HITACHI
Inspire the Next

Hitachi NAS Platform

Replication and Disaster Recovery Administration Guide

Release 12.1

Hitachi Data Systems

MK-92HNAS009-03

Hitachi, Ltd., reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. This document contains the most current information available at the time of publication. When new or revised information becomes available, this entire document will be updated and distributed to all registered users.

Some of the features described in this document might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Data Systems Corporation at https://portal.hds.com.

Notice: Hitachi, Ltd., products and services can be ordered only under the terms and conditions of the applicable Hitachi Data Systems Corporation agreements. The use of Hitachi, Ltd., products is governed by the terms of your agreements with Hitachi Data Systems Corporation.

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Contact the Hitachi Data Systems Legal Department for any export compliance questions.

Contents

Preface

In PDF format, 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.

Document Revision Level

Revision	Date	Description
MK-92HNAS009-00	August 2012	First publication
MK-92HNAS009-01	June 2013	Revision 1, replaces and supersedes MK-92HNAS009-00.
MK-92HNAS009-02	April 2014	Revision 2, replaces and supersedes MK-92HNAS009-01.
MK-92HNAS009-03	September 2014	Revision 3, replaces and supersedes MK-92HNAS009-02.

Contacting Hitachi Data Systems

2845 Lafayette Street
Santa Clara, California 95050-2627
U.S.A.
https://portal.hds.com
North America: 1-800-446-0744

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.

Administration Guides

- *System Access Guide* (MK-92HNAS014)—In PDF format, this guide 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)—In PDF format, this guide provides information about administering servers, clusters, and server farms. Includes information about licensing, name spaces, upgrading firmware, monitoring servers and clusters, the backing up and restoring configurations.
- *Storage System User Administration Guide* (MK-92HNAS013)—In PDF format, this guide 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)—In PDF format, this guide 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)—In PDF format, this guide 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)—In PDF format, this guide provides information about the Data Migrator feature, including how to set up migration policies and schedules.
- *Snapshot Administration Guide* (MK-92HNAS011)—In PDF format, this guide provides information about configuring the server to take and manage snapshots.
- *Replication and Disaster Recovery Administration Guide* (MK-92HNAS009)—In PDF format, 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.
- *Antivirus Administration Guide* (MK-92HNAS004)—In PDF format, this guide 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)—In PDF format, this guide provides information about configuring the server to work with NDMP, and making and managing NDMP backups. Also includes information about Hitachi NAS Synchronous Image Backup.
- *Command Line Reference*—Opens in a browser, and describes the commands used to administer the system.

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 G1 Hardware Reference* (MK-92HNAS016)—Provides an overview of the second-generation server hardware, describes how to resolve any problems, and replace potentially faulty parts.

- *Hitachi NAS Platform 3080 and 3090 G2 Hardware Reference* (MK-92HNAS017)—Provides an overview of the first-generation server hardware, describes how to resolve any problems, and replace potentially faulty parts.
- *Hitachi NAS Platform 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 High-performance NAS Platform* (MK-99BA012-13)—Provides an overview of the NAS Platform 3100/NAS Platform 3200 server hardware, and describes how to resolve any problems, and replace potentially faulty parts.

Best Practices

- *Hitachi USP-V/VSP Best Practice Guide for HNAS Solutions* (MK-92HNAS025)—The HNAS practices outlined in this document describe how to configure the HNAS system to achieve the best results.
- *Hitachi Unified Storage VM Best Practices Guide for HNAS Solutions* (MK-92HNAS026)—The HNAS system is capable of heavily driving a storage array and disks. The HNAS practices outlined in this document describe how to configure the HNAS system to achieve the best results.
- *Hitachi NAS Platform Best Practices Guide for NFS with VMware vSphere* (MK-92HNAS028)—This document covers VMware best practices specific to HDS HNAS storage.
- *Hitachi NAS Platform Deduplication Best Practice* (MK-92HNAS031) —This document provides best practices and guidelines for using HNAS 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.
- *Brocade VDX 6730 Switch Configuration for use in an HNAS Cluster Configuration Guide* (MK-92HNAS046)—This document describes how to configure a Brocade VDX 6730 switch for use as an ISL (inter-switch link) or an ICC (inter-cluster communication) switch.
- *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 NAS Platform Storage Pool and HDP Best Practices* (MK-92HNAS048)—This document details the best practices for configuring and using HNAS storage pools, related features, and Hitachi Dynamic Provisioning (HDP).

1

Hitachi NAS File & Object Replicator data replication

Hitachi NAS File & Object Replicator 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, CIFS shares, permissions and all other file-level metadata. Administrators can use Web 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. After the replication, the file-level replication can operate on file systems formatted using WFS-1 or WFS-2.
- 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. Files and directories are made up of objects, such as files, directories, security descriptors, snapshot lists, root directory, and many others. 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.

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 Hitachi NAS Platform 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 can operate only on file systems formatted using WFS-2.

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 is only available for use with WFS-2 file systems. WFS-1 formatted file systems can not be configured for object replication.
- Object replication works at the file system level only; entire file systems may be replicated using object replication, but individual files or directories cannot.
- 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 file system if CNS is used with object replication.

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 Web 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.

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. Pre-replication and post-replication scripts can also be set up in the **Policy** page.
- **Replication Rules**: Optional configuration parameters that allow tuning of 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, Hitachi Data Systems Support Center 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:** 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 replications, you should consider this limitation.

Incremental replication

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

- Upon establishing a replication policy, the SMU 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 replication*, the system copies to the target, in full, those files that have been changed since the last scheduled replication.
- 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. Note that, in order to use block-level replication, a Replication license is required.

A replication policy defines the properties of a replication, including a *replication rule* (source and target), and a *replication schedule*. Replication rules can be expanded to include optional settings. Pre-replication and post-replication scripts can also be configured.

Incremental data (file-level) replication

The server supports incremental data replication, performed under control of the System Management Unit (SMU). 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 SMU schedules and starts all replications, replicated data flows directly from source to target without passing through the SMU.

Incremental block-level replication

By default, *incremental data replication* copies files that have changed since the last replication. With the Block-Level Replication feature enabled, only data blocks in large files that have been written since the last replication are copied. Depending on the use of files within the source volume, this could substantially reduce the amount of data copied.

⚠ **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 should also alleviate performance problems caused by high speed WAN connections with high latencies. Connections with high latencies limit the throughput of a single TCP connection, because no data is sent during the time spent waiting for acknowledgments. These pauses in the sending of data result in an under-utilization of high speed WAN links. By using multiple TCP connections (one per stream), multi-stream replication addresses the problem of under utilization of the high speed WAN connections.

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 Web 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, CIFS 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*.

Supported replication applications

In addition to the built-in replication tools, Hitachi NAS Platforms and High-performance NAS Platforms support Hitachi Data Systems (HDS) replication applications (when used with HDS 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 HDS 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 HDS 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 HDS software interfaces, because the replication occurs between or within the HDS storage subsystems. Contact your HDS technical representative for assistance configuring these HDS features. For more information about these applications, contact Hitachi Data Systems Support Center or your HDS 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 Data Systems has added to the standard SCSI inquiry.
- Hitachi Data Systems 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 Data Systems 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 an HDS storage subsystem with another storage server (or another cluster), the system configuration must use HDS 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 HDS storage systems to determine their status.

Replication and disaster recovery

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

File system 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 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 (CIFS 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.

File system 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 synchronization.
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 (to maintain the replication, which will now go in the opposite direction).
11. Verify that the replication runs successfully.
12. Allow user access to the new primary file system.

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.

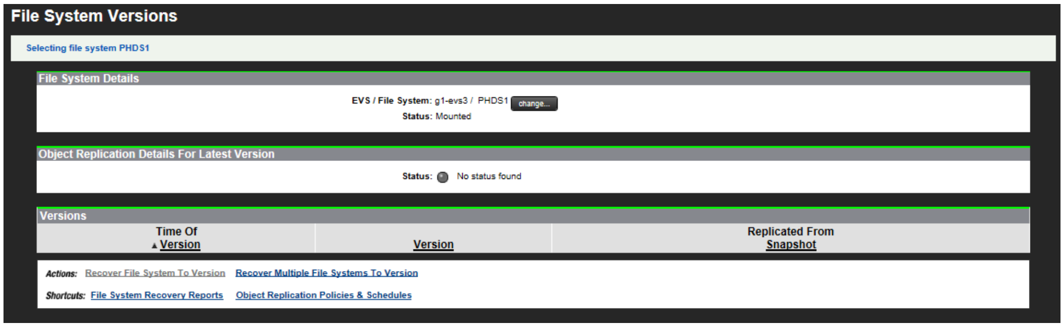

File System Versions

Selecting file system PHDS1

File System Details

EVS / File System: g1-evs3 / PHDS1

Status: Mounted

Object Replication Details For Latest Version

Status: ● No status found

Versions

Time Of ▲Version	Version	Replicated From Snapshot

Actions: Recover File System To Version | Recover Multiple File Systems To Version

Shortcuts: File System Recovery Reports | Object Replication Policies & Schedules

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 unmounted, mounted, or mounted as a replication target.
Object Replication Details	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 SMU), the Source File System, Source Server, and Source Server fields are not displayed.
Object Replication 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 "Not an object replication target" is displayed. The status indicator is green if a replication is currently running, or if the most recent object replication completed successfully. 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".
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 unmounted or mounted. 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.

Field/Item	Description
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 Versions** 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 Progress	Opens the file system **Recovery Progress** page.
Object Replication Policies & Schedules	Opens the main object replication page.

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, roll back the file system to the last successful replication snapshot. You must check the available snapshots so that you can choose the most recent (in order to reduce the amount of data lost by promoting 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 destination. 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 in which you select a recovery method, 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(s) 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 the file system *name* and**, select one of the following mount options for the recovered file system:

    - **mount read write**
    - **mount read only**
    - **not mounted**
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 CIFS shares or NFS exports for the recovered file system.

    - CIFS **shares**
    - NFS **exports**
      For NFS exports, 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 Progress** page.
7. Monitor the file system recovery.

    The following table describes the fields in this page:

Field/Item	Description
EVS/File System	Displays the name of the currently selected EVS and file system. Click **change** to select a different file system.
Time Started	Displays the date and time of the file system recovery operation.
File System	Displays the name of the file system being recovered.
Recovered to Snapshot	The snapshot to which the file system was recovered.
Recovery Option	Displays how the recovered file system will be mounted. One of the following messages is displayed:

Field/Item	Description
	• Mounting file system in read write • Mounting file system in read only • File system not mounted
Status	Displays the current status of the recovery operation.
details	Displays the recovery progress in more detail.
delete	Removes the selected progress report.
File System Versions	Opens the **File System Versions** page.

For more information about the recovery, click **details** to display the next **File System Recovery Progress** details page.

This page displays much more detailed information, and you can display recovery details about:

- The overall recovery status
- Mount request status
- CIFS Shares recovered, failed, and skipped
- NFS Exports recovered, failed, and skipped

This page also allows you to display the file system recovery log.

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

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 in which you select a recovery method, 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:** Syslock the file system as soon as the file system is recovered. 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	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 SMU), the **Source File System**, **Source Server**, and **Source Server** fields are not displayed.
Object Replication 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 "Not an object replication target" is displayed. The status indicator is green if a replication is currently running, or if the most recent object replication completed successfully. 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".
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 might be unmounted or mounted. If the currently selected file system is not a replication target, this field is not displayed.
The following steps will be taken	This section lists the recovery steps, and allows you to specify recovery options. **1.** Displays the file system to be recovered. **2.** 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

Field/Item	Description
	select a different snapshot by the date and time the snapshot was taken. 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. **3.** Displays the recovery goal to confirm that the file system you have chosen will be demoted to an object replication target. **4.** Check boxes allow you to specify if you want the access points (CIFS 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. ⚠ **Note:** In general, when demoting a file system as a part of disaster recovery, you should remove both CIFS shares and NFS exports.

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 Progress** page.
   - Click **cancel** to return to the **File Versions** page.
5. Monitor the file system recovery.

Field/Item	Description
EVS/File System	Displays the name of the currently selected EVS and file system. Click **change** to select a different file system.
Time Started	Displays the date and time of the file system recovery operation.
File System	Displays the name of the file system being recovered.
Recovered to Snapshot	The snapshot to which the file system was recovered.
Recovery Option	Displays how the recovered file system will be mounted. One of the following messages is displayed: • Mounting file system in read write • Mounting file system in read only • File system not mounted
Status	Displays the current status of the recovery operation.
details	Displays the recovery progress in more detail.
delete	Removes the selected progress report.

Field/Item	Description
File System Versions	Opens the **File System Versions** page.

For more information about the recovery, click **details** to display the next **File System Recovery Progress** page.

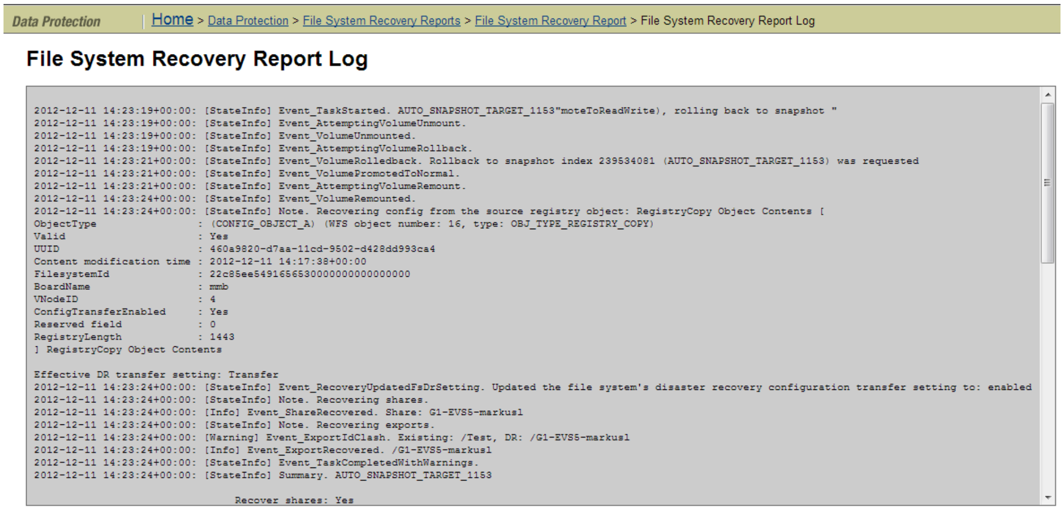

Data Protection
Home > Data Protection > File System Recovery Reports > File System Recovery Report > File System Recovery Report Log

File System Recovery Report Log

```
2012-12-11 14:23:19+00:00: [StateInfo] Event_TaskStarted. AUTO_SNAPSHOT_TARGET_1153"moteToReadWrite), rolling back to snapshot "
2012-12-11 14:23:19+00:00: [StateInfo] Event_AttemptingVolumeUnmount.
2012-12-11 14:23:19+00:00: [StateInfo] Event_VolumeUnmounted.
2012-12-11 14:23:19+00:00: [StateInfo] Event_AttemptingVolumeRollback.
2012-12-11 14:23:21+00:00: [StateInfo] Event_VolumeRolledback. Rollback to snapshot index 239534081 (AUTO_SNAPSHOT_TARGET_1153) was requested
2012-12-11 14:23:21+00:00: [StateInfo] Event_VolumePromotedToNormal.
2012-12-11 14:23:21+00:00: [StateInfo] Event_AttemptingVolumeRemount.
2012-12-11 14:23:24+00:00: [StateInfo] Event_VolumeRemounted.
2012-12-11 14:23:24+00:00: [StateInfo] Note. Recovering config from the source registry object: RegistryCopy Object Contents [
ObjectType                : (CONFIG_OBJECT_A) (WFS object number: 16, type: OBJ_TYPE_REGISTRY_COPY)
Valid                     : Yes
UUID                      : 460a9820-d7aa-11cd-9502-d428dd993ca4
Content modification time : 2012-12-11 14:17:38+00:00
FilesystemId              : 22c85ee549165653000000000000000
BoardName                 : mmb
VNodeID                   : 4
ConfigTransferEnabled     : Yes
Reserved field            : 0
RegistryLength            : 1443
] RegistryCopy Object Contents

Effective DR transfer setting: Transfer
2012-12-11 14:23:24+00:00: [StateInfo] Event_RecoveryUpdatedFsDrSetting. Updated the file system's disaster recovery configuration transfer setting to: enabled
2012-12-11 14:23:24+00:00: [StateInfo] Note. Recovering shares.
2012-12-11 14:23:24+00:00: [Info] Event_ShareRecovered. Share: G1-EVS5-markusl
2012-12-11 14:23:24+00:00: [StateInfo] Note. Recovering exports.
2012-12-11 14:23:24+00:00: [Warning] Event_ExportIdClash. Existing: /Test, DR: /G1-EVS5-markusl
2012-12-11 14:23:24+00:00: [Info] Event_ExportRecovered. /G1-EVS5-markusl
2012-12-11 14:23:24+00:00: [StateInfo] Event_TaskCompletedWithWarnings.
2012-12-11 14:23:24+00:00: [StateInfo] Summary. AUTO_SNAPSHOT_TARGET_1153

                               Recover shares: Yes
```

This page displays much more detailed information, and you can display recovery details about:

- The overall recovery status
- Mount request status
- CIFS Shares deleted or that failed to delete
- NFS Exports deleted or that failed to delete

This page also allows you to display the file system recovery log.

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

6. Put the recovered file system into Syslock as soon as the file system is recovered.
   When Syslock is enabled for a file system, NDMP has full access to the file system and can write to it during a backup or replication, but the file system remains in read-only mode to clients using the file service protocols (NFS, CIFS, FTP, and iSCSI).

   Refer to the *File Services Administration Guide* for information on how to put the file system into Syslock mode.

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.

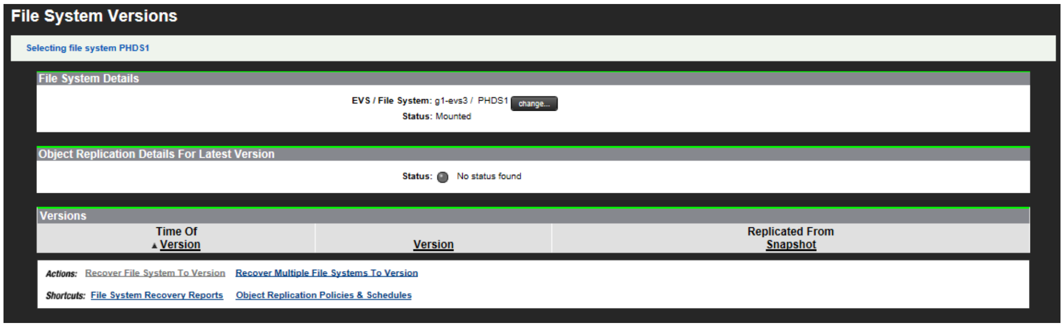

File System Versions

Selecting file system PHDS1

File System Details

EVS / File System: g1-evs3 / PHDS1

Status: Mounted

Object Replication Details For Latest Version

Status: ● No status found

Versions

Time Of ▲Version	Version	Replicated From Snapshot

Actions: Recover File System To Version Recover Multiple File Systems To Version

Shortcuts: File System Recovery Reports Object Replication Policies & Schedules

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 unmounted, mounted, or mounted as a replication target.
Object Replication Details	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 SMU), the Source File System, Source Server, and Source Server fields are not displayed.
Object Replication 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 "Not an object replication target" is displayed. The status indicator is green if a replication is currently running, or if the most recent object replication completed successfully.

Field/Item	Description
	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".
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 unmounted or mounted. 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 Versions** 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 Progress	Opens the file system **Recovery Progress** page.
Object Replication Policies & Schedules	Opens the main object replication page.

2. Click **Recover Multiple File System to Version** to display the **File System Versions** page.
3. Select the file systems to recover from the **Available File Systems** list, and click the right arrow to the add the file systems to the **Selected File Systems** list.
4. Specify how the recovered file system is to be mounted.

   - **read write**

- **read only**
- **not mounted**

5. 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 CIFS shares or NFS exports for the recovered file system.
   - CIFS **shares**
   - NFS **exports**
6. Click **next** to display the **Recover File System 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 Progress** page.
8. Monitor the file system recovery.

   The following table describes the fields in this page:

Field/Item	Description
EVS/File System	Displays the name of the currently selected EVS and file system. Click **change** to select a different file system.
Time Started	Displays the date and time of the file system recovery operation.
File System	Displays the name of the file system being recovered.
Recovered to Snapshot	The snapshot to which the file system was recovered.
Recovery Option	Displays how the recovered file system will be mounted. One of the following messages is displayed: • Mounting file system in read write • Mounting file system in read only • File system not mounted
Status	Displays the current status of the recovery operation.
details	Displays the recovery progress in more detail.
delete	Removes the selected progress report.
File System Versions	Opens the **File System Versions** page.

For more information about the recovery, click **details** to display the next **File System Recovery Progress** details page.

This page displays much more detailed information, and you can display recovery details about:
- The overall recovery status
- Mount request status
- CIFS Shares recovered, failed, and skipped
- NFS Exports recovered, failed, and skipped

This page also allows you to display the file system recovery log.

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

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.

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, CIFS shares, permissions and all other file-level metadata. Administrators can use Web 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

Prerequisites

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 SMU 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 SMU. 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.

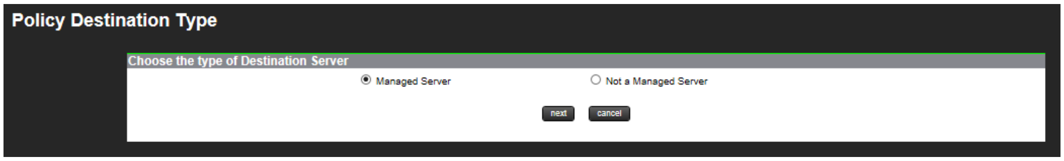

2. Select the policy destination type:
   - Select **Managed Server** to create a policy to replicate to a server that is managed by the SMU.
   - Select **Not a Managed Server** to create a policy to replicate to a server that is not managed by the SMU.
3. Click **next** to display a destination type-specific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

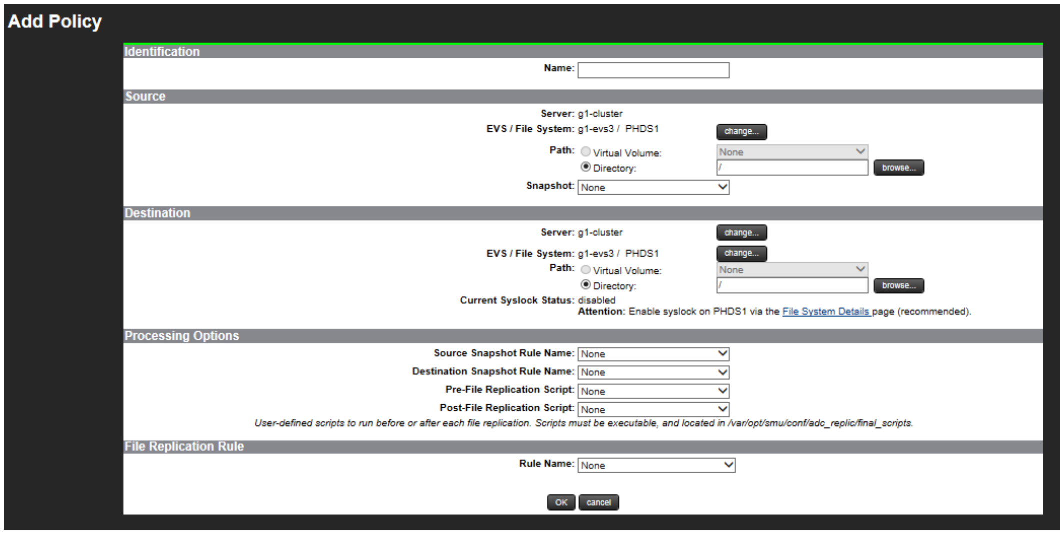

Add Policy

Identification

Name:

Source

Server: g1-cluster

EVS / File System: g1-evs3 / PHDS1

Path: ○ Virtual Volume: None
● Directory: /

Snapshot: None

Destination

Server: g1-cluster

EVS / File System: g1-evs3 / PHDS1

Path: ○ Virtual Volume: None
● Directory: /

Current Syslock Status: disabled
Attention: Enable syslock on PHDS1 via the File System Details page (recommended).

Processing Options

Source Snapshot Rule Name: None

Destination Snapshot Rule Name: None

Pre-File Replication Script: None

Post-File Replication Script: None

User-defined scripts to run before or after each file replication. Scripts must be executable, and located in /var/opt/smu/conf/adc_replic/final_scripts.

File Replication Rule

Rule Name: None

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

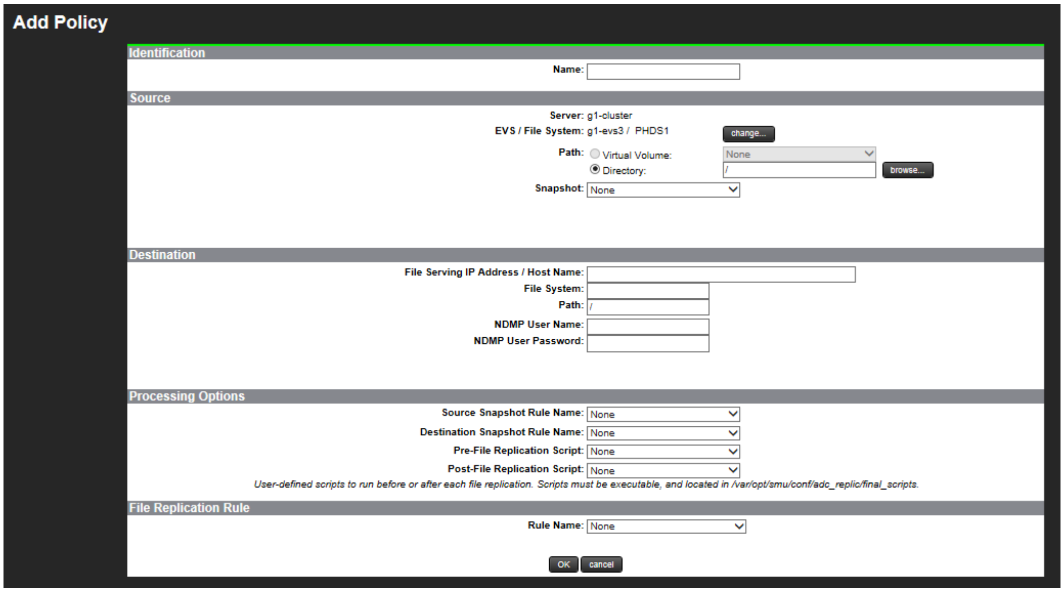

Add Policy

Identification

Name:

Source

Server: g1-cluster

EVS / File System: g1-evs3 / PHDS1

Path: ○ Virtual Volume: None
● Directory: /

Snapshot: None

Destination

File Serving IP Address / Host Name:

File System:

Path: /

NDMP User Name:

NDMP User Password:

Processing Options

Source Snapshot Rule Name: None

Destination Snapshot Rule Name: None

Pre-File Replication Script: None

Post-File Replication Script: None

User-defined scripts to run before or after each file replication. Scripts must be executable, and located in /var/opt/smu/conf/adc_replic/final_scripts.

File Replication Rule

Rule Name: None

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	**Name**: Allows you to specify the name of the replication policy. The name may not contain spaces or any of the following characters: `\/<>"'!@#£$%^%&*(){}[] +=?:;,~`|.'`

Field/Item	Description
Source	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: • **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)	Destination of the replication (managed server): • **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. • **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, CIFS, 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 *File Services Administration Guide*.
Destination (for unmanaged replication destinations)	Destination of the replication (non-managed server): • **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.

Field/Item	Description
Processing Options	• **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. • **Pre-/Post-Replication Script**: A user-defined script to run before or after each replication. Scripts must be located in `/opt/smu/adc_replic/final_scripts`. The permissions of the scripts must be set to "executable".
Replication Rule	Optional configuration parameters that allow tuning of replications to enable and disable specific functions or to optimize performance.

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

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 SMU. 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 SMU 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 SMU 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 SMU access by way of the management ports using the **`ndmp-management-ports-set`** command. The SMU connections then bypass the allowip/blockip checks.

The SMU 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 SMU. When this is the case, the **`ndmp-management-ports-set`** command can be used to request that the SMU 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 SMU 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 will affect 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 databases or other applications that require a consistent state at the time of a snapshot, best practices indicate 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.

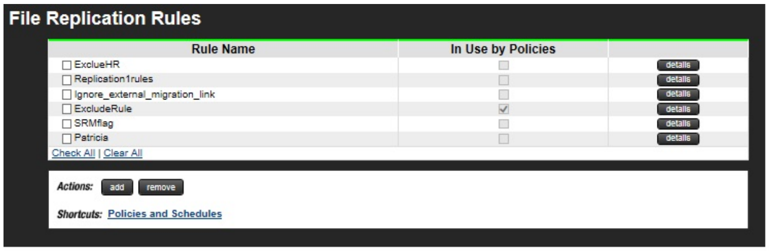

File Replication Rules

	Rule Name	In Use by Policies	
□	ExclueHR	□	details
□	Replication1rules	□	details
□	Ignore_external_migration_link	□	details
□	ExcludeRule	☑	details
□	SRMflag	□	details
□	Patricia	□	details

Check All | Clear All

Actions: add remove

Shortcuts: Policies and Schedules

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.
details	Displays the properties for a selected rule.
Shortcut	
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.

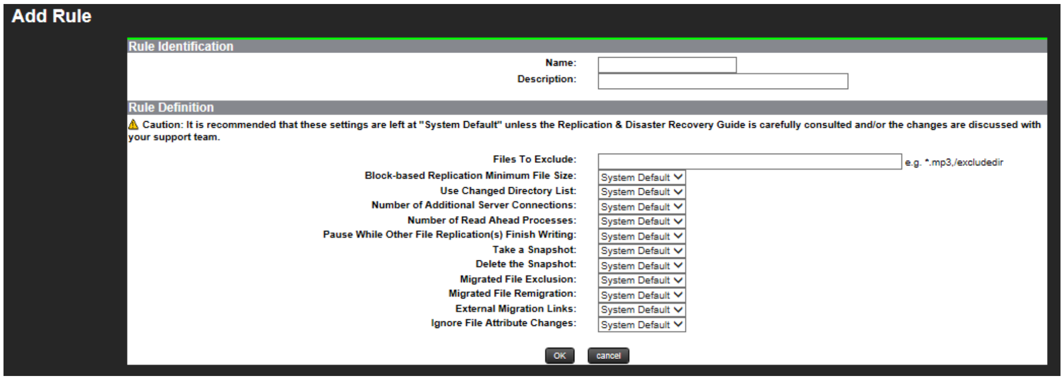

Add Rule

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: e.g. *.mp3,/excludedir

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

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: • 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.	None are excluded.

Field/Item	Description	Default
	• 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.	
Block Replication Minimum File Size	Block replication minimum file size controls the minimum file size that is used for block replication. The list options available are: 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: • 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. ⚠ **Note:** Requires a Replication license to function.	Minimum file size used for block replication is 32 MB.
Use Changed Directory List	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. Options: • **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). ⚠ **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. 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	**Use Changed Directory List** is disabled.

Field/Item	Description	Default
	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.	
Number of Additional Server Connections	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. **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. See Configuring NDMP performance options on page 56 for more information on server connections.	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. 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. 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.	If a value is not set, the default value is set by the application (depending on the number of read-ahead processes set in **Number of Additional Server Connections**).
Pause While Replication(s) Finish Writing	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	Set to **no**.

Field/Item	Description	Default
	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. 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: • 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.	
Take a Snapshot	Overrides the Backup configuration option `Automatic Snapshot Creation`. 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 live file system while it is changing may give an inconsistent copy. 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. • Enable this option to support incremental replication copies. • Disable only for full replication copies or when making a complete copy of a directory. 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. **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.	Snapshots are taken and backed up automatically.
Delete the Snapshot	Determines when snapshots are deleted. The setting for this option should be left as the	If the replication is an incremental replication, the application

Field/Item	Description	Default
	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 taken for the replication. The deletion options are: • **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. **Caution:** As changing these settings can adversely effect the replication process, Hitachi Data Systems Support Center recommends that this option be changed only at the direction of your Hitachi Data Systems representative.	automatically selects the correct setting.
Migrated File Exclusion	Indicates if the replications will include files whose data has been migrated to secondary storage. • **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	Controls the action at the destination when the source file had been migrated. • **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 file that has been migrated to an external server). • 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	**Remigrate** and **re-create link** are enabled.

Field/Item	Description	Default
	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. ⚠ **Note:** For externally migrated files, to make sure that the file or link is replicated properly, you should either: • Specify that the replication operation should remigrate files and the destination should test before recreating links (using the `migration-recreate-links-mode` command). • Specify that the replication operation should re-create links and the destination should always recreate links (using the `migration-recreate-links-mode` 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
OK	Configures the rule as defined, and returns to the **Replication Rules** page.	
cancel	Discards the entered values, and returns to the **Replication Rules** page.	

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: Hitachi Data Systems Support Center 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 than 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.

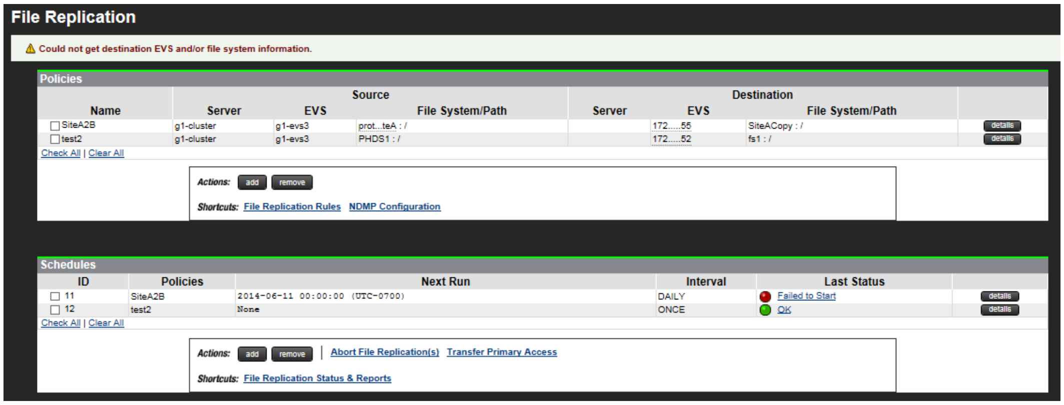

File Replication

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

Policies

	Source			Destination			
Name	Server	EVS	File System/Path	Server	EVS	File System/Path	
□ SiteA2B	g1-cluster	g1-evs3	prot...teA : /		172.....55	SiteACopy : /	details
□ test2	g1-cluster	g1-evs3	PHDS1 : /		172.....52	fs1 : /	details

Check All | Clear All

Actions: add remove

Shortcuts: File Replication Rules NDMP Configuration

Schedules

ID	Policies	Next Run	Interval	Last Status	
□ 11	SiteA2B	2014-06-11 00:00:00 (UTC-0700)	DAILY	● Failed to Start	details
□ 12	test2	None	ONCE	● OK	details

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: • **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): • **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 a policy that is selected by filling it check box.
Schedules	
ID	ID assigned to the replication policy.
Policies	Name of the replication policy.

Field/Item	Description
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	• Green indicates that a successful replication job has completed. • Red indicates a failed replication job and lists the reason for failure. ⚠ **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 one or more schedules that are selected by filling the check box next to the schedules.
Abort Replication(s)	Aborts one or more running replication operations by filling the check box next to the policy you want to abort.
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.

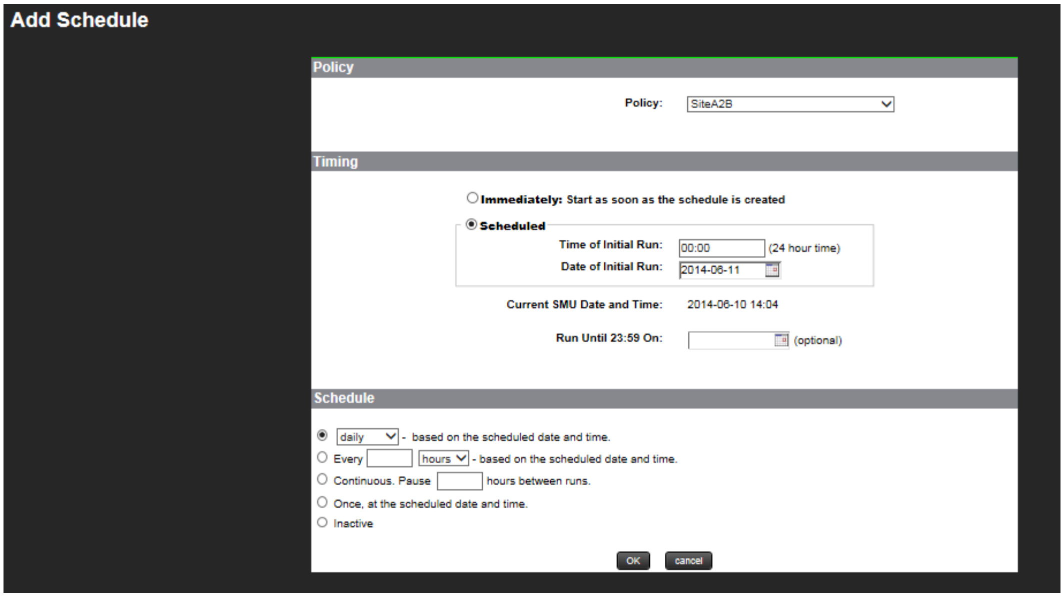

Add Schedule

Policy

Policy: SiteA2B

Timing

○ **Immediately:** Start as soon as the schedule is created

● **Scheduled**

Time of Initial Run: 00:00 (24 hour time)

Date of Initial Run: 2014-06-11

Current SMU Date and Time: 2014-06-10 14:04

Run Until 23:59 On: (optional)

Schedule

● daily - based on the scheduled date and time.

○ Every hours - based on the scheduled date and time.

○ Continuous. Pause hours between runs.

○ Once, at the scheduled date and time.

○ Inactive

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	• **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 `MM/DD/YYYY` (month/day/year), or select the date by clicking the calendar icon to display the calendar. 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.
Current SMU Date and Time	Provided for reference.
Date of Final Run	If you do not specify a Date of Final Run, the policy will run at the interval you specify in the Schedule section.

Field/Item	Description
Schedule	Select one of the options: • `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. • **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. ⚠ **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.
OK	Saves the configuration, and returns the **File Replication** page.
cancel	Closes the page without saving configuration changes.

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.

Field/Item	Description
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	Click **Run now** to run the replication policy immediately, regardless of schedule. ⚠ **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. Click **Abort** to stop an in-progress replication.
Recovery Actions	Click **restart** to restart the replication if the previous replication attempt failed. 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. ⚠ **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). 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 Data Systems 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."
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.

Field/Item	Description
Reschedule	This section allows changing the schedule of the next or final replication specified by this schedule: • 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	This section allows you to specify how often the replication policy is to be executed. Select one of the radio buttons: • 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: `Continuous. Pause quantity hours between runs`. The new replication job can start immediately or after a specified number of hours. • Selecting `Once, at the scheduled date and time` guarantees that the policy is scheduled to run only once. • Selecting `Inactive` causes the replication schedule to be placed on pause.
Actions	
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.

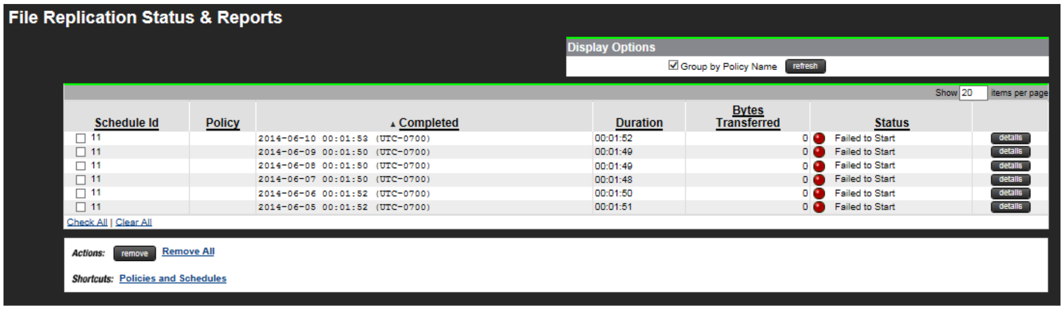

File Replication Status & Reports

Display Options

☑ Group by Policy Name

Show 20 items per page

	Schedule Id	Policy	▲ Completed	Duration	Bytes Transferred	Status	
□	11		2014-06-10 00:01:53 (UTC-0700)	00:01:52	0	🔴 Failed to Start	details
□	11		2014-06-09 00:01:50 (UTC-0700)	00:01:49	0	🔴 Failed to Start	details
□	11		2014-06-08 00:01:50 (UTC-0700)	00:01:49	0	🔴 Failed to Start	details
□	11		2014-06-07 00:01:50 (UTC-0700)	00:01:48	0	🔴 Failed to Start	details
□	11		2014-06-06 00:01:52 (UTC-0700)	00:01:50	0	🔴 Failed to Start	details
□	11		2014-06-05 00:01:52 (UTC-0700)	00:01:51	0	🔴 Failed to Start	details

Check All | Clear All

Actions: Remove All

Shortcuts: Policies and Schedules

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
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.

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

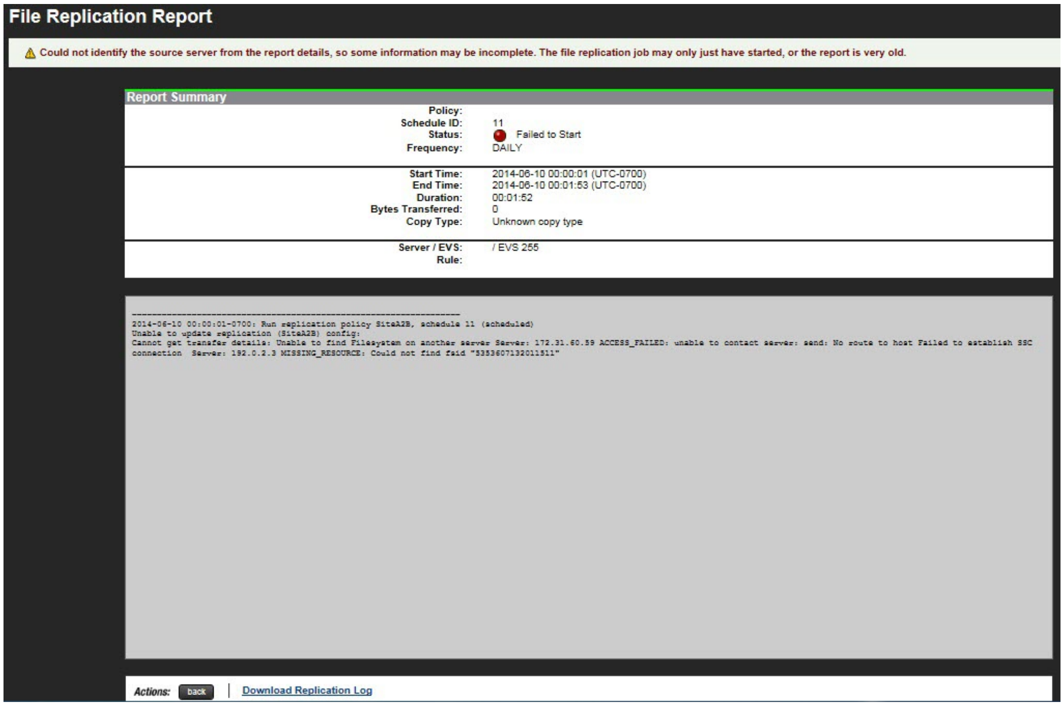

File Replication Report

⚠ 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:	
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) config:
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: 192.0.2.3 MISSING_RESOURCE: Could not find fsid "5353607132011511"
```

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: • 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. • 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.

Field/Item	Description
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 CIFS 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 CIFS 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 Web 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 CIFS), 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 CIFS 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.
- If you are using a high speed cross-continental or inter-continental link, then using multiple connections may help utilize more of the bandwidth of the connection. In cases where latency is high due to the connection in use, it might be useful to increase the number of connections used, even when the average file size is large.

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.

For instance, to get TCP connections capable of delivering 30 MB/sec from a file system containing mostly large files, 10 additional connections and 12 read-ahead processes might be suitable settings. Note that, in this case, it is not necessary to increase the number of read-ahead processes, because reading from the source file system will not be a limiting factor.

- 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.

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.
- SMU 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 Data Systems 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".

3

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. Hitachi NAS Platform supports object-based file system replication.

When using object-based file system replication and the transfer of primary access feature, you can replicate the file system, and the associated CNS links, CIFS shares, permissions and all other file-level metadata. Administrators can use Web 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 your 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. ⚠ **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.
apply	Applies the configuration changes, and closes the page.
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.

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 to be 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. 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 replicated on the object replication target. To change the source EVS/file system, click **change...**. If you viewed the object replication policy information by clicking the **details** button, the **change...** does not appear because the source file system cannot be altered.
EVS IP Address	Select the IP address for the source EVS from the list.
Target	
Server	Select a server from the list as the target of the object replication policy. After selecting a server, click **select a target...** to select an EVS and file system. 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	Select the EVS IP/host name in this field.
EVS IP Address	Enter the EVS IP address in this field.
File System	Enter the name of the target file system of the object replication policy. If the SMU 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. **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. **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.
Object Replication Port	The port on which the destination server is listening. The port on which the destination server listens is configured on the **Object Replication Configuration** page of Web Manager, or through the appropriate CLI command.

Field/Item	Description
	⚠ **Note:** To change the listening port for the target server, you must make that server the currently managed server of the SMU, then use the **Object Replication Configuration** page of Web Manager to change the listening port.
next	Advances to the next page to continue the policy configuration.
cancel	Closes the page without saving configuration changes.

3. Specify the details for the policy identification, replication source and the replication target.
4. Click **next** to continue with the policy configuration.
   The SMU checks if the source file system is configured to allow the access points (CIFS shares and NFS exports) on it to transfer with the object replication. If it is not, a GUI 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 depends on 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	Options for taking snapshots of the source file system: • **Snapshot source file system using automatic snapshot rule**, which allows the replication to use its default snapshot rule to take and manage snapshots. **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. • **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. **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.
Target File System	Options for taking snapshots of the target file system: • **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.
next	Advances to the confirmation page.

Field/Item	Description
back	Returns to the previous page.
cancel	Closes the page without saving configuration changes.

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 (CIFS 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.

An object replication policy requires the WFS-2 file system format for both the source and target file systems. If you selected a WFS-1 file system format as the source, this page is also displayed.

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.

⚠ **Note:** An object replication policy requires the WFS-2 file system format for both the source and target file systems. If you selected a WFS-1 file system format as the source, and you click this link, you are returned to the previous page, which now displays the error message `Failed to set the file system to transfer access points to an object replication target : The format on the file system is unsupported`.

- *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 SMU **Home** Page. As 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.

⚠ **Note:** An object replication policy requires the WFS-2 file system format for both the source and target file systems. If you selected a WFS-1 file system format as the source, this page is displayed. If you click this link, you may continue with the configuration; however, an error is generated and displayed when you run a test of the object replication policy.

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.
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. • **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.

Field/Item	Description
Current SMU 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. You can select any one of the following options: • **Every X minutes, hours, days, weeks, or months**: 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 completes. • **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. • **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. ⚠ **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.
OK	Saves configuration changes, and closes the page.
cancel	Closes the page without saving configuration changes.

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.
2. Fill the check box next to the schedule to modify 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	Fill this check box to change the schedule of the next replication specified by this schedule: • **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. The current SMU 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 `Run Until Date`.
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.

Field/Item	Description
Schedule Type	Specifies the interval, if any, between the repeated execution of the policy to which this schedule will apply. • **Every *X* minutes, hours, days, weeks, or months** 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 completes. • **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. • **Once, at the scheduled date and time** schedules the policy to run only once, at the date and time specified by the `Initial Run` 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 `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. ⚠ **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.
OK	Saves configuration changes, and closes the page.
cancel	Closes the page without saving configuration changes.

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: • 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: • 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: • 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. ⚠ **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.
details	Displays the details of the selected policy.
add	Advances to **Add Object Replication Policy** page.
delete	Fill the check box next to each policy you want to remove, and then click **delete**.
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.
Schedules	
ID	ID assigned to the replication policy.
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.

Field/Item	Description
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**.

Displaying object replication status and reports

This page displays detailed information about the status of object replication operations. This page displays 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
Policy	Displays the name of the replication policy for which the report was created.
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 all instances of one policy, selected in the Policy list.
Source	The source file system and its associated snapshot rule.
Target	The target file system and its associated snapshot rule. You can sort and reverse the order of the snapshot list by clicking snapshot. Start 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	Selects 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.

Field/Item	Description
download all reports	Downloads a `.csv` file containing data and status of all the object replication policies on the server. The downloaded file is displayed as a spreadsheet.

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

4

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:
  - CIFS Shares (if within replicated path)
  - NFS Exports (if within replicated path)
  - FTP Initial Directories/Users (if within replicated path)
  - Snapshot 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

□ Process of transferring primary access

How a transfer of primary access moves CNS links

Using the diagram below as a sample file system:

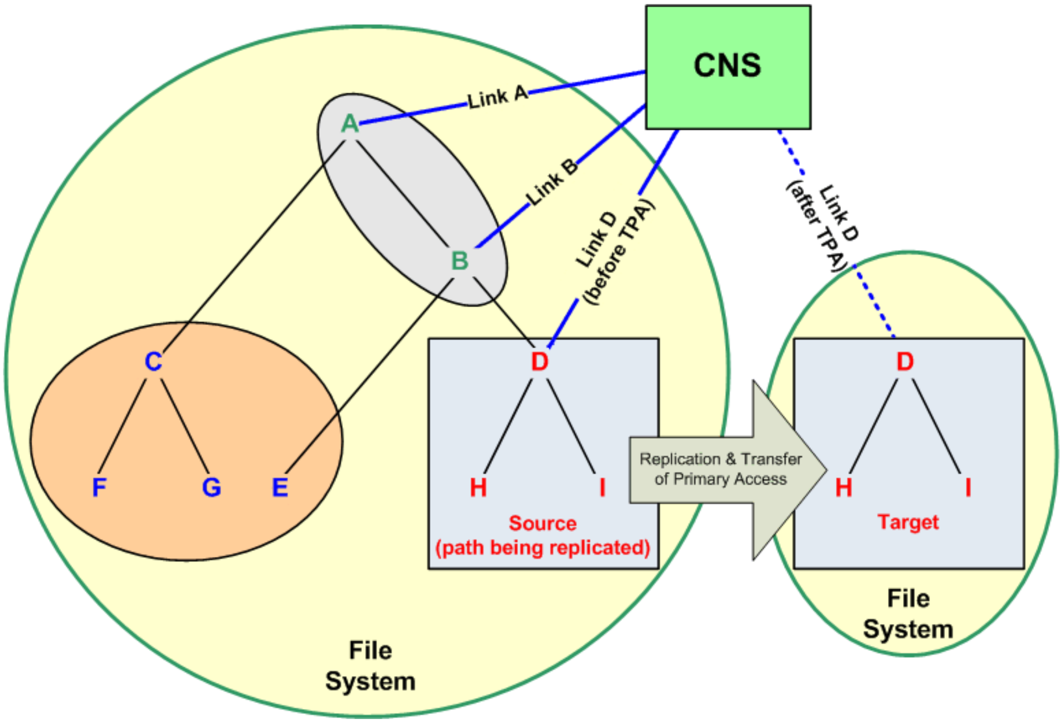

When replicating the file system path beginning at “D,” CNS links are transferred as follows:

- 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

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 SMU monitors the replication to determine when it is complete. When the replication is complete, the SMU 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**
CIFS 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.	
Snapshot Rules	Not moved if replicating only part of a file system. Moved only if file system will be a standalone file system when replicating data and access points to root (/), meaning that the target file system will be a standalone file system when the transfer of primary access is complete.		
CNS Links	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. After a transfer of primary access, network clients will not be able to access the file		If CNS links exist, the relocation is not allowed to proceed, and a message advises the administrator to remove the links before proceeding.

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
	system through the a CNS name space if any of the following are true: • 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. To access the file system in its new location, network clients must reconnect through CIFS 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. **Note:** If clients will not access the relocated file system using CNS links, they must access it using new IP addresses.		
iSCSI Targets	Not moved.		
Global Symlinks	Not moved.		

Note: For CIFS shares and NFS exports: if possible, a text file backup of the moved shares and export will be left on the SMU.

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 SMU 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 SMU is rebooted during a transfer of primary access, the source file system may not be returned to its original online state. If the SMU 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 original server data should be removed or made inaccessible.

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 SMU 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 CIFS 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.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.
www.hds.com

Regional Contact Information

Americas
+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa
+44 (0) 1753 618000
info.emea@hds.com

Asia Pacific
+852 3189 7900
hds.marketing.apac@hds.com

Hitachi Data Systems