

Data Protector NDMP Integration Quickstart Guide

How to setup the Data Protector software NDMP integration and the HP StoreOnce Backup Systems to protect data stored in the HP X9000 and NetApp FAS3070 NAS filers

Technical white paper

Table of contents

Executive Summary	2
Audience	2
Disclaimer.....	2
NAS Concept.....	3
NAS NDMP 2-way or local backup concept.....	3
NDMP Backup	4
Preparation - Prerequisite	4
Configuration Steps.....	5
Creating a NDMP backup specification	16
Performance and Implementation Considerations.....	18
Configuring Data Protector block size and the NDMP blocking factor	19
Deduplication Performance Test Results	21
Licensing Considerations.....	22
Conclusions	22
For more information.....	24



Executive Summary

This white paper provides quick start information on integrating Data Protector with the HP X9000 Network Storage System or NetApp FAS3070 Universal Storage System, for NDMP (Network Data Management Protocol) based backup tasks where the backup target is a HP StoreOnce Backup System.

Key take-away:

- When using HP Data Protector for direct NDMP backups of HP X9000 and NetApp filers the HP StoreOnce Backup System can be used as deduplicating disk backup target
- Additional Data Protector licenses are needed to enable NDMP backups
- The StoreOnce Backup System must be configured to emulate tape devices (Virtual Tape) for the backup targets
- Backing up to the HP StoreOnce Backup System delivers backup storage capacity savings of 90%+ with deduplication in the range 12:1 to 16:1 after around 30 backups. These deduplication ratios are similar to the best deduplication ratios achieved using regular file system backup, that is file system backups using the largest Data Protector configurable block size.
- StoreOnce Backup System replication can be used to create DR copies of virtual tapes at a remote location

Audience

This document is intended for solution architects, project managers, engineers, and support personnel involved in planning, designing and configuring NDMP backup/recovery solutions using HP Data Protector software with HP X9000 NAS technology, NetApp FAS filer technology and HP StorageWorks Backup to Disk technology utilizing HP StoreOnce deduplication for backup data reduction.

Familiarity with the following topics is recommended:

- HP X9000 architecture
- HP Data Protector software
- HP StoreOnce Backup System (formerly: HP StorageWorks D2D) architecture
- NetApp FAS3070 architecture
- NDMP (Network Data Management Protocol)

Disclaimer

The configurations in this document are HP recommended configurations. They are provided as a reference only, as configurations vary with specific customer needs. If mentioned - Memory, processor count and speed, and I/O storage recommendations - should be considered as a minimum recommendation.

NAS Concept

Network-attached storage (NAS) is file-level computer data storage connected to a computer network providing data access to heterogeneous network clients.

Although it may technically be possible to run other software on a NAS unit, it is not designed to be a general purpose server. For example, NAS units usually do not have a keyboard or display, and are controlled and configured over the network, often using a browser.

A fully-featured operating system is not needed on a NAS device, so often a stripped-down or customized standard operating system is used e.g.: Linux.

NAS systems contain one or more hard disks, often arranged into logical, redundant storage containers or RAID arrays (redundant arrays of inexpensive/independent disks) often referred to "Primary Storage".

NAS removes the responsibility of file serving from other servers on the network.

NAS uses file-based protocols such as NFS (popular on UNIX systems), SMB/CIFS (Server Message Block/Common Internet File System) (used with MS Windows systems), AFP (used with Apple Macintosh computers), WebDAV, HTTP/HTTPS and FTP/FTPS to serve data to applications or users.

NAS units rarely limit clients to a single protocol.

The backup of data residing on filer to what often is called "Secondary Storage" can now be done in two different ways:

1. on the Application Server
2. direct from filer to tape using NDMP

The first case would be a "normal" backup, using the Data Protector DA (disk agent), which is installed on each Application Server. In case of a backup data would be transferred via LAN to the system where a tape device is connected to.

The second way uses the NDMP integration of HP Data Protector and performs the backup locally on the NDMP server.

This has two major advantages:

1. The backup data has not to be transferred over LAN
2. No performance degradation on the application server

NAS NDMP 2-way or local backup concept

The NDMP architecture uses a client-server model in which HP Data Protector is the NDMP client (DMA-Data Management Application) to the NDMP-Host Data Mover (NDMP server or DSP-Data Service Provider).

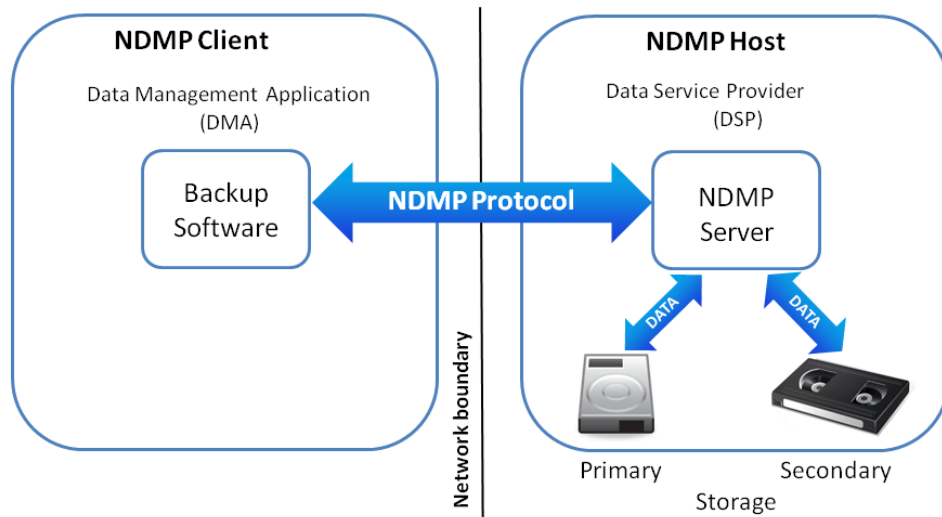
In the event of a NDMP backup, backup data flows from the Primary Storage system to the Data Mover and further on to the Secondary Storage system which can be an attached tape library backup device, without traversing the Ethernet network. Only the backup software's control data, catalog data, and tape library commands are send over the network utilizing the NDMP protocol.

In a 2-way or local backup scenario (Figure 1) both, Primary and Secondary Storage is directly attached to the filer using either SCSI or Fiber channel connections.

The Data Mover maintains a state machine for each NDMP client connection that executes and maintains the backup and/or restore processes.

Note: Please refer to the filer vendor documentation to find supported NDMP session concurrency scheme. HP Data Protector software up to version 6.2 only supports the 2-way or local backup concept.

Figure 1: NDMP 2-way backup architecture



NDMP Backup

Preparation - Prerequisite

StoreOnce Backup System (secondary Storage):

- Setup StoreOnce system and create at least one virtual tape library with the appropriate number of tape drives. Select desired tape emulation and number of tape slots. Make sure library is mapped to an available fiber channel port
- Connect the configured library to your SAN segment.

SAN Router:

- Create Alias for each WWN (secondary (StoreOnce backup device), primary storage (NAS appliance)).
- Create zone and include Aliases
- Enable zone

NAS appliance (primary Storage):

- Setup HP X9000/NetApp FAS3070 system and create at least one file system. Apply LTU were applicable.
- Connect the NAS appliance to your SAN segment (via FC).
- Detect the newly attached tape library/device.
- Configure DSP (NDMP) on NAS appliance.

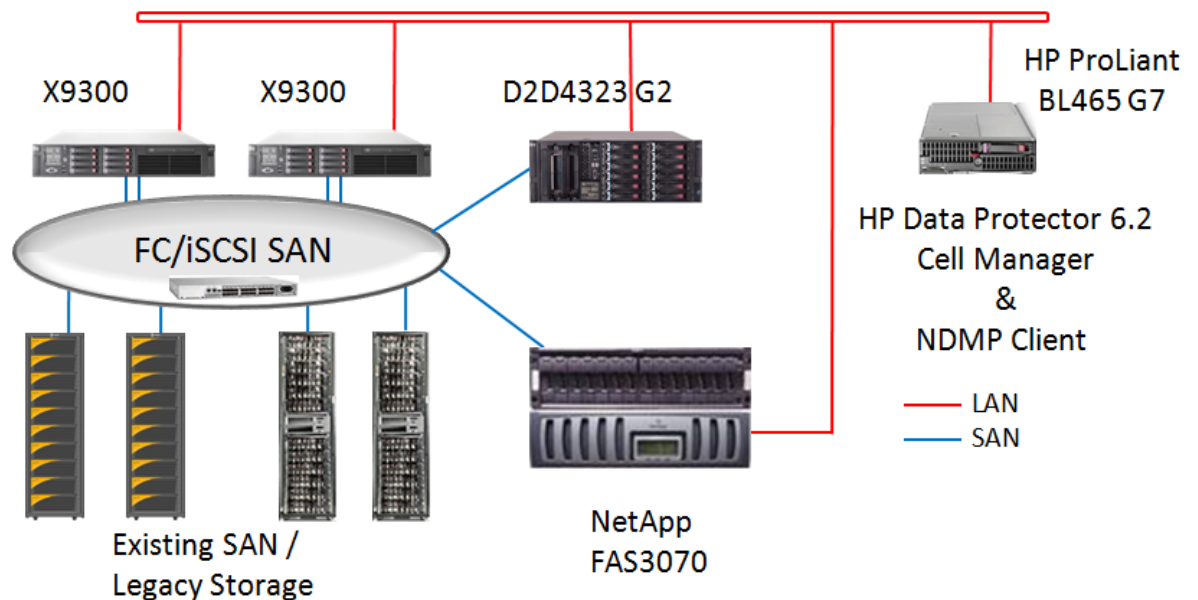
Data Management Application (DMA) & NDMP Client

- Setup HP Data Protector 6.2 cell manager
- Install NDMP Media Agent, which promotes cell manager to a NDMP client
 - Note: The NDMP media agent can be installed on any client within the cell on which operating system it is supported. Always consult the latest support matrix relevant for the used components to find any potential incompatibilities.
- Ensure proper HP Data Protector software licensing for NDMP backups
 - NDMP LTU for all platforms: 1xTB B7022BA, 10xTB B7022DA, 100xTB TD186AA
 - Appropriate number of library slot LTU : 1x 61-250/unlimited slots B6957BA/B6958BA, upgrade to unlimited slots B6958CA
 - Appropriate number of tape drive LTU for SAN,UNIX,NAS: B6953AA

Alternative licensing to tape drive LTU

- Advanced Backup to Disk LTU for all platforms: 1xTB B7038AA, 10xTB B7038BA, 100xTB B7038CA
- Configure NDMP client
- Configure backup library on the HP Data Protector cell manager.

Figure 2: NDMP 2-way backup topology



Note: HP Data Protector software supports library sharing in a SAN environment.

Configuration Steps

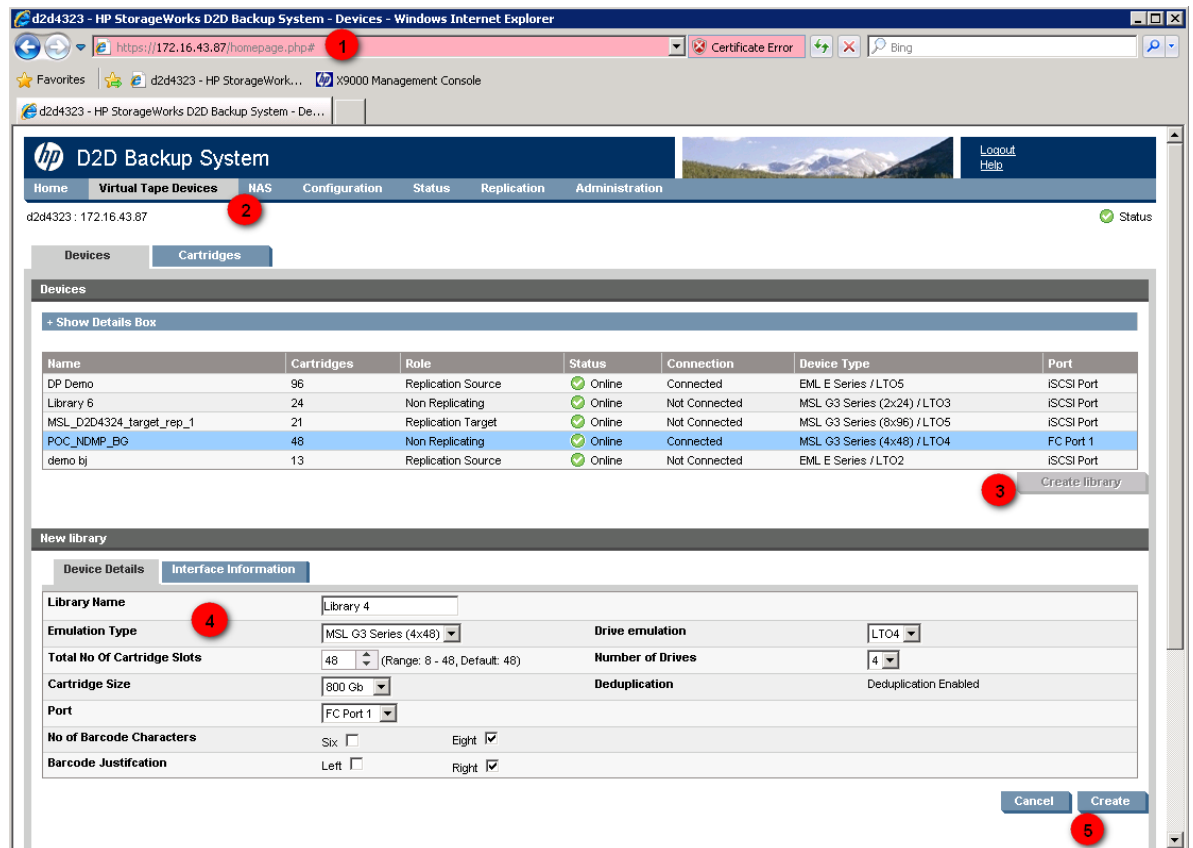
The following pages will highlight the necessary steps to properly configure the backup target (StoreOnce), the DSP (NAS appliance) and the DMA (NDMP client & Library configuration on the cell manager). It will not cover SAN, LAN configuration aspects. Furthermore all DMA (NDMP client) settings are non-modified standard settings.

StoreOnce Backup System (secondary Storage):

The StoreOnce backup system requires just a few steps in order to create and configure a new virtual tape library (VTL). Figure 3 depicts the StoreOnce management GUI. The following are the five steps to create and configure a new library:

1. Logon to the StoreOnce web administration GUI.
Enter the administrator user credentials and password.
2. Select "Virtual Tape Devices" from the menu list
3. Select "Create Library"
This will change the screen in the lower part allowing you to make changes to the library characteristics.
4. Change library characteristics according to your needs. This includes:
 - o Library Name
Robotic emulation type, tape drive emulation
 - o Total number of cartridge slots, number of tape drives
 - o Cartridge size, FC port through which this library will be made accessible
 - o Tape media labeling (barcode) details
5. Select "Create" in order to create the virtual tape library according to your specification

Figure 3: StoreOnce Backup System Administration GUI




Note: Make sure the robotic emulation type and the drive emulation is supported by the DSP as well as the DMA. Always consult the latest support matrix relevant for the used components to find any potential incompatibilities.

The creation of the virtual library will make the new library appear on the StoreOnce management GUI. To verify the interface information you may proceed as following:

6. Select "Interface Information" in the lower part of the screen

Figure 4 depicts a snapshot of the interface information screen.

Figure 4: Interface Information



Device Name	Port	Device Serial Number	FC Address	World Wide Node Name	World Wide Port Name
Media Changer	FC Port 1	CZJ1030J0R	0x0A0405	50014380119CCAB4	50014380119CCAB5
Drive 1	FC Port 1	CZJ1030HXK	0x0A0401	50014380119CC9EE	50014380119CC9EF
Drive 2	FC Port 1	CZJ1030HXL	0x0A0402	50014380119CC9F0	50014380119CC9F1
Drive 3	FC Port 1	CZJ1030HXM	0x0A0403	50014380119CC9F2	50014380119CC9F3
Drive 4	FC Port 1	CZJ1030HXN	0x0A0404	50014380119CC9F4	50014380119CC9F5

Note: "Device Serial Number" and "World Wide Node Name" are configuration details which are used in further configuration tasks such as setting up FC SAN zones on the SAN router or identifying tape devices attached to the DSP (NAS appliance).

DSP on NAS appliance (primary Storage):

HP X9000 models

The HP X9000 Network Storage System requires just a few steps in order to configure the cluster for the NDMP backup feature. Figure 5 depicts a management GUI of an X9000 Network Storage Gateway consisting of two nodes. The following are the steps:

1. Logon to the HP X9000 web management GUI.
Enter the administrator user credentials and password.
2. Select "Cluster Configuration" in the navigator list
3. Select "NDMP Backup" in the Cluster Configuration
4. Select "Modify" and make the necessary changes in the appearing web form. This includes the following details:
 - o NDMP Sessions = Yes
 - o Listener port (typically port: 10000)
 - o Username / Password (must be a user with appropriate backup user rights)
 - o DMA (NDMP client) IP Addresses
 - o Enable NDMP Sessions = 1

Note: The remaining configuration items can be left unmodified.

Select "Synchronize" to distribute the configuration change to all nodes of the cluster

Once the configuration change got distributed make sure all backup libraries which have been configured in the previous step are discovered and visible on the HP X9000 Network Storage System.

5. Select "Tape Devices" to verify discovered backup components. Refer to Figure 6 for details.

6. Select "Rescan Devices" if newly attached devices have not shown up.

Figure 5: HP X9000 Network Storage Gateway Administration GUI

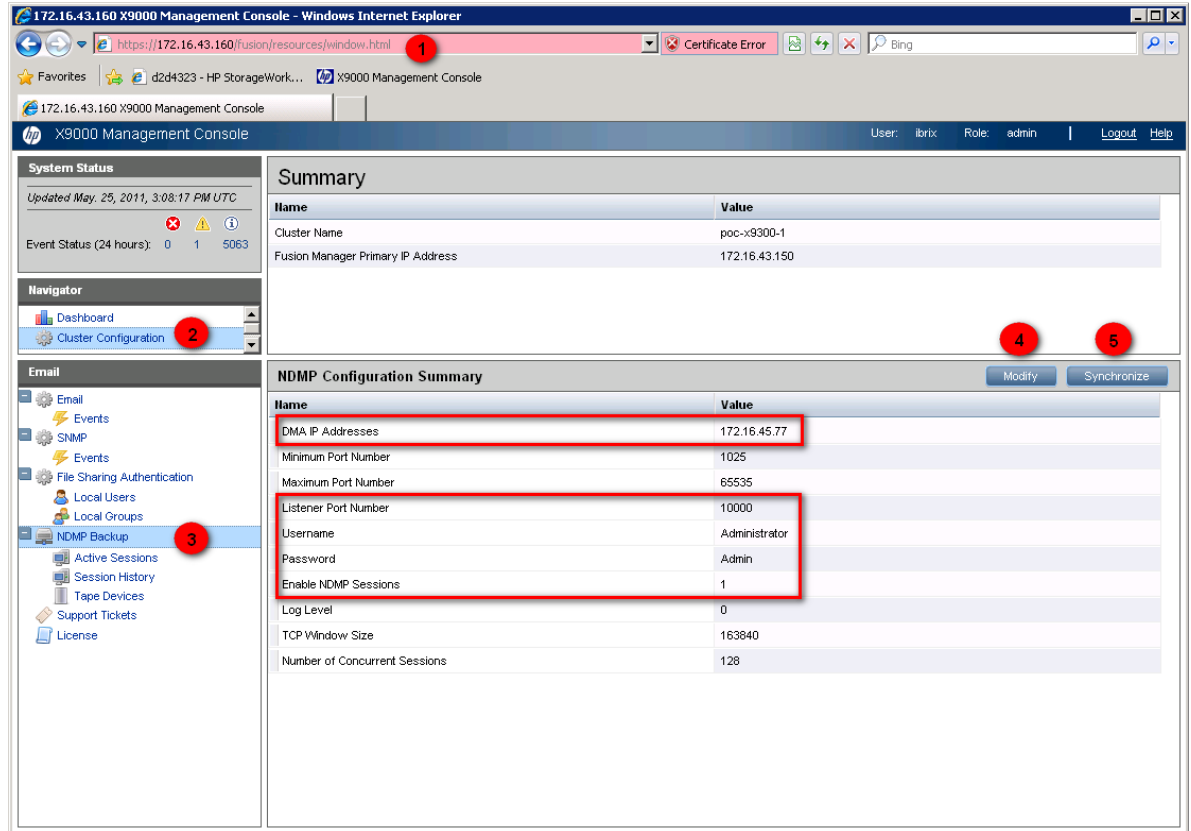
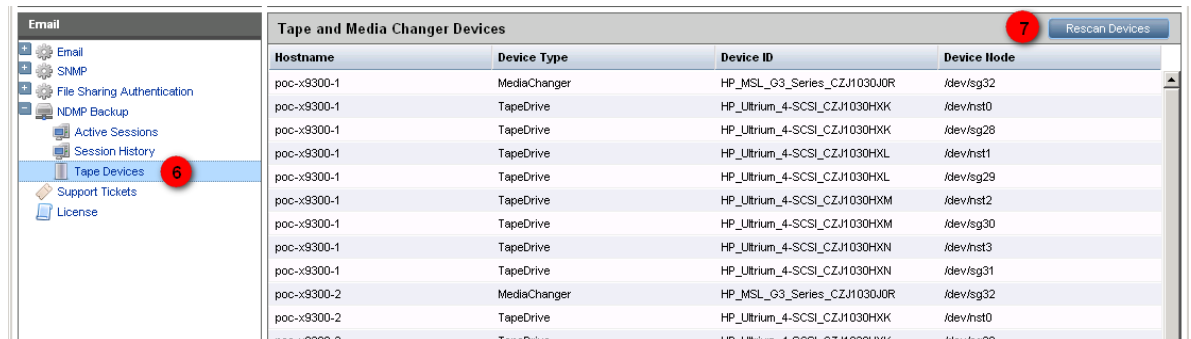


Figure 6: Discover Tape Devices



Note: "Hostname" and "Device ID" are configuration details which are used in further configuration tasks within the DMA (NDMP client) configuration. Each segment server (Hostname) of the X9000 cluster might discover the attached backup devices (Device ID).

NetApp FAS3070

The NetApp FAS3070 platform has very limited access to NDMP configuration items through either the web GUI or through the NetApp System Manager plug-in.

Note: Interaction to NDMP commands through the web GUI is limited to the following:

- Enable/Disable NDMP - (The Enable/Disable NDMP Services page enables you to enable or disable the NDMP service on the storage system).
- Terminate NDMP sessions - (The Terminate an NDMP Session page enables you to terminate a current active NDMP session).

There is no NDMP management functionality available through the NetApp System Manager plug-in.

It is recommended to use a FTP/SSH client such as PUTTY in order to manage the NDMP functionality using the CLI of the FAS3070 controller.

To verify newly attached tape library components such as tape libraries (medium changers) or tape devices which are recognized and configured on the FAS3070 controller proceed as following:

Logon to the FAS3070 controller using the telnet/ssh client. Provide administrator credentials and password to get the system prompt.

To verify the supported tape device types enter:

```
tpc008> storage show tape supported [-v]
```

If no options are given, the list of supported tape drives is displayed. The following option is supported:

-v

The -v option displays all the information about the supported tape drives including their supported density and compression settings.

Note: Make sure to only select backup tape device types in your StoreOnce backup system configuration which appear on the appearing list (e.g.: LTO4)

To display information about attached tape drives and tape libraries (medium changers), and their associated serial numbers enter the following command:

```
tpc008> storage show {tape | mc} [{alias | PPN | WWN}]
```

alias is the logical name of the tape drive or medium changer.

PPN is the physical path name.

WWN is the worldwide name.

The following paragraph shows the result of a medium (mc) changer query:

```
Medium Changer:      fcs009:0.129
Description:         HP          MSL G3 Series
Serial Number:       CZJ04002H3
WWNN:                5:001:438007:5a2878
WWPN:                5:001:438007:5a2879
```

Alias Name(s) : mc4
Device State: available

Note: The "Serial Number" must match the device serial number of the configured library on the StoreOnce backup system. The "Alias Name" is the device file to be used in the DMA (NDMP client) configuration

The following paragraph shows the result of a tape (tape) device query:

Tape Drive: fcs009:0.124
Description: HP Ultrium 4-SCSI
Serial Number: CZJ04002D1
WWNN: 5:001:438007:5a27ba
WWPN: 5:001:438007:5a27bb
Alias Name(s) : st41
Device State: available

Note: The "Serial Number" must match the device serial number of the configured tape device on the StoreOnce backup system.

The "Alias Name" is the device file to be used in the DMA (NDMP client) configuration and can be verified by issuing the following command:

```
tpc008> sysconfig -t
```

The sysconfig -t command displays device and configuration information for each tape drive.

If you have a tape device that Network Appliance has not qualified, the sysconfig -t command output for that device is different from that for qualified devices. If the filer has never accessed this device, the output indicates that this device is a non-qualified tape drive, even though there is an entry for this device in the /etc/clone_tape file. Otherwise, the output provides information about the qualified tape drive that is being emulated by this device.

The following paragraph shows the result of a sysconfig -t device query:

```
Tape drive (fcs009:0.124) HP Ultrium 4-SCSI
rst41l - rewind device, format is: LTO-2(ro)/3 2/400GB
nrst41l - no rewind device, format is: LTO-2(ro)/3 2/400GB
urst41l - unload/reload device, format is: LTO-2(ro)/3 2/400GB
rst41m - rewind device, format is: LTO-2(ro)/3 4/800GB cmp
nrst41m - no rewind device, format is: LTO-2(ro)/3 4/800GB cmp
urst41m - unload/reload device, format is: LTO-2(ro)/3 4/800GB cmp
rst41h - rewind device, format is: LTO-4 800GB
nrst41h - no rewind device, format is: LTO-4 800GB
urst41h - unload/reload device, format is: LTO-4 800GB
rst41a - rewind device, format is: LTO-4 1600GB cmp
nrst41a - no rewind device, format is: LTO-4 1600GB cmp
urst41a - unload/reload device, format is: LTO-4 1600GB cmp
```

Note: The device file for the DMA (NDMP client) configuration can now be verified. It is advisable to choose a "no rewind device" such as "nrst41a" or "nrst41h".

To enable NDMP services on the NDMP host/server issue the following command:

```
tcp008> ndmpd on
```

Note: The NDMP listener is on port 10000

To create a NDMP user password which is used in the DMA (NDMP client) configuration, issue the following command:

```
tcp008> ndmpd password <username>
```

This will create a NDMP user password similar to:

```
tpc008> ndmpd password Administrator  
password ZVC5D0JU0JUw0BhV
```

Note: Cut and paste the generated password into your "Import Client" task used on the DMA.

NDMP Client & DMA

The following are high-level steps which are required to setup the NDMP backup environment from a HP Data Protector perspective:

1. Push NDMP Media Agent component to designated NDMP Client.
2. Import NDMP host/server
3. Create dedicated NDMP media pool.
4. Create backup library configuration
5. Scan library slots
6. Format media in NDMP media pool

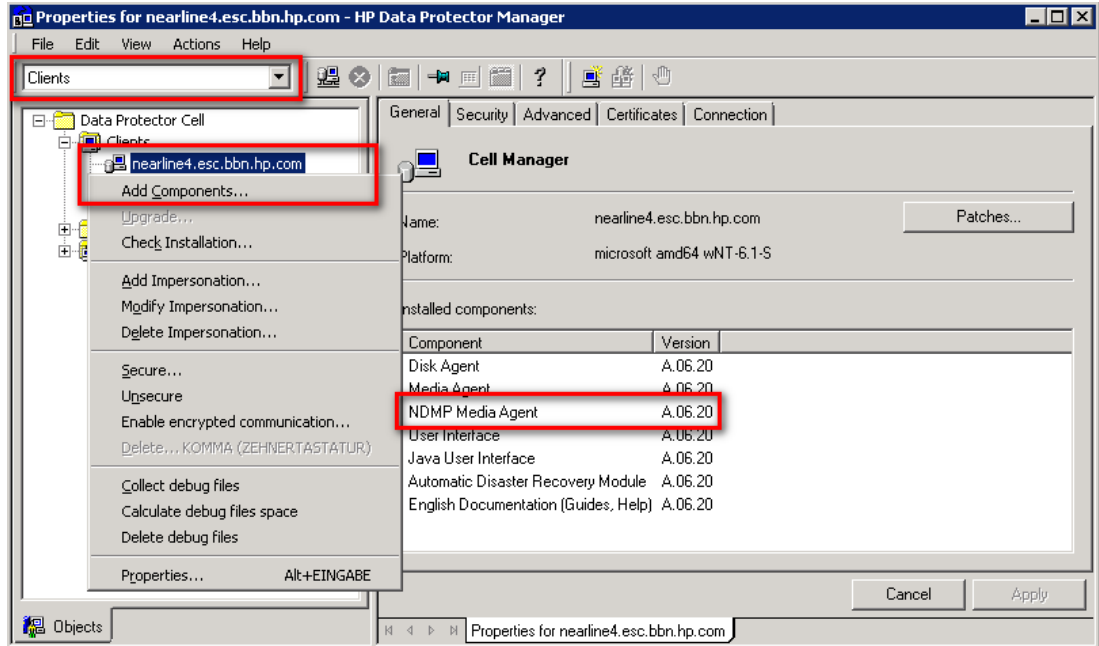
The following screenshots depict the sequence of tasks to complete the above listed high-level steps.

Step 1: Push NDMP Media Agent component to NDMP client.

You can install/upgrade Data Protector clients remotely, by distributing the software using the Installation Server.

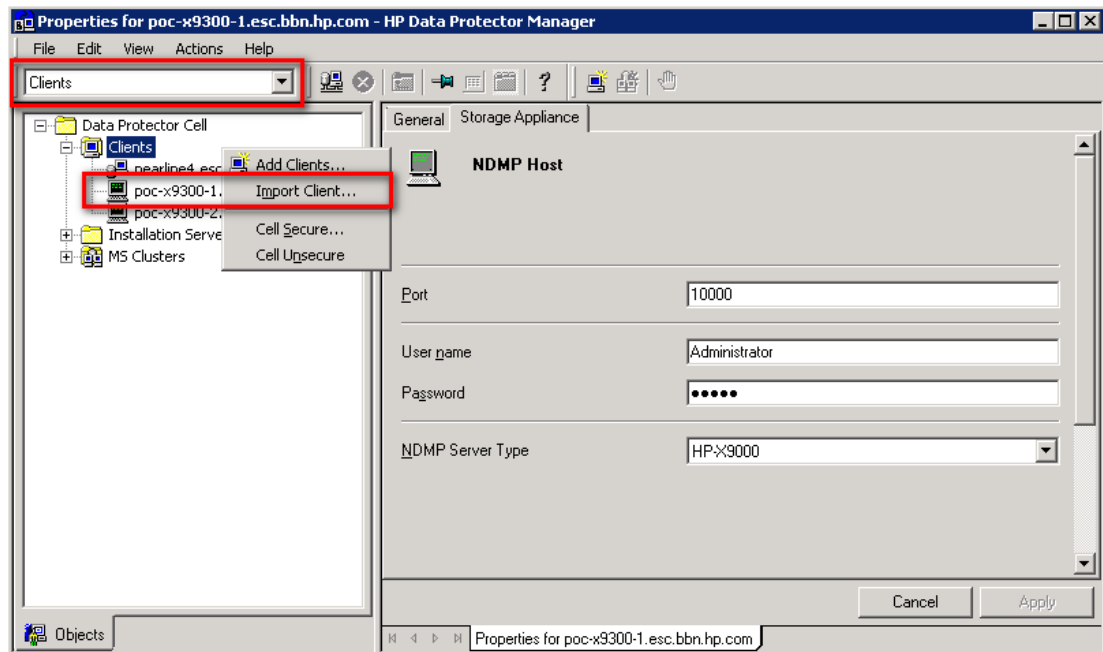
Note: The NDMP client can be any client in the cell as long as the NDMP media agent is supported on the OP/Sys platform this client runs on. Even though the installation process reports that the standard media agent got removed it will still be listed on the client's component list because the NDMP media agent provides the same functionality plus NDMP backup/restore functionality. Hence the NDMP media agent fulfills the legacy role of that client too.

Figure 7: Push NDMP Media Agent component to NDMP client



Step 2: Import NDMP host/server.

Figure 8: Import NDMP host/server



Note: Importing a client means manually adding a computer to a DP cell. On NDMP hosts/servers there will no Data Protector component be installed. When added to a Data Protector cell, the system

becomes a Data Protector client. Once the system is a client of the cell, information about the new client is written to the IDB, which is located on the Cell Manager.

Specify information about the NDMP Server.

NDMP Server Type

Select one of the available NDMP host types that matches your setup: NetApp, Celerra, BlueArc, Hitachi or HPX9000.

Port

Specify the TCP/IP port number of the NDMP Server. The default is 10000.

User Name

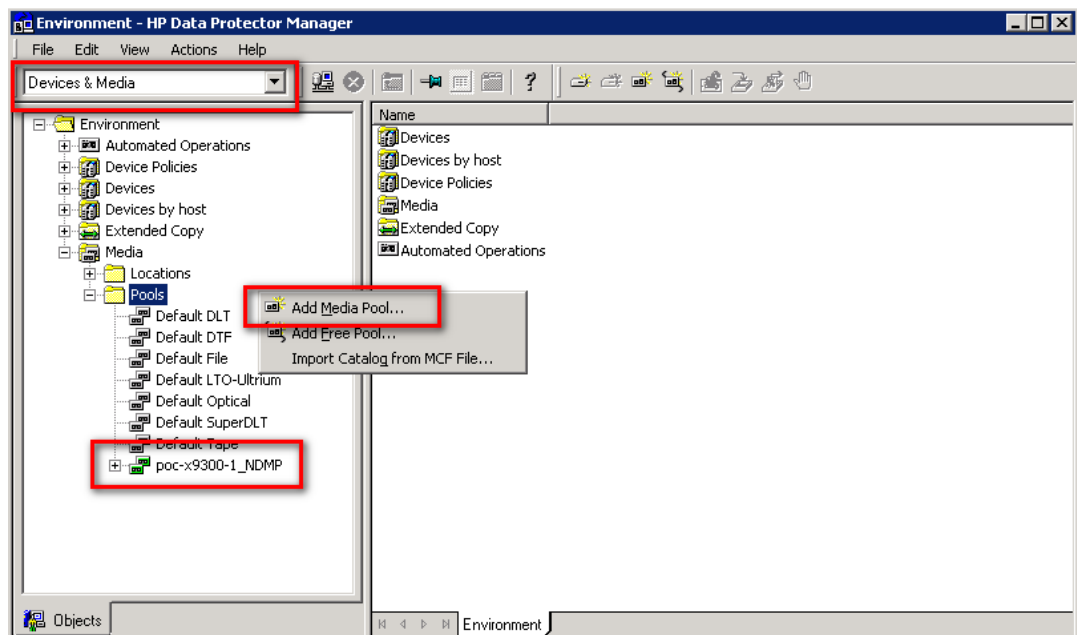
Type the user name that Data Protector will use to establish the connection to the NDMP Server. Default might be “ndmp”. Check vendor documentation for a valid “ndmp” user name and password.

Password

Type the password of the user. For NetApp FAS3070 use the generated password shown in one of the previous steps.

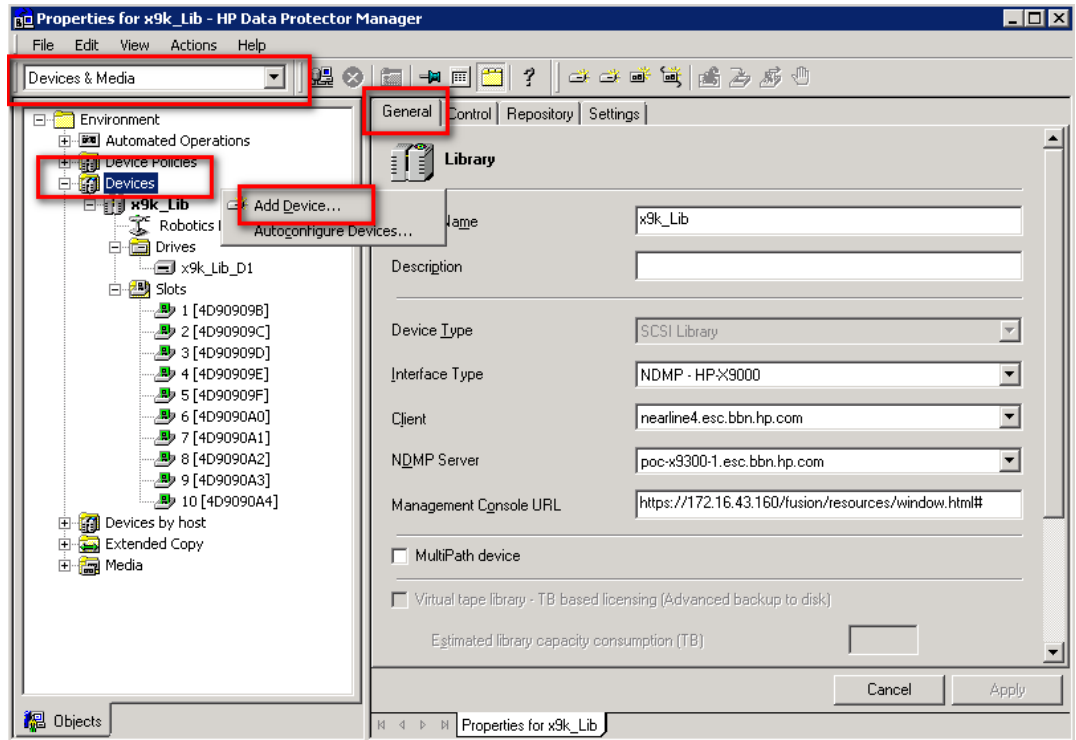
Step 3: Create NDMP media pool

Figure 9: Create NDMP media pool



Note: Data Protector provides default media pools, but you must create your own media pool to suit your needs. You must select the “Media Type” that matches the tape drive technology on your StoreOnce backup device. E.g.: select LTO-Ultrium if you have selected LTO4 on the backup device.

Figure 10: Create a Backup library configuration

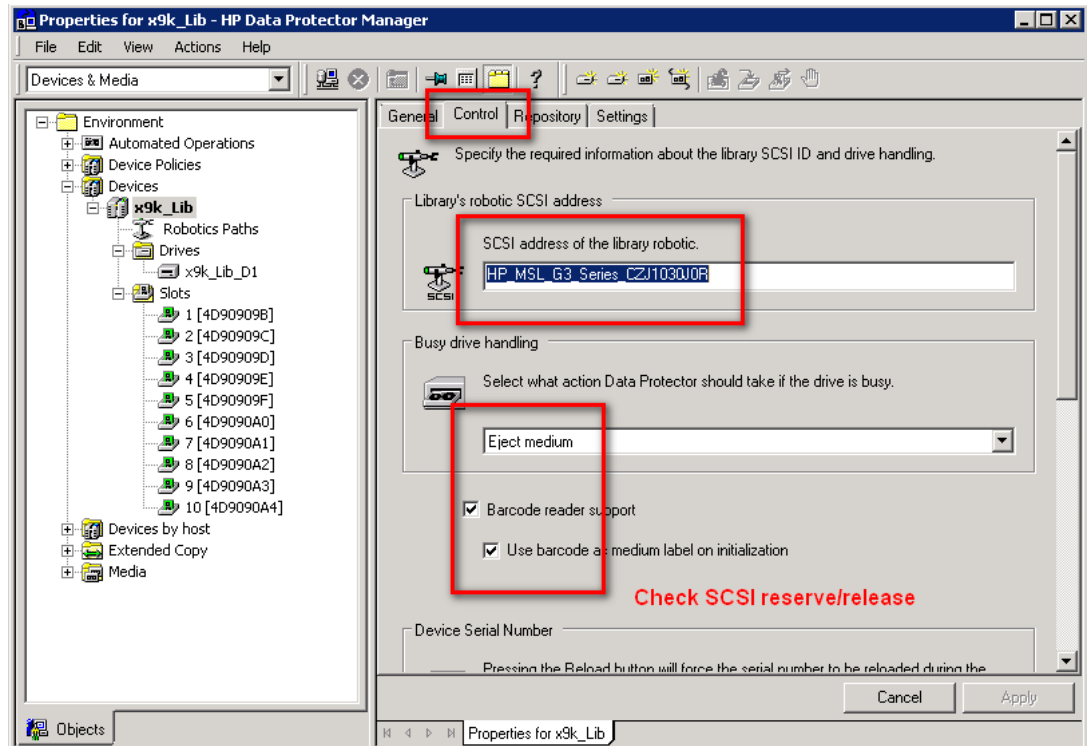


Preparation of a backup device consists of connecting the device to the system or in a SAN environment to the SAN segment and knowing which of the (working) associated device files (SCSI address) is to be used in the HP Data Protector configuration.

The library robotic as well as any of the tape drive units requires the following common configuration items:

1. In the Device Name text box, enter the unique name of the device.
2. In the Description text box, enter a description (optional).
3. In the Device Type list, select the SCSI Library device type.
4. In the Interface Type list, select the SCSI interface type.
Note: Depending on which host, the NDMP host or NDMP client should control the library robotic you select the proper host which controls the library robotic. Useful for library sharing purposes.
5. In the Client list, select the client which handles the NDMP communication (NDMP client).
6. In the NDMP server list, select the NDMP server.
7. Enter a valid URL of the library management console in the Management Console URL text box. (Optional for the robotic configuration)

Figure 11: Enter SCSI and Media Handling details



SCSI library devices are large backup devices, also called autoloaders. They consist of a number of media cartridges in a device's repository and can have multiple drives handling multiple media at a time.

A typical library device has a SCSI ID (Windows) or a device file (UNIX) for each drive in the device and one for the library's robotic mechanism, which moves media from slots to drives and back again. (For example, a library with four drives has five SCSI IDs, four for the drives and one for the robotic mechanism).

A medium is stored in a slot in the device's repository. Data Protector assigns a number to each slot, starting from one. When managing a library, you refer to the slots using their numbers.

Continue with the following steps:

8. Enter the SCSI address of the library robotics.
Note: Use the SCSI address as it appears on the X9000 GUI (Device ID; e.g.: HP_MSL_G3_Series_CZJ1030J0R) or as device file as it appears on the NetApp CLI (e.g.: mc4)
9. In the Busy Drive Handling list, select the action Data Protector should take if the drive is busy
10. Select SCSI Reserve/Release (robotic control).
11. Specify the slots for the device. Use a dash to enter slot ranges and then click Add. For example, enter 1-3 and click Add to add slot 1, 2, and 3 at once. Do not use letters or leading zeros.
12. In the Media Type drop-down list, select a media type for the device that you are configuring.
13. Click Finish to exit this wizard. You are prompted to configure a library drive. Click Yes and the drive configuration wizard appears.

Note: Creating a tape drive configuration requires the same steps as for the library robotic except that you have to care about the following details:

- Enter the correct drive index
- Assign the previously created NDMP media pool to every drive configuration you create
- Verify the correct tape block size
- Use “Lock Names” when appropriate

The SCSI address follows the same convention as described for the library robotic.

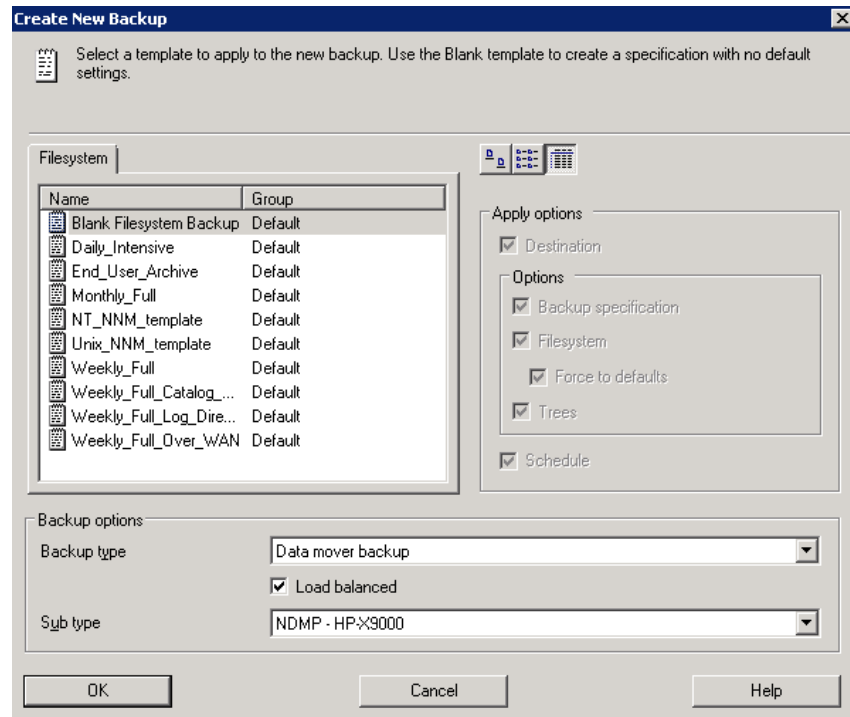
Once the device configuration is finished you must scan the library slots and format the tape media to prepare the media pool for NDMP backups.

Creating a NDMP backup specification

Creating a NDMP backup specification in the Data Protector GUI follows a similar approach as for regular file system backups using the “normal” Disk agent (DA) except for a view details.

In the backup context expand “Backup Specifications” and right-click on “Filesystem” to start the backup specification wizard (see Figure 12) by selecting “Add Backup”.

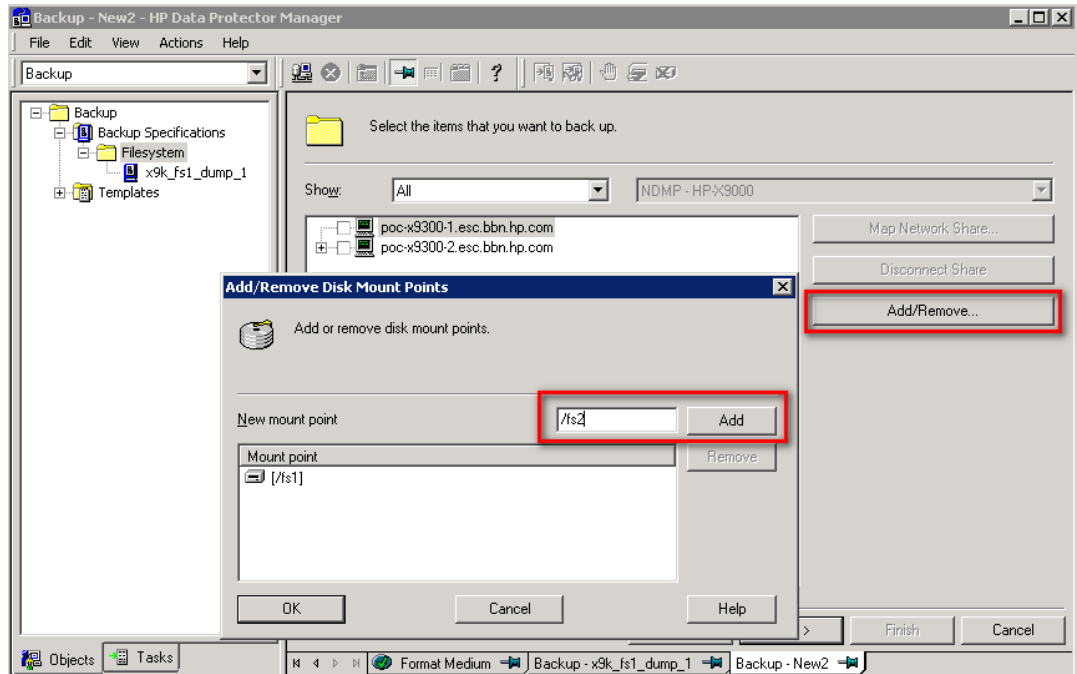
Figure 12: Adding a Backup specification using the GUI Wizard



Select “Data mover backup” as the Backup type and your NDMP interface type in the Sub type field.

Continue in the wizard (see Figure 13) and start adding the backup source to the new backup specification.

Figure 13: Adding the Backup source for HP X9000

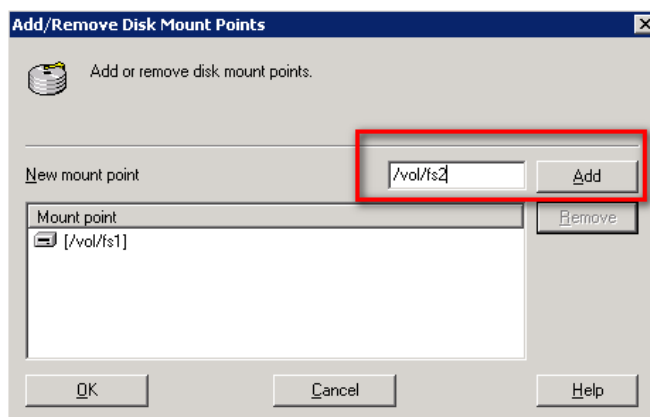


As there is no Data Protector DA (Disk Agent) installed on any of the NDMP server platforms you must manually specify the mount point to be included into the backup specification. You cannot browse existing file systems and have them added to the backup specification.

Each filer vendor has a specific syntax on how to add mount points to the backup specification.

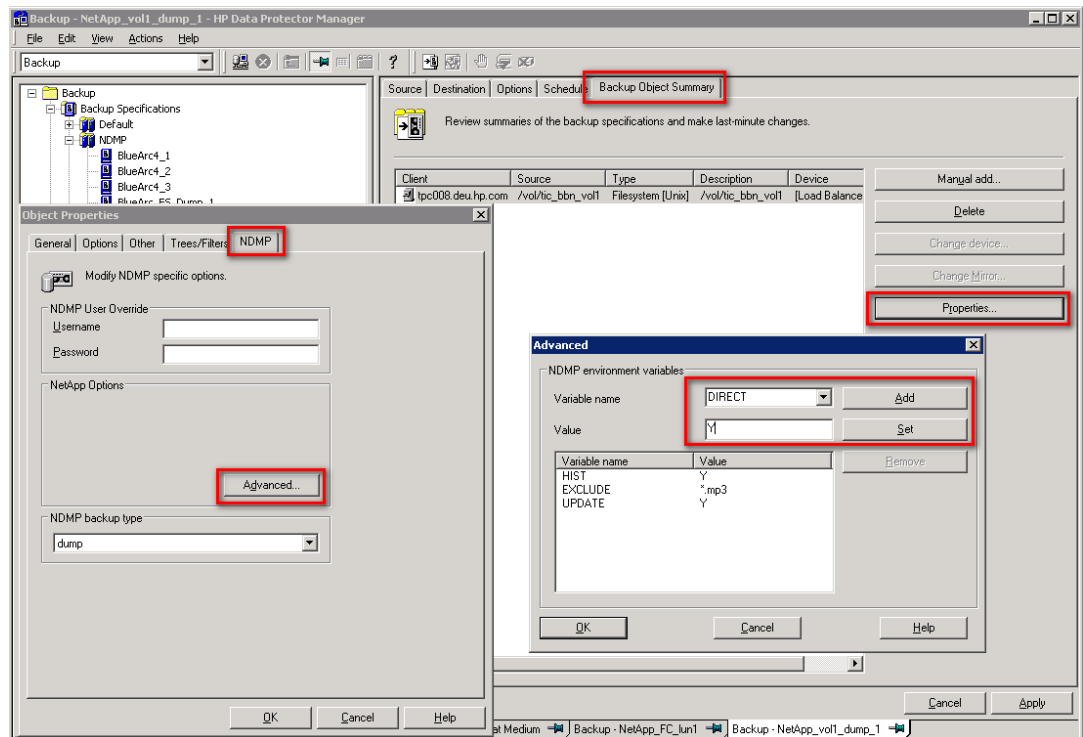
For HP X9000 you must enter the file system name right after the root slash for e.g.: /fs2, whereas for NetApp FAS3070 the syntax is as following: /vol/fs2 (see Figure 14).

Figure 14: Add Disk mount point syntax for NetApp FAS3070



Continue in the wizard and make your selection for the backup destination, global file system options and backup schedule. On the Backup Object Summary tab of the wizard you may configure each backup object individually (see Figure 15).

Figure 15: Backup Summary



On the Advanced NDMP Object Properties you can now specify your backup specific environment variable such as EXCLUDE to exclude specific file types or directories.

Note: Consult the NAS filer administrator handbook for supported environment variables and their usage.

Performance and Implementation Considerations

As described earlier in this paper there are two methods of backing up the data on the NAS filers. The first method uses a regular file system backup of application server. The NAS share is backed up using the same process as the backup of a file system on a local disk. This method is well understood and simple but not optimized for NAS filer protection.

The second method uses NDMP to enable a direct backup from the filer to the HP StoreOnce Backup System. This requires additional setup and licensing, as covered in detail in this paper, but is optimized for NAS filer protection. The performance and implementation considerations for both methods are summarized in Table 1.

Table 1: Summary of performance and implementation differences between the two methods of protecting NAS filer data

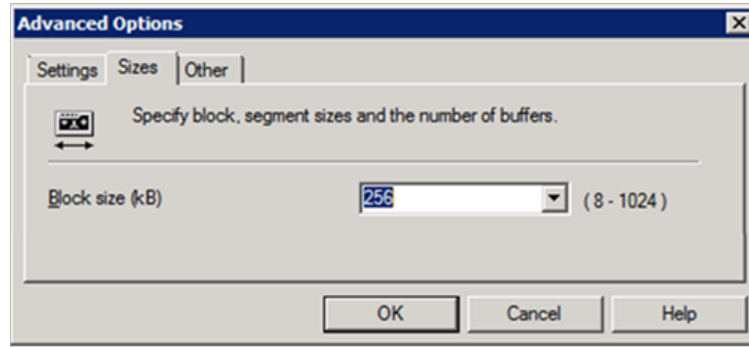
	Regular file system backup	NDMP controlled backup
Impact on production servers	Higher, the Data Protector Disk Agent runs on each protected machine. Requires a backup window.	None, all processing is done on the filer. No impact on data network.
Impact on data network	Higher, all backup data is transmitted over the data network contending for bandwidth with the production data. Additional impact if the StoreOnce Backup System is connected using the data network.	None, the backup data is written over the storage network to the StoreOnce Backup System.
Backup flexibility	Simple process to include and/or exclude files and directories for backup through the Data Protector GUI.	Standard dump format enables control over what is backed up via file and directory exclusion policies implemented using vendor specific environment variables.
Backup throughput	Depends upon 'front end' performance of filer to the clients: protocol used and filer connections to the data network (10GbE) and design of backup infrastructure. Protection can be configured using multiple backup objects to enable multiple parallel backup streams to be written to StoreOnce Backup System for best performance.	Depends on 'back end' performance of filer to the StoreOnce Backup System (8Gb Fibre Channel) and performance of backup format (dump). Multiple virtual tape drives (one per backup object) can be used to enable multiple parallel backup streams during backup to be written to StoreOnce Backup System for best performance. Data Protector supports load balancing for NDMP backups for optimal virtual drive usage.
Backup data deduplication (for a given data set and change rate)	Key variable is the block size set in Data Protector – larger block sizes have higher deduplication ratio.	Test results showed block size set in Data Protector and NDMP BLOCKING_FACTOR have no significant effect of deduplication ratio.
StoreOnce Backup System backup targets supported	All: Virtual tape drives via FC or iSCSI and Data Protector File Library on NAS Share.	Requires a virtual tape target. Connectivity over FC because no iSCSI support on filers.
Restore flexibility	All or selected files/directories restored to original or alternate location.	All or selected files/directories restored to same or alternate location on same or same model of filer if HISTORY environmental variable = "YES".
Data Protector licensing (assuming this is a new installation)	For backup and recovery management: <ul style="list-style-type: none"> • Cell Manager • Disk Agents (no cost) • Media Agent (no cost) For StoreOnce Backup System <ul style="list-style-type: none"> • Advanced Backup To Disk license for usable capacity of appliance --- or --- • Drive and Slot licenses for virtual tape library(ies) used 	For backup and recovery management: <ul style="list-style-type: none"> • Cell Manager • NDMP Media Agent (no cost) • Direct Backup Using NDMP license for capacity of protected data on filer For StoreOnce Backup System <ul style="list-style-type: none"> • Advanced Backup To Disk license for usable capacity of appliance --- or --- • Drive and Slot licenses for virtual tape library(ies) used

Configuring Data Protector block size and the NDMP blocking factor

The Data Protector block size determines, within the limits of the backup device, the size of the data units processed. These data units are called blocks. The backup device processes data in the specified block sizes. The block size is specified as backup device parameter as shown in Figure 16.

The default block size value for all devices is 64kB. Increasing the block size can improve backup performance and deduplication ratio for regular file system backup. For NDMP backup changing the block size has no significant effect on deduplication ratios.

Figure 16: Specifying the Data Protector block size used to write the backup to the virtual tape drive.



For a NDMP controlled backup the data stream is not in a Data Protector format. The default format is dump (others are possible). This can use different blocking factors. The blocking factor determines the size of the units of backup data moved and processed by the NDMP Server.

The blocking factor is specified in Data Protector as a NDMP environment variable in the backup object properties as shown in Figure 17. The syntax and default values for the blocking factor parameter are different for the HP X9000 and NetApp FAS3070 as shown in Table 2.

Figure 17: Specifying the blocking factor environment variable in the backup object properties.

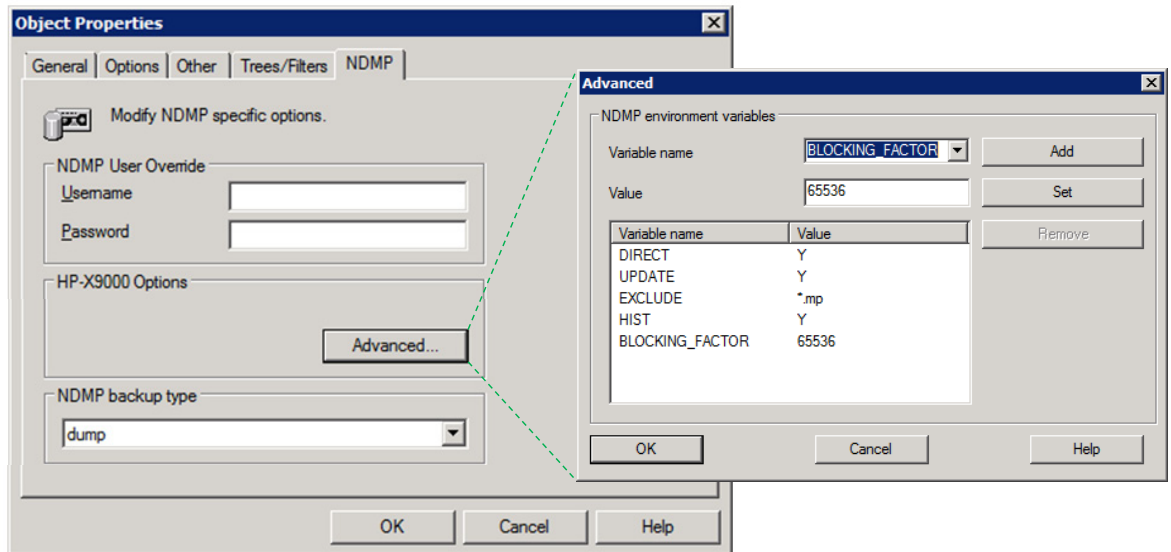


Table 2: HP X9000 and NetApp FAS3070 vendor specific variables and values for the NDMP Blocking Factor.

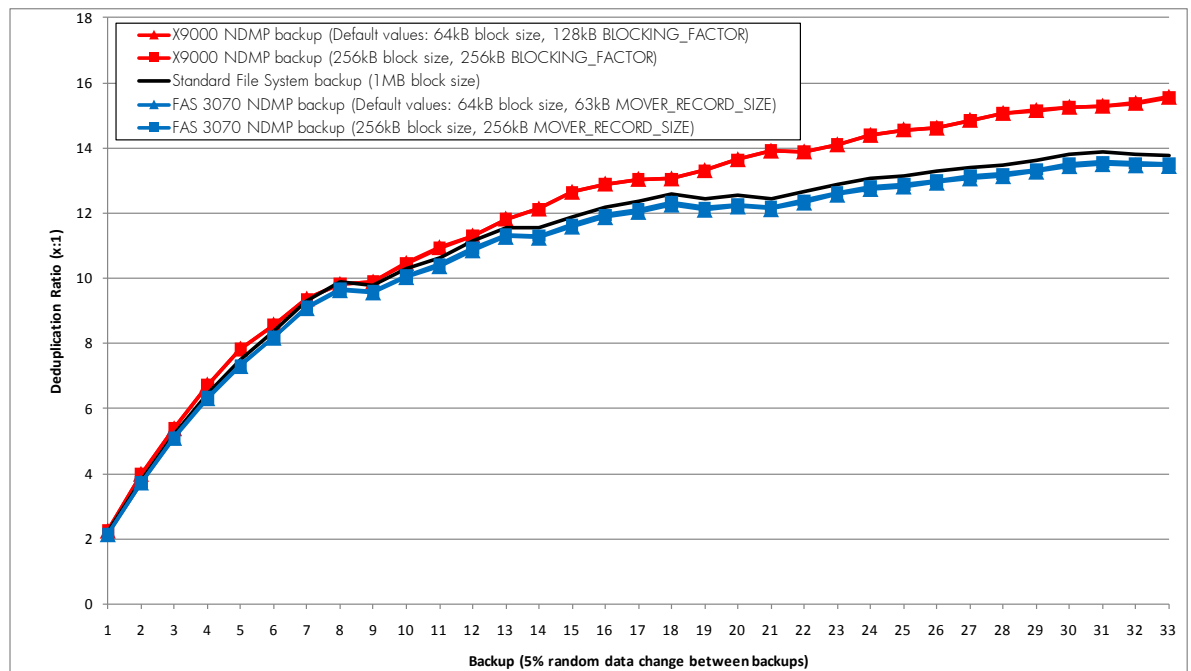
NAS Filer	Blocking Factor Parameter	Default Value	Other Available Settings
HP X9000	BLOCKING_FACTOR	128kB	64kB - 256kB
NetApp FAS3070	MOVER_RECORD_SIZE	63kB	4kB - 256kB

Deduplication Performance Test Results

As part of the validation of Data Protector managed protection of HP and NetApp NAS filers using a HP StoreOnce Backup System as the backup storage the backup data deduplication rates were measured. The test data set was made up of unstructured file data from a production environment. A backup scheme that stimulated daily full backups was used. Between each backup a random selection of 5% of the data was changed.

The tests were run multiple times from the HP X9000 and the NetApp FAS3070 to investigate the effects of varying the Data Protector block size and the NDMP blocking factor. The results are shown in Figure 18.

Figure 18: Comparison of deduplication ratio using different tape block sizes and different NDMP blocking factors versus a regular file system backup of the same data to a HP StoreOnce Backup System. Backup configuration was for full backups with a random 5% of the data changed between each backup.



Licensing Considerations

As shown in Table 1, the file system and NDMP controlled methods of protecting the NAS filer require different Data Protector licenses. Both require a Data Protector Cell Manager and either Advanced Backup To Disk or Drive and Slot licenses for the StoreOnce Backup System.

The difference is that the NDMP controlled backup requires a Direct Backup Using NDMP license. The licensing for this is calculated based on the configured capacity of the NAS filer. The result is the NDMP based method will have higher Data Protector licensing costs.

Conclusions

- The main benefits of NDMP direct backups are delivered when:
 - There is a need to minimize the backup load on the data network, application servers and backup media servers.
 - The filer is required to be backed up as a complete entity. (Although files can be excluded using exclude commands in the NDMP environmental variables set in the NDMP backup job.)
- The benefits of NDMP direct backups should be traded off against the fact that NDMP backups can be more expensive to implement (require additional licensing such as HP Data Protector Direct NDMP backup license) and more complex to manage/configure compared to regular file system backups.
- Backing up to a HP StoreOnce Backup System enables a deduplication ratio of between 12:1 and 16:1 over a period of 20+ backups with a simulated typical daily change rate of 5%. This provides

approximately 90% + increase in backup storage utilization compared to backing up the same data to a physical tape device or non- deduplicating disk.

- NDMP direct backups achieve similar deduplication ratios as file system backups using the largest Data Protector block size. (NDMP backups achieve higher deduplication ratios than file system backups using smaller Data Protector block sizes).
- The NDMP direct backup testing was conducted with various HP Data Protector tape block sizes along with HP X9000 NDMP BLOCKING_FACTORS and NetApp FAS3070 NDMP MOVER_RECORD_SIZES. The results show that changing tape block size and NDMP blocking factor had no appreciable effect on the backup data deduplication ratio. The recommendation is to use the default values in both HP Data Protector and the NAS filer.
- Testing showed that the HP X9000 implementation of NDMP backup delivers an approximate 10% + better backup data deduplication ratio compared to the NetApp FAS3070.

For more information


To read more about:

HP Data Protector software, go to www.hp.com/go/dataprotector

HP X9000 Network Storage System, go to www.hp.com/go/x9000

HP StoreOnce Backup System, go to www.hp.com/go/d2d

HP Data Deduplication, go to www.hp.com/go/storeonce



Get connected
www.hp.com/go/getconnected

Current HP driver, support, and security alerts
delivered directly to your desktop

© Copyright 2011 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

HP Data Protector is a registered trademark of Hewlett-Packard.

4AA3-5894ENW, Created August 2011

