

Backup Administration Guide

Hitachi Virtual Storage Platform Gx00 and Fx00 with NAS
Modules

VSP N series

Hitachi NAS Platform

Release 14.2

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

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

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

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

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

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

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

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

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

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

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

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

Copyright and license information for third-party and open source software used in Hitachi products can be found at <https://www.hitachivantara.com/en-us/company/legal.html>.

Contents

Preface	5
Related Documentation.....	5
Accessing product documentation.....	8
Getting help.....	8
Comments.....	8
About NDMP backup support.....	9
LAN-free NDMP configuration.....	9
LAN-based NDMP configuration.....	10
Storage management applications.....	11
Enabling and disabling NDMP.....	11
About NDMP user name, password, and version.....	14
Specifying an NDMP user name and password.....	14
Enabling and disabling NDMP devices.....	15
Modifying NDMP device access configuration	17
About NDMP with snapshots.....	20
Configuring NDMP snapshot options.....	21
Configuring backup and restore of migrated data.....	24
Locally migrated data.....	25
Externally migrated data.....	26
NDMP environment variables.....	29
Direct.....	29
EXCLUDE.....	30
EXTRACT.....	31
FILESYSTEM.....	31
FUTURE_FILES.....	31
HIST.....	31
LEVEL.....	32
NDMP_BLUEARC_AWAIT_IDLE.....	32
NDMP_BLUEARC_EMBEDDED_HARDLINKS.....	32
NDMP_BLUEARC_EXCLUDE_MIGRATED.....	34
NDMP_BLUEARC_EXTERNAL_LINKS.....	35
NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED.....	35
NDMP_BLUEARC_TAKE_SNAPSHOT.....	36
NDMP_BLUEARC_USE_CHANGE_LIST.....	36

NDMP_BLUEARC_USE_SNAPSHOT_RULE..... 36

Preface

This guide provides information about configuring the server to work with NDMP, and making and managing NDMP backups. This guide also includes information about Hitachi NAS Synchronous Image Backup. Note that some features apply only to individual platforms and may not be applicable to your configuration.

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

Related Documentation

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

Command Line References

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

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

Administration Guides

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



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

Hardware References

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

Best Practices

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

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

Accessing product documentation

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

Getting help

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

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

Comments

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

Thank you!

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.

LAN-free NDMP configuration

Important: The NAS module architecture and the Hitachi NAS Platform Series 5000 do not support the direct attachment of tape drives or autochangers. LAN-free backup is available only on the NAS Platform Series 3000 and NAS Platform Series 4000.


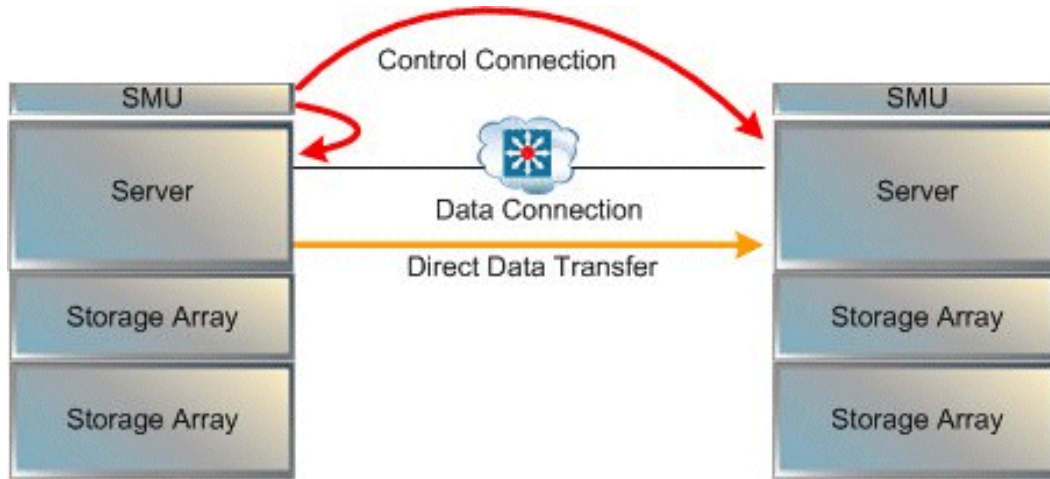


Figure 1 LAN-free NDMP configuration

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:

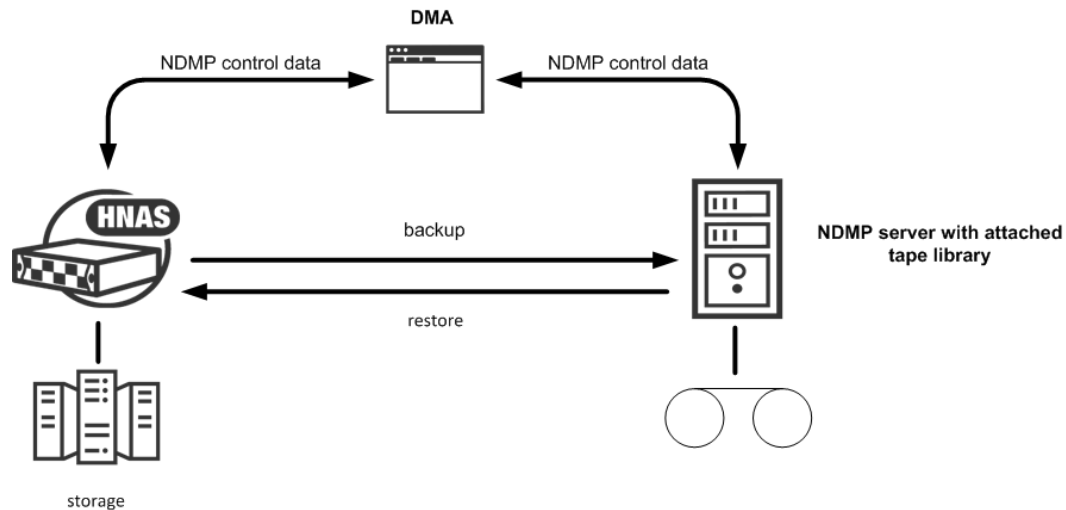
- Backing up (or recovering) data on a server to (or from) a FC-attached NDMP tape library.
- Backing up (or recovering) data on a server without a tape library to (from) a second storage server that has a tape library attached.
- Using a utility, such as Accelerated Data Copy (ADC) or Data Replication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

While the server supports backups done over network protocols such as NFS or SMB, 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, data can be backed up without using snapshots.

LAN-based NDMP configuration

This type of configuration is useful when the tape library is **not** directly attached to the NAS server. For example, you can use this configuration for the NAS Platform Series 5000 or NAS module as these do not support LAN-free backup.



In the diagram, the storage management application (DMA) sends backup instructions to the NAS server which then sends a backup copy of the data from its storage to the server which is attached to the tape library. The data is then stored onto tapes in the tape library.

Storage management applications

The server acts as an NDMP server, operating with leading storage management applications. It supports NDMP Version 2, 3 and 4. The 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.

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.

Data Protection [Home](#) > [Data Protection](#) > NDMP Configuration

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 storage server's NDMP configuration settings.

Field/Item	Description
NDMP Settings	
User Name	The user name cannot be 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. <div style="border: 1px solid black; background-color: #e0f2f1; padding: 5px;">  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. </div>
Port	The NDMP port number. By default, port 10000 is used.
apply	Apply the changes made to the NDMP Settings.
NDMP Server Status	
Current Status:	The options are: <ul style="list-style-type: none"> ▪ start - Enable NDMP server. ▪ stop - Disable NDMP server and terminate any NDMP operations in progress.
Enable NDMP Server At Boot:	The options are: <ul style="list-style-type: none"> ▪ enable - Enable NDMP server at boot time. ▪ disable - Disable NDMP server at boot time.

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. To help avoid this, you can terminate 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 Customer Support recommends taking measures to keep the information secure.

An administrator can specify two types of users:

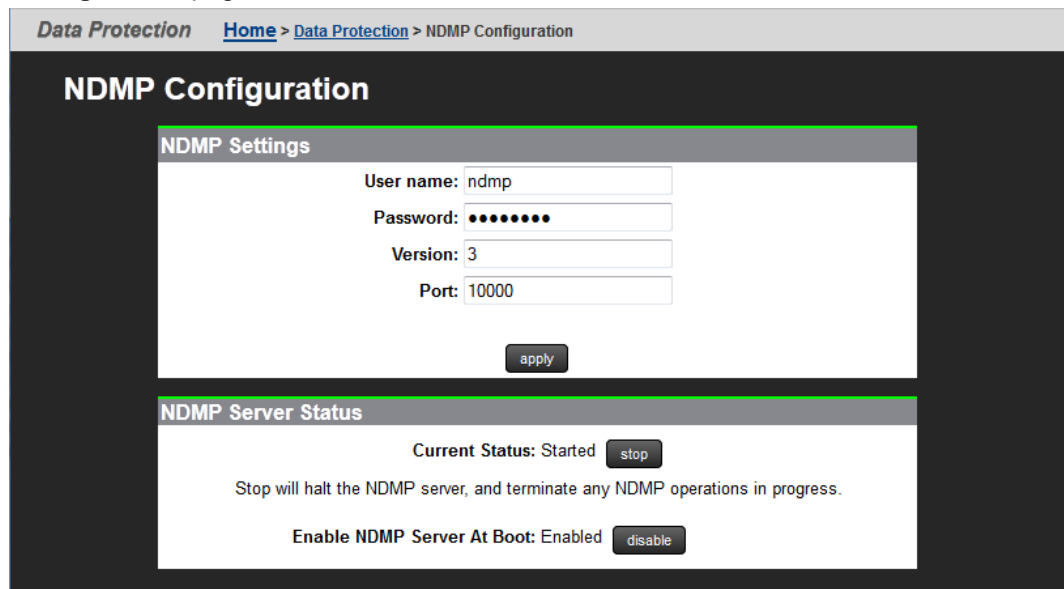
- *NDMP Primary User.* For an NDMP primary user, an account user name and password provide full access to the files on the system, supporting most backup, recovery and replication activities.
- *Restricted NDMP Users.* The SSC command `ndmp-ruser` can create less trusted NDMP Restricted Users with access 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
NDMP Settings	
User Name	The user name cannot be 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.
apply	Apply the changes made to the NDMP Settings.
NDMP Server Status	
Current Status:	The options are: <ul style="list-style-type: none"> ▪ start - Enable NDMP server. ▪ stop - Disable NDMP server and terminate any NDMP operations in progress.
Enable NDMP Server At Boot:	The options are: <ul style="list-style-type: none"> ▪ enable - Enable NDMP server at boot time. ▪ disable - Disable NDMP server at boot time.

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.

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


NDMP Device List

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

[Check All](#) | [Clear All](#)

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

Shortcuts: [NDMP Configuration](#)

Item/Field	Description
EVS:Device Name	<p>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.</p> <p>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.</p>
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	<p>Displays if access is allowed to the device. If access is not allowed, then NDMP will not attempt to use the corresponding device.</p> <div style="border: 1px solid #ccc; background-color: #e0f2f1; padding: 5px; margin: 5px 0;"> <p> Note: An NDMP device must be assigned to an EVS before access can be allowed to the device.</p> </div> <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** on the Device List page 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

- Navigate to **Home > Data Protection > NDMP Device List** to display the **NDMP Device List** page.
- Click on the **details icon** for the device for which you wish to modify access.



3. 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> <div style="background-color: #e0f2f1; padding: 10px; border: 1px solid #ccc;"> <p> Note: An NDMP device must be assigned to an EVS before access can be allowed to the device.</p> <p>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> </div>
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 server. ▪ The tape device is not shared with another storage device. <div style="background-color: #ffe0b2; padding: 10px; border: 1px solid #ccc;"> <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, which can be enabled using the <code>ndmp-option reserve_devices</code> CLI command. ▪ 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. </div>

Item/Field	Description
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 NAS 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. <div style="background-color: #e0f2f1; padding: 5px; margin-top: 10px;">  Note: Devices will not be available or visible if access to them has not been enabled. </div>

4. 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** on the Device List page 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**.

Data Protection [Home](#) > [Data Protection](#) > NDMP History & 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> <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, but backup from the live file system - This option causes a backup to be performed from the live file system (no snapshot is taken).</p> <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> <p> Note: If a backup path explicitly contains a snapshot reference, the system does not take a new snapshot, regardless of this setting.</p>
Automated Snapshot Deletion	<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.

Field/Item	Description
	<ul style="list-style-type: none"> ▪ 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. <div data-bbox="753 474 1393 743" style="background-color: #e0f2f1; padding: 10px; border: 1px solid #ccc;"> <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> </div>
Automated Snapshot Retention	<p>Determines number of days (1 to 80) 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> <div data-bbox="716 1066 1393 1434" style="background-color: #e0f2f1; padding: 10px; border: 1px solid #ccc;"> <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 replication copies and the time required to make the copies.</p> </div>

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

Configuring backup and restore of migrated data

There are several ways to configure backups of data that has been migrated by Hitachi NAS data migration features. For more information on the Data Migrator to Cloud cloud providers and features, see the *Data Migrator Administration Guide*. The configuration varies depending on the data migration mechanism used and the desired backup behavior.

Locally migrated data

Backup and restore of data sets that contain data migrated to other local Hitachi NAS file systems are affected by the following NDMP environment variables:

- NDMP_BLUEARC_EXCLUDE_MIGRATED
- NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED
- NDMP_BLUEARC_REMIGRATE

Environment variable	D*	S*	Backup	Restore
NDMP_BLUEARC_EXCLUDE_MIGRATED	n	y	Migrated files are not included in the backup data.	Migrated files cannot be restored.
		n	Migrated files are included in the backup data	Depends on NDMP_BLUEARC_REMIGRATE.
NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED	n	y	Only migrated files are included in the backup data.	Depends on NDMP_BLUEARC_REMIGRATE.
		n	Normal files and migrated files are included in the backup data.	Depends on NDMP_BLUEARC_REMIGRATE.
NDMP_BLUEARC_REMIGRATE	y	y	Migrated files are marked as such in the backup data.	The recovery target file system is checked for a configured migration path. If one exists, the file is immediately migrated. The file data is never written to the primary file system. If no migration path is configured, a normal file is created.
		n	Migrated files appear as normal files in the backup data.	Migrated files are converted to normal files.



Note: In the table, D* - Default and S* = Setting

Note that in most cases, NDMP Data Management Applications automatically use the environment variables that were specified on backup when restoring. It is recommended that you consider the desired restore behavior when specifying the backup variables. Other than in exceptional circumstances, use the same environment variables on restore as those specified at backup.

The variables to use for standard use cases are:

- Only backup non-migrated data
 - NDMP_BLUEARC_EXCLUDE_MIGRATED = y
 - Other variables are ignored
- Only backup migrated data, automatically migrating on recovery
 - NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED = y
 - NDMP_BLUEARC_EXCLUDE_MIGRATED = n
 - NDMP_BLUEARC_REMIGRATE = y
- Backup all data, migrated files automatically migrated on recovery
 - NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED = n
 - NDMP_BLUEARC_EXCLUDE_MIGRATED = n
 - NDMP_BLUEARC_REMIGRATE = y

Externally migrated data

Backup/restore of data sets that contain data migrated to external servers are affected by the following NDMP environment variables:

- NDMP_BLUEARC_REMIGRATE
- NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED
- NDMP_BLUEARC_EXTERNAL_LINKS
- The CLI command: **migration-recreate-links-mode**

Note that NDMP_BLUEARC_EXCLUDE_MIGRATED has no effect on externally migrated files.

Environment variable	Default	Setting	Effect on backup	Effect on restore
NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED	n	y	Only externally migrated files are included in the backup data.	Depends on NDMP_BLUEARC_EXTERNAL_LINKS/migration-recreate-links-mode.

Environment variable	Default	Setting	Effect on backup	Effect on restore
		n	Normal files and externally migrated files are included in the backup data.	Depends on NDMP_BLUEARC_EXTERNAL_LINKS/migration-recreate-links-mode.
NDMP_BLUEARC_REMIGRATE	y	y	Externally migrated files are marked as such in the backup data.	Depends on NDMP_BLUEARC_EXTERNAL_LINKS/migration-recreate-links-mode..
		n	Externally migrated files appear as normal files in the backup data.	Externally migrated files are converted to normal files.
NDMP_BLUEARC_EXTERNAL_LINKS	remigrate	ignore	Externally migrated files are not included in the backup data.	Externally migrated files cannot be restored.
		recreate_link	Only the metadata for the externally migrated file is included in the backup data.	The recovery target file system is checked for a migration path. If a path exists, an external link is created (see the CLI migration-recreate-links-mode command). There is no restore of user data. If a migration path does not exist, a link is not created.

Environment variable	Default	Setting	Effect on backup	Effect on restore
		remigrate	The metadata and user data of the externally migrated file are included in the backup data.	The recovery target file system is checked for a migration path. If a path exists, an external link is created (see the CLI migration-recreate-links-mode command), and the user data is written to the external link. The user data is never written to the local file system. If a migration path does not exist, a normal file is created.
migration-recreate-links-mode	test-before-recreate	always-recreate-links	No effect	Creates a link. Does not check if the link is pointing at anything.
		test-before-recreate	No effect	If an existing file (migrated or otherwise) exists, delete it, and then create a new link to the migration target and write the user data to it.

- Note that in most cases, NDMP Data Management Applications automatically uses the environment variables that were specified on backup when restoring. It is recommended that you consider the desired restore behavior when specifying the backup variables. Other than in exceptional circumstances, use the same environment variables on restore as those specified at backup.
- **migration-recreate-links-mode** is set per EVS and not per NDMP operation. It applies to restore only; therefore, it must be set on the EVS that hosts the recovery target file system.
- Do not use `NDMP_BLUEARC_EXTERNAL_LINKS = recreate_link` with **migration-recreate-links-mode = test-before-recreate**. This configuration could lead to data loss.
- When using `NDMP_BLUEARC_EXTERNAL_LINKS = remigrate`, the operation is unsuccessful if the migration path at restore is the same path as when the backup was made. This prevents accidental data loss when copying data between filesystems. It is therefore not possible to recover migrated data to the same filesystem that is the source of the backup using this method. An alternative is to restore to a different location (if no migration path is configured there, migrated files are converted to normal files) and then copy the data into the original location.

The settings to use for standard use cases are:

- Only backup non-migrated data:
 - NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED = n
 - NDMP_BLUEARC_EXTERNAL_LINKS = ignore
 - Other variables are ignored
- Only backup migrated data, automatically migrating on recovery:
 - NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED = y
 - NDMP_BLUEARC_REMIGRATE = y
 - NDMP_BLUEARC_EXTERNAL_LINKS = remigrate
 - migration-recreate-links-mode = test-before-recreate
- Backup all data, automatically migrating data on recovery:
 - NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED = n
 - NDMP_BLUEARC_REMIGRATE = y
 - NDMP_BLUEARC_EXTERNAL_LINKS = remigrate
 - migration-recreate-links-mode = test-before-recreate
- Backup non-migrated data, preserve links to migrated data (requires recovery of data at migration target by another mechanism):
 - NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED = n
 - NDMP_BLUEARC_REMIGRATE = y
 - NDMP_BLUEARC_EXTERNAL_LINKS = recreate_link
 - migration-recreate-links-mode = always-recreate

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	Used on recovery to request Direct Access Recovery (DAR).

Possible value	Notes
	<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 ▪ 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.</p> <p>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 y 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 y to enable, or n to disable. For backups, inline hard linked file support is set to n (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>

Possible value	Notes
	<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> <div data-bbox="553 562 1393 722" style="background-color: #e0f2f1; padding: 10px; border: 1px solid #ccc;"> <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> </div> <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 have been migrated because they are less useful, so they are not replicated in order to save time. 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. <div style="border: 1px solid black; background-color: #e0f2f1; padding: 5px; margin: 10px 0;"> <p> Note: Do not use the default <code>test-before-create</code> setting of the <code>migration-create-links-mode</code> CLI command if you use this option.</p> </div> <ul style="list-style-type: none"> 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.

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_TAKE_SNAPSHOT


Possible value	Notes
y or n	<p>This variable is used to override the default behavior with respect to taking a file system-level snapshot during NDMP backup.</p> <p>By default, HNAS server will take the file system-level snapshot during NDMP backup.</p> <p>This option also applies to the file-level NDMP replication, which is configured through the File Replication Rules of the NAS Manager.</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>

Possible value	Notes
	 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.

Hitachi



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact