

Replication and Disaster Recovery Administration Guide

Hitachi Virtual Storage Platform Gx00 and Fx00 with NAS
Modules

VSP N series

Hitachi NAS Platform

Release 14.5 or higher

© 2011, 2019 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at <https://www.hitachivantara.com/en-us/company/legal.html> or https://knowledge.hitachivantara.com/Documents/Open_Source_Software.

Contents

Preface.....	6
Related Documentation.....	6
Accessing product documentation.....	9
Getting help.....	9
Comments.....	9
Chapter 1: File and Object data replication.....	10
File replication and object replication.....	11
File replication overview.....	11
Incremental file replication.....	11
Incremental block replication.....	12
Multiple stream replication.....	12
Relocating file systems.....	13
Object replication overview.....	14
Incremental replication.....	15
Transferring XVLs as links during object replication.....	15
Single migration destination for replication source and target.....	16
Separate migration destinations for replication source and target.....	17
Configuring transfer XVLs as links during object replication.....	18
Recovering a file system.....	19
Recovering and promoting a file system.....	22
Recovering and demoting a file system.....	24
Replication and disaster recovery.....	27
Replication process for planned promotions.....	28
Recovering multiple file systems.....	29
Policy-based replication.....	30
Replication schedules.....	30
Supported replication applications.....	31
TrueCopy and ShadowImage considerations.....	31
Chapter 2: Using file replication.....	33
Configuring policy-based file replication.....	33
Specifying IP addresses for an NDMP replication policy manually.....	37
Connection errors.....	38
Understanding snapshot rules.....	39

Understanding custom replication scripts.....	40
Using file replication rules.....	41
Displaying file replication rules.....	41
Adding a file replication rule.....	42
Modifying a file replication rule.....	50
Understanding files-to-exclude statements.....	51
Using file replication schedules.....	51
Displaying scheduled file replications.....	51
Adding a file replication schedule.....	53
Modifying a file replication policy.....	55
Understanding incremental replications.....	58
Displaying file replication status and reports.....	58
Enabling multiple replication streams.....	61
Configuring NDMP performance options.....	62
Maximum concurrent replications.....	64
Troubleshooting replication failures.....	64
Manually restarting a failed replication.....	65
Rolling back an incomplete replication.....	65
Chapter 3: Transferring primary access.....	67
How a transfer of primary access moves CNS links.....	68
Process of transferring primary access.....	68
Handling a failure during a transfer of primary access.....	73
Chapter 4: Using object replication.....	74
Configuring object-based replication.....	74
Read only access to replication target file systems.....	75
NDMP access to object replication target file systems.....	76
Using object replication policies.....	78
Adding object replication policies.....	78
Correcting access point problems in an object replication policy.....	83
Using object replication schedules.....	83
Adding an object replication schedule.....	83
Modifying an object replication schedule.....	85
Displaying object replication policies	88
Asynchronous replication from a secondary to a tertiary target.....	90
Benefits of Asynchronous Replication.....	90
Configuring tertiary replication targets.....	91
Displaying object replication status and reports.....	96
Chapter 5: Using storage-based mirroring.....	99
Storage-based mirroring overview.....	99

Storage-based mirroring terminology.....	99
Configuring storage-based mirrors.....	101
Updating a mirror.....	102
Performing failover on a mirror.....	103
Recovering from loss of production SDs.....	103
Copy-on-Write copyback.....	104
Removing a single mirror relationship.....	104
Removing all mirror relationships.....	105
Mounting two copies of the data on different clusters.....	106
Tolerating primary-primary relationships.....	106
Secondary SD warnings.....	107
Dealing with SSWS.....	107

Preface

This guide provides information about replicating data using file-based replication and object-based replication, provides information on setting up replication policies and schedules, and using replication features for disaster recovery purposes. Note that some features apply only to individual platforms and may not be applicable to your configuration.

Virtual Storage Platform G400, G600, G800 and Virtual Storage Platform F400, F600, F800 storage systems can be configured with NAS modules to deliver native NAS functionality in a unified storage platform. The term 'NAS module' in this document also applies to VSP F series, VSP G series, and VSP N series. The unified VSP Gx00 models, VSP Fx00 models, and VSP N series models automatically form a two-node cluster in a single chassis upon installation, with no external cabling required.

Related Documentation

Release Notes provide the most up-to-date information about the system, including new feature summaries, upgrade instructions, and fixed and known defects.

Command Line References

The Command Line Reference provides information on the commands used to manage your system, and includes relevant information on the operation of your hardware and software. Depending on the model of your server or cluster node, refer to the Command Line Reference that is appropriate for your system.

- *NAS Module Server Command Line Reference*
- *Command Line Reference for models 4060, 4080, and 4100*
- *Command Line Reference for models 3080, 3090 and 4040*
- *Command Line Reference for models 5200 and 5300*

Administration Guides

- *System Access Guide* (MK-92HNAS014)—Explains how to log in to the system, provides information about accessing the NAS server/cluster CLI and the SMU CLI, and provides information about the documentation, help, and search capabilities available in the system.
- *Server and Cluster Administration Guide* (MK-92HNAS010)—Provides information about administering servers, clusters, and server farms. Includes information about licensing, name spaces, upgrading software, monitoring servers and clusters, and backing up and restoring configurations.
- *Storage System User Administration Guide* (MK-92HNAS013)—Explains user management, including the different types of system administrator, their roles, and how to create and manage these users.
- *Network Administration Guide* (MK-92HNAS008)—Provides information about the server's network usage, and explains how to configure network interfaces, IP addressing, name and directory services.
- *File Services Administration Guide* (MK-92HNAS006)—Explains about file system formats, and provides information about creating and managing file systems, and enabling and configuring file services (file service protocols).
- *Data Migrator Administration Guide* (MK-92HNAS005) —Provides information about the Data Migrator feature, including how to set up migration policies and schedules.
- *Storage Subsystem Administration Guide* (MK-92HNAS012)—Provides information about managing the supported storage subsystems (RAID arrays) attached to the server/cluster. Includes information about tiered storage, storage pools, system drives (SDs), SD groups, and other storage device related configuration and management features and functions.
- *Snapshot Administration Guide* (MK-92HNAS011)—Provides information about configuring the server to take and manage snapshots.
- *Replication and Disaster Recovery Administration Guide* (MK-92HNAS009)—Provides information about replicating data using file-based replication and object-based replication, provides information on setting up replication policies and schedules, and using replication features for disaster recovery purposes.
- *Antivirus Administration Guide* (MK-92HNAS004)—Describes the supported antivirus engines, provides information about how to enable them, and how to configure the system to use them.
- *Backup Administration Guide* (MK-92HNAS007)—Provides information about configuring the server to work with NDMP, and making and managing NDMP backups.



Note: For a complete list of Hitachi NAS open source software copyrights and licenses, see the *System Access Guide*.

Hardware References

- *Hitachi NAS Platform 3080 and 3090 G2 Hardware Reference* (MK-92HNAS017) — Provides an overview of the second-generation server hardware, describes how to resolve any problems, and replace potentially faulty parts.
- *Hitachi NAS Platform and Hitachi Unified Storage Series 4000 Hardware Reference* (MK-92HNAS030)—Provides an overview of the Hitachi NAS Platform Series 4000 server hardware, describes how to resolve any problems, and how to replace potentially faulty components
- *Hitachi NAS 5000 Series Hardware Reference* MK-92HNAS089—Provides an overview of the Hitachi NAS Platform Series 5000 server hardware, describes how to resolve any problems, and how to replace potentially faulty components.
- *Hitachi NAS Platform System Manager Unit (SMU) Hardware Reference* (MK-92HNAS065)—This document describes the usage and replacement instructions for the SMU 300/400.

Best Practices

- *Hitachi NAS Platform Best Practices Guide for NFS with VMware vSphere* (MK-92HNAS028)—This document covers best practices specific to using VMware vSphere with the Hitachi NAS platform.
- *Hitachi NAS Platform Deduplication Best Practice* (MK-92HNAS031)—This document provides best practices and guidelines for using deduplication.
- *Hitachi NAS Platform Best Practices for Tiered File Systems* (MK-92HNAS038)—This document describes the Hitachi NAS Platform feature that automatically and intelligently separates data and metadata onto different Tiers of storage called Tiered File Systems (TFS).
- *Hitachi NAS Platform Data Migrator to Cloud Best Practices Guide* (MK-92HNAS045)—Data Migrator to Cloud allows files hosted on the HNAS server to be transparently migrated to cloud storage, providing the benefits associated with both local and cloud storage.
- *Best Practices for Hitachi NAS Universal Migrator* (MK-92HNAS047)—The Hitachi NAS Universal Migrator (UM) feature provides customers with a convenient and minimally disruptive method to migrate from their existing NAS system to the Hitachi NAS Platform. The practices and recommendations outlined in this document describe how to best use this feature.
- *Hitachi Data Systems SU 12.x Network File System (NFS) Version 4 Feature Description* (MK-92HNAS056)—This document describes the features of Network File System (NFS) Version 4.
- *Hitachi NAS HDP Best Practices* (MK-92HNAS057)—This document lists frequently asked questions regarding the use of Hitachi Dynamic Provisioning.
- *Hitachi Multi-tenancy Implementation and Best Practice Guide* (MK-92HNAS059)—This document details the best practices for configuring and using Multi-Tenancy and related features, and EVS security.
- *Hitachi NAS Platform HDP Best Practices* (MK-92HNAS063)—This document details the best practices for configuring and using storage pools, related features, and Hitachi Dynamic Provisioning (HDP).

- *Hitachi NAS Platform System Manager Unit (SMU) Hardware Reference* (MK-92HNAS065)—This document describes the usage and replacement instructions for the SMU 300/400.
- *NAS Platform ICC with Cisco Nexus Reference Guide* (MK-92HNAS085)—This document describes how to configure Cisco Nexus series switches for use as an ICC (intra-cluster communication) switch.
- *File System Snapshots Operational Best Practice* (MK-92HNAS068)—This document provides operational guidance on file system snapshots.
- *Virtual Infrastructure Integrator for Hitachi Storage Platforms Operational Best Practice* (MK-92HNAS069)—This document provides operational guidance on Hitachi Virtual Infrastructure Integrator for the HNAS platform.
- *Hitachi NAS Platform Replication Best Practices Guide* (MK-92HNAS070)—This document details the best practices for configuring and using HNAS Replication and related features.
- *Hitachi NAS Platform to Hitachi Virtual Storage Platform Unified Gx00 Models Migration Guide* (MK-92HNAS075)—This best practice guide describes how to perform a data-in-place migration of the Hitachi NAS Platform and Virtual Storage Platform (VSP) Gx00 File solution to the VSP Gx00 platform.

Accessing product documentation

Product user documentation is available on the Hitachi Vantara Support Website: <https://knowledge.hitachivantara.com/Documents>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

The [Hitachi Vantara Support Website](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to the Hitachi Vantara Support Website for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](https://community.hitachivantara.com) is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Comments

Please send comments to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Chapter 1: File and Object data replication

File and Object data replication allows you to copy or relocate both file data and file system metadata. Storage servers provide manual and automatic mechanisms for data replication, supporting the replication of data and, when using the transfer of primary access feature, also supporting the replication of file system settings. When using replication with the transfer of primary access feature, you can relocate file system data and CNS links, SMB shares, permissions and all other file-level metadata. Administrators can use NAS Manager to configure policy-based replication jobs independently from other backup strategies.

Replication is a licensed feature, and the *Replication* license must be installed before replications can be performed. Refer to the *Storage Subsystem Administration Guide* for more information about licenses.

File replication and object replication

There are two basic methods used to replicate file system contents (data and metadata): file-level replication and object-level replication.

- File-level replication operates by copying file system structures, such as files, directories, and the metadata for those structures.

In file-based replication operations, to determine which files or directories to replicate, the metadata for the files and directories must be retrieved (often from disk) and examined, a process that is resource intensive.

- Object-level replication operates by copying the objects that make up the files, directories, and metadata for the files and directories in the file system. Object replication replicates these objects natively, regardless of which file or directory that they may belong to, negating the need to assemble all of the objects associated with a file or directory before transfer, making the overall transfer more efficient.



Note: Object replication operates only on file systems, not on individual directories or files.

Both replication methods can be used with policies and schedules to automate data replication based on criteria you specify, and both replication methods can be initiated and managed manually, through the NAS Manager or the CLI.

A *Replication* license is required to use either file replication or object replication, and a single *Replication* license enables both features.



Caution: Care should be taken when configuring systems with a single migration destination for both replication source and target (known as a triangular arrangement). Such arrangements should not be considered a valid solution in any disaster recovery (DR) or backup scenario, as there is only a single copy of the user data pointed to by XVLs at each end of the replication policy.

File replication overview

File-level replication supports incremental data replication and multiple replication streams.

Incremental file replication

The server supports incremental data replication, performed under control of the replication process. Incremental replication means that, after the initial copy, only changes in the source volume or directory are actually replicated on the target. Snapshots ensure the consistency of the replication.



Note: If the snapshot that was copied by the last successful replication copy is deleted, an incremental copy cannot be performed, so the full data set is replicated.

Incremental data replication uses the same data management engine as NDMP to copy:

- The contents of an entire file system,
- A virtual volume, or
- An individual directory tree to a replication target.

Upon configuration of a replication policy and schedule, the incremental data replication process takes place automatically at the specified interval. The replicated data can be left in place (and used as a standby data repository). In addition, the replicated file system or directory can be backed up through NDMP to a tape library system for long-term storage (which can also be automated).

Incremental data replication supports the following targets for replication:

- A file system or directory within the same server.
Tiered storage technology ensures that replications taking place within a server are performed efficiently, without tying up network resources.
- A file system, virtual volume, or directory on another server.
- A file system, virtual volume, or directory on another server model.

Although the replication process schedules and starts all replications, replicated data flows directly from source to target without passing through the replication process.

Incremental block replication

For file replication only, to replicate large files more efficiently, the server also supports incremental block replication. With incremental block replication, only the changes in files are replicated and not the whole file, thus reducing the amount of data replicated and the overall replication time.



Note: Block-level replication copies the entire file if the file has multiple hard links.



Note: The Block-Level Replication feature is automatically enabled if the Replication license is present.

Multiple stream replication

Multiple replication streams are created by adding TCP connections between the source and target systems of a replication or ADC copy operation. Each additional connection is used for an additional data stream by the replication application.

Multi-stream replication helps to alleviate some latency problems found with single-stream replication by running multiple independent streams in parallel. When latency from sequentially executed functions limits performance, multiple independent streams can produce a significant performance improvement.

Multi-stream replication is only supported if both the source and destination systems are using software release 6.1 or later.

For policy-based replication operations, multi-stream replication is controlled using the replication Add Rule or Modify Rule pages of the NAS Manager.

For ADC copies, multi-stream support is enabled by setting the number of additional connections requested as the value of the environment variable `NDMP_BLUEARC_MULTI_CONNECTION` (refer to the *Backup Administration Guide* for more information).

Note the following:

- When using software release 6.1 or later, and using multi-stream replication or embedded inline hard linked files, if a replication fails part way through, it will not be possible to restart replication if the server is downgraded to an earlier release. Refer to the *Backup Administration Guide* for more information about NDMP support for embedded hard links.
- Multi-stream replication features are not enabled using the `ndmp-option` CLI command; instead, the invoking NDMP command must request multiple streams.

For policy based replications the multi-stream feature is configured using replication rules.

For individual ADC copies, multiple streams are specified by adding the `NDMP_BLUEARC_MULTI_CONNECTION` environment variable to the ADC script (refer to the *Backup Administration Guide* for more information).

- NDMP has two ways of copying data from files with hard links. The `NDMP_BLUEARC_EMBEDDED_HARDLINKS` environment variable controls this behavior (refer to the *Backup Administration Guide*) for more information on this variable.



Note: When multiple connections/streams are used, the data from files with hard links is embedded within the hierarchical path data, regardless of the setting of the `NDMP_BLUEARC_EMBEDDED_HARDLINKS` variable.

Relocating file systems

Storage servers support relocation of file systems, or parts of file systems, *including both file system data and file system metadata* from one server to another. Metadata refers, for example, to CNS links, SMB shares, NFS mount points, FTP users, Snapshot rules, backup files, and other file system-level settings.



Note: Unlike other file system metadata, iSCSI configuration settings remain with the original EVS, as an iSCSI target may contain Logical Units (LUs) from multiple file systems. In this instance, the **Relocation** page displays a message, reminding the Administrator to properly configure iSCSI Domains, Targets, iSNS, and Initiator Authentications on the new EVS.

Allowable destinations for a relocation may be:

- Another EVS on the same cluster node,
- Another node in the cluster, or
- An EVS on another server or cluster.

The following list includes some examples of file system relocations:

- Moving data to a new storage system.
- Dividing a single large file system into several smaller file systems within a storage pool.

- Load balancing, by moving data from one file system to another, or by moving a file system from one EVS to another.
- Moving an EVS (and all its file systems) to another server to gain access to other storage devices or to change the structure of the data.

From a high level, relocating file systems requires two steps:

1. Replicate online data while the system is live and in normal use. This may require several incremental replications, to synchronize the data on the source and the target as much as possible. Synchronizing the data shortens the amount of time required for the next step.
2. Perform a final replication with source data (file system) in Syslocked mode. When in *Syslocked mode*, the data is write-protected, so the data can be accessed and read, but data cannot be changed or added. At the end of this stage, the target is brought online in place of the source. For more information on Syslock mode, refer to the *File Services Administration Guide*.

Object replication overview

Object-based replication operations are based on snapshots.

The first time a replication is performed, a snapshot is taken (the initial snapshot), and the first replication operation replicates all objects on the source to the target. All following (incremental) replications take a snapshot of the changes to the file system and replicate only the objects that have changed.

Detecting and copying objects from a source to a target requires fewer system resources than detecting files and directories (which include directory structures and metadata). Object-level replications detect and replicate only those objects that have changed on the source file system, thereby using minimal system resources. Object replication is the fastest method for performing replications using the NAS server storage system.

In an object replication, a snapshot of a file system is replicated to another server, typically remote, to provide backup and recovery of the source data. The replicated files are immediately available for use in a disaster recovery situation. Additionally, the roles of the source and target servers can be reversed, allowing the target server to quickly take over the responsibilities of the source server.

Object-level replication has the following benefits:

- Higher performance than file-based replication. The greatest performance improvements are seen with incremental replication, especially dense file systems (many small files) or those file systems with a high rate of change. Larger file systems will achieve even greater improvements in incremental replication performance than smaller file systems.
- Object replication enables the ability to quickly failover in the event of a disaster.
- Object replication maintains the replication status on both the source and target file systems. If the replication relationship is broken, such as during a system shut-down or move, when the relationship is re-established, incremental replication can continue, rather than requiring a full re-sync of the file system.

Object-level replication has the following limitations:

- Object replication works at the file system level only; entire file systems may be replicated using object replication, but individual files or directories cannot.
- The target for object replication is read-only by file-serving protocols.
- During disaster recovery failover, target file systems are not accessible until promoted to primary. As the file system is being replicated as its constituent objects, the file system is in an inconsistent state until all objects have been replicated.
- CNS tree structures are not replicated; they must be manually replicated on the target system if CNS is used with object replication.

Incremental replication

Storage servers can also perform incremental data replication, which works as follows:

- Upon establishing a replication policy, the replication process performs an *initial copy* of the source file system (or directory) to a destination / replication target file system.
- Once a successful *initial copy* has occurred, the system performs *incremental copies* (replications) at scheduled intervals. During an *incremental data file replication*, the system copies to the target, in full, those files that have been changed since the last scheduled replication. During an *incremental data object replication*, the system copies to the target, in full, those objects that have been changed since the last scheduled replication.
- For file replication only, to replicate large files more efficiently, the server also supports incremental block level replication. With incremental block-level replication, only the changes in files are replicated and not the whole file, thus reducing the amount of data replicated and the overall replication time.

Transferring XVLs as links during object replication

When using object replication, the default behaviour is to 'rehydrate' data that has been migrated using either external migration or the Data Migrator to Cloud (DM2C) feature.

Files that have been converted to External Volume Links (XVLs) are copied in full to the replication target. This requires system administrators to ensure that there is extra space on the replication target. This is because the source filesystem only requires enough space to store a link to the externally stored data, whereas the replication target filesystem requires space for all the data.

If required, it is possible to copy XVLs to the replication target as links using the **Transfer XVLs as links during object replication** option and therefore remove the additional space overhead on the target system. This functionality can be enabled using the `replication-xvl-as-links` CLI command or by selecting the option on the **File System Details** page on the NAS Manager.



Note: To use this functionality, the NAS server containing the replication target file system must be able to access the migration destination, or a synchronised replica of it, in the same way as the source file system.

Considerations when transferring XVLs as links

- Virtual Volume level migration paths - you can configure external migration paths for a Virtual Volume on a replication source file system. It is not possible to configure external migration paths on a replication target filesystem, because it is not possible to apply a path to a Virtual Volume. To transfer XVLs as links on this type of replication source filesystem, configure a file system level migration path on the replication target with the same name as the Virtual Volume level path on the replication source.
- This option can be enabled on a filesystem that has existing migration and replication policies, but it only applies to migrated files that are replicated after the option is enabled. When the option is enabled on an existing replication source filesystem that has a cloud or external migration policy, you can re-format the replication targets to reclaim the space on the replication target that was previously occupied by rehydrated files. However, this requires a time-consuming full replication.
- Once the option is enabled on the source, it cannot be disabled. This is to prevent the possibility of a broken configuration.

Using migration targets with replicated XVLs

There are two main methods of configuring migration targets for use with replicated XVLs:

- Single migration destination for replication source and target
- Separate migration destinations for replication source and target

Single migration destination for replication source and target

This solution requires only one migration target. It is most appropriate to use this solution when the single migration target is an external cloud entity provided by a third-party cloud provider.



Note: As there is only one instance of the migrated data, this arrangement is not suitable for use as a backup or disaster recovery solution (as it does not represent a 'complete' backup).

At any time, there is only one 'correct' view of the data - the replication source live file system. Migrated and non-migrated data in other locations may be inconsistent with that view. This is because the migrated data is modified immediately, but the non-migrated data is modified only when the replication has taken place. This is of concern when migrated files are deleted or reverse-migrated (either explicitly or as a result of being auto-recalled). When the **Transfer XVLs as links during object replication** option is not enabled, the deletion or reverse migration of a migrated file would result in the data for that file on the migrated target being deleted.

As the data that has been deleted is now removed from the single migration target, the link to that data from the replication target(s) is now broken. To prevent this, two new CLI command options are available:


```
set delete-target-on-unlink-even-if-xvl-as-links
```

```
set delete-target-on-reverse-migration-even-if-xvl-as-links
```


These settings apply on the cluster hosting the replication source filesystem, but they can be set on all clusters hosting replication target filesystems in anticipation of one of the replication targets being promoted to read/write at a future date. The default value for these settings is false.

When the **Transfer XVLs as links during object replication** option is enabled for a filesystem and the above options are set to false, a deletion or reverse migration does not cause the migrated data to be deleted (or a retention policy to be set on HCP if configured), and the access to the migrated data at the replication target(s) is not broken.

After protected reverse migration



These settings preserve access to data at replication targets, but the data for reverse migrated or deleted files is preserved forever on the migration target unless the deleted files are manually removed from it.

After protected reverse migration and subsequent replication



As a consequence, with this setting, even when all files have been reverse migrated, the migration path is reported as in use.

If there is a requirement to promptly delete data for storage or regulatory reasons, the options should be set to true, but with the expectation that access to the migrated data from replication targets is imperfect until it is manually cleaned up.

For non-cloud externally migrated data, a list of migrated files for a filesystem can be generated using the **Report Migrated Files** migration type in the schedule configuration. This can then be compared to the data on the migration target to determine which files can be safely removed. This option does not exist for Data Migrator to Cloud.

Separate migration destinations for replication source and target

This solution requires a different migration target for each replication target filesystem.

It also requires a method of copying the migrated data between the migration targets. As there is more than one copy of the migrated data, this arrangement is suitable for use in a disaster recovery or backup scenario and is appropriate for internally provided cloud targets such as HCP. As there is no risk of deleting the link targets relied upon by the replication target filesystems, there is no requirement to protect the link targets from deletion or reverse migration.

If the replication has been configured in this way then the following CLI command options should be set to true to avoid having to manually remove the non-deleted link targets:

```
set delete-target-on-unlink-even-if-xvl-as-links true
```

```
set delete-target-on-reverse-migration-even-if-xvl-as-links true
```

These settings apply on the cluster hosting the replication source filesystem, but they can be set on all clusters hosting replication target filesystems in anticipation of one of the replication targets being promoted to read/write at a future date. The default value for these settings is false.

A time window still exists when the replication target and its migration target may be inconsistent, when one or other of the replications (either the NAS object replication or the copy of data between migration targets) has completed and the other has not. Scheduling of the replications should be configured to reduce this window to a minimum and minimize the opportunity for client access during the window.

Configuring transfer XVLs as links during object replication

You can enable this option using the CLI or the NAS Manager.



Note: Both source and target systems involved in the replication relationship must be running NAS server release v13.4 or later.

You can enable the **Transfer XVLs as links during object replication** option as follows:

- Run the `replication-xvl-as-links` CLI command for each replication source file system

or

- Select the enable button for the **Transfer XVLs as links during object replication** option on the **File System Details** page on the NAS Manager.



Note: The replication source file system must be unmounted for this option to be available.

File System Details

Settings/Status

Label:

Capacity

53% Total Used Space

Capacity: 3.97 GiB (47%)

Free: 1.85 GiB (47%)

Total Used: 2.12 GiB (53%)

Expansion Limit: -

Live file system: 2.12 GiB (53%)

Snapshots: 0 Bytes (0%)

Legend: ■ Live file system ■ Snapshots ■ Usage Warning ■ Usage Severe

Configuration

Status: Mounted

Deduplication: [Not converted](#)

Thin Provisioning: Disabled

EV: [HNAS-EVS1 \(Online\)](#)

Security Mode: [Mixed \(Windows and Unix\) \(Inherited\)](#)

Block Size: 4 KiB

Read Cache: No

WFS Version: WFS-2

Syslock: Disabled

Object Replication Target: Disabled

Transfer Access Points During Object Replication: Disabled

Transfer XVLs as Links During Object Replication: Disabled

Configuring the migration path for the replication target file system

To enable access to the migrated data on the replication target filesystem, configure a migration path using the CLI or the NAS Manager (see the Data Migrator Administration Guide for details).

Migration paths are configured on a replication target filesystem in the same way as they are configured on a normal file system except for the following constraints:

- The migration path must have the same logical name as that on the replication source filesystem. For cloud paths, the logical name of the migration path is the name of the cloud destination in use by the path, and the cloud destination for the replication target must also have the same name as the one on the replication source. The cloud destination on the replication target also needs to have the same UUID as one on the replication source - creating a cloud destination with a specific UUID is only possible using the `migration-cloud-destination-create` CLI command and not through the NAS Manager.
- The cloud account on which the replication target's cloud destination depends has no constraints but must exist.
- It is not possible to configure multiple cloud destinations on the same cluster with the same name and/or UUID. If the replication source and target file systems are both on the same cluster, they must share the same migration target and cannot be configured to use different ones.

Example when source and target are on different clusters

Consider a replication source file system that has a cloud migration path configured as follows:

- Migration destination name: `repln_src_mign_dest`
- Migration destination UUID: `a5de129e-8f75-11d3-9010-18de0fceb7d`
- Migration path name (same as the destination name): `repln_src_mign_dest`

To create the migration path for the replication target:

- Create a cloud migration account on the cluster that hosts the replication target file system using either the NAS Manager or the `migration-cloud-account-create` CLI command.
- Create a cloud destination with the name `repln_src_mign_dest` and UUID `a5de129e-8f75-11d3-9010-18de0fceb7d` using the `migration-cloud-destination-create` CLI command.
- Create a cloud path for the replication target file system to the destination configured above using either the NAS Manager or the `migration-add-external-path` CLI command.




Note: Other details of the migration, for example, account, destination or path, can be different or the same as the replication source but the names and UUIDS must match.

Recovering a file system

To recover a file system from a snapshot:

Procedure

1. Navigate to **Data Protection > File System Versions** to display the **File System Versions** page.



The following table describes the fields in this page:

Field/Item	Description
File System Details	This section displays the name of the EVS hosting the file system, and the currently selected file system that can be recovered from the snapshots listed in the Versions section.
EVS/File System	Displays the name of the currently selected EVS and file system. Click change to select a different file system.
Status	Displays the current mount status of the file system. The file system status may be Syslocked, Checking, Unmounted, Mounted, Mounted as read only or Mounted as a replication target.
Object Replication Details For Latest Version	If the file system is a replication target, this section displays the status of the most recent replication and information about the replication source. If the currently selected file system is not a replication target, or the replication information cannot be retrieved (if the source server is not known to the replication process), the Source File System, Source Server, and Source File System Status fields are not displayed.
Status	If the currently selected file system is a replication target, this field displays a status indicator and a message about the most recent replication. If the file system is not a replication target, the status indicator is greyed out and the message reads "Latest version is not an object replication target". The status indicator is green if the most recent object replication completed successfully. If the first full replication is currently running, this status indicator is grey. If the currently selected file system is not involved in a replication policy (neither the source or the target), then the message reads "No status found".

Field/Item	Description
	If a replication associated with this file system has not yet run, the light is grayed out, and the message reads, "No status found".
Source File System	If the currently selected file system is a replication target, this field displays the name of the source EVS and file system. If the currently selected file system is not a replication target, this field is not displayed.
Source Server	If the currently selected file system is a replication target, this field displays the name of the server hosting the EVS/replication source file system. If the currently selected file system is not a replication target, this field is not displayed.
Source File System Status	If the currently selected file system is a replication target, this field displays the current mount status of the replication source file system. The file system status may be Not mounted, Mounted, Mounted as read-only, Syslocked, or Checking. If the currently selected file system is not a replication target, this field is not displayed.
Versions	This section lists versions of this file system that are in available snapshots, and identifies the snapshot copied to the replication target and the replication source snapshot.
Time of Version	The date and time the object replication policy last ran. "Time" refers to when the snapshot was taken.
Version	Identifies the specific snapshot copied to the replication target.
Replicated From Snapshot	Identifies the replication's source snapshot.
Recover File System To Version	Opens the File System Recovery Selection page, on which you choose the type of file system recovery you want.
Recover Multiple File Systems To Version	Opens the Recover File Systems page, on which you choose options for recovering multiple file systems.
File System Recovery Reports	Opens the file system File System Reports page.
Object Replication	Opens the main Object Replication page.

Field/Item	Description
Policies & Schedules	

2. Click **Recover File System to Version**.
3. Use this page to begin the recovery process and either:
 - Promote the file system to a normal file system (and optionally, mount it as read-write or read-only).
 - Demote the file system to an object replication target (and mount it as an object replication target).

In order to promote or demote a file system, the process rolls back the file system to the last successful replicated snapshot. You must check the available snapshots so that you can choose the most recent (in order to minimise data loss due to promoting a previous version of the file system). Using your browser, go back to the **File System Versions** page and note the time and version of the snapshots, so you can choose the most recent successful replication snapshot.

In a disaster recovery scenario, your primary system will probably be unavailable, so you must access the file system version using the **File System Versions** page of the backup system at the recovery site. Note the time of the replication snapshot and the versions of the source and target. You will use the latest version when promoting the file system to a normal file system.

Recovering and promoting a file system

Procedure

1. From the **File System Versions** page, select the "Recover File System To Version" option and then click "**Promote the file system to a normal file system (and, optionally, mount as read-write or read-only)**" to display the **Recovery File System** page.
2. In step 2, **Recover file system to version created at**, on the **Recover File System** page, verify that the file system version, the snapshot name, and the snapshot source are correct.
3. In step 3, **Promote file system 'name' and**, select one of the following mount options for the recovered file system:
 - **mount read write**
 - **mount read only**
 - **do not mount** (selecting this option disables the "Recover access points" option - no access points are moved)

4. In step 4, **Recover access points**, fill in the check boxes to specify which types of file system access points to recover. Leave the check boxes empty to specify not to recover SMB shares or NFS exports for the recovered file system.

- **shares**
- **exports**

For NFS exports, when promoting a replication target file system on the same EVS as the original source file system, there are two additional options to determine which file system the NFS clients will access after the recovery process:

- **Clients will continue to access source file system**

Enable this option so that the exports that are recovered on the target file system are named `<export name>_<task id>` to avoid any conflicts with `<export name>` on the source. Both file systems will have their own set of exports.

- **Clients will access target file system without interruption**

Enable this option so that exports are moved from the source file system to the target without the need to remount NFS v2/3 clients. The source file system will no longer have the recovered exports assigned to it.

5. Click **next** to display the **Recover File System Confirmation** page.
6. Verify the file system recovery settings, and click **OK** to proceed with the file system recovery, and display the **File System Recovery Report** page.
7. Monitor the file system recovery.

The following table describes the fields in this page:

Field/Item	Description
File System Details	Display the name of the currently selected EVS and file system, along with its status.
Recovery Details	Displays the progress of the current or last run recovery, and status of the recovery.
Progress	Displays the progress of the current or last run recovery, including start and end times.
Request Summary	Displays the recovery options in use, including the name of the snapshot used for rollback and the share/export handling options.
Source File System "Transfer Access Point" Setting	After Object Replication completes, it promotes the file system to a normal file system on the target EVS. That is, Object Replication transfers the access point to the target file system. By default, it uses the source file system.
Recovery Statistics	Displays NFS Exports and SMB Shares import statistics.
abort	Stops the active recovery operation.
View Log	Opens the File System Recovery Report Log page.

Use your browser's back button to return to the previous page, or click **abort** to abort an active recovery operation.

Result

After the file system has been recovered, and is "live," you may want to create a replication policy and schedule to replicate the new primary file system and use the old primary file system as the replication target.

Recovering and demoting a file system


Procedure

1. From the **File System Versions** page, select the "Recover File System To Version" option and then click "**Demote the file system to an Object Replication Target (and mount as an Object Replication Target)**" to display the **Demote File System To Object Replication Target** page.
2. Specify recovery options.



Note: After you have selected all the options on this page, it will take a few minutes for the server to roll back the file system and remove all the access points. You can monitor the server's progress on the **File System Recovery Progress** page.

Field/Item	Description
File System Details	Displays the name of the EVS hosting the file system, and the currently selected file system that can be recovered from the snapshots listed in the Versions section.
EVS/File System	Displays the name of the currently selected EVS and file system.
Status	Displays the current mount status of the file system. The file system status might be unmounted, mounted, or mounted as a replication target.
Object Replication Details For Latest Version	If the file system is a replication target, this section displays the status of the most recent replication and information about the replication source. If the currently selected file system is not a replication target, or the replication information cannot be retrieved (if the source server is not known to the replication process), the Source File System , Source Server , and Source Server fields are not displayed.
Status	<p>If the currently selected file system is a replication target, this field displays a status indicator and a message about the most recent replication. If the file system is not a replication target, the status indicator is greyed out and the message "Not an object replication target" is displayed.</p> <p>The status indicator is green if a replication is currently running, or if the most recent object replication completed successfully.</p> <p>If a replication associated with this file system has not yet run, the light is grayed out, and the message reads, "Not an object replication target".</p>
Source File System	<p>If the currently selected file system is a replication target, this field displays the name of the source EVS and file system.</p> <p>If the currently selected file system is not a replication target, this field is not displayed.</p>
Source Server	<p>If the currently selected file system is a replication target, this field displays the name of the server hosting the EVS/ replication source file system.</p> <p>If the currently selected file system is not a replication target, this field is not displayed.</p>

Field/Item	Description
Source File System Status	<p>If the currently selected file system is a replication target, this field displays the current mount status of the replication source file system. The file system status might be unmounted or mounted.</p> <p>If the currently selected file system is not a replication target, this field is not displayed.</p>
The following steps will be taken	<p>This section lists the recovery steps, and allows you to specify recovery options.</p> <ul style="list-style-type: none"> a. Displays the file system to be recovered. b. Displays the date and time the snapshot to be used to recover the file system was taken. You can use the drop-down list to select a different snapshot by the date and time the snapshot was taken. <p>For the selected snapshot, the name of the corresponding snapshot on the target file system and the name of the snapshot on the source file system are also displayed.</p> <ul style="list-style-type: none"> c. Displays the recovery goal to confirm that the file system you have chosen will be demoted to an object replication target. d. Check boxes allow you to specify if you want the access points (SMB shares and NFS exports) of the demoted file system to be removed. Fill the check boxes of the access points you want to remove. Leave the check boxes empty to keep the access points for the demoted file system. <div style="background-color: #e0f2f1; padding: 10px; border: 1px solid #ccc;"> <p> Note: In general, when demoting a file system as a part of disaster recovery, you should remove both SMB shares and NFS exports.</p> </div>

3. Click **next** to display the **Demote File System to Object Replication Target Confirmation** page.
4. Verify the file system recovery settings and proceed with the file system recovery.
 - Click **back** to return to the **Demote File System To Object Replication Target** page.
 - Click **OK** to begin the recovery, and display the **File System Recovery Report** page.
 - Click **cancel** to return to the **File System Versions** page.
5. Monitor the file system recovery.

Field/Item	Description
File System Details	Display the name of the currently selected EVS and file system, along with its status.
Recovery Details	Displays the progress of the current or last run recovery, and status of the recovery.
Progress	Displays the progress of the current or last run recovery, including start and end times.
Request Summary	Displays the recovery options in use, including the name of the snapshot used for rollback and the share/export handling options.
Recovery Statistics	Displays NFS Exports and SMB Shares import statistics.
abort	Stops the active recovery operation.
View Log	Opens the File System Recovery Report Log page.

Use your browser's **back** button to return to the previous page, or click **abort** to abort an active recovery operation.

Replication and disaster recovery

Object replication provides replication of file systems, including the replication of related access points, such as SMB shares and NFS exports, as well as tools to automate disaster recovery.

Object replication has several important concepts:

- **Primary file system:** The "primary file system" is the file system that network clients access. The primary file system is the "live" file system and the source of the replication.
- **Direction of replication:** When the primary file system from site "A" is replicated to another server at site "B," the direction of the replication determines which file system is designated the "source" and the "target." The primary file system from site "A" is always the replication source. The target file system is the replica on the server at site "B" (which may be located at the same physical site as "A" or at a remote location). Note that the target file system at site "B" may also be used as a replication source to a third site (site "C").
- **Swapping roles:** Moving the network client access from the file system at site "A" to the replicated file system at site "B." Roles are swapped when the replication either stops (for example, if site "A" goes offline for some reason) or the direction of the replication is intentionally reversed (a planned role swap).
- During the role swap, the file system at site "B" is promoted to primary, and the file system at site "A" is demoted. (If site "A" is accessible, the file system at site "A" typically becomes the target file system.)

- As a part of the role swap, access point (SMB share or NFS export) settings are deleted from the server at site "A" and, along with other configuration settings, are applied to the server at site "B" so that network clients now access the primary file system, which is now physically located at site "B." Clients accessing the file system now communicate with the server at site "B" and read from and write to the file system at that site, which has become the primary file system.
- Note that a single server can host many file systems, and could be providing "primary" access to several file systems while other file systems hosted by the same server could be target file systems. Primary access for any file system can be moved independently of any other file system on the same server.

Object replication is most often used in the following situations:

- A planned promotion of the file system at site "B" to primary. In this case, it is possible to ensure that the file systems at sites "A" and "B" are exact replicas (though this would require a period of read-only access at site "A"). If both sites are functional, it may be possible for the server at site "B" to access the server at site "A" to retrieve information such as configuration settings.

In the case of a planned promotion, the administrator puts the primary file system into syslocked-mode, then schedules an incremental replication to the target to ensure both file systems are synchronized. Once the replication is complete, the promotion can proceed, and after the transfer of primary access, the clients access the newly promoted primary file system.

- An unplanned promotion of the file system at site "B" to primary (also known as "disaster recovery". If, for any reason, the primary file system at site "A" becomes inaccessible, the file system at site "B" is promoted (becomes primary). In this case it is unlikely that the file system at site "B" will be an exact copy of the file system at A at the time of the outage (because the replication is asynchronous). Server B must already have access to all the information necessary to function as the primary.

Replication process for planned promotions

In general, the process followed for planned promotions is:

Procedure

1. On the primary file system, create a replication policy to synchronize the primary and target file systems.
2. Create and enable a replication schedule to perform the replication.
3. Synchronize the source and target file systems.
4. Syslock the primary (source) file system and perform the final replication.
5. Swap roles, which promotes the target file system to become the primary file system.
6. Verify that all access points have been created.
7. Redirect network clients to the new primary file system.
8. Verify that clients can access the newly promoted primary file system.
9. Demote the original source file system to become a replication target.
10. If the replication policy schedule was disabled, reactivate it.
11. Verify that the replication runs successfully.

12. Allow user access to the new primary file system.


Recovering multiple file systems



Note: Each of the file systems will be recovered to its most recent version, and each recovered file system will be promoted to a normal file system.

Procedure

1. Navigate to **Home > Data Protection > File System Versions** to display the **File System Versions** page.
2. Click **Recover Multiple File System to Version** to display the **Recover File Systems** page.



3. Select the file systems to recover from the **Available File Systems** list, and click the right arrow to add the file systems to the **Selected File Systems** list.
4. Specify how the recovered file systems are to be mounted:
 - **mount read write**
 - **mount read only**
 - **do not mount** (selecting this option disables the 'Recover access points' options)
5. In step 4, **Recover access points**, select the check boxes to specify which types of file system access points to recover. Leave the check boxes empty to specify not to recover SMB shares or NFS exports for the recovered file systems.
 - **shares**
 - **exports**
6. Click **next** to display the **Recover File Systems Confirmation** page.

7. Verify the file system recovery settings, and click **OK** to proceed with the file system recovery, and display the **File System Recovery Reports** page.
8. Monitor the file system recovery.

Result

After the file system has been recovered, and is "live," you may want to create a replication policy and schedule to replicate the new primary file system and use the old primary file system as the replication target.

Policy-based replication

Policies can be used for both file replications and object replications. Policy-based replication comprises:

- **Replication Policy:** A replication policy identifies the data source, the replication target, and optionally a replication rule. For file replication, pre-replication and post-replication scripts can also be set up in the **Policy** page.
- **Replication Rules:** Optional configuration parameters that allow tuning of file replications to enable and disable specific functions or to optimize performance.
- **Replication Schedule:** Defines all aspects of automated timing.

Replication schedules

After a replication policy has been defined, it must be scheduled to run. Replications can be scheduled and rescheduled at any time and with any of the available scheduling options.

Replication schedules overview:

- **Periodic replication:** Replications occur at preset times. Periodic replications can be set up to run daily, weekly, monthly or at intervals specified in numbers of hours or days.
- **Continuous replication:** When a replication policy specifies continuous replication, as soon as the replication job completes, the same replication job starts again.
- **One time replication:** A new replication job starts after the previous job has ended. The new replication job can start immediately or after a specified number of hours.

When planning replication schedules, customer support recommends scheduling during off-peak times such as nights or weekends. After a replication has started, additional replications for the same policy cannot start until the current replication has completed; however, multiple concurrent replications are allowed for replications by different policies.



Note: For file replication only, when the replication operation begins, the destination file system should be placed into syslock mode. If the destination file system is not in syslock mode during a replication operation, clients may write to the file system, creating inconsistencies between the source and target of the replication. When scheduling file replications, you should consider this limitation.

Supported replication applications

In addition to the built-in replication tools, Hitachi NAS Platform supports Hitachi Vantara replication applications (when used with Hitachi Vantara storage subsystems).

- **TrueCopy Synchronous** provides synchronous data replication for disaster recovery or data migration. TrueCopy Synchronous software is a continuous, nondisruptive, host independent remote data replication solution for use between Hitachi Vantara storage subsystems within a data center or within the same metropolitan area.
- **ShadowImage** provides a nondisruptive, host-independent data replication solution for copying data within a single Hitachi Vantara storage subsystem. The original data, and each copy of the data, remain RAID-protected by the storage subsystem to ensure the availability of the data.

TrueCopy and ShadowImage considerations

When using TrueCopy and ShadowImage, keep the following in mind:

- TrueCopy and ShadowImage functionality are managed using the Hitachi Vantara software interfaces, because the replication occurs between or within the Hitachi Vantara storage subsystems. Contact your Hitachi Vantara technical representative for assistance configuring these features. For more information about these applications, contact Hitachi Vantara Support Center or your Hitachi Vantara technical representative.
- Storage servers rely on SCSI commands to determine that a system drive is simplex, primary mirror, secondary mirror, TrueCopy, or ShadowImage. The information is obtained using a proprietary extension that Hitachi Vantara has added to the standard SCSI inquiry.
- Hitachi Vantara storage subsystems support the following volume states:
 - Unmirrored (simplex)
 - Mirrored primary (p-vol)
 - Mirrored primary (p-vol)
 - Unknown

- Neither the storage server nor the TrueCopy application can automatically cause failover to secondary storage. An external agent (a person or application) must make the decision on the failover and execute the required commands. Once the storage failover is initiated, the storage server automatically starts accessing the new primary storage.

An external utility controls TrueCopy can be used to issue commands to the storage server. This functionality must be provided by an external utility that checks the health of the system and makes appropriate decisions. This same utility can be used to initiate an EVS migration in conjunction with the storage failover.

Switching mirror roles involves a brief changeover period when the system drives do not have a normal primary/secondary relationship. During that period, I/O is impossible. Therefore, the storage server unmounts the file systems on the storage pool. Changing mirror roles always involves unmounting and remounting file systems, and, if possible, users should unmount the file systems before the change and remount after the changeover is complete.

Note that, in the case of automated storage failover, steps must be taken to ensure that the original primary volume does not come back online after a failover, because the storage server would then see two different primary volumes for the same data.

- Hitachi Vantara requires the usage of ShadowImage for synchronous TrueCopy configurations to help recover in the case of data corruption in the file system. With this configuration, the storage server does not have to be configured to see the ShadowImage volume (in other words, the ShadowImage volumes should remain unlicensed). For recovery, data can be copied back to the TrueCopy primary mirror, leaving the ShadowImage copy intact until it is determined that the file system is back up and running.
- Hitachi NAS Platforms and High-performance NAS Platforms support only the “Never” mirror fence level of TrueCopy.
- Data Migrator can be used with TrueCopy, as the FS IDs in the mirrored file systems are identical.
- Whenever a storage server (or cluster) shares a Hitachi Vantara storage subsystem with another storage server (or another cluster), the system configuration must use Hitachi Vantara host groups to prevent the storage servers from seeing system drives that they are not intended to use. Host groups are a LUN-mapping mechanism that controls which servers see the different system drives and LUNs. This type of configuration prevents potential problems caused by the storage server periodically sending commands to unlicensed Hitachi Vantara storage systems to determine their status.

Chapter 2: Using file replication

File replication provides a mechanism, manual or automatic, for copying or relocating both file data and file system metadata. Hitachi NAS Platforms support replication of data and, when using the transfer of primary access feature, of file system settings. When using replication with the transfer of primary access feature, you can relocate file system data and CNS links, SMB shares, permissions and all other file-level metadata. Administrators can use NAS Manager to configure policy-based replication jobs independently from other backup strategies.

This section provides a deeper conceptual understanding of the components of data replication and instructions for configuring and implementing replication.

Configuring policy-based file replication

Before you begin


Before administrators can add a replication policy, the type of server that will be used for storing the replicated data must be determined. You can choose from one of the following policy destination types:

- **Managed Server:** For a server to be considered as managed server, it needs to be entered in the replication process configuration.
- **Not a Managed Server:** A non-managed server is one where the IP Address and user name/password of the server is not known by the replication process. Administrators can still select a non-managed server as the target by specifying the IP address along with the user name and password

To configure policy-based data replication:

Procedure

1. Navigate to **Home > Data Protection > File Replication**, and click add to display the **Policy Destination Type** page.



The screenshot shows a web interface for configuring a policy destination type. The title is "Policy Destination Type". Below the title is a form with the heading "Choose the type of Destination Server". There are two radio button options: "Managed Server" (which is selected) and "Not a Managed Server". At the bottom of the form are "next" and "cancel" buttons.

2. Select the policy destination type:
 - Select **Managed Server** to create a policy to replicate to a server that is managed by the replication process.
 - Select **Not a Managed Server** to create a policy to replicate to a server that is not managed by the replication process.
3. Click **next** to display a destination type-specific **Add Policy** page.
The **Add Policy** page for a managed server replication destination displays as:

The **Add Policy** page for an unmanaged server replication destination is similar, with only the fields in the Destination section being different:



Note: Administrators should be authorized to use the external server to access and store replication data.

4. Enter the requested information:

Field/Item	Description
Identification	<p>Name: Allows you to specify the name of the replication policy. The name may not contain spaces or any of the following characters: \<>"' !@#\$%^&* () {} [] +=?:; , ~ ` . '</p>
Source	<p>Source of the replication. Set this field only if you want to make a simple copy of a specific snapshot. Do not set this field if you are intending to run incremental replications. The source is identified using the following fields:</p> <ul style="list-style-type: none"> ▪ Server: Name of the server/cluster that has the source file system for this replication policy. ▪ EVS/file system: Name of the EVS and file system to which the replication source is mapped. Click change to change the EVS or file system. ▪ Path: Select a virtual volume from the drop-down list. Or select Directory and enter the path. ▪ Snapshot: Select a snapshot to migrate a file system from a snapshot taken at a specific point in time. Using a snapshot as a source allows you to replicate the snapshot rather than the live file system, eliminating the possibility of file changes during the replication.
Destination (for managed replication destinations)	<p>Destination of the replication (managed server):</p> <ul style="list-style-type: none"> ▪ Server: Name of the server/cluster that hosts the destination file system for this replication policy. ▪ EVS/file system: Name of the virtual server and file system to which the replication is mapped. Click change to change the EVS/file system.

Field/Item	Description
	<ul style="list-style-type: none"> ▪ Path: Specify the directory path. Note that you may not specify a virtual volume as a path. ▪ Current Syslock status: Indicates if the file system is in Syslocked mode. When System Lock is enabled for the destination file system, a warning icon is displayed. NDMP has full access to the file system and can write to the syslocked file system during a backup or replication, but the file system remains in read-only mode to clients using the file service protocols (NFS, SMB, FTP, and iSCSI). If the destination file system is not in syslock mode during a replication operation, clients may write to the file system, creating inconsistencies between the source and target of the replication. During transfer of primary access operations, both the source file system and the destination file system are put into System Lock mode. To manually enable or disable the Syslock mode for a file system, you must navigate to the File System Details page for the file system. For more information on Syslocked mode, see the <i>File Services Administration Guide</i>.
Destination (for unmanaged replication destinations)	Destination of the replication (non-managed server): <ul style="list-style-type: none"> ▪ File Serving IP Address / Host Name: Name of the server containing the target EVS/ file system. Click change to change the destination to a different server. ▪ File System: Name of the file system to which the replication is mapped. Click change to change the file system. ▪ Path: Specify the directory path. Note that you may not specify a virtual volume as a path. ▪ NDMP User Name: Name of the NDMP user for which the replication target was created. ▪ NDMP User Password: Password for the selected NDMP user.
Processing Options	<ul style="list-style-type: none"> ▪ Source Snapshot Rule Name: The snapshot rule for replication of the source file system. ▪ Destination Snapshot Rule Name: The snapshot rule to use for the snapshot of the destination file system following a successful replication.

Field/Item	Description
	<ul style="list-style-type: none"> Pre-/Post-Replication Script: A user-defined script to run before or after each replication. Scripts must be located in <code>/opt/smu/adc_replic/final_scripts</code>. The permissions of the scripts must be set to "executable".
File Replication Rule	Optional configuration parameters that allow tuning of replications to enable and disable specific functions or to optimize performance.

- Verify your settings, click **OK** to save, or **cancel** to decline.

Specifying IP addresses for an NDMP replication policy manually

You can manually specify the IP addresses to be used by the source and destination servers for the control and data connections of a replication policy.

A file replication policy requires the Service EVS IPs to be able to route data to each other. To specify exactly which EVS IP addresses to use, you can create a per-policy address configuration file.



Note: You must have root access on the SMU to create the file.

The configuration file must be placed on the SMU in the following location:

`/opt/smu/adc_replic/conf/replication_policies`

Each of the policies has its own sub-directory. The configuration text file must be located in the parent directory above and called **<policy name>_addresses**. The policy names must match exactly.



Note: Create this file **before** creating the policy.

The file contents consist of IP addresses and port numbers as follows:

```
SRC_CTRL_ADDRESS=IP-Address
DEST_CTRL_ADDRESS=IP-Address
SRC_DATA_ADDRESS=IP-Address
DEST_DATA_ADDRESS=IP-Address
DEST_DATA_PORT=Port-Number
```

IPv4 addresses should be specified in conventional dot notation: X.X.X.X

IPv6 address should be specified in conventional notation, for example:

```
ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
```

You do not need to specify all values. However, if you specify one DATA address, you must also specify the other DATA address and they must have a common IP family type (both IPv4 or both IPv6).



Note: The CLI command `ndmp-option data_port_range` does not apply when using a policy address file.

Connection errors

When attempting to add a new replication policy, a connection error may be indicated by “Unable to connect to <IP address>” or “Error accessing <source/destination> server”.

The “Unable to connect to” message means one of the following:

- The server is not currently powered up or is temporarily disconnected from the network. The server must be available and properly connected when creating a replication policy.
- The NDMP service may be disabled. The replication uses the NDMP service which must be enabled when adding or running replications. Please use the NDMP configuration page (or the `ndmp-status` command) to enable and start the NDMP service.
- The gigabit Ethernet port providing access to the EVS which hosts the file system is not accessible from the NAS Manager. This may be the case if the network is set up with private subnetworks as commonly used with VLANs. In this case, the server may have been configured so that NAS Manager access is through the management ports instead of the ports set using the `ndmp-management-ports-set` command.

The “Error accessing server” message may occur as a result of restricting NDMP access using the `ndmp-option` command. The `allowip` and `blockip` options can be set such that the NAS Manager is not allowed to access the NDMP services using the standard routes. If the NDMP connection restrictions are definitely required, change the configuration of the server to allow NAS Manager access by way of the management ports using the `ndmp-management-ports-set` command. The NAS Manager connections then bypass the `allowip/blockip` checks.

The NAS Manager replication and data migration features use the NDMP service on the NAS server. The NDMP service is usually accessed by way of the IP address of the EVS which hosts the file system, this access usually happens through a gigabit Ethernet port. In some cases, the IP address is within a private subnetwork and is not accessible from the NAS Manager. When this is the case, the `ndmp-management-ports-set` command can be used to request that the NAS Manager access goes through the management ports and is then relayed to the NDMP service.

The `ndmp-management-ports-set` command takes two parameters which are the TCP ports. One is used to accept the incoming connection on the management port and one to pass the requests to the NDMP code. These must be ports that are not in use by any other service. In particular, these ports must not be the standard NDMP service port. The port numbers 10001 and 10002 usually work and, being next to the standard NDMP port 10000, can be useful in identifying the port usage.

Having set up the NDMP management ports this way, all NAS Manager replication and data migration NDMP accesses will be routed by way of the management port. Note that the actual data transfer connections involved are between the NAS server EVSs and do not run over the management connections. In particular, a replication between two NAS servers passes the data over a TCP connection between EVS IP addresses through the gigabit Ethernet ports. Therefore, the two EVSs must have IP addresses that can communicate with each other.

Understanding snapshot rules

By default, replications automatically create and delete the snapshots they require to complete consistent copies. That being the case, snapshot rules are not usually required. However, there are cases where the snapshots must be taken or used by external software. In these cases, snapshot rules are used so that the external software and the replication can be sure they are using the same snapshot.



Note: Snapshot creation is normally synchronized with a specific event. The snapshot is explicitly created at this time, so the snapshot rule should not have an associated snapshot schedule.

Specific instances where snapshot rules may be used include:

- Replications which copy databases or iSCSI LUNs. A snapshot taken automatically at the start of a replication will not capture a consistent image of a database or an iSCSI LUN that is actively in use. In order to capture a consistent image, the database/iSCSI LUN needs to be brought into a quiescent state before the snapshot is taken. These actions are normally be executed by a script, which then takes a snapshot within the snapshot rule so that the replication can identify which snapshot to copy.

The script could be invoked as part of a pre-replication script. Alternatively the script could be independently scheduled. If scheduled independently, however, the schedule must allow the script to complete before the replication starts.

- Linked, two-stage replications, which copy a file system from server A to server B and then copy on from server B to server C. These types of replications can use snapshot rules to synchronize the copies.

The replication from server B to server C may start while a copy from server A to server B is running. If a snapshot was taken at this point, an inconsistent file system state would be captured. One way to avoid this is to use a specific snapshot rule as both the destination snapshot rule of the server A to server B copy and the source snapshot rule of the copy from B to C. Then the B to C copy will always copy a snapshot taken at the end of the last complete copy from A to B.

Two kinds of rules define snapshot use during replication:

- **Source Snapshot Rules** determine which snapshot to use as the replication source.

For Replication Policies configured to use a source snapshot rule, the most recent snapshot associated with the rule becomes the replication source.

Source snapshot rules are particularly useful when the replication includes a database or other system that must be stopped in order to capture a consistent copy. Based on an external command (perhaps issued by a pre-replication script), the data management engine expects that a snapshot will be taken.

To perform *incremental replications*, the data management engine requires that the snapshot used during the previous successful replication still exist when a new replication is made. If you are using the snapshot rule queue length to control the deletion of snapshots, you must take this requirement into account and set the queue length long enough to allow for keeping the snapshot used during the previous successful replication. Also, you must take into account the possibility of intermediate failed replications, which may also create snapshots.

The following actions are taken if the required snapshots do not exist:

- *If no snapshot exists in the rule*, then the data management engine issues a warning message and performs a full replication, using an automatically created snapshot that it deletes immediately after the copy.
 - *If the snapshot taken during the previous replication has been deleted*, the data management engine cannot take an incremental snapshot and therefore performs a full copy.
- **Destination Snapshot Rules** govern the snapshot taken after a successful replication operation.



Note: Disabling snapshot usage affects the ability to run incremental replications. Snapshots must be enabled in order to make incremental replication copies, and snapshots should only be disabled if the rule is for a one-time, full replication.

Understanding custom replication scripts

Under normal conditions, pre- and post-replication scripts are not required. Where required to perform specific functions (for example, to stop an application to facilitate a snapshot of its files in a quiescent, consistent state), these custom scripts can be run before or after each instance of a replication.

In the case of a database or other application that requires a consistent state at the time of a snapshot, best practices suggest using scripts and snapshot rules together:

- **Pre-replication scripts** are executed to completion before the replication is started.
- **Post-replication scripts** are executed after a successful replication.

Potential uses of scripts are illustrated in the following examples:

- **Database replication.** A pre-replication script can be used to enable the replication of a consistent copy of the database. Typically, this pre-replication script will need to:
 1. Shut down the database to bring it into a consistent or quiescent state.
 2. Take a snapshot of the file system using a snapshot rule.
 3. Restart the database.
- **Backing Up Data from the Replication Target.** A post-replication script can initiate incremental (or full) backups from a replication target after each incremental replication has completed. Backing up from the replication target (rather than the original volume or directory) minimizes the performance impact on network users.

Using file replication rules


The **File Replication Rules** page displays all existing file replication rules and allows creation of new rules. Replication rules comprise optional configuration parameters that allow replications to be tuned to enable/disable specific functions or to optimize performance.

Replication Rules control values like the number of read-ahead processes, minimum file size used in block replication, when snapshots are deleted and whether replications will include migrated files. The server's default values should be optimal in most cases; however, these values can be changed to customize replication performance characteristics based on the data set.

Displaying file replication rules

Procedure

1. Navigate to **Home > Data Protection > File Replication Rules** to display the **File Replication Rules** page.



Field/Item	Description
Rule Name	Displays the name given to the rule when created, and referenced when creating or configuring replication policies.
In Use by Policies	Select to indicate that the rule is being used by one or more policies.
Details	Click details for a rule to display its complete details. Select a rule, and click remove to delete it.
Actions	
add	Adds a new rule.
remove	Deletes a selected rule.
Shortcuts	
Policies and Schedules	Displays the Policies and Schedules page.

Adding a file replication rule

Procedure

1. Navigate to **Home > Data Protection > File Replication Rules**, and then click **add** to display the **Add Rule** page.

Add Rule

Rule Identification

Name:

Description:

Rule Definition

⚠ Caution: It is recommended that these settings are left at "System Default" unless the Replication & Disaster Recovery Guide is carefully consulted and/or the changes are discussed with your support team.

Files To Exclude:

Block-based Replication Minimum File Size: System Default ▼

Use Changed Directory List: System Default ▼

Number of Additional Server Connections: System Default ▼

Number of Read Ahead Processes: System Default ▼

Pause While Other File Replication(s) Finish Writing: System Default ▼

Take a Snapshot: System Default ▼

Delete the Snapshot: System Default ▼

Migrated File Exclusion: System Default ▼

Migrated File Remigration: System Default ▼

External Migration Links: System Default ▼

Ignore File Attribute Changes: System Default ▼

OK Cancel

2. Enter the requested information.




Important: In general, the system default settings for this page are correct for handling most replication policies; however, in specific cases, the default values for some of the fields on this page are set when configuring the Network Data Management Protocol (NDMP). In these cases, the value might be specified by an `ndmp-option` setting on the server that overrides the system default. The `ndmp-option` command sets global system default values for certain NDMP options. These options apply to NDMP operations unless they are overridden by explicit settings sent by the NDMP client, including settings in the **Replication Rule** page.



When applicable, exceptions to the system defaults are noted in the following table.




Caution: Particular caution should be exercised when setting snapshot options (if the intent is to use the replication for incremental copies).


Field/Item	Description	Default
Name	Name of replication rule. The rule name is may include only alphanumeric characters, hyphens, and underscores.	
Description	Free-form description of what the replication rule does.	
Files to Exclude	Specifies files or directories to exclude from a replication. When specifying a file or directory, enter either: <ul style="list-style-type: none"> A full path name, relative to the top-level directory specified in the replication path. The path name must begin with a forward slash (/); at the end, an asterisk (*) can be entered as a wildcard character. A terminal file or directory name, which is simply the last element in the path. The name may not contain a character, but it may start or end with a wildcard character *. A list of files or directories to exclude from a replication. When listing files or directories to exclude from a replication, all items in the list must be separated by a comma. 	None are excluded.
Block-based Replication	Block replication minimum file size controls the minimum file size that is used for block replication. The list options available are:	Minimum file size used for block replication is 32 MB.

Field/Item	Description	Default
Minimum File Size	<p>256 or 512 K, and 1, 2, 4, 8, 16, 32, 64 or 128 MB. For instance, if this option is set to 64 MB:</p> <ul style="list-style-type: none"> ▪ For a source data file of 63 MB, for which the system determines that only 1 MB has changed, the entire source file (63 MB) will be replicated. ▪ For a source data file of 65 MB, for which the system determines that only 1 MB has changed, only the delta will be replicated. <div style="border: 1px solid black; background-color: #e0f2f1; padding: 5px; margin-top: 10px;">  Note: Requires a Replication license to function. </div>	
Use Changed Directory List	<p>Indicates if incremental replications will search for changed files in directories that only contain changed files. Processes not using the changed directory list must search the entire directory tree looking for changed files. When using the changed directory list, however, the search only is limited to those directories that contain changed files.</p> <p>Options:</p> <ul style="list-style-type: none"> ▪ system default: uses the currently specified system default. ▪ Enabled: uses the changed directory list. ▪ Disabled: always searches the entire directory tree for changed files (a full hierarchical search). 	<p>Use Changed Directory List is disabled.</p>


Field/Item	Description	Default
	<p> Note: Using the change object list is likely to improve performance in some cases; for example, where there are sparse changes. However, it can degrade performance where there are many changes throughout the directory structure.</p> <p>The calculation of the change list might take a long time as there can be a long delay between replications. Use Changed Directory List should only be selected if a large part of the directory tree will be unchanged between replication copies. Also, the list can include up to one million directories that contain changed files. If this limit is exceeded the replication reverts to a full hierarchical scan.</p>	
Number of Additional Server Connections	<p>Controls the number of additional server connections that are established during a replication operation. Ranges from 0 to 30. Increasing the number of additional server connections might improve performance by allowing multiple transfers in parallel.</p> <p> Note: Each additional server connection consumes system resources, and best practices indicate limiting the number of additional server connections to situations in which they improve performance. Also, as the number of additional server connections is increased, more read-ahead processes are required.</p>	Number of additional server connections that are established during a replication operation is four.
Number of Read Ahead Processes	Controls the number of read-ahead processes used when reading directory entries during a replication.	If a value is not set, the default value is set by the application (depending on the

Field/Item	Description	Default
	<p>Each additional read-ahead process uses system resources, so it is best to limit the number of additional processes unless it makes a significant difference in performance.</p> <p>While the default number of read-ahead processes is suitable for most replications, file systems made up of many small files increase the amount of time spent reading directory entries proportionately. In such cases, adding additional read-ahead processes may speed up the replication operation.</p>	number of read-ahead processes set in Number of Additional Server Connections).
Pause While Replication(s) Finish Writing	<p>By default, the data management engine imposes an interlock to stop NDMP backups and accelerated data copies (ADCs) from the destination of a replication during active replication writes. This function supports installations that replicate to a particular volume, then back up from that volume. However, as the lock is held at the volume level, it may be useful to override this action in the case of directory-level replication.</p> <p>To make use of this replication interlock, specify this rule option on both the replication that waits and the replication that is waited upon. As a best practice:</p> <ul style="list-style-type: none"> ▪ Create one rule with this option enabled and have each participating replication policy enable the same rule. ▪ Then, schedule the replication policy that waits to run after the replication policy that is waited upon. 	Set to no .
Take a Snapshot	<p>Overrides the Backup configuration option <i>Automatic Snapshot Creation</i>. The setting for this option should be left as the system default in almost all cases. The only case in which it might be useful is when taking a single, non-incremental copy of a file system or a directory. If there is insufficient space on the file system to take a snapshot, the copy may be taken from the live file system by selecting Disable. However, it should be noted that copying the</p>	Snapshots are taken and backed up automatically.

Field/Item	Description	Default
	<p>live file system while it is changing may give an inconsistent copy.</p> <p>Disabling snapshot usage will affect the ability to run incremental replications. This option should only be set to No if the rule is going to be used for a one-off full replication.</p> <ul style="list-style-type: none"> ▪ Enable this option to support incremental replication copies. ▪ Disable only for full replication copies or when making a complete copy of a directory. <p>Different files will be copied at different times, so if the source file system is changing and there are dependencies between different files on the system, then inconsistencies may be introduced.</p> <div style="background-color: #e0f2f1; padding: 5px;"> <p> Note: Snapshots are an integral part of the algorithm for incremental replication, and disabling snapshot usage will affect the ability to run incremental replications. This option must be enabled in order to make incremental replication copies.</p> </div>	
Delete the Snapshot	<p>Determines when snapshots are deleted. The setting for this option should be left as the system default in almost all cases. The only case in which it might be useful is when taking a single, non-incremental copy of a file system or directory. If the file system is short on space, it may be useful to request the immediate deletion of the snapshot</p>	<p>If the replication is an incremental replication, the application automatically selects the correct setting.</p>

Field/Item	Description	Default
	<p>taken for the replication. The deletion options are:</p> <ul style="list-style-type: none"> ▪ IMMEDIATELY gives the same effect as Delete snapshot after replication is done. ▪ LAST preserves snapshot for use with incremental replications. ▪ OBSOLETE deletes an automatically created snapshot when the next backup of the same level is taken. <div style="background-color: #fff9c4; padding: 5px; border: 1px solid #ccc;"> <p> Caution: As changing these settings can adversely effect the replication process, customer support recommends that this option be changed only at the direction of your Hitachi Vantara LLC representative.</p> </div>	
Migrated File Exclusion	<p>Indicates if the replications will include files whose data has been migrated to secondary storage.</p> <ul style="list-style-type: none"> ▪ Enabled: the replication will not include files whose data has been migrated to another volume using the Data Migrator facility. ▪ Disabled: migrated files and their data are replicated as normal files. 	Disabled
Migrated File Remigration	<p>Controls the action at the destination when the source file had been migrated.</p> <ul style="list-style-type: none"> ▪ Enabled: the file will be remigrated on recovery provided the volume or virtual volume has a Data Migrator path to indicate the target volume. ▪ Disabled: all the files and their data will be written directly to the recovery or replication destination volume. 	Remigration of the files is attempted.
External Migration Links	Controls when a replication operation encounters a cross volume link (a link to a	Remigrate and re-create link are enabled.

Field/Item	Description	Default
	<p>file that has been migrated to an external server).</p> <ul style="list-style-type: none"> ▪ If set to system default, the replication operation uses the default setting, which is remigrate. ▪ If set to remigrate, the replication operation copies the file contents but marks the file as having been externally migrated. The destination remigrates to secondary storage if there is an existing data migration path. This is the default behavior. Use this setting when the replication is between a main site and a disaster recovery site, in which the disaster recovery site includes a similar data migration configuration. ▪ If set to ignore, the replication operation copies only the files on the primary (migrated files are not copied). Use this setting when files have been migrated because they are less useful, so they are not replicated in order to save time. ▪ If set to re-create link, the replication operation copies only the details of the cross volume link. The cross volume link is recreated on the destination if the relevant external migration data path is in place and the migrated file is accessible. Use this setting when the replication is between storage servers or clusters on the same site, and there is a single external migration target server. 	

Field/Item	Description	Default
	<p> Note: For externally migrated files, to make sure that the file or link is replicated properly, you should either:</p> <ul style="list-style-type: none"> ▪ Specify that the replication operation should remigrate files and the destination should test before recreating links (using the <code>migration-recreate-links-mode</code> command). ▪ Specify that the replication operation should re-create links and the destination should always recreate links (using the <code>migration-recreate-links-mode</code> command). 	
Ignore File Attribute Changes	Specifies that files in which the only change is an attribute change, are not included in a replication. Only enable this option if you are certain that you do not want to replicate files with only attribute changes.	Disabled

Modifying a file replication rule

Procedure

1. Navigate to **Home > Data Protection > File Replication Rules**, select the rule you want to modify, and click **details** to display the **Modify Rules** page.
2. Enter the requested information.
The fields on this page are the same as those on the **Add Rule** page.
3. After you complete making changes, click **OK**.

Understanding files-to-exclude statements

Files-to-exclude statements contain expressions identifying directories or files to exclude from the replication. They can be written using the following guidelines:

- The asterisk "*" can be used as a wildcard character to qualify path and file name values.
In a path, "*" is only treated as a wildcard if it appears at the end of a value, for example: /path*.
In a file name, a single * can appear at the beginning and or at the end of the value; for example, *song.mp*, *blue.doc, file*.
- Parentheses (), spaces, greater than (>), and quotation marks (") are allowed around a file name or path list, but they will be treated as literal characters.
- Path and file name can be defined together but must be separated by a comma (,); for example, subdir/path*, *song.doc, newfile*, /subdir2
- The forward slash (/) is used as a path separator. As such, it must not be used in a file name list.



Note: customer support recommends creating the files-to-exclude list before the initial replication copy, and not changing it unless necessary. When running incremental updates, changes in the list do not act retroactively. For example, if a list initially excludes *.mp3 files, and the list is changed to remove this exclusion, new or changed mp3 files will now be replicated; however, any .mp3 files that have not changed since the previous replication copy will not be replicated.

Using file replication schedules

Displaying scheduled file replications

Procedure

1. Navigate to **Home > Data Protection > File Replication** to display the **File Replication** page.

File Replication

⚠ Could not get destination EVS and/or file system information.

Policies		Server	EVS	Source	File System/Path	Server	EVS	Destination	File System/Path
<input type="checkbox"/>	SiteA2B	g1-cluster	g1-evs3	prot...teA : /		172...55		SiteACopy : /	
<input type="checkbox"/>	test2	g1-cluster	g1-evs3	PHDS1 : /		172...52		fs1 : /	

Check All | Clear All

Actions: [add](#) [remove](#)

Shortcuts: [File Replication Rules](#) [NDMP Configuration](#)


Schedules		Policies	Next Run	Interval	Last Status
<input type="checkbox"/>	11	SiteA2B	2014-06-11 00:00:00 (UTC-0700)	DAILY	Failed to Start
<input type="checkbox"/>	12	test2	None	ONCE	OK

Check All | Clear All

Actions: [add](#) [remove](#) | [Abort File Replication\(s\)](#) [Transfer Primary Access](#)

Shortcuts: [File Replication Status & Reports](#)


Field/Item	Description
Policies	
Name	Identifies the replication policy.
Source	Source of the replication. The source is identified using the following fields: <ul style="list-style-type: none"> ▪ Server: Name of the server/cluster that has the source file system for this replication policy. ▪ EVS: Name of the EVS to which the replication source is mapped. ▪ File System:Path: Name of the file system and the path to which the replication source is mapped.
Destination	Destination of the replication (managed server): <ul style="list-style-type: none"> ▪ Server: Name of the server/cluster that hosts the destination file system for this replication policy ▪ EVS: Name of the virtual server hosting the file system to which the replication is mapped. ▪ File System:Path: Name of the file system and the path to which the replication is mapped.
Actions	
add	Creates a new policy.
remove	Deletes any selected policies.
Schedules	
ID	ID assigned to the replication policy.
Policies	Name of the replication policy.
Next Run	Year, month, day, and time for the next scheduled replication run for this policy.
Interval	Frequency at which the replication has been scheduled to run.
Last Status	<ul style="list-style-type: none"> ▪ Green indicates that a successful replication job has completed. ▪ Red indicates a failed replication job and lists the reason for failure.

Field/Item	Description
	 Note: In case of a replication failure, the next time a replication starts, the data management engine attempts to restart the failed replication instead of starting a new replication.
Actions	
add	Creates a new schedule.
remove	Deletes any selected schedules.
Abort Replication(s)	Aborts any selected running replication operations.
Transfer Primary Access	After one full replication and at least one incremental replication have succeeded, this starts the transfer of primary access.

Adding a file replication schedule

Procedure

1. Navigate to **Home > Data Protection > File Replication**, in the Schedules area, click **add** to display the **Add Schedules** page.




2. Enter the requested information.



Note: After a data replication job has begun, additional replications for the same policy cannot be started until the current job has completed. However, it is possible to start multiple concurrent replications, each for its own policy.

Field/Item	Description
Policy	Allows you to identify the policy to which this schedule will apply.
Replication Policy	Selects a replication policy.
Timing	Allows you to specify when the policy should run.
Immediately: Start as soon as the schedule is created	Runs the associated policy as soon as the schedule is successfully created.
Scheduled	<ul style="list-style-type: none"> ▪ Time of Initial Run: Specify the time, using the 24-hour format (such that 11:59 PM will be entered as 23:59). ▪ Date of Initial Run: Specify the date for the first run of the policy. Use the format YYYY/MM/DD (year/month/day), or select the date by clicking the calendar icon to display the calendar. <p>When using the calendar control, select the desired day by clicking the link on the day in the calendar. You can change the month and year displayed by clicking the next button or the previous button to move forward or back in one month increments.</p>
Current SMU Date and Time	Provided for reference.
Run until 23:59 on	If you do not specify a date for a final run, the policy runs at the interval specified in the Schedule section.
Schedule	<p>Select one of the options:</p> <ul style="list-style-type: none"> ▪ Daily/Weekly/Monthly - based on the scheduled date and time.: From the list, select daily, monthly, or weekly based on the scheduled date and time. ▪ Every x hours/days - based on the scheduled date and time.: Enter a quantity, then from the list, select hours or days based on the scheduled date and time. ▪ Continuous. Pause x hours between runs.: Starts a new replication job x hours after the previous job ends. The new replication job can start immediately (0 hours), or after pausing a specified number of hours.

Field/Item	Description
	<ul style="list-style-type: none"> ▪ Once, at the scheduled date and time.: Schedules the policy to run only once, at the scheduled Time and Date of Initial Run ▪ Inactive: Pauses the replication schedule. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note: If an excess amount of time elapses between replication runs, snapshots may take up a larger amount of space. By default, replication-defined snapshots are purged after 7 days (configurable to 40 days). Waiting 8 or more days between replication runs could result in a full replication.</p> </div>

3. Verify your settings, and click **OK** to save, or **cancel** to decline.

Modifying a file replication policy



After defined, schedules can be easily modified to meet the changing requirements of the replication policies. When modifying a schedule, the scheduled date and time, as well as the interval in which the schedule will run, can be changed.

Procedure

1. Navigate to **Home > Data Protection > File Replication**, select a schedule, and click **details** to display its properties in the **Modify Schedule** page.
2. Enter the requested information.



Note: After a data replication job has begun, additional replications for the same policy cannot be started until the current job has completed. However, it is possible to start multiple concurrent replications, each for its own policy.

Field/Item	Description
Policy	Displays information about the replication policy being scheduled.
Replication Policy	Displays the name of the replication policy being scheduled.
Next Run	Displays the date and time of the next replication run specified by this schedule.
Last Status	Displays the status of the last run of this schedule. Click the View Latest Report to display the replication report for the last replication run according to this schedule.
Immediate Actions	<p>Click Run now to run the replication policy immediately, regardless of schedule.</p> <div data-bbox="683 871 1393 1024" style="border: 1px solid #add8e6; padding: 5px; margin: 10px 0;">  Note: A replication job cannot be started if a previous instance of the same policy is still in progress. In this case, the replication is skipped, and an error is logged. </div> <p>Click Abort to stop an in-progress replication.</p>
Recovery Actions	<p>Click restart to restart the replication if the previous replication attempt failed.</p> <p>Click rollback to roll back a failed or aborted replication. The target file system is rolled back to the last good snapshot. Note that a snapshot is taken after every successful replication.</p> <div data-bbox="683 1318 1393 1577" style="border: 1px solid #add8e6; padding: 5px; margin: 10px 0;">  Note: Rollback should only be used when the target will be used as the live file system. If the replication's source file system cannot be used as the live file system (either permanently or temporarily), users can access the latest available data on the replication target (the file system created by the last successful replication). </div> <p>If the target file system will be used as the live file system permanently, delete the replication policy and all related schedules (since the source will not be used for this replication again). You can then create new replication policies and schedules.</p>

Field/Item	Description
	If the target file system will be used as the live file system temporarily, contact customer support for assistance in synchronizing the "old" (source) and the "new" (target) file systems before transferring access and resuming replication operations as implemented prior to the "rollback."
Timing	This section displays information about the execution timing for the replication policy, and it allows you to reschedule the next (or final) execution of the replication policy.
Schedule Time	Time of the next replication specified by this schedule.
Schedule Date	Date of the next replication specified by this schedule.
Date of Final Run	Date of the final replication specified by this schedule.
Reschedule	<p>This section allows changing the schedule of the next or final replication specified by this schedule:</p> <ul style="list-style-type: none"> ▪ To change the schedule of the next replication, fill the Reschedule box, then enter the new values for the Time and/or Date. ▪ To change the schedule of the final replication, fill the Reschedule box, and enter the new value for the Final Run in the appropriate fields.
Current SMU Date and Time	Current date and time as set on the SMU.
Schedule	<p>This section allows you to specify how often the replication policy is to be executed. Select one of the radio buttons:</p> <ul style="list-style-type: none"> ▪ From the list, select daily, monthly, or weekly based on the scheduled date and time. ▪ Enter a quantity, and from the list select hours or days based on the scheduled date and time. ▪ Enter a quantity to complete the label: <i>Continuous. Pause quantity hours between runs. The new replication job can start immediately or after a specified number of hours.</i> ▪ Selecting <i>Once</i>, at the scheduled date and time guarantees that the policy is scheduled to run only once. ▪ Selecting <i>Inactive</i> causes the replication schedule to be placed on pause.
Actions	

Field/Item	Description
OK	Saves changes to the replication policy schedule, and returns to the File Replication page.
cancel	Returns to the File Replication page without saving changes to the replication policy schedule.

3. Verify your settings, and click **OK** to save, or **cancel** to decline.

Understanding incremental replications

Incremental replications rely on the existence of the snapshot taken during the previous replication. If this snapshot no longer exists, the data management engine performs a full replication. The data management engine automatically preserves the snapshots it needs for replication. However, there is an age limit applied to snapshots that are automatically taken by the NDMP system (including during a replication).

Snapshots older than the age limit are automatically purged from the system. The default limit is 7 days, but the limit can be configured through the **NDMP History and Snapshots** page. If the replication copy time is very long, or the interval between replications is long, then the default age limit must be extended.

Displaying file replication status and reports

The **Replication Status & Reports** page displays a list of replication jobs in progress or completed. It also includes reporting details on files replicated, amount of data replicated, and success or failure status. If a schedule is deleted, the reports associated with it are also deleted.

Procedure

1. Navigate to **Home > Data Protection > File Replication Status & Reports** to display the **File Replication Status & Reports** page.

The screenshot shows the 'File Replication Status & Reports' page. At the top, there is a breadcrumb trail: 'Data Protection > Home > Data Protection > File Replication Status & Reports'. Below this is a 'Display Options' section with a checkbox for 'Group by Policy Name' and an 'apply' button. The main content is a table with the following columns: 'Schedule Id', 'Policy', 'Completed', 'Duration', 'Bytes Transferred', and 'Status'. The table contains six rows of data, all with a status of 'OK'. Below the table are links for 'Check All' and 'Clear All'. At the bottom, there is an 'Actions' section with a 'remove' button and a 'Remove All' link, and a 'Shortcuts' section with a link to 'Policies and Schedules'.

Schedule Id	Policy	Completed	Duration	Bytes Transferred	Status
1	Test	2016-12-09 00:00:12 (UTC-0800)	00:00:11	245,808	OK
1	Test	2016-12-08 00:00:12 (UTC-0800)	00:00:11	245,344	OK
1	Test	2016-12-07 00:00:12 (UTC-0800)	00:00:11	244,880	OK
1	Test	2016-12-06 00:00:12 (UTC-0800)	00:00:11	244,416	OK
1	Test	2016-12-05 00:00:12 (UTC-0800)	00:00:11	243,952	OK
1	Test	2016-12-04 00:00:12 (UTC-0800)	00:00:11	243,488	OK

The replication report Status column displays results of a replication job (green for OK, red for failed). Reports can also be beneficial for analyzing the effects of a particular incremental replication policy. The information in the **Report Summary** page provides a detailed view of the replication job results. This information can be used to make performance adjustments to the replication policy and schedule.

Field/Item	Description
Display options: Group by Policy Name	Displays the file systems in groups based on the policy name.
Schedule ID	ID number for the completed replication.
Policy	Policy name.
Completed	Month, date, year and time when the replication was completed.
Duration	Duration of a replication schedule run.
Bytes Transferred	Volume of replicated data in bytes.
Status	Status of replication completion.
details	Opens the File Replication Report page for the selected file replication.
remove	Delete the selected reports.
Remove All	Delete all the present reports on the screen.

2. Click **details** for a selected replication to display its properties.

File Replication Report

⚠ Could not identify the source server from the report details, so some information may be incomplete. The file replication job may only just have started, or the report is very old.

Report Summary

Policy:	11
Schedule ID:	11
Status:	Failed to Start
Frequency:	DAILY
Start Time:	2014-06-10 00:00:01 (UTC-0700)
End Time:	2014-06-10 00:01:53 (UTC-0700)
Duration:	00:01:52
Bytes Transferred:	0
Copy Type:	Unknown copy type
Server / EVS:	/ EVS 255
Rule:	

```

2014-06-10 00:00:01-0700: Run replication policy SiteA2B, schedule 11 (scheduled)
Unable to update replication (SiteA2B) stats:
Cannot get transfer details: Unable to find filesystem on another server Server: 172.31.60.59 ACCESS_FAILED: unable to contact server: send: No route to host Failed to establish SSC
connection Server: 192.0.2.9 MISSING_RESOURCE: Could not find fsid "9393607122010511"
    
```

Actions: [back](#) | [Download Replication Log](#)

Field/Item	Description
Replication Policy	Completed replication policy name.
Schedule ID	Completed replication schedule ID.
Status	Indicates whether the replication was successfully completed.
Frequency	How often the Policy is scheduled to run.
Start Time	Date and time when the replication began.
End Time	Date and time when the replication ended.
Duration	Duration of replication.
Bytes Transferred	Volume of data replicated, in bytes.
Copy Type	Type of replication performed. May be any of the following: <ul style="list-style-type: none"> Full Copy: A complete initial replication of the entire source to the target. Incomplete Copy: The replication did not complete. Incremental Copy: A replication of the changes on the source file system to the target.

Field/Item	Description
	<ul style="list-style-type: none"> ▪ Restart Copy: The replication started from the point of failure of the previous replication. ▪ Rollback Copy: After a failed replication run, the target file system was rolled back to its state following the last successful replication.
Server/EVS	EVS on which the source and destination file systems reside.
Rule	The name of the rule used by the policy.
Transfer Primary Access Summary	This section appears in the replication report only after a transfer of primary access.
Status	Indicates whether the transfer of primary access was successfully completed, and indicates any actions that should now be taken.
CIFS	Number of SMB shares that were successfully transferred to the new location.
NFS	Number of NFS exports that were successfully transferred to the new location.
FTP	Number of FTP initial directories that were successfully transferred to the new location.
FTP Users	Number of FTP users that were successfully transferred to the new location.
Snapshot Rules	Number of snapshot rules successfully transferred to the new location.
CNS Links	Number of CNS links successfully transferred to the new location.
Backup Files	List of SMB shares backup files and NFS exports backup files that were successfully transferred to the new location.
View Failures	Click View Failures to display a list of items not transferred during the transfer of primary access.

Enabling multiple replication streams

You can add additional server connections to a replication rule using the Number of Additional Server Connections field of the replication **Add Rule** page or the **Modify Rule** page.

Select the number of additional connections to add for use by the replication/accelerated data copy (ADC) utility operation. You can specify between 0 and 30 additional server connections. Note that these are additional server connections; if the number of additional connections is set to 0, the replication operation will have a single connection. The default is 4 additional connections, along with 12 read-ahead processes.

If the number of additional server connections has been set to non-default and more than zero, then the number of read-ahead processes must also be set to a non-zero value that is appropriate for the specified number of additional server connections.

Configuring NDMP performance options

NDMP Performance options are set using the **Add Rule** or **Modify Rule** page of the NAS Manager. On these pages, you can set the number of additional server connections and the number of read-ahead processes (which are the options with the biggest effects on replication performance), as well as other replication options.



Note: The `readahead_procs` setting of the `ndmp-option` command is no longer used for replications.

The number of additional server connections and the number of read-ahead processes should be coordinated to get the best performance. Each additional connection causes the creation of a separate process at the source and one at the destination, and these processes are connected by their own separate TCP connection. These two processes work together as an independent replication stream which can process subdirectories in parallel with other replication processes. Read-ahead processes are used only at the replication source; these processes pre-read directory entries and file details from the storage media (typically disks) so that the main replication processes can use them immediately without being delayed by disk read latencies.

Although allocating more processes to a replication can improve its performance, the extra processes take up system resources. Using these resources for replication operations may negatively impact the performance of other processes for protocols (such as NFS or SMB), features, or even other replications. Also, the performance improvement per additional process reduces as the number of processes increases, and at some point there will be no further improvement (there may be a reduction in performance if too many processes are used). With these points in mind, you should not request a very high number of processes, except in very special cases.

The optimal settings for these values depend on many factors, including:

- File system size and layout of files. Typically, to get best performance when replicating file systems with smaller files and fewer files per directory, you should dedicate more read-ahead processes and connections to the replication.
- The number of replications that will run simultaneously. When running many replications simultaneously, each replication should be allocated fewer read-ahead processes so that the overall load from replication processes is not too high.

- The relative priority of the replication and other uses of the system (such as NFS and SMB access). If replications appear to be adversely affecting the performance of system functions or user access, then reducing the number of read-ahead processes and connections used by replications should alleviate the problem.
- The number of physical disks containing data for the file system (which can be found by looking at the number of physical disks in the system drives used by the file system). If the data of the file systems being replicated is stored on relatively few physical disks, then increasing the number of connections (data streams) used in the replication operation will not improve performance. Refer to the *Storage Subsystem Administration Guide* for information on system drives.
- The properties of the network route between the source and destination machines. When the connection between source and destination machines has high bandwidth available, long latency connections (high speed cross-continental or intercontinental links), then the long latency may impose an artificially low data rate over a single TCP connection. Using parallel connections (data streams) for the replication operation can improve performance in this case.

The following notes give more specific indications of how to choose settings. However, because of the many factors involved these recommendations may not be absolutely optimal. If it is vital to get the very highest performance from your system, some experimentation will be needed to find the most suitable values for the installation.

- The default settings are 4 additional connections and 12 read-ahead processes. These settings should be adequate for a wide range of applications, including files systems with mixed small and medium size files (average file size 250 KB), up to file systems containing very large files. Also, these settings are appropriate for file systems with data on a relatively small number of physical disks (such as 32). A single replication operation with these settings should not severely impact other system users.
- If many replication operations are running simultaneously, you may want to reduce the number of read-ahead processes and connections each replication uses.

For example, if you are running eight simultaneous replications, the settings might be one additional connection and six read-ahead processes.

- Where the files in the file systems being replicated are mostly small, increasing the number of connections and read-ahead processes used will usually produce better performance.

For example:

- For a file system with an average file size of less than 64 KB, you may want to set 8 additional connections and 20 read-ahead processes.
- For a file system with an average file size of 32KB, you may want to set 12 additional connections and 24 read-ahead processes.

If the number of files per directory is also very low (less than 4 or 5 files per directory), even more connections and read-ahead processes might improve performance.

- The default TCP window size used by the server is 256 KB. If the latency (round trip time) of the link is 70ms, then the maximum realistic throughput on a single TCP connection is about 3 MB per second.
- If the file systems involved have relatively few physical disks, increasing the number of connections and read-ahead processes will gain relatively little performance improvement. For instance, for a small source file system with data on only 32 physical disks, there will not be much to gain by increasing the values above the defaults.

Maximum concurrent replications

There are two variables which limit the number of concurrent Object Replication sessions per node on the replication source and target (destination) respectively:

- replication-max-source-replications
- replication-max-destination-replications

These variables are intended to limit the replication source and replication target pools. They default to 128 and apply independently i.e. you can have 128 concurrent sessions on the replication source and target. If the number of concurrent replications on the replication source has been reached, then for instance a replication started from the console outputs an error message similar to:

```
hnas:$ vn 1 replication-start ObjRepPolicy ObjRepSource
Could not start object replication using policy 'ObjRepPolicy' on file system
'ObjRepSource': Node has reached maximum number of concurrent object replications
```

This will abort the replication and if necessary cleanup any source auto snapshot.

If the number of concurrent replications on the replication target has been reached, the replication target returns a status of NoResources which on the replication source, if the replication has been started from the console, causes an error message similar to:

```
hnas:$ vn 1 replication-start ObjRepPolicy ObjRepSource
...
Could not start object replication using policy 'ObjRepPolicy' on file system
'ObjRepSource': Insufficient resources on target to perform object replication
```

Also, the replication history log contains the reason for not starting the replication.

Troubleshooting replication failures

The following are some scenarios in which a replication job can fail:

- The destination volume is offline.
- The destination volume was full.

- One of the volumes involved may have been unmounted.
- The replication process was rebooted while a replication job was in progress.



Note: Without any further action upon a replication failure, the replication will continue as expected on its next scheduled run. However, this will recopy any changes already copied during the failed replication attempt.

Clicking Restart will cause the failed replication to be restarted immediately and will avoid recopying most of the data.

Manually restarting a failed replication

If a replication has failed, the replication will be started normally at its next scheduled run time, rather than "picking up where it left off." To restart the replication from the point of failure (before its next scheduled time), you must restart it manually:

Procedure

1. Navigate to **Home > Data Protection > File Replication** to display the **File Replication** page.
2. Click **details** for the failed replication to display its **Replication Schedule** page, and click **restart**.

Rolling back an incomplete replication

Upon successful completion of a replication, the system takes a snapshot to preserve the state of the target file system. With this snapshot, if an offline source leads to failure of a subsequent replication, the target file system can be rolled back to the state of the last successful replication.

To rollback the target file system to the state of the last successful replication:

Procedure

1. Navigate to **Home > Data Protection > File Replication** to display the **File Replication** page.
2. Click **details** for the failed replication to display its **File Replication Schedule** page, and click **rollback**.



Note: Rollback should only be used when the target will be used as the live file system. If the replication's source file system cannot be used as the live file system (either permanently or temporarily), users can access the latest available data on the replication target (the file system created by the last successful replication). There are two possible approaches:

- If the target file system will be used as the live file system permanently, delete the replication policy and all related schedules (since the source will not be used for this replication again). You can then create new replication policies and schedules.
- If the target file system will be used as the live file system temporarily, contact Hitachi Vantara Support Center for assistance in synchronizing the "old" (source) and the "new" (target) file systems before transferring access and resuming replication operations as implemented prior to the "rollback".

Chapter 3: Transferring primary access

A transfer of primary access copies data from a portion of a file system and relocates the access points for that data, or relocates an entire file system and its access points (copying the data and metadata), with very little down time, while the file system is live and servicing file read requests. For a short period, access is limited to read-only.

A transfer of primary access cannot move all attributes/relationships for a file system, but it can move most of them. For example:

- The following can be moved:
 - SMB Shares (if within replicated path)
 - NFS Exports (if within replicated path)
 - FTP Initial Directories/Users (if within replicated path)
 - Snapshot Rules (Moving snapshot rules requires their deletion and recreation. Deletion of snapshot rules automatically leads to the deletion of snapshots created with those rules)
 - CNS Links
- The following cannot be moved:
 - iSCSI Targets
 - Global Symlinks

A transfer of primary access can be performed on any replication policy as long as the following conditions are met:


1. A full replication has completed. Preferably an incremental replication should also have completed.
2. The snapshot required to support another incremental replication must still be available.



Note: For any given replication policy, only one transfer of primary access operation may be in progress at any time.

How a transfer of primary access moves CNS links

Using the diagram below as a sample file system:



When replicating the file system path beginning at “D,” CNS links are transferred as follows:

- If the file system is linked to the CNS tree at “A” or “B,” the CNS link is not moved and is listed in the replication report as an error. The CNS link is not moved because doing so would deny access to the file system at point “E.”



Note: In this situation, users will be able to access the “old” data after the replication and transfer of primary access are complete. After a successful transfer of primary access, the original source data should either be removed or made inaccessible by network clients (permissions should be changed).

- If the replication is within the cluster, and the file system is linked to the CNS tree at “D” or below, then the CNS link is moved and is listed in the replication report as having been successfully moved.
- If the file system is linked to the CNS tree at “C” or “E,” the CNS link is not moved and it is not listed in the replication report, because it is not relevant to the path being replicated.

Process of transferring primary access



Caution: When performing a transfer of primary access, the snapshot rules may be deleted on the source file system. This will lead to the snapshots and snapshot schedules also being deleted on the source file system.

A single transfer of primary access operation may be in progress at any time for any given replication policy, and the process for the transfer of primary access is as follows:

Procedure

1. Put the source and destination (target) file systems into "syslock" mode.


When a file system is in Syslock mode, the storage server allows read-only access to the file system, but not write access.

The storage server ensures that the target file system data is consistent with the source file system data before primary access is transferred. This involves making the source and destination file systems read-only for a short time. Although any arbitrary directory can be relocated, the entire source file system is set to syslocked mode while the final replication operation is in progress. For more information on syslock mode, refer to the *File Services Administration Guide*.


2. Notify clients that a short period of read-only access will be necessary while data and file system settings are relocated.
3. Replicate the data and file system settings to the new location.

After a transfer of primary access has been started, the replication process monitors the replication to determine when it is complete. When the replication is complete, the replication process starts moving configuration information automatically. The following table describes how network access points on the source file system are moved or deleted:

Source File System Setting/ Network Access Point Being Moved	Destination		
	Within the EVS	Another EVS in the Same Cluster	An EVS on Another Server or Cluster
SMB Shares (if within replicated path)	Moved (path is modified). Clients that had the share mounted before the transfer of primary access do not have to remount the share after the transfer.	Moved (deleted from source EVS then added on target EVS). Clients that had the share mounted before the transfer of primary access must remount the share after the transfer only if the share was not to a directory in the CNS.	Moved (added to target EVS then deleted from source EVS). Clients that had the share mounted before the transfer of primary access must remount the share after the transfer.
NFS Exports (if within replicated path)	Moved (path is modified).	Moved (deleted from source EVS then added on target EVS).	Moved (added to target EVS then deleted from source EVS).

Source File System Setting/ Network Access Point Being Moved	Destination		
	Within the EVS	Another EVS in the Same Cluster	An EVS on Another Server or Cluster
	 Note: Clients that had the export mounted before the transfer of primary access must mount the export again after the transfer (the NFS mount becomes stale after the transfer).		
FTP Initial Directories/Users (if within replicated path)	If all users within an initial directory can be moved, the initial directory is also moved. If all users within an initial directory cannot be moved, no users are moved and the initial directory is not moved.	If all users within an initial directory can be moved, the initial directory is also moved. If all users within an initial directory cannot be moved, users are moved where possible, and the initial directory is duplicated.	

Source File System Setting/ Network Access Point Being Moved	Destination		
	Within the EVS	Another EVS in the Same Cluster	An EVS on Another Server or Cluster
Snapshot Rules	<p>Where a file system is replicated from root (/) and all data and access points transferred to a new standalone file system, Snapshot Rules related to the source file system are deleted and recreated for the target file system. Deletion of snapshot rules automatically leads to the deletion of all snapshots associated with the rule on the original, source file system. This behaviour is intentional and unavoidable.</p> <p>Where a file system is only replicated partially, from sub-directories off the root of the file system (/), and only a subset of data and access points are transferred to the target file system, snapshot rules are not deleted and recreated on the target file system.</p>		
CNS Links	<p>If CNS entries already point to the replication source, then the CNS link is removed and a link to the new file system is added at the corresponding path. Note, however, that if the file system is linked to the cluster name space at a point higher in the directory structure than the root directory for the file system path being replicated, moving the CNS link is not possible. In such cases, the CNS link is reported as an error in the list of successful/failed transfers, and the administrator must manually create a CNS link to the file system in the new location.</p> <p>After a transfer of primary access, network clients will not be able to access the file system through the a CNS name space if any of the following are true:</p> <ul style="list-style-type: none"> ▪ The file system did not have CNS links. ▪ The file system's CNS links were not moved. ▪ The file system was replicated to another server or cluster. 		<p>If CNS links exist, the relocation is not allowed to proceed, and a message advises the administrator to remove the links before proceeding.</p>

Source File System Setting/ Network Access Point Being Moved	Destination		
	Within the EVS	Another EVS in the Same Cluster	An EVS on Another Server or Cluster
	<p>To access the file system in its new location, network clients must reconnect through SMB shares or NFS exports pointing to the relocated file system or to a CNS name space into which the file system is linked. NFS clients pointing to a CNS name space will not experience any interruption.</p> <p> Note: If clients will not access the relocated file system using CNS links, they must access it using new IP addresses.</p>		
iSCSI Targets	Not moved.		
Global Symlinks	Not moved.		



Note: For SMB shares and NFS exports: if possible, a text file backup of the moved shares and export is retained.

4. Bring the target file system online.

The system administrator receives instructions to bring the target file system on-line, by allowing read/write access. Read/write access is re-enabled on the entire source file system unless it was syslocked originally).

The replication process tracks/records the progress of the final replication. Status of the network access point relocation is available through the **Status and Reports** page; replication failures are logged and can be viewed by following a link from the replication report.

5. Begin servicing file service requests from the relocated file system.

6. If the source file system was online when the transfer of primary access was started, put it back online by taking it out of syslocked mode.



Note: If the replication process is rebooted during a transfer of primary access, the source file system may not be returned to its original online state. If the replication process is rebooted during a transfer of primary access, you may have to take the file system out of syslock manually, from the **File System Details** page. For more information on syslock mode, refer to the *File Services Administration Guide*.

After the final replication has completed, the original source data is still present on the source. This data can be accessed (and modified) through access points configured higher up in the directory tree, and should be deleted manually.



Note: After the successful completion of a transfer of primary access, the source file system should be deleted or made inaccessible manually.



WARNING: It is especially important to perform the above in the case where data from the source file system had previously been migrated to cloud. Do not delete, or reverse migrate, the data from the source file system as both the source and the target (relocated) file systems refer to the same copy of data migrated to cloud.

All replication schedules configured for the replication policy are set to inactive once the transfer of primary access is completed, and these inactive policies should then be deleted manually.

Handling a failure during a transfer of primary access

If a failure occurs during a transfer of primary access:

- The target file system is not brought online in place of the source.
- The source remains accessible and usable to network clients.
- There is no attempt to rollback after a failure.
- The replication process performs as many actions as possible, but leaves the replication policy in place.
- A partially failed final replication does not remove the replication policy/schedule.
- The system administrator can usually resolve the issue that caused the failure, then run transfer primary access again.

For example, when replicating several SMB shares, one share fails to be replicated, but the others are replicated successfully:

- The share that failed is logged to a simple text file (which is viewable from the **File Replication Report** page).
- All other shares that were successfully recreated are brought online, and deleted from the source.
- When complete, the system administrator sees the error message and then views the text file using the **File Replication Report** page).
- Viewing the text file, the administrator sees that the share could not be created on the target, perhaps because the name is already in use. The system administrator can delete the named share, either from the source or the target, and then transfer primary access again, this time successfully.

Chapter 4: Using object replication

Object based file system replication provides a mechanism, manual or automatic, for copying or relocating both file systems and the metadata associated with those file systems (such as access points, security descriptors, and other file system related data). The source file system may be replicated to one or more target file systems. When configured correctly, object replication can mirror file systems at different physical locations, which can be used as a disaster recovery configuration. Object-based replication operates on the entire file system, not at the individual file or directory level. NAS Platform supports object-based file system replication.

When using object-based file system replication, you can replicate the file system, and the associated CNS links, SMB shares, permissions and all other file-level metadata. Administrators can use NAS Manager to configure policy-based object replication jobs independently from other replication strategies.

Object replication, like file replication, uses policies and schedules to determine which file systems get replicated, where they are replicated, and when replication operations are run. Policies specify the replication source and the target, and schedules specify the timing and the interval of repetition, if any.


Configuring object-based replication

Replication (file or object) is a licensed feature, and the *Replication* license must be installed before replications can be performed. Refer to the *Server and Cluster Administration Guide* for more information about licenses.

After the replication license is installed, and the listening port is specified, you can create file systems to serve as replication targets, and then create the policies and schedules to control the replication of the file systems.

Procedure

1. Navigate to **Home > Data Protection > Object Replication Configuration** to display the **Object Replication Configuration** page.

Field/Item	Description
Object Replication Listening Port	Specify the port on which a cluster is to listen for object replication connections from other servers/clusters. The default value is 59550. <div style="background-color: #e0f2f1; padding: 10px; border: 1px solid #ccc;">  Note: When you change the port, the object replication service is restarted. Under some circumstances the service cannot be restarted at a specific time. For example, when the port is in use because it has been randomly allocated to another service, or when the service is handling an active object replication process. If the service cannot be restarted at that time, manually restart file serving on that node after changing the port. </div>
apply	Applies the configuration changes.
restore default	Restores the default configuration values.

2. Enter the listening port.
If you need to restore the default port, click **restore default**.
3. Click **apply**.
After the configuration has been successfully saved, a confirmation message is displayed on the **Object Replication Configuration** page.

Read only access to replication target file systems

Read only access to replication target file systems is possible using NFS, SMB, FTP and NDMP protocols. Previously, only object replication could access and update target file systems. Any manual or client-based attempt to access a replicated filesystem required that the filesystem be changed to a read-write file system first, preventing the file system from being a replication target. Read only access by clients is possible without any adverse effect on file systems.

When using read only access, files and directories on a replication target:

- Are read only.
- Match the most recent completed replication.
- Are updated as soon as a new replication completes.
- Are unavailable until the first completed replication. The file system can be accessed while replications are active, however there must first be an initial replication to the target to provide a valid file system to clients.

Note the following changes from normal filesystem behavior when using read only access:

- NFS locks have no effect and are ignored.
- SMB Oplocks are not issued.
- SMB3 transparent failover is not available.
- When virus scanning is underway, unverified files may be re-scanned.

Shares and exports with read-only access

Using shares and exports with read-only access can cause a transfer to fail at the disaster recovery site if a shared name is in use on the replication target, therefore:

- Names must be defined so transfers work.
- Incoming shares and exports can overlap existing ones and be renamed.

NDMP access to object replication target file systems

Read only access to a target file system is possible using the NDMP protocol, but note that some restrictions apply:

- Only backups are allowed.
- Snapshot Rules must be used.
- Object replications must use snapshot rules at the target.
- Setting the snapshot rule queue depth is critical.
- Snapshot rules for backups use Hitachi Vantara-specific settings.

The following areas of functionality are affected by the file system being an object replication target:

- Incremental backup mechanisms.
- Direct access recovery.
- Rollback of the object replication target.

The restrictions and considerations are explained in detail here:

Only backups are allowed

When using NDMP, only backups work on a replication target. Besides backups, other NDMP operations need to be able to write to the file system, for example file replications need to be able to create snapshots. All other NDMP operations besides backups will fail on an object replication target. Recovery from backed up data would be made to a different file system, or the object replication target would need to be promoted. The backup need not be of the whole file system, and sub-directories or virtual volumes can be backed up normally.

Snapshot rules must be used

The file system state of a replication target is created only by the object replication, and nothing else. NDMP backups must reference a past state of the file system in order to perform incremental backups. A snapshot is created by the snapshot rule defined for the replication policy. There is an existing mechanism by which NDMP operations can already specify a snapshot rule. The latest snapshot in the specified rule becomes the dataset that is backed up. Only backups that specify a snapshot rule are allowed. All others will fail.

Setting the snapshot rule queue depth is critical

If more object replications are run using the snapshot rule than the queue depth, the oldest snapshot in the queue is destroyed. If the backups are run less frequently than the replications, it is possible that the destroyed snapshot was one that the incremental backups were using as a base. The frequency and timing of the object replications and backups must be very carefully planned so the base snapshot is not destroyed.

Incremental backup mechanisms

There are two ways that NDMP backups can specify the base from which an incremental backup is made:

1. The standard practice of using backup levels.
2. Use token-based backups.

Both work correctly as long as a snapshot rule is specified. The level or token defines the base file system state that the incremental backup refers back to and the snapshot rule defines the state of the file system at the point of backup.

Direct Access Recovery

Although it is not possible to recover to an object replication target, you can use backups made of the target to recover to a different file system, via Direct Access Recovery (DAR). Per typical practices, the NDMP environment variables must be set to support such a recovery.

Rollback of the object replication target

Rollback of the object replication target affects incremental backups. No data is lost, but backups that immediately follow the rollback must re-establish themselves to lower levels, if they depended on a snapshot that is lost in the rollback. This could cause the backup to take longer than might normally be expected.



Note: Specifying snapshot rules for backups is a Hitachi Vantara proprietary setting. The existing mechanism for specifying a snapshot rule is to set the NDMP environment variable `NDMP_BLUEARC_USE_SNAPSHOT_RULE` to the name of the rule that you wish to use. This is not possible on all backup applications. To permit the use of backup applications where it isn't possible, use **`ndmp-backup-snapshot-rules`**. This command allows the per-file system mapping of a snapshot rule to a particular level of backup. Allowing the use of different snapshot rules for different levels allows more flexibility in setting the queue size of the rules e.g. low numbered backup levels happen rarely and require their snapshots to be kept for a long time

Please see the NAS man pages for more detailed information on using commands for backups and snapshots.

Using object replication policies

In the same way that file replication policies specify the details about file replication operations, object replication policies specify the details about object replication operations.

Adding object replication policies

The **Add Object Replication Policy** pages allow you to define the properties of a new object replication policy, including its source file system and its replication target (server/cluster, EVS, and file system).

If this is the first policy created, and if the file system is not configured to transfer access points with a replication, then the option to manually configure the file system to transfer access points will be given. You are provided with more resolution options on the second **Add Object Replication Policy** page.



Note: Selecting a source or target file system that already has an object replication policy defined will not cause the old policy (or policies) to be overwritten. You can have multiple replication policies on a source file system, each pointing to a different target file system. You could also have multiple policies with a single source/target pair; however, it is not recommended.

Procedure

1. Navigate to **Home > Object Replication** to display the **Object Replication** page.
2. In the Policies section, click **add** to display the **Add Object Replication Policy** page.






Caution: If the file system is not configured to transfer access points with replications, and if the target file system must later be recovered as a read-write file system (for example, to replace the source file system in a disaster recovery situation), the shares and exports configured on the source file system are not copied to the target file system. Network clients that relied on those shares and exports to access their data on the source file system are longer able to access the recovered data on the target file system. The shares and exports must be individually configured to be transferred with replications, via their respective share and export details pages, in order to allow network clients to access their data on the source file system. You are provided with more resolution options on the next page, after clicking **next**.



Note: It is recommended that the replication target is at least as big as the source file system, to ensure that all data can be replicated on the target. This is especially important if you intend to keep multiple snapshots on the target, as they require more storage space.

Field/Item	Description
Identification	
Name	The name of the object replication policy.
Source	
EVS/file system	The name of the source EVS and file system that is to be replicated to the object replication target. To change the source EVS/file system, click change...
EVS IP Address	The default value for this field is 'Automatically selected'. If required, select the IP address for the source EVS from the list.
Target	
Server	Select a server from the list to be the target of the object replication policy. After selecting a server, click select a target... to select an EVS and file system. For managed servers, the EVS Name and File System fields are automatically populated when you select a file system using select a target... Or, alternatively, type the EVS IP/host name and file system name in the corresponding fields.
EVS DNS Name or IP	Select the EVS IP/host name in this field.
EVS IP Address	Enter the EVS IP address in this field.
File System	Select the target file system of the object replication policy. If the replication process is not managing the server on which the target file system is hosted, or if the destination file system does not yet exist, select Specify EVS IP/host name and file system from the list and type the details in the appropriate fields.

Field/Item	Description
	<p> Note: The replication target file system should be at least as large as the source file system to ensure that all data can be replicated on the target. If you intend to keep multiple snapshots on the target, it is especially important that the target be larger than the source, because the additional snapshots on the target file system will require storage space.</p> <p> Note: The tiering of both the replication source and target file systems must agree; you cannot have a tiered source file system replicated to an untiered file system, and you cannot have an untiered source file system replicated to a tiered file system.</p>
Object Replication Listening Port	<p>The port on which the destination server is listening. The default is 59550. The port on which the destination server listens is configured on the Object Replication Configuration page of the NAS Manager, or through the <code>replication-listening-port</code> CLI command.</p> <p> Note: To change the listening port for the target server, you must make that server the currently managed server of the replication process, then use the Object Replication Configuration page of the NAS Manager to change the listening port.</p>

3. Specify the details for the policy identification, replication source and the replication target.
4. Click **next** to continue with the policy configuration.
 The replication process checks if the source file system is configured to allow the access points (SMB shares and NFS exports) on it to transfer with the object replication. If it is not, a NAS Manager page is displayed that gives you the option of configuring the source file system so it can transfer shares and exports with the replication. If another policy already exists for this file system, the check for access point transferability is skipped, and the second part of the **Add Object Replication Policy** page is displayed:

 The policy can result in two types of replications of the source file system: an initial replication, and an incremental replication. A replication of an initial snapshot results in a full copy of the source file system being replicated on the target, while subsequent replications use an incremental snapshot that only copies objects that have changed since the initial snapshot was taken.

 When setting snapshot rules, you can choose either an automatic snapshot rule or an existing, named rule.

In order to avoid snapshots being taken of inconsistent data, please carefully note the following recommendations:

- If you choose the automatic snapshot rule option, the snapshot is taken whenever the replication first runs. In order for the replication engine to take a snapshot with consistent data (that is, data that is not actively being modified on the source file system), it is recommended that the replication be run when the file system is not being actively accessed. For example, the replication can be run manually, when the source file system is inactive, to obtain a snapshot with consistent data. Or, the replication policy should be scheduled to run when the file system is not being accessed.
- If you choose a named snapshot rule, the snapshot is taken when specified by the rule; however, it is recommended that the snapshot is taken manually, when the file system is not being actively accessed, or scheduled to be taken when such a time is anticipated. Then the replication should be scheduled at a time when the server is minimally active, or run manually at such a time.



The following table describes the fields on the second **Add Object Replication Policy** page:



Note: If a database application is writing to its database when a snapshot is taken, the snapshot can contain only some of the writes to the database, rather than all of them. In such a case, the database might contain only partially written records and is therefore not consistent.



Caution: Setting the snapshot options for the source and target file systems must be done very carefully, to ensure that replications provide a good copy of the data on the target. It is recommended that you consult with support representative if you are unsure of how to correctly set snapshot options for an object replication.

Field/Item	Description
Source File System	<p>Options for taking snapshots of the source file system:</p> <ul style="list-style-type: none"> ▪ Snapshot source file system using automatic snapshot rule, which allows the replication to use its default snapshot rule to take and manage snapshots. <div data-bbox="639 842 1393 1100" style="border: 1px solid #ADD8E6; padding: 5px; margin: 5px 0;">  Note: If you choose this option, each incremental snapshot is deleted when the next replication runs. Therefore, because the snapshot queue only contains one snapshot, it is recommended that replications are not scheduled too closely together in order to prevent an existing snapshot from being removed before the next replication starts. </div> ▪ Use snapshot rule, which means that the source snapshot retention policy can be customized to retain a different number of snapshots on the source file system. <div data-bbox="639 1234 1393 1564" style="border: 1px solid #ADD8E6; padding: 5px; margin: 5px 0;">  Note: If you choose this option, set the schedule for the snapshot rule so a snapshot is created before the replication runs, to ensure that a new snapshot is available for the replication. A snapshot of the source file system is only taken if the replication policy is configured to use an automatic snapshot rule. If it is using a named rule, the replication will use the latest snapshot created by that rule; it does not take one automatically. </div>
Target File System	<p>Options for taking snapshots of the target file system:</p> <ul style="list-style-type: none"> ▪ Snapshot target file system using automatic snapshot rule, which allows the replication to use its default snapshot rule to take and manage snapshots on the object replication target. ▪ Use snapshot rule, which allows the snapshot retention policy to be customized to retain a different number of snapshots on the source and destination.

5. Enter the processing options.
6. Click **next** to advance to the **Add Object Replication Policy** page.



Note: The source and target file systems are evaluated and, if any issues are detected, a page appears listing them and offering options to help you either correct the issue or to go back to the previous page.

7. Review policy settings, and then click **create**.
8. To add a schedule for this policy, click **yes** to display the **Add Object Replication Schedule** page.

Correcting access point problems in an object replication policy

Before a new object replication policy is created, it is evaluated for potential problems, such as the source not being configured to allow the transfer of access points (SMB shares or NFS exports). If an issue is discovered, a page appears, offering options to correct the issue or to go back to a previous page, where you can make changes to resolve the issue.

Typically, when this page appears, the following links are displayed:

- **Configure:** Click configure to automatically configure the selected source file system and continue creating the policy.

Clicking this link configures the selected file system to allow the access points on it to be transferred with an object replication. After clicking this link, you are returned to the previous page to continue the configuration.

- **Change:** Click change to return to the previous page, in which you can select another file system as the source of the object replication.
- **Continue:** Click continue to continue configuring the object replication policy using the selected file system, even though that file system is not configured to allow the access points to be transferred with an object replication. You must later explicitly set the shares and exports to be transferred with the object replication at a later time.

To manually override the default, you must manually reset the shares and exports from the File Services link on the replication process **Home** Page. Since there may be many individual settings for shares and exports on that page, it is recommended that you choose the **Configure** link on this page instead, which automatically configures the file system to allow shares and exports without further steps.

Using object replication schedules

Adding an object replication schedule


Schedules when an object replication runs.

Procedure

1. Navigate to **Home > Data Protection > Object Replication** to display the **Object Replication** page.

2. Under the Schedules section, click **add** to display the **Add Object Replication Schedule** page.

Field/Item	Description
Policy	Allows you to identify the policy to which this schedule will apply.
EVS / File system	Displays the current source EVS and file system for the object replication.
Replication Policy	Selects the object replication policy to which this schedule will apply.
Initial Run	Allows you to specify when the policy should run for the first time.
Immediately: Start as soon as the schedule is created	Runs the associated policy as soon as the schedule is successfully created.
Scheduled	Runs the associated policy at the date and time specified in this section. <ul style="list-style-type: none"> ▪ Time of Initial Run: Specify the time, using the 24-hour format (such that 11:59 PM will be entered as 23:59). ▪ Date of Initial Run: Specify the date for the first run of the policy. Click the calendar next to the field, then select the start date for the policy's initial run.
Current Date and Time	Provided for reference.
Run Until (Optional)	Allows you to specify a date and time after which the policy should no longer run.
Run Until Time	In this edit box, specify the last time (in 24-hour format) that the policy should be run. If you specify a time, you must also specify a Run Until Date .
Run Until Date	The date (year, month, and day) the replication runs for the last time. Click the calendar next to the field, then select the end date for the policy's final run. The selected date appears on the field. This is an optional setting.
Schedule Type	Allows you to specify the interval, if any, between the repeated execution of the policy to which this schedule will apply.

Field/Item	Description
	<p>You can select any one of the following options:</p> <ul style="list-style-type: none"> ▪ Every X minutes, hours, or days: Schedules the replication to run at the specified interval. For example, if you set it to every 4 days, the policy will run again automatically 4 days after it last started. ▪ Continuous. Pause X minutes, hours, or days between runs: Schedules the replication to run continuously, but the replication will pause between runs for the specified duration. For example, if you set it to pause for 1 day between runs, after the policy completes one cycle, it will pause for 1 day. ▪ Once, at the scheduled date and time: The policy is scheduled to run only once, at the date and time specified in the Initial Run settings. ▪ Test Only - at the scheduled date and time causes the replication policy to be tested: The object replication policy runs once, as a test only, at the time scheduled in the Initial Run field. During a test run, the system assesses if the object replication policy will run successfully as currently configured. The test also calculates the amount of data to be replicated. The results should be checked in the object replication Status & Reports page before scheduling an actual run. <div style="background-color: #e0f2f1; padding: 10px; margin-top: 10px;"> <p> Note: A test may take a long time to run, depending on the size of the files system being replicated. Additionally, the results of a test run are not displays on the Status page. Only actual replication results are shown on the Status page; however, if you schedule and run the policy as a test only, error messages appear if the test fails.</p> </div>

3. Enter the requested information.
4. Click **OK**.

Modifying an object replication schedule


Procedure

1. Navigate to **Home > Data Protection > Object Replication** to display the **Object Replication** page.

- Click **details** next to the schedule that requires modification, in order to display the **Modify Object Replication Schedule** page.

Field/Item	Description
Details	
Policy	Displays the name of the replication policy with the schedule that is being modified.
Policy Status	Displays the status of the last run of this policy.
Schedule Enabled	Indicates if the policy schedule is currently enabled or disabled. If the schedule is disabled, the policy will not be run automatically. If the schedule is disabled, click enable to reactivate (enable) the policy. If the schedule is enabled, click disable to deactivate (but not delete) the policy.
Next Run	
Reschedule	<p>Fill this check box to change the schedule of the next replication specified by this schedule:</p> <ul style="list-style-type: none"> ▪ Immediately: Start as soon as the schedule is created runs the associated policy as soon as the schedule is successfully created. ▪ Scheduled schedules the next run of the associated policy for the date and time specified in this section. Specify the time, using the 24 hour format (such that 11:59 PM will be entered as 23:59). Specify the date for the first run of the policy. Click the calendar next to the field, then select the start date for the policy's initial run.

Field/Item	Description
	The current date and time are provided at the bottom of the section for reference.
Run Until	This optional section allows you to specify a date and time after which the policy should no longer run.
Run Until Time	Specifies the last time (in 24-hour format) that the policy should be run. If you specify a time, you must also specify a <code>Run Until Date</code> .
Run Until Date	The date (day and month) the replication runs for the last time. Click the calendar next to the field, then select the end date for the policy's final run. The selected date appears on the field. This is an optional setting.
Schedule Type	<p>Specifies the interval, if any, between the repeated execution of the policy to which this schedule will apply.</p> <ul style="list-style-type: none"> ▪ Every X minutes, hours, or days schedules the replication to run at the specified interval. For example, if you set it to every 4 days, the policy will run again automatically 4 days after it last started. ▪ Continuous. Pause X minutes, hours, days, weeks, or months between runs schedules the replication to run continuously, but the replication will pause between runs for the specified duration. For example, if you set it to pause for 1 day between runs, after the policy completes one cycle, it will pause for 1 day.

Field/Item	Description
	<ul style="list-style-type: none"> ▪ Once, at the scheduled date and time schedules the policy to run only once, at the date and time specified by the <code>Initial Run</code> settings. ▪ Test Only - at the scheduled date and time causes the replication policy to be tested runs the policy once, as a test only, at the time scheduled in the <code>Initial Run</code> field. During a test run, the system assesses if the object replication policy will run successfully as currently configured. The test also calculates the amount of data to be replicated. The results should be checked in the Object Replication Status & Reports page before scheduling an actual run. <div style="border: 1px solid #add8e6; padding: 10px; margin-top: 10px;"> <p> Note: A test can take a long time to run, depending on the size of the files system being replicated. Additionally, the results of a test run are not displays on the status page. Only actual replication results are shown on the status page; however, if you schedule and run the policy as a test only, error messages appear if the test fails.</p> </div>


3. Modify the schedule as necessary.
4. Click **OK**.

Displaying object replication policies

Displays the replication policies and schedules you have created, and allows you manage those policies and schedules.

Procedure

1. Navigate to **Home > Data Protection > Object Replication** to display the **Object Replication** page.

Field/Item	Description
Policies	
Name	Identifies the replication policy.
Source	Source of the replication, identified using: <ul style="list-style-type: none"> ▪ EVS: Name of the EVS on the source server that the replication source file system is owned by. ▪ File System: Name of the replication source file system.
Target	Destination of the replication: <ul style="list-style-type: none"> ▪ EVS: Name of the EVS on the target server that the replication target file system is bound to. ▪ File System: Name of the target (destination) file system.
Status	Light indicator and short status message for successful and failed replication jobs. For the indicators: <ul style="list-style-type: none"> ▪ Green indicates that a successful replication job has completed. ▪ Red indicates a failed replication job and lists the reason for failure. ▪ Gray indicates replication job for which no status information could be found. Either the replication has never been run, or status information is not available. <div style="border: 1px solid #ccc; background-color: #e0f2f1; padding: 5px; margin-top: 10px;">  Note: In the case of a replication failure, the next time a replication starts, the data management engine attempts to restart the failed replication instead of starting a new replication. </div>
details	Displays the details of the selected policy.
add	Advances to Add Object Replication Policy page.
remove	Fill the check box next to each policy you want to remove, and then click remove .
run now	Fill the check box next to each policy you want to run, and then click run now .
abort	To abort one or more running replication operations, fill the check box next to the policy or policies you want to abort, and then click abort .
Object Replication Status & Reports	Advances to the Object Replication Status & Reports page.


Field/Item	Description
Schedules	
ID	ID assigned to the replication policy. The ID is not unique, it is per policy. For example, the first schedule for each policy will both have an ID of 1, the second schedule for each policy will both have an ID of 2 etc...
Policy	Name of the replication policy.
Next Run	Month, date, year and time for the next scheduled replication run for this policy.
Interval	Frequency at which the replication has been scheduled to run.
details	Displays the details of the selected schedule.
add	Advances to the Add an Object Replication Schedule page.
remove	Fill the check box next to the schedule you want to remove, and then click remove .

Asynchronous replication from a secondary to a tertiary target

Asynchronous Replication from a secondary to a tertiary target allows a replication target to be used as a source for another replication (the tertiary target). Its file systems can also be browsed in a read-only state and be replicated to the tertiary target.

Benefits of Asynchronous Replication

This feature has the benefit of performing replications to a remote target via a WAN link without making heavy demands on the resources of the main server. Asynchronous Replication can offload the network demands of the slower WAN link away from the main server. Typically, replications done over a WAN link can cause frequent interruptions to the main read/write server and tax its resources. In fact, backing up larger file systems to a remote target over a WAN link can take hours and make heavy demands on the main server's resources. In Asynchronous Replication, however, a chain is created wherein the backup from the main server to the initial target (called the secondary target) is completed first. Then, the backed up data from the secondary target is replicated via a WAN link to the tertiary target, without making frequent demands on the CPU of the main read/write server. With fewer resource demands on the main server it is possible to replicate to the secondary target more often than to the tertiary replica.



Configuring tertiary replication targets

Tertiary replication targets are configured using the **Object Replication** page of the NAS Manager. This procedure enables you to configure an additional (tertiary) replication target using an existing replication target as the source.

Definitions

- Primary - the original data source
- Secondary - the first replication target
- Tertiary - the final replication target

Pre-requisites

This procedure requires an existing (primary to secondary) replication policy which uses manual snapshot rules at both ends. Ensure that a snapshot has been created using the rule on the source and that at least one replication has completed using the policy before starting this procedure. Automatic snapshot rules are not supported with tertiary replication.

Procedure

1. Navigate to **Home > Data Protection > Object Replication**, and click **add** to open the **Add Object Replication Policy** page.




Data Protection [Home](#) > [Data Protection](#) > [Object Replication](#) > Add Object Replication Policy

Add Object Replication Policy

Identification	
Name:	Tertiary
Source	
EVS / File System:	HNAS-EVS2 / hsr-igt <input type="button" value="change..."/>
EVS IP Address:	Automatically selected
Target	
Server:	NFS-Auto-4060
<small>Click "select a target..." to choose an EVS and file system.</small>	
EVS:	HNAS-EVS3 <input type="button" value="select a target..."/>
EVS IP Address:	172.16.19.33
File System:	hsr-igt-3
Object Replication Listening Port:	59650
<input type="button" value="next"/> <input type="button" value="cancel"/>	

The following table describes the fields on this page:

Field/Item	Description
Identification	
Name	The name of the object replication policy.
Source	
EVS/file system	The name of the source EVS and file system that is to be replicated to the object replication target. To change the source EVS/file system, click change....
EVS IP Address	The default value for this field is 'Automatically selected'. If required, select the IP address for the source EVS from the list.
Target	
Server	Select a server from the list to be the target of the object replication policy. After selecting a server, click select a target... to select an EVS and file system. For managed servers, the EVS Name and File System fields are automatically populated when you select a file system using select a target... Or, alternatively, type the EVS IP/host name and file system name in the corresponding fields.
EVS DNS Name or IP	Select the EVS IP/host name in this field.
EVS IP Address	Enter the EVS IP address in this field.

Field/Item	Description
File System	<p>Select the target file system of the object replication policy. If the replication process is not managing the server on which the target file system is hosted, or if the destination file system does not yet exist, select Specify EVS IP/host name and file system from the list and type the details in the appropriate fields.</p> <p> Note: The replication target file system should be at least as large as the source file system to ensure that all data can be replicated on the target. If you intend to keep multiple snapshots on the target, it is especially important that the target be larger than the source, because the additional snapshots on the target file system will require storage space.</p> <p> Note: The tiering of both the replication source and target file systems must agree; you cannot have a tiered source file system replicated to an untiered file system, and you cannot have an untiered source file system replicated to a tiered file system.</p>
Object Replication Listening Port	<p>The port on which the destination server is listening. The default is 59550. The port on which the destination server listens is configured on the Object Replication Configuration page of the NAS Manager, or through the <code>replication-listening-port</code> CLI command.</p> <p> Note: To change the listening port for the target server, you must make that server the currently managed server of the replication process, then use the Object Replication Configuration page of the NAS Manager to change the listening port.</p>

2. In the **Identification** section, enter a name for the replication policy in the **Name** field.
3. In the **Source** section, select the secondary EVS/File System and the EVS IP Address.
4. In the **Target** section, select the tertiary server from the **Server** drop down list.
5. Click the **select a target** button. Use the **Select a File System** drop down menu to select the file system/EVS that you want to use as the tertiary replication target. The **File System** field of the **Target** section now shows the selected file system. Click **Next**.


EVS:

EVS IP Address:

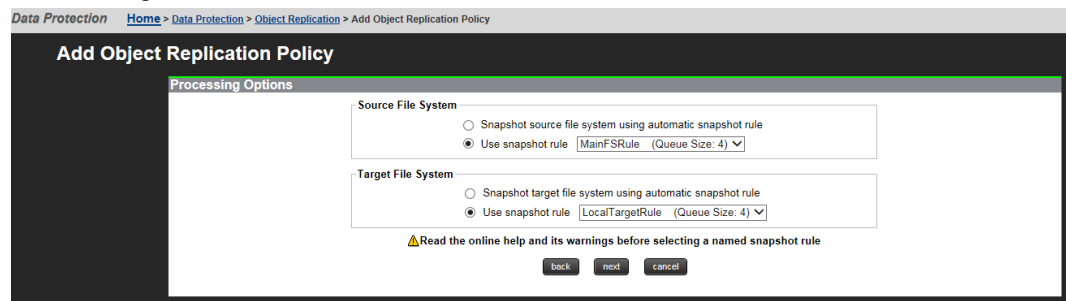
File System:

6. If access point transfer was not enabled on the source file system (or set during a replication policy creation), an information page appears which explains that the source

file system is not configured to allow its access points to be transferred with an object replication. For the tertiary policy, no action is necessary. Click **Continue**.




- You are now prompted to select the snapshot rules to be used by the source (secondary) and target (tertiary) file systems. Check the **Use Snapshot rule** buttons in the **Processing Options** section and select a manual snapshot rule for both the source and the target. Click **Next**.




Warning: Never use automatic snapshot rules when setting up secondary to tertiary targets. When using multiple replications in a chain, a snapshot of the source is not automatically propagated to the secondary and tertiary targets at the same time. While the replication from primary to secondary is underway, the snapshots on the tertiary do not yet reflect the source. Once the primary to secondary replication is complete, and the subsequent secondary to tertiary replication is complete, at that point the snapshots on the three file systems will match, but not before. If an updated snapshot appears in a rule during replication, this snapshot is not used until the next scheduled replication. If you are uncertain on how to proceed, it is recommended that you consult with your Hitachi Vantara technical support representative.


- If you haven't yet created an initial snapshot for the selected manual rule, an information page appears which explains that a baseline snapshot could not be established. Click **Continue** and then **OK**. After completing the setup of this policy, run a manual snapshot between the primary source and secondary target.



9. A summary page appears containing the details of the tertiary object replication policy. Click **create** to confirm the creation of the policy.




10. A message appears explaining that the policy has been successfully created. You are now prompted to create a policy schedule for the new policy. You can proceed without creating a schedule at this time, but note that you must use policy schedules when using the tertiary target for disaster recovery. You can create schedules later. If you want to create one now, click **yes** to create a policy schedule. For detailed information on schedules, see the sections *Using object replication policies* and *Using object replication schedules* in this guide.




Click **no** to return to the **Object Replication** page.

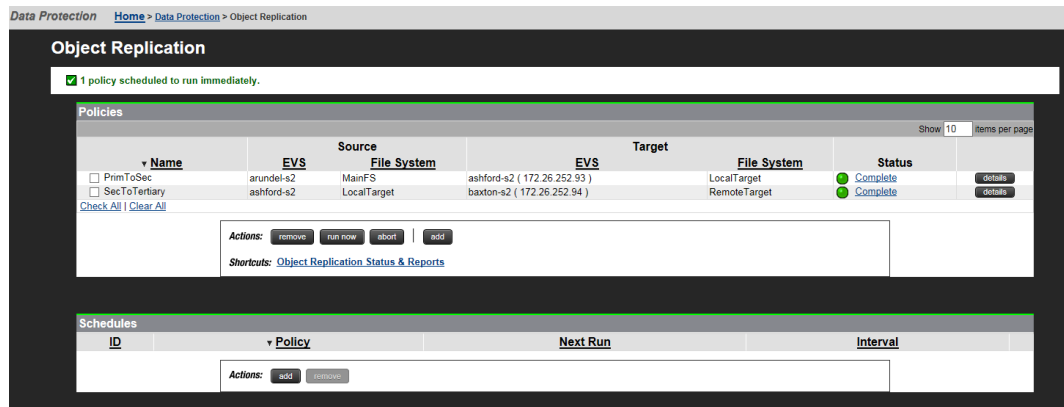
11. On the **Object replication** page, run and test the new replication policies as described in the example below. You can run the replications manually, without policy schedules.



Test and run the primary to secondary replication. After the policy has run successfully, the icon in the Status column is green and the status shows as "Complete".



Test and run the secondary to tertiary replication. After the replication has run successfully, both icons in the Status column are green.



Displaying object replication status and reports

This page displays detailed information about the status of object replication operations, as well as a list of reports for all policies for the selected file system. You may view reports for policies in all file systems on that EVS, or you can select one source file system to see all policy details associated with that file system only. Additionally, this page displays the status of each run of a policy, and whether the run was full or incremental.

Procedure

1. Navigate to **Home > Object Replication Status & Reports** to display the **Object Replication Status & Reports** page.

Field/Item	Description
File System Details	Displays the current source EVS or file system for the object replication. Click change in order to select a different EVS or file system.
Filter	Displays the name of the replication policy for which the report was created. To change the policy, select a policy name from the drop-down list in the Policy field and click filter .
Policy	The list of all object replication policies associated with the EVS and file systems shown in the EVS/file system field or the list of all instances of one policy, selected in the Policy list.
Source	The source file system and its associated snapshot rule.
Target	The target EVS, file system and its associated snapshot rule. You can sort and reverse the order of the snapshot list by clicking snapshot. The Start column describes the time that the object replication policy ran. You can sort and reverse the order of the list by clicking Time .
Start Time	Displays the time that the object replication policy was started.
Status	Displays the status of a run of the policy, showing whether an incremental or full replication has completed.
details	Displays a detailed log of all status and data for a particular replication policy that is running or has completed running.
delete reports	Select the reports to delete. A dialog box displays with a list of all object replication policies associated with the EVS. You may select All policies , or one specific policy. All reports in the selected policy's history are deleted. This option is only available when you select a file system. If all file systems are selected, the option is disabled.
download all reports	Downloads a <code>.csv</code> file containing data and status of all the object replication policies on the server. The downloaded file is displayed as a spreadsheet.
Object Replication Policies and Schedules	This shortcut link takes you directly to the Object Replication page.

2. Click **details** to display the details of a report.

Chapter 5: Using storage-based mirroring

Storage-based mirroring enables the NAS server to maintain two or more copies of its data. This can be useful when performing maintenance and also provides resiliency.

Storage-based mirroring overview

Some storage can improve resiliency by maintaining two or more copies of the data written by the server. The Administrator uses a storage configurator application such as Storage Navigator or Hitachi Command Suite to make mirrored copies of all the system drives (SDs) in a span. During the initial configuration, the Administrator licenses the mirrored SDs, then manually sets up the mirror relationships. Alternatively, it is possible for the server to detect the mirror relationships automatically.

Once this initial configuration has been carried out, the server can respond to storage failover without manual intervention. If a primary SD fails or becomes secondary, the file systems are unmounted. As soon as the server sees a complete set of primary SDs and no unexpected primary-primary relationships, the file systems become mountable.

The NAS server can, optionally, warn if any secondary SD fails and failover becomes impossible. The Administrator has the option of allowing or forbidding partial failovers, where the server mounts file systems on a mixture of production and backup storage. The NAS server also allows other configurations, such as mounting two copies of the same file system on different clusters.

Storage-based mirroring terminology

When mirror relationships are in place, the storage automatically mirrors all data from SDs that it calls P-Vols to SDs of the same capacity called S-Vols, causing no overhead on the server.

In most circumstances, a primary SD is a P-Vol and a secondary SD is an S-Vol. The two exceptions are:

- If storage-based snapshots are in use, V-Vols (SDs on which snapshots of spans reside) are also considered to be primary SDs, because the server mounts file systems on them. However, storage-based snapshots and storage-based mirrors cannot be used together on the same span. For further information, see the `storage-based-snapshot` CLI man page.
- If the P-Vols are destroyed, the Administrator can place the mirror relationships into a special SSWS state. In this situation, the Administrator configures the server to treat the surviving S-Vols as primary SDs and perform I/O to them.

The server always does I/O to primary SDs and never to secondaries, although it is capable of reporting and warning about secondary SDs' statuses.

There are various types of mirrors as follows:

- A **synchronous** mirror, such as TrueCopy, copies data as soon as the server writes it. In a healthy system, P-Vols and S-Vols have identical contents. If the P-Vols are lost, the server can fail over to the secondaries with a break in service but with no loss of data.
- An **asynchronous** mirror, such as TrueCopy Extended Distance, copies data from P-Vols to S-Vols at fixed times or on demand. Failover discards all writes since the data was last mirrored. However, an asynchronous mirror needs less bandwidth than a synchronous mirror and may offer improved performance.
- A **near-synchronous** mirror, such as Hitachi Universal Replicator (HUR), copies data from P-Vols to S-Vols as soon as possible, but not instantaneously. Failover typically discards some recently written data, but less than with an asynchronous mirror. A near-synchronous mirror offers the same bandwidth and performance advantages as an asynchronous mirror. ShadowImage works in this way if mirrored pairs are kept synchronized.
- A **copy-on-write** mirror, such as Hitachi CoW, efficiently stores one or more snapshots of each SD. The server does not usually fail over to these snapshots; instead, snapshots can efficiently be copied back to the SDs used by the server. The advantages of CoW are its efficient use of disk space and the speed with which a snapshot can be created. ShadowImage works in this way when its point-in-time snapshot capability is in use. For information on using copy-on-write mirrors in which multiple snapshots of a span are visible at once and multiple snapshots of a filesystem are mounted simultaneously, see the `storage-based-snapshot` CLI man page. Storage-based mirrors cannot be set up on a span that is in snapshot mode.
- A **Global Access Device** mirror does not require the Administrator to configure GAD mirror relationships into the server. GAD failovers are non-disruptive and do not require file systems to be unmounted.

To distinguish clearly between SDs during failover, we refer not only to primary and secondary SDs but also to production and backup SDs. Production SDs are used for everyday work; backup SDs are used if the production SDs should ever fail or whenever the administrator wishes to test the failover procedure. In normal use, production SDs are P-Vols and treated as primaries and backup SDs are S-Vols and treated as secondaries. Typically, if a problem occurs, the Administrator makes the NAS server treat the backup SDs as primaries and perform I/O to them. This involves either swapping mirror roles so that the production SDs become P-Vols or moving the mirror relationships into the SSWS state and temporarily configuring the S-Vols as primaries.

Configuring storage-based mirrors

This section describes how to set up and use a storage-based mirror. Mirrors can be configured at any time: before the file systems are first mounted, or on a mature system whose resiliency needs to be increased. Any mounted file systems can remain mounted.



Note: A 'storage configurator' refers to a product such as Storage Navigator or Hitachi Command Suite. If you are setting up a Copy-on-Write mirror and you do not plan to mount file systems on snapshot SDs, or if your storage does not offer that facility, skip steps 2 to 5.

To set up a storage-based mirror

Procedure

1. Use your storage configurator to select a mirroring type (synchronous, asynchronous, etc...) and set up P-Vols and S-Vols. If mirrors are asynchronous, near-synchronous or CoW, place all the new mirror relationships into a single consistency group per span, so that the S-Vols always form a coherent point-in-time snapshot of the P-Vols. Wait until mirrors have completely synchronized. This can take many hours. Your storage configurator can report progress.
2. If the S-Vols are intended for use on the same cluster as the P-Vols, use the NAS server `sd-list` command to check that both primary and secondary SDs are healthy and visible to the server. Use 'sd-allow-access' to license (grant the server permission to use) all the SDs, including the S-Vols.
 - a. If the P-Vols are intended for use on one cluster and the S-Vols on another, license on each cluster only the SDs to which that cluster is connected.
 - b. In an advanced setup with more than two copies of the data, license only two copies: the primary SDs and the SDs to which you want this server to fail over.
 - c. In advanced setup with two or more copies of the data, where one copy is mounted by a different cluster, license just two copies: the primary SDs and the SDs on which the other cluster mounts file systems. The other cluster's SDs will be unlicensed in a separate step.
 - d. Do not license more than two copies of the data. This system does not permit more than two-way failover. There can be any number of copies of the data, but only two copies of each span can be loaded and only two copies of each file system can be mounted.

3. Provided that only the P-Vols and one set of S-Vols have been licensed and both are visible to the same server, use the `sd-mirror-prepare` and `sd-mirror-detect` commands to instruct the server to detect the mirror relationships set up in step 1. If the P-Vols are visible to one cluster and the S-Vols to another cluster, use the `sd-mirror-remotely` command instead. In exceptional circumstances, you may need to configure mirror relationships into the server manually. The `sd-metadata` command can help to identify relationships; the `sd-list --raid-names` command displays the same identifiers that storage configurators use to identify SDs. In complex configurations, set up relationships with the S-Vols that this or another cluster can eventually access -- even if those are not the S-Vols with an immediate mirror relationship to the P-Vols. The two SDs in a mirror relationship must be in the same tier, or both must be in no tier at all.
4. If mirrors are not synchronous, or if you want to forbid partial failovers for any other reason, use the `span-set-site-id` command to place the production SDs at site 1 and the backup SDs at site 2. There is no need for the two copies of the data to be geographically separate. Although it is not mandatory, we recommend the use of SD sites even with synchronous mirrors, because the server can respond to SD failovers faster than Storage Navigator can fail over all the SDs in a large span. Without SD sites, the result is a series of partial failovers, as the server repeatedly tries to mount file systems on a mixture of production and backup SDs.
5. If the filesystems on the S-Vols are to be mounted by another cluster:
 - a. On the production cluster, which is connected to the P-Vols, use the `span-add-cluster-uuid` command to add the remote cluster's UUID to the span, so that each cluster can mount its own copy of the filesystems.
 - b. Mounting the same copy of the file systems on two clusters can cause data loss. Therefore, on the production cluster, unlicense the S-Vols using the `span-deny-access --secondaries` command. Zone your switches to ensure that each cluster can see only its own SDs. No SD should be visible to both clusters.
 - c. If necessary, synchronize mirrors, so that the secondary SDs contain on-disk configuration that includes both clusters' UUIDs.
 - d. Break the mirror relationship nearest to the remote cluster, making its SDs into simplex P-Vols.
 - e. Walk to the second cluster. Use the `sd-allow-access` command to license the second cluster's SDs. That cluster can now mount its copy of the file systems. Ensure that, on the backup cluster, you do not license the production cluster's SDs.

Updating a mirror

Updating a mirror depends on whether it is synchronous or asynchronous.

Updating an asynchronous or Copy-on-Write mirror

There is no need to unmount file systems for this type of mirror. After ensuring that all the mirror relationships in the span are in a single consistency group, use your storage configurator to take or update a copy of the data. This action is usually scripted and run from a cron job.

Updating a synchronous mirror

Synchronous mirrors are automatically updated every time the server writes to primary storage. No manual procedure is required.

Performing failover on a mirror

Failing over is useful as a test procedure or in order to minimize disruption while production storage is under maintenance. This procedure does not apply to GAD. GAD failover is invisible to the NAS server.



Caution: Failover always involves a brief loss of service.

To fail over a mirror

Procedure

1. Unmount all file systems on the span you want to fail over.
2. Use your storage configurator to ensure that the S-Vols contain an up-to-date copy of the data on the P-Vols. If relationships are synchronous, this means ensuring that the relationships are healthy; otherwise, it can mean running an update command or script.
3. Use your storage configurator to promote the current S-Vols to the primary role and demote the current P-Vols to the secondary role.
4. Run the `pn all span-list` command until it shows that the span is healthy again throughout the cluster.
5. Remount the filesystems. To fail back, repeat the procedure.

Recovering from loss of production SDs

To recover production SDs which are lost while file systems are mounted

Procedure

1. If you do not wish to mount file systems at once, use the 'automount' command to disable automount.
2. Use your storage configurator to break the mirror relationships if the production SDs no longer exist, or to promote the backup SDs to primaries if the production SDs exist but have failed. If, as a result, mirror relationships move into the SSWS state and the surviving SDs remain as S-Vols, use the `sd-peg` command to make the server treat them as primaries. Remember to unpeg them (by pegging them back to the default role) as soon as the SSWS state is resolved.
3. Run the `pn all span-list` command to confirm that the span is healthy. If automount is enabled, file systems mount automatically. If they do not, or if automount is disabled, use the `mount` command to mount them. Before doing so, you can use the `check-directory` command to confirm that the backup file systems are intact.

4. This step depends on whether the mirrors are synchronous:
 - If mirrors are synchronous then, as soon as you mount the file systems on your backup SDs, the server replays any uncommitted writes held in NVRAM, and the data on the production SDs must be considered out of date. If at all possible, avoid manually discarding NVRAM data during your recovery procedure, because it results in the loss of writes that the server has acknowledged to clients but not yet written to disk.
 - If mirrors are near-synchronous or asynchronous, the data on disk is too old to permit NVRAM data to be replayed. In addition, recent changes to on-disk configuration (such as filesystem expansions) may not have been copied to the backup SDs. Start by rescanning the on-disk configuration using the `sd-rescan-cod <span-instance-name>` command. Next, try to mount the filesystems. If they do not mount in the usual way, discard the NVRAM data and then mount by force. Some recent writes will be lost.
5. Once file systems are mounted, either repair your existing production SDs or create new mirrored SDs for eventual use as production SDs; synchronize, then schedule an orderly failback. Re-enable automount.

Copy-on-Write copyback

CoW systems differ from the other mirroring types. Instead of failing over from primary SDs to secondaries, the administrator copies the contents of an old snapshot to the production SDs. If the P-Vols are lost, S-Vols cannot be used for recovery, because S-Vols are stored as deltas against the P-Vols. Recovery is, therefore, used for recovery from an application-level problem, such as a virus attack or an unwanted database update, rather than to recuperate from physical damage to the storage.

To recover production SDs

Procedure

1. Unmount all file systems on the span.
2. Use your storage configurator to copy all the SDs in one snapshot over all the SDs in the span. If the source SDs are not in a single point-in-time snapshot of the span, or if some production SDs are updated and some not, the data will be lost.
3. On the server, use the `sd-rescan-cod <span-instance-name>` command to reload the on-disk configuration, which may have changed.
4. Use the `mount` command to mount the filesystems.

Removing a single mirror relationship

This procedure is useful for replacing a failing secondary SD. Coupled with a clean failover and failback, it can also be used to replace a failing primary SDs without data-loss.

To remove a single mirror relationship

Procedure

1. Use the `sd-span-cod-reads` command to instruct the server not to load the on-disk configuration. This prevents severe events being logged when the mirror relationship is broken.
2. Use the `sd-mirror S P=none` command (where S is the span's instance name and P is the device ID of the P-Vol), to instruct the server to remove the S-Vol from the span.
3. Use your storage configurator to break the mirror relationship.
4. Do one of the following:
 - Use your storage configurator to delete the failing SD, and then use the `sd-forget` command to remove it from the registry.
 - Use the `sd-deny-access` command to unlicense the failing SD.
 - Use the `sd-wipe-cod-signature` command to erase the Cod signature on the failing SD.
5. Use the `sd-span-cod-reads --resume` to tell the server to resume loading the on-disk configuration. It is usual to set up a new mirror SD and then use either the `sd-mirror-prepare` and `sd-mirror-detect` commands or the `sd-mirror` command in order to incorporate it into the span, as detailed above.

Removing all mirror relationships

This procedure is useful for replacing all the backup SDs in a span.

To remove all mirror relationships**Procedure**

1. To prevent the server from trying to load the on-disk configuration from the S-Vols in a later step, use the `span-deny-access --secondaries` command.
2. To remove the mirror relationships from the span, use the `span-break-mirrors` command.
3. Use your storage configurator to break the mirror relationships.
4. Do one of the following:
 - Use your storage configurator to delete or reinitialize the former S-Vols, leaving the production SDs intact.
 - Use the `sd-wipe-cod-signature --danger-ignore-sd-licensing` command to erase the on-disk configuration signatures on the former S-Vols, so that they do not confuse the server.
 - Use the `sd-span-cod-reads --resume` command to tell the server to resume loading the on-disk configuration.
5. To reuse the former S-Vols on this cluster, use the `sd-allow-access` command to relicense them.

Mounting two copies of the data on different clusters

This configuration is useful for performing off-line data mining on recent copies of live data.

The configuration appears as follows:

- Site A contains Cluster A and Storage A.
- Site B has Cluster B, Storage B1 and Storage B2. Between B1 and B2 is a very high-bandwidth mirror link (typically, ShadowImage or TrueCopy); B1 and B2 can be in the same storage subsystem. The mirror relationship between B1 and B2 is broken during office hours. Cluster B loads spans from, and mounts file systems on, Storage B2. No cluster is connected to Storage B1.
- Between Storage A and Storage B1 is a mirror relationship over a medium- or high-bandwidth link: typically, TrueCopy, TrueCopy ED or HUR.

To bring Cluster B up to date with Cluster A

Procedure

1. If the mirror between Storage A and Storage B1 is asynchronous, or near-synchronous, synchronize it. Only changed blocks are copied across the WAN, and so the time taken is acceptable. If the mirror between A and B1 is synchronous, this step is unnecessary, but more bandwidth is required between the two sites in order to maintain acceptable performance at Site A.
2. On Cluster B, use the `span-prepare-for-resync` command.
3. Remake the broken mirror relationship between Storage B1 and Storage B2. Because of the high-bandwidth link between these two sets of storage, the time taken to perform the update is acceptable, but synchronization is not instantaneous.
4. Once synchronisation is complete, break the mirror link between B1 and B2, bringing B1 and B2 back into the simplex state. B2 now has a copy of the data on B1, which in turn is a copy of the data that was on A at the start of step 1.
5. On Cluster B, use the `span-reload-after-resync` command. If Cluster B mounted file systems on Storage B1, it would write to the storage, even if file systems were mounted read-only. These writes would make it impossible to do an incremental update over the link between Sites A and B. However, the fast link between B1 and B2 makes a complete resynchronization realistic. This is why three copies of the data are required.

Tolerating primary-primary relationships

If you use your storage configurator to break one or more mirror relationships, the server sees both production and backup SDs as primary. Not knowing which SDs to use, it fails the span.

On most systems, this behaviour is desirable. However, on certain complex systems such as the one described above, where the 'backup' SDs belong to another cluster and can legitimately be primary, the main cluster must not fail the span. Use the `span-deny-access --secondaries` command to make the server ignore backup SDs even if they become primary. It also makes it impossible for the server to fail over to those backup SDs.

Secondary SD warnings

On most configurations, the NAS server warns (using the event log and the `trouble` command) if S-Vols become unhealthy or uncontactable, such that failover is impossible. On some advanced systems, such as when the S-Vols belong to another cluster, these warnings are undesirable. The following commands instruct the NAS server not to warn about S-Vols in the span on which you run them:

- `span-deny-access --primaries`
- `span-deny-access --secondaries.`

The following commands instruct the NAS server to warn about S-Vols in the span on which you run them:

- `span-allow-access`
- `span-break-mirrors`
- `sd-mirror` (if used to break all the mirror relationships in the span)

To license all the SDs in a span but not warn about the health of backup SDs, use the `sd-allow-access` command instead of the `span-allow-access` command.

A mirrored span is always healthy as long as it has a complete set of licensed, primary, healthy SDs and no unexpected primary-primary relationships (a primary-primary relationship is expected and ignored if only one of the SDs is licensed). If you have set up SD sites (see above), then all the primary SDs must be in the same site. There is specifically no requirement for the secondary SDs to be healthy; they need not even still exist.

Dealing with SSWS

If production SDs are lost, recovery of the data on the backup SDs can leave them as S-Vols and put the mirror relationship into a state called SSWS. Uniquely, SSWS causes the S-Vols to become writable. But the server is unable to detect this state, and treats the S-Vols as secondary SDs in the usual way, refusing to perform I/O to them.

For the duration of the SSWS state, use the `sd-peg` command to override the SDs' normally-determined roles and make the NAS server treat S-Vols as primary SDs. Be sure to remove the peg using the `sd-peg --default-role` command as soon as the SSWS state ends. In all states not involving SSWS, the server must be allowed to detect SDs' mirror roles for itself.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA



HitachiVantara.com/contact