> <p>Output:</p> <pre>done</pre>

Ensure Data Mover eligibility

This section describes:

- ◆ Data Mover conditions that apply during RecoverPoint/SE initialization.
- ◆ How to gather additional information when a Data Mover is not eligible to participate in a remote DR configuration.
- ◆ How to verify the Data Mover relationships during initialization.

To ensure Data Mover eligibility and conditions, you must:

- ◆ [Ensure Data Mover network device compatibility on page 153](#)
- ◆ [Get additional information during initialization on page 154](#)
- ◆ [Verify configuration during initialization on page 155](#)

[Table 10 on page 151](#) summarizes the source Data Mover conditions that appear when you initialize RecoverPoint/SE.

Table 10. Source Data Mover conditions during initialization

Condition	Can you select it?	Description
[not eligible for remote DR]	No	This Data Mover is not eligible to participate in a remote disaster recovery configuration. If you see this condition during the initialization process, type the d option to get more information about why the Data Mover is ineligible.
[is remote standby for server_x]	No	This Data Mover is configured as a remote standby. This applies to an active/active configuration.

Table 10. Source Data Mover conditions during initialization *(continued)*

Condition	Can you select it?	Description
[remote standby is server_x]	Yes	This Data Mover is configured with a remote standby.
[local standby]	Yes	This Data Mover is configured as a local standby; remote standby can be configured to activate in a disaster recovery situation.
[unconfigured standby]	Yes	This Data Mover is configured as a standby; however, no primary Data Movers are configured to use it.

[Table 11 on page 152](#) summarizes the destination Data Mover conditions that appear when you initialize RecoverPoint/SE.

Table 11. Destination Data Mover conditions during initialization

Condition	Can you select it?	Description
[is remote standby for server_x]	No	This Data Mover is configured as a remote standby.
[remote standby is server_x]	No	This Data Mover is configured with a remote standby. Applies to an active/active' configuration.
[remote standby]	No	This Data Mover is configured as a remote standby but the source Data Mover cannot be determined.
[local standby]	No	One or more of the Data Movers configured to use this local standby are not remote standbys.
[unconfigured standby]	Yes	This Data Mover is configured as a standby; however, no source Data Movers are configured to use it.
[local standby for remote standbys]	Yes	This Data Mover is configured as a local standby. All Data Movers configured to use this local standby are remote standbys.
[non-root file system mounted]	No	This Data Mover has one or more user file systems mounted and cannot be a standby.

Table 11. Destination Data Mover conditions during initialization *(continued)*

Condition	Can you select it?	Description
[not compatible]	No	The source Data Mover has one or more network devices not available on this Data Mover. If you see this as a Data Mover condition, the network devices in the two cabinets do not have the same configuration. This could happen if you are mixing an NS series gateway cabinet with an NSX cabinet and you have a different number of network ports in use on both sides (for example, 6 cge ports on a source NS versus 5 on the destination NSX). Ensure Data Mover network device compatibility on page 153 provides more information about resolving Data Mover network device incompatibility.

Ensure Data Mover network device compatibility

The RecoverPoint/SE initialization procedures check for and enforce Data Mover compatibility (including network device configuration compatibility) between the source Data Movers configured for disaster recovery and the standby Data Movers at the destination. A source Data Mover and remote standby Data Mover must appear to have the same network device configuration. To ensure network device compatibility and prevent a destination Data Mover condition of not compatible during initialization, you can edit the file `/nas/site/nas_param` to specify a system parameter called `hidden_interfaces`. This parameter enables you to specify a list of device names that need to be masked to make the Data Movers in each cabinet appear to have the same network device configuration.

The E-Lab Interoperability Navigator, available on EMC Online Support, provides detailed information about Data Mover compatibility based on the different types of cabinets. Consult this tool before attempting to resolve Data Mover incompatibility problems.

Note: Ensure that you do not edit anything in `/nas/sys`.

Step	Action
1.	Log in to the source Control Station.
2.	Open the file <code>/nas/site/nas_param</code> with a text editor. A short list of configuration lines appears.

Step	Action
3.	<p>Edit the file to specify the <code>hidden_interfaces</code> parameter with a list of devices to be hidden. If you need to specify multiple devices, use a comma-separated list (for example, <code>cge6, cge5</code>). The list you supply applies to all Data Movers in the system.</p> <p>For example, with a source NS702G, which has six <code>cge</code> (copper-wire Ethernet) ports, and a destination NSX, which has five <code>cge</code> ports, the <code>hidden_interfaces</code> parameter can be specified to mask (hide) <code>cge6</code>, if it is unused on all source Data Movers: <code>hidden_interfaces:cge6:</code></p> <hr/> <p>Note: Ensure you do not add a blank line to this file.</p> <hr/> <p>To the <code>/nas/site/nas_param</code> file on the NS702G, this logically hides the <code>cge6</code> network port on all Data Movers from all VNX for File user interfaces, including Unisphere and the <code>server_sysconfig server_x -pci <device></code> command.</p>
4.	<p>Save and close the file.</p> <hr/> <p>Note: Changing this value does not require a Data Mover or Control Station reboot.</p> <hr/>

Get additional information during initialization

If a Data Mover is not selectable during the initialization process, and you see the condition (not eligible for remote DR), you can display additional information about why the Data Mover is ineligible.

The following sample excerpt from the initialization of source `new_york` from destination `new_jersey` shows ineligible Data Movers.

Action
<pre>[root@new_jersey nasadmin]# /nas/sbin/nas_rp -cabinetdr -init new_york</pre>

Output for getting information about an ineligible Data Mover

```

...
Contacting new_york for remote storage info
Gathering server information...
Contacting new_york for server capabilities...
Analyzing server information...

Source servers available to be configured for remote DR
-----
server_2:new_york [ not eligible for remote DR ]
server_3:new_york [ not eligible for remote DR ]
d.   Display details for servers not eligible for remote DR
v.   Verify standby server configuration
q.   Quit initialization process
c.   Continue initialization
Select a new_york server: d

Info 26306805889: server_2:new_york file system "local_fs" is utilizing
storage which is not mirrored to new_jersey.
Info 26306805891: server_2:new_york file system "pfs" is involved in an
IP Replicator environment.
Info 26306805888: server_2:new_york file system "local_fs" has Automatic
File System Extension enabled.
Info 26306805891: server_3:new_york file system "sfs" is involved in an
IP Replicator environment.

Press <Enter> continue

```

Verify configuration during initialization

During the initialization process, you can verify the Data Mover configuration as follows:

- ◆ You can perform the verification after each RecoverPoint/SE source-to-destination Data Mover relationship that you define.
- ◆ You can perform the verification after defining all the relationships but before you type `c` to continue and complete the initialization.
- ◆ You can rely on the software to perform the verification automatically after you type `c` to continue and complete the initialization.

The following excerpt from the initialization of source `new_york` from destination `new_jersey` shows how Data Mover verification can help clarify configuration requirements.

Action
<code>[root@new_jersey nasadmin]# /nas/sbin/nas_rp -cabinetdr -init new_york</code>

Output for Data Mover verification

Output for Data Mover verification

```

...
Contacting new_york for remote storage info
Gathering server information...
Contacting new_york for server capabilities...
Analyzing server information...

Source servers available to be configured for remote DR
-----
1.      server_2:new_york
2.      server_3:new_york [ local standby ]
v.      Verify standby server configuration
q.      Quit initialization process
c.      Continue initialization
Select a new_york server: v

Warning 17716871681: No remote standby servers specified
Press <Enter> continue
-----
Note: The warning indicates that you have not configured any remote standby Data Movers yet.
-----

The following shows the initialization process for configuring the Data Mover relationship between source server_2 and
remote server_2:

Source servers available to be configured for remote DR
-----
1.      server_2:new_york
2.      server_3:new_york [ local standby ]
v.      Verify standby server configuration
q.      Quit initialization process
c.      Continue initialization
Select a new_york server: 1

Destination servers available to act as remote standby
-----
1.      server_2:new_jersey
2.      server_3:new_jersey
b.      Back
Select a new_jersey server: 1

Source servers available to be configured for remote DR
-----
1.      server_2:new_york [ remote standby is server_2:new_jersey ]
2.      server_3:new_york [ local standby ]
v.      Verify standby server configuration
q.      Quit initialization process
c.      Continue initialization
Select a new_york server: v

Error 13421904938: new_york:server_2 has slot 3 configured as a local
standby.  new_jersey:server_3 slot 3 is not configured as a standby,
therefore local failover for new_york:server_2 will not be available after
activation

Error 13421904936: new_york:server_2 has a local standby, new_york:server_3,
which does not have a remote standby configured.

Press <Enter> continue

```

Output for Data Mover verification

Note: The first error reports that the source server_2 is configured with slot 3 (server_3) as its local standby, but the destination slot 3 is not a standby. Therefore, local failover is not available after activation. The second error reports that because source server_2 has slot 3 (server_3) as a local standby, server_3 should have a remote standby configured. For either error, you must either resolve the problem or quit. You cannot complete the initialization.

The following configures and validates the relationship between source server_3 and remote server_3:

```
Source servers available to be configured for remote DR
-----
1.      server_2:new_york [ remote standby is server_2:new_jersey ]
2.      server_3:new_york [ local standby ]
v.      Verify standby server configuration
q.      Quit initialization process
c.      Continue initialization
Select a new_york server: 2
Destination servers available to act as remote standby
-----
          server_2:new_jersey [ is remote standby for server_2:new_york ]
2.      server_3:new_jersey
b.      Back
Select a new_jersey server: 2

Source servers available to be configured for remote DR
-----
1.      server_2:new_york [ remote standby is server_2:new_jersey ]
2.      server_3:new_york [ remote standby is server_3:new_jersey ]
v.      Verify standby server configuration
q.      Quit initialization process
c.      Continue initialization
Select a new_york server: v

Standby configuration validated OK

Press <Enter> continue
Source servers available to be configured for remote DR
-----
1.      server_2:new_york [ remote standby is server_2:new_jersey ]
2.      server_3:new_york [ remote standby is server_3:new_jersey ]
v.      Verify standby server configuration
q.      Quit initialization process
c.      Continue initialization
Select a new_york server: c

Standby configuration validated OK
```

Second example of verification with local standby error

In this example of an invalid configuration, server_5 is configured as a local standby for server_3 (intended RecoverPoint/SE-protected Data Mover) and server_4 (local-only Data Mover). Both cannot share server_5, because server_3 would be owned by the remote administration account and server_4 is local-only and owned by nasadmin.

```
Source servers available to be configured for remote DR
-----
      server_2:new_york [ is remote standby for server_2:new_jersey ]
2.    server_3:new_york [ remote standby is server_3:new_jersey]
3.    server_4:new_york
4.    server_5:new_york [ remote standby is server_5:new_jersey ]
v.    Verify standby server configuration
q.    Quit initialization process
c.    Continue initialization
Select a new_york server: v

Error 13421904926: new_york:server_5 is a local standby for a mix of local-
only and remote standby Data Movers.
```

Change the Data Mover configuration

This section describes how to change a Data Mover configuration by rerunning the `/nas/sbin/nas_rp -cabinetdr -init` command as root at the destination.

The basic steps for changing the configuration are as follows:

Step	Action
1.	<p>To rerun the initialization and change the remote standby Data Mover configuration, run the <code>nas/sbin/nas_rp-cabinetdr -init</code> from the destination, as root.</p> <p>Example:</p> <pre>[root@new_jersey nasadmin]# /nas/sbin/nas_rp -cabinetdr -init new_york</pre> <p>Culham with RecoverPoint Disaster Recovery</p> <p>Initializing new_york --> new_jersey</p> <p>Contacting new_york for remote storage info</p> <p>Local storage system: FNM00094700042 Remote storage system: FNM00093600019</p> <p>Discovering storage on new_york (may take several minutes) Setting security information for FNM00094700042</p> <p>Discovering storage on new_jersey (may take several minutes)</p> <p>Contacting new_york for remote storage info Gathering server information... Contacting new_york for server capabilities... Analyzing server information...</p> <hr/> <p>Note: If the global account information is already known, you are not prompted for it.</p>
2.	<p>From the source server side, specify the selection number of the source Data Mover for which you want to change the configuration.</p> <p>Example:</p> <pre>Source servers available to be configured for remote DR ----- 1. server_2:new_york [remote standby is server_2:new_jersey] 2. server_3:new_york v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: 1</pre>
3.	<p>From the destination server side, type <code>r</code> to remove the remote standby relationship of a destination Data Mover currently serving as the remote standby for the source Data Mover.</p> <p>Example:</p> <pre>Destination servers available to act as remote standby ----- r. -> Remove server_2:new_jersey [is remote standby for serv- er_2:new_york] 2. server_3:new_jersey [unconfigured standby] b. Back Select a new_jersey server: r</pre>

Step	Action
4.	<p>From the source side, specify the selection number of the source Data Mover for which you want to define the remote standby.</p> <p>Example:</p> <pre>Source servers available to be configured for remote DR ----- 1. server_2:new_york 2. server_3:new_york v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: 1</pre>
5.	<p>From the destination server side, specify the selection number of the destination Data Mover to serve as the remote standby for the source Data Mover.</p> <p>Example:</p> <pre>Destination servers available to act as remote standby ----- 1. server_2:new_jersey [unconfigured standby] 2. server_3:new_jersey [unconfigured standby] b. Back Select a new_jersey server: 2</pre>
6.	<p>When you are done making changes, type c to continue with the initialization. If the change is valid, a message confirms the configuration is OK.</p> <p>Example:</p> <pre>Source servers available to be configured for remote DR ----- 1. server_2:new_york [remote standby is server_3:new_jersey] 2. server_3:new_york v. Verify standby server configuration q. Quit initialization process c. Continue initialization Select a new_york server: c Standby configuration validated OK _____ Note: To abort the initialization instead, type q and verify that you want to quit. A message then confirms that you have aborted the operation. _____</pre>

Step	Action
7.	<p>After the remote administration account information is detected, type yes to continue with the initialization. You then see messages reporting the update.</p> <p>Example:</p> <pre>Using administrative user "dradmin" Initializing new_york-->new_jersey Do you wish to continue? [yes or no] yes Setting up server_2 on new_york Setting acl for server_3 on new_jersey Updating the VNX for File domain information Creating consistency group on new_jersey 45 cg_new_jersey rpal Prod_98 Remote_46 True Creating replication sets on new_jersey Replication Sets 314 RSet 1 114 18 11811160064 315 RSet 2 115 19 11811160064 316 RSet 3 118 16 2147483648 317 RSet 4 25 90 10737418240 318 RSet 5 26 115 10737418240 done</pre> <p>Note: The remote administrative account you defined with the first initialization is used, so you are not prompted for the information.</p>

Modify VNX for Block security information after a failover

To perform RecoverPoint/SE cabinet disaster recovery operations, the system relies on a VNX for Block username and password to access a global administrative VNX for Block account. You specify this global VNX for Block account information as part of the initialization procedure. A change to the existing global VNX for Block password, for example, requires an update of the VNX for File storage security information on the Control Station. After initialization, this change can be captured by rerunning `/nas/sbin/nas_rp -cabinetdr -init`; however, if a failover is activated (and the source system is unavailable), you must capture the change by issuing the `nas_storage` command with the `-modify` option, followed by the security (username and password) information.

Note: Ensure that you perform this procedure when logged in as `nasadmin`, and `su` to root.

If a failover is activated, perform this procedure from the destination VNX for File to set the VNX for Block global account security information.

```

Action
-----
To modify the VNX for Block account security information, log in as nasadmin, su to root, and use the following syntax:

# nas_storage -modify {<name>|id=<storage_id>} -security [-username <username>]
  -password <password>]

where:

<name> = name of the consistency group
<storage_id> = ID assigned to the storage
<username> = username for the global VNX for Block administrative account
<password> = password for the global VNX for Block administrative account

Example (with recommended commands to get information and validate the storage):

[root@new_york nasadmin]# nas_storage -list

id    acl    name                serial_number
1     0     APM00041700549     APM00041700549
2     0     APM00042000817     APM00042000817

[root@new_york nasadmin]# nas_storage -sync -all

done

[root@new_york nasadmin]# nas_storage -modify id=2 -security -username nasadmin -password
nasadmin

Setting security information for APM00042000817    id        = 2
serial_number      = APM00042000817
name               = APM00042000817
acl                = 0

```

Change file system configuration

In the course of managing the RecoverPoint/SE configuration, you might need to change a local file system to a replicated file system, or to change a replicated file system to an unreplicated local file system. This section addresses both scenarios:

- ◆ [Change a local file system to a replicated file system on page 164](#)
- ◆ [Change a replicated file system to a local file system on page 165](#)

When you perform this kind of change, you will notice that the associated disk types change. When a disk is replicated, it has the disk type CMSTD or CMATA. When a disk is not replicated, it has the disk type CLSTD or CLATA. All disks in a storage pool must be all replicated or all unreplicated, but you cannot have a mix of both. If a mix exists based on a storage system configuration change, the `/nas/sbin/nas_rp -cabinetdr -init` command fails.

Change a local file system to a replicated file system

Use the following procedure to change a local file system to a replicated file system:

Step	Action
1.	Work with your local EMC Customer Support Representative to configure replication on the VNX for Block.
2.	<p>Run the <code>nas_storage -sync</code> command on the source system and the destination system by using the following syntax:</p> <pre>\$ nas_storage -sync {-all <name> id=<storage_id>}</pre> <p>Example:</p> <p>To synchronize the Control Station's view with that of the storage, run the following command on both the source and destination systems:</p> <pre>\$ nas_storage -sync -all</pre>
3.	<p>Unmount the file system associated with the local disk changed to replicated from the local Data Mover, if you keep the Data Mover as local. Use the following syntax:</p> <pre>\$ server_umount <movername> -perm <fs_name> [<mount_point>]</pre> <p>where:</p> <ul style="list-style-type: none"> <movername> = name of the local Data Mover <fs_name> = name of the file system <mount_point> = mount point, optional <p>Example:</p> <pre>\$ server_umount server_3 -perm ufs3</pre>
4.	<p>Run the <code>/nas/sbin/nas_rp -cabinetdr -init</code> command from the destination system.</p> <p>Example:</p> <pre>[root@new_jersey nasadmin]# /nas/sbin/nas_rp -cabinetdr -init new_york</pre>
5.	<p>Mount the new replicated file system on the replicated Data Mover, if needed, using the following syntax:</p> <pre>\$ server_mount <movername> <fs_name> [<mount_point>]</pre> <p>where:</p> <ul style="list-style-type: none"> <movername> = name of the local Data Mover <fs_name> = name of the file system <mount_point> = mount point <p>Example:</p> <pre>\$ server_mount server_3 ufs3 /ufs3</pre>

Change a replicated file system to a local file system

Use the following procedure to change a replicated file system to a local (unreplicated) file system:

Step	Action
1.	<p>Work with your local EMC Customer Support Representative to do the following:</p> <ul style="list-style-type: none"> ◆ Identify the replication sets with disk volumes of the file system as the source LUNs by using the Unisphere RecoverPoint Plugin. ◆ Remove the identified replication sets from the NAS consistency group by using the Unisphere RecoverPoint Plugin. ◆ Remove all the consistency group changes from the Unisphere RecoverPoint Plugin. ◆ Delete the remote copy LUN(s) by using Unisphere or Navicli. <hr/> <p>Note: Ensure that you stop all I/O before proceeding. Also, if you intend to change the configuration of disks in a storage pool, all disks in the storage pool must be either replicated or unreplicated, but you cannot have a mix. A mix of disks in a storage pool causes /nas/sbin/nas_rp -cabinetdr -init to fail.</p> <hr/>
2.	<p>Unmount the file system from the replicated Data Mover on the source system by using the following syntax:</p> <pre>\$ server_umount <movername> -perm <fs_name> [<mount_point>]</pre> <p>where:</p> <p><movername> = name of the local Data Mover</p> <p><fs_name> = name of the file system</p> <p><mount_point> = mount point, optional</p> <p>Example:</p> <pre>\$ server_umount server_2 -perm ufs1</pre>
3.	<p>Run the server_devconfig ALL -create -scsi -all command.</p> <p>Example:</p> <p>To save the device configuration into the Data Mover's database:</p> <pre>\$ server_devconfig ALL -create -scsi -all</pre>

Step	Action
4.	<p>Mount the file system on the local Data Mover on the source system, if needed, by using the following syntax:</p> <pre>\$ server_mount <movername> <fs_name> [<mount_point>]</pre> <p>where:</p> <p><movername> = name of the local Data Mover</p> <p><fs_name> = name of the file system</p> <p><mount_point> = mount point</p> <p>Example:</p> <pre>\$ server_mount server_2 ufs1 /ufs1</pre>
5.	<p>Run the <code>nas_storage -sync</code> command on the source system and the destination system by using the following syntax:</p> <pre>\$ nas_storage -sync {-all <name> id=<storage_id>}</pre> <p>Example:</p> <p>To synchronize the Control Station's view with that of the storage, run the following command on both the source and destination systems:</p> <pre>\$ nas_storage -sync -all</pre>
6.	<p>Run the <code>/nas/sbin/nas_rp -cabinetdr -init</code> command from the destination system.</p> <p>Example:</p> <pre>[root@new_jersey nasadmin]# /nas/sbin/nas_rp -cabinetdr -init new_york</pre>

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

If a product does not function properly or does not function as described in this document, contact your EMC Customer Support Representative.

Topics include:

- ◆ [EMC E-Lab Interoperability Navigator on page 168](#)
- ◆ [Known problems and limitations on page 168](#)
- ◆ [Error messages on page 181](#)
- ◆ [EMC Training and Professional Services on page 182](#)

EMC E-Lab Interoperability Navigator

The EMC E-Lab™ Interoperability Navigator is a searchable, web-based application that provides access to EMC interoperability support matrices. It is available on the EMC Online Support website at <http://Support.EMC.com>. After logging in, locate the applicable Support by Product page, find **Tools**, and click **E-Lab Interoperability Navigator**.

Known problems and limitations

[Table 12 on page 169](#) summarizes common troubleshooting scenarios for RecoverPoint/SE.

This section provides information on how to:

- ◆ [Retrieve information from log files on page 170](#)
- ◆ [Resolve initialization failures on page 172](#)
- ◆ [Resolve failover failures on page 174](#)
- ◆ [Resolve failback failures on page 177](#)
- ◆ [Additional situations involving Data Movers on page 180](#)

Overview of RecoverPoint/SE troubleshooting scenarios

Table 12. RecoverPoint/SE troubleshooting scenarios

Scenario	Description	Actions
Planned testing, maintenance failover activated and then failback	Handles routine maintenance for testing purposes only; the consistency group is functioning properly.	<p>Check the RecoverPoint/SE consistency group information, then manually fail over to the destination site VNX for File/Block pair while the source site is still operational. Then, perform a failback operation to recover the original configuration. When the failover occurs, the destination site assumes the primary role for all LUNs in the consistency group and the source takes the backup role.</p> <hr/> <p>Note: For testing purposes, ensure that the source VNX for Block and RecoverPoint/SE links are up when you do a test failover. If you perform a test failover with the links down, a full synchronization is automatically performed as part of the restore to reconstruct the consistency group and replication sets, which is time-consuming.</p> <hr/>
Temporary loss of one or both RecoverPoint/SE links	A loss of one or both RecoverPoint/SE links causes a consistency group data transfer state to change to Paused. After the link is recovered, the consistency group changes the data transfer state to Active and the two sites are synchronized.	When the data transfer between the production and remote site is paused, the RPA uses delta marking to cache host writes. When the data transfer is restored, the RPA uses the cached write to synchronize the two sites.

Table 12. RecoverPoint/SE troubleshooting scenarios *(continued)*

Scenario	Description	Actions
Source-site failure	A source site failure causes a consistency group data transfer state to change to Paused.	For source-site failures other than a single source storage processor failure, perform a forced failover activation. (A forced failover of the backup images is required to make them visible to the destination VNX for File.) The source images cannot be converted to backup images. Fix the source-site problem and then perform the restore, which automatically reconstructs the consistency group and replication sets with a full synchronization.
Destination-site failure	A destination site failure causes a consistency group data transfer state to change to Paused.	Repair the failure at the destination and see if the group automatically recovers from the failure.

Retrieve information from log files

System messages are reported to the system log files. To retrieve information from log files:

- ◆ Check the system log (sys_log) by using the server_log command
- ◆ Check the command error log (cmd_log.err) for message information

To retrieve substantial RecoverPoint/SE logging information:

- ◆ Use the /nas/tools/collect_support_materials script, which collects data, such as disaster recovery information, from the following log files:
 - /nas/log/dr_log.al
 - /nas/log/dr_log.al.rll
 - /nas/log/dr_log.al.err
 - /nas/log/dr_log.al.trace*
 - /nas/log/symapi.log*
 - /nas/log/rplici.log*

These log files can also be viewed individually.

- ◆ To monitor these logs while the /nas/sbin/nas_rp command is running, check the file in the /tmp directory. After the command completes, the logs appear in the /nas/log directory.

When a failover is activated and the destination system is in an activated state, the log files reside in the `/nas/rdf/<dradmin_user_id>/log` directory. Although the `/nas/rdf/<dradmin_user_id>/log/cmd_log` file currently points to `/nas/log/nas_log.al` in an activated state, view the `cmd_log` file for the remote administration account using `/nas/rdf/<dradmin_user_id>/log/nas_log.al`.

The disaster recovery (dr*) files provide state changes as well as other key informational messages. The `/nas/log/symapi.log` file logs storage-related errors.

Note: In the `/nas/log/dr_log.al` file, if you see a "Clearing orphaned inode" message, which is associated with a file system check (fsck), do not be concerned. The clearing of orphaned inodes is considered harmless.

In general, the log files can help you gather additional information after a failure (for example, a failed restore).

To alert you to consistency group condition changes on the source VNX for File, NaviEventMonitor events (ID 200) are generated and the state changes are logged in the `sys_log`. For example:

```
Mar  9 09:11:00 2006 NaviEventMonitor:3:200 Consistency group cg_new_york
on APM00042000817 is System Fractured
Mar  9 10:30:59 2006 NaviEventMonitor:5:200 Consistency group cg_new_york
on APM00042000817 is Active
```

[Table 13 on page 171](#) identifies these events, which are associated with the NaviEventMonitor facility (facility ID138). *Configuring Events and Notifications on VNX for File* provides more information about viewing and managing the handling of events and notifications.

Table 13. NaviEventMonitor facility events for RecoverPoint/SE

Event ID	Event	Event description/corrective action
200	Consistency group <group> on <system> is System Fractured	For a severity of Warning or higher, indicates that the RecoverPoint/SE consistency group is in the condition System Fractured (for example, because a RecoverPoint/SE CRR link is down).
200	Consistency group <group> on <system> is Active	As an informational event, indicates that the RecoverPoint/SE consistency group is Active, indicating that communication has resumed on the RecoverPoint/SE CRR link.

In Unisphere, the link-down event appears as an alert under the Status for a severity of Warning or higher. Unisphere online help provides more information about Unisphere elements.

Resolve initialization failures

To resolve initialization failures, consider:

- ◆ Initialization failure due to the source control LUNs not being configured as part of the consistency group on page 172
- ◆ Initialization failure due to the consistency group failing over on page 173
- ◆ Initialization failure due to a replication set being configured with LUNs of different sizes on page 173
- ◆ Initialization failure due to a replication set name having invalid characters on page 174

Initialization failure due to the source control LUNs not being configured as part of the consistency group

In this initialization failure scenario, which uses a source called `new_york` and a destination called `new_jersey`, initialization fails because the source control LUNs are not configured as part of the consistency group.

Initialization failure scenario

```
# nas_rp -cabinetdr -init new_york

Culham with RecoverPoint Disaster Recovery
Initializing new-york --> new-jersey
Contacting new-jersey for remote storage info
Local storage system: APM00044700306
Remote storage system: APM00050601161
Enter the Global Culham Block Storage account information
Username: nasadmin
Password: *****          Retype your response to validate
Password: *****
Discovering storage on new-jersey (may take several minutes)
Setting security information for APM00044700306
Discovering storage APM00050601161 (may take several minutes)
Discovering storage (may take several minutes)
Contacting new-jersey for remote storage info

Error 13431865361: No control LUNs are configured in the consistency groups.
```

To resolve this initialization failure, configure the source LUNs as part of the consistency group.

Initialization failure due to the consistency group failing over

In this initialization failure scenario, which uses a source called `new_york` and a destination called `new_jersey`, initialization fails because of the consistency group failing over.

```
Initialization failure scenario

# nas_rp -cabinetdr -init new_york

Culham with RecoverPoint Disaster Recovery
Initializing new-york --> new-jersey
Contacting new-jersey for remote storage info
Local storage system: APM00044700306
Remote storage system: APM00050601161
Enter the Global Culham Block Storage account information
Username: nasadmin
Password: *****                               Retype your response to validate
Password: *****
Discovering storage on new-jersey (may take several minutes)
Setting security information for APM00044700306
Discovering storage APM00050601161 (may take several minutes)
Discovering storage (may take several minutes)
Contacting new-jersey for remote storage info

Error 13431865361: No control LUNs are configured in the consistency groups.
```

To resolve this initialization failure, fallback the consistency group.

Initialization failure due to a replication set being configured with LUNs of different sizes

In this initialization failure scenario, which uses a source called `new_york` and a destination called `new_jersey`, initialization fails because a replication set is configured with LUNs of different sizes.

```
Initialization failure scenario

# nas_rp -cabinetdr -init new_york

Culham with RecoverPoint Disaster Recovery
Initializing new-york --> new-jersey
Contacting new-jersey for remote storage info
Local storage system: APM00044700306
Remote storage system: APM00050601161
Enter the Global Culham Block Storage account information
Username: nasadmin
Password: *****                               Retype your response to validate
Password: *****
Discovering storage on new-jersey (may take several minutes)
Setting security information for APM00044700306
Discovering storage APM00050601161 (may take several minutes)
Discovering storage (may take several minutes)
Contacting new-jersey for remote storage info

Error 13431865366: The consistency group is configured with replication
sets that have different source and destination LUN sizes.
```

To resolve this initialization failure, configure the replication set with LUNs of the same size.

Initialization failure due to a replication set name having invalid characters

In this initialization failure scenario, which uses a source called `new_york` and a destination called `new_jersey`, initialization fails because a replication set name has invalid characters.

Initialization failure scenario
<pre># nas_rp -cabinetdr -init new_york Culham with RecoverPoint Disaster Recovery Initializing new-york --> new-jersey Contacting new-jersey for remote storage info Local storage system: APM00044700306 Remote storage system: APM00050601161 Enter the Global Culham Block Storage account information Username: nasadmin Password: ***** Retype your response to validate Password: ***** Discovering storage on new-jersey (may take several minutes) Setting security information for APM00044700306 Discovering storage APM00050601161 (may take several minutes) Discovering storage (may take several minutes) Contacting new-jersey for remote storage info Error 13431865367: The consistency group is configured with replication set that has a name with invalid characters.</pre>

To resolve this initialization failure, remove invalid characters such as "." from the replication set name.

Resolve failover failures

Use the following guidelines to recover from various failover failures:

- ◆ If the source site is unreachable during the failover, check the source site and review the information in [Failover failure scenario when the source site is unreachable on page 175](#):
 - If both source storage processors are down, you need to continue with a forced failover.
 - If one source storage processor is down, do not continue with the failover. Instead, run the `nas_rp -cg -suspend` command from the source to halt consistency group operations and use the `nas_rp -cg -resume` command to restart consistency group operations. This suspend/resume trespasses the LUNs to the other storage processor automatically and you can perform a normal failover activation after the resume.
- ◆ If a RecoverPoint/SE CRR link is down, do not proceed with the failover until the link problem is resolved. The consistency group can recover automatically from a temporary link down failure as long as the recovery policy for the group is set to Automatic (the

default when the group is created as part of storage system configuration), not Manual. If the recovery policy for the group is set to Manual, you must run the `nas_rp -cg -resume` command to recover from a link down failure.

- ◆ If the consistency group configuration cannot be validated during failover, the failover fails. If there has been a storage system configuration change, follow the corrective action listed for errors 13431865362, 13431865361, and 13431865360. In a source-site failure scenario in which the storage system is damaged and storage processors A and B booted up with no configuration information, contact your local EMC Service Provider or EMC Customer Service to shut down the source storage processors or hide the source system from the destination side completely (that is, put the RecoverPoint/SE CRR link down, and shut down the IP network to the source side to simulate a power-off activate scenario) before retrying the failover. After successful failover, your local EMC Service Provider or EMC Customer Service can rebuild the configuration (which involves rebuilding the LUNs) and then run failback.

To resolve failover failures, consider:

- ◆ [Failover failure scenario when the source site is unreachable on page 175](#)
- ◆ [Resolution for failover failure scenario when the source site is unreachable on page 176](#)

Failover failure scenario when the source site is unreachable

In this failover failure scenario, the failure occurs during failover when the source site is unreachable. You are instructed to check the source site. Also, a warning message appears after a successful failover. It notifies you that the source site was down during failover and recommends that you set up the source site and start data transfer manually.

Activation failure scenario

```
# /nas/sbin/nas_rp -cabinetdr -failover
```

Activation failure scenario

```

Sync with CLARiiON backend ..... done
Validating consistency group configuration ..... done

Is source site new_york ready for complete shut down (power OFF)? [yes
or no]yes

Contacting source site new_york, please wait... done

Shutting down remote site new_york
.....
done

Failing over ... consistency group : cg_new_york

Failing over Devices ... done

Adding NBS access to local server server_2 ..... done
Adding NBS access to local server server_3 ..... done
Activating the target environment ... done

server_2 : going offline
rdf : going active
replace in progress ...done
failover activity complete
commit in progress (not interruptible)...done
commit in progress (not interruptible)...done

Warning 17726832673: The source site is down. Cannot start data transfer
to source site.

done
    
```

Resolution for failover failure scenario when the source site is unreachable

The following shows the failback (run from the destination) after the source-side storage system hardware problem is resolved and the storage system is back to normal operation.

Activation failure scenario

```
# /nas/sbin/nas_rp -cabinetdr -failback
```


Activation failure scenario

```

Sync with CLARiiON backend ..... done
Validating consistency group configuration ..... done

Contacting source site new_york, please wait... done

Running restore requires shutting down source site new_york.
Do you wish to continue? [yes or no] yes

Shutting down remote site new_york ..... done

Is source site new_york ready for storage restoration ? [yes or no] yes
Is source site ready for network restoration ? [yes or no] yes

Restoring local servers ..... done
Waiting for local servers to reboot ..... done

Removing NBS access from local server server_2 .. done
Removing NBS access from local server server_3 .. done

Failing back ... consistency group : cg_new_york

Restoring remote site new_york, please wait... done

done

```

Note: After this failback process completes, you should wait 5–10 minutes for Control Station 0 to come back up before logging in to the source system (new_york) and managing it directly from the source nasadmin account. If you have a dual Control Station environment, keep in mind that CS1 remains powered off after the failback. [Additional situations involving Data Movers on page 180](#) describes what to do after the failback if a Data Mover was replaced after failover at the destination.

Resolve failback failures

This section provides sample failback failure scenarios. [Additional situations involving Data Movers on page 180](#) describes how to handle problems involving a reboot or replacement of a Data Mover.

To resolve failback failures, consider:

- ◆ [Failed failback from destination on page 178](#)
- ◆ [Failed failback caused by state of source Data Movers on page 179](#)
- ◆ [Failback failure during fsck on page 180](#)

Failed failback from destination

In this failback failure scenario, an active/passive restore operation started by root from dradmin on destination new_jersey by using the /nas/sbin/nas_rp -cabinetdr -failback command fails. The output shows the events leading to the error message on the destination system, followed by the output from the source after taking corrective action (that is, completion of the failback from the source).

Note: When you perform a source-side restore from the source Control Station, you must use the local VNX for File administration account and perform the failback by using the path /nasmcd/sbin/nas_rp -cabinetdr -failback (not the path used for a routine destination-side failback). If you use the incorrect path and the command completes but the NAS service does not start, you can remove the /nasmcd/lock/dr.lck file and then make sure that the Control Station reboots.

```
Restore failure scenario
#nas/sbin/nas_rp -cabinetdr -failback

Sync with CLARiiON backend ..... done
Validating consistency group configuration ..... done

Contacting source site new_york, please wait... done

Running restore requires shutting down source site new_york.
Do you wish to continue? [yes or no] yes

Shutting down remote site new_york ..... done

Is source site new_york ready for storage restoration ? [yes or no] yes

Activating the target environment ... done
Waiting for consistency group is synchronized ... done

Synchronizing consistency group ... done

Is source site ready for network restoration ? [yes or no] yes

Restoring local servers ..... done
Waiting for local servers to reboot ..... done

Removing NBS access from local server server_2 .. done
Removing NBS access from local server server_3 .. done
Failing back ... consistency group : cg_new_york

Restoring remote site new_york, please wait... failed
WARNING: Cannot restore new_york. Please run restore on site new_york.
done
```

Restore failure scenario

The following shows sample output from the failback operation performed on the source system (new_york) after the failure.

```
Restore failure scenario

[root@new_york nasadmin]# /nasmcd/sbin/nas_rp -cabinetdr -failback

Trying to unmount /nas... done

Powering on servers ( please wait ) ..... done

server_2 : going standby
rdf : going active
replace in progress ...done
failover activity complete
commit in progress (not interruptible)...done
commit in progress (not interruptible)...done

done
done
```

Note: After the failback process completes, you should wait 5–10 minutes for Control Station 0 to come back up before logging in to the source system (new_york) and managing it directly from the source nasadmin account. If you have a dual Control Station environment, keep in mind that CS1 remains powered off after the failback. If the failback is still not successful, collect the following log files listed in [Retrieve information from log files on page 170](#). [Additional situations involving Data Movers on page 180](#) describes what to do after the restore if a Data Mover was replaced after failover at the destination.

Failed failback caused by state of source Data Movers

A failback operation can fail because of the source Data Mover state. [Table 14 on page 179](#) summarizes the source Data Mover states and requirements for a failback.

Table 14. Source Data Mover requirements during a failback

Source Data Mover	State	Description
Minimally, one source RecoverPoint/SE-protected Data Mover	Up (reason code 4 or 5)	To prevent an error, at least one source RecoverPoint/SE-protected Data Mover must be available during the failback.
An additional source RecoverPoint/SE-protected Data Mover whose remote standby use the same slot number	Does not matter	As long as at least one source RecoverPoint/SE-protected Data Mover is available, any other source RecoverPoint/SE-protected Data Mover can be down if both it and its remote standby use the same slot number (for example, server_2 in slot 2 at both source and destination).

Table 14. Source Data Mover requirements during a failback *(continued)*

Source Data Mover	State	Description
Any source RecoverPoint/SE-protected Data Mover whose remote standby uses a different slot number	Up (reason code 4 or 5)	A mismatch of slot numbers between a source RecoverPoint/SE-protected Data Mover and its remote standby mandates a state of Up for that source Data Mover. For example, if source Data Mover server_2 is in slot 2 and remote standby server_4 is in slot 4, then server_2 must be Up.
Any local source Data Mover (non-RecoverPoint/SE protected)	Does not matter	A local source Data Mover can be down during the failback; you can reboot it on the source VNX for File after the failback completes.

Failback failure during fsck

If the command `nas_rp -cabinetdr -failback` fails or does not respond during `fsck`, perform the following steps to recover.

Step	Action
1.	Log in to the Control Station with an appropriate administrator account.
2.	Check each Data Mover to determine whether it has remote device access permission set to read/write (rw). Use the command syntax: <pre>.server_config <mover_name> -v "nbs_list"</pre> where: <pre><mover_name> = name of the Data Mover</pre>
3.	If none of the Data Movers has remote device permission set to rw, use the <code>server_cpu</code> command to reboot at least one Data Mover that is owned by the administrator account.
4.	If the failback operation was stopped, restart it by using the <code>nas_rp -cabinetdr -failback</code> command.

Additional situations involving Data Movers

The following summarizes additional situations concerning the management of Data Movers during or after a RecoverPoint/SE failover or failback:

- ◆ If you shut down or reboot any Data Movers (RecoverPoint/SE-protected or non-protected Data Movers) at the destination while the `/nas/sbin/nas_rp -cabinetdr -failover` or the `/nas/sbin/nas_rp -cabinetdr -failback` command is running, the Control Station does not find a path to the storage system. With the communication of the VNX for File system to the storage system interrupted, the command fails. You must then rerun the

`/nas/sbin/nas_rp -cabinetdr -failover` or `-failback` command after the Data Mover is operational. Do not shut down or reboot any Data Movers at the destination while these commands are running.

- ◆ If a Data Mover needs to be replaced after a successful failover at the destination, the setup procedure performed for the hardware at the destination must also be performed at the source site after a successful RecoverPoint/SE failback. Contact your local EMC Service Provider or EMC Customer Service to perform the hardware setup procedures at the destination and source and prevent a subsequent initialization failure when you need to change the RecoverPoint/SE Data Mover configuration.
- ◆ If a failover is activated at the destination and you reboot a local Data Mover by using the wrong account (that is, you are logged in using the remote administration account instead of the VNX for File administration account at the destination and you su to root to perform the reboot), a subsequent RecoverPoint/SE failback might appear suspended while attempting to reboot the local servers, and associated fsck log messages might appear in `/tmp/dr_log.al`. In this situation, contact your local EMC Service Provider or EMC Customer Service. These procedures involve stopping the RecoverPoint/SE failback processes, rebooting the local Data Mover using the correct VNX for File administration account, and rerunning the failback as usual from the remote administration account.
- ◆ If an fsck error occurs on the source, waiting for a Data Mover to boot as part of the failback, the failback fails, and the fsck read block error appears in `dr_log.al`. In this situation, perform the failback from the source side, as shown by the sample output in [Failed failback from destination on page 178](#).
- ◆ When you run the `-init` command, it changes the ACL for the local Data Movers to 1000. This prevents RP administrators from inadvertently accessing the local Data Movers in the failed over state. However, this also prevents Global Users in Unisphere from accessing this Data Mover in the normal state.

To resolve this problem, the ACL for the local Data Movers is no longer changed when you run the `-init` command. Instead, the ACL is changed to 1111 during failover when you run the `-failover` command. This prevents RP administrators from accessing the Data Movers after failover and allows Global Users to access them in the normal state. During failback, when you run the `-failback` command, the ACL is changed to 0. If you initially use ACL 1111 intentionally for the local Data Movers on the source side, ensure that you change the ACL from 0 back to 1111 after failback. Alternately, you can change the ACL for the local Data Movers on the source side to some other value, for example, 1000, to avoid this manual change.

Error messages

All event, alert, and status messages provide detailed information and recommended actions to help you troubleshoot the situation.

To view message details, use any of these methods:

- ◆ Unisphere software:

- Right-click an event, alert, or status message and select to view Event Details, Alert Details, or Status Details.
- ◆ CLI:
 - Type `nas_message -info <MessageID>`, where `<MessageID>` is the message identification number.
- ◆ *Celerra Error Messages Guide*:
 - Use this guide to locate information about messages that are in the earlier-release message format.
- ◆ EMC Online Support website:
 - Use the text from the error message's brief description or the message's ID to search the Knowledgebase on the [EMC Online Support](#) website. After logging in to EMC Online Support, locate the applicable **Support by Product** page, and search for the error message.

EMC Training and Professional Services

EMC Customer Education courses help you learn how EMC storage products work together within your environment to maximize your entire infrastructure investment. EMC Customer Education features online and hands-on training in state-of-the-art labs conveniently located throughout the world. EMC customer training courses are developed and delivered by EMC experts. Go to the EMC Online Support website at <http://Support.EMC.com> for course and registration information.

EMC Professional Services can help you implement your system efficiently. Consultants evaluate your business, IT processes, and technology, and recommend ways that you can leverage your information for the most benefit. From business plan to implementation, you get the experience and expertise that you need without straining your IT staff or hiring and training new personnel. Contact your EMC Customer Support Representative for more information.

A

active/active'

For EMC® Symmetrix® Remote Data Facility (SRDF®) or EMC MirrorView™/Synchronous configurations, a bidirectional configuration with two production sites, each acting as the standby for the other. Each VNX for file has both production and standby Data Movers. If one site fails, the other site takes over and serves the clients of both sites. For SRDF, each Symmetrix system is partitioned into source (production) and remote destination volumes. For MirrorView/S, each VNX for block is configured to have source and destination LUNs and a consistency group.

active/passive

For SRDF or MirrorView/S configurations, a unidirectional setup where one VNX for file, with its attached system, serves as the source (production) file server and another VNX for file, with its attached storage, serves as the destination (backup). This configuration provides failover capabilities in the event that the source site is unavailable. An SRDF configuration requires Symmetrix systems as backend storage. A MirrorView/S configuration requires specific VNX for block series systems as backend storage.

C

Common Internet File System (CIFS)

File-sharing protocol based on the Microsoft Server Message Block (SMB). It allows users to share file systems over the Internet and intranets.

Consistency group

A logical entity that comprises a collection of replication sets grouped together to ensure write order consistency across all the replication sets' primary volumes. Consistency groups are used to configure protection policies, and set RPO and RTO policies according to specific resource allocation and prioritization. They are also used to fail over to, recover production from, and test any replica defined in the group.

D***destination VNX for file***

Term for the remote (secondary) VNX for file in an SRDF or MirrorView/S configuration. The destination VNX for file is typically the standby side of a disaster recovery configuration. Symmetrix configurations often call the destination VNX for file: the target VNX for file.

L***logical unit number (LUN)***

Identifying number of a SCSI or iSCSI object that processes SCSI commands. The LUN is the last part of the SCSI address for a SCSI object. The LUN is an ID for the logical unit, but the term is often used to refer to the logical unit itself.

R***remote replication***

Replication of a file system from one VNX for file to another. The source file system resides on a different system from the destination file system.

Replication set

A collection of production source volume and the replica volume/s to which it replicates. Every SAN-attached storage volume in the production storage must have a corresponding volume at each copy. Each consistency group contains as many replication sets as there are volumes in the production storage to replicate.

RPA

Hardware that is used to run the RecoverPoint software, which is responsible for most of the product's intelligence. These intelligent data protection appliances manage all aspects of reliable data replication at all sites.

S***storage processor (SP)***

Storage processor on VNX for block. On VNX for block, a circuit board with memory modules and control logic that manages the system I/O between the host's Fibre Channel adapter and the disk modules.

V***VNX***

EMC network-attached storage (NAS) product line.

VNX FileMover

Policy-based system used to determine where files should be physically stored. In most cases, policies are based on file size or last access time (LAT) or both and are used to identify data that can be moved to slower, less-expensive storage.

-info argument of `nas_rp -cabinetdr` command 142

A

absolute path requirement 101, 135
 accounts
 establishing remote DR 80
 global 17
 specifying remote DR for failover 92
 ACL and remote admin account 110
 activating RecoverPoint/SE failover 53, 88, 89, 128
 Active (Synchronized state) 146
 Active state 146
 active/active'
 failback 135
 failover 128
 initializing 108, 117
 overview of 28
 recommendations for remote admin account 45, 110
 active/passive
 configuration 73
 failback 100
 failover, activating 89
 overview of 28
 adding
 RPA entries to the NAS database 140
 adding RPA information 52, 140
 with `nas_rp -rpa` 140
 Automatic File System Extension, restriction 18

B

backend configuration
 NAS consistency group 75
 rules about changing 47, 139
 summary of 46

backend unreachable 175

C

changing
 Control Station IP address 139
 Data Mover configuration 160, 161
 checking cabinet-level RecoverPoint/SE information 142
 checklist for Data Mover configuration 49
 command
 `nas_cel -create` 70, 71
 commands
 for managing RecoverPoint/SE with VNX 139
 for managing RPA information 140
 `nas_cel -create` 52
 `nas_rp -cabinetdr` 142
 `nas_rp -cg` 55, 83, 123
 `nas_rp -rpa` 54
 `nas_rp -rpa -add` 52
 `nas_storage` 162, 164
 compatibility, Data Mover network device 153
 configuration tracking sheet 50
 configuring
 Data Movers 47, 48
 consistency group
 creation 47
 getting information about 139
 listing 83, 123
 managing 139
 resuming 150
 states 146
 suspending 149
 consistency group data transfer
 paused state 149
 consistency group data transfer state
 Paused 169, 170
 Control Station

Control Station (*continued*)
 distance 16
 preinitialization 70, 71
 rules for redundant 16, 90

D

Data Movers
 changing configuration 159, 160, 161
 checking configuration 139
 checklist for 49, 50
 conditions during initialization 151
 ensuring network device compatibility 153
 failure scenarios after failover and/or failback 181
 replacement at destination after failover 137
 replicating 47, 48
 rules for configuring 47, 48
 source DM state causing failback failure 179
 validating compatibility 151
 validating during initialization 155
 dedicated RecoverPoint/SE CRR links 26
 destination LUNs
 backend configuration of 46
 destination system
 activating failover from 89
 destination VNX
 performing failback from 101
 tasks performed on 54
 different IP subnets 99
 disk types for RecoverPoint/SE 163
 DR administrative account 80, 92
 dr_log files 170
 DR, Data Mover not eligible for 151

E

EMC E-Lab Navigator 168
 enabling
 CS-RPA communication 141
 error messages 181
 error scenarios 169
 errors
 failback 177, 179
 events 170

F

failback
 active/active 134
 active/active' 135
 active/passive 100
 system postfailover 101

failback (*continued*)
 VNX postfailover 54
 failover
 failback VNX after a 54
 initiating active/active' 128
 initiating active/passive 53, 89
 initiating RecoverPoint/SE 53, 88, 89, 128
 requirement to use remote admin account 17
 restoring system after a 101, 135
 restoring VNX after a 100
 restrictions 17
 SNMP or email for event notification after 18
 testing active/active' 128
 failure scenarios
 failback 177, 179
 overview of 168
 fibre channel links 10
 file system configuration, changing 163, 165
 files
 log 170
 firewalls 50

G

gateway configuration 13
 getting
 detailed RPA information 141
 getting information
 with nas_rp -cabinetdr 142
 with nas_rp -cg 145, 146, 147
 getting information about an ineligible Data Mover 154
 getting RPA information
 with nas_rp -rpa 141
 global account information 17
 graceful failover 128
 guidelines
 for change backend configuration 47
 for file systems 45
 for local standbys 47
 for remote admin account 45, 110

H

halting Data Movers 92, 129
 hidden_interfaces parameter 153
 host visibility of destination storage at promotion 26
 HTTP communication 80

I

initializing

- initializing (*continued*)
 - active/active' 108, 117
 - active/passive failover 53, 89
 - RecoverPoint/SE relationship 52, 76, 114, 118
 - to capture changes 47, 159, 162
- IP subnets 99
- L**
- license 13
- limits
 - nas_cel passphrase 70, 71
- link down error 169, 175
- links, RecoverPoint/SE 26
- listing
 - consistency group information 145, 146
 - consistency groups 83, 123
 - RPA information 140
- listing RPAs in the system
 - with nas_rp -rpa 140
- local file systems versus DR-protected file systems 45
- local standby configuration guidelines 47
- log files 170
- login account 80
- LUNs
 - backend configuration of 46
 - control LUN mapping 46
 - number of supported 10
- M**
- managing RecoverPoint/SE with VNX 139
- managing RPA information 140
- messages, error 181
- modifying VNX for Block security information 47, 162
- MPFS
 - restriction for 18
- N**
- nas_acl command 110
- nas_cel -create command 52
- nas_fs command
 - listing snapshot information 87
- nas_mview command
 - init 52
- nas_rp -cabinetdr command
 - failback 54, 135
 - failover 53, 92
 - info 142
 - init 52, 76, 114, 118
- nas_rp -cabinetdr command (*continued*)
 - restore 101
- nas_rp -cg command 55, 145, 146, 147, 149, 150
 - info 146, 147
 - list 145
 - resume 150
 - suspend 149
- nas_rp -rpa -add command 52
- nas_rp -rpa command 54, 140, 141
 - add 140
 - info 141
 - update 141
- nas_rp command
 - list 140
- nas_storage command
 - modify 162
 - sync 164
- network restoration phase 134
- not eligible for remote DR condition 151, 154
- P**
- passphrase 71, 80, 92
 - for nas_cel preinitialization 71
 - remote administrative 92
- paused state 149
- Paused state 146
- ports, IP network 50
- preinitializing RecoverPoint/SE relationship 52
- R**
- RecoverPoint
 - conceptual overview 24
- RecoverPoint/SE
 - disk types 163
 - links 26, 169
 - restrictions 11
 - tracking sheet 50
 - troubleshooting 169
- RecoverPoint/SE with VNX
 - configuration 26
- remote administration account
 - recommendations for active/active' 45, 110
 - specifying for failover 92
- remote administration account:
 - creating 80
- repairing SSH RSA key issues 141
- replicating Data Movers 47, 48
- Replicator 18
- restoring
 - path requirement for 101, 135
 - phases for 134
- restrictions for RecoverPoint/SE 11

resuming consistency group operation 150
RPA
 getting information about 140
rpcli.log file 170

S

SnapSure, restriction 18
source LUNs
 backend configuration of 46
source system 114, 118
source VNX 76
source-site failback 178
source-site failure 170
standby Data Movers 47, 48, 75
 sample configuration for active/passive 75
state of source Data Movers during failback 179
states of consistency group 146
storage configuration 139
storage processor failure 170, 174
storage restoration phase 134
su to root, care when using 16
suspending consistency group operation 149
symapi.log file 170

T

testing

testing (*continued*)
 active/active' failover 128
testing a failover 169
troubleshooting RecoverPoint/SE 169

U

Unisphere
 restriction after failover 16
 viewing MirrorView disk types and storage pools 19
 viewing VNX with RecoverPoint/SE disk types and storage pools 15
updating
 RPA 141
user interface 19

V

verifying
 active/active' RecoverPoint/SE 121, 132
 active/passive RecoverPoint/SE 94
VNX configuration 13
VNX for Block security information 47, 162
VNX for File cabinet configurations supported 10
volumes, configuring sites 47, 48