



Hitachi Unified Storage File Module

Backup Administration Guide

Release 12.1

© 2011-2014 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd.

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.

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

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

Archivas, Dynamic Provisioning, Essential NAS Platform, HiCommand, Hi-Track, ShadowImage, Tagmaserve, Tagmasoft, Tagmasolve, Tagmastore, TrueCopy, Universal Star Network, and Universal Storage Platform are registered trademarks of Hitachi Data Systems Corporation.

AIX, AS/400, DB2, Domino, DS8000, Enterprise Storage Server, ESCON, FICON, FlashCopy, IBM, Lotus, OS/390, RS6000, S/390, System z9, System z10, Tivoli, VM/ESA, z/OS, z9, zSeries, z/VM, z/VSE are registered trademarks and DS6000, MVS, and z10 are trademarks of International Business Machines Corporation.

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

Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>). Some parts of ADC use open source code from Network Appliance, Inc. and Traakan, Inc.

Part of the software embedded in this product is gSOAP software. Portions created by gSOAP are copyright 2001-2009 Robert A. Van Engelen, Genivia Inc. All rights reserved. The software in this product was in part provided by Genivia Inc. and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the author be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

The product described in this guide may be protected by one or more U.S. patents, foreign patents, or pending applications.

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	7
Document Revision Level	8
Contacting Hitachi Data Systems.....	8
Related Documentation.....	8
1 About NDMP backup support.....	11
Standard NDMP configuration.....	12
Storage management applications.....	13
Enabling and disabling NDMP.....	14
About NDMP user name, password, and version.....	15
Specifying an NDMP user name and password.....	16
Enabling and disabling NDMP devices.....	17
Modifying NDMP device access configuration	18
About NDMP with snapshots.....	21
Configuring NDMP snapshot options.....	22
NDMP environment variables.....	25
Direct.....	25
EXCLUDE.....	25
EXTRACT.....	26
FILESYSTEM.....	26
FUTURE_FILES.....	26
HIST.....	26
LEVEL.....	26
NDMP_BLUEARC_AWAIT_IDLE.....	27
NDMP_BLUEARC_EMBEDDED_HARDLINKS.....	27
NDMP_BLUEARC_EXCLUDE_MIGRATED.....	28
NDMP_BLUEARC_EXTERNAL_LINKS.....	28
NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED.....	29
NDMP_BLUEARC_USE_CHANGE_LIST.....	29
NDMP_BLUEARC_USE_SNAPSHOT_RULE.....	30
A About Synchronous Image Backup.....	31



Preface

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.

- [Document Revision Level](#)
- [Contacting Hitachi Data Systems](#)
- [Related Documentation](#)

Document Revision Level

Revision	Date	Description
MK-92USF012-00	July 2012	First publication
MK-92USF012-01	September 2012	Revision 1, replaces and supersedes MK-92USF012-00.
MK-92USF012-02	June 2013	Revision 2, replaces and supersedes MK-92USF012-01.
MK-92USF012-03	August 2013	Revision 3, replaces and supersedes MK-92USF012-02.
MK-92USF012-04	April 2014	Revision 4, replaces and supersedes MK-92USF012-03.
MK-92USF012-05	September 2014	Revision 5, replaces and supersedes MK-92USF012-04.

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-92USF002)—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-92USF007)—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-92USF011)—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-92USF003)—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-92USF004)—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-92USF005)—In PDF format, this guide provides information about the Data Migrator feature, including how to set up migration policies and schedules.
- *Storage Subsystem Administration Guide* (MK-92USF006)—In PDF format, this guide provides information about managing the supported storage subsystems (RAID arrays) attached to the server/cluster. Includes information about tiered storage, storage pools, system drives (SDs), SD groups, and other storage device related configuration and management features and functions.
- *Snapshot Administration Guide* (MK-92USF008)—In PDF format, this guide provides information about configuring the server to take and manage snapshots.
- *Replication and Disaster Recovery Administration Guide* (MK-92USF009)—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-92USF010)—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-92USF012)—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 Unified Storage File Module 3080 3090 G2 Hardware Reference* (MK-92USF001)—Provides an overview of the second-generation server

hardware, describes how to resolve any problems, and replace potentially faulty parts.

- *Hitachi Unified Storage File Module 4000 Hardware Reference* (MK-92HNAS030) —Provides an overview of the Hitachi Unified Storage File Module 4000 server hardware, describes how to resolve any problems, and how to replace potentially faulty components.

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* —This document details the best practices for configuring and using HNAS storage pools, related features, and Hitachi Dynamic Provisioning (HDP).

About NDMP backup support

The storage server supports Network Data Management Protocol (NDMP), an open standard protocol for network-based backups, with two significant advantages:

- It enables a storage management application to control backup and recovery on another device without transfer of the backup data across the network.
- NDMP backups can preserve security settings in a mixed protocol environment, including virtual volume and quota information.

- [Standard NDMP configuration](#)
- [Storage management applications](#)
- [Enabling and disabling NDMP](#)
- [About NDMP user name, password, and version](#)
- [Specifying an NDMP user name and password](#)
- [Enabling and disabling NDMP devices](#)
- [Modifying NDMP device access configuration](#)
- [About NDMP with snapshots](#)
- [Configuring NDMP snapshot options](#)
- [NDMP environment variables](#)

Standard NDMP configuration





Figure 1-1 Standard NDMP configuration

In the diagram, the storage management application sends backup instructions to the server, which makes a backup copy of data onto tapes in the tape library. The data travels through the Fibre Channel (FC) network, not the Ethernet network. Details of the backup data are sent to the storage management application, which initiates recovery of the data if necessary.

NDMP transfers data between disks and tapes attached to the same server. Data can also be transferred between two separate NDMP servers over an Ethernet connection (in NDMP this is known as a three-way backup or recovery):



Some common applications of NDMP include:

- Backing up (or recovering) data on a server to (or from) a FC-attached NDMP tape library.
- Backing up (or recovering) data on a server without a tape library to (from) a second storage server that has a tape library attached.
- Using a utility, such as Accelerated Data Copy (ADC) or Data Replication to copy file systems between storage servers.

While the server supports backups done over network protocols such as NFS or CIFS, only NDMP will preserve security settings in a mixed protocol environment, including virtual volume and quota information.

When using NDMP, the server uses snapshots to backup data consistently and without being affected by ongoing file activity. Snapshots also facilitate incremental backups. However, if so desired, data can be backed up without using snapshots.

Storage management applications

The storage server acts as an NDMP server, operating with leading storage management applications. It supports NDMP Version 2, 3 and 4. The storage server implementation of NDMP can back up and restore:

- Both Windows and UNIX files from a single storage management application.
- The full attributes of each Windows and UNIX file (including Windows ACLs), saving and restoring whole volumes and preserving all file attributes. Usig NDMP Backups

The server supports recovery of single files or subdirectories, associated lists, or complete backup images. The Direct Access Recovery (DAR) mechanism can be used, provided the Storage Management Application supports it. DAR

allows NDMP to go directly to the correct place in the tape image to find the data, rather than reading the whole image. This can dramatically reduce recovery times.

Enabling and disabling NDMP


The **NDMP Configuration** page allows you to specify NDMP configuration information for a cluster or for the currently managed server, including NDMP user name, password, version, and port. NDMP processing status can be started or stopped at any time.

Procedure

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

2. Enter the storage server's NDMP configuration settings.

Field/Item	Description
User Name	The user name cannot more than 20 characters long and cannot contain the following characters: \ / < > " ' .
Password	By default, the password is "ndmp". The password cannot be more than 20 characters long.
Version	By default, the storage server uses NDMP version 4 for NDMP backup and recovery; if required, it can be configured for version 2 or 3 of the NDMP protocol.

Field/Item	Description
	 Note: Both incremental data replication and ADC require NDMP version 3 or 4. Set NDMP to version 2 only if required by your backup software.
Port	The NDMP port number. By default, port 10000 is used.

3. Start or stop the NDMP process.



Caution: Read this caution before following instructions to start and stop! Clicking stop terminates all NDMP processes immediately, leaving any tapes in use in an untidy state. It may also confuse the storage management application. Therefore, Hitachi Data Systems Support Center recommends terminating NDMP transfers using the storage management application before clicking stop.

- To stop NDMP processing, click **stop**. If any NDMP operations are in progress when you click stop, those operations will be aborted.
- To start NDMP processing, click **start**.

4. Enable or disable the NDMP process at Boot.

- To automatically enable NDMP processing at Boot, click **enable**.
- To automatically disable NDMP processing at Boot, click **disable**.

About NDMP user name, password, and version

A storage management application must successfully authenticate a configured NDMP user before starting a backup or recovery.



Note: Any user with NDMP user name and password knowledge can access an NDMP-enabled storage management application to access data on the system. Therefore, Hitachi Data Systems Support Center recommends taking measures to keep the information secure.

An administrator can specify two types of users:

- *NDMP Primary User*. For an NDMP primary user, an account user name and password provide full access to the files on the system, supporting most backup, recovery and replication activities.
- *Restricted NDMP Users*. The SSC command `ndmp-ruser` can create less trusted NDMP Restricted Users with access to a restricted set of files (and possibly devices). An administrator could assign these user names to various users to allow them to use the accelerated data copy (ADC) utility to copy data within limited areas of the file systems. The SSC command `ndmp-ruser-pwd` can also change the password for a selected restricted user.

For more information about `ndmp-ruser` and `ndmp-ruser-pwd`, see the *Command Line Reference*.

Specifying an NDMP user name and password

Procedure

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

NDMP Configuration

NDMP Settings

User name:

Password:

Version:

Port:


NDMP Server Status

Current Status: Started

Stop will halt the NDMP server, and terminate any NDMP operations in progress.

Enable NDMP Server At Boot: Enabled

2. Enter the required information NDMP settings:

Field/Item	Description
User Name	The user name cannot more than 20 characters long and cannot contain the following characters: \ / < > " ' .
Password	By default, the password is "ndmp". The password cannot be more than 20 characters long.
Version	By default, the storage server uses NDMP version 4 for NDMP backup and recovery; if required, it can be configured for version 2 or 3 of the NDMP protocol.  Note: Both incremental data replication and ADC require NDMP version 3 or 4. Set NDMP to version 2 only if required by your backup software.
Port	The NDMP port number. By default, port 10000 is used.

3. Click **Apply** to save your changes.

Enabling and disabling NDMP devices

NDMP backup devices, such as tape libraries and auto-changers, require special configuration. The server monitors its Fibre Channel (FC) links periodically and automatically detects the presence of backup devices. Because the server may be connected into a Storage Area Network (SAN) shared with other servers, it does not automatically make use of backup devices it detects on its FC links.


Procedure

1. Navigate to **Home > Data Protection > NDMP Device List** to display the **NDMP Device List** page.

EVS:Device Name	WWN Node (LUN)	Manufacturer (Model)	Serial Number	Allow Access	Status
<any>-Drive1	20:01:00:0e:11:14:74:7a (0)	IBM (ULT3580-HH5)	1068007372	Allowed	OK
<any>-Drive2	20:04:00:0e:11:14:74:7a (0)	IBM (ULT3580-HH5)	1068012502	Allowed	OK
<any>-Drive3	20:07:00:0e:11:14:74:7a (0)	IBM (ULT3580-HH5)	1068020340	Allowed	OK
<any>-Drive4	20:0a:00:0e:11:14:74:7a (0)	IBM (ULT3580-HH5)	1068008669	Allowed	OK
<any>-Robot	20:01:00:0e:11:14:74:7a (1)	IBM (3573-TL)	00X4U78P4127_LL0	Allowed	OK
<none>-Unknown	20:07:00:17:a4:fd:c0:af (0)	HP (Ultrium 3-SCSI)	HU10635LLG	Deny	OK

Actions: [allow access](#) [deny access](#) [forget](#) | [Refresh Status](#)

Shortcuts: [NDMP Configuration](#)

Item/Field	Description
EVS:Device Name	Displays the EVS or EVSs allowed to use the device, and the ID of the device. This ID is generated by the system and cannot be changed. To configure your storage management application to work with an NDMP device, in the storage management application you must specify the device name of each autochanger/tape drive you want the application to use.
WWN Node (LUN)	Displays the WWN (World Wide Name) and LUN ID of the Fibre Channel node.
Manufacturer (Model)	Displays the manufacturer and model of the device, if detected.
Serial Number	Serial number of the device.
Allow Access	Displays if access is allowed to the device. If access is not allowed, then NDMP will not attempt to use the corresponding device.  Note: An NDMP device must be assigned to an EVS before access can be allowed to the device.

Item/Field	Description
	<p>If access is not allowed to a device, fill the checkbox next to the device, and click allow access.</p> <p>To deny access to a device, fill the checkbox next to the device, and click deny access. A request to deny access will be rejected if an NDMP client has opened the device. The backup application configuration should be changed to avoid use of the device before denying access.</p>
Status	Current status of the selected device.

2. To enable/disable access to devices:

- Click **deny access** to disable access to a device, which prevents NDMP from attempting to use the device



Note: While an NDMP server has the device open, a deny access request will be rejected. Therefore, the storage management application configuration should be changed to avoid use of the device before the current configuration process.

- Click **allow access** to enable access to a device, which allows NDMP to use the device.



Note: Before using an NDMP device, you must first allow access to it, then it must be assigned to an EVS. NDMP Devices are assigned to an EVS using the **NDMP Device Access Details** page described in

- Click **forget** to remove the selected device from the list (only available for devices that have been disconnected from the FC).
- Click **Refresh Status** to discover any changes in the Fibre Channel connection; that is, to find any newly attached devices and discover whether any previously discovered devices that are no longer accessible. If new devices are plugged into the Fibre Channel, use Refresh to identify them.

Modifying NDMP device access configuration

NDMP backup devices, such as tape libraries and auto-changers, require special configuration. The server monitors its Fibre Channel (FC) links periodically and automatically detects the presence of backup devices. Because the server may be connected into a Storage Area Network (SAN) shared with other servers, it does not automatically make use of backup devices it detects on its FC links.

Procedure

1. Navigate to **Home > Data Protection > NDMP Device List** to display the **NDMP Device List** page.

Data Protection [Home](#) > [Data Protection](#) > [NDMP Device List](#) > NDMP Device Access Details

NDMP Device Access Details

ID: 12

Allow Access: Allowed

EVS: Any EVS reassign

Hardware Details



Device Type: Tape Robot
Manufacturer (Model): IBM (3573-TL)
Version: C.30


Device Identification

NDMP Device Name: Robot
Location: N/A
Serial Number: 00X4U78P4127_LL0
Fibre Channel Address: 20:01:00:0e:11:14:74:7a
LUN: 1

Actions: allow access deny access forget

2. The following table describes the fields in this page.

Item/Field	Description
ID	Displays the server-assigned device identifier.
Allow	<p>Indicates if device access is allowed (Allow) or denied (Deny).</p> <hr/> <p> Note: An NDMP device must be assigned to an EVS before access can be allowed to the device. If access is not allowed to a device, click allow access to enable access.</p> <p>To deny access to a device, click deny access. A request to deny access will be rejected if an NDMP client has opened the device. The backup application configuration should be changed to avoid use of the device before denying access.</p>
EVS	<p>Indicates the specific EVS to which the device is assigned, or indicates that the device is assigned to all EVSs.</p> <p>To change the device assignment, select the EVS to which you want to assign the device, or select All EVS to assign the device to all EVSs hosted by the server/cluster, and click reassign.</p> <p>Tape devices can be shared among EVSs under the following conditions:</p> <ul style="list-style-type: none"> The EVSs must be within the same cluster. The tape device is not shared with another Hitachi Data Systems server. The tape device is not shared with another storage device. <hr/> <p> Notes:</p> <ul style="list-style-type: none"> If the device is to be shared between this server and another non-clustered server, additional sharing logic is required. Some backup applications automatically allow such sharing without any extra configuration. For other backup applications it is necessary to use the SCSI Reserve/Release protocol,

Item/Field	Description
	<p>which can be enabled using the <code>ndmp-option reserve_devices</code> CLI command.</p> <ul style="list-style-type: none"> • When one EVS is currently using a tape device, any attempt to use it through a different EVS will prompt a notification that the device is currently in use (that is, the operation will not be queued). • When a tape device is currently assigned to a specific EVS (but not to All EVS), any attempt to access it through a different EVS will prompt notification that the device has not been found.
Hardware Details	<p>This section displays hardware-related details about the device, including:</p> <ul style="list-style-type: none"> • Device Type, which can be either tape drive or autochanger. • Manufacturer (Model), which are the device manufacturer and model detected when the device is discovered. • Version, which indicates the version of the firmware currently on the device, if it was detected when the device was discovered.
Device Identification	<p>This section displays identification information about the device, including:</p> <ul style="list-style-type: none"> • NDMP Device Name, which displays the name by which the device can be addressed by the server. To configure your storage management application to work with an NDMP device, in the storage management application you must specify the device name of each autochanger/tape drive you want the application to use. • Location, which displays the name of the autochanger that holds the drive and the position of the drive in the autochanger. For example, the location of the first drive in autochanger <code>/dev/mc_d010</code> is <code>/dev/mc_d010 : 1</code>. • Serial Number, which indicates the device's serial number, if it was detected when the device was discovered. • Fibre Channel Address, which indicates the device's Fibre Channel node name. • LUN, which indicates the LUN identifier for the device. When the Web Manager cannot determine the location of a tape drive, it displays <i>*unknown*</i>. When this occurs, check for the following conditions and follow the troubleshooting instructions: <ul style="list-style-type: none"> • The tape library is offline. • The autochanger does not support the server's mechanism for querying the tape drive location, or the autochanger has not been set up to accept this query. Where this is the case, compare the serial numbers of the tape drives with displays available in the tape library to verify the drive locations. • The autochanger and a tape drive within it are attached to different servers. In this case, use the tape drive serial numbers to match the device name shown by one server with the location shown on the other. <hr/> <p> Note: Devices will not be available or visible if access to them has not been enabled.</p>

3. The following **Actions** are available:

- Click **deny access** to disable access to a device, which prevents NDMP from attempting to use the device



Note: While an NDMP server has the device open, a deny access request will be rejected. Therefore, the storage management application configuration should be changed to avoid use of the device before the current configuration process.

-
- Click **allow access** to enable access to a device, which allows NDMP to use the device.



Note: Before using an NDMP device, you must first allow access to it, then it must be assigned to an EVS. NDMP Devices are assigned to an EVS using the **NDMP Device Access Details** page described in

-
- Click **forget** to remove the selected device from the list (only available for devices that have been disconnected from the FC).
 - Click **Refresh Status** to discover any changes in the Fibre Channel connection; that is, to find any newly attached devices and discover whether any previously discovered devices that are no longer accessible. If new devices are plugged into the Fibre Channel, use Refresh to identify them.

About NDMP with snapshots

The server uses snapshots to backup data consistently and without being affected by on-going file activity.

The following options should be considered when planning a backup strategy:

- Back up automatically created snapshots.

When backing up a file system that is being actively updated, a snapshot of the file system is much more likely to produce a fully consistent image than backing up the live file system. As a result, NDMP is configured by default to automatically create a snapshot for backup.

- Back up pre-created snapshots

A backup can be taken from a specific snapshot that has been *created by a rule or created spontaneously by user request*:

- To back up the latest snapshot created under a snapshot rule, use the environmental variable `NDMP_BLUEARC_USE_SNAPSHOT_RULE`.
- To back up the latest snapshot created spontaneously by user request, request a specific snapshot by explicitly including the snapshot name in the path to back up. Where the path is based on a CIFS share name, indicate the snapshot using `/~snapshot/snapshot_name`; for paths based on an NFS export name, indicate the snapshot using `/.snapshot/snapshot_name`. CIFS shares and NFS exports may also include a snapshot name.

- Backing up databases and iSCSI Logical Units

The internal structures of Databases and iSCSI LUs are tightly coupled with the state of the client software (database manager/iSCSI Initiator) that is controlling the files. For example, backing up such files during a client operation may produce inconsistencies in the backup that would prevent recovery.

Therefore, any backup of databases and iSCSI LUs must ensure that files are in a consistent state at the time of back up. Snapshots can be used to achieve this. Snapshot rules provide the most convenient mechanism, as this avoids having to explicitly specify the name of the snapshot used.



Note: When configuring snapshot rules, ensure that snapshots have a sufficiently long shelf life, and before initiating a backup, verify that the snapshot is not scheduled to be replaced during the anticipated time of the backup, as such replacement would cause the backup to fail.

For more information on backing up and restoring iSCSI LUs, refer to the *File Services Administration Guide*.

Configuring NDMP snapshot options

To configure NDMP snapshot options:

Procedure

1. Navigate to **Home > Data Protection > NDMP History and Snapshots**.

NDMP History & Snapshots

NDMP Backup History

Clear NDMP backup records on all EVSes.

Note:

- These settings apply to tape backups and ADC, but not file replication.
- Changes will result in a full backup, not an incremental one.

Snapshot Options

Automated Snapshot Use

Do not automatically create snapshots, but backup from the live file system.

Automatically create snapshots. (This option does not affect file replication snapshot usage.)

Automated Snapshot Deletion

Delete snapshot after use


Delete snapshot after next backup




Delete snapshot when obsolete

Automated Snapshot Retention

Set Retention Maximum To: Days

Note: This setting will affect file replication

Field/Item	Description
clear All	<p>When necessary, you can clear the records of completed tape-based backups and either scripted or command line-based incremental accelerated data copies (ADCs). Clearing the history does not affect replication operations (replication history is managed separately) or data migration operations (migration is not an incremental operation). When performing incremental backups, the server uses the records of old backups to determine the date and time after which it must back up modified files. If you have lost a backup for any reason, you can clear the records, which forces the next backup to be a full backup instead of an incremental backup.</p> <hr/> <p> Note: To force a full backup for replication, delete the snapshot that was automatically created at the start of the last replication</p>
Automated Snapshot Use	<p>Do not automatically create snapshots, backup from the live file system - This option causes a backup to be performed from the live file system (no snapshot is taken).</p>

Field/Item	Description
	<p>Automatically Create Snapshots (recommended) - This option causes a snapshot to be taken, then the backup is performed from that snapshot. Note that this option does not affect replication snapshot usage.</p> <hr/> <p> Note: If a backup path explicitly contains a snapshot reference, the system does not take a new snapshot, regardless of this setting.</p> <hr/>
<p>Automated Snapshot Deletion</p>	<p>By default, NDMP keeps the snapshot to make incremental backups more accurate. In the Automated Snapshot Deletion section, select whether to delete the snapshot:</p> <ul style="list-style-type: none"> • Delete snapshot after use deletes an automatic snapshot after completion of the backup for which it was taken. To prevent accumulation of unneeded snapshots, select this option for full backups or if the file system is changing very rapidly. • Delete snapshot after next backup deletes an automatic snapshot after it has been used as the basis of a new incremental backup. With an exception for full backups, this option supports "incremental" backup schedules based on the immediately preceding backup. • Delete snapshot when obsolete deletes an automatic snapshot upon next backup at the same level. For example, a snapshot taken for a full backup will only be deleted when the next full backup is completed. This option supports "differential" backup schedules based on a common base backup. <hr/> <p> Note: Snapshots initiated by the Microsoft Volume Shadow Copy Service (VSS) may not be deleted by rule. These snapshots should be managed through the application that requested the snapshot. You can, however, delete these snapshots through the Snapshots page.</p> <hr/>
<p>Automated Snapshot Retention</p>	<p>Determines number of days (1 to 40) to keep snapshots before auto-deletion.</p> <p>Usually, the system deletes automatically created snapshots according to the rule selected in the previous step; however, after a sequence of backups using automatically created snapshots is stopped, snapshots may be left over. The maximum retention time provides a way of tidying up in these circumstances.</p> <hr/> <p> Note: This setting applies to snapshots automatically taken by replications. Set the retention time to be long enough to make sure that a snapshot from a replication copy is not deleted until after the next successful copy is complete. This means that the maximum time set here must be longer than the time taken to run two replication copies, including the interval between the</p>

Field/Item	Description
	replication copies and the time required to make the copies.

2. Click **Apply** to save your changes.

NDMP environment variables

You can use NDMP environment variables to modify backup actions. The storage management application generates most of these variables and supports configuration of additional variables. They are invoked from the Replication Rules: **Add Rules** page.

Direct

Possible value	Notes
y or n	<p>Used on recovery to request Direct Access Recovery (DAR).</p> <p>May be used to recover a subset of a full backup. If the storage management application supports DAR, the recovery will position the tape to the start of the required data rather than reading a complete backup image to find the data. This saves time in recovery of single files and similar operations.</p> <p>The Storage Management Application may control the setting of this variable, based on either the setting of a user interface option, or on an assessment of the likely efficiency of using DAR; however, in some cases, it may be necessary to explicitly set <code>DIRECT=y</code>.</p>

EXCLUDE

Possible value	Notes
Comma-separated list of files or directories	<p>Specifies files or directories to exclude from a backup. By default, none are excluded.</p> <p>When specifying a file or directory, type either:</p> <ul style="list-style-type: none"> • A full path name, relative to the top-level directory specified in the backup path. The path name must start with a forward slash (/). An asterisk (*) can be typed at the end as a wildcard character. • A terminal file or directory, which is simply the last element in the path. The name must not contain any forward slash (/) characters, but it may start or end with the wildcard character *. <p>For example:</p> <pre>ENVIRONMENT EXCLUDE "/dir1/tmp*,core*,*.o"</pre> <p>This command excludes all files and directories that:</p> <ul style="list-style-type: none"> • Start with the letters tmp in the directory /dir1

Possible value	Notes
	<ul style="list-style-type: none"> • Are called core • End with the characters .o <p>The command is case-sensitive if backing up an NFS export but not if backing up a CIFS share.</p>

EXTRACT

Possible value	Notes
y or n	The default value y causes a recovery operation to extract files from a file list rather than recovering the whole backup.

FILESYSTEM

Possible value	Notes
Name of directory to back up	The Storage Management Application sets the FILESYSTEM variable to the name of the path to be backed up.

FUTURE_FILES

Possible value	Notes
y or n	Enables back up of files created after the start of the current backup. With NDMP version 2, the inode number that identifies a file can be reused during a backup, thereby causing the backup to fail. By default, therefore, only files created before the start of the backup are backed up. To override this behavior, set FUTURE_FILES=y.

HIST

Possible value	Notes
y or n	The default value y causes file history information to be sent to the storage management application. This enables the display and recovery of the contents of a backup.


LEVEL

Possible value	Notes
0 – 9, or i	The default value is 0 (full backup). If the value is set to 4, an incremental backup is taken based on the most recent previous backup of the same FILESYSTEM with level 0, 1, 2, or 3. If the value is set to i, an incremental backup is taken based on the most recent previous backup of the same FILESYSTEM of any level.

NDMP_BLUEARC_AWAIT_IDLE

Possible value	Notes
y (default) or n	<p>By default, the data management engine imposes an interlock to prevent NDMP backups and accelerated data copies (ADCs) from a replication destination while a replication copy is actively writing data. This is intended for installations that replicate to a particular volume, then back up from that same volume. However, as the lock is held at a volume level, it may be desirable in the case of directory-level replication to override this action.</p> <p>To make use of this replication interlock, specify <code>y</code> on both the replication that is intended to do the waiting and the replication that is waited upon.</p>

NDMP_BLUEARC_EMBEDDED_HARDLINKS

Possible value	Notes
y or n	<p>Used to enable or disable inline hard linked file support. Set the value to <code>y</code> to enable, or <code>n</code> to disable. For backups, inline hard linked file support is set to <code>n</code> (disabled) by default, but for multi-stream operations, such as replications and accelerated data copies (ADCs) between servers, the default is overridden and inline hard linked file support is enabled. By default, replication and ADC operations use multiple data streams, so for those operations, inline hard linked file support is used by default.</p> <p>When enabled, inline hard linked file support causes NDMP to back up hard linked files with both file data and file metadata inline (in a single data stream), which reduces the amount of memory the server needs to manage the data.</p> <p>Set to <code>n</code> to disable inline hard linked file support, which causes file metadata and file data to be sent in two data streams. Disabling inline hard link file support maintains backup compatibility with older systems or releases.</p> <p>Inline hard linked file support may not be enabled using the <code>ndmp-option</code> command. Rather, the command used to invoke NDMP must request inline hard linked file support.</p> <hr/> <p> Note: Existing programs that can read NDMP data streams for releases prior to release 6.1 will not be able to read backups or recover from backups created using inline hard linked file support.</p> <hr/> <p>If a replication fails part way through, it will not be possible to restart replication if the server is downgraded to a release prior to release 6.1.</p>

Using this option with replications and ADCs

When multi-streamed replication or ADC operations are started, this option is enabled. Starting in release 6.1, replication and ADC operations are multi-streamed by default, meaning that this option will be enabled by default for those operations.

Using backups

When backing up a file system:

- When the embedded hard link option is enabled, the data for each hard linked file is included in the data stream wherever a path to that file is included.

When enabled, the embedded hard link option increases the amount of data backed up, because multiple copies of the hard linked file data are included. However, it reduces the complexity of managing the backup.

Also, note that enabling the embedded hard link option reduces the memory requirements needed to keep track of all the hard links.

- When the embedded hard link option is disabled, paths to hard linked files are included without any data in the main part of the backup and a single copy of the hard link file data is included at the end of the backup.

This reduces the amount of data backed up, because only a single copy of the hard linked file data is included.

Recommendations for usage with backups


- If the backup contains many (more than a few hundred thousand) hard linked files, you should enable this option, because it reduces the memory overhead. Note that, where the backup includes many millions of hard linked files, enabling this option may allow the backup to complete where it would not complete if the option is disabled.
- If the backup contains a relatively small number of hard linked files each containing a large amount of data, you should disable the option.
- If there is a chance that the backup may need to be restored on an older version of software, you should disable this option.

NDMP_BLUEARC_EXCLUDE_MIGRATED

Possible value	Notes
y or n	<p>Indicates whether backups or replications will include files whose data has been migrated to secondary storage.</p> <p>If set to y, the backup or copy will not include files whose data has been migrated to another volume using the Data Migrator facility.</p> <p>The default n specifies that migrated files and their data will be backed up as normal files. The backup/copy retains the information that these files had originally been migrated.</p>

NDMP_BLUEARC_EXTERNAL_LINKS

Possible value	Notes
ignore, recreate_link, or remigrate	<p>Controls what happens when a replication operation encounters a cross volume link (a link to a file that has been migrated to an external server).</p> <ul style="list-style-type: none">• If set to <code>ignore</code>, the replication operation copies only the files on the primary (migrated files are not copied). Use this setting when files

Possible value	Notes
	<p>have been migrated because they are less useful, so they are not replicated in order to save time.</p> <ul style="list-style-type: none"> • If set to <code>recreate_link</code>, 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. • If set to <code>remigrate</code>, the replication operation copies the file contents but marks the file as having been externally migrated. The destination re-migrates to secondary storage if there is an existing data migration path. Use this setting when the replication is between a main site and a disaster recovery site, where the disaster recovery site includes a similar data migration configuration. This is the default. <hr/> <p> Note: If this option is explicitly set, it overrides the setting of the <code>backup_ignore_external_links</code> option of the <code>ndmp-option</code> CLI command.</p>


NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED

Possible value	Notes
y or n	<p>Indicates if backups or ADC copies will exclude files whose data has not been migrated to secondary storage.</p> <p>If set to <code>y</code>, the backup or copy includes only those files whose data has been migrated to another volume using the Data Migrator facility.</p> <p>The default <code>n</code> indicates that files whose data has not been migrated be backed up like normal files.</p>

NDMP_BLUEARC_USE_CHANGE_LIST

Possible value	Notes
y or n	<p>Indicates whether incremental backups or replications will use a <i>changed object list</i> to direct the search for changed files; otherwise, it will have to search the entire directory tree looking for changed files. When using the changed object list, the search only passes through those directories that contain changed files.</p> <p>Where a relatively small proportion of the file system includes directories containing changed files, the use of <i>changed object lists</i> may significantly reduce incremental backup and replication time; however, processing of the <i>changed object list</i> itself may take considerable time. Therefore, where file changes exist in many directories, its use is not recommended.</p> <p>The default setting for this option can be set using the CLI <code>ndmp-option change_list_incr</code> command.</p>

NDMP_BLUEARC_USE_SNAPSHOT_RULE

Possible value	Notes
Snapshot rule name	<p>Causes NDMP to back up the latest snapshot created under a specified snapshot rule. This can be used to backup a snapshot taken at a specific time; for example, for databases.</p> <p>If set, NDMP does not create or delete snapshots.</p> <hr/> <p> Note: Following a successful backup, the snapshot should not be deleted until after the operation has completed. In addition, the snapshot should be kept around long enough to support incremental backups.</p> <hr/>

About Synchronous Image Backup

Synchronous Image Backup enables fast, full file system, high-speed tape backup and restoration. This tool is designed specifically for situations where there is a need to backup multiple petabytes of data to tape in a short time period. However, if there is a need to restore files and directories on a more granular level, NDMP file backup is the correct tool.

Image backup is a key component for Synchronous Image Backup. It enhances the performance and scalability of HDS's NDMP backup functionality. The HDS file system is an object based file system, in which core file system structures and user data are stored as objects, rather than files or blocks. Synchronous Image Backup takes advantage of this structure and takes a snapshot of the complete file system image. The system then creates the NDMP backup stream by moving sequentially through the file system copying objects as they are found regardless of the file or directory they belong to. This negates the need to assemble all file objects associated with a file before transfer, making the creation of the NDMP data stream more efficient.

When backing up to tape, one of the primary performance considerations is the ability to keep the tape streaming at speed (no stopping, rewinding or repositioning). By increasing the efficiency of the NDMP backup stream creation, the ability to keep the backup tape streaming at speed is maintained. This extends the ability to keep multiple tapes streaming at speed, improving overall NDMP backup throughput.

Backup Vendor—Only Quest (Bakbone) NetVault is supported with Synchronous Image Backup. If the system is using any other backup software vendor, then the existing NDMP file backup must be used.



Note: A server will always be able to restore a Synchronous Image Backup backup if it is running a later software version than was used when the backup was generated.

A server will not necessarily be able to restore a Synchronous Image Backup backup if it is running an earlier software version than was used when the backup was generated. If the restore cannot proceed for this reason, the

application managing the restore (in general a third-party backup application) should log the reason why the restore has failed.

A NAS Platform Series 3000 server cannot restore a backup generated using Synchronous Image Backup from a HNAS 3080 or HNAS 3090 server running a NAS File OS build from 10.0 onwards. If this is required then Synchronous Image Backup must not be used to generate the backup.

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



MK-92USF012-05