



Technical report:
**Symantec Veritas NetBackup 6.0 and IBM
System Storage N series with NearStore
and SnapVault**

Deployment and implementation guide

• • • • • • • • •

Document NS3465-0

October 17, 2007



Table of contents

Abstract	4
Introduction	4
Intended audience	4
Purpose.....	4
Prerequisites and assumptions.....	4
Solution overview	5
File system export.....	7
Space-optimized image	7
Concepts and How It Works	7
Dense volumes	7
SV-NBU	8
Media server to IBM Storage System N series protocol	8
Processing of data from media server.....	9
Creation of directories and files in the AFS	9
Quick start: Installing and configuring	10
Requirements overview	10
NetBackup	10
Data ONTAP.....	12
Licensing	12
NDMP authentication	13
Enabling NDMP on IBM N series	13
NetBackup 6.0 NDMP Graphical User Interface	14
NetBackup 6.0 NDMP Command Line Interface.....	18
SV-NBU configuration.....	21
IBM N series SV-NBU Configuration.....	21
NetBackup SV-NBU Configuration.....	22
SV-NBU Restores.....	33
Drag-and-Drop via FSE.....	38
Operating characteristics	40
SV-NBU target environment	40
SV-NBU performance	41
SV-NBU storage savings	42
SV-NBU storage overhead	46
SV-NBU limitations	47
SV-NBU best practices	48
Common problems and troubleshooting	49
SV-NBU	49
Verify backup	51
Kernel settings for NetBackup	52
NetBackup tuning parameters	53
Tape and disaster recovery scenarios	54



Tape	55
SV-NBU	55
SnapMirror	56
SV-NBU and SnapMirror	56
Technology primer	57
Symantec NetBackup	57
Architecture	57
Storage units	59
IBM N series with Data ONTAP	59
IBM N series with WAFL	59
IBM N series with Snapshot and SnapRestore	60
Appendix 1: Acronyms	61
Trademarks and special notices	62



Abstract

This deployment guide describes and educates the reader about optimized Symantec Veritas NetBackup disk-to-disk backup solutions using IBM System Storage N series with NearStore feature as the storage and target destination and using IBM System Storage N series with SnapVault for NetBackup technology as the advanced enabling tool. Refer to other technical publications for updates on processes, command syntax and the latest requirements, issues and limitations.

Introduction

Intended audience

This technical report is designed for system engineers and Professional Services engineers who seek education on the joint IBM® System Storage™ N series with NearStore® feature and SnapVault® for NetBackup solution that IBM and Symantec offer. It also prepares the audience for deploying the technology in customer environments.

It will be most beneficial to those who are already familiar with IBM N series hardware and software and Symantec NetBackup, although a high-level “Technology Primer” section is included for readers to educate themselves or get a quick review of the basic technologies involved. Individuals at Symantec, IBM, and the channel partners will all benefit.

Purpose

The purpose of this paper is to present a guide for implementing the SnapVault for NetBackup (SV-NBU) disk-to-disk backup solution, in an IBM N series-NearStore storage environment, addressing step-by-step configuration examples as well as introducing known caveats and recommendations to assist the reader in designing an optimal solution. Its use is three-fold:

1. Provide detailed information to all interested parties.
2. Educate prior to performing deployments.
3. Serve as a reference for resolving issues that could arise.

(It is also a sort of catch-all repository for relevant technical topics that are not covered elsewhere.)

This document is not:

- A sales guide (although some high-level thoughts are covered in the “Solution Overview” section)
- A competitive comparison
- A complete product design document.

Prerequisites and assumptions

For the details and procedures listed in this document to be useful, the following assumptions are made:

The reader has general knowledge of backup and disaster recovery (DR) solutions.

The reader has general knowledge of IBM N series with NearStore feature platforms and products, particularly in the area of data protection.

The reader has general knowledge of Symantec NetBackup 6.0.

The features described here are specific to IBM System Storage N series with Data ONTAP® 7G.



Solution overview

This section provides an overview of the IBM N series with NearStore feature and Symantec SnapVault for NetBackup solution.

At the highest level:

Deep NetBackup integration with IBM N series with NearStore secondary storage system and deduplication technology

Referred to as “context-based block sharing.”

This solution integrates the IBM N series with NearStore secondary storage system as an optimized backup repository for heterogeneous (not IBM N series) primary storage. Its value is based on the assumption that a file in the same dataset and path but in different backups is likely to have a lot of blocks in common. Integration features include:

Integrates IBM System Storage N series with NearStore disk storage unit (DSU).

Tightly coupled integration with NetBackup, *initially* delivered in NetBackup 6.0 and Data ONTAP 7.1, and then moving onto 7.2, etc.

Supports all NetBackup UNIX[®] and Windows[®] clients.

A modified network protocol has been created to increase data transfer performance between the NetBackup media server and the IBM N series. The IP-based protocol allows synchronous data transfers to occur without the delays associated with interlocked packet acknowledgement that comes with NFS and CIFS.

Creates an IBM N series with WAFL[®] (Write Anywhere File Layout) file system structure from backup data stream.

Reduced storage with redundant data elimination: Backups written to an IBM N series storage unit utilize less disk space when compared to traditional DSUs. After an initial client backup is performed, the WAFL file system saves only changed blocks when subsequent backups are performed for the same client, providing single-instance storage.

To NetBackup, the backup on the IBM N series system looks like a standard NetBackup TAR image backup, allowing most normal NetBackup operations to be performed (duplication, synthetics, vaulting, etc.).

To end users, the backup on the IBM N series system looks like a standard WAFL file system, accessible through NFS and CIFS.

The NetBackup master server schedules, manages, and catalogs open systems client backups as usual. But media servers send data to a new NearStore DSU, where it is deduplicated by removing redundant 4KB blocks of data from the client backup image. This results in substantial disk-to-disk backup space savings because it eliminates the need to store multiple copies of the same backed-up data. The overall benefit is that many more backups can be performed and maintained online for rapid recovery and a more frequent recovery point objective (RPO).

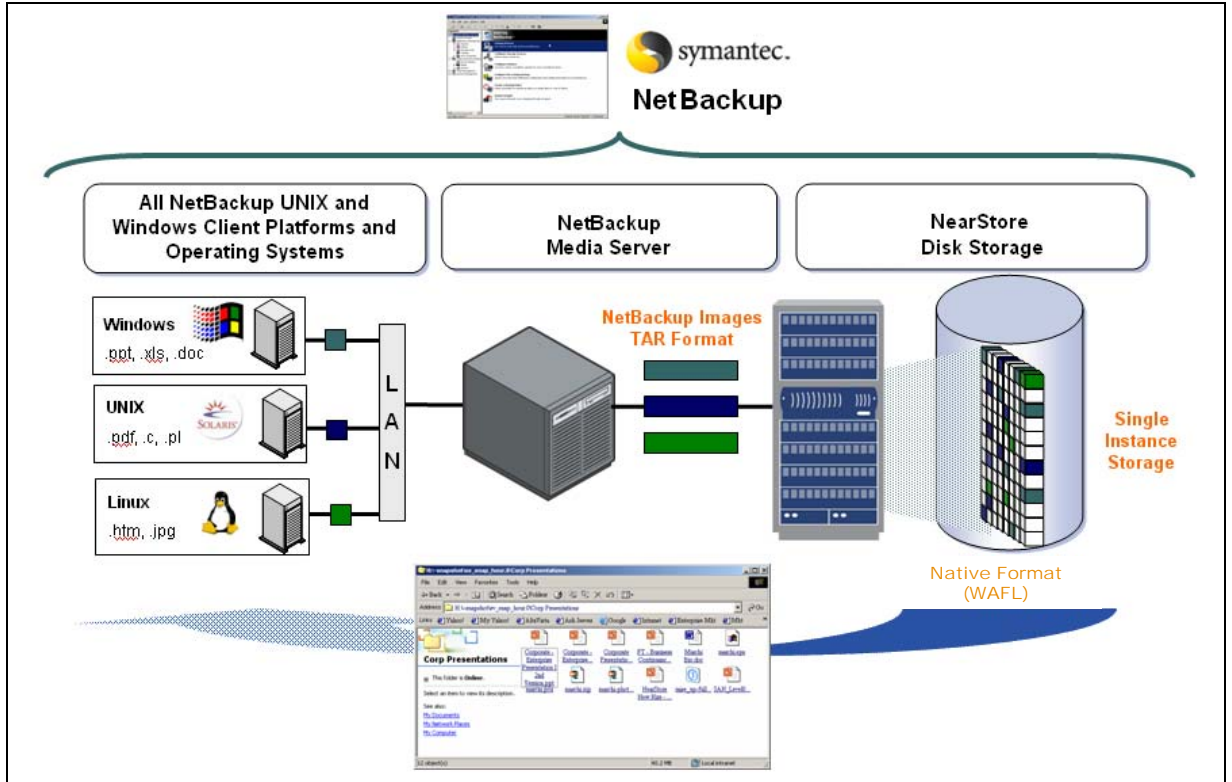


Figure 3-1) SV-NBU space-optimized heterogeneous backups.

Figure 3-1 shows how open systems clients' data is sent from the media server to the IBM N series with NearStore feature storage system. When it is stored on NearStore, individual files are unpacked into standard WAFL file system format and redundant data is eliminated. The backup images are thus also exportable to clients via NFS/CIFS, allowing users to perform non-NetBackup restores.

SV-NBU single instance store (SIS) processing takes place on the IBM N series. This means that all of the data (files) is sent from the client to the NetBackup media server and then to the IBM N series, but only changed blocks are saved by the IBM N series.

SIS is performed on a per client-policy basis.

NetBackup images are converted to WAFL.

Redundant data elimination is performed on-the-fly during the backup, serially comparing blocks 1, 2, ... of the new version of the file with previous versions.

If the block has already been saved, then the new block is eliminated and linked to the block from a previous backup. Only additional blocks take extra space.

This means that although a NetBackup incremental backup sends changed files to the IBM N series, only the changed blocks are stored on disk. In addition, if another full backup is done, although the entire backup is sent to the IBM N series, only the changed blocks are saved.

This also applies to synthetic full backups. The synthetic is generated by NetBackup, but then IBM N series removes duplicate blocks.

For more detailed discussions on data movement, WAFL file system creation, and single-instance storage, see the "Concepts and How It Works" section later in this document.

File system export

File System Export (FSE) is the ability for the WAFL file systems that were created from the backups to be shared via NFS and/or CIFS so that users and administrators can browse the backup images (as snapshot copies) and do restores outside of NetBackup (for example, using Windows Explorer). This ability is turned on by a new check box, “Enable file system export,” on the storage unit GUI. The FSE feature became available in Data ONTAP 7.2.

Space-optimized image

Space-optimized image (SOI) provides the ability for the backup image to be aligned on WAFL 4K boundaries but not deduplicated real-time and not exported. This allows the backup to run faster up front and defers the removal of redundant data to background processing later (through processing such as ASIS). The SOI feature became available in Data ONTAP 7.2. In future versions, SOI will be the primary mechanism for protecting application data with SV-NBU.

Concepts and How It Works

This section presents technical information about the solution and how it works. The content is not strictly necessary for installing and deploying these solutions, but it is highly recommended reading.

Dense volumes

Despite the introduction of less expensive ATA disk drives, one challenge for disk-based backup today continues to be the storage cost. There is a desire to reduce storage consumption (and therefore storage cost per MB) by eliminating duplicated data through sharing across files.

The core technology to accomplish this goal is the dense volume, a volume that contains shared data blocks. The IBM N series with Data ONTAP file system, WAFL, is a file system structure that supports shared blocks in order to optimize storage space consumption. Basically, within one file system tree there is the ability for multiple references to the same data block.

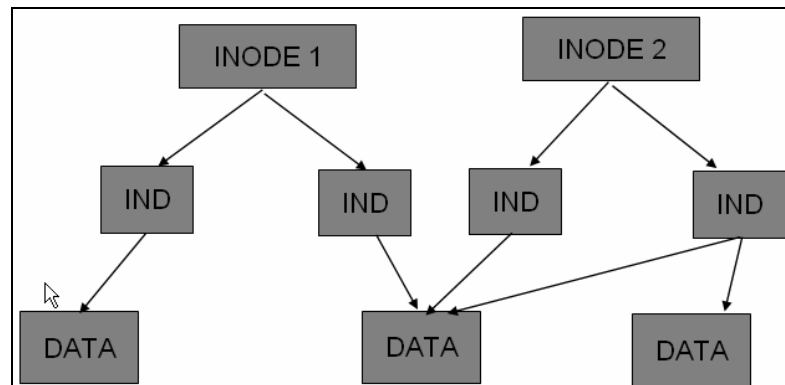


Figure 4-1) Dense volumes.

In SV-NBU, this concept is utilized to allow duplicate 4KB blocks for subsequent backups of the same file to be deleted (for example, INODE 1 might represent last week’s full backup of a file, and INODE 2 would represent this week’s full backup of the same file).



SV-NBU

In NetBackup 6.0, SV-NBU storage units have the ability to store data as clearfile, image, or space-optimized image (SOI).

Clearfile is the preferred method to protect file-system data because it exposes the file structure of the TAR data stream and thereby enables user-invoked drag-and-drop restores.

Image sends the standard NetBackup TAR file image to be stored in its entirety, but without exposing the contents via the WAFL file system. This is equivalent to a NetBackup basic disk backup.

Space-optimized image will be used in the future to protect application-generated data, such as databases.

Media server to IBM Storage System N series protocol

With NetBackup 6.0 and Data ONTAP 7.1, a new proprietary protocol was introduced for efficient streaming of data communication between the NetBackup media server and the IBM N series.

The media server sets up two TCP connections to the IBM N series:

TAR data stream

Standard NetBackup TAR image: header, file, header, file, etc.

Metadata stream describing the TAR contents

Set of RPCs

Directory contents

File location and length

UID/GID, permissions

mtime, ctime, atime.



Processing of data from media server

This section describes the processing that occurs on the IBM N series as and after it receives the two streams.

TAR stream is processed on the fly as it is received

- TAR stream is written to disk (to allow NetBackup-initiated restores)

- Each TAR header and file data is aligned on a 4KB block boundary

If FSE was selected, context-based block sharing is also performed on the fly

- Each file is unpacked into the WAFL Active File System (AFS)

- Unpacked AFS file data is accessible via NFS mounts or CIFS shares

- Removes redundant blocks between the TAR images and AFS data

- Removes redundant blocks from multiple versions of a backup set

Archival snapshot is taken of the SV-NBU volume

- Preserves the TAR image and point-in-time state of AFS data

[Optional] ASIS/context-independent block sharing scan is initiated

- Occurs if ASIS is enabled on the volume

- Removes common blocks from independent unpacked file data

- See TR-3505, *IBM Storage System N Series Advanced Single Instance Storage (ASIS)*

- Deployment and Implementation Guide*, for details

IBM System Storage N series with Snapshot™ management removes redundant snapshots

- Archival snapshot of the ASIS-enabled volume is taken, which replaces the original SV-NBU archival snapshot

Creation of directories and files in the AFS

Backups in clearfile mode have data aligned to a 4KB boundary for unpacking and deduplication in the active file system. Once the data arrives at the IBM N series, there are three threads that handle the aligning and unpacking:

directory creator: Makes sure that every directory is created

file creator: Makes sure that every file is created (0 length files)

data writer: Writes the data.



Quick start: Installing and configuring

This section walks you through the steps that you need to accomplish to make the solution work. Although it discusses some basic setup, in general this document assumes that both the IBM N series with NearStore storage systems and NetBackup are already installed and running.

Requirements overview

The following table specifies the software and hardware required for each solution. Version numbers listed are the minimum required.

IBM N series Hardware	N Series with NearStore feature N5200, N5500 or higher
Data ONTAP	Data ONTAP 7.2 or later
IBM N series Software	sv_ontap_sec license nearstore_option (for N5x00 or later) license
NetBackup	NetBackup Enterprise 6.0MP4 or later (NDMP Option—installed, but need not be licensed) Disk Optimization Option
Protocols	N/A
Client	(Whatever UNIX and Windows platform/OS NetBackup supports) NetBackup 5.0 or later
Media Server	(Whatever UNIX and Windows platform/OS NetBackup supports) NetBackup 6.0MP4 or later
Applications	File Services

Table 5-1) Solution requirements overview.

NetBackup

All solutions need to make sure that NetBackup Enterprise Server is installed on the master server and any media servers involved in the solution.

New to NetBackup 6.0 is that before installing the master and media servers, you need to install Infrastructure Core Services (ICS).

The Network Data Management Protocol (NDMP) for NetBackup Option must be installed, but does not require licensing because it is not used for data movement. If you have a Windows master server, then NDMP is already installed. If you are running UNIX, then the NDMP package must be installed separately.

When NetBackup is installed, you manage and configure it via the NetBackup Administration Console. If you're running the Java™ version, first you need to log in to the master server (Figure 5.1); this can be accomplished from the master server, media server, or, usually, an administrator's client desktop.

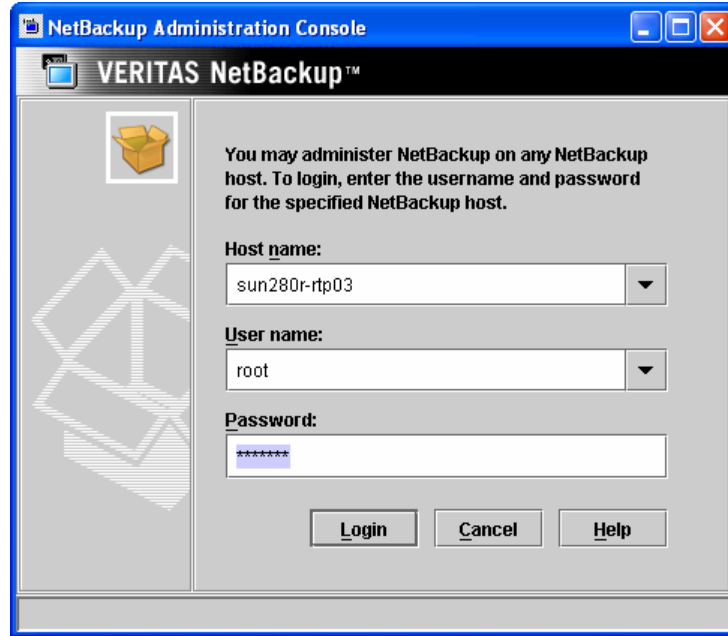


Figure 5-1) NetBackup login.

When you've logged in, you see the NetBackup Administration Console.

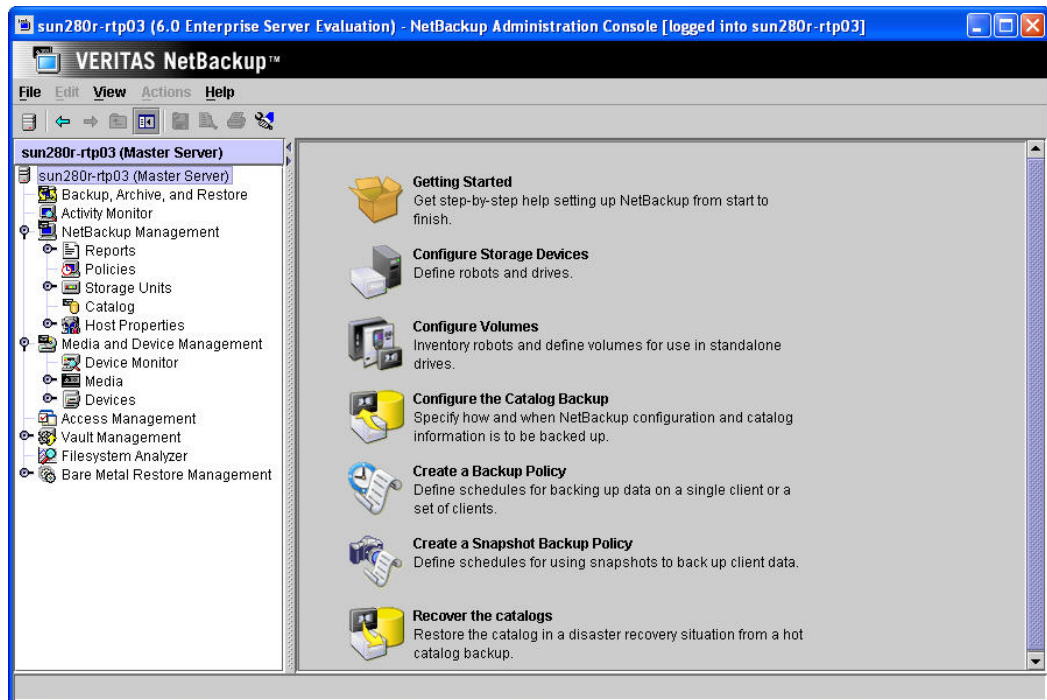


Figure 5-2) NetBackup Administration Console (Java).

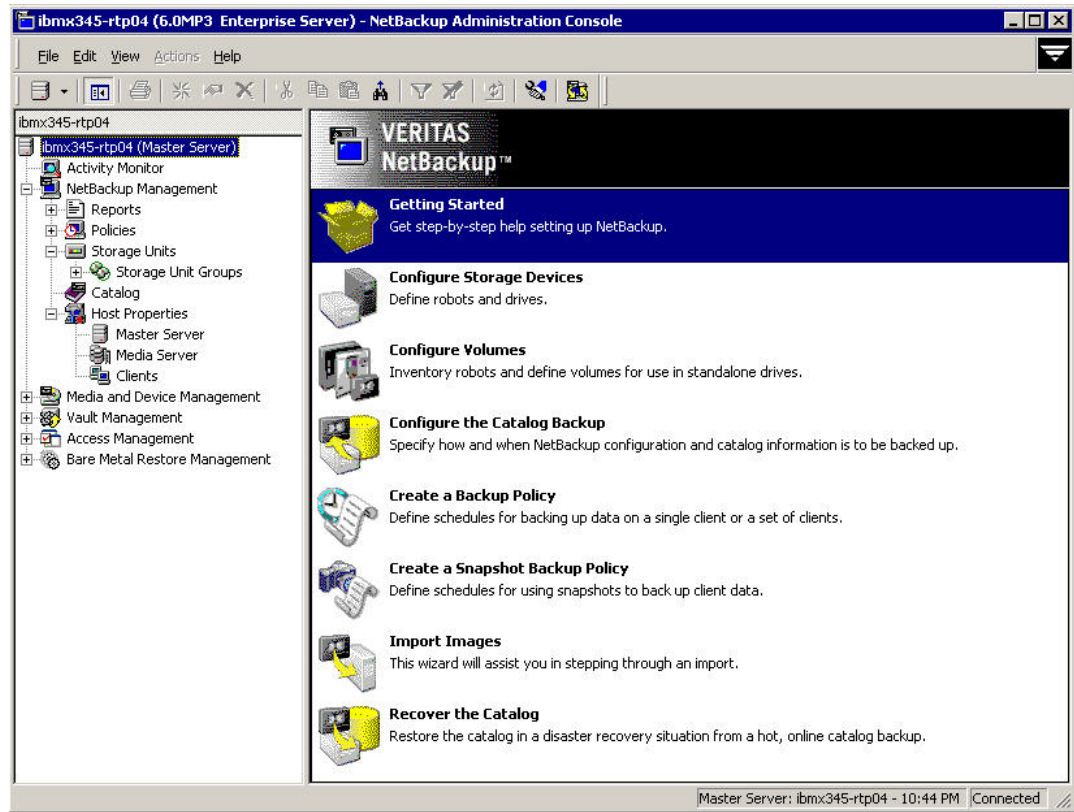


Figure 5-3) NetBackup Administration Console (Windows).

The NetBackup Administration Console is typically used for configuring, managing, and reporting of the NetBackup environment. The Windows version is used in the remainder of this document.

Data ONTAP

NDMP needs to be enabled for SV-NBU.

A SnapVault secondary license is required for SV-NBU.

Licensing

See the table in section 5.1 for required software, which must be appropriately licensed.

NDMP authentication

NDMP is required for authentication for SV-NBU. If NDMP is not set up correctly, SV-NBU will not work, so this section discusses installing and configuring it in some detail.

If you have a Windows master server, then NDMP is already installed. If you are running UNIX, then the NDMP package must be installed separately.

Enabling NDMP on IBM N series

Before NetBackup can communicate with IBM N series as part of the joint disk-to-disk backup solution, NDMP needs to be turned on and configured on one or more IBM N series devices. For SV-NBU, this means the IBM N series NearStore secondary storage system.

To determine whether NDMP is running on the IBM N series, use the following command:

```
ndmpd status
```

If NDMP is not running, start it by using the `ndmpd on` command or by using IBM System Storage N series with FilerView® (a web-based management tool) to enable it.

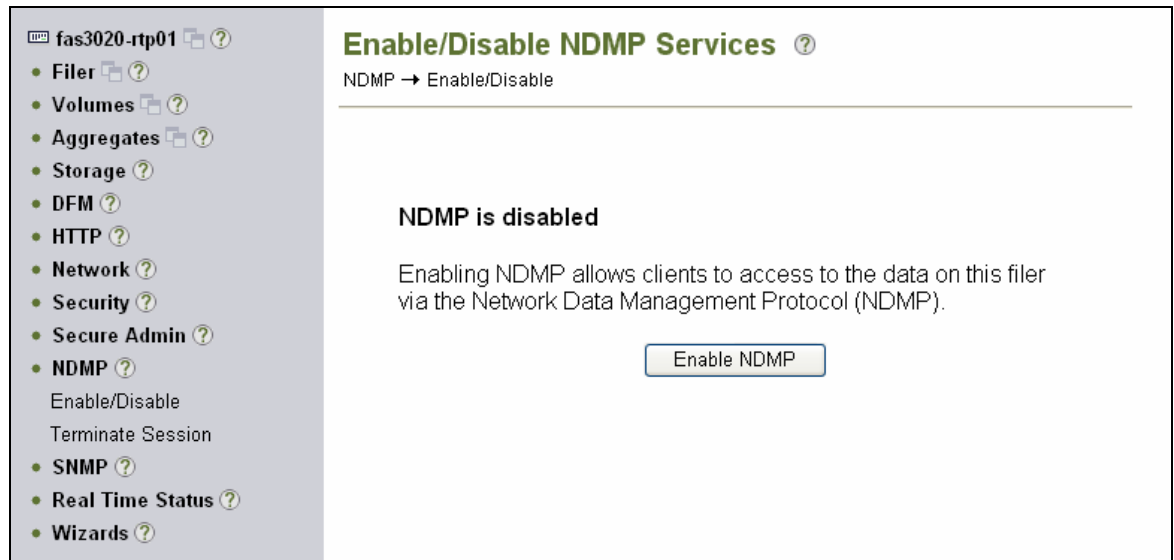


Figure 5-4) Turning on NDMP using FilerView.

This paper uses the root user for NDMP authentication. In production environments, that might pose a security risk, so use the following commands to specify additional NDMP credential information:

```
useradmin user add
ndmpd password [username].
```

NetBackup 6.0 NDMP Graphical User Interface

New to NetBackup 6.0 is the ability to configure NDMP authentication via a graphical user interface (GUI). This section walks through setting up those credentials.

To add an NDMP host (i.e., an IBM N series), select Media and Device Management > Devices > NDMP Hosts in the navigation pane. Then select Actions > New > New NDMP Host...

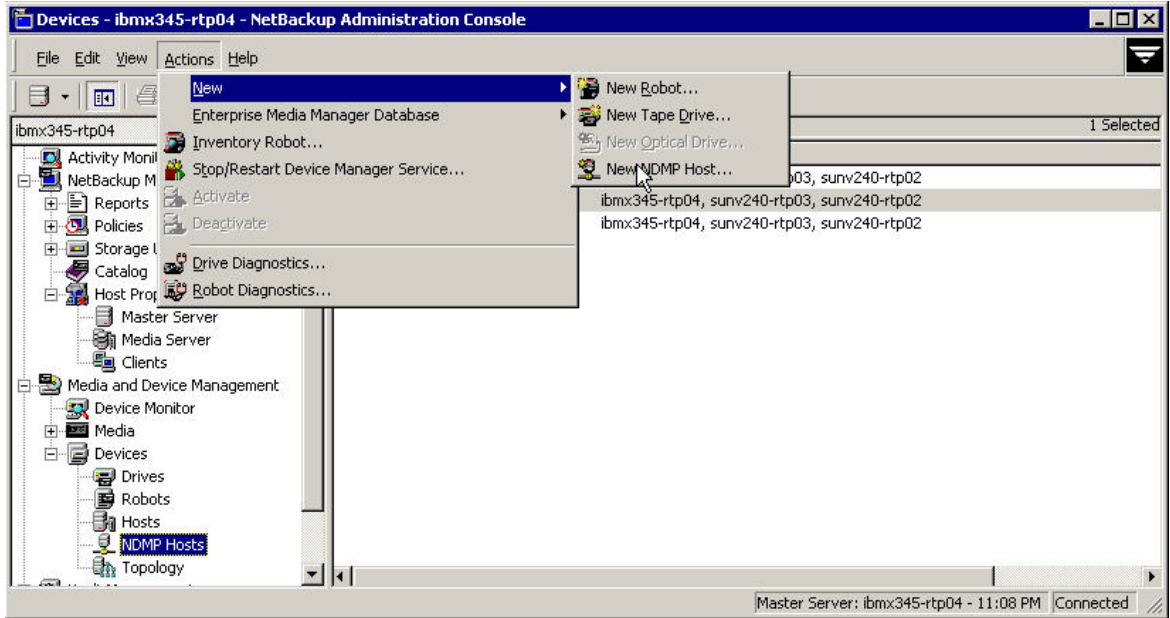


Figure 5-5) Adding a new NDMP host to NetBackup.

Enter the name of the IBM N series as the NDMP host name.




Figure 5-6) Specifying the name of the new NDMP host in NetBackup.

When you click OK a window opens in which you set credentials for the NDMP host. There are several options, but for a specific NDMP-authenticated host, typically one of two is used.

The first option is to set specific NDMP authentication attributes for this newly added NDMP host and all the existing media servers.

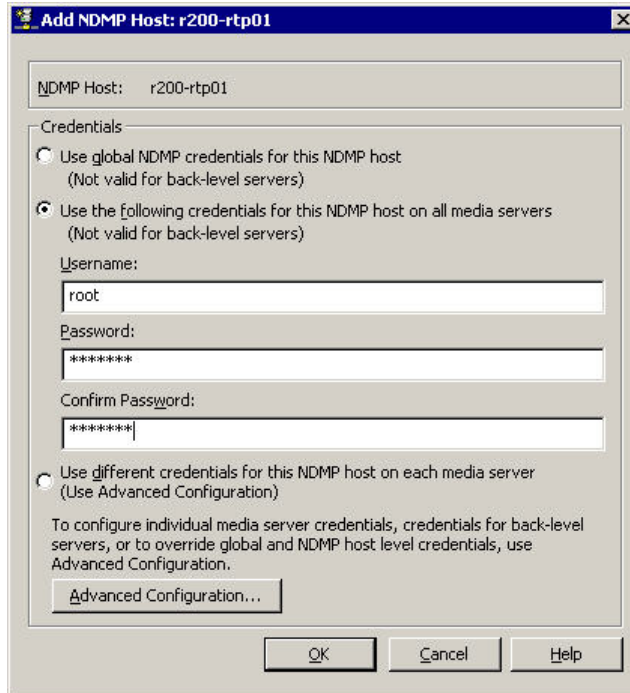


Figure 5-7) Specifying NDMP host credentials in NetBackup.

The second option is to have the NDMP host use global NDMP authentication. Although global authentication is easier to configure and manage, it is also somewhat less secure.

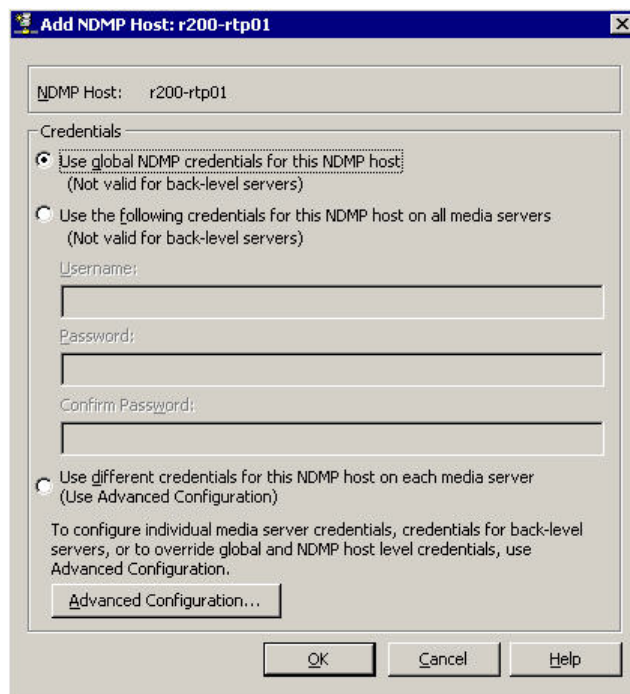


Figure 5-8) Specifying NDMP host to use global credentials in NetBackup.

To set the global NDMP authentication, select NetBackup Management > Host Properties > Master Server in the navigation pane, highlight the master server name, and then select Actions > Properties...

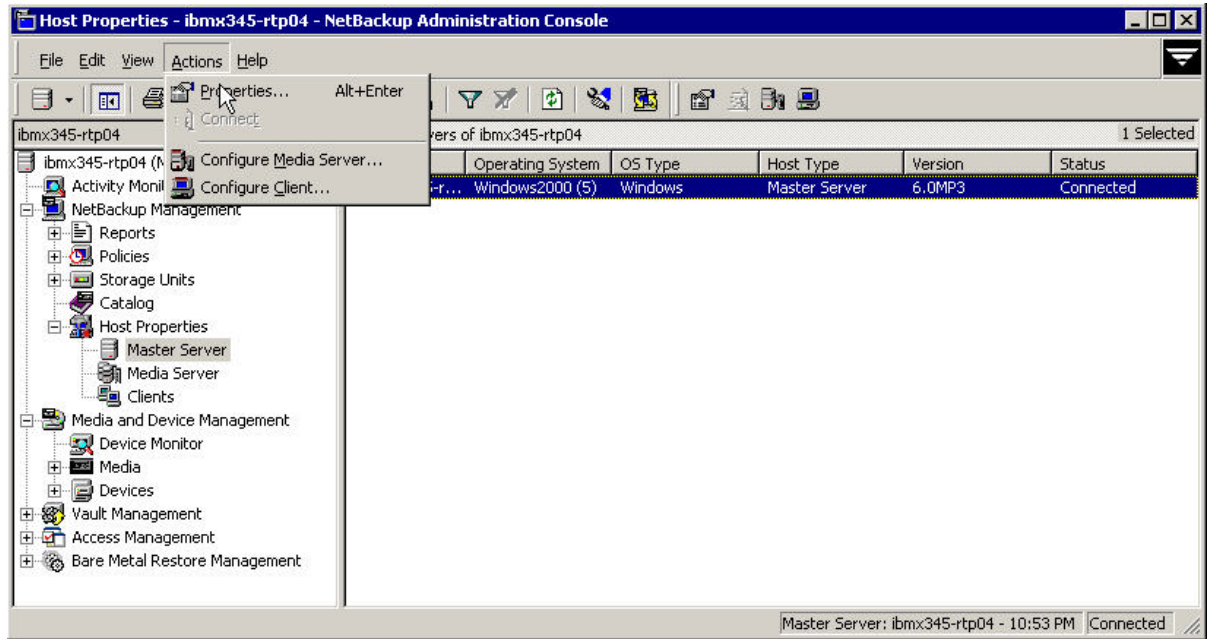


Figure 5-9) Accessing NetBackup master server host properties.

Select the NDMP property for the master server and enter the global user name and password. (As mentioned earlier, this paper uses the *root* user for NDMP authentication; in production environments, that could be considered a security risk.)

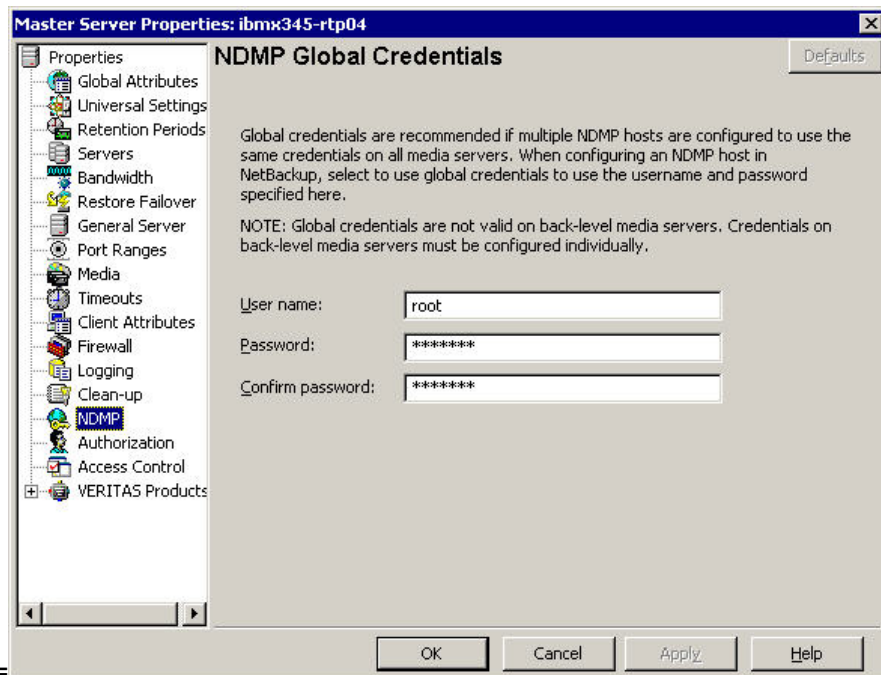


Figure 5-10) Specifying NDMP global credentials.

When you are finished, the configured NDMP hosts should look similar to the following figure.

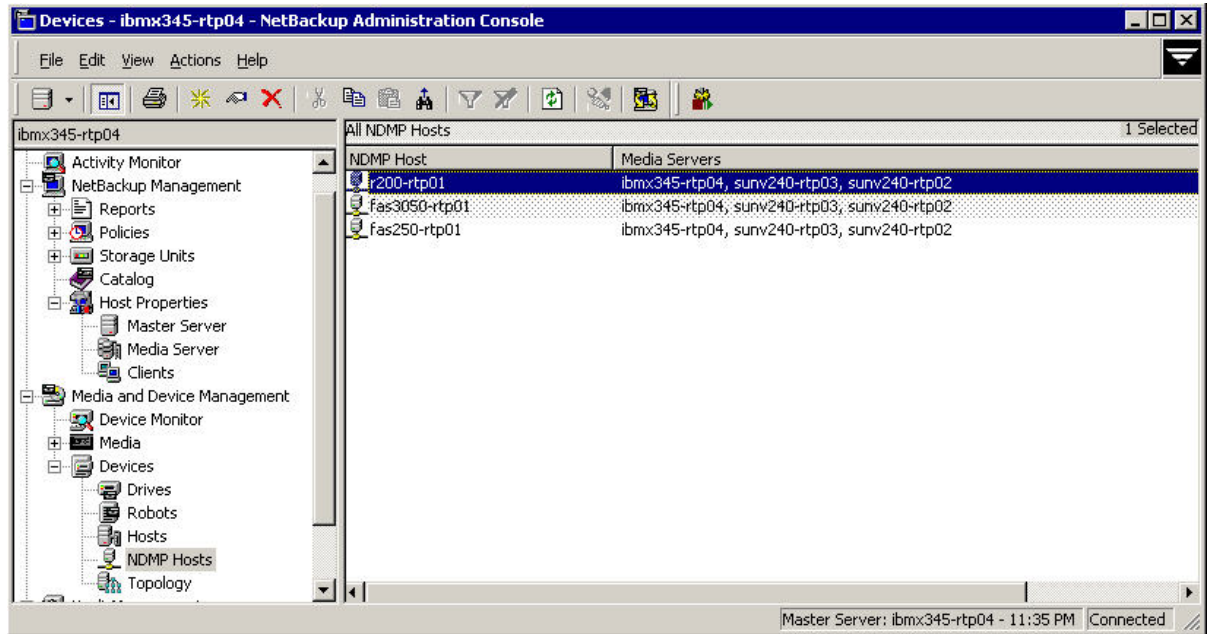


Figure 5-11) NDMP hosts authenticated with master and media servers.



NetBackup 6.0 NDMP Command Line Interface

You can use a command line interface (CLI) to configure NetBackup NDMP authentication. `tpconfig` is the utility used to create, manage, and verify the NDMP authentication relationships. The following example uses `tpconfig` on UNIX to configure the NDMP username and password.

```
NDMP Host Credentials Configuration
1) Add
2) Delete
3) Update
4) List Configuration
5) Configure Default Authentication Credentials
6) Verify Host's Authentication Credentials
7) Help
8) Quit - return to previous menu
```

```
Enter option: 3
Enter NDMP host name: r200-rtp01
Enter a User ID: root
Enter the NDMP host's password for User Id root:
Re-enter the NDMP host's password to confirm it:
```

```
This user ID and Password are for:
1. Just this Media Server
2. All Media Servers on this NDMP host
Please enter a number, 1 or 2: 2
```

```
NDMP host r200-rtp01 is SnapVault/NearStore capable.
```

Press any key to continue

The following example shows using `tpconfig` on Windows to configure the same NDMP username and password as in the previous example:

```
tpconfig -add -nh r200-rtp01 -user_id root -password <password>
```

For the old-school NetBackup administrator (NetBackup 5.1 and earlier releases), there was no NetBackup GUI for managing NDMP, and a different CLI was used to configure NetBackup NDMP credentials. `set_ndmp_attr` is a legacy alternative way to configure NetBackup NDMP.

```
cd <netbackup directory>/volmgr/bin
set_ndmp_attr -auth ndmp-server-host username
```

The following example shows using `tpconfig` on UNIX to verify that NDMP credentials are configured correctly. (`tpconfig` on Windows cannot be used to verify NDMP credentials.)

```
sunv240-rtp02# which tpconfig
/usr/opensv/volmgr/bin/tpconfig
sunv240-rtp02# tpconfig
```



Device Management Configuration Utility

- 1) Drive Configuration
- 2) Robot Configuration
- 3) NDMP Host Credentials Configuration
- 4) Print Configuration
- 5) Help
- 6) Quit

Enter option: 3

NDMP Host Credentials Configuration

- 1) Add
- 2) Delete
- 3) Update
- 4) List Configuration
- 5) Configure Default Authentication Credentials
- 6) Verify Host's Authentication Credentials
- 7) Help
- 8) Quit - return to previous menu

Enter option: 4

```
=====
Media Server:                sunv240-rtp02
NDMP Host Name:              r200-rtp01
User Id:                     root <Default>
=====
```

Press any key to continue

NDMP Host Credentials Configuration

- 1) Add
- 2) Delete
- 3) Update
- 4) List Configuration
- 5) Configure Default Authentication Credentials
- 6) Verify Host's Authentication Credentials
- 7) Help
- 8) Quit - return to previous menu

Enter option: 6

Enter NDMP host name: r200-rtp01



```
Connecting to host "r200-rtp01" as user "root"...
Waiting for connect notification message...
Opening session--attempting with NDMP protocol version 4...
Opening session--successful with NDMP protocol version 4
  host supports MD5 authentication
Getting MD5 challenge from host...
Logging in using MD5 method...
Host info is:

  host name "r200-rtp01"
  os type "DTAP"
  os version "DTAP Release 7.1RC3"
  host id "0050409813"

Login was successful
Host supports LOCAL backup/restore
Host supports 3-way backup/restore
Host has SnapVault Secondary license installed
```

Press any key to continue

Although the menu-driven `tpconfig` example shown above provides the most capabilities, to verify most quickly you can simply use the following command on UNIX or Windows:

```
<installpath>/volmgr/bin/tpautoconf -verify <ndmp host>
```

Here is a sample of running the command:

```
root@sunv240-rtp03 # cd /usr/opensv/volmgr/bin
root@sunv240-rtp03 # ./tpautoconf -verify r200-rtp01
Connecting to host "r200-rtp01" as user "root"...
Waiting for connect notification message...
Opening session--attempting with NDMP protocol version 4...
Opening session--successful with NDMP protocol version 4
  host supports MD5 authentication
Getting MD5 challenge from host...
Logging in using MD5 method...
Host info is:

  host name "r200-rtp01"
  os type "DTAP"
  os version "DTAP Release R7.2xN_060406_0030"
  host id "0050409813"

Login was successful
Host supports LOCAL backup/restore
Host supports 3-way backup/restore
Host has SnapVault Secondary license installed
```



SV-NBU configuration

This section contains steps for getting the SV-NBU solution running. It is assumed that in addition to installing the IBM N series with NearStore secondary storage system and the proper version of Data ONTAP, the following tasks have already been done:

Install and license NetBackup Enterprise Server.

Install and configure NDMP on both the IBM N series with NearStore secondary storage system and via NetBackup. A NetBackup NDMP license is *not* required to create a NearStore storage unit. However, NDMP should be enabled on the IBM N series because this enables the IBM N series NDMP credentials to be entered using the NetBackup Administration Console.

The SV-NBU environment looks like any other NetBackup environment, with the usual clients and media servers existing.

Note in the configuration example used throughout the remainder of this section:

The master server is also the client

Master server = ibmx345-rtp04

NearStore secondary = r200-rtp01

Client data is C:\testdata on the master server

The remainder of this section steps through the following:

Configuring SnapVault on the IBM N series

Making a volume on the IBM N series

Configuring a NetBackup SV-NBU storage unit

Configuring NetBackup SV-NBU policy

Running a NetBackup SV-NBU backup job

Performing a restore from an SV-NBU backup

IBM N series SV-NBU Configuration

Ensure that NDMP authentication has been configured between the NetBackup master server (and any necessary media servers) and the IBM N series. (Note: `ndmpd password` is *only* for authentication; NetBackup NDMP option is *not* required; port 10000 is not used.)

To make the NearStore storage unit available for SV-NBU backups, you add, enable, and configure SnapVault on the IBM N series secondary by executing the following at the secondary command line.

Add the secondary SnapVault license:

```
license add <sv_secondary_license>
```

Enable SnapVault:

```
options snapvault.enable on
```

Grant access to media servers authorized to access the storage system with the following command:

```
options snapvault.access host=nbu_master_server,nbu_media_server1...
```

The following setup results from the previous steps:

```
r200-rtp01> options snapvault
snapvault.access    host=ibmx345-rtp04,sunv240-rtp03
snapvault.enable    on
```

First create a volume on the IBM N series using IBM N series with FilerView. (Note that it must be an IBM System Storage N series with FlexVol™ volume.) When creating the volume, adhere to the following maximum usable volume size guidelines. This size includes any `snap reserve` space. (Refer to the size limitations in the “Operating Characteristics” section of this document for the technical reason for the restrictions.)

Table 5-2) Maximum SV-NBU volume sizes.

N5200	N5500	IBM N series with NearStore
1 TB	2 TB	4 TB

SVNBUvol is the volume created for this example.

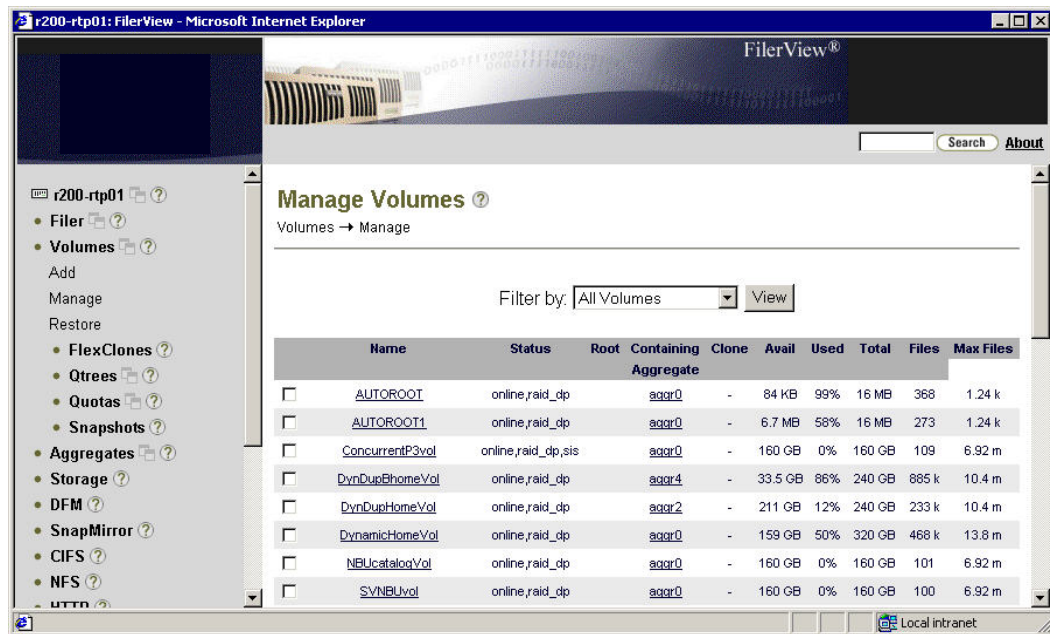


Figure 5-12) Creating an SV-NBU DSU volume via FilerView.

After the volume is configured, turn off the Snapshot schedule for it, since Snapshot management will be handled by SV-NBU. This is not a requirement, but the additional Snapshot copies will reduce storage savings. Use one of the following commands to do this:

```
snap sched <volume> 0 0 0
vol options <volume> nosnap true
```

The IBM N series part of the SV-NBU configuration is now complete.

NetBackup SV-NBU Configuration

First, configure a storage unit that utilizes the volume that was created on the IBM N series. In the navigation pane, right-click Storage Units and select New Storage Unit.

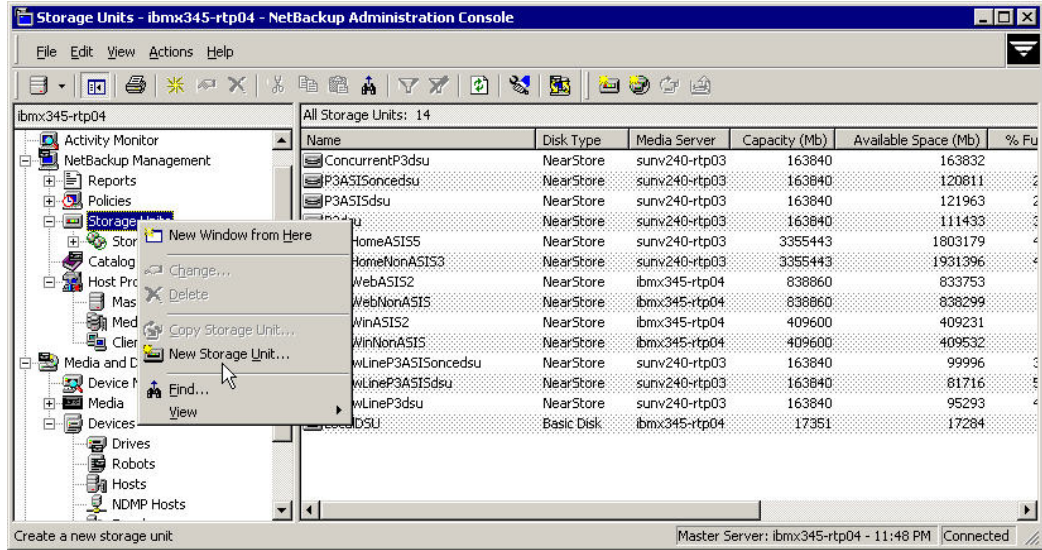


Figure 5-13) Creating a new storage unit.

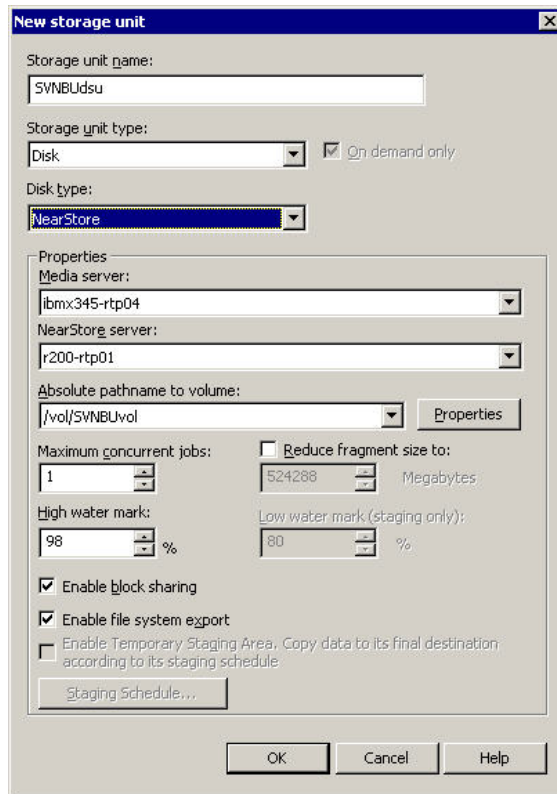


Figure 5-14) Naming and configuring a NearStore storage unit.

First, enter a name for the storage unit. (Note that storage unit names are case sensitive.)

Set the storage unit type to Disk and set the disk type to NearStore.

Select a media server (in this example, this is the master server). The drop-down list contains all the NetBackup media servers in the environment.



Specify the NearStore server as the name of the IBM N series. The drop-down list contains all devices that are NDMP authenticated with the previously selected media server.

Specify the absolute pathname to volume as the volume that was created on the IBM N series. The drop-down list contains all volumes that are on the NearStore server specified in the previous step.

Important: It is critical to understand the functions of two check boxes. To turn on deduplication, you must select both the “Enable block sharing” check box and the “Enable file system export” check box.

Selecting the “Enable file system export” check box has the added benefit of making it possible to share the backups so that users and administrators can browse the backups via NFS or CIFS.

Selecting different combinations of the boxes provides the following functionality:

If neither box is checked, the backup goes into “image mode”; no deduplication is done and the file system is not exported.

If “Enable block sharing” is checked but “Enable file system export” is not checked, this is referred to as SOI (Space Optimized Image) mode. This enables the backup to have deduplication performed by ASIS but does not export the backup as a file system and does not make use of snapshots.

If both boxes are checked, then the backup has deduplication performed, has its content exported as a file system, and makes use of snapshots to ensure the consistency of the export.

“Enable file system export” cannot be checked unless “Enable block sharing” is checked.

The next step is to configure a policy to use the newly created NearStore storage unit. Use the NetBackup Administration Console to create a new policy by selecting Policies and then Actions > New > New Policy.

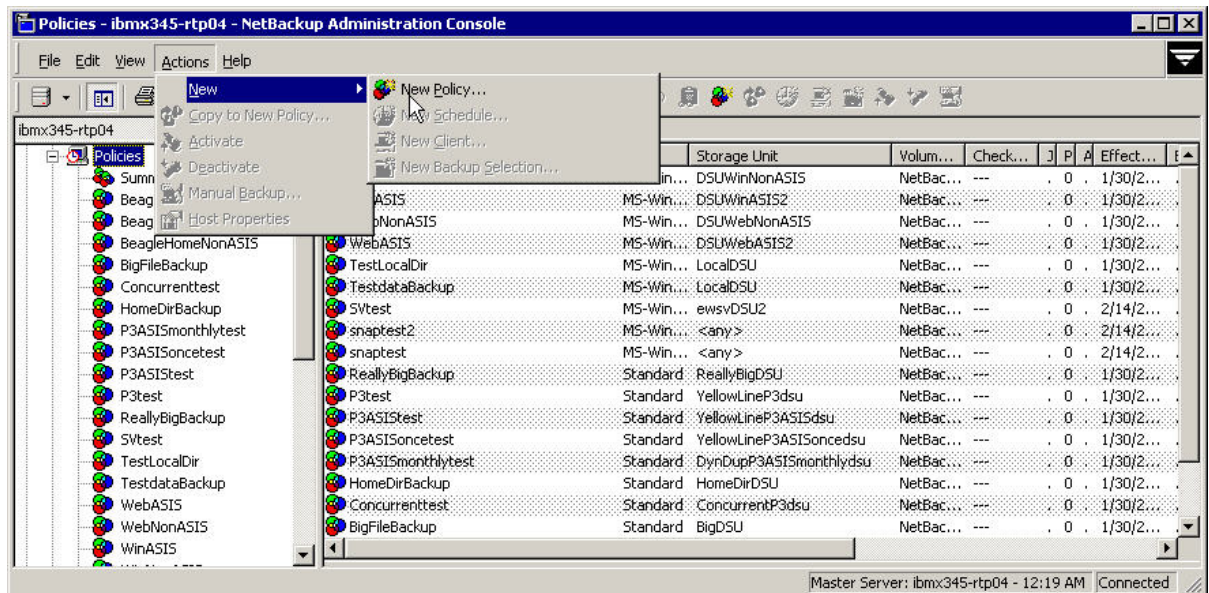


Figure 5-15) Creating NetBackup SV-NBU policy.

Name the policy.




Figure 5-16) Naming the NetBackup SV-NBU policy.

After specifying the new name, you modify the policy attributes.

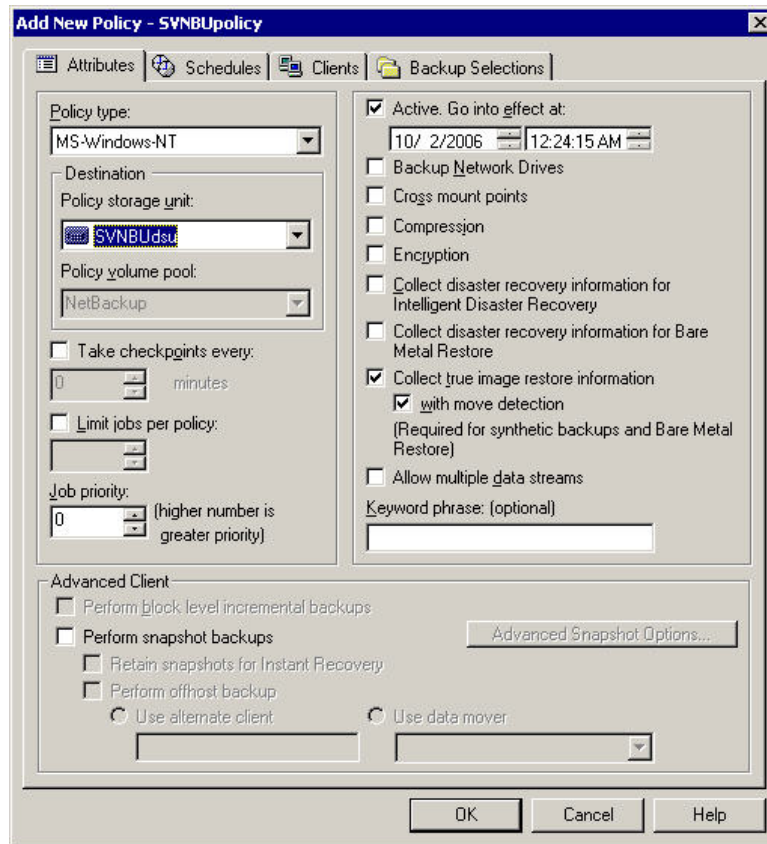


Figure 5-17) Configuring NetBackup SV-NBU policy attributes.

Most of the default policy attributes can be utilized. The exceptions are that you need to:

Specify the NearStore storage unit created earlier as the Policy storage unit.

Check the “Collect true image information” and “with move detection” check boxes so that the snapshot copies and FSE exports/shares are accurate.

Do not select NetBackup Compression. Compressed backups are always sent from the media server to the IBM N series as an *image* data stream instead of a *clearfile* stream. Because of this, the data stream is not aligned, unpacked, and deduplicated. Also, even if FSE is enabled on the storage unit, if NetBackup compressed backups are sent the backup is not exportable and accessible via CIFS/NFS.

Now you must create at least one schedule for the policy. Begin by selecting the Schedules tab for the policy.

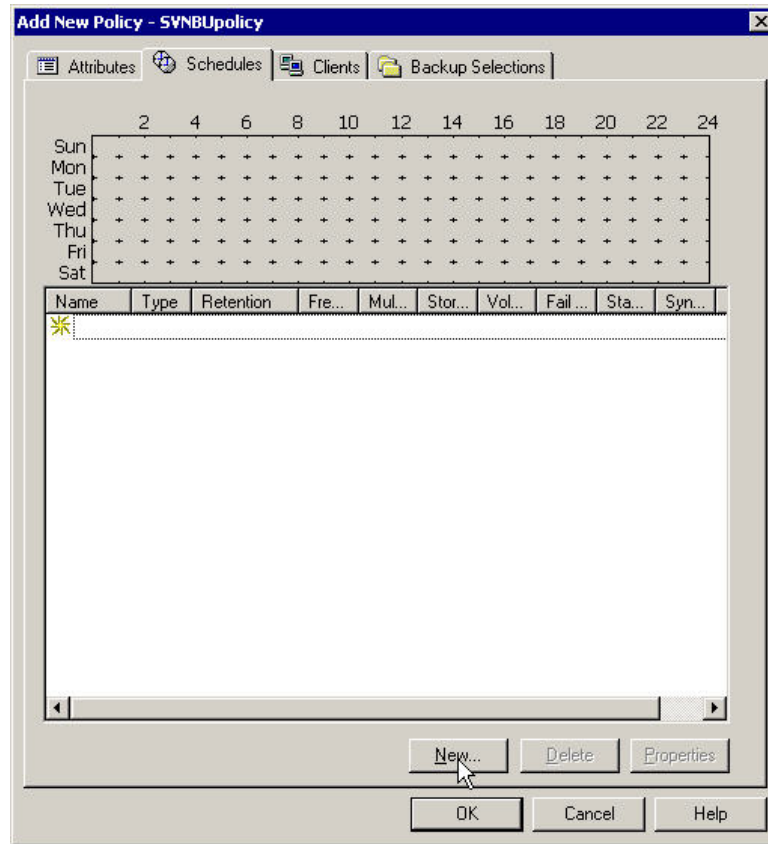


Figure 5-18) Configuring NetBackup SV-NBU policy schedules.

Then click New to create a new policy.

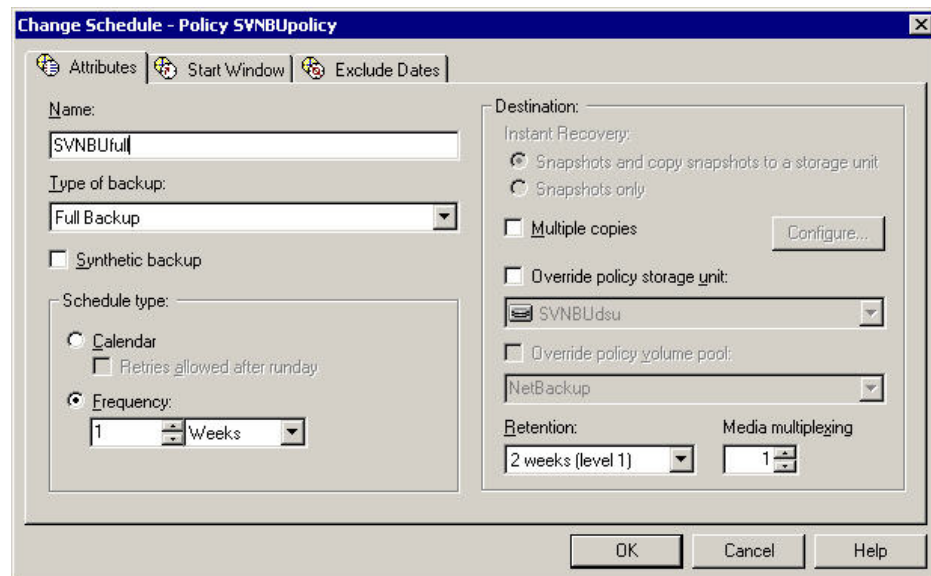


Figure 5-19) Creating a new NetBackup full schedule.

In the following figure, a single full schedule named SVNBUfull has been created. (Note that you must do at least one full backup before any incrementals.) Click OK to return to the policy tabs, where you see a summary of the schedule that was just created.

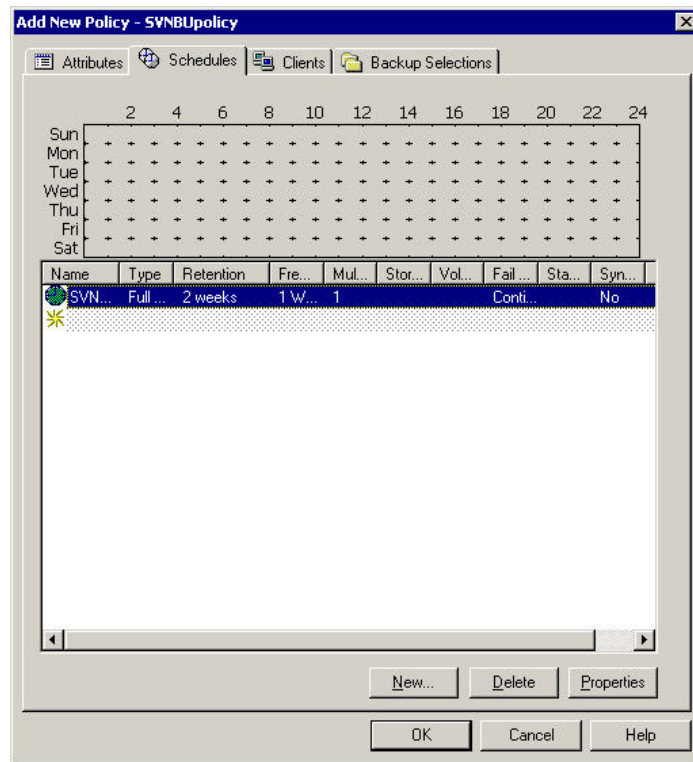


Figure 5-20) NetBackup SV-NBU policy schedules complete.

Click the Clients tab to specify one or more clients to be backed up by this policy.

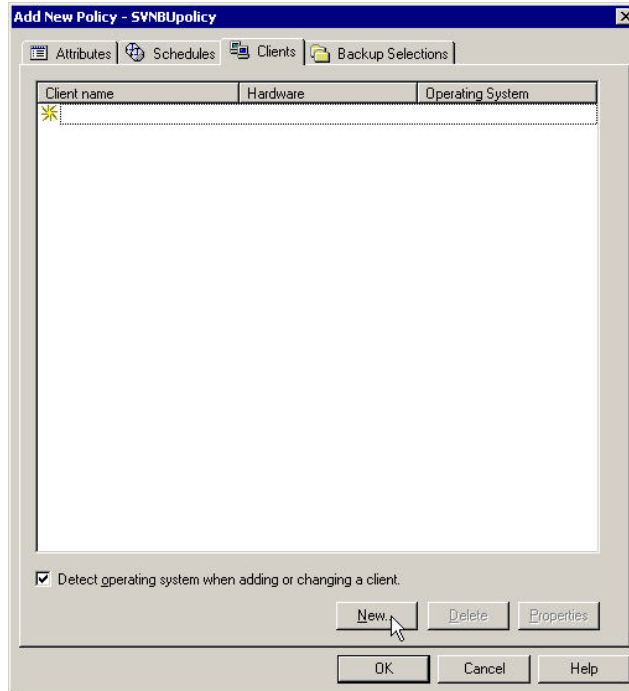


Figure 5-21) Configuring NetBackup SV-NBU policy clients.

Click New and enter the name of the client you want to add.

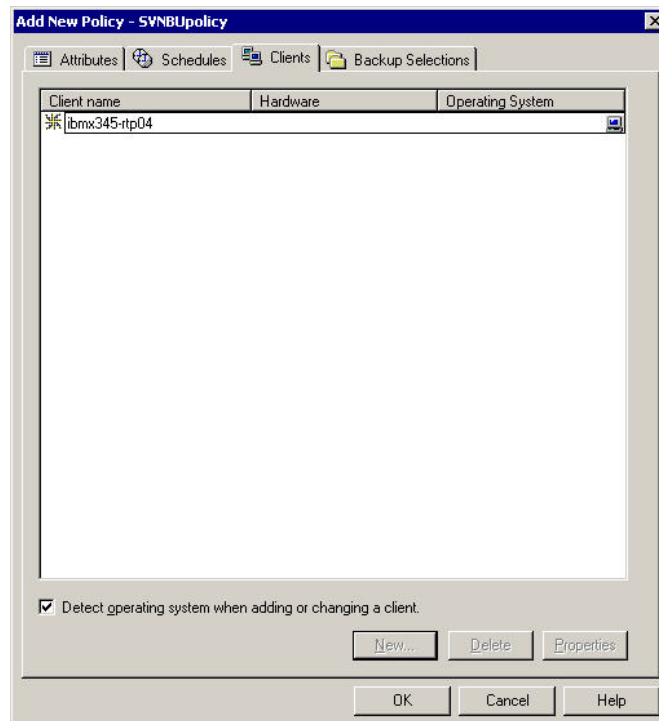


Figure 5-22) Adding a new client to the policy.

Press the Tab or Enter key. The hardware and operating system are automatically detected.

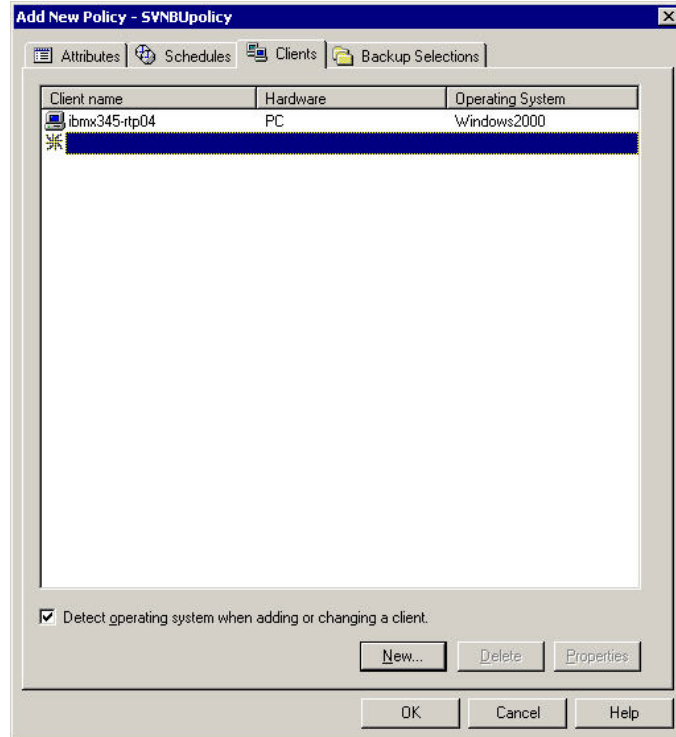


Figure 5-23) NetBackup SV-NBU policy clients complete.

Finally, on the policy's Backup Selections tab, you specify what is to be backed up.

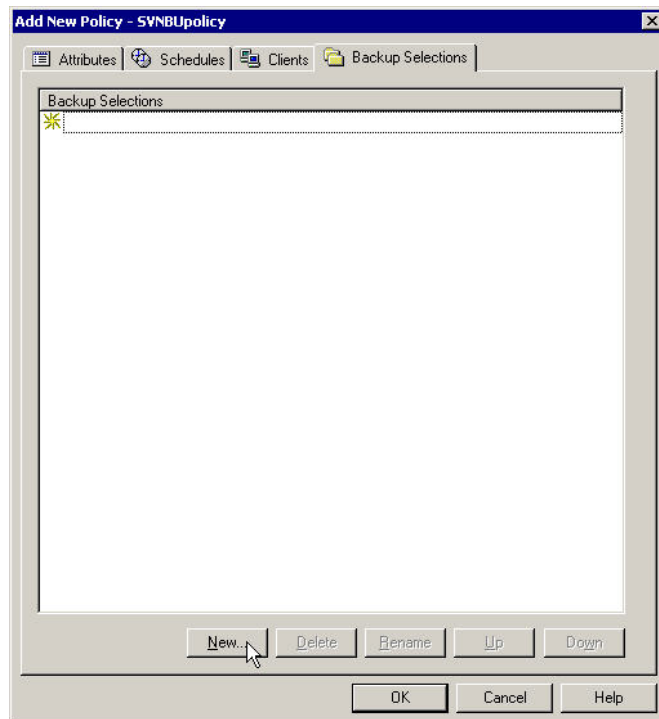


Figure 5-24) Configuring NetBackup SV-NBU policy backup selections.

You can either browse to specify each desired directory or file, or type it in.




Figure 5-25) NetBackup SV-NBU policy backup selections complete.

Click OK. The SV-NBU policy creation is complete!

The SV-NBU backup policy can be run manually or when scheduled to execute. In this document, they are run manually for demonstration purposes.

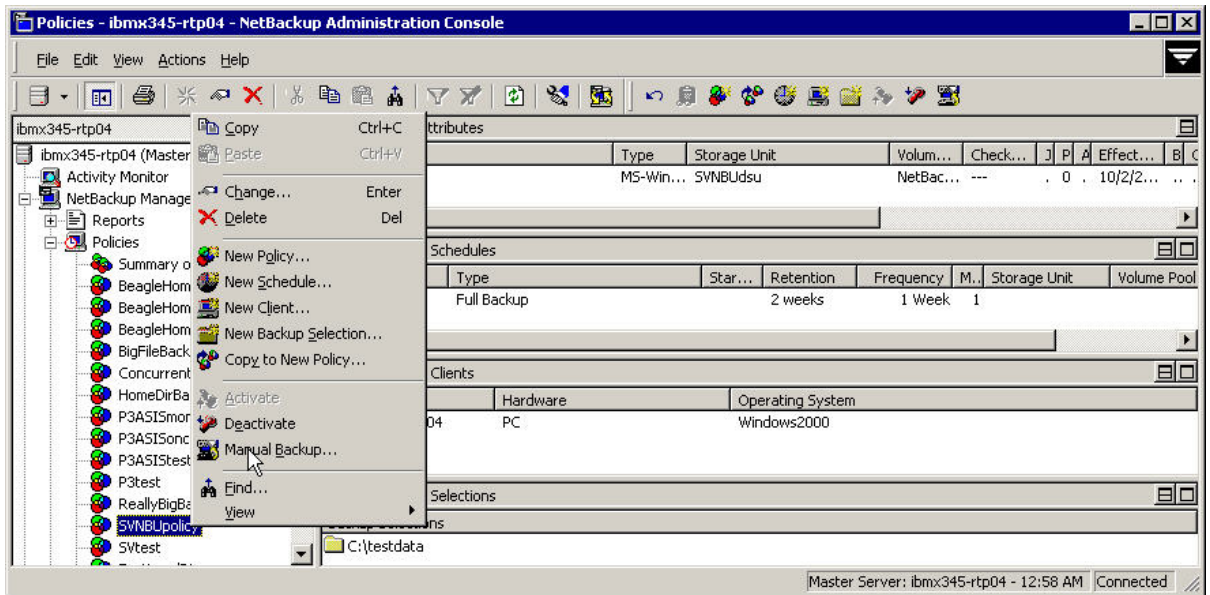


Figure 5-26) Manually running NetBackup SV-NBU policy.

In the Manual Backup window, you can simply click OK, because there is only one schedule for the policy. If there were several schedules, you would select one and then click OK.

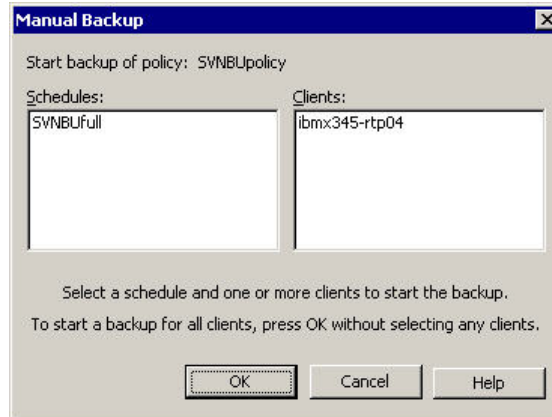


Figure 5-27) Selecting the NetBackup SV-NBU policy schedule to run.

Use the Activity Monitor to watch the progress of the backup job. The following figure shows the job that was just started running as Job ID #1074.

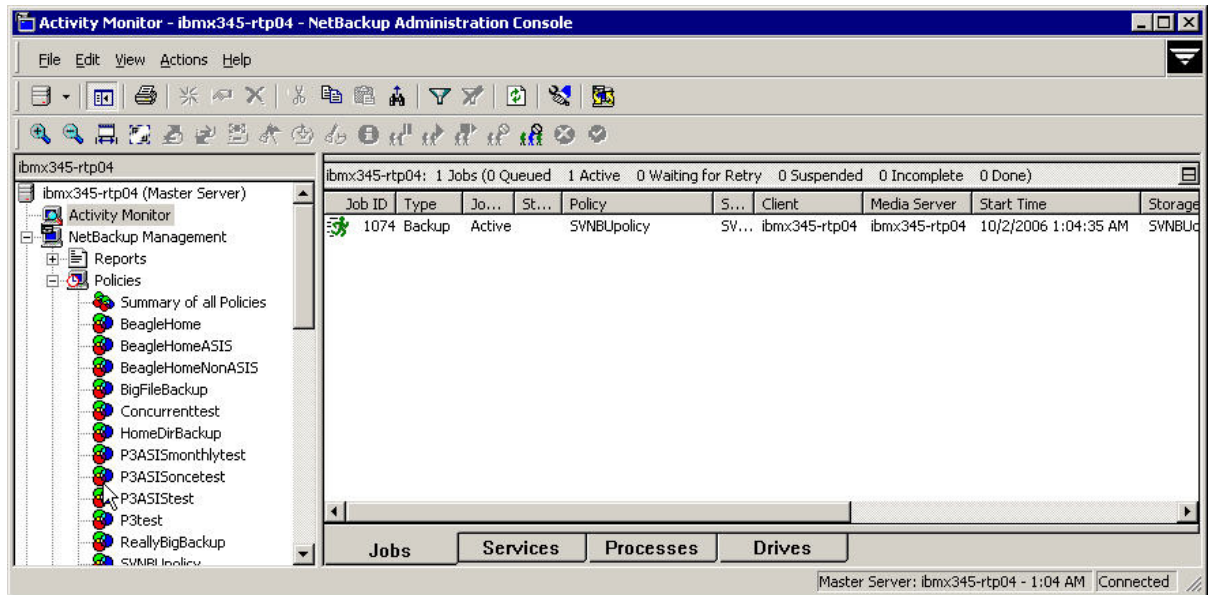


Figure 5-28) NetBackup Activity Monitor – active SV-NBU job.

Double-click the job to see additional information about it, and select the Detailed Status tab to see all the details. In the following example, the job has completed successfully.

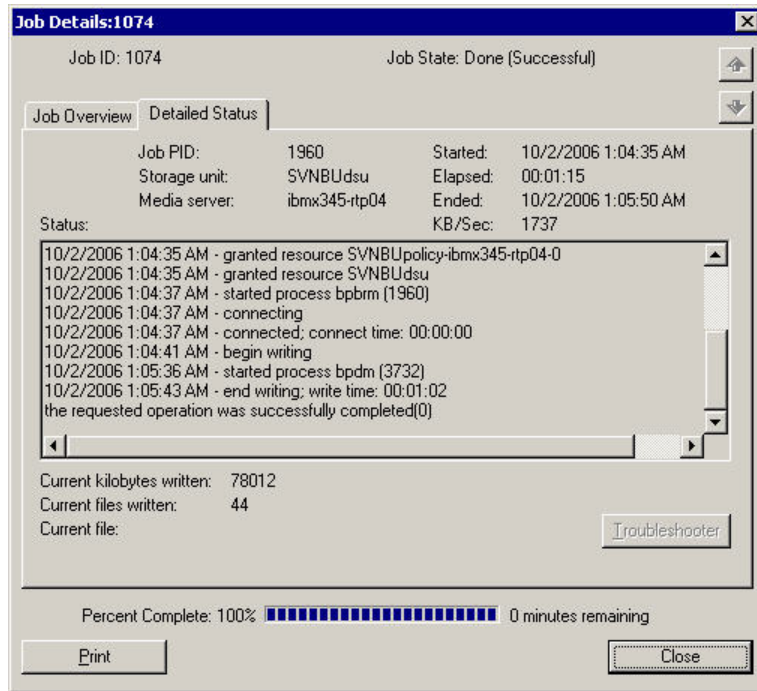


Figure 5-29) NetBackup Activity Monitor – SV-NBU detailed status.

After the first time an SV-NBU backup is accomplished to a volume on the IBM N series, the volume shows the `sis` (Single Instance Storage) attribute; use the `volume status` command to verify this.

```
r200rtp01> vol status SVNBUvol
      Volume State      Status      Options
      SVNBUvol online   raid_dp, flex  nosnap=on, create_ucode=on
                        sis
      Containing aggregate: 'aggr0'vol status
```

You can use the following command to monitor the increasing space savings.

`snapvault status -b`

After the first backup occurs, you will not see any space savings.

```
r200-rtp01> snapvault status -b SVNBUvol
Snapvault secondary is ON.
Volume      actual      used      saved      %saved      ratio
-----
SVNBUvol    76MB       94MB      0MB        0%          1:1
```

High-level descriptions of the fields:

- `actual` — Total data sent from NetBackup and “stored”
- `used` — Total disk space consumed on the volume
- `saved` — Total storage savings on the volume
- `%saved` — Percentage of storage saved
- `ratio` — The effective compression ratio that storage savings is providing

After several SV-NBU backups have been accomplished, you will observe increasing space savings:

```
r200-rtp01> snapvault status -b SVNBUvol
Snapvault secondary is ON.
```

Volume	actual	used	saved	%saved	ratio
-----	-----	-----	-----	-----	-----
SVNBUvol	311MB	134MB	177MB	57%	2.31:1

For the above examples, four backups were performed.

SV-NBU Restores

SV-NBU restores are accomplished using either the standard NetBackup Backup, Archive, and Restore window, or drag-and-drop from the Snapshot copies created via FSE.

NetBackup Backup, Archive, and Restore Window

Restore via the NetBackup Backup, Archive, and Restore (BAR) interface is performed like any other NetBackup Standard or MS-Windows-NT restore. This section presents a quick demonstration of a restoration.

You can start the NetBackup BAR interface from the Administration Console by selecting File > Backup, Archive, and Restore.

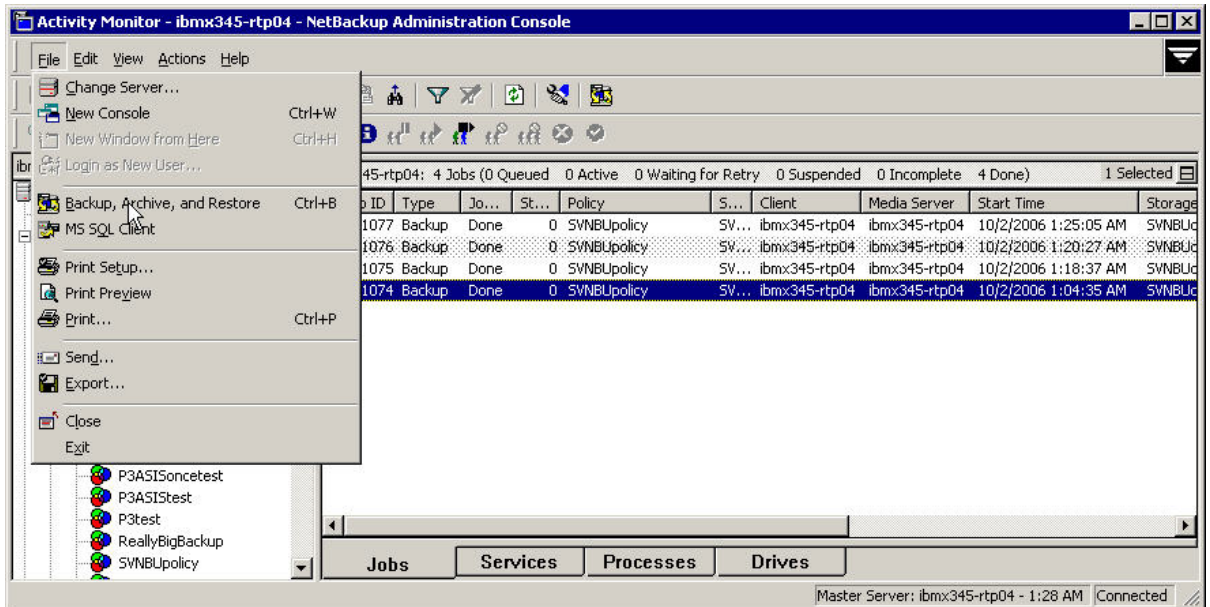


Figure 5-30) Starting the NetBackup Backup, Archive, and Restore (BAR) interface.

If you have a large environment, when BAR starts you may need to specify the client that you want to restore.

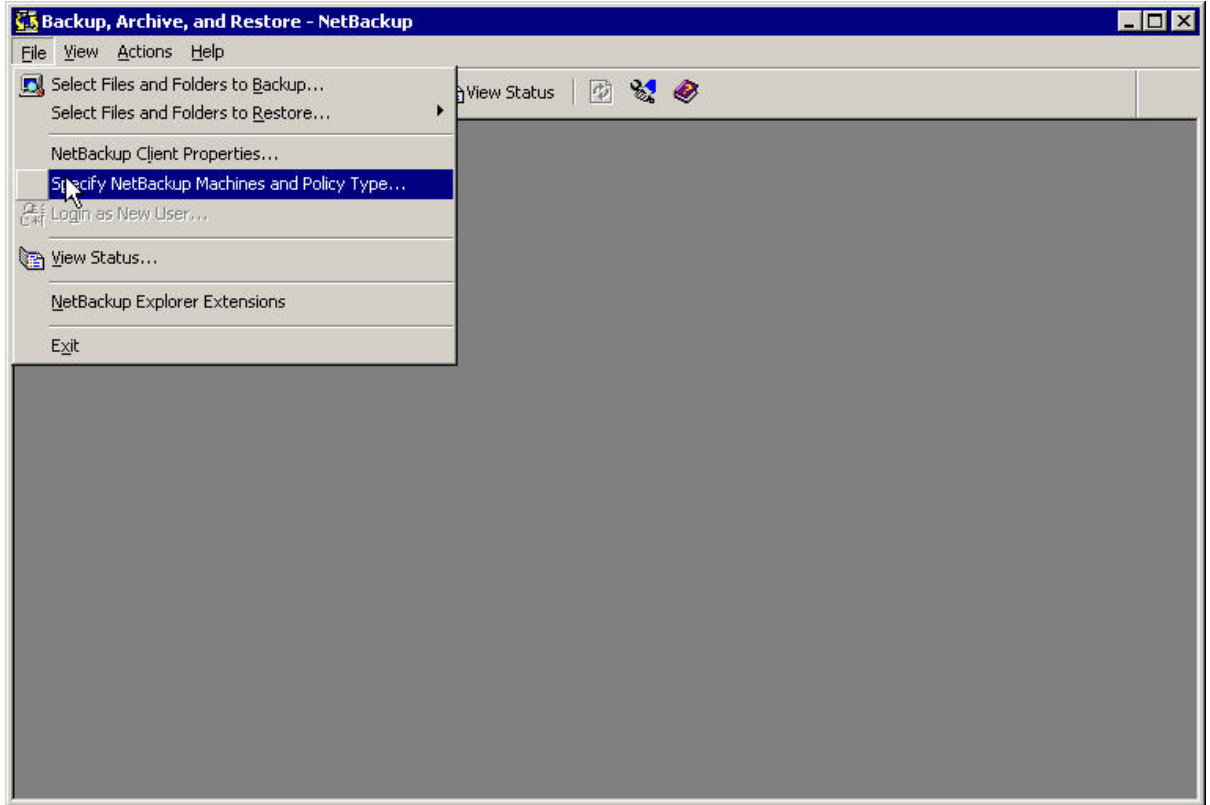


Figure 5-31) Initial NetBackup BAR display.

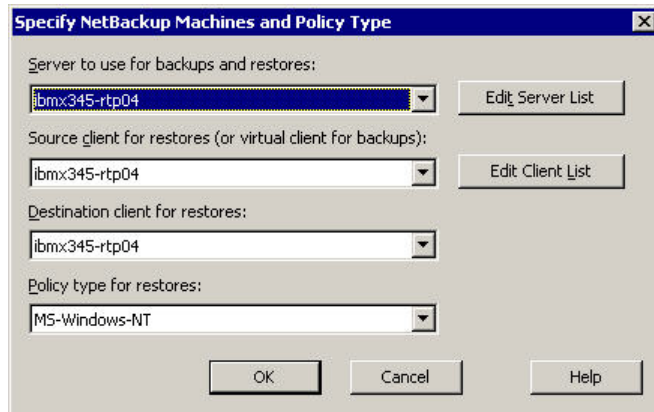


Figure 5-32) NetBackup BAR – specifying machines and policy type.

After selecting the desired machine and policy information, click OK to display the existing backups.

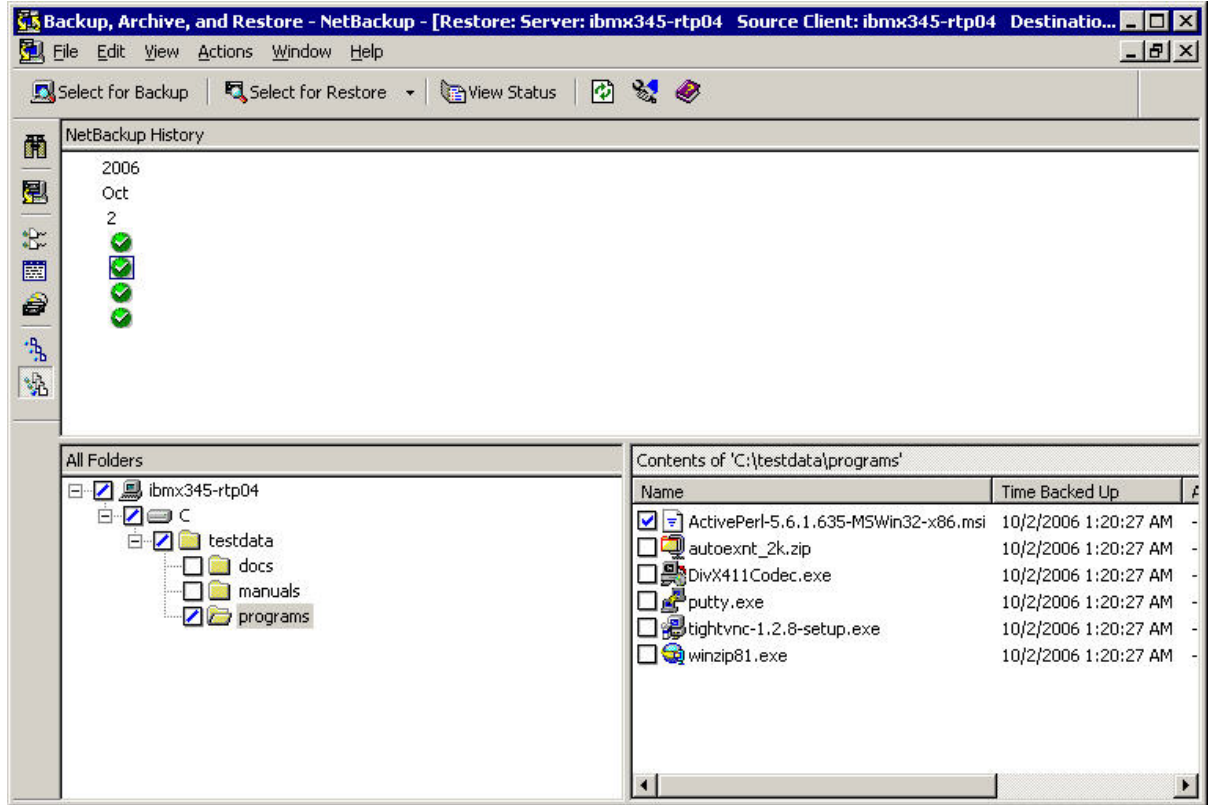


Figure 5-33) NetBackup BAR – browsing selections for restore.

You can highlight the backup from which you want to restore files and navigate to the information that you want to restore.

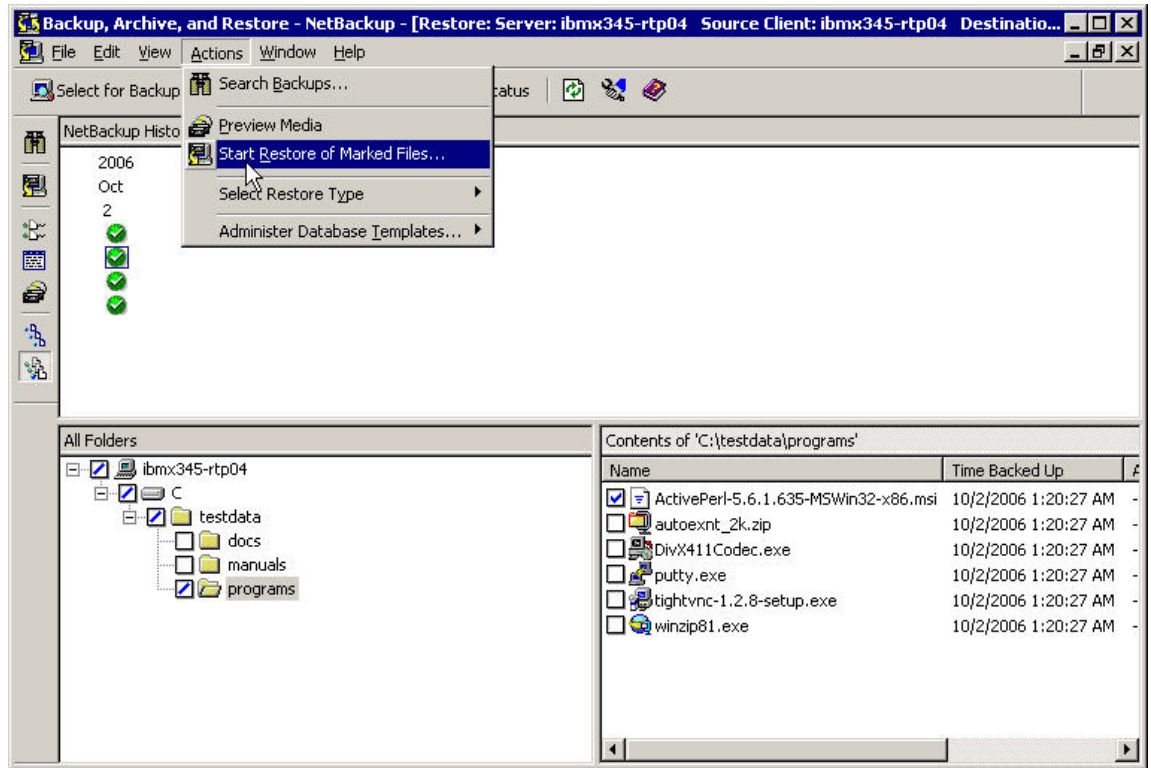


Figure 5-34) NetBackup BAR – starting restore of marked files.

When you have selected the files and directories to restore, select Actions > Start Restore of Marked Files to open the Restore Marked Files window.

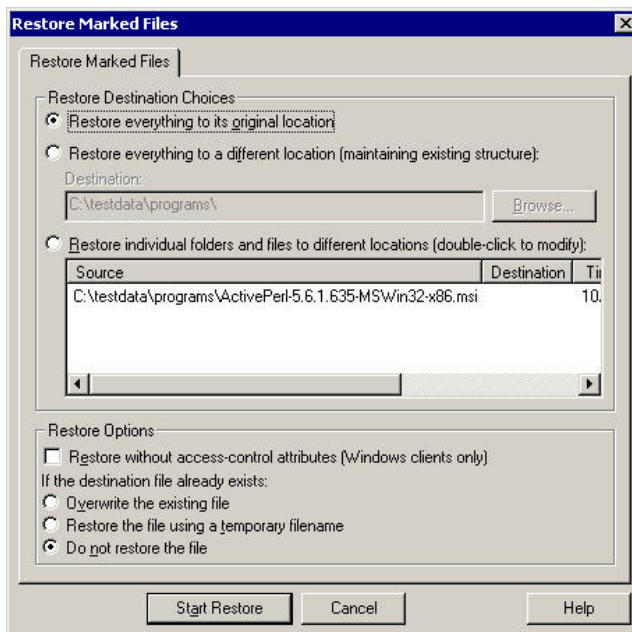


Figure 5-35) NetBackup BAR – Restore Marked Files window.

The Restore Marked Files window provides numerous options for how to do the restore. See the *NetBackup System Administrator's Guide* for details.

After clicking Start Restore, you are prompted to view the status. Either click Yes at that point or later click the View Status button in the BAR window.

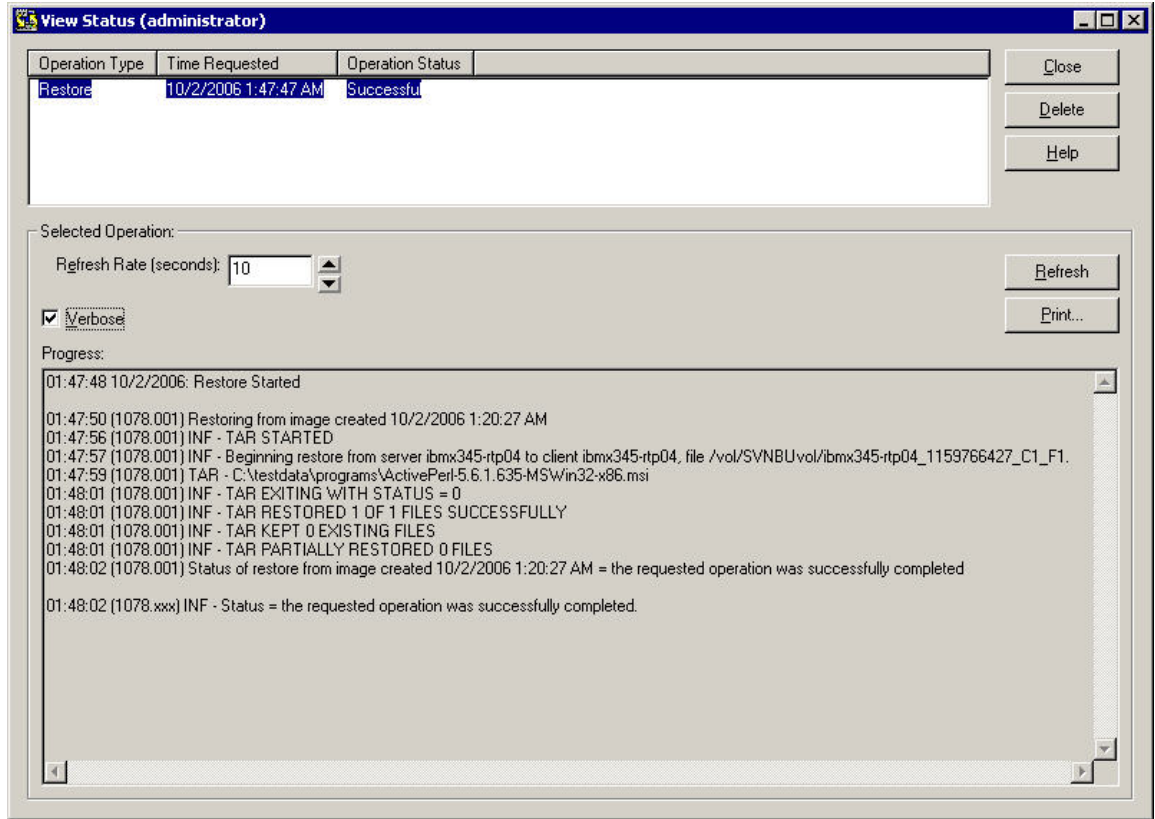


Figure 5-36) NetBackup BAR – viewing restore status.

The restore job is also displayed in the Activity Monitor.

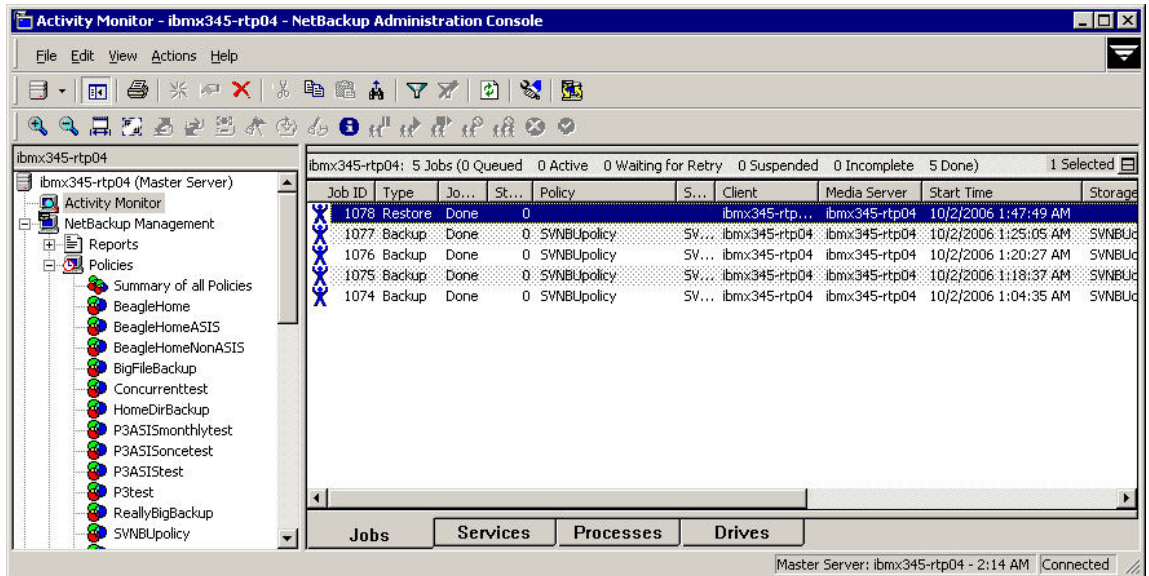


Figure 5-37) NetBackup Activity Monitor – restore job.

Drag-and-Drop via FSE

One of the nice features of SV-NBU is that the backups can be stored in such a way that they can be accessed as a standard Windows or UNIX file system in Data ONTAP. Therefore drag-and-drop or copy-and-paste restores can be accomplished from them without using NetBackup.

For Windows to access the backup data, a CIFS share must be created for the volume to which the backups were done. To do this, select CIFS > Shares > Add in FilerView.

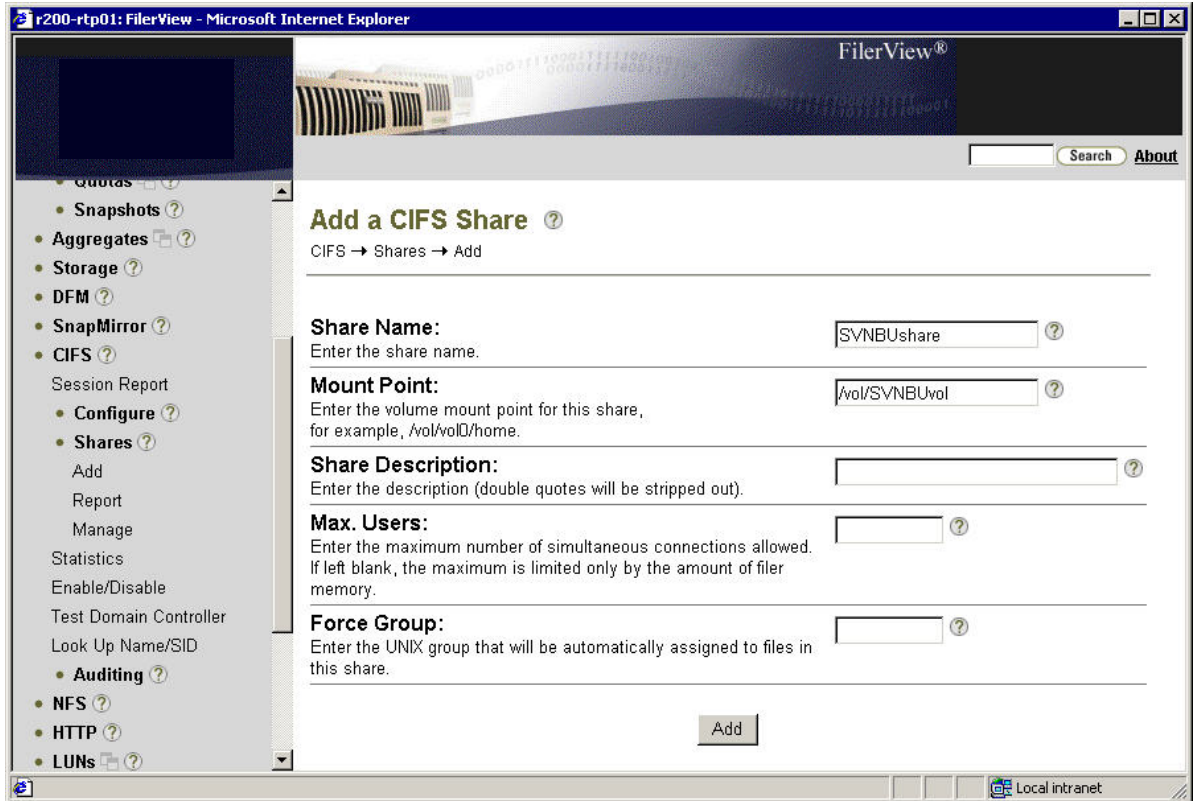


Figure 5-38) Creating a CIFS share of the SV-NBU volume in FilerView.

Enter a volume and give the share a name, then click Add.

The share can now be mapped to a network drive on a Windows box.

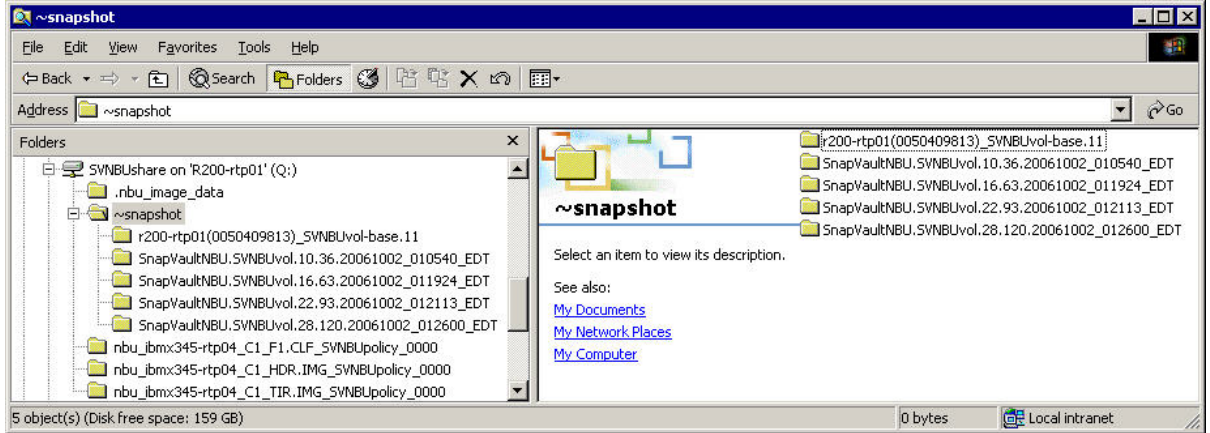


Figure 5-39) SV-NBU restores via FSE.

The four Snapshot copies with the date and time stamps contain the four backups that were sent to this volume. Navigate to the one you want and grab the required files and directories.

For UNIX to access the backup data, the UNIX machine needs to simply mount the NFS export of the volume from the IBM N series, which is probably already there. A Snapshot copy exists for each of the backups.

```

root@sunv240-rtp03 # pwd
/svnbufsemount/.snapshot
root@sunv240-rtp03 # ls
SnapVaultNBU.SVNBUvol.10.36.20061002_010540_EDT
SnapVaultNBU.SVNBUvol.16.63.20061002_011924_EDT
SnapVaultNBU.SVNBUvol.22.93.20061002_012113_EDT
SnapVaultNBU.SVNBUvol.28.120.20061002_012600_EDT
r200-rtp01(0050409813)_SVNBUvol-base.11

```

You navigate to the desired location and copy the desired files and directories.

Here is the information on exactly where the information is located in each Snapshot directory:

/svnbufsemount/	< Export/share
.snapshot/	< Snapshot directory
SnapVaultNBU.SVNBUvol.10.36.20061002_010540_EDT /	< Date/Time of backup
nbu_ibmx345-rtp04_C1_F1.CLF_SVNBUpolicy_0000/	< The "ClearFile" directory
nbu_online_data/	< Where the data lives
testdata	< Policy "backup selection"



Operating characteristics

This section discusses the behavior that you can expect of the solution. Information in this section comes from testing, observations, and knowledge of how the solution is built.

The Performance section discusses the performance understandings of the solution. Continued and additional testing is being conducted as additional platforms are supported and features announced.

The Storage Overhead section discusses storage savings that the solution can be expected to deliver. Again, more testing is ongoing, and results will continue to be delivered and enhanced in updates to this document.

The Limitations section discusses what is not supported, what you shouldn't do, and what may work but still needs to be tested (perhaps that could occur in the field). Some of this information may be covered elsewhere, but it bears reiterating here.

The Best Practices section contains items that might not have been covered elsewhere in this document, as well as (*in future revisions*) lessons learned from real-world implementations.

SV-NBU target environment

SV-NBU in Data ONTAP 7.2 is ideally suited and supported for file services backups, such as home directories, file shares, etc.

SV-NBU in Data ONTAP 7.2 is not supported for the backup of application data, such as databases and e-mail.

Specifically, SV-NBU is a good fit for environments where the following are true:

Customer wants to back up non-IBM N series primary storage and uses NetBackup 6.0.

Customer desires easy online access to a long history of backups. The SV-NBU space-saving technology provides a cost-effective way of achieving this.

A majority of the data to be backed up remains unchanged between backups—typical of home directories and file shares.

SV-NBU in Data ONTAP 7.2 is *not* recommended for environments where any of the following is true:

Raw backup throughput is the primary concern.

Data to be backed up consists mostly of large (> 500MB) Windows files

Data to be backed up consists mostly of small (< 10KB) files.

The remaining sections give details that support these points.



SV-NBU performance

SV-NBU in Data ONTAP 7.2 has been optimized for storage efficiency of file system backups, while providing backup throughput comparable to a NetBackup Basic DSU solution.

SV-NBU in Data ONTAP 7.2 can achieve the following levels of throughput when backing up a typical file system data set with Gigabit Ethernet network connectivity. (Restore performance is comparable.)

Platform	Single Stream	Aggregate
IBM N series with NearStore	35 MB/s	60 MB/s
N5200	30 MB/s	60 MB/s
N5500	35 MB/s	80 MB/s

Table 6-1) SV-NBU performance.

When determining what a specific environment can expect in terms of performance, several factors should be considered:

Overall backup throughput is determined by the composition of the data set, including the number of files and directories and the size of individual files.

Single-stream backup throughput is negatively affected for data streams consisting predominantly of small files or deep directory structures.

Single-stream backup throughput is negatively affected for Windows data streams consisting primarily of large files (typically greater than 500MB). This limitation does not apply to UNIX data streams. The limitation is due to the structure of the Windows “backup read” formatted files delivered by NetBackup. For full backups, the greater the number of modified files, the lower the overall throughput. This limitation is due to increased SV-NBU processing required for modified files.



SV-NBU storage savings

SV-NBU offers outstanding storage savings via block-sharing and deduplication technology. All duplicate data blocks residing in multiple backup images of a given client-policy-path are eliminated. The actual space savings and storage overhead that one could expect or compute are a function of:

- The size of files
- The number of files and directories
- The number of backup copies kept (retention)
- How much the data is changing (data change rate).

In general, greater storage savings are achieved with:

- Larger files
- Fewer files and directories
- More copies kept (longer retention)
- Less data changed.

The following table shows some examples of storage savings achieved in various environments. When examining and using these numbers, it is important to note the following caveats:

- The Web Server, Windows, and Large Files Library tests only included performing full backups.
- The home directory was a live real-world system and had weekly fulls and daily incrementals run.
- The Windows C\$ data set was a live NetBackup master server.

Data Set	Size	# Files	# Backups	Compression	Storage Savings
Home Directory	1 TB	2 M	5	3.6:1	72%
Web Server	150 GB	250 K	12	7.7:1	87%
Windows C\$	10 GB	25 K	111	11:1	91%
Large Files Library	16 GB	920	162	20:1	95%

Table 6-2) SV-NBU storage savings.

After each backup, specific results were kept for each of the first three data sets. The following figure shows the storage savings achieved over a month.




Figure 6-1) SV-NBU storage savings over a month.

Key points:

Storage savings graphed are roughly those of a month.

Space savings are fairly consistent across data sets—that is, they follow a predictable path, rarely varying more than 5% regardless of data.

The following figures show more specific graphed examples of storage savings. These figures demonstrate the effect of change rate, average file size, and expiring backups.

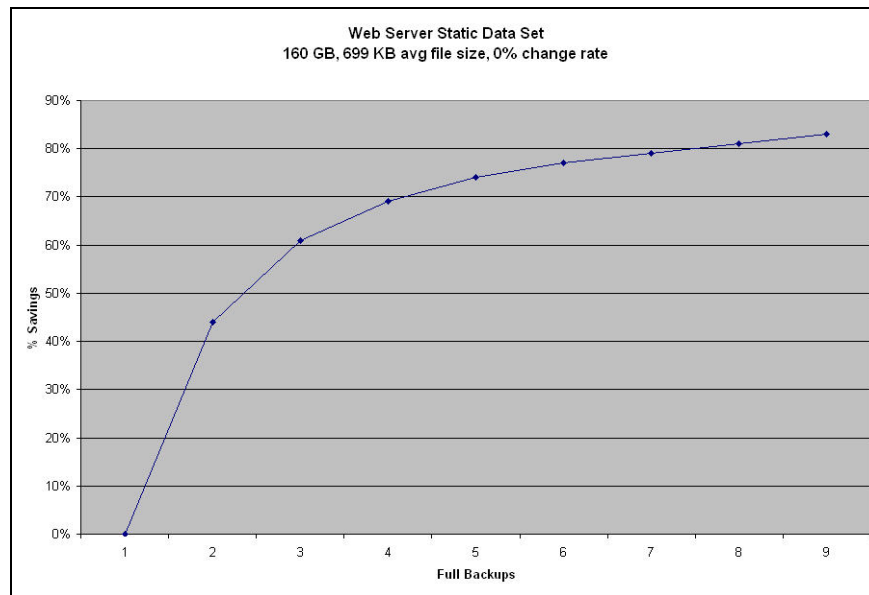


Figure 6-2) Static Web server storage savings.

Both the static nature of the Web server data being backed up and the large average file size lead to increasing storage savings, as shown in Figure 6-2. Being static means that there is no new unique data that needs to be stored. The larger files mean that there's less overhead. As more backups are performed, it's easy to see that storage savings will continue to grow.

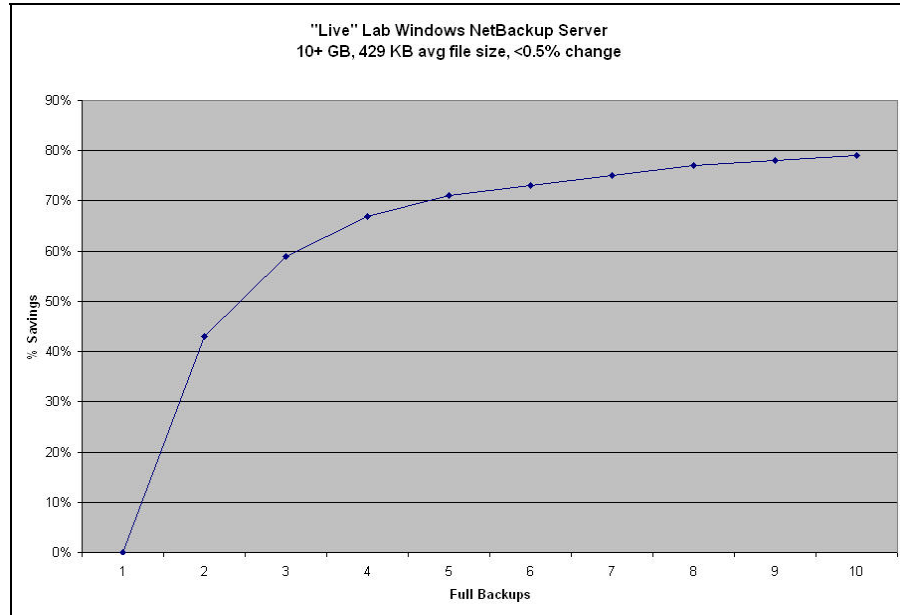


Figure 6-3) "Live" Lab Windows NetBackup master server storage savings.

In Figure 6-3, with a change rate of only 0.5%, storage savings are still very good, but the smaller average file size means that the storage savings don't increase quite as quickly as those seen in the static example of Figure 6-2. The changes are predominantly due to new data (files) in the NetBackup catalog. As more backups are performed, the storage savings continue to grow.

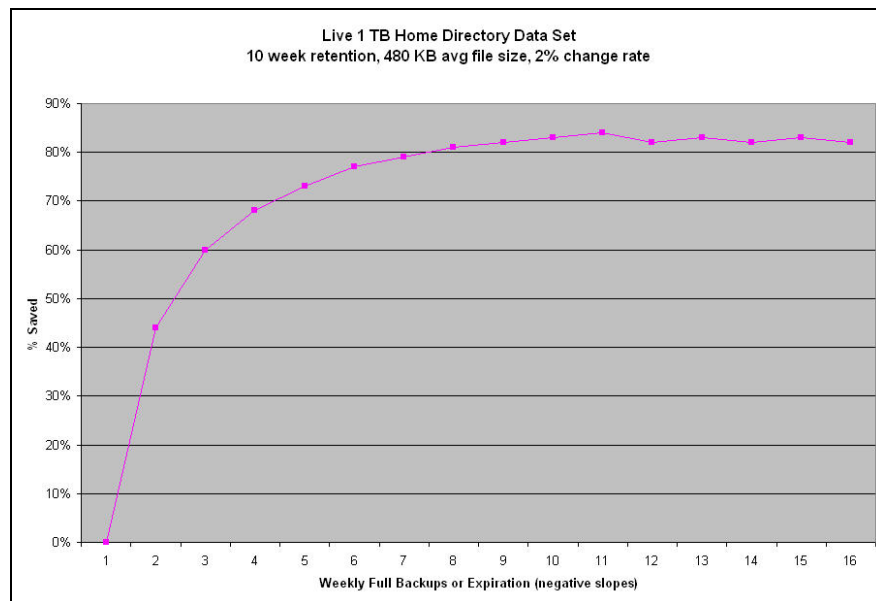


Figure 6-4) Live 1 TB home directory storage savings.

In Figure 6-4, storage savings grow typically in the same way as other data sets. The changes are due to both new data and modified data of users. This graph also shows what happens with storage savings when backup images expire: The storage savings actually decrease slightly. The total savings peak at around 83% in this case.




Figure 6-5) Fairly static home directory storage savings.

Figure 6-5 shows a subset of the home directory data set with smaller average file size (about a third of that shown in Figure 6-4). As a result, the savings grow more slowly and also peak lower.




Figure 6-6) Dynamic home directory storage savings.

Figure 6-6 demonstrates that by retaining fewer backup images, the maximum storage savings achieved are less. The higher change rate (20%) also decreases the storage savings.

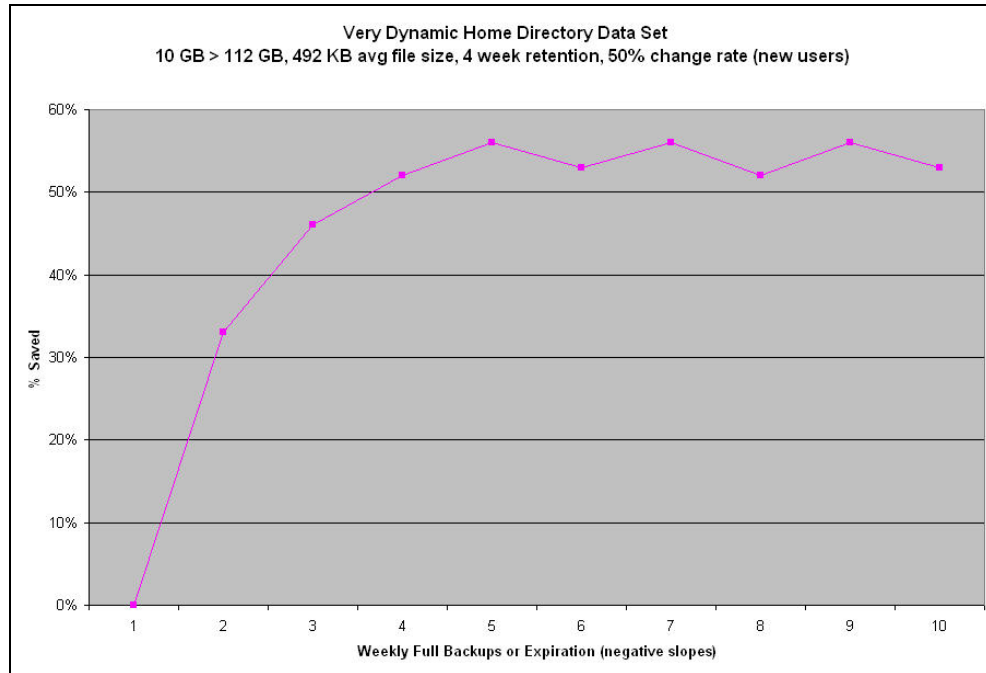


Figure 6-7) Very dynamic home directory storage savings.

Figure 6-7 demonstrates an extreme case of a 50% change rate (all new data in this case) and keeping backups for only a month. Although the savings are the lowest of all the data sets examined, note that over 50% storage savings are still achieved.

SV-NBU storage overhead

As stated, SV-NBU offers outstanding storage savings through the use of block-sharing and deduplication technology. All duplicate data blocks residing in multiple backup images of a given client-policy-path are eliminated. However, there is a small amount of storage overhead associated with the SV-NBU solution.

Total overhead for typical home directory environments is less than 10% of backup image size.

When determining what a specific environment can expect for storage overhead, there are two categories of overhead that should be considered: mapping and alignment.

Mapping overhead is the result of data structures required to translate the NetBackup 256 byte blocks to WAFL 4KB blocks, and is fixed at 3.2% of the backup image size.

Alignment overhead is the result of data block padding when converting from NetBackup 256 byte blocks to WAFL 4KB blocks and is variable as a function of the number of files and directories. The overhead is typically 3.5KB per directory, 3.5 to 7KB per UNIX file, and 7 to 10KB per Windows file. Storage efficiency is negatively affected by small files. This effect is exacerbated if small files make up a significant percentage of the total data.

Using the 1TB home directory example in the Storage Savings section, only 250GB of total overhead was required for the approximately 3.3TB of stored backup images, roughly 7%.



SV-NBU limitations

The following table shows exactly what is supported for the SV-NBU solution.

IBM N series Hardware	N series with NearStore; N5200, N5500 or higher
Data ONTAP	Data ONTAP 7.2 or later
IBM N series Software	sv_ontap_sec nearstore_option (for N5x00)
NetBackup	NetBackup Enterprise 6.0MP4 or later (NDMP Option – installed, but need not be licensed) Disk Optimization Option
Protocols	N/A
Client	(Whatever UNIX and Windows platform/OS NetBackup supports) NetBackup 5.0 or later
Media Server	(Whatever UNIX and Windows platform/OS NetBackup supports) NetBackup 6.0MP4 or later
Applications	File Services

Table 6-3) SV-NBU solution requirements overview.

Replication of the SV-NBU volume is currently not supported in any form. IBM N series cluster services are not supported and the cluster license should not be installed. If the cluster license is present, you will not be able to connect to the SV-NBU volume.

Backup of the SV-NBU volume to tape via NDMP or native dump is not supported. (Use the NetBackup Inline Tape Copy feature or the NetBackup Vault option to create redundant copies of the NetBackup backup images.) There is no space optimization when the data is written to tape, and you can't recover the SV-NBU volume from tape and then resume SV-NBU backups to that volume.

Only FlexVol volumes are supported; no traditional and no WORM (write once read many) volumes are supported. SV-NBU should not share a volume with anything else, and cannot share one with native SnapVault or NetBackup SnapVault Management (NSVM).

There is no checkpoint restart. If an SV-NBU backup fails mid-transfer, a new transfer starts at the beginning again. All data transferred prior to the failure is discarded. FlexVol volumes must be 4TB or smaller, based on platform type. Maximum FlexVol size is a function of platform memory:

Platform memory affects volume sizing because the `refcount` file for the SV-NBU volume must fit in memory to ensure consistency when `WAPL_check` is run. (These processes run serially, so you can have multiple volumes.) See the following maximum usable volume sizes (including `snap reserve` space) for the associated IBM N series.

N5200	N5500	IBM N series with NearStore
1TB	2TB	4TB

Table 6-4) Maximum SV-NBU volume sizes.

This is important if the volumes move to a platform with a smaller maximum volume size.

NearStore storage units cannot be used as part of a NetBackup storage unit group. On the NetBackup policy Attributes tab, do not select NetBackup Compression. Compressed backups are always sent from the media server to the IBM N series NearStore secondary storage system as an image data stream instead of a clearfile stream. Because of this the data stream is not aligned, unpacked, and deduplicated. Also, even if FSE is enabled on the storage unit, if NetBackup compressed backups are sent, the backup will not be exportable and accessible via CIFS/NFS.



SV-NBU best practices

The use of Gigabit Ethernet between the media server and the IBM N series is strongly recommended in order to maximize backup performance. Use of 100base-T will probably result in unacceptable performance. NetBackup policy considerations:

For data consistency, full and incremental backups for a given NetBackup policy should go to the same SV-NBU storage unit.

Policies with similar backup schedules (frequency and retention) should go to the same SV-NBU storage unit.

Policies with varying backup schedules must go to separate volumes. For example, if policy1 has a weekly full schedule and policy2 has an hourly full schedule, they should go to separate volumes.

Policies with different retention periods should go to separate volumes. If policy1 has a retention period of 2 months and policy2 has a retention period of 6 months, they should be in separate volumes.

Don't direct more than 100 NetBackup policies to the same SV-NBU storage unit.

Don't use NetBackup compression. Compressed backups are sent as image and are not deduplicated.

For UNIX client backups, the `DO_NOT_RESET_ATIME` option should be set on the NetBackup client.

This provides performance benefits.

NetBackup storage unit considerations:

Separate SV-NBU storage units should be used for Windows and UNIX NetBackup clients, as well as for backup clients configured for different languages.

Do not rename the SV-NBU storage unit volume.

Once a storage unit is created with "Enable block sharing" and backed up, do not change it.

An IBM N series volume should not be configured as an SV-NBU storage unit on more than one media server. But if it is done, block sharing for the SV-NBU storage units should be enabled or disabled the same on both.

Because Data ONTAP limits total snapshots to 255 per volume, consider disabling any unneeded scheduled Data ONTAP snapshots (configured with the `snap sched` command). On an SV-NBU storage unit, use one of the following:

```
snapsched MYVOL 0 0 0
vol options MYVOL nosnap true
```

If a volume is no longer needed for SV-NBU, do not destroy it without first using NetBackup to expire all the backup images contained on it.

Do not rename the SV-NBU storage unit volume.

For Windows backups, exclude `pagefile.sys`.

Open port 10571 if it is behind the firewall.



Common problems and troubleshooting

This section covers known issues in installing, configuring, and deploying SV-NBU.

SV-NBU

One of the most common problems is incorrectly configuring NDMP authentication. When troubleshooting, begin by ensuring that NDMP is installed, licensed, and configured properly between the master server (and any media servers involved) and the IBM N series.

```
tpautoconf-verify -nh <hostname>
```

For troubleshooting possible NDMP-related problems on the IBM N series, issue `ndmpd debug 70`, reproduce the problem, and examine `/etc/log/ndmpdlog.*`

Make sure that NetBackup is 6.0 MP4 or later.

You must do at least one full backup before any incrementals.

If the drop-down list for the NearStore volumes doesn't appear, ensure that you issued the `options snapvault.access all` (or specify the names of the master and media servers) command.

Check the version of Data ONTAP; it must be version 7.2RC3 or later.

There is a new log file: `/etc/log/nbu_snapvault`.

```
msg Sun Mar 27 00:07:07 PST [172.29.19.90:55810] msgtype=0x200
msg Sun Mar 27 00:07:07 PST [172.29.19.90:55810] msgerr=0
msg Sun Mar 27 00:07:07 PST [172.29.19.90:55810] msgtype=0x300
msg Sun Mar 27 00:07:07 PST [172.29.19.90:55810] msgerr=0
```

Debugging and diagnostics related to new features:

If a problem is experienced, set the NDMP trace level higher with `ndmpd debug 70`, reproduce the problem, and examine `/etc/log/ndmpdlog.*`

If a problem is experienced, set the NetBackup SnapVault trace level higher with `options snapvault.nbu.trace_level`, reproduce the problem, and examine `/etc/log/nbu_snapvault*`.

A reasonable number to use for this parameter for debugging is 100.

Note that this parameter is set to zero by default. Higher numbers provide increased debug information (in `/etc/log/nbu_snapvault`), but backups are slower. Therefore, when you are finished troubleshooting, set it back to zero.

If jobs are failing to write to the NearStore system, make sure that the space reserved for Snapshot copies on the NearStore system is not completely full. When the reserved space is full, NetBackup uses the active file system space as needed.

In the case of a disk-full condition on the NearStore system, make sure that no WAFL Snapshot copies are consuming disk space unnecessarily.

The maximum number of concurrent backup and/or restore connections is 128. If the maximum number of transfers allowed to a single NearStore system is exceeded, the Data ONTAP kernel reports the following error:



```
Inf Wed Jul 6 07:28:27 CDT [10.80.106.36:58645] Maximum active \
transfers reached.
```

Common problem: Running out of snapshots.

Policies with varying backup schedules must go to separate volumes. For example, if policy1 has a weekly full schedule and policy2 has an hourly full schedule, they should go to separate volumes. Policies with different retention periods should go to separate volumes. If policy1 has a retention period of 2 months and policy2 has a retention period of 6 months, they should be in separate volumes.

EMS error messages (/etc/messages):

```
Snapshot creation failed
Open/Write failure to the protocol log file
Platform mismatch error
Qtree deletion error
UTF translation error
Volume language warning.
```

Never delete archival snapshots manually or you will lose data! Use NetBackup to delete them.

```
harte> snap list p3vol
```

%/used	%/total	date	name
0% (0%)	0% (0%)	Dec 24 20:15	SnapVaultNBU.p3vol.207.12479.20051224_201500_PST (snapvault)

The "snapvault" in parentheses indicates that this is a *soft-locked* snapshot used by P3 and should not be manually deleted.

```
snap list -q tells you, for each snapshot, which qtrees are captured in that snapshot.
snap list -o tells you, for each qtree, which snapshots contain this qtree.
```

If you see unusual STU errors such as 174 in the NetBackup Activity Monitor, there is a chance that clustering has been licensed or enabled on the IBM N series. This is not supported, and the only way to resolve the problem is to remove the license and reboot the storage system.



Verify backup

The SV-NBU solution's performance optimization is based on the assumption that `mtime` always changes if a file changes. If the assumption may not be correct for a particular customer environment, bypass this performance enhancement by turning on Verify Backup; use the following command:

```
options snapvault.nbu.verify_backup
```

When you turn on Verify Backup, a `bcompare` is performed on every file that it is identified as not being changed.

Key things to understand about using this command:

Off by default.

If set to On, disregards the performance optimizations that use the file's `mtime`.

If set to On, backups are *significantly* slower.

Should be set to On if `mtime` of the file does not change when the file's contents change.

If the customer is not sure about these `mtime` assumptions, use Verify Backup on a one-time basis to verify that deduplication is working properly.

If the customer knows that the apps don't change the `mtime`, but the data can change, Verify Backup should be set to On.



Kernel settings for NetBackup

There can be issues with NetBackup if kernel settings aren't at least at some minimal settings. In existing NetBackup production environments, they should be fine, but if you are setting up a new NetBackup environment, the following are reasonable settings to use on the master server:

```
-----  
* Message queues  
set msgsys:msginfo_msgmap=500  
set msgsys:msginfo_msgmax=8192  
set msgsys:msginfo_msgmnb=65536  
set msgsys:msginfo_msgmni=256  
set msgsys:msginfo_msgssz=32  
set msgsys:msginfo_msgtql=500  
set msgsys:msginfo_msgseg=8192  
  
* Semaphores  
set semsys:seminfo_semmap=64  
set semsys:seminfo_semmni=1024  
set semsys:seminfo_semmns=1024  
set semsys:seminfo_semmnu=1024  
set semsys:seminfo_semmsl=300  
set semsys:seminfo_semopm=32  
set semsys:seminfo_semume=64  
  
* Shared memory  
set shmsys:shminfo_shmmax=16777216  
set shmsys:shminfo_shmmin=1  
set shmsys:shminfo_shmmni=230  
set shmsys:shminfo_shmseg=100  
-----
```

See the NetBackup Planning and Performance Tuning Guide for details on kernel tuning.



NetBackup tuning parameters

For optimal NetBackup performance, some tuning parameters can be configured that usually provide substantially better performance. This section shows typical values.

Create the following files and put the indicated values into them:

```
<install_path>/netbackup/NET_BUFFER_SZ 1048576
<install_path>/netbackup/db/config/NUMBER_DATA_BUFFERS 16
<install_path>/netbackup/db/config/SIZE_DATA_BUFFERS_DISK 1048576
```

Where:

UNIX, install_path = /usr/opensv

Windows, install_path = C:\Program Files\VERITAS

See the *NetBackup Backup Planning and Performance Tuning Guide* for additional specifics on tuning NetBackup performance.

On UNIX clients, add the following to the `/usr/opensv/netbackup/bp.conf` file for improved performance:

```
DO_NOT_RESET_FILE_ACCESS_TIME
```

Tape and disaster recovery scenarios

Although there are substantial benefits to be achieved with the SV-NBU solution, the complete end-to-end data protection story doesn't typically end there.

Most customers will still want to, as a minimum, put their backup data onto tape at some point. And more advanced customers will want to look at ways to replicate their backup data to a remote location for DR purposes (and from there perhaps put it on tape). These two situations can be seen by looking again at a figure presented earlier in this document.

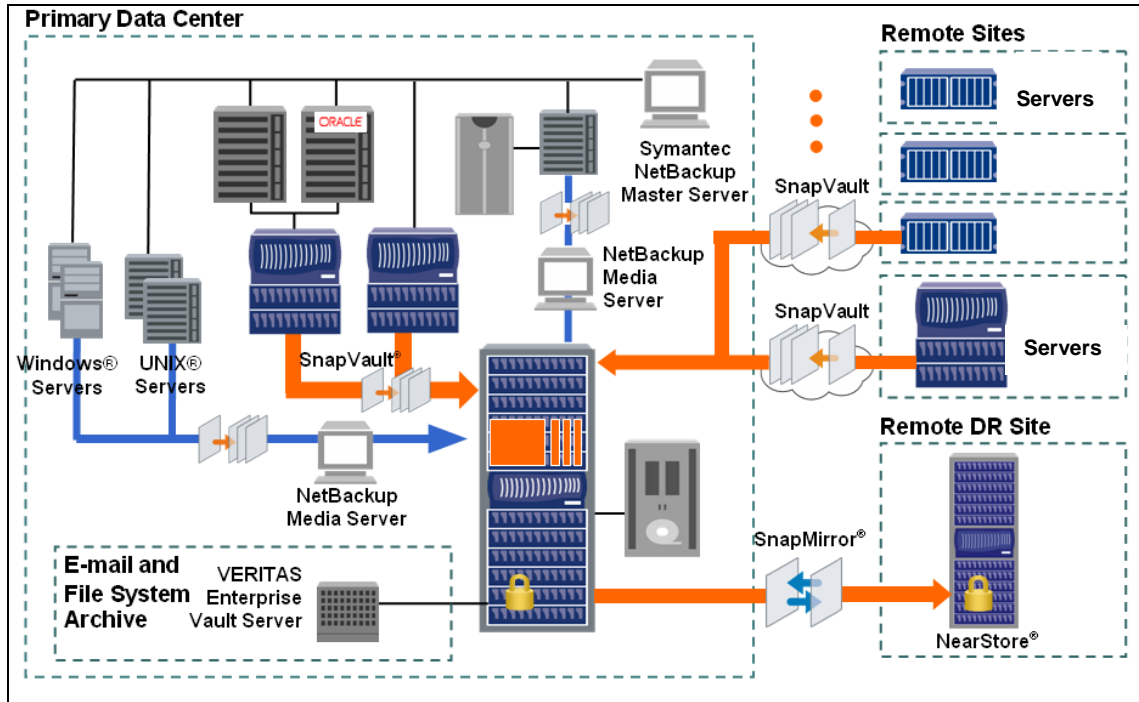


Figure 8-1) Joint solutions and disaster recovery.

Tape

Once backup images reside on the IBM N series, you can use other methods to put the images on tape for offsite storage. The following sections cover specifics for each solution.

SV-NBU

For SV-NBU, NDMP is not supported, for the following reasons:

- When you copy a whole SV-NBU volume using NDMPcopy, only the contents outside of the SV-NBU qtrees in the volume (non-SV-NBU data) are copied.
- When you copy an SV-NBU qtree using NDMPcopy, nothing gets copied.
- When you do a subtree dump of an SV-NBU directory, the content does get copied, but none of the SV-NBU volume metadata that is kept in the registry residing on the root volume of the IBM N series is written, so you wouldn't be able to use NetBackup to restore the data to the IBM N series.

To move backup images to tape, there are two choices:

1. Use Inline Tape Copy (ITC) for the policy schedules. This feature of NetBackup sends a copy of the backup stream to the tape and to the NearStore DSU simultaneously.

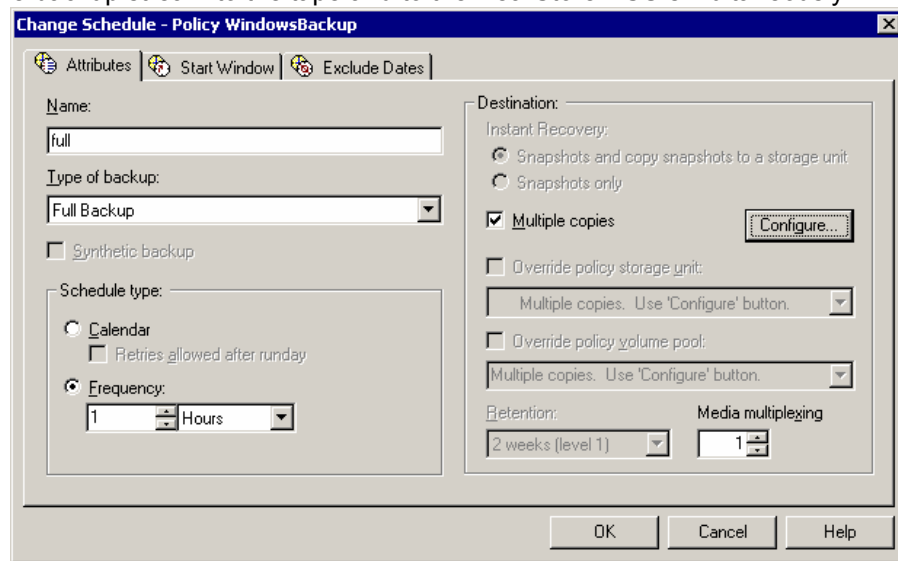


Figure 8-2) Enabling Inline Tape Copy (ITC).

After selecting “Multiple copies,” you can configure the tape storage unit to which to send the copies and the retention desired.

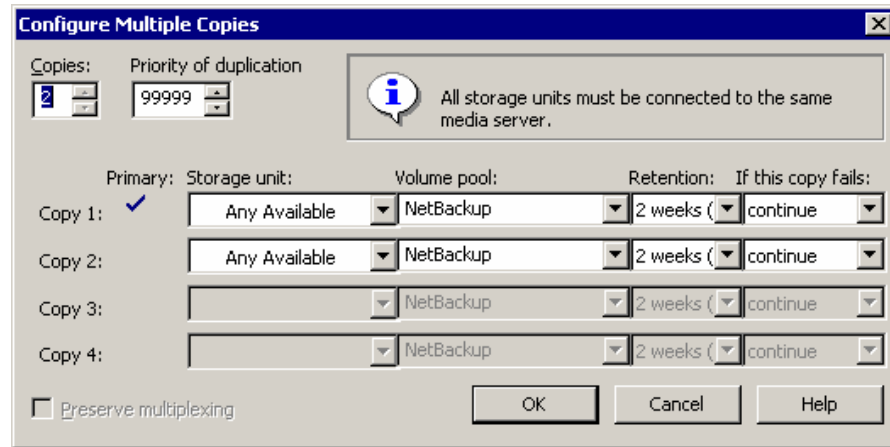


Figure 8-3) Specifying ITC destinations.

2. The other choice is to use the NetBackup Vault Option. (See the *NetBackup Vault System Administrator's Guide* for specifics on tuning NetBackup performance.)

The important thing about both of these methods is that they are completely integrated into NetBackup, and the client restore can be done from either the NearStore storage unit or the tape. The downside is that the data moves from the IBM N series through a media server and then to tape, so it is likely that performance will be decreased.

Note that there is no space optimization when the data is written to tape, and you can't recover the SV-NBU volume from tape and then resume SV-NBU backups to that volume.

SnapMirror

Once backup images reside on the IBM N series, typically IBM System Storage N series with SnapMirror® technology can be used to replicate the backup images to an online DR site. However, NetBackup 6.0MP4 and Data ONTAP 7.2 don't provide an integrated automated solution. Similarly, for DR of the backup environment itself, you can use SnapMirror to replicate the NetBackup catalogs to a DR site, and bring up a master server there in the event of a disaster. Refer to the NetBackup documentation for information about bringing this NetBackup DR environment online. (For a complete discussion of SnapMirror, refer to TR-3390, *SnapMirror Deployment and Implementation Guide*.)

SV-NBU and SnapMirror

SnapMirror with SV-NBU is not supported.

When the TAR data stream from NetBackup is converted to 4KB blocks by Data ONTAP and turned into a real WAFL file system, block mapping and block reference count maps are created and maintained. Some of this information is included in the registry residing on the root volume of the IBM N series. Because of this, replication of the SV-NBU volume is currently not supported in any form.

Technology primer

Because the joint solutions described in this paper involve both Symantec NetBackup and IBM N series hardware and Data ONTAP technologies, this section provides a brief high-level primer on Symantec and IBM N series products that are key components of the joint disk-to-disk solutions described in this document.

Symantec NetBackup

This section provides a general overview of NetBackup.

Architecture

NetBackup consists of both the server and client software:

- Server software resides on the computer that manages the storage devices.
- Client software resides on the computer whose data you want to back up. A server also has client software and can be backed up like other clients.
- The storage devices that client data is backed up to are called storage units. They are virtual representations of physical tape drives or disk drives.




Figure 9-1) NetBackup storage domain.

NetBackup supports both master and media servers. The master server manages the backups, archives, and restores. Media servers provide additional storage by allowing NetBackup to use the storage devices that they control. Media servers can also increase performance by distributing the network load.




Figure 9-2) NetBackup architecture.

During a backup or archive, the client sends backup data across the network to a NetBackup server that has the type of storage specified for the client. The storage requirement is specified during NetBackup configuration (for example, LTO tape or disk).

During a restore, users use the NetBackup GUI to browse to and select the files and directories that they want to recover. NetBackup finds the selected files and directories and restores them to the disk on the client.



Storage units

A NetBackup storage unit is a storage device attached to a NetBackup server. In order to send backups to a storage device, the administrator needs to define storage units by using the Storage Units utility.

There are three types of storage units:

- **Media Manager Storage Units:** This type encompasses the tape robots, standalone tape drives, and optical disk devices, all of which are under the control of Media Manager. Media Manager controls the allocation and mounting of media (called *volumes*) in the storage devices.
- **NDMP Storage Units:** NDMP storage units are controlled by Media Manager but attach to NDMP hosts and require that the NetBackup for NDMP option be installed.
- **Disk Storage Units:** A disk type storage unit consists of a directory on a disk that stores data. NetBackup permits an unlimited number of DSUs. There are three types of DSUs:

Basic disk, used for traditional NetBackup DSUs. These simply specify directories on the media server in which to store the backups.

NearStore, used for any heterogeneous client. This is the DSU type that is used for SnapVault for NetBackup. NearStore appears as a selection only when the NetBackup Disk Optimization Option is licensed. (Note that despite “NearStore” being displayed in the GUI, other IBM N series are supported if properly licensed with the NearStore Option.)

SnapVault, used for network-attached storage (NAS). SnapVault appears as a selection only when the NetBackup IBM N series with SnapVault Option is licensed.

IBM N series with Data ONTAP

This section provides a general overview of various components of Data ONTAP, the “operating system” of the IBM N series. There are many other features and options not discussed here; the items covered in this section are those that are pertinent to the integrated solution with NetBackup.

Data ONTAP is a powerful enterprise storage environment that delivers a flexible storage infrastructure providing high performance, massive and simple scalability, and the intelligence and automation necessary to minimize management overhead. Version 7.1 and beyond have features such as integration with NetBackup 6.0 and multiple iSCSI connections.

IBM N series with WAFL

WAFL stands for “Write Anywhere File Layout.” It is, essentially, the file system that exists on an IBM N series. WAFL is write optimized and always writes new disk blocks to available locations on disk. Preexisting data blocks are never overwritten. This is true whether a new file is being created or an existing file is being updated. This process minimizes disk drive seeks, which improves performance. Thus preexisting blocks still exist and can be accessed in the form of Snapshot copies until they are deleted. It’s important to point out, however, that a Snapshot copy is *not* a copy of data; rather a Snapshot copy records the state of the blocks in the file system at a given point in time and provides read-only access to that *image* of the file system.



IBM N series with Snapshot and SnapRestore

A Snapshot copy is a locally retained point-in-time image of data. IBM N series with Snapshot technology is a feature of the WAFL storage virtualization technology that is a part of Data ONTAP. A Snapshot copy is a "frozen," read-only view of a WAFL volume that provides easy access to old versions of files, directory hierarchies, and/or LUNs (logical unit numbers).

The high performance of Snapshot also makes it highly scalable. A Snapshot copy takes only a few seconds to create—typically less than 1 second, regardless of the size of the volume or the level of activity on the IBM N series. After a Snapshot copy has been created, changes to data objects are reflected in updates to the current version of the objects, as if Snapshot copies did not exist. Meanwhile, the Snapshot version of the data remains completely stable. A Snapshot copy incurs no performance overhead; users can comfortably store up to 255 Snapshot copies per WAFL volume, all of which are accessible as read-only and online versions of the data.

IBM N series with Snapshot technology makes extremely efficient use of storage by storing only block-level changes between each successive Snapshot copy. Because the Snapshot process is automatic and virtually instantaneous, backups are significantly faster and simpler.

System administrators use Snapshot copies to facilitate frequent, low-impact, user-recoverable backups of files, directory hierarchies, LUNs, and/or application data. Snapshot copies vastly improve the frequency and reliability of backups, because they incur minimal performance overhead and can be safely created on a running system.

Snapshot copies provide near-instantaneous, secure, user-managed restores. Users can directly access Snapshot copies to recover from accidental deletions, corruptions, or modifications of their data. Because the security of the file is retained in the Snapshot copy, the restoration is both secure and simple.

IBM System Storage N series with SnapRestore® software uses Snapshot technology to perform near-instantaneous data restoration.

SnapRestore software allows an enterprise to recover almost instantly from any number of disaster scenarios. In seconds, SnapRestore software can recover anything from an individual file to a multi-terabyte volume so that operations can be quickly resumed. From a single home directory to a huge production database, SnapRestore does the job in seconds, regardless of file or volume size.

With SnapRestore, data can be restored from any one of the Snapshot copies stored on the file system. This allows an application development team, for example, to revert to Snapshot copies from various stages of their design, or test engineers to quickly and easily return data to a baseline state. Restoring to the base environment takes only seconds, and the restored environment is identical to the point at which the Snapshot copy was created.



Appendix 1: Acronyms

AFS	Active File System
ASIS	Advanced Single Instance Storage
BAR	Backup, Archive, and Restore
CIFS	Common Internet File System
CLI	Command Line Interface
DB	Database
DR	Disaster Recovery
DSU	Disk Storage Unit
GUI	Graphical User Interface
ICS	Infrastructure Core Services
ITC	Inline Tape Copy
LUN	Logical Unit Number
NAS	Network-Attached Storage
NDMP	Network Data Management Protocol
NFS	Network File System
PSE	Professional Services Engineer
QSM	Qtree SnapMirror
RPO	Recovery Point Objective
SE	Systems Engineer
SIS	Single Instance Storage
SLA	Service Level Agreement
SM	SnapMirror
SOI	Space-Optimized Image
SV	SnapVault
SV-NBU	SnapVault for NetBackup
TCE	Total Customer Experience
TR	Technical Report
VSM	Volume SnapMirror
WAFL	Write Anywhere File Layout
WORM	Write Once Read Many



Trademarks and special notices

© International Business Machines 1994-2007. IBM, the IBM logo, System Storage, and other referenced IBM products and services are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. All rights reserved

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Network Appliance, the Network Appliance logo, Data ONTAP and FilerView, FlexVol, NearStore SnapMirror, SnapRestore, Snapshot, SnapVault and WAFL are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.