Technical report:

Symantec NetBackup and IBM System Storage N series

Disk-based backup solutions deployment and implementation guide

Document NS3429-0

January 22, 2008

Table of contents

Abstract

This deployment guide describes optimized Symantec NetBackup disk-to-disk backup solutions using IBM N series storage systems as the destination and utilizing IBM System Storage N series with Snapshot, IBM System Storage N series with SnapVault and Single Instance Storage technologies. It is useful for sales and services field personnel requiring assistance in understanding, deploying and architecting a NetBackup disk-based backup solution with IBM N series storage systems in a customer environment. Please refer to the latest technical publications for specific updates on processes, command syntax, and the latest requirements, issues and limitations.

Introduction

Intended audience

This technical report is designed for System Engineers (SEs) and Professional Services Engineers (PSEs) who seek education on the joint solutions IBM® System Storage™ N Series and Symantec™ offer and who additionally may need to prepare for performing deployments in customer environments. It is most beneficial to those who are already familiar with IBM N series hardware and software and Symantec Veritas NetBackup™, although high-level sections are included that attempt to quickly educate the reader on the basic technologies involved. Those individuals at Symantec, IBM and in the channel all will benefit.

Purpose

The purpose of this paper is to present a guide for implementing joint IBM N series and Symantec NetBackup disk-based backup solutions, 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:

Provide detailed information to all interested parties,
Educate prior to performing deployments, and
Serve as a reference for resolving issues which could arise.

(It has also somewhat become a catch-all repository for relevant technical topics not covered elsewhere.)

This document is not:

A sales guide (although some high-level thoughts are covered in the *Business Applications* and *Overall Benefits* sections),
A competitive comparison, or
A complete product design document.

Prerequisites and assumptions

For the details and procedures described in this paper 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 platforms and products, particularly in the area of data protection.
The reader has general knowledge of NetBackup.

Technology primer

As the joint solutions described herein involve both NetBackup and IBM System Storage N series hardware with the Data ONTAP® operating system and various complementary software 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.

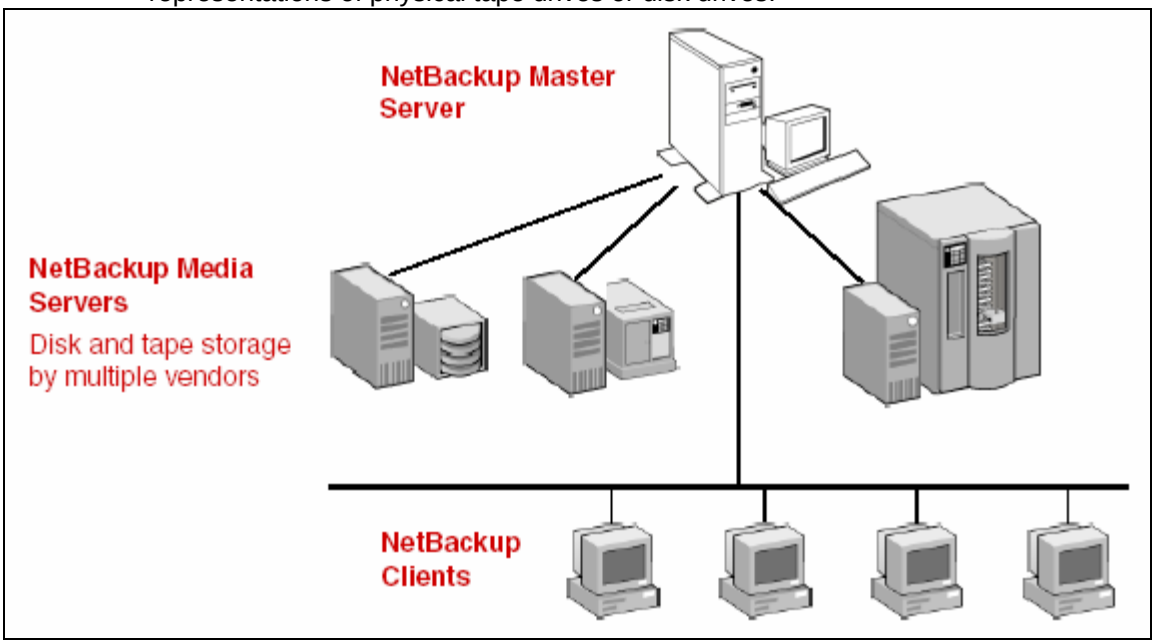

Figure 1. NetBackup Storage Domain.

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.

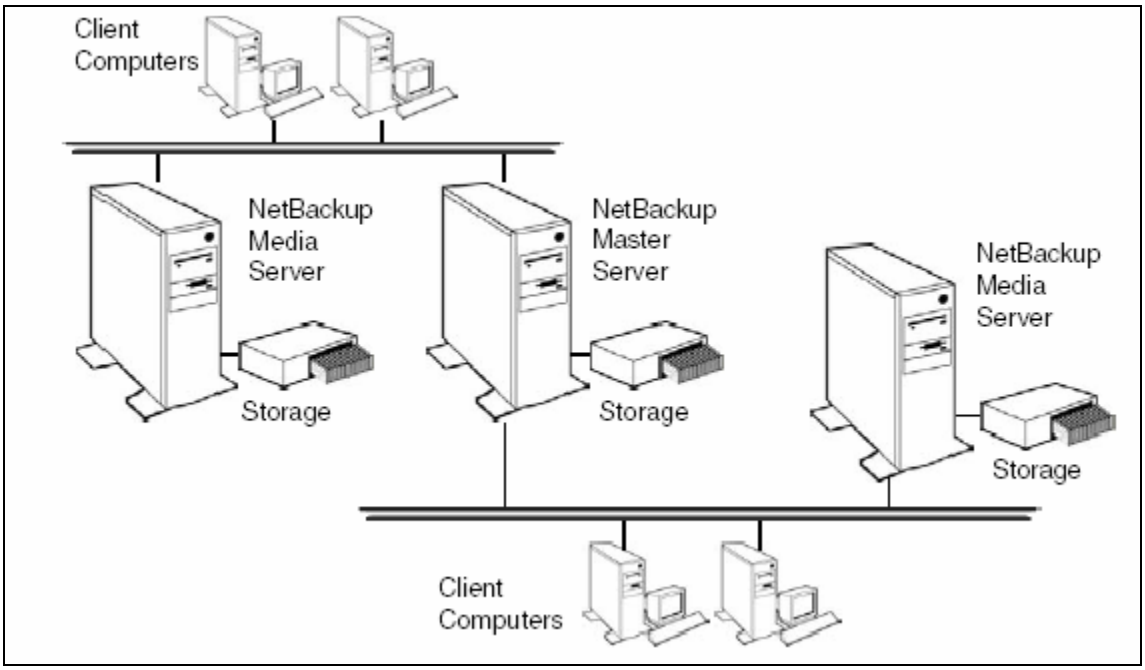

Figure 2. NetBackup Architecture.

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 graphical user interface (GUI) to browse and then 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. To send backups to a storage device, the administrator needs to define storage units using the Storage Units utility. There are three types of storage units:

Media Manager (MM) Storage Units: This type encompasses the tape robots, standalone tape drives, and optical disk devices—all of which are under the control of MM. MM controls the allocation and mounting of media (called Volumes) in the storage devices.

Network Data Management Protocol (NDMP) Storage Units: NDMP storage units are controlled by MM but attach to NDMP hosts and require installation of the NetBackup for NDMP option.

Disk Storage Units: A disk type storage unit consists of a directory on a disk that stores data. NetBackup permits an unlimited number of disk storage units. There are three types:

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

IBM System Storage N series with NearStore®, used for any heterogeneous client. IBM N series featuring NearStore appears as a GUI selection when the NetBackup Disk Optimization Option is licensed.

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

Any disk storage unit—except SnapVault—can be used for disk staging. In disk staging, the storage unit provides the first storage location in a two-stage process. In this process, client data is backed up to a disk staging storage unit, then, in the second stage, the data is relocated to another storage unit.

As this technical report covers the topic of joint disk-to-disk backup solutions, it will focus almost exclusively on disk storage units.

Disk staging

Disk staging provides a method for administrators to create images on disk initially, then later copy the images to another media type (as determined in the disk staging schedule). The media type for the final destination is typically tape, but could be disk.

This two-stage process allows the NetBackup administrator to leverage the advantages of disk-based backups in the near term, while preserving the advantages of tape-based backups for long term.

Disk staging meets the following objectives:

Allows backups when tape drives are scarce.
Allows for faster restores from disk.
Facilitates streaming to tape without image multiplexing.

Disk staging is conducted in two separate operations:

A backup creates an image on the storage unit acting as the disk staging storage unit.
A relocation schedule determines when the image from the disk staging storage unit should be relocated to the destination storage unit.

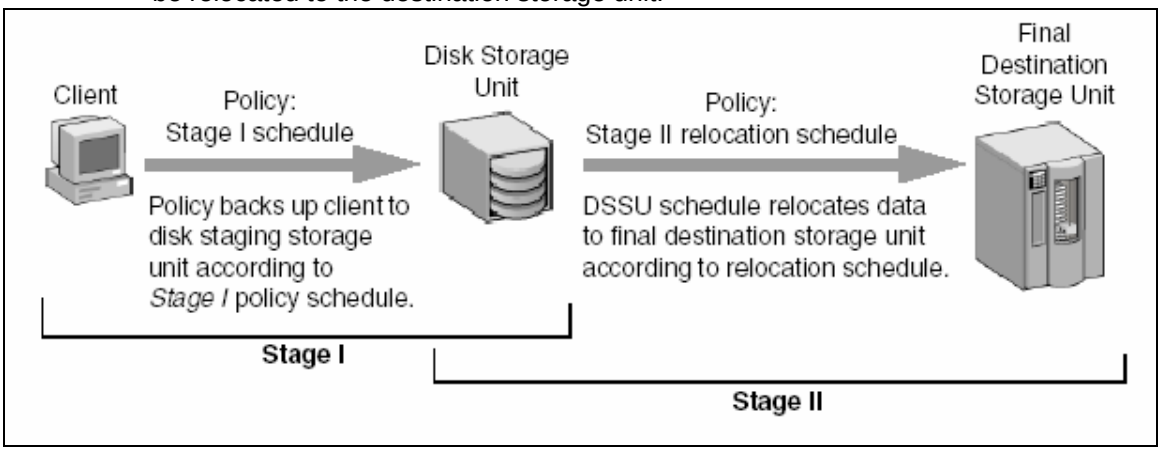

Figure 3. NetBackup Disk Staging Storage Unit.

The image continues to exist on both the disk staging storage unit and the destination storage unit. File restores are done from the disk staging storage unit copy, while the destination storage unit copy can be considered the long term copy.

The image copy continues to exist on the disk staging storage unit until either the copy expires based on the copy's retention period, or until another Stage I process needs space on the disk storage unit.

When a Stage I process detects a full disk staging storage unit, it pauses the backup, finds the oldest image that has been successfully copied to the destination storage unit, and expires this image copy.

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 storage system. There are many other features and options not discussed here; the items covered in this section are those pertinent to the integrated solutions 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. At the time of this report writing, the latest version, 7.1, has new features such as integration with NetBackup 6.0, LockVault™ and multiple internet small computer system interface (iSCSI) connections.

WAFL

IBM System Storage N series with WAFL® (Write Anywhere File Layout) is, essentially, the file system that exists on a IBM N series storage system. 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.

Snapshot and SnapRestore

A snapshot copy is a locally retained point-in-time image of data. 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 technology also makes it highly scalable. A snapshot copy takes only a few seconds to create—typically less than one second, regardless of the size of the volume or the level of activity on the IBM N series storage system. 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. Since 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, since 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. Since the integrity 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.

SnapVault

SnapVault protects data on a SnapVault primary system by maintaining a number of read-only versions of that data on a SnapVault secondary system. First, a complete copy of the data set is pulled across the network to the SnapVault secondary. This initial, or baseline, transfer may take some time to complete, as it is duplicating the entire source data set on the secondary much like a level-zero backup to tape. Each subsequent backup transfers only the data blocks that have changed since the previous backup.

When the initial full backup is performed, the SnapVault secondary stores the data in a WAFL file system, and creates a snapshot image of that data. A snapshot copy is a read-only, point-in-time version of a data set. A new snapshot copy is created each time a backup is performed, and a large number of snapshot copies can be maintained according to a schedule configured by the backup administrator. Each snapshot copy consumes an amount of disk space equal to the differences between it and the previous snapshot copy.

SnapVault ensures backup reliability by storing the backups on disk in a WAFL file system; they are protected by redundant arrays of inexpensive disks (RAID), block checksums, and periodic disk scrubs, just like all other data on an IBM N series storage system. Restores are simple because each incremental backup is represented by a snapshot copy, which is a point-in-time copy of the entire data set, and can be restored in a single operation. For these reasons, only the incremental changes to a data set ever need to be backed up once the initial baseline copy has completed. This reduces load on the source, network bandwidth consumption, and overall media costs.

One of the unique benefits of SnapVault is that users do not require special software or privileges to perform a restore of their own data. Any users who wish to perform a restore of their own data may do so without the intervention of a system administrator, saving time and money. Restoring a file from a SnapVault backup is simple. Just as the original file was accessed via a network file system (NFS) mount or commom internet file system (CIFS) share, the SnapVault secondary may be configured with NFS exports and CIFS shares. So long as the destination qtrees are accessible to the users, restoring data from the SnapVault secondary is as simple as copying from a local snapshot copy.

Restoration of an entire data set can be performed the same way if the user has appropriate access rights, however SnapVault provides a simple interface to restore an entire data set from a selected snapshot copy using the `snapvault restore` command. Details on syntax and procedures for performing such a restore can be found in applicable Data ONTAP guides.

Managing Via IBM System Storage N series with FilerView®

All IBM N series storage systems support an exhaustive Command Line Interface (CLI) to manipulate every aspect of it. As the perhaps half of the audience of this document are not experienced IBM N series administrators, most operations are demonstrated using FilerView, the web-based GUI for managing all IBM N series devices.

NOTE: There are some items which can only be accessed via the CLI and they will be shown.

Access FilerView by typing the following:

http://filer-name/na_admin

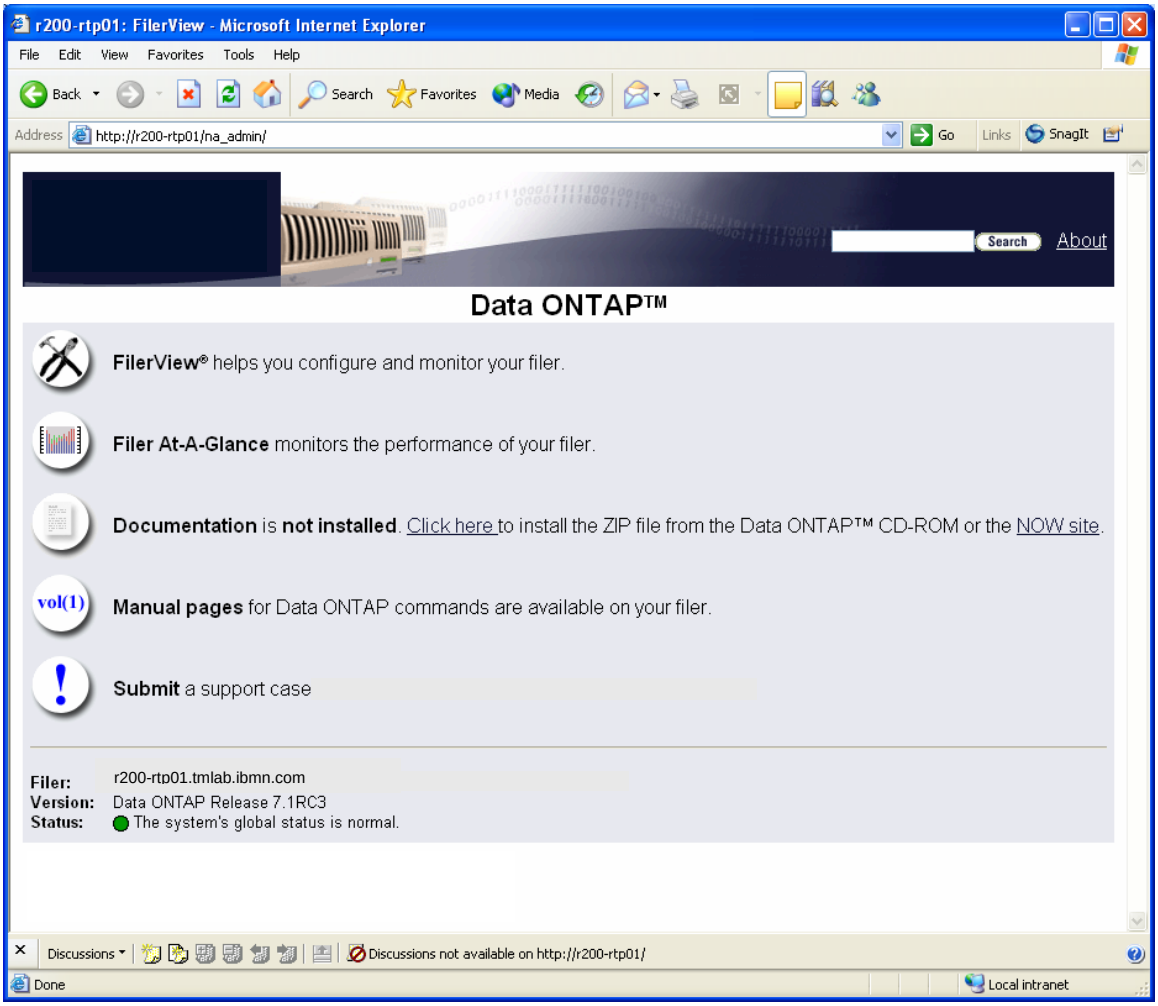

Figure 4. Web-Based Management of IBM N series Storage System.

There is plenty of reference and reading material available there, but to begin managing the box simply click on the “FilerView” link. You’ll first be prompted to login:

Figure 5. Login to FilerView.

FilerView opens and from here you can perform the rest of the operations described throughout this document.

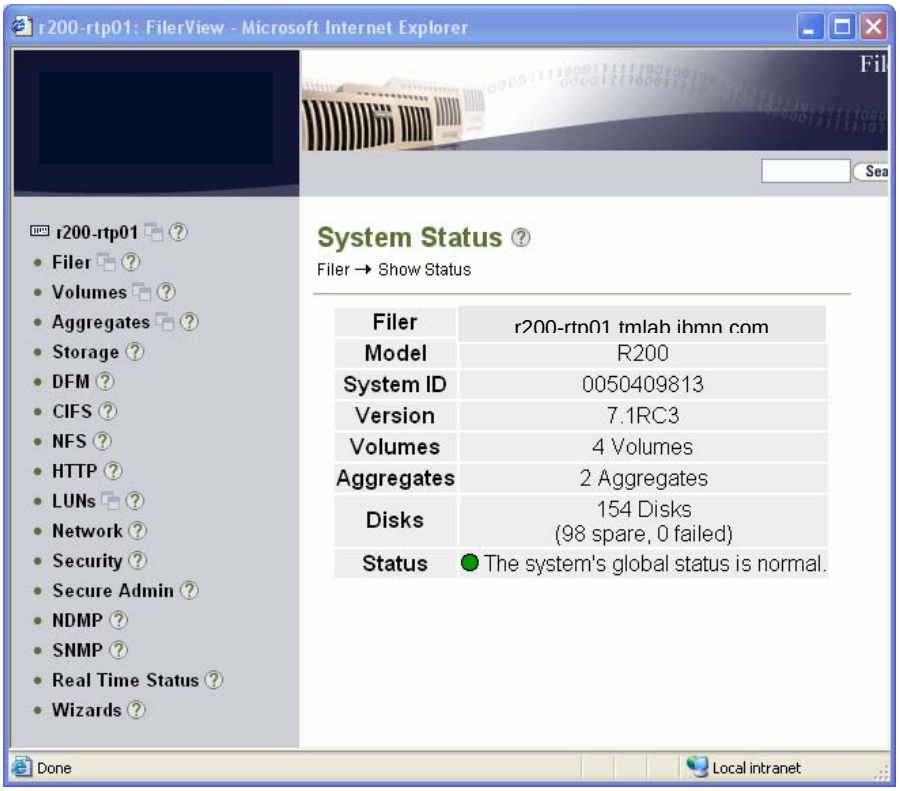

Figure 6. FilerView Web-Based Management GUI.

Solution overviews

This section provides an overview of the IBM N series and Symantec disk-to-disk backup solutions and the customer challenges that they address.

Customer backup challenges

The following are some of the issues facing virtually all customers today:

Unable to make backup windows. Customer data continues to grow at phenomenal rates making it more and more difficult to backup in existing backup windows. The problem is compounded with more applications going 7x24 and more stringent Service Level Agreements (SLAs) being put in place – the result being decreasing backup windows.

Unmanageable tape infrastructure "sprawl." With the constantly and rapidly growing amount of data, administrators are naturally and necessarily throwing more and more tape drives into their backup environments. As customers manage tens and hundreds of terabytes, and approach the petabyte mark, this traditional brute-force method of backups falls short.

Lack of scalability and manageability in traditional backup/recovery processes.

Recovery speeds are slow, labor intensive, absorb administrator resources and immobilize users.

Inability to recover business-critical data, due to inadequate backup processes and faulty backup media that cause administrators to recover stale data, no data or the wrong type of data.

Features and benefits

The IBM N series and Symantec disk-to-disk backup solutions discussed in this report solve the above pain points by delivering the following benefits:

Better manageability and simplified data protection. Users can now perform highly scalable, centralized backups on disk and eliminate redundant data, infrastructure and management.

More reliable, rapid recovery. Users can quickly create frequently updated, online disk backups, via integrated point-in-time snapshot copies and added intelligence about the backup data.

Significant backup-related cost savings. Centralized backup and consolidation reduce administrator overhead and maximize the use of existing tape infrastructures. Increased scalability and elimination of redundant backup data minimize the amount of disk backup storage required by up to 90%.

Increased business value via optimized recovery. Productivity increases on the part of both users and administrators, thanks to rapid recovery or up-to-date data, and simple user-driven or administrator-driven recovery options.

The solutions

Integrated IBM N series and Symantec solutions enable effective data protection across the enterprise.

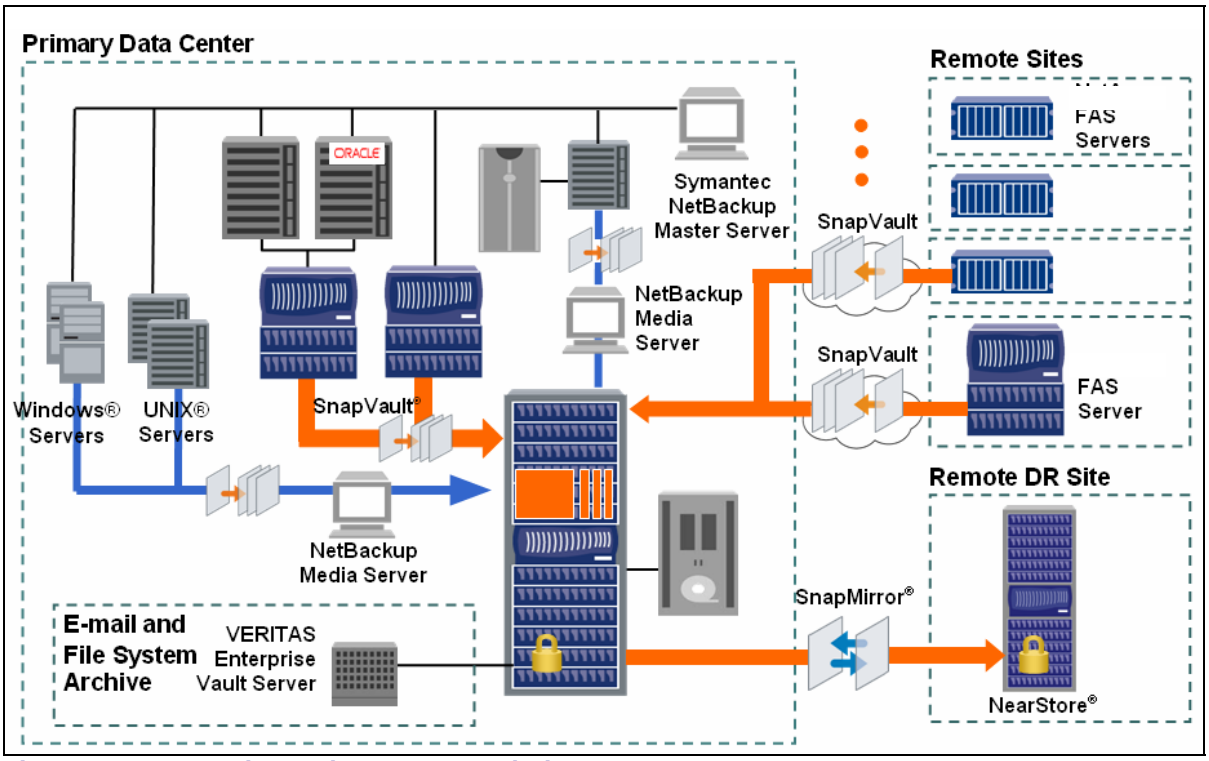

Figure 7. IBM N series and Symantec Solutions.

There are three jointly engineered and integrated solutions:

NSM --- NetBackup Snapshot Management
NetBackup integration with Snapshot and SnapRestore
Protects IBM N series primary storage systems

NSVM --- NetBackup SnapVault Management
NetBackup integration with SnapVault
Protects IBM N series primary storage systems

SV-NBU --- SnapVault for NetBackup
(Previously referred to as "**NDO** --- NetBackup Disk Storage Unit (DSU) Optimizer")
Deep NetBackup integration with IBM N series storage system secondary storage and de-duplication technology
Protects "any" primary storage.

These are described in the following subsections and in the remainder of this document the various sections will be broken up by each solution.

NetBackup Snapshot Management (NSM)

This solution integrates Snapshot and SnapRestore with NetBackup, and is for data which is NFS-mounted and/or CIFS-mapped from IBM N series primary storage. It offers nearly instantaneous backup of IBM N series storage system primary data, and restores which range from nearly instantaneous to very fast.

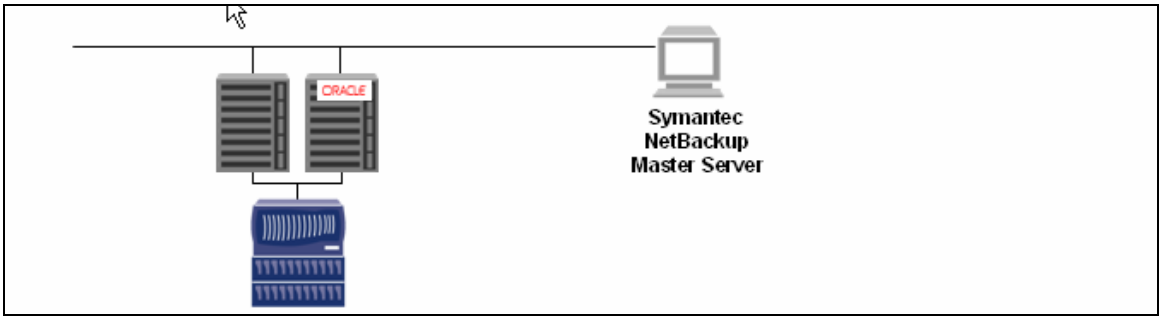

Figure 8. Snapshot and SnapRestore Integration (NSM) Solution.

Integration overview:

- IBM N series primary storage system
- Snapshot and SnapRestore Integration
- Snapshot copy creation and management – allow NetBackup to initiate Snapshot technology mechanism through a standard scheduled policy
- Delivered in NetBackup 5.1 and Data ONTAP 7.1+
- Tightly coupled integration with NetBackup / Advanced Client
- Supports Microsoft® Windows® with CIFS and Solaris™ with NFS file systems
- NetBackup Oracle® Database Agent integrated as well on Solaris
- NetBackup will catalog the image (instant recovery) allowing restores to be driven via the NetBackup backup / restore interface
- Client-direct restores via NFS and CIFS can also be done if the administrator allows it.

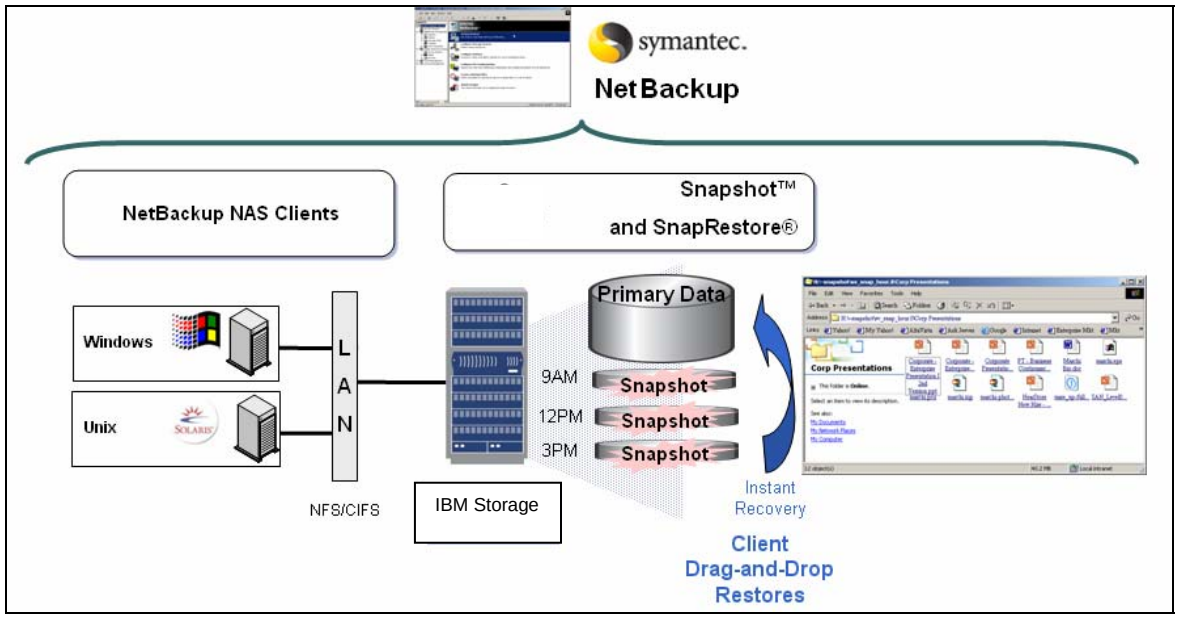

Figure 9. NSM Instant Recovery via SnapRestore.

Instant recovery from user error:

- Recovery from local disk snapshot copies
- Rollback to point-in-time
- Optimized for space
- Block-level changes
- Simplified restores
- End user view of data
- File, volume, or file system
- Single management interface (NetBackup).

IBM N series with Snapshot software:

- Point-in-time images
- Locally retained, read-only, point-in-time image of a volume
- Up to 255 snapshot images per volume
- No performance overhead on the application or storage system
- Security of files retained in snapshot read-only images
- Facilitates frequent, low-impact, user-recoverable backups of files, directories, LUNs, and application data
- Provides near-instantaneous, secure, user-managed restores.

IBM N series with SnapRestore:

- Tape-less application data recovery and testing
- Revert from a snapshot copy to the active file system
- Ability to rollback to a previous consistent state of the file system within minutes
- Single-file restores
- Intra-storage system operation is near-instantaneous
- Especially useful to recover from a corrupted database (Oracle)
- Much reduced dependency on tape
- Useful for software testing situations requiring frequent returns to a baseline state.

While an incredibly fast and intuitive backup solution, the one downside is that the backup data remains on same spindles as primary data.

NetBackup SnapVault Management (NSVM)

This solution takes the NSM solution one step further by integrating SnapVault with NetBackup. The data is still NFS-mounted or CIFS-mapped from IBM N series primary storage, but now the backup data additionally resides on an IBM N series secondary storage device.

Backups thus reside in a separate storage platform / location, so the solution is great for both data centers and remote sites.

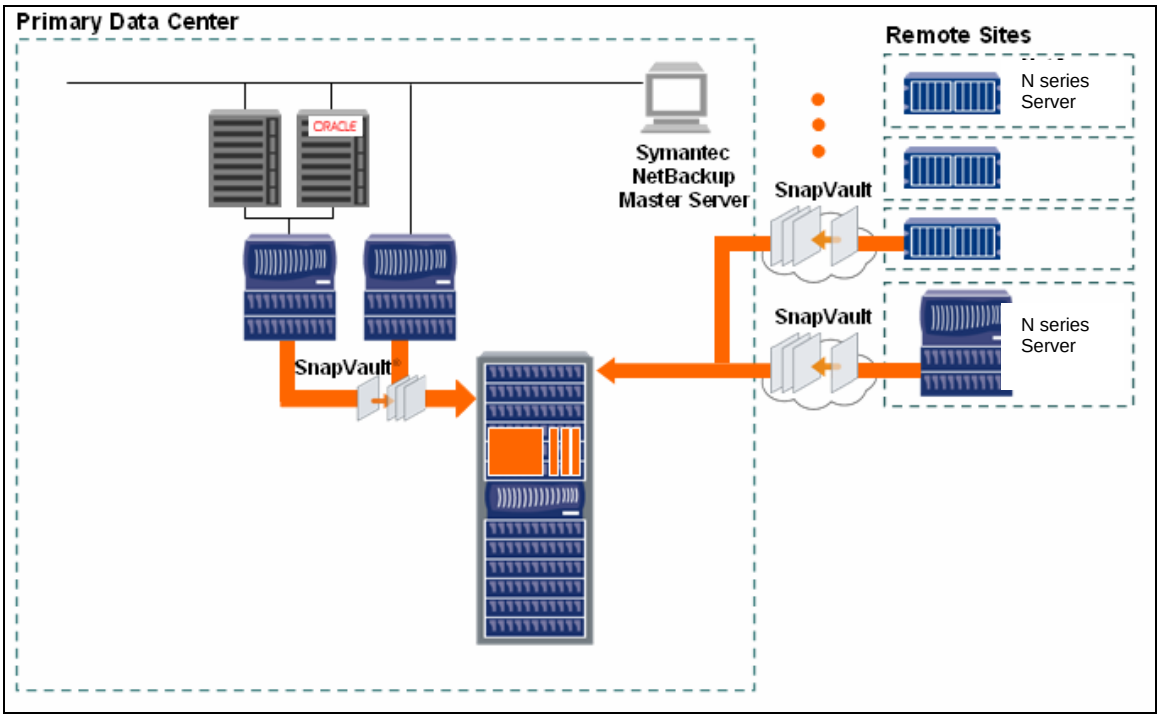

Figure 10. SnapVault Integration (NSVM) Solution.

Integration overview:

- IBM N series primary storage system
- SnapVault integration
- Allows NetBackup to initiate a SnapVault backup = a snapshot copy is created on the primary IBM N series storage system and then is replicated to an IBM N series secondary storage system.
- Fast backups and restores, as only changed blocks are transferred = “data de-duplication”
- Delivered in NetBackup 6.0 and Data ONTAP 7.1+
- Tightly coupled integration with NetBackup / Advanced Client
- Supports Windows with CIFS and Solaris with NFS file systems
- NetBackup Oracle Database Agent integrated as well
- Backups in the form of snapshot copies on the primary and secondary IBM N series storage systems are managed (creation, retention, deletion) with NetBackup
- Consolidated snapshot copies from distributed NAS filers can reside on same IBM N series secondary storage system
- Client-direct restores via NFS and/or CIFS can also be accomplished if the administrator wants to allow it.

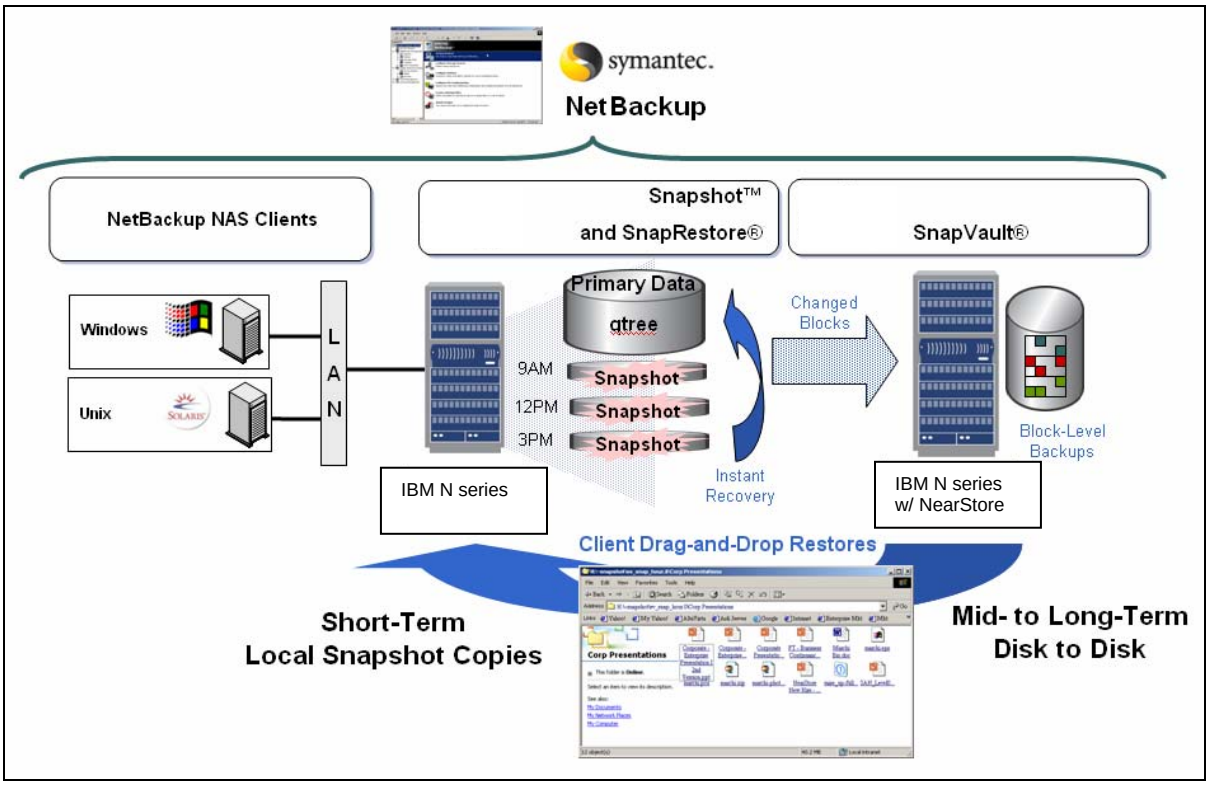

Figure 11. NSVM Changed-Block Backups.

After initial full backup of all systems is done to IBM N series secondary storage system, all subsequent backups are incremental backups. This means that only those 4K blocks that changed since the last backup will be sent, which means backups are faster (unless all of the data changed).

It also means that only changed 4K blocks are stored on disk. This dramatically reduces the amount of information stored on disk when compared to traditional backup where entire files are stored in an incremental backup anytime a file changes. For IBM N series to IBM N series backup, only changed blocks are sent across the network and stored on the secondary storage system.

What is nice about this approach is that the data that is stored is in file format, including all incremental backups and each incremental backup can be viewed as a full backup image by any administrator. So, if you want to go to the backup that was 4 hours ago or 4 days ago you can quickly locate that backup and have a full view into what the entire environment looked like at that time. You don't need to go back step-by-step to get a view or locate the information you need.

This "file system" view of backups can be accessed by end-users as well if administrators choose to do so, but permissions and security of the original data are maintained in the backup.

SnapVault for NetBackup (SV-NBU)

This solution integrates the IBM N series storage system as an optimized backup repository for heterogeneous (not IBM N series) storage.

Integrated (disk-based) server protection:

NetBackup leverages an IBM N series storage system as an optimized disk target (disk storage unit).

Enhanced Data Transfer Protocol – Symantec and IBM N series have created a modified network protocol to increase data transfer performance between the NetBackup Media Server and the IBM N series storage system. The IP-based protocol allows synchronous data transfers to occur without delays associated with interlocked packet acknowledgement that comes with NFS and CIFS.

Reduced Storage with Redundant Data Elimination – Backups written to an IBM N series storage unit may utilize less disk space when compared to traditional disk storage units. After an initial client backup is performed, the WAFL file system will save only changed blocks when subsequent backups are performed for the same client. The quantity of disk space saved using data de-duplication technology will vary depending on client data change rates.

All NetBackup clients are supported.

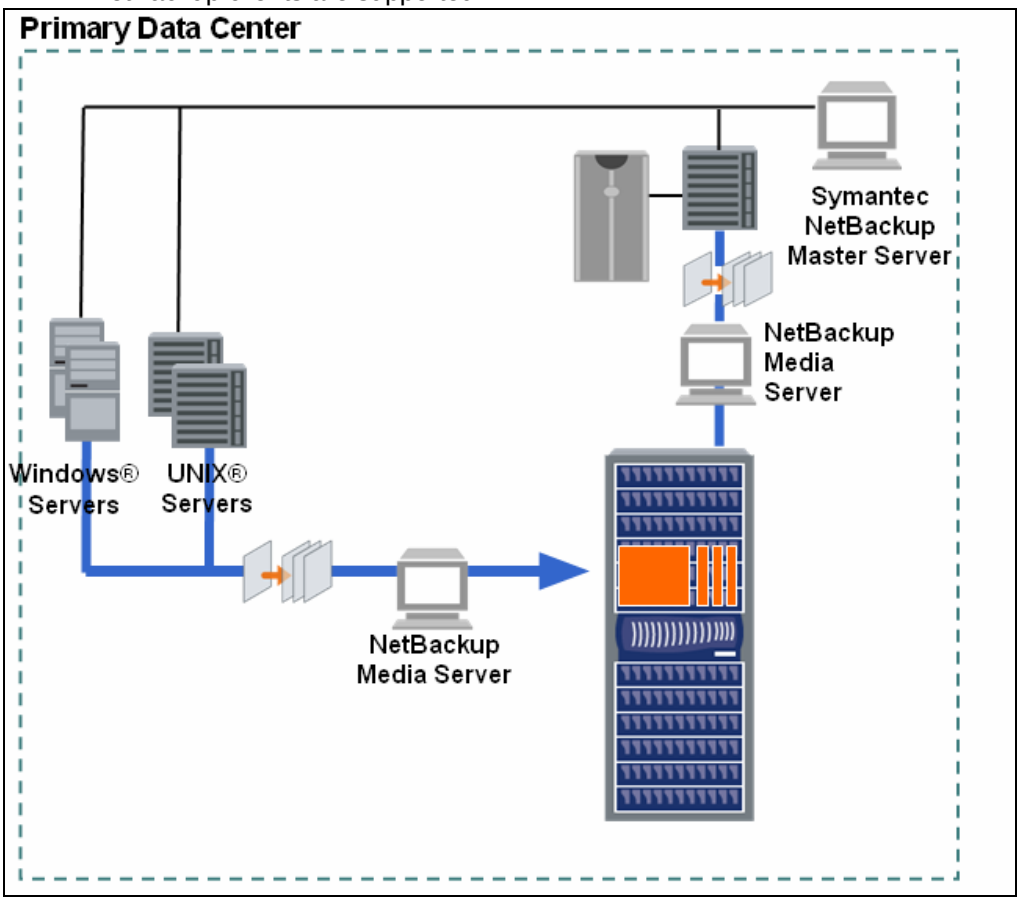

Figure 12. SnapVault for NetBackup (SV-NBU) Solution.

Integration overview:

Most seamless to integrate… it's just a new storage unit to point policies to
Heterogeneous primary storage
Integrates NearStore feature Disk Storage Unit
Delivered in NetBackup 6.0 and Data ONTAP 7.1+
Tightly coupled integration with NetBackup
Supports all NetBackup clients[1]
Proprietary streaming protocol between Media Server and IBM N series storage system
IBM N series storage system is able to inspect the files within the standard TAR image
Creates a WAFL file system structure from backup data stream
Through the WAFL file system, the IBM N series storage system will perform block level change backup and single instance storage
To NetBackup the backup on the IBM N series system looks like a standard NetBackup backup allowing all the normal NetBackup operations to be performed (duplication, synthetics, vaulting etc)
The IBM N series storage system can also act as staging area - Disk Staging Storage Unit (DSSU) - for backup before backups go to a final storage unit.

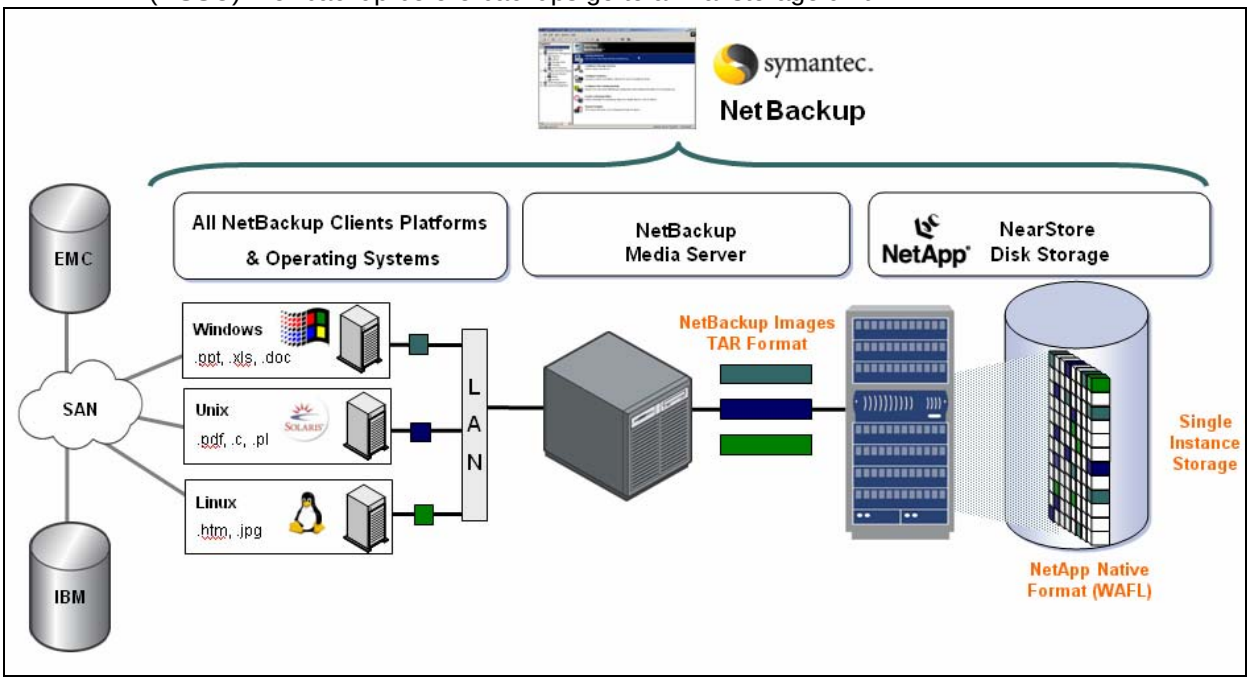

Figure 13. SV-NBU Space-Optimized Backups.

[1] All will work with the solutions, but there are some caveats regarding exactly which ones enjoy the benefits the solution has to offer. See the "Application Specific Integration" section of this document for more information.

Backup streams leave the client through a NetBackup Media Server. At the IBM N series storage system (through Data ONTAP), the NetBackup images are translated to native WAFL format. This enables single instance store for reduced storage requirements.

Single Instance Store (SIS) summary: Occurs per NetBackup client. The SIS processing occurs on the IBM N series storage system. This means that all of the data (files) will be sent from the client to the NetBackup Media Server and then to the IBM N series storage system, but only changed blocks will be saved by the IBM N series storage system. This applies to incremental and synthetic backups.

Data ONTAP 7.1 will perform SIS such that on a per-client basis, any NetBackup images that are converted to WAFL will only have changed blocks saved. If the block has already been saved, then a map of pointers will keep track of that information. This means that although a NetBackup incremental backup will send changed files to the IBM N series storage system, only the changed blocks will be stored on disk. In addition, if another full backup is done, although the entire backup will be sent to the IBM N series storage system, only the changed blocks will be saved. This also applies to a synthetic full backup. The synthetic will be generated, but then IBM N series will remove duplicated blocks.

For more detailed discussions on data movement, WAFL file system creation and Single Instance Storage, see the SV-NBU section in the Concepts / How It Works section of this document.

Concepts: How it works

This section presents technical information on each solution and how it works. It isn't really necessary information for installing/deploying these solutions, but is good (and highly recommended) reading.

NSM

NSM tightly, seamlessly and invisibly integrates IBM N series with Snapshot technology with NetBackup. For additional, specific reports on the integration, refer to "IBM System Storage N series with NearStore and SnapVault—Disk-Based Backup & Recovery" and "Symantec NetBackup and IBM System Storage N series with Snapshot and SnapRestore in an Oracle Environment—Integration for Backup and Recovery."

The remainder of this section will provide the reader with a better understanding of how Snapshot technology works on an IBM N series storage system.

Snapshot copies are a benefit of the WAFL write anywhere approach. A snapshot copy is an online, read-only copy of the entire file system. Typically, a snapshot copy only takes a few seconds to create – usually less than one second, regardless of the volume's size or the level of activity on the IBM N series storage system. 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 didn't exist. Meanwhile, the snapshot version of the data remains completely stable. A snapshot copy incurs no performance overhead.

A snapshot copy can be used for online backup capability, allowing users to recover their own files. A snapshot copy also simplifies backup; since a snapshot copy is a read-only copy of the entire file system, it allows self-consistent backup from an active system. Instead of taking the system offline, the system administrator can make a backup to disk or tape of a recently created snapshot copy.

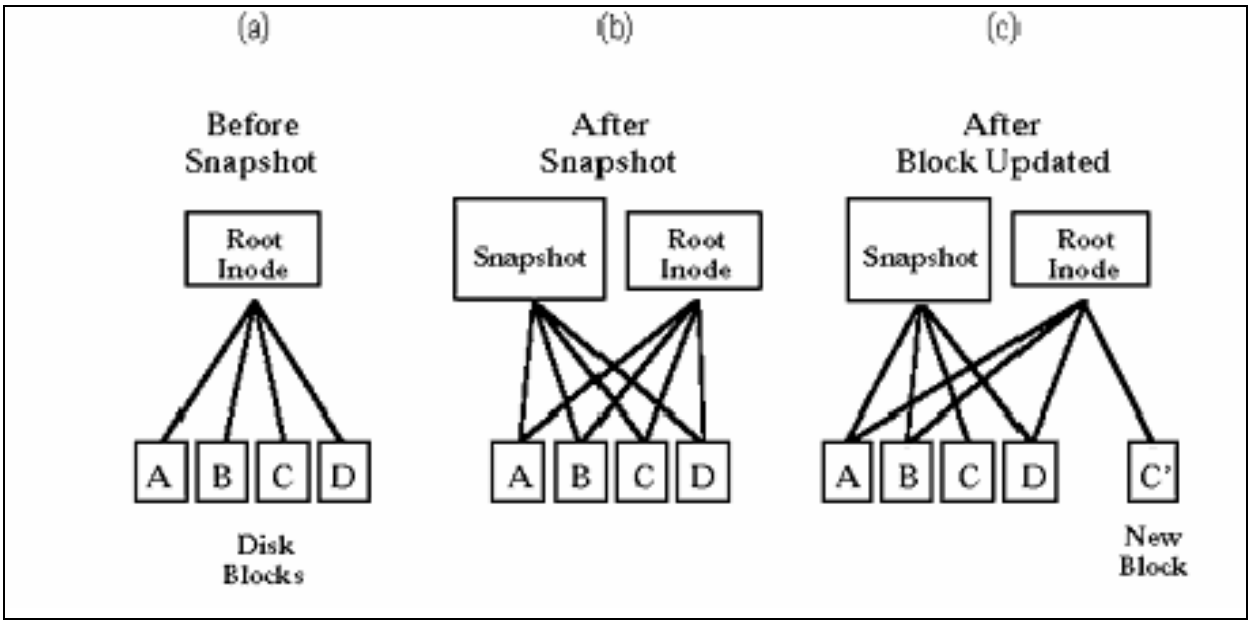

Figure 14. Snapshot Copy Creation.

The previous figure shows how a snapshot copy works. A WAFL file system can be thought of as a tree of blocks rooted by a data structure that describes the inode file. The inode file in turn contains the inodes that describe the rest of the files in the system, including meta-data files such as the free block bitmap and the free inode bitmap. (a) is a simplified representation of a complete file system with a root data structure at the top pointing to disk blocks.

WAFL creates a new snapshot copy by making a duplicate copy of the root data structure, as shown in (b). Since the root data structure is only 192 bytes, and since no other data blocks need to be copied on disk, a new snapshot copy does not actually consume any additional disk space until a user deletes or modifies data in the active file system. WAFL creates a snapshot copy in just a few seconds. (c) shows what happens when a user modifies data block C. WAFL writes the new data to block C' on disk, and changes the root structure for the active file system to point to the new block. The snapshot copy still references the original block C, which is unmodified on disk.

The storage appliance can keep up to 255 snapshot copies per volume. The copies can be created manually or the storage appliance can create and delete them automatically according to a user-defined schedule. In the case of NSM, snapshot creation and deletion is handled by NetBackup.

NSVM

NSVM tightly, seamlessly and invisibly integrates IBM N series with SnapVault technology with NetBackup. For additional, specific information on SnapVault, refer to the SnapVault deployment and implementation, configuration, and best-practice performances guides.

How does the NSVM SnapVault function differently in this solution from “regular” SnapVault?

Internally, essentially rewrites SnapVault snapshot management.

“Native” SnapVault updates all qtrees on a secondary volume, and then creates a snapshot copy. All primaries for the volume are contacted at the same time.

NSVM updates qtrees individually, with dataset names defined according to client view. Snapshot copy usage is individual but consumption is minimized via snapshot management (see the NSVM Snapshot Management section below).

Qtrees

A difference between NSM and NSVM is that NSVM works ONLY with qtrees (‘quota trees’).

A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within either a traditional volume or a flexible volume.

You can have a maximum of 4,995 qtrees on any volume.

Every qtree and volume has a security style setting. This setting determines whether files in that qtree or volume can use Windows NT® or UNIX® security.

Using quotas, you can apply a tree quota to a qtree: the qtree is similar to a disk partition, except that you can change its size at any time.

NSVM Snapshot Management

This section describes how snapshot management, also known as snapshot coalescing, works. It will help to describe how many snapshot copies are really consumed after numerous backups occur and will be key when it comes to understanding just how many backups you can do.

Start with a volume which contains qtree1 (Q1) which is /home. After running the first backup of a policy for /home, Snapshot1 (S1) will be created as shown below.

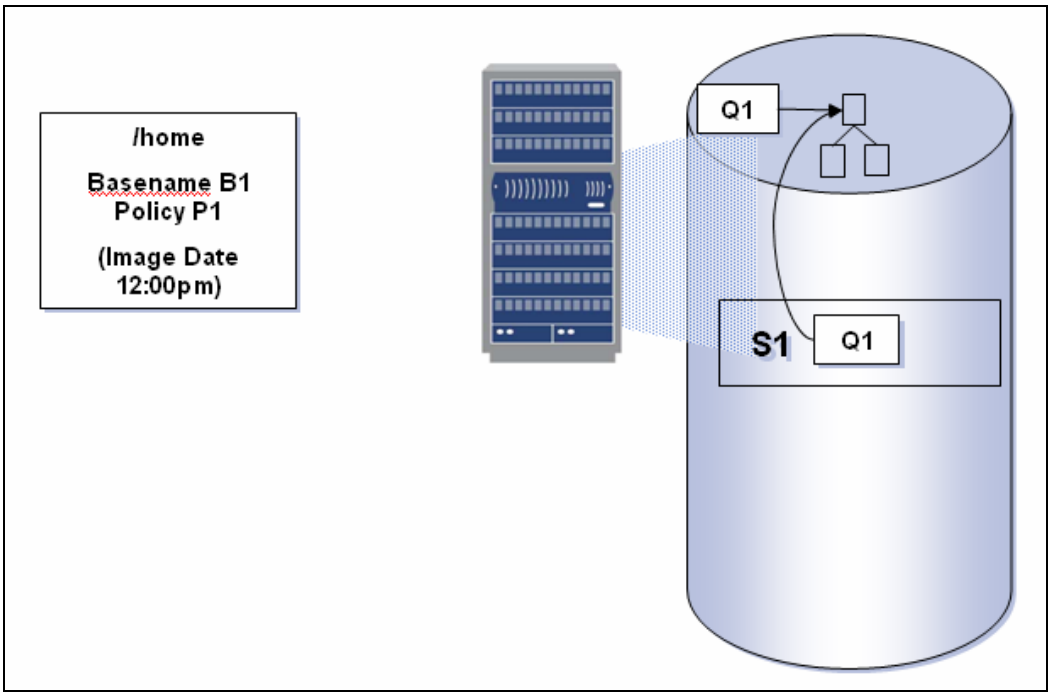

Figure 15. NSVM Snapshot management (A).

Next another qtree, Q2, for /extra is created on the same volume. After running the first backup of a policy for /extra, snapshot1 (S2) will be created as shown below.

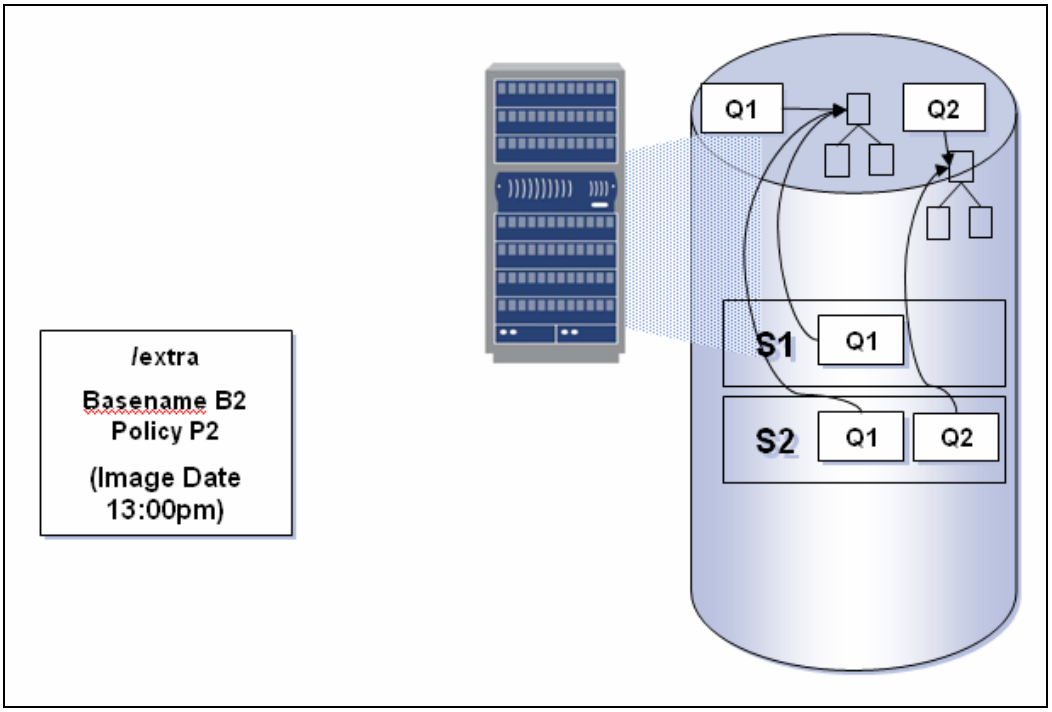

Figure 16. NSVM Snapshot Management (B).

Snapshot S1 can safely be removed.

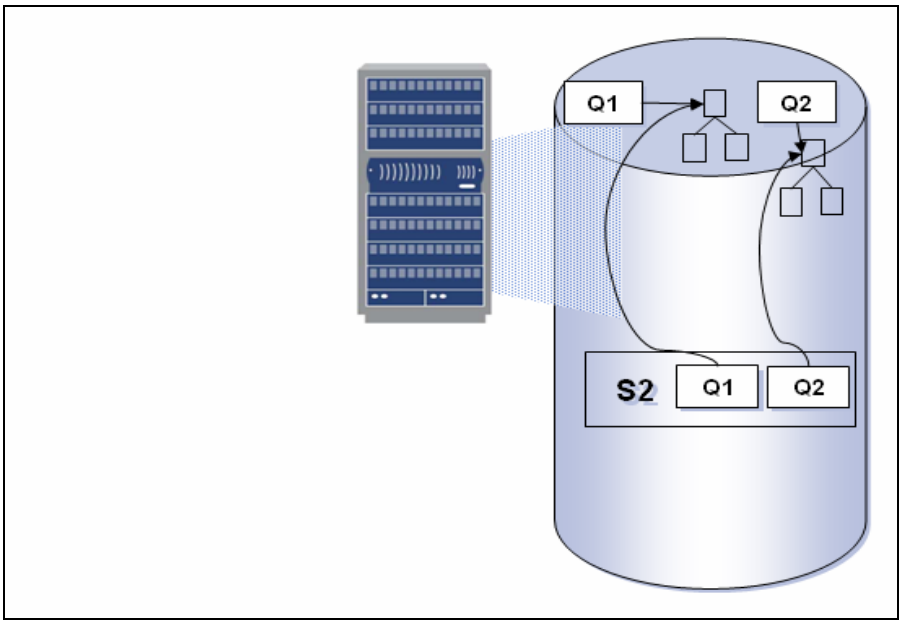

Figure 17. NSVM Snapshot Management (C).

Client modifies a file in /home.

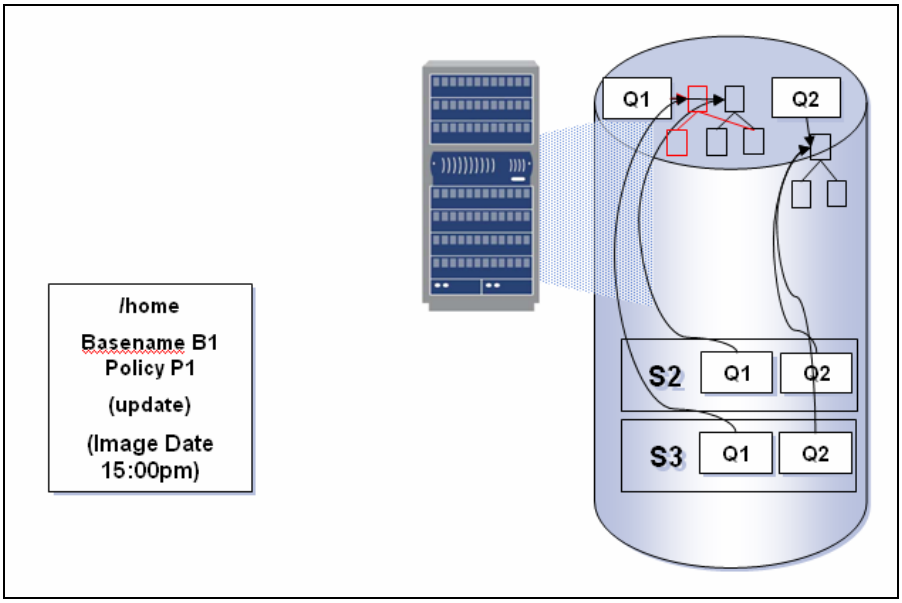

Figure 18. NSVM Snapshot Management (D).

At this point we need to keep both S2 and S3.

Based on the above, now we look at a real backup schedule for several qtrees. For example:

/vol/280er/users (perf280er /users): 0500, 1300, 2100
/vol/c38/home (fsr-c38 /home): 0500, 1200, 1700, 2300
/vol/c38/extra (fsr-c38 /extra): 1300, 2000
Legend: X = updated and snapshot copy created
tmp = updated, but qtree “moved”
>X = snapshot copy replaces a “tmp” one

The following tables demonstrate which snapshot copies are maintained for the above policy schedules.

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users							
/home							
/extra							

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X						
/home	X						
/extra							

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X						
/home	X	X					
/extra							

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X		X				
/home	X	X					
/extra			X				

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X		X				
/home	X	tmp	>X				
/extra			X				

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X		X				
/home	X	tmp	>X	X			
/extra			X				

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X		X				
/home	X	tmp	>X	X			
/extra			X		X		

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X		X				
/home	X	tmp	>X	tmp	>X		
/extra			X		X		

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X		X			X	
/home	X	tmp	>X	tmp	>X		
/extra			X		X		

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X		X			X	
/home	X	tmp	>X	tmp	tmp	>X	
/extra			X		tmp	>X	

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X		X			X	
/home	X	tmp	>X	tmp	tmp	>X	X
/extra			X		tmp	>X	

Figure 19. Snapshot Demonstration (A).

What if we delete last night's 5pm backup of /home? Answer: it will coalesce, i.e.:

1 - Remove lock for /home at 2100.

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X		X			X	
/home	X	tmp	>X	tmp	tmp	deleted	X
/extra			X		tmp	>X	

Figure 20. Snapshot Demonstration (B).

2 - Move /users and /extra to 23:00.

Snapshot copy of 2100 deleted.

	05:00	12:00	13:00	17:00	20:00	21:00	23:00
/users	X		X			tmp	>X
/home	X	tmp	>X	tmp	tmp	deleted	X
/extra			X		tmp	tmp	>X

Figure 21. Snapshot Demonstration (C).

NSVM Snapshot specifics

When NSVM backups occur, snapshot copies result on the secondary storage system with obvious, but somewhat cryptic names.

```
# snap list
  %/used       %/total  date          name
----------  ----------  ------------  --------
  1% ( 0%)    0% ( 0%)  Dec 01 08:16
SnapVaultBackup.vol2.125.4163.20051201_081630_MYT (snapvault)
  1% ( 0%)    0% ( 0%)  Nov 30 11:55
SnapVaultBackup.vol2.111.3875.20051130_115539_MYT (snapvault)
```

The breakdown of the naming convention is as follows:

`SnapVaultBackup` identifies NSVM uniquely.
The next field is volume name.
The next two fields are `snapid` and `cpcount,` and represent how the snapshot is identified in ONTAP (internally). Both together are guaranteed to be unique per volume and will help identify to some extent which snapshot is the latest in case date/time stamp is wrong.
The next field is the date/time stamp.

SV-NBU

SV-NBU backups achieve de-duplication of data through Data ONTAP SIS technology. When discussing the functionality and features, one of the biggest challenges is to understand exactly what processing and data movement are occurring behind the scenes to deliver the benefits touted.

In NetBackup 6.0 storage units have the ability to transfer data as either “image” or “clearfile.” *Image* is used for all traditional NetBackup backups and basically just sends the standard NetBackup gtar file image to be stored in its entirety. *Clearfile* is used for SV-NBU backups and essentially exposes the file structure of the gtar data stream.

A NetBackup SV-NBU storage unit, which has a disk type of “NearStore,” will set the transfer mode to *clearfile* automatically when the storage unit is created via the GUI (by default the “Enable Block Sharing” check-box is selected). You can additionally examine the parameter using CLI on the master server with “`Block Sharing:  yes`” indicating that *Clearfile* data transfer is turned on.

```
sun280r-rtp03# ./bpstulist –L
…
Label:                NearStoreDSU1
Storage Unit Type:    Disk
Media Subtype:        NearStore (2)
Host Connection:      sun280r-rtp03
NearStore Server:     r200-rtp01
Concurrent Jobs:      1
On Demand Only:       yes
Path:                 "/vol/NearStoreDSU1"
Robot Type:           (not robotic)
Max Fragment Size:    524288
Max MPX:              1
Stage data:           no
Block Sharing:        yes
High Water Mark:      98
Ok On Root:           no
```

Dense volumes

Despite the introduction of less expensive ATA disk drives, one of the biggest challenges 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 the above stated 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 the storage space consumption. Basically within one file system tree there is the ability for multiple references to the same data block.

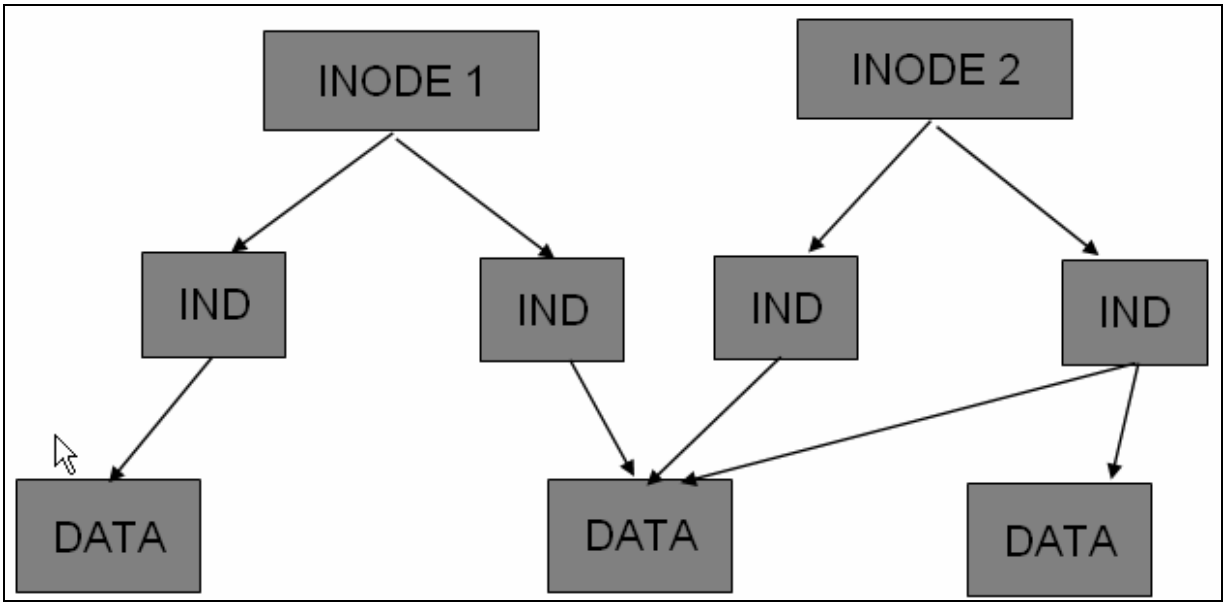

Figure 22. Dense Volumes.

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

Media Server to IBM N series storage system protocol

With NetBackup 6.0 and Data ONTAP 7.1, a new proprietary protocol is introduced to efficiently stream data communication between the NetBackup Media Server and the IBM N series storage system.

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

- Data stream
  - NetBackup gtar image
  - Aligned for our purposes
- Meta data
  - Set of RPCs
  - Directory contents
  - File location and length
  - UID/GID, permissions
  - Mtime, ctime, atime.

Processing of data from Media Server

Backups in “clearfile” mode have data aligned to a 4Kbyte boundary for unpacking and de-duplication. Once the data arrives at the IBM N series storage system, three threads handle alignment and unpacking:

directory creator - makes sure every directory gets created
file creator - makes sure every file is created (0 length files)
data writer – writes the data.

A checksum is kept for each block as it comes in from the NetBackup data stream and is maintained in a sorted list. The next time a backup data stream comes in, it is unpacked and compared to the list. If there is a match **for the same named file in the same directory** the data will be de-duplicated.

File system is unpacked into qtree: `/vol/MYVOL/nbu_BASE_POLICY_SUF`

e.g., fsr-pc7_C1_F1:win unpacks into /vol/vp3/nbu_fsr-pc7_C1_F1_win_0000.

The tape image is in a hidden subdirectory of `/vol/MYVOL/.nbu_image_data` (not visible in “`snapvault status`”).

Quick start: Installing and configuring

This section walks the reader through steps they need to accomplish to specifically make these solutions work. This document is not designed to make the reader a NetBackup or IBM N series guru. Although it discusses some basic set-up below, in general it assumes both the IBM N series storage system(s) and NetBackup are already installed and running.

Requirements overview

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

	NSM	NSVM	SV-NBU
IBM N series Hardware	Any N series, V-Series, or R-Series	Any platform combination SnapVault supports	NearStore R200, R150, R100 N5200, N5500
Data ONTAP	Data ONTAP 7.1	Primary: Data ONTAP 7.1 Secondary: Data ONTAP 7.1	Data ONTAP 7.1
IBM N series Software	Snapshot SnapRestore	Snapshot, SnapRestore2 Primary: sv_ontap_pri Secondary: sv_ontap_sec	sv_ontap_sec nearstore_option (for N5xxx)
NetBackup	NetBackup Enterprise 5.1MP2 NDMP Option Advanced Client Option3	NetBackup Enterprise 6.0 NDMP Option Advanced Client Option4 IBM N series SnapVault Option	NetBackup Enterprise 6.0 (NDMP Option – installed, but not licensed) Disk Optimization Option
Protocols	NFS, CIFS	NFS, CIFS	N/A
Client	Solaris Windows NetBackup 5.1	Solaris Windows NetBackup 6.0	(Whatever platform/OS NetBackup supports) NetBackup 5.0
Media Server	N/A	N/A (but Master must be NetBackup 6.0)	(Whatever platform/OS NetBackup supports) NetBackup 6.0
Applications	File Sharing Oracle (8i or later) on Solaris	File Sharing Oracle (8i or later) on Solaris	File Services

Table 1. Solution Requirements Overview.

The following subsections will provide a bit more substance on things which affect all three solutions.

NetBackup

All solutions need to make sure NetBackup Enterprise Server is installed on the Master Server and any Media Servers involved in the solution. Version 5.1 is required for NSM and 6.0 is required for NSVM and SV-NBU.

New to NetBackup 6.0 is that before installing the Master and Media Servers, you need to install Infrastructure Core Services (ICS).

---

[2] SnapRestore not needed, but recommended for NSM restores.
[3] Advanced Client requires NetBackup Enterprise Server.
[4] Advanced Client requires NetBackup Enterprise Server.

The NDMP for NetBackup Option must be installed for all three solutions, but only licensed for NSM and NSVM. If you have a Windows Master Server then NDMP is already installed. If you are running UNIX then the NDMP package must be separately installed.

The Advanced Client Option must be installed for NSM and NSVM, and optionally for SV-NBU if database backups are to be accomplished. If you have a Windows Master Server then Advanced Client is already installed. If you are running UNIX then the Advanced Client package must be separately installed.

When NetBackup is installed, you manage and configure it via the Administration Console. First you'll need to login to the master server (shown below); this can be accomplished from the Master Server, Media Server or, usually, an administrator's Client desktop.

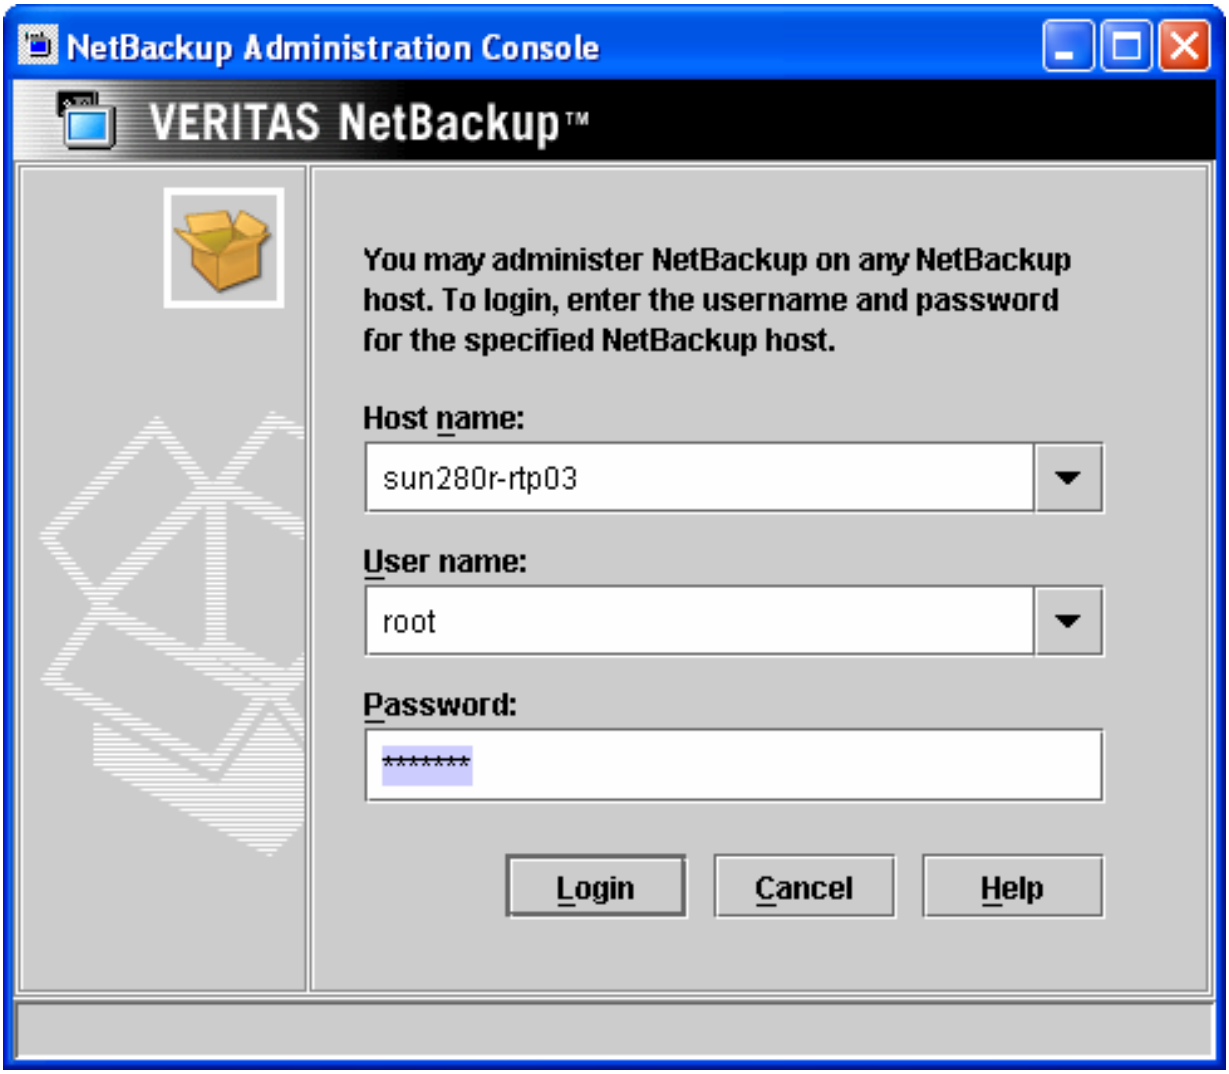

Figure 23. NetBackup Login.

After you've successfully logged in, you'll see the NetBackup Administration Console.

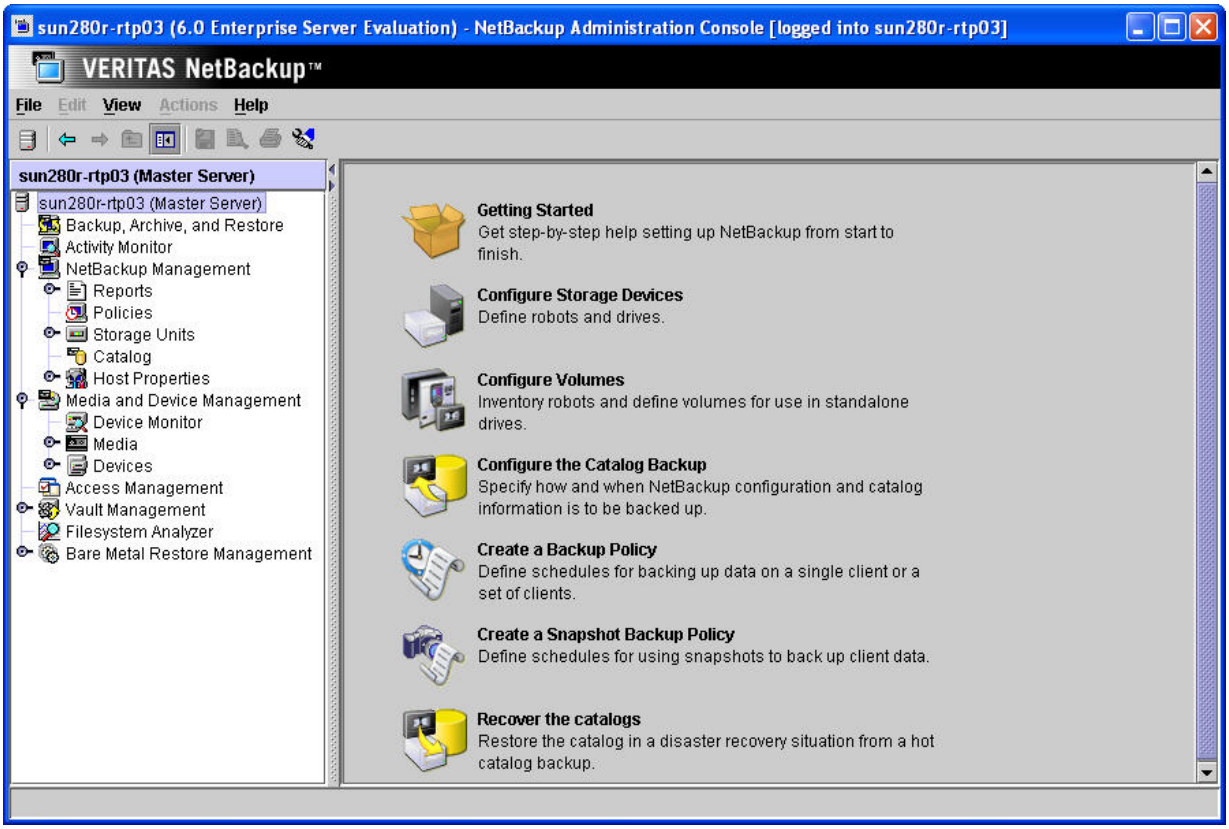

Figure 24. NetBackup Administration Console.

The NetBackup Administration Console is typically utilized for all configuring, managing and reporting of the NetBackup environment (and what is depicted throughout this document).

Data ONTAP

NDMP needs to be enabled for all the solutions.

SnapRestore license required for NSM.

SnapVault Primary and SnapVault secondary licenses required for NSVM.

SnapVault secondary license required for SV-NBU.

Licensing

See table in section 6.1 for required software licenses.

NDMP authentication

NDMP is required for authentication for all three solutions. If NDMP is not set-up correctly the solutions will not work, so this section discusses installing and configuring it in fairly good detail.

If you have a Windows Master Server then NDMP is already installed. If you are running UNIX then the NDMP package must be separately installed.

Enabling NDMP on IBM N series

Before NetBackup can communicate with IBM N series storage systems as part of the joint disk-to-disk backup solutions, NDMP needs to be turned on and configured on the IBM N series device(s). For NSM that means the primary Filer. For NSVM it means the Primary and Secondary Filers. And for SV-NBU it means the IBM N series secondary storage system.

To determine if NDMP is running on the IBM N series storage system, use the following command:

ndmpd status.

If NDMP is not running, start it by using the "`ndmpd on`" command or using FilerView to enable it.

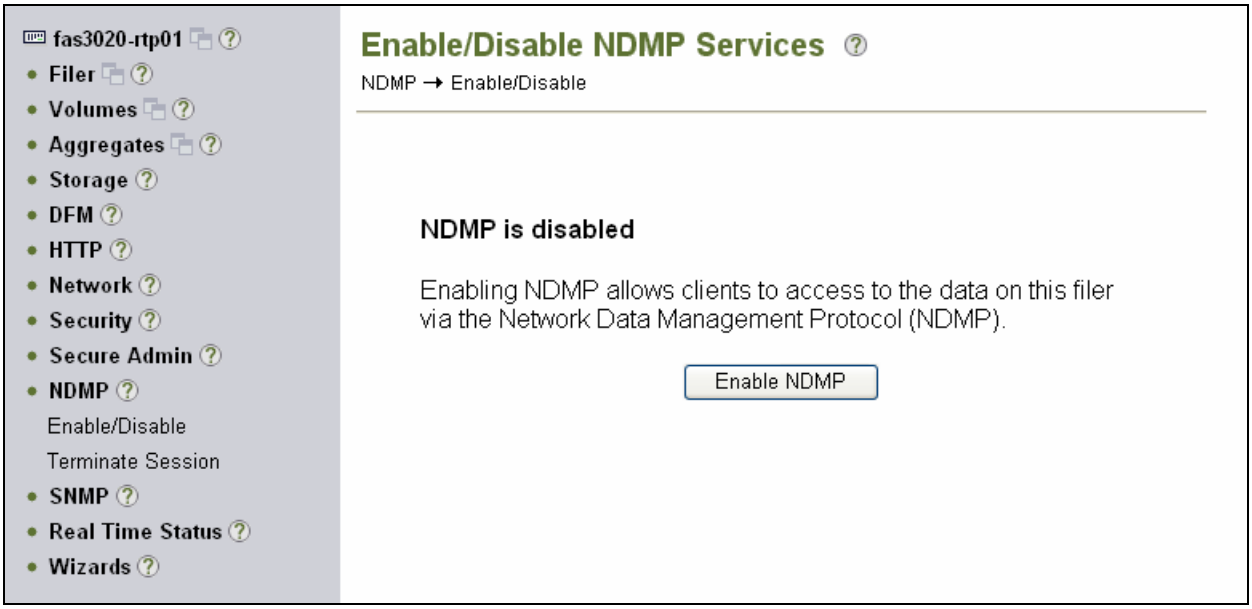

Figure 25. Turning on NDMP Using FilerView.

This paper uses the root user for NDMP authentication. In production environments that could be considered a security risk, so you can 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 GUI. This section walks through setting up those credentials. It first depicts setting up authentication for the Primary Filer (sunv240-rtp02), necessary for NSM and NSVM, and then adds the same for the secondary (r200-rtp01), required for NSVM and SV-NBU.

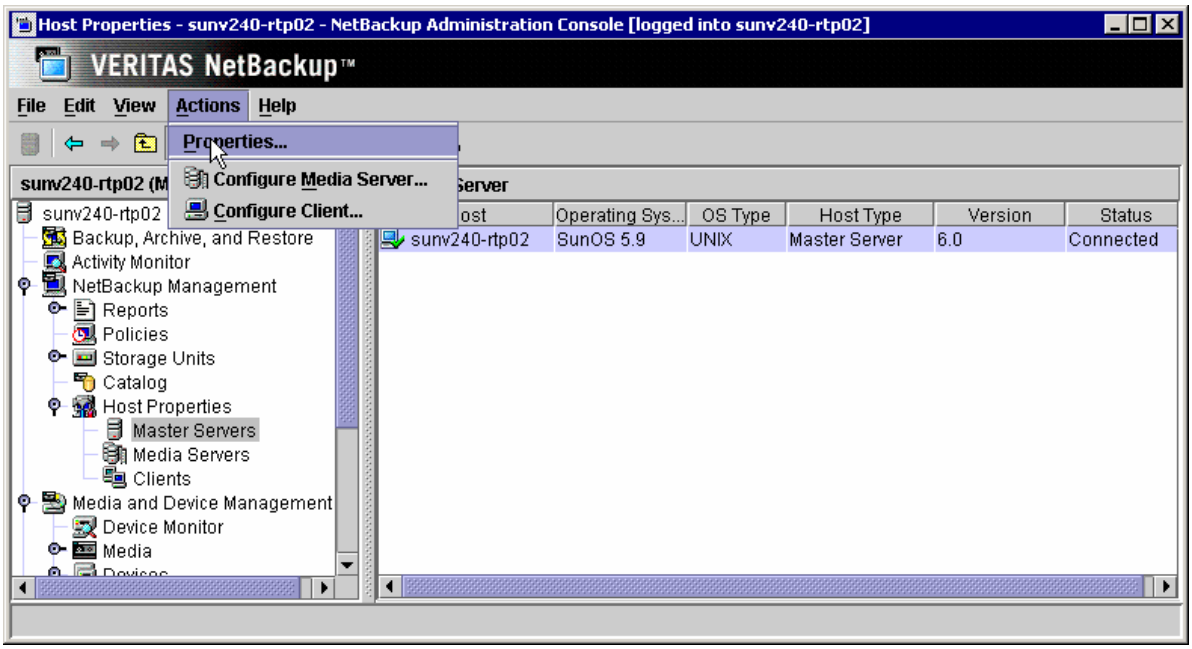

Figure 26. Accessing NetBackup Master Server Host Properties.

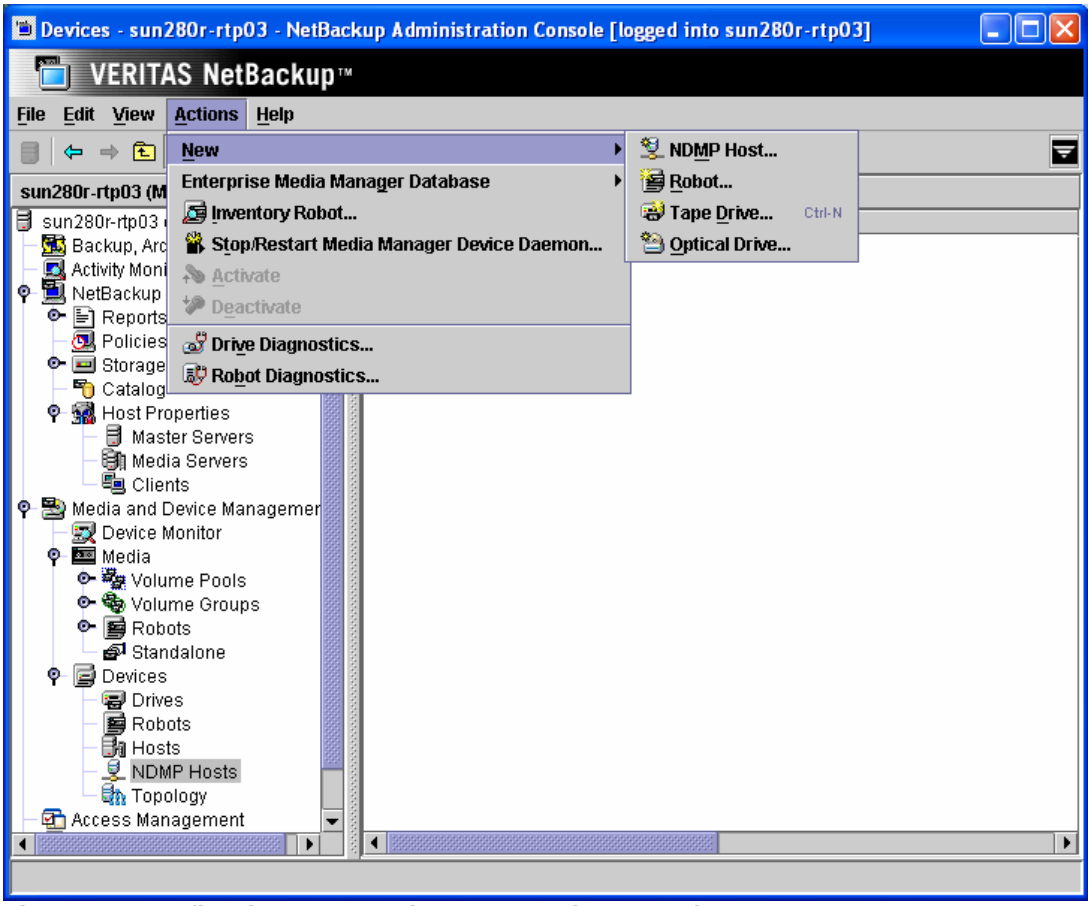

Figure 27. Configuring NDMP Primary Host Via NetBackup GUI.

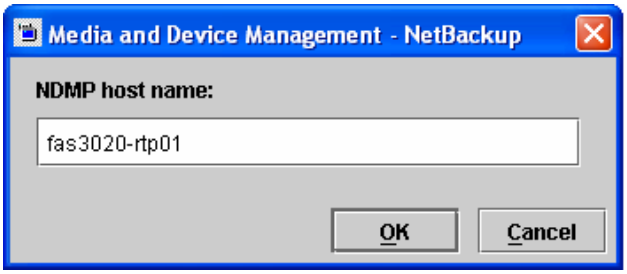

Figure 28. Specifying NDMP Primary Host Name.

For a specific NDMP authenticated host, you can either specify specific credentials just for it…

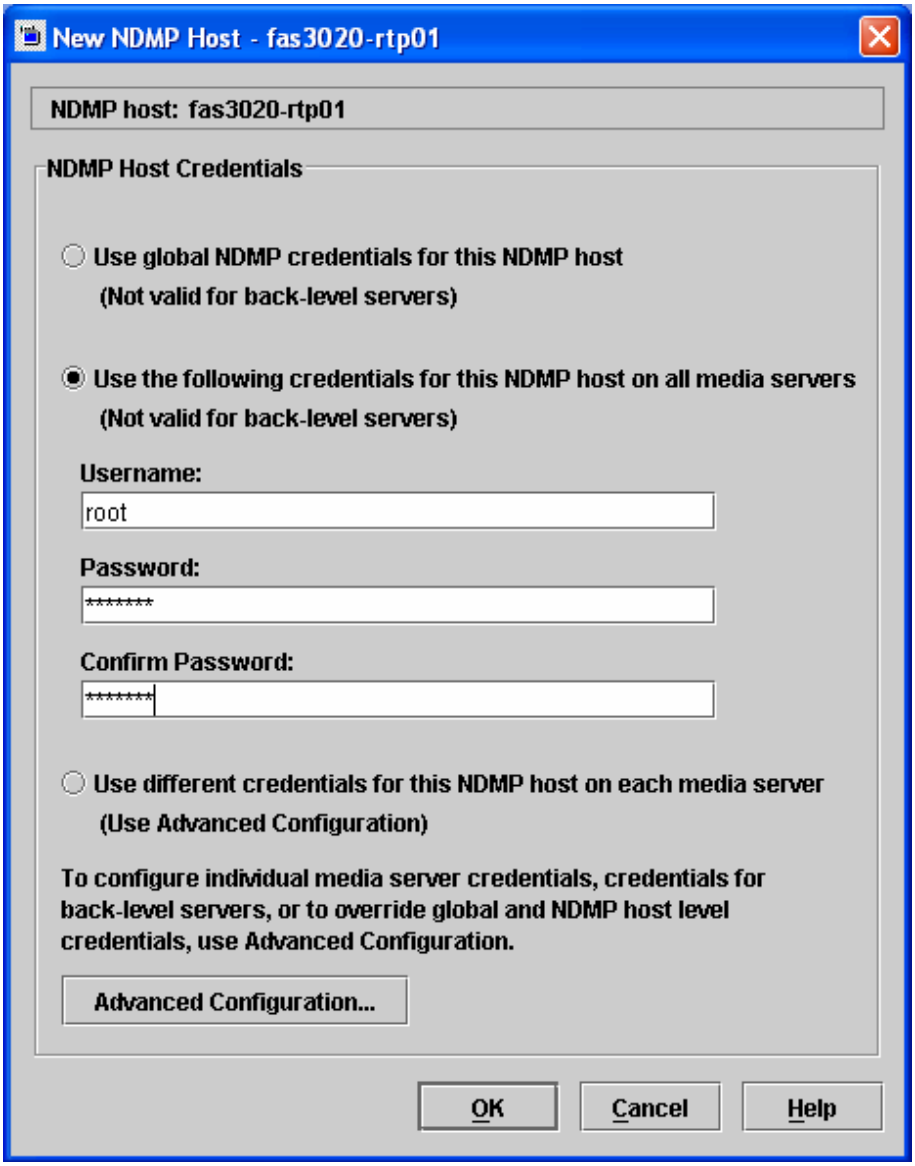

Figure 29. Specifying Specific NDMP Primary Host Credentials.

Or set Global NDMP Authentication. While Global authentication is easiest, it is also less secure.

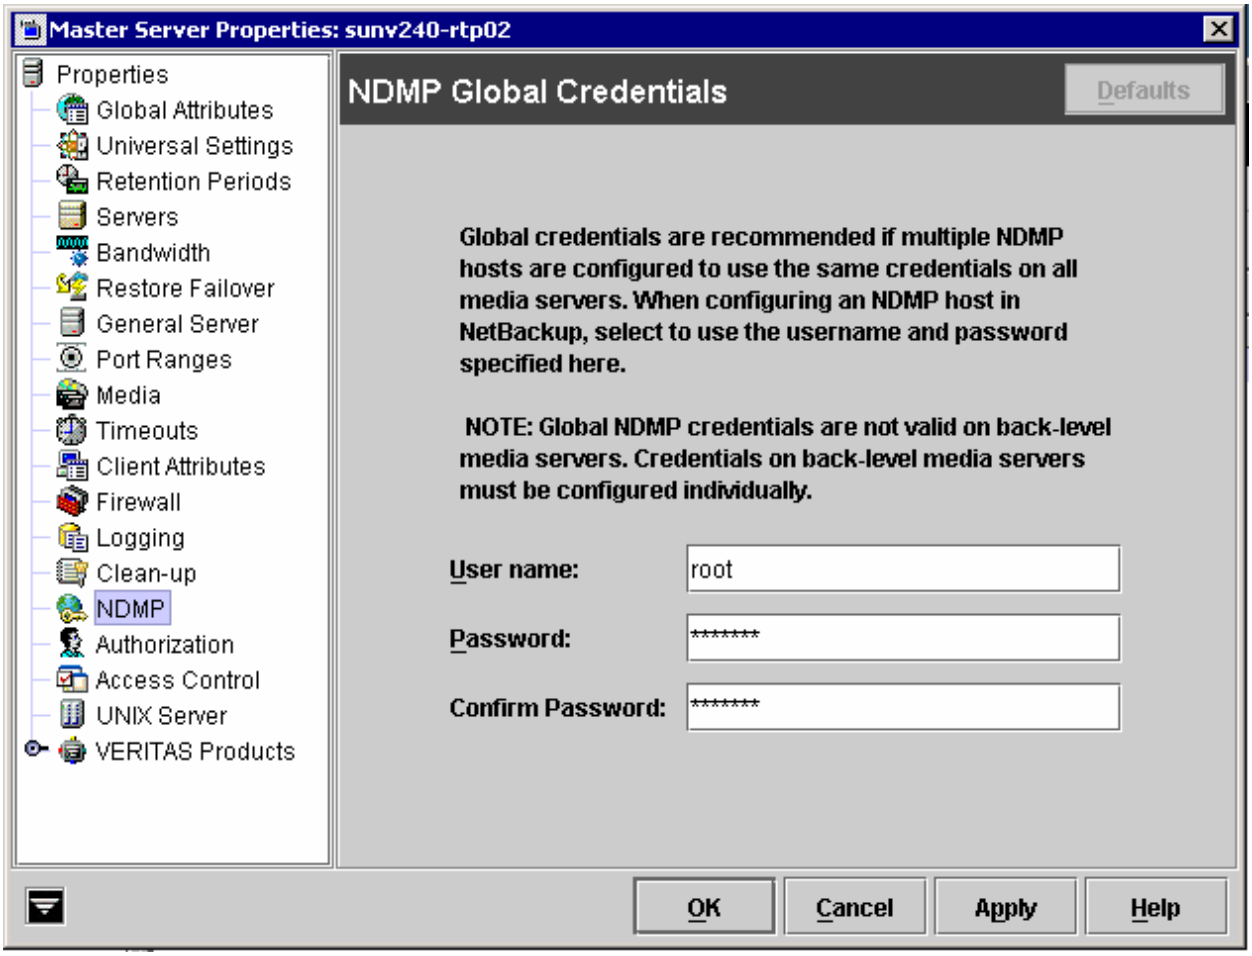

Figure 30. Specifying NDMP Global Credentials.

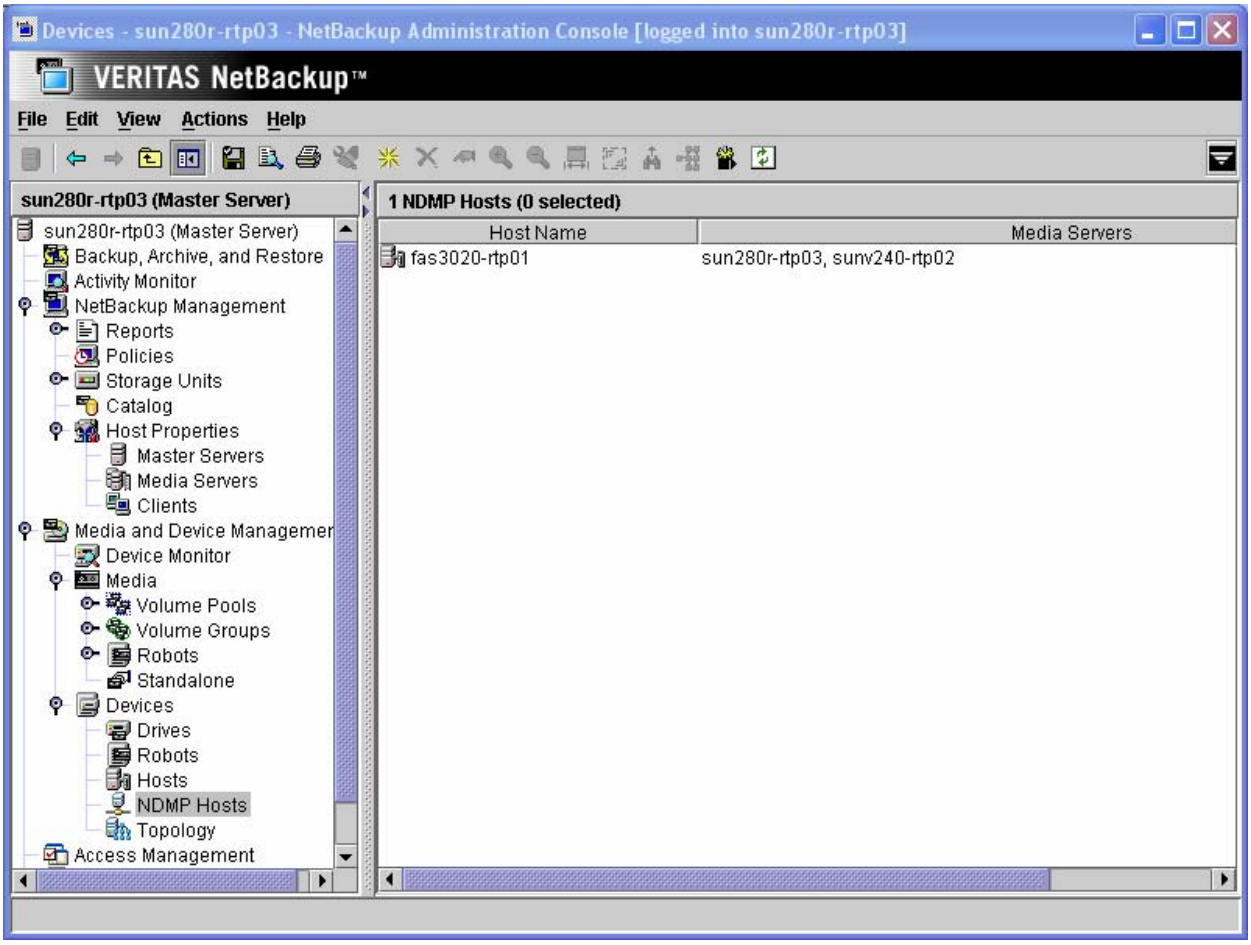

Figure 31. NDMP Host (Primary) Authenticated with Master and Media Servers.

Now we set-up the NDMP credentials for the Secondary.

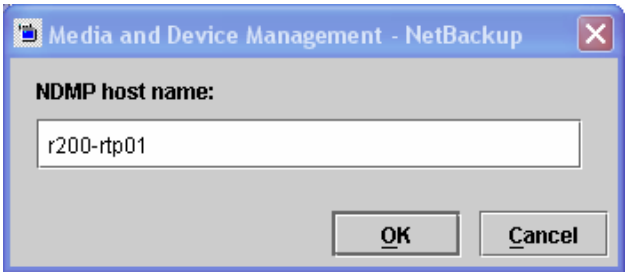

Figure 32. Specifying NDMP Host Name for Secondary.

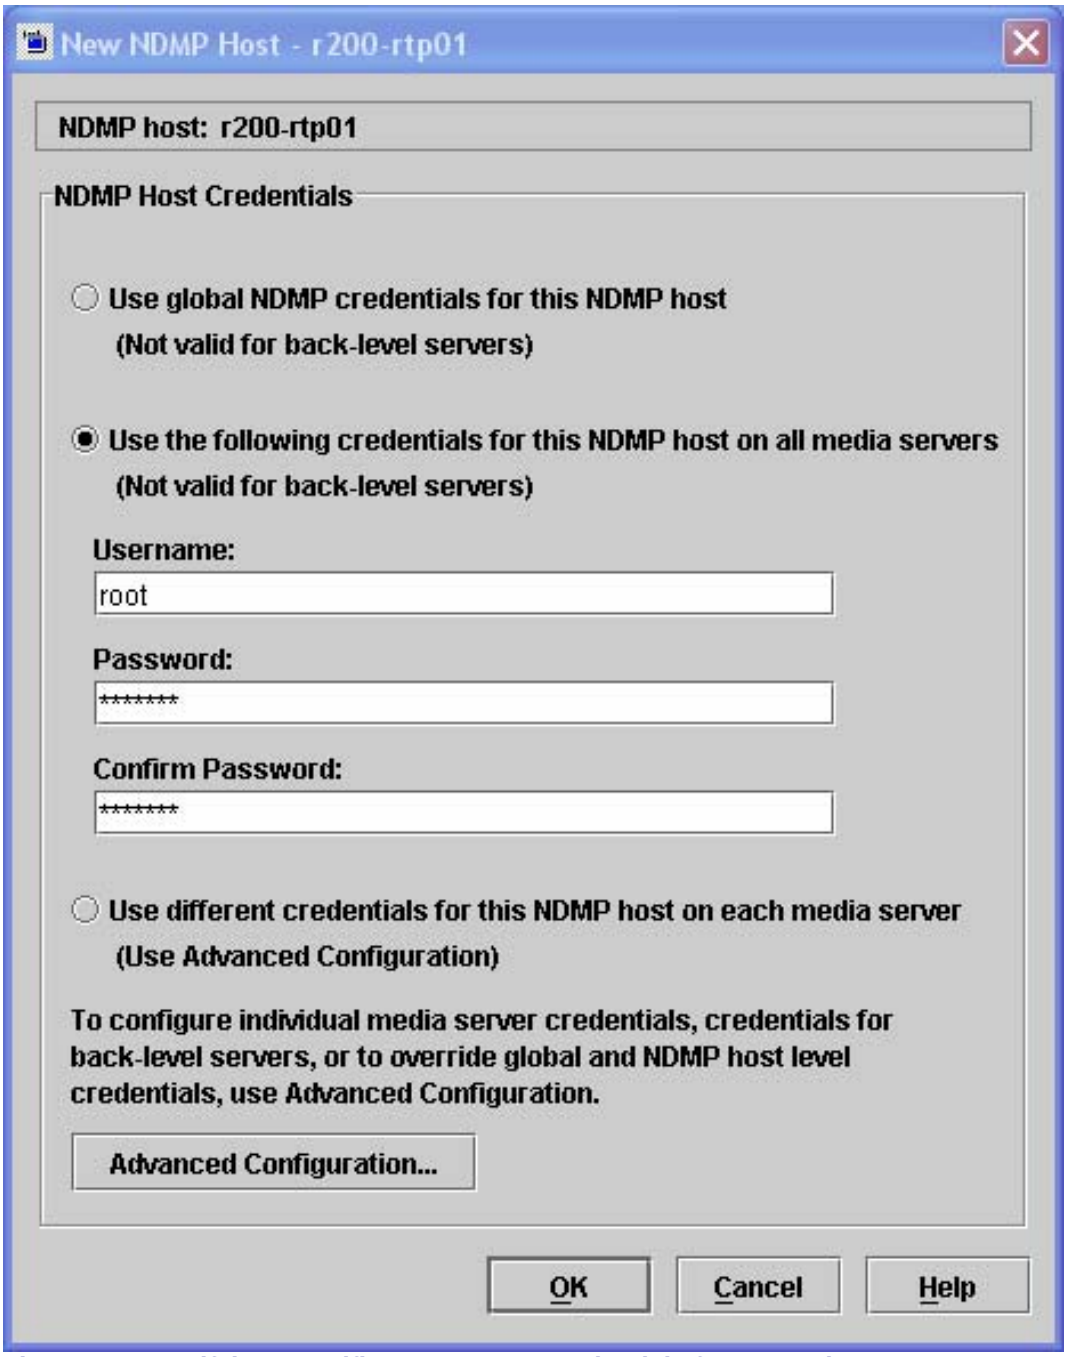

Figure 33. Specifying Specific NDMP Host Credentials for Secondary.

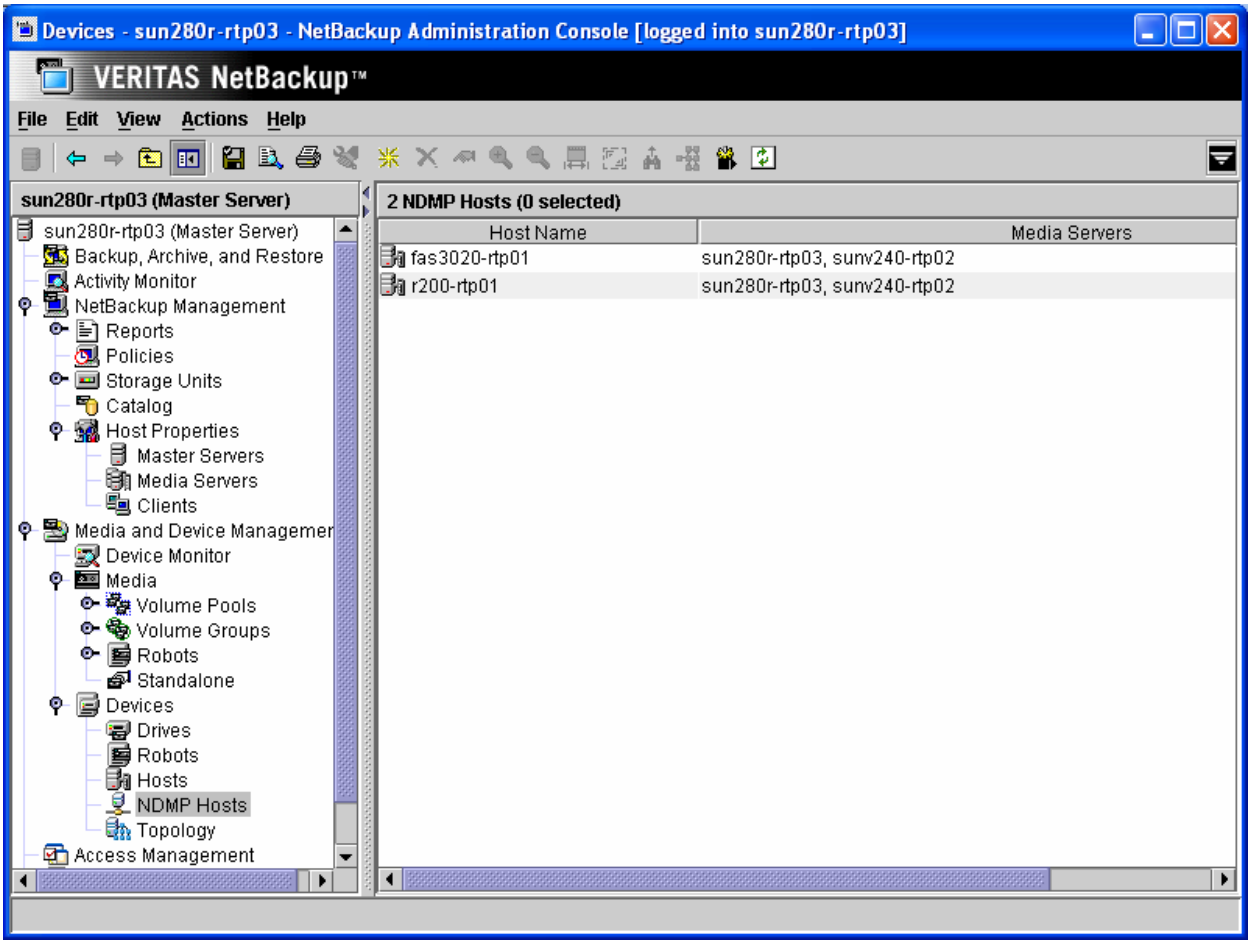

Figure 34. NDMP Hosts (Primary Filer and Secondary NearStore System) Authenticated with Master and Media Servers.

NetBackup 6.0 NDMP command line interface

If you like to type, you can use a CLI to configure NetBackup NDMP Authentication. "`tpconfig`" is the utility used to create, manage and verify the NDMP authentication relationships.

The example below shows using "`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:
Please 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 example below shows using "`tpconfig`" on Windows to configure the same NDMP username and password as the previous example.

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

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

```
sunv240-rtp02# which tpconfig
/usr/openv/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 "IBM N series DataONTAP"
  os version "DataONTAP 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` 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>.

NetBackup 5.1 command line Interface

With 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:

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

NSM configuration

This section contains specific steps about getting the NSM solution up and running. It is assumed that in addition to the IBM N series storage array and the proper version of Data ONTAP being installed, the following items have also already been done:

Install / License NetBackup Enterprise Server
Install / License NetBackup Advanced Client
Install / License / Configure NDMP (on both IBM N series storage system and via NetBackup).

The basic NSM environment looks like the following:

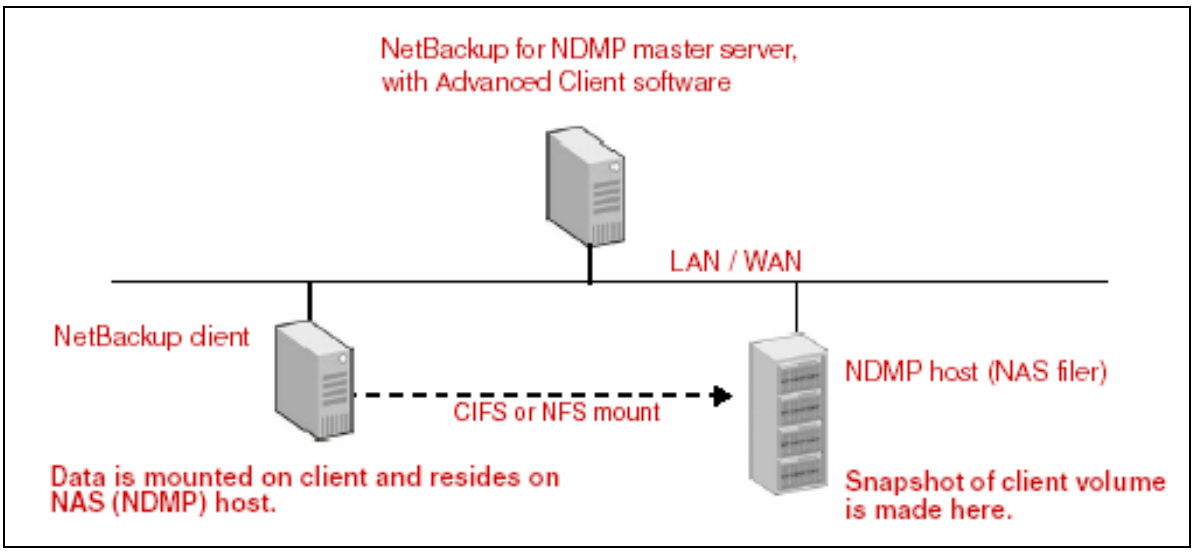

Figure 35. Basic NSM Environment.

Note in the ensuing example:

The master server is also the client.
Master Server = sun280r-rtp03
Filer = fas3020-rtp01.

The remainder of this section steps through the following:

Making and sharing a volume on IBM N series filer
Configuring NetBackup NSM policy.

IBM N series NSM configuration

Ensure NDMP authentication has been configured between the NetBackup master server (and any necessary media servers) and the IBM N series storage system.

Using FilerView, first create a volume on the IBM N series storage system (NAS filer) by selecting *Volume* > *Add*. You can see what volumes have been created by selecting *Volume* > *Manage*, as shown in the following figure. *vol1code* is the volume we've created for this example.

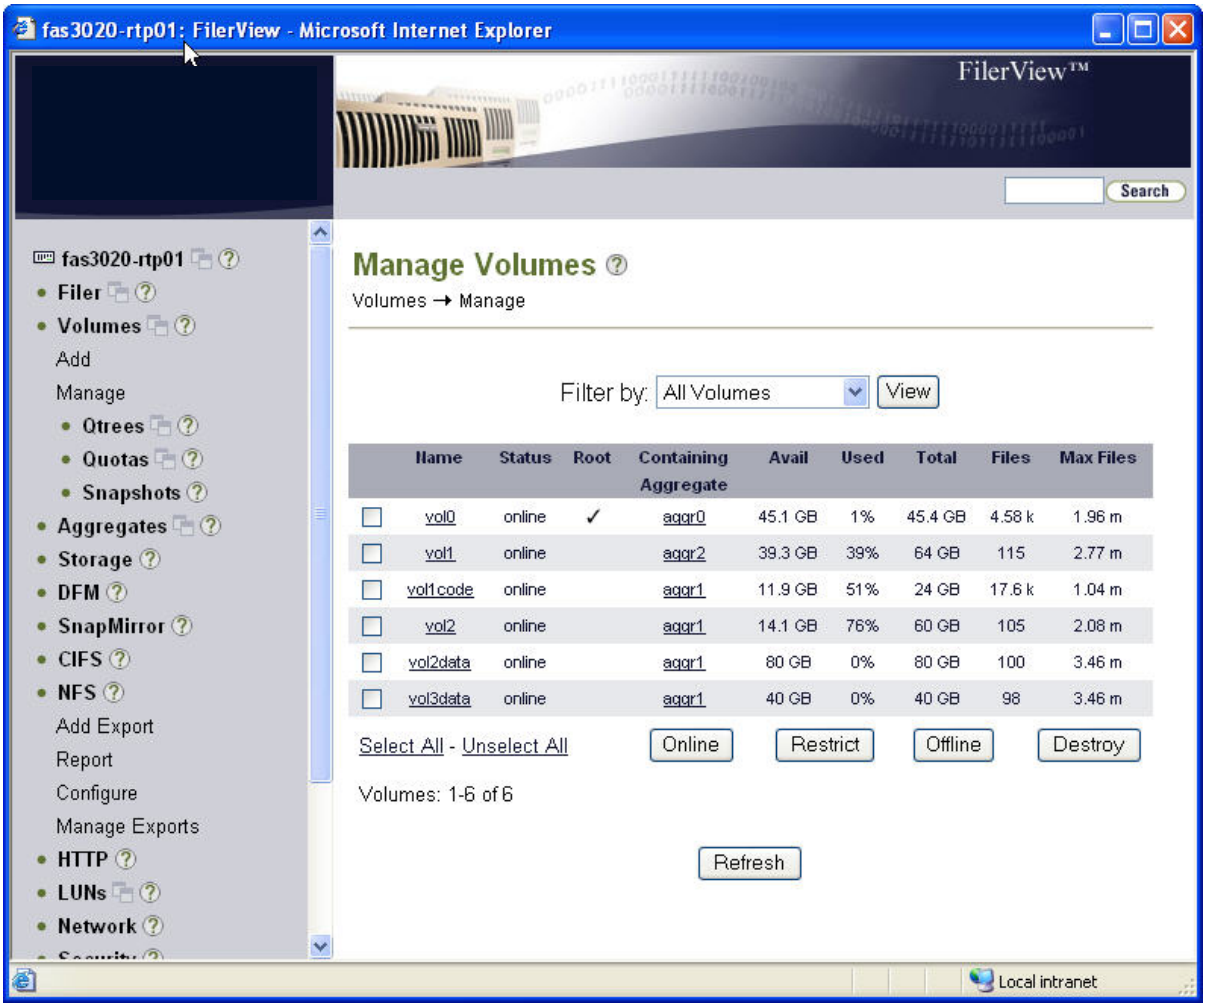

Figure 36. Creating a Volume Using FilerView.

Now create an NFS export on the IBM N series storage system (NAS filer) by selecting *NFS* > *Add Export*. You can see what exports have been created by selecting *NFS* > *Manage Exports*, as shown in the following figure. (It could also be a CIFS share for a Windows client to map for its data.) In the example below, */vol/vol1code* is what we're exporting.

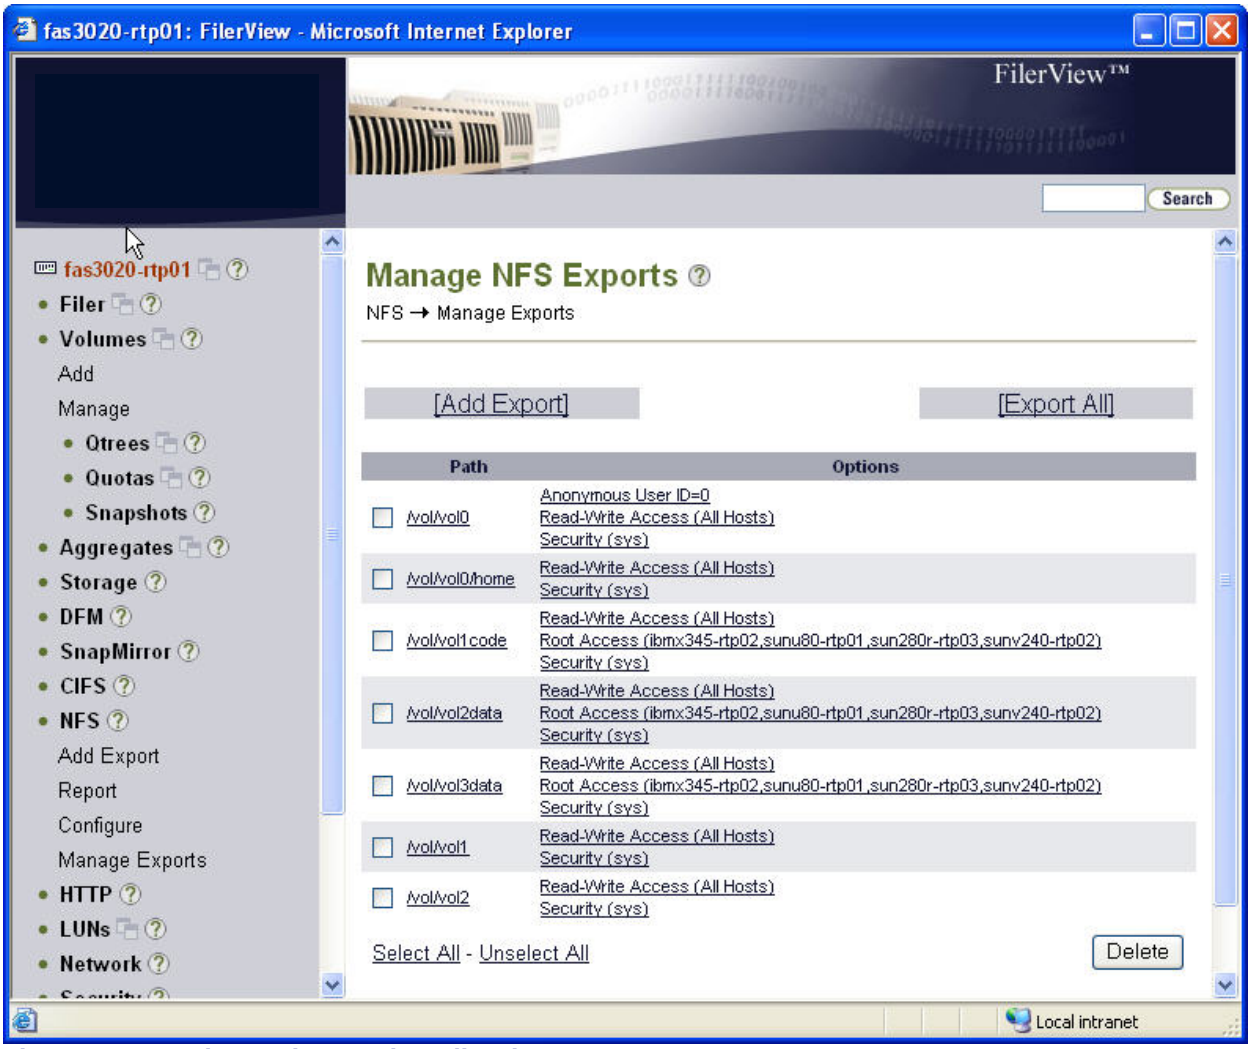

Figure 37. Creating a Share Using FilerView.

We'll next create a Qtree, a quota tree on the volume being exported. **Although this is not required for NSM backups, it IS necessary for NSVM backups so we'll go ahead set it up and use it**. In the example below, */vol/vol1code/Codeqtree* is what the UNIX client will mount.

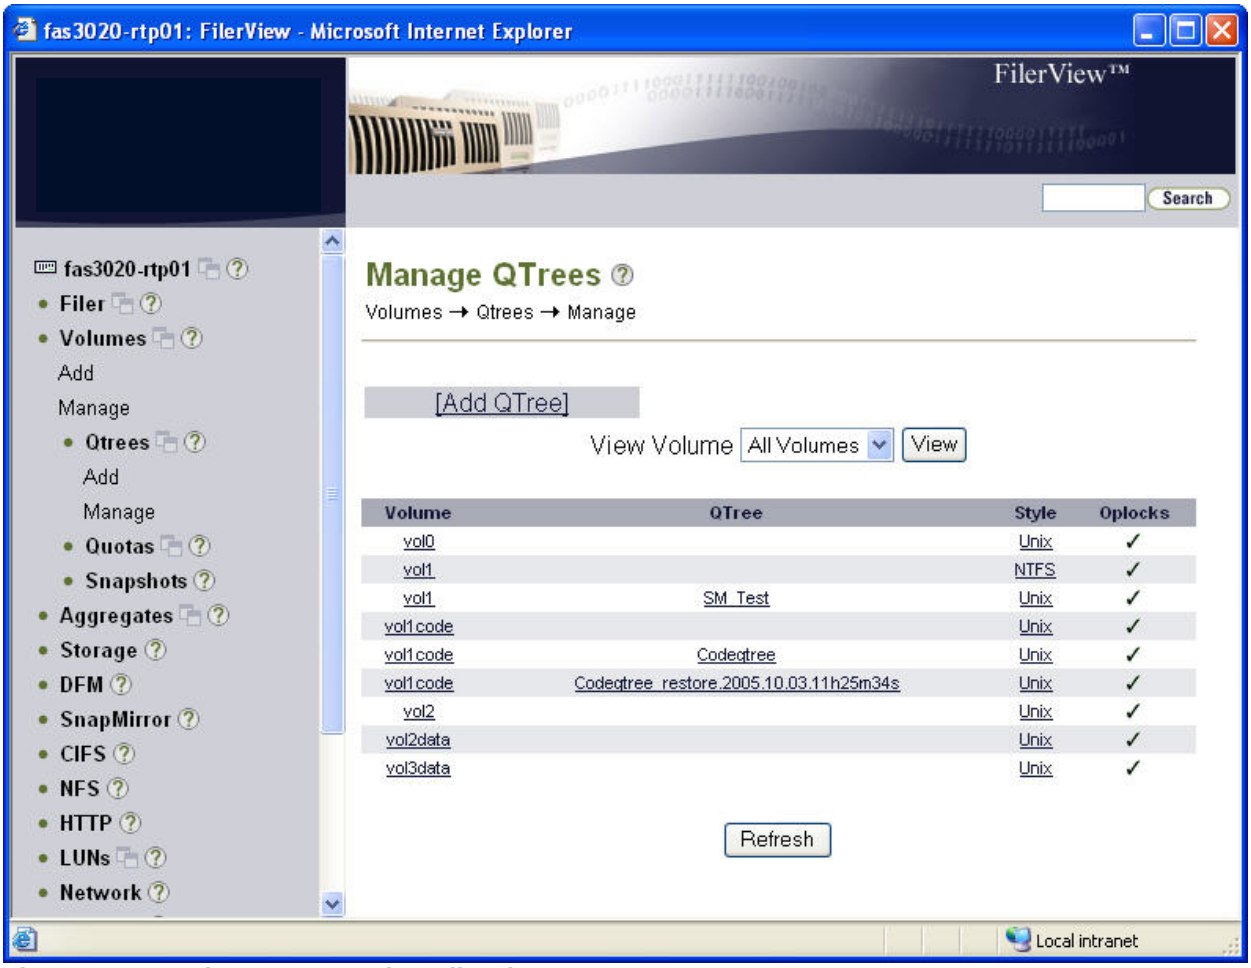

Figure 38. Creating a Qtree Using FilerView.

Quickly looking at the client, sun280r-rtp03, you can see what is mounted:

```
sun280r-rtp03# df -kF nfs
Filesystem            kbytes    used   avail capacity  Mounted on
fas3020-rtp01:/vol/vol1code/Codeqtree
                     25165824 12736928 12428896    51%    /Codemount
fas3020-rtp01:/vol/vol2data
                     83886080    2360 83883720     1%    /Datamount2
fas3020-rtp01:/vol/vol3data
                     41943040 5898268 36044772    15%    /Datamount3
```

As NSM is now effectively managing the snapshots for the volume there, is no need to keep Data ONTAP scheduled snapshots too so a final step is to turn off scheduled snapshots for the volume. While this can be accomplished from FilerView, it is typically accomplished using the CLI as it's quite simple:

r200-rtp01> vol options <volume> nosnap on

The IBM N series NSM configuration is now complete.

NetBackup NSM configuration

Use the NetBackup Administration Console to create a new policy, by selecting *Policies* and then *Action > New Policy*.

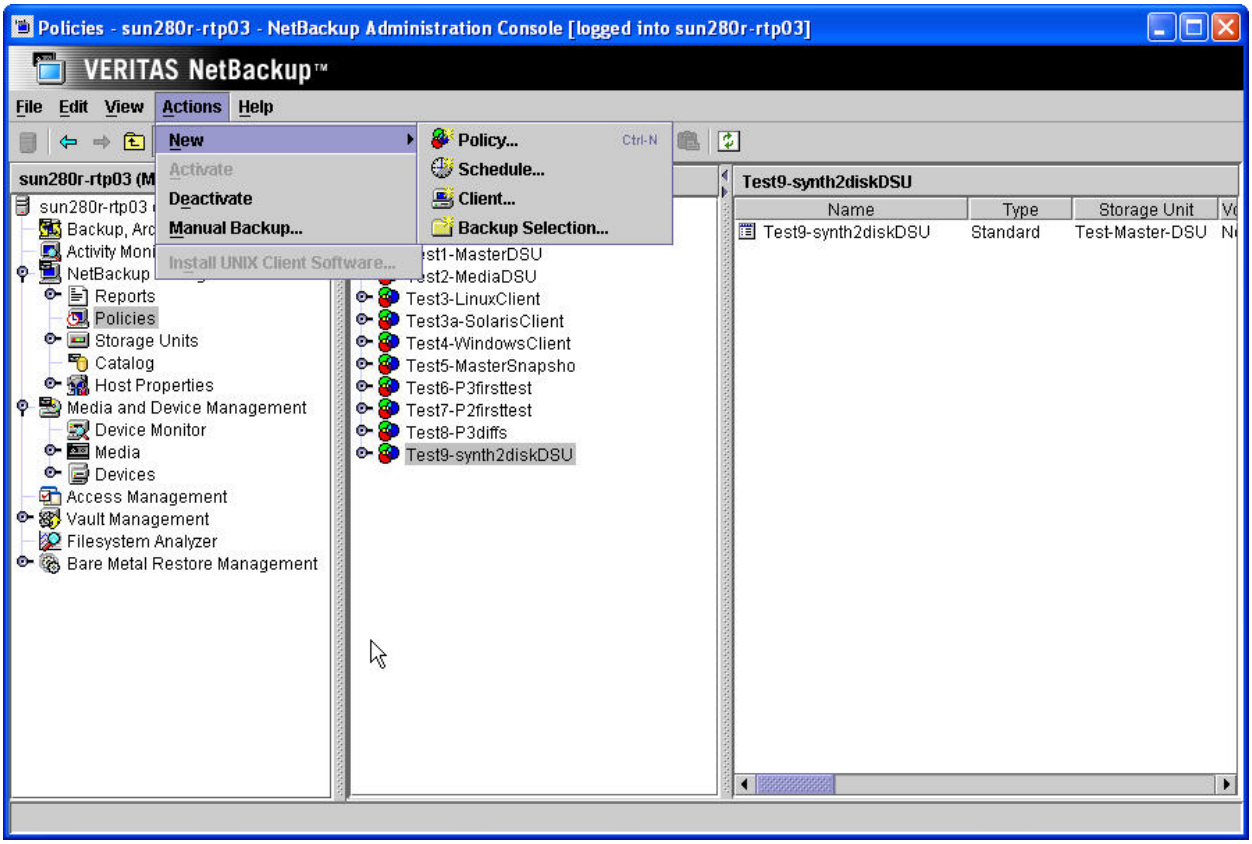

Figure 39. Creating a New NetBackup Policy.

Then enter the name of the new policy.

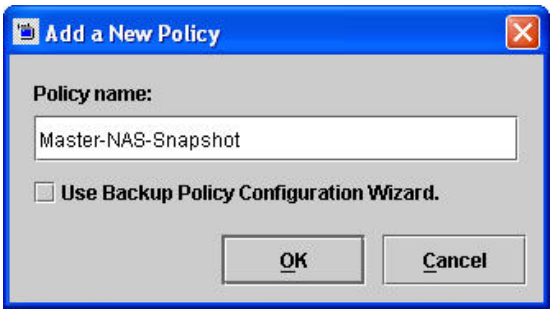

Figure 40. Naming a New NetBackup Policy.

The change policy GUI comes up, with the attributes tab displayed.

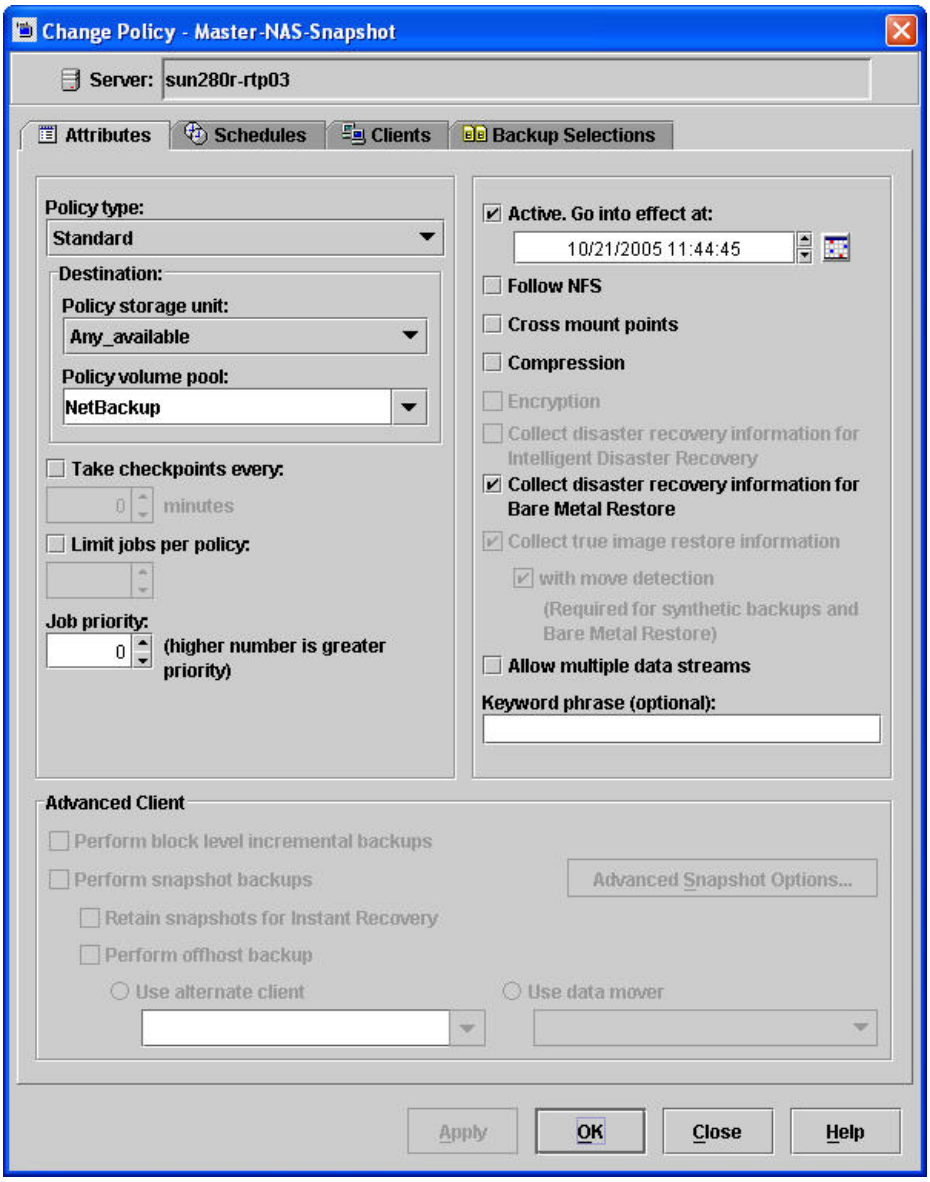

Figure 41. NetBackup Change Policy GUI.

You'll need to turn off "*Collect disaster recovery information for bare metal restores*" before the Advanced Client options are accessible.

Then select:

Perform snapshot backups
Retain snapshots for Instant Recovery
Perform offhost backup

Use data mover (pull-down and specify Network Attach Storage).

Note that the Policy Storage Unit field is ignored with NSM backup policies, since NetBackup creates dnapshot copies on the NAS-attached disk only, not on storage devices attached to a NetBackup server or client.

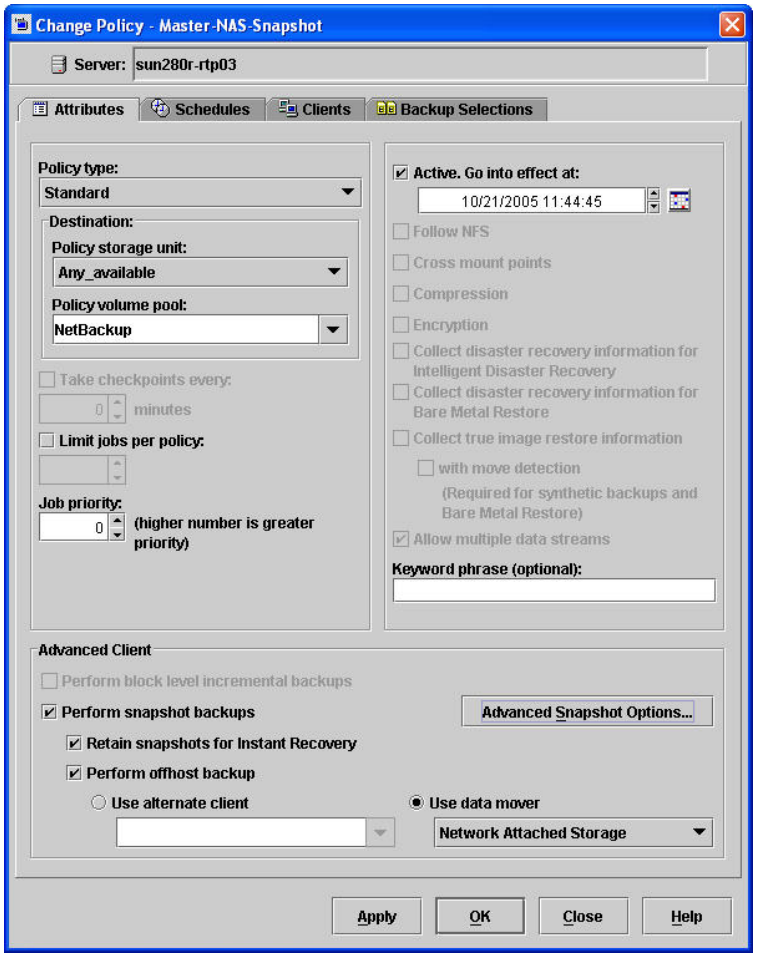

Figure 42. NetBackup Change Policy for Advanced Client.

Access the *Advanced Snapshot Options* to specify the retention of snapshot copies via the *Maximum Snapshots (Instant Recovery only)* parameter. This, not the schedule *Retention* field, is what determines how long backups are kept.

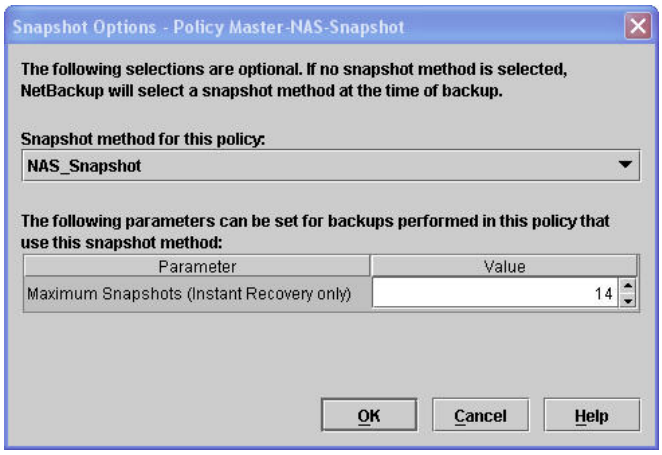

Figure 43. Setting the Maximum Number of Snapshot Copies to Keep.

(Note that in the above example after the 15th time the policy is run, the first backup - snapshot copy - will be deleted.)

Now you need to configure the schedules for the policy via the *Schedules* tab.

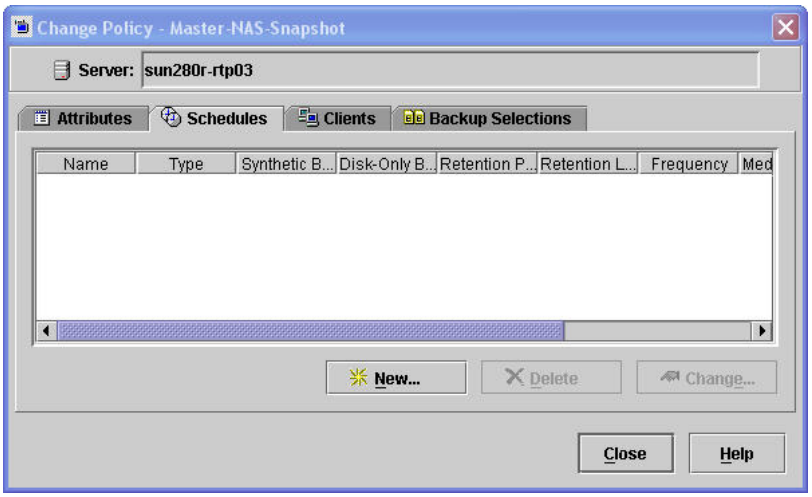

Figure 44. NetBackup Policy Schedules Tab.

Click on *New…* to create a schedule for this policy. Special things to note for a NSM backup policy are:

Snapshots only must be selected and

Retention should be set to Infinity (as this is effectively specified by the “Maximum Snapshot Copies to Keep” parameter configured via the “Advanced Snapshot Options…” on the Attributes tab)

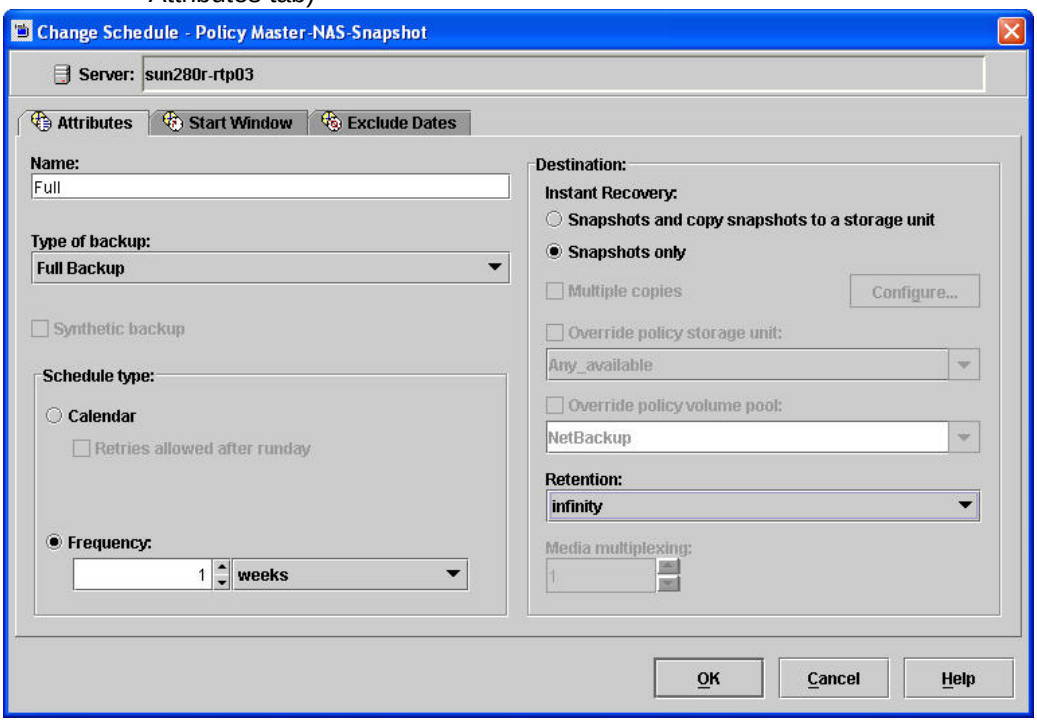

Figure 45. Creating a New Policy Schedule.

Now configure the clients which will be backed up with this policy by accessing the *Clients* tab.

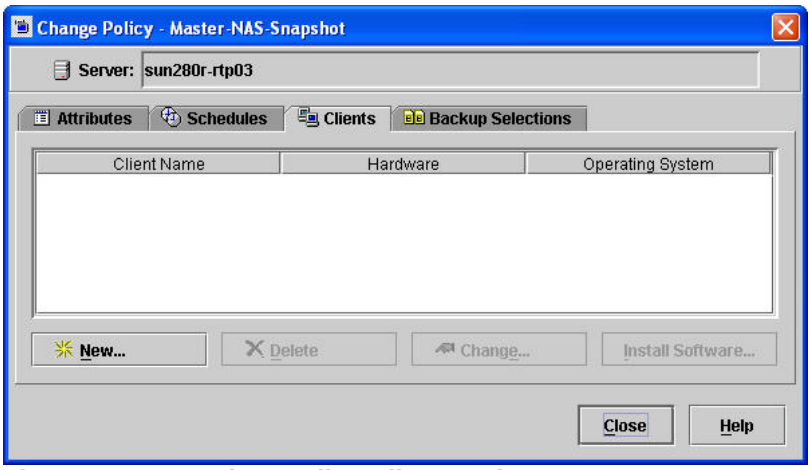

Figure 46. NetBackup Policy Clients Tab.

Click on *New…* to add a client for this policy.

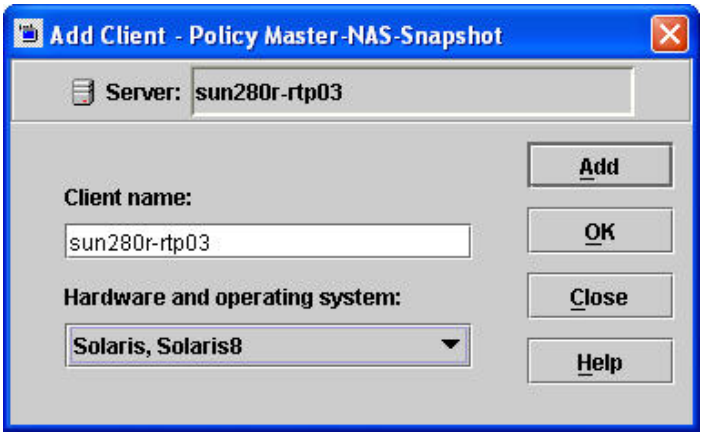

Figure 47. Adding a Client to a Policy.

Finally, access the *Backup Selections* tab to specify what you want to backup.

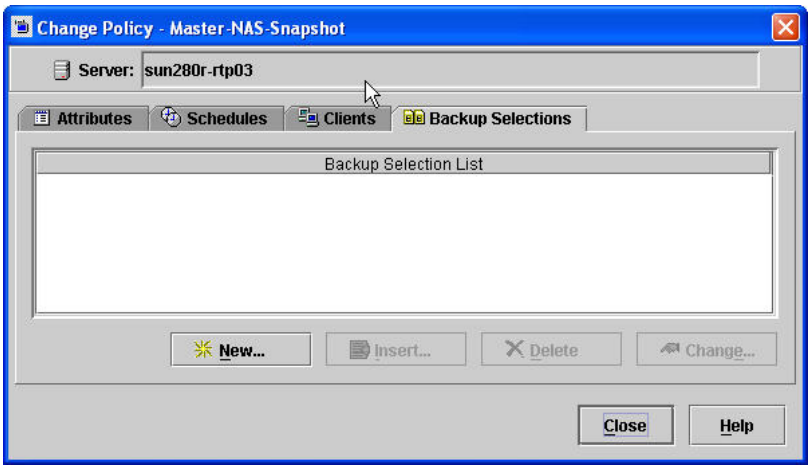

Figure 48. NetBackup Backup Selections Tab.

It is important to note that the way the backup selections are specified is from the Client's perspective. And it is subtly different between UNIX and Windows clients:

For UNIX client: `/NFS_mountpoint`
For Windows client: `\\ndmp_hostname\share`
Note: Windows pathnames must use the Universal Naming Convention (UNC).

On our Solaris client, sun280r-rtp03, we have the following file systems NFS-mounted:

```
sun280r-rtp03# df -kF nfs
Filesystem            kbytes    used   avail capacity  Mounted on
fas3020-rtp01:/vol/vol1code/Codeqtree
                     25165824 12736928 12428896    51%    /Codemount
fas3020-rtp01:/vol/vol2data
                     83886080    2360 83883720     1%    /Datamount2
fas3020-rtp01:/vol/vol3data
                     41943040 5898268 36044772    15%    /Datamount3
```

So you specify */Codemount* as the backup selection. (Note that the backup selection must be a volume – in this case it isn't, it's a qtree – or *Point-In-Time Rollback* will not work.)

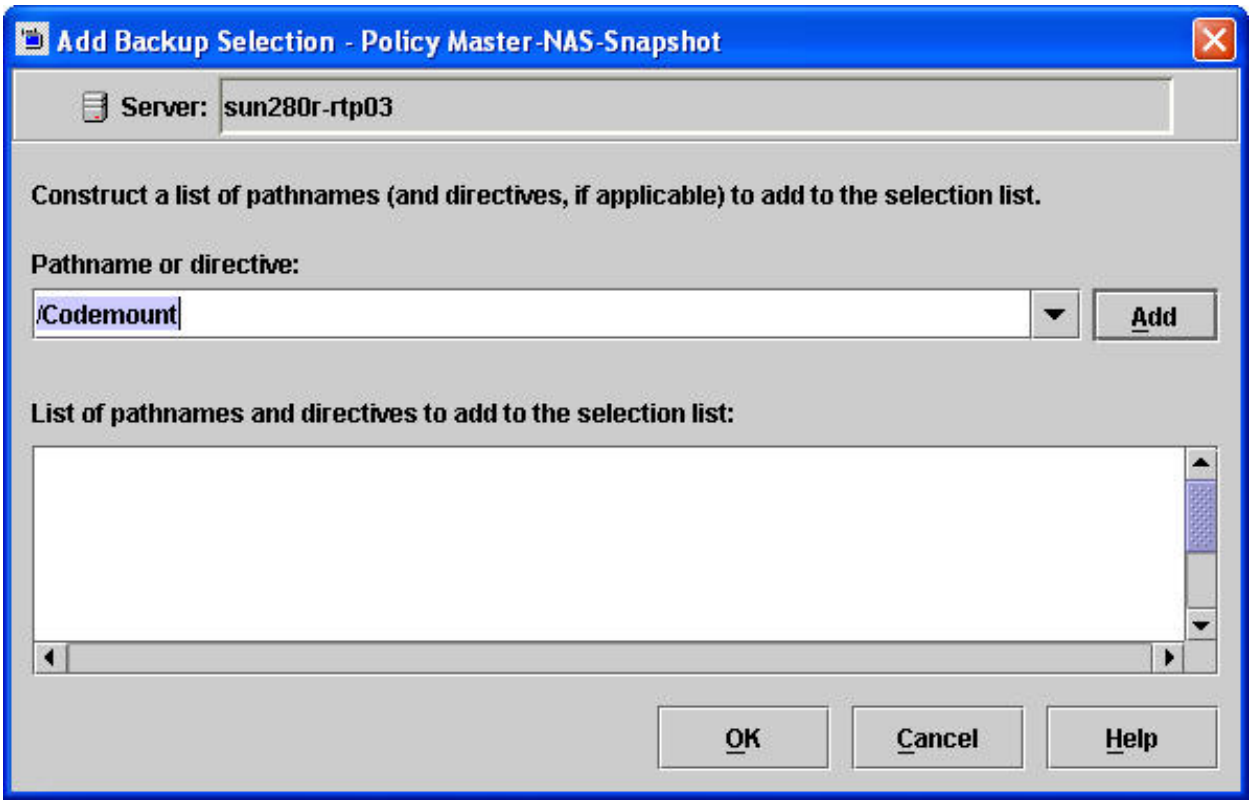

Figure 49. Specifying a Backup Selection.

When finished with all the above steps, click *Close*, and NetBackup will verify the Advanced Client policy is set-up correctly. For NSM backups the pathname specified can be for either a volume or a qtree which is mounted to the client.

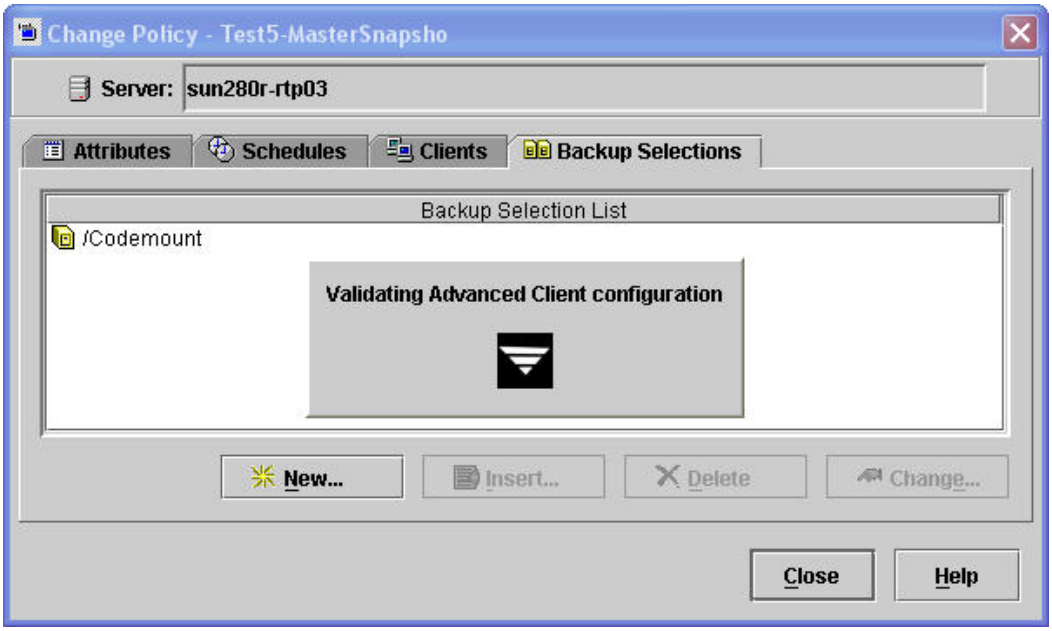

Figure 50. NetBackup Verifies an Advanced Client Policy.

The NSM NetBackup policy configuration is now complete.

Running a NSM backup

To run the NSM backup policy configured above, highlight it and select *Actions* > *Manual Backup*.

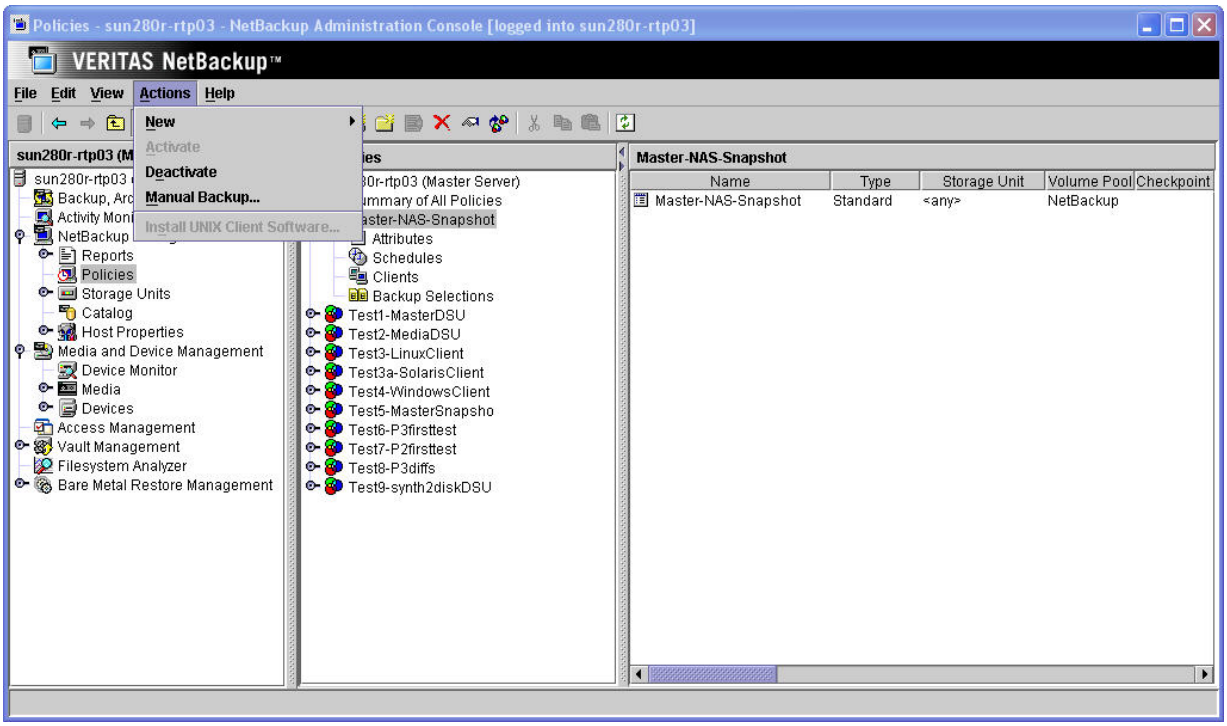

Figure 51. Manually Running NSM Backup Policy.

Then select either the specific schedule to run or click on *Okay* to let the scheduler decide which one is due to run. (In this example there is only one schedule so that one will run no matter what.) A backup job is then initiated.

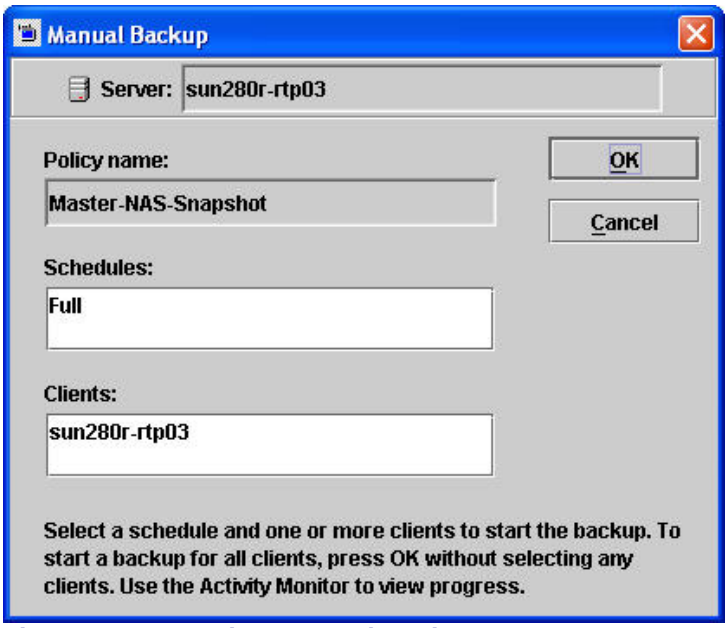

Figure 52. NetBackup Manual Backup GUI.

Use the Activity Monitor to watch the progress of the backup job – Job Id #64 shown below.

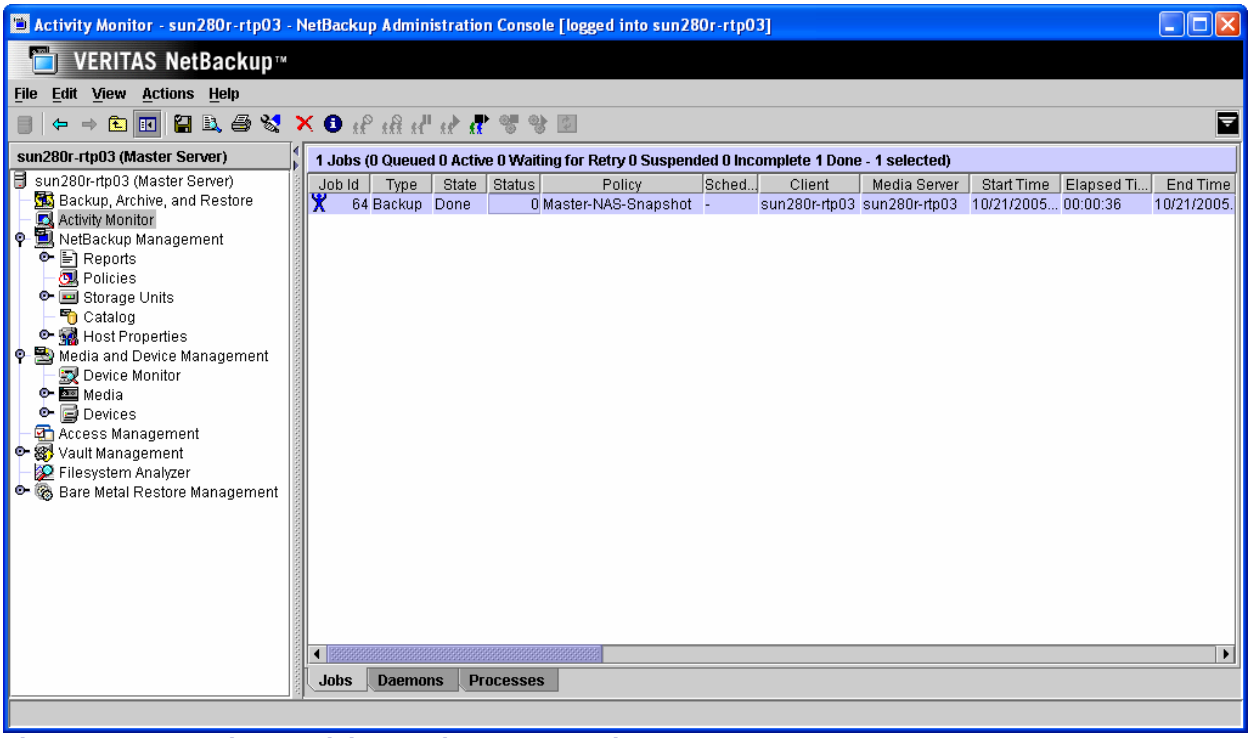

Figure 53. NetBackup Activity Monitor – NSM Job.

Double-click on the job to see additional information on it, and select the *Detailed Status* tab to see all the details.

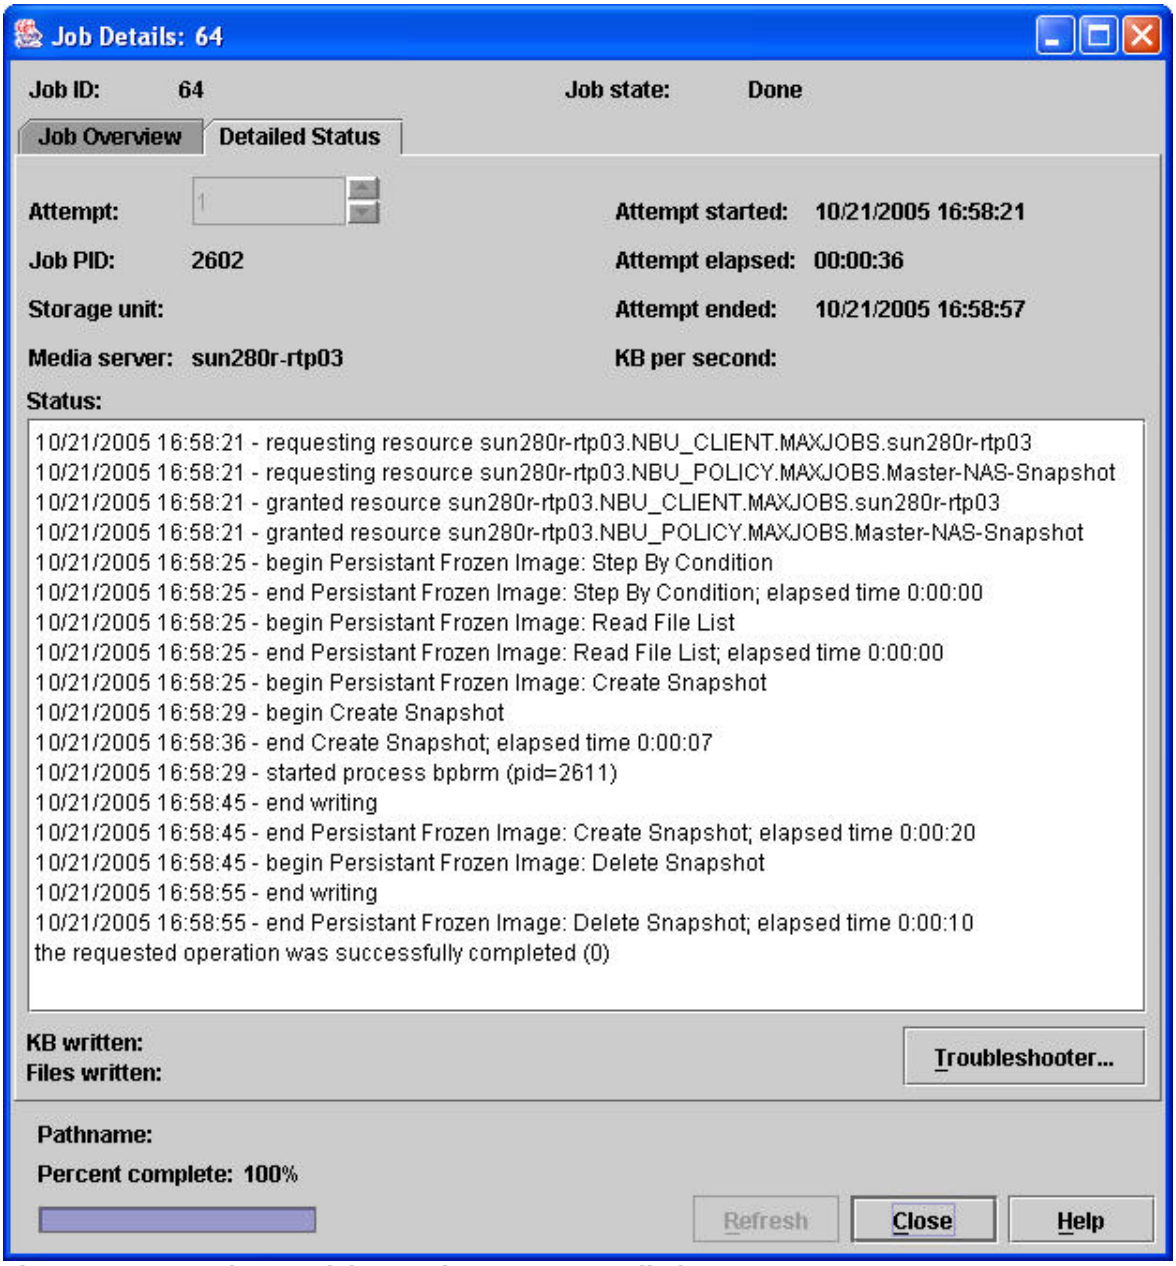

Figure 54. NetBackup Activity Monitor – NSM Detailed Status.

NSM restores

Restores are accomplished via the NetBackup *Backup Archive Restore* GUI, which is selected from the NetBackup Administration Console.

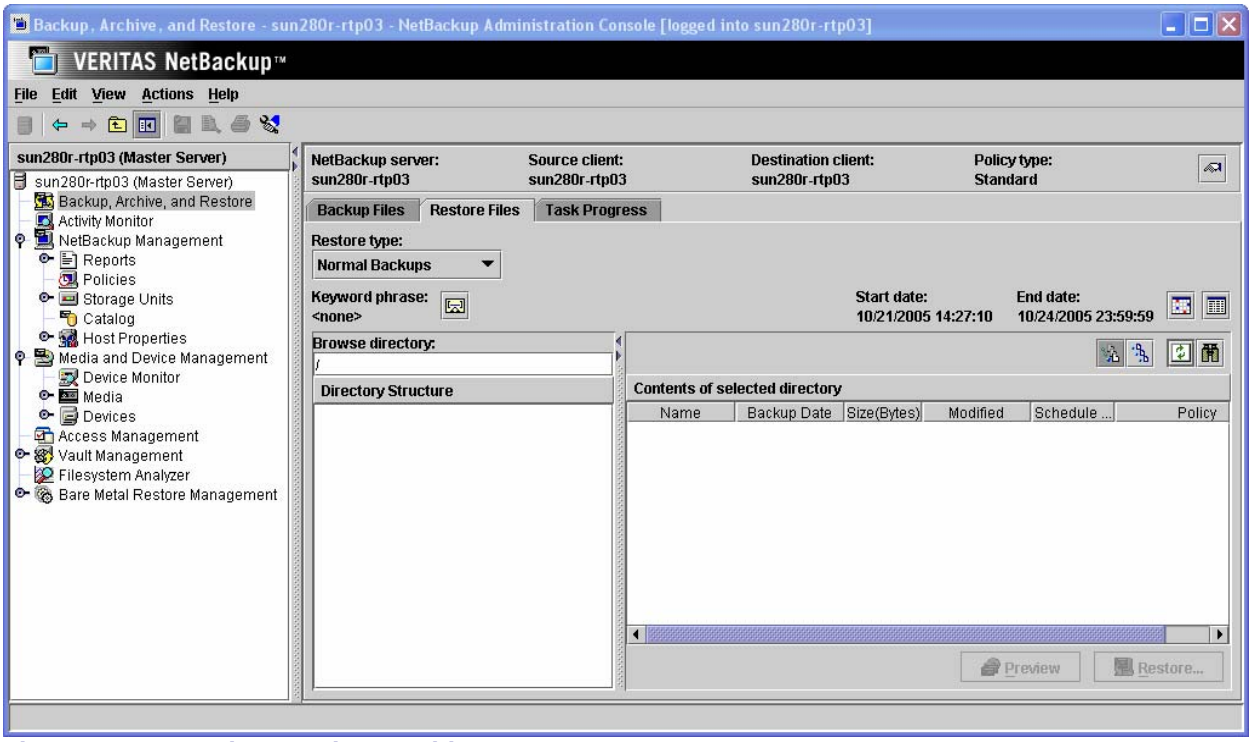

Figure 55. NetBackup Backup Archive Restore GUI.

To restore specific files and directories, set the *Restore Type* to *Normal Backups*. Verify the client whose data you want to restore, select a date range of interest and click on the refresh icon to search for the backups.

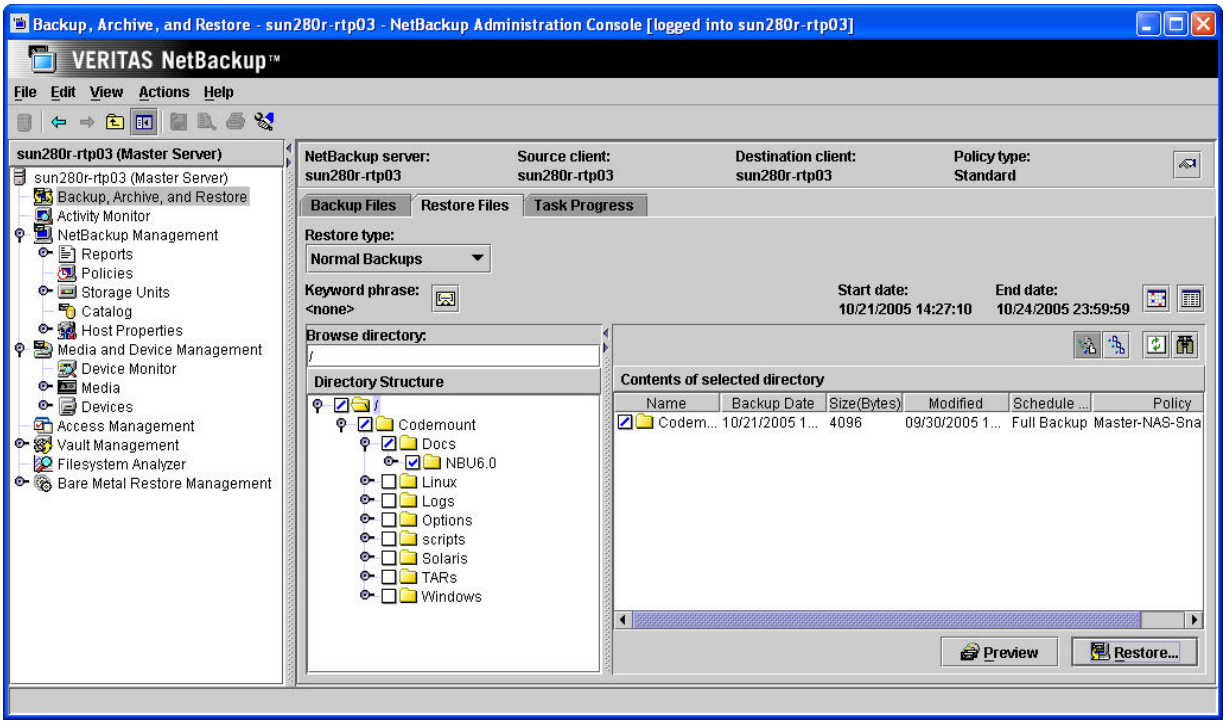

Figure 56. NetBackup Backup Archive Restore – NSM Normal Backups Restore.

After selecting the files/directories you want to restore, click on the *Restore…*button to access the GUI which will initiate the restores.

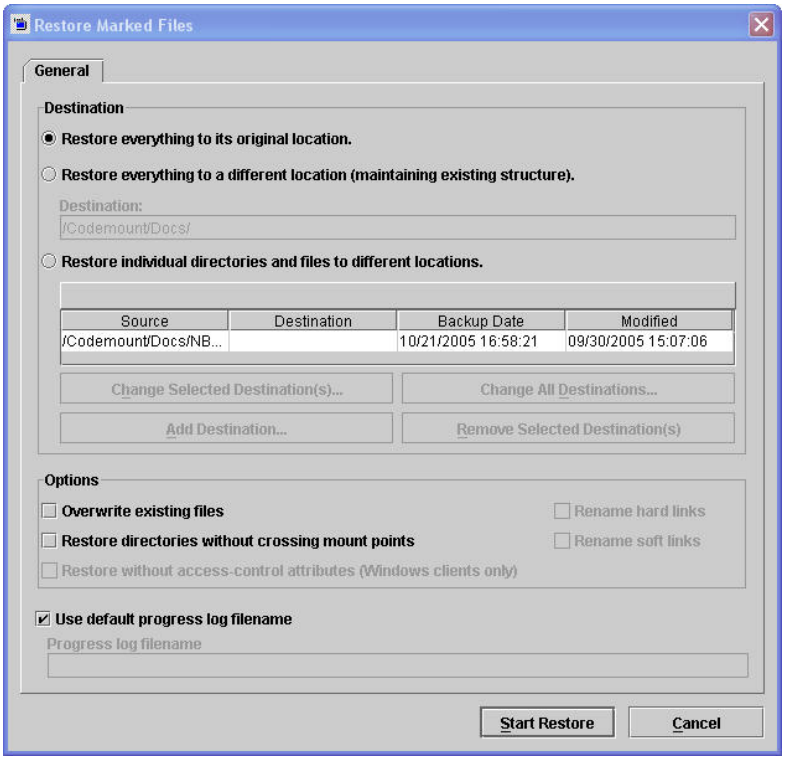

Figure 57. NetBackup Restore Marked Files GUI – Normal Backups.

File promotion (UNIX clients only)

Since the NSM backup utilized NAS_Snapshot as the method and Instant Recovery was selected, large files that have had many changes since the backup can be recovered more quickly by means of file promotion. File promotion optimizes single-file restore by using a minimum of I/O to recover the files.

Notes on file promotion:

- Only regular files can be promoted, not file links or directories.
- File promotion is available when restoring to the original volume on the original client.
- File promotion can be done from older snapshot copies, but any newer NAS snapshot copies are deleted after the file promotion takes place.
- The file system requirements depend on the NAS vendor.

To use file promotion

The procedure for restoring individual files with file promotion is the same as the standard restore procedure for NetBackup. No special settings or choices are required when using the Backup, Archive, and Restore interface.

NetBackup pelects file promotion on a file-by-file basis

If the above requirements are met (see “Notes on File Promotion”), NetBackup automatically attempts file promotion for the file. Otherwise, the restore of the file takes place in the standard manner, without file promotion: all file data is copied from the snapshot copy to the primary file system. The NetBackup progress log indicates how many files were promoted, and how many files that could not be promoted were restored by means of tar (intra-storage system ndmpcopy?).

If the entire volume which was backed up via NSM needs to be restored, select the *Point-In-Time Rollback Restore type*. Note that when using point-in-time rollback, specific directories and/or files “beneath” the backup volume cannot be specifically selected. Also note that if the backup selection was a qtree instead of a volume then *Point-In-Time Rollback* will not work.

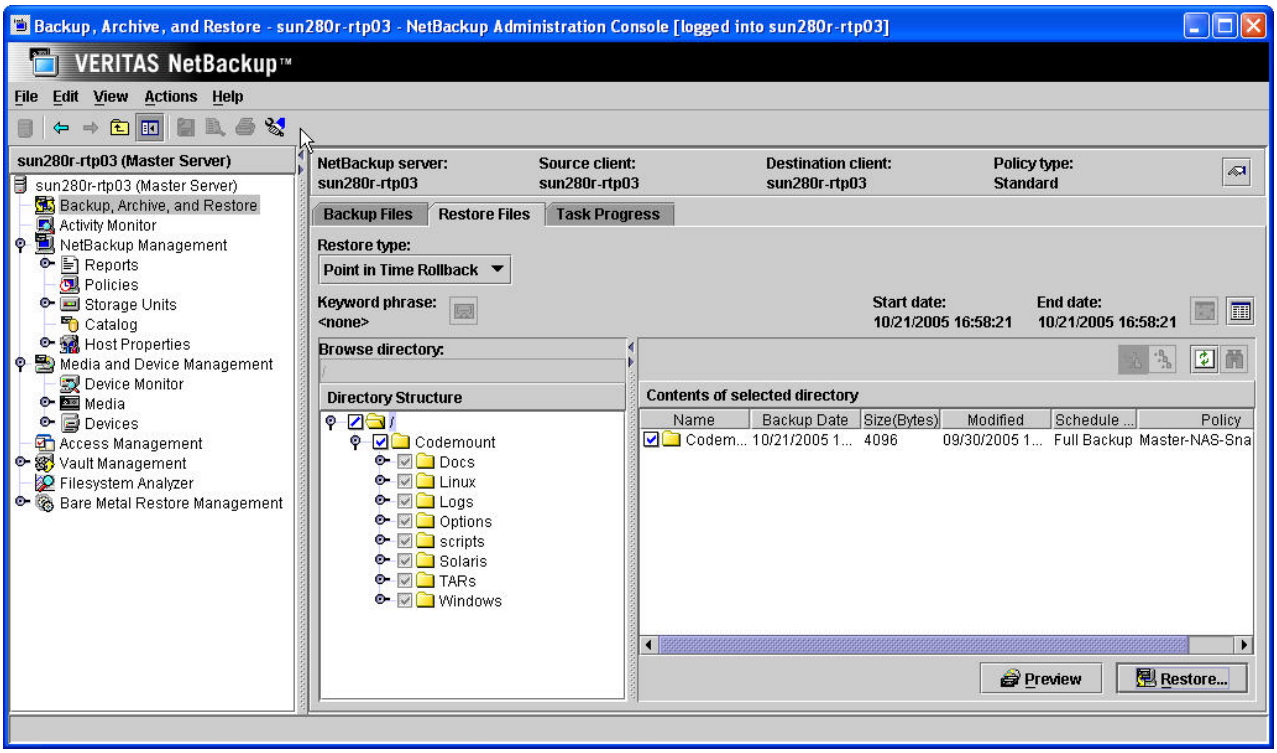

Figure 58. NetBackup Backup Archive Restore – Point In Time Rollback.

This uses SnapRestore for the point-in-time rollback, so despite “restoring” the entire volume it is very fast.

Note that upon completion of the restore, NSM backups which occurred after the snapshot copy to which the active file system was restored are deleted from the NetBackup catalog and the corresponding snapshot copies deleted.

NSVM configuration

This section contains specific steps about getting the NSVM solution up and running. It is assumed that in addition to the IBM N series storage array and the proper version of Data ONTAP being installed, the following items have also already been done:

Install / License NetBackup Enterprise Server
Install / License NetBackup Advanced Client
Install / License / Configure NDMP (on both IBM N series storage system and via NetBackup).

Additionally, it is assumed that the client data is the same as that being backed up in the previous example on NSM. I.e., the client data is the same and the Advanced Client set-up is nearly identical (where there is deviation it will be discussed/shown).

The basic NSVM environment looks like the following:

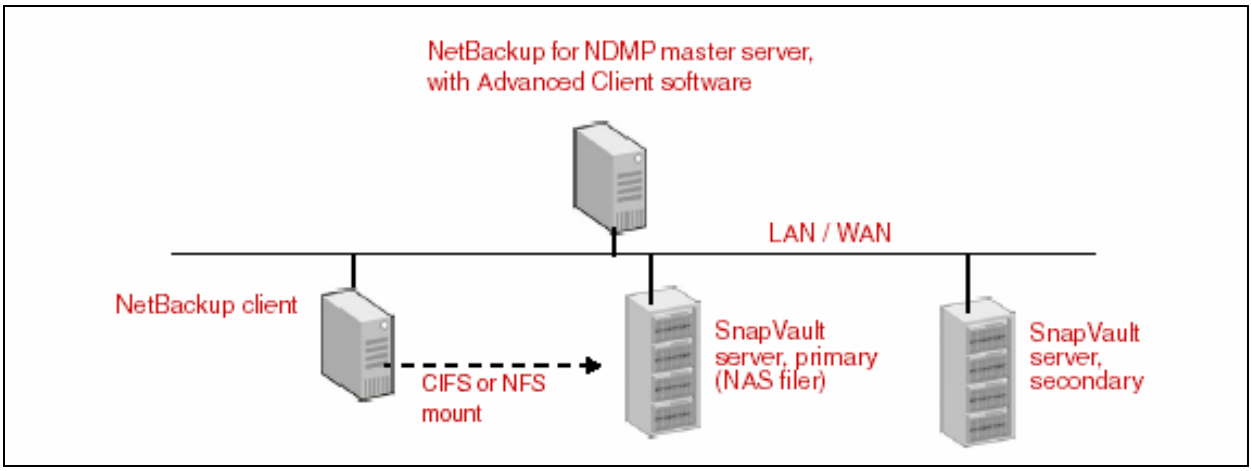

Figure 59. Basic NSVM Environment.

Note in the ensuing example:

The master server is also the client.
Master Server = sun280r-rtp03
SnapVault server, primary = fas3020-rtp01
SnapVault server, secondary = r200-rtp01
The client data is the same as that configured in the previous NSM example.

The remainder of this section steps through the following:

Configuring SnapVault on the primary and secondary servers
Making (and sharing a volume optionally) on IBM N series secondary server
Configuring NetBackup NSVM policy.

IBM N series NSVM configuration

Ensure NDMP authentication has been configured between the NetBackup master server (and any necessary media servers) and the IBM N series storage systems.

SnapVault on Primary

To make the IBM N series primary storage able to be backed up with NSVM, add, enable and configure SnapVault by executing the following at the primary command line.

Add the primary SnapVault license:

license add sv_primary_license

Enable SnapVault:

options snapvault.enable on

Grant access to media servers authorized to access the Filer by entering the following command:

options snapvault.access host=secondary,nbu_master_server,\ nbu_media_server1...

In this example the following is the set-up after the above is accomplished:

fas3020-rtp01> options snapvault
snapvault.access host= r200-rtp01,sun280r-rtp03
snapvault.enable on

SnapVault on Secondary

To make the SnapVault storage unit available for NSVM backups, add, enable and configure SnapVault on the IBM N series secondary 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 NearStore system by entering the following command:

options snapvault.access host=primary,nbu_master_server,\ nbu_media_server1...

In this example the following is the set-up after the above is accomplished:

r200-rtp01> options snapvault
snapvault.access host= fas3020-rtp01,sun280r-rtp03
snapvault.enable on

First create a volume on the IBM N series storage systems using FilerView. *SnapVaultDSU1* is the volume we've created for this example.

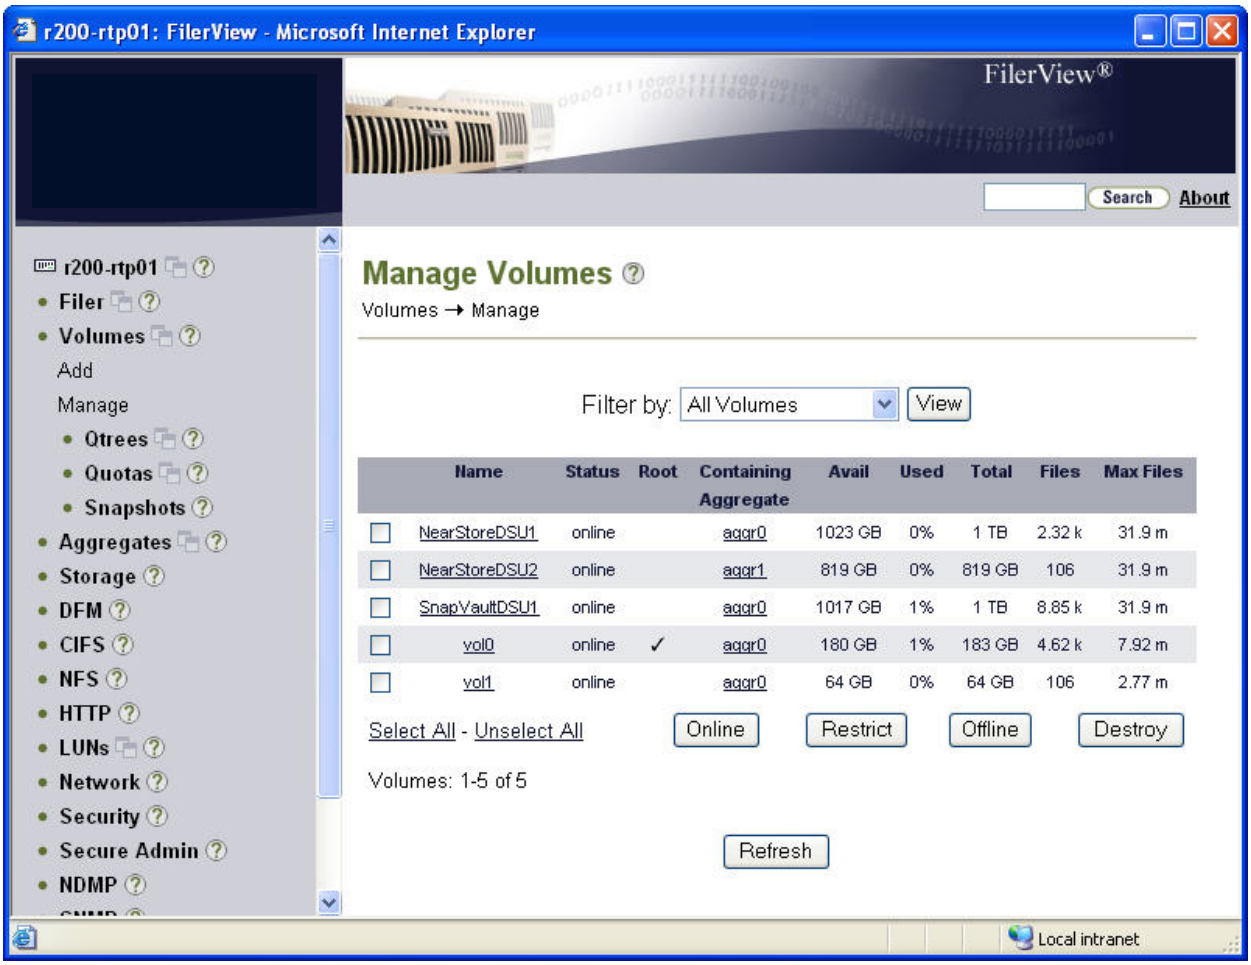

Figure 60. Creating a NSVM DSU Volume Via FilerView.

All the NSVM SnapVault backups are managed via NetBackup for both backups and restore. If desired, you can also make the backups available for drag-and-drop restores by exporting (NFS) or sharing (CIFS) the SnapVault volume on the secondary. In the example below, */vol/SnapVaultDSU1* is what we're exporting.

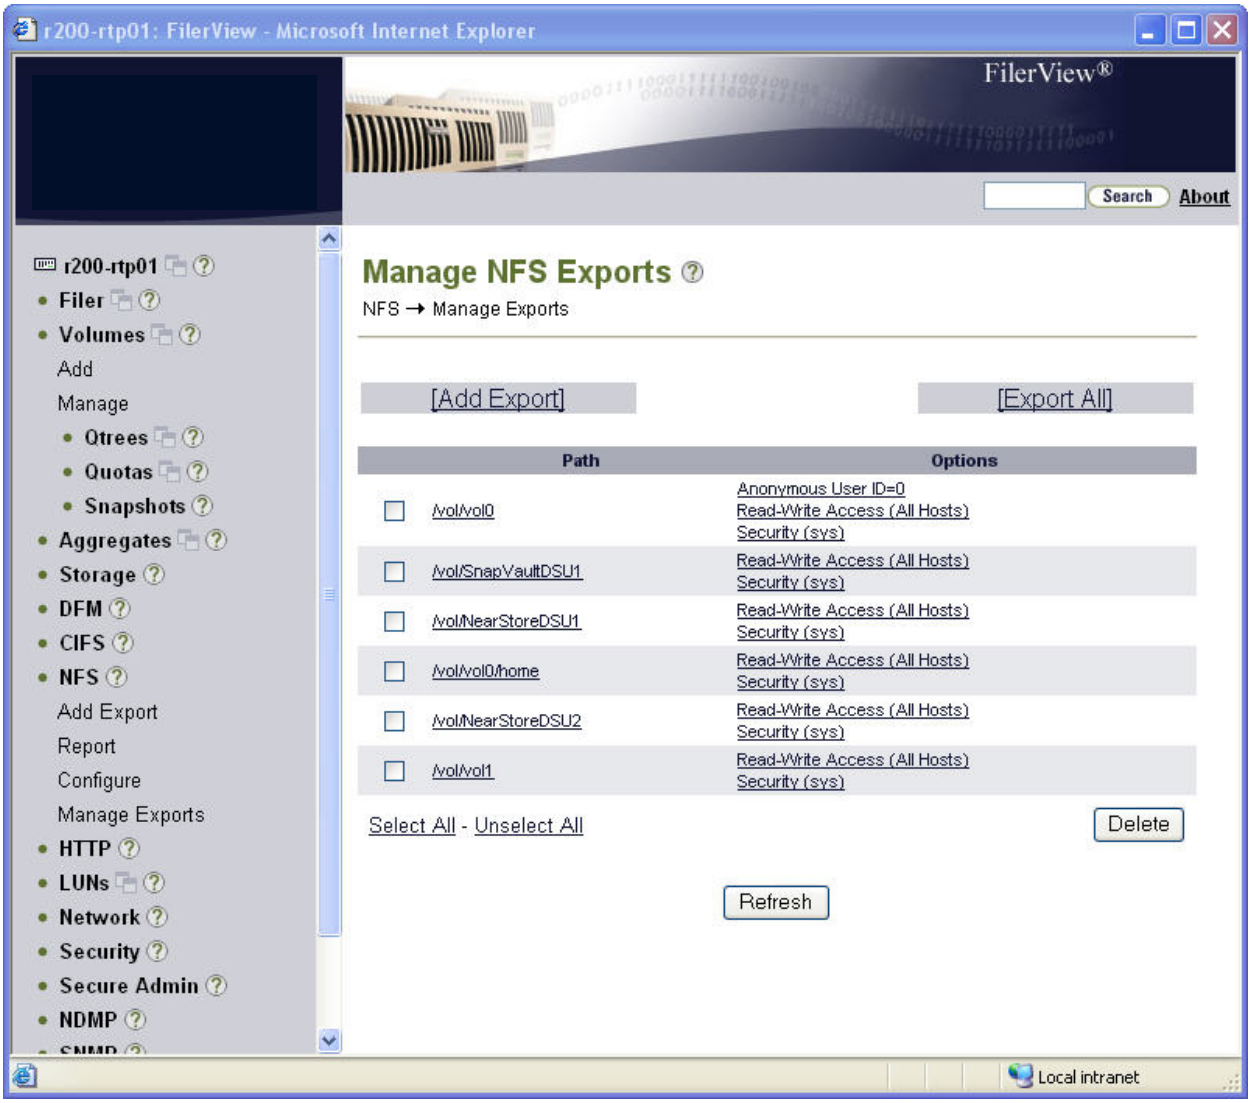

Figure 61. Sharing a NSVM DSU Volume for Drag-and-Drop Restores.

As NSVM is now effectively managing the snapshots for the volume (on the primary and secondary) there is no need to keep Data ONTAP scheduled snapshots too so a final step is to turn off scheduled snapshots for the volume on both the primary and secondary IBM N series storage system. While this can be accomplished from FilerView, it is typically accomplished using the CLI as it's quite simple:

r200-rtp01> vol options <volume> nosnap on

The IBM N series NSVM configuration is now complete.

NetBackup NSVM Configuration

The first thing to do is configure a storage unit which utilizes the volume which was created on the IBM N series secondary. Right-click on *Storage Units* and select *New Storage Unit*….

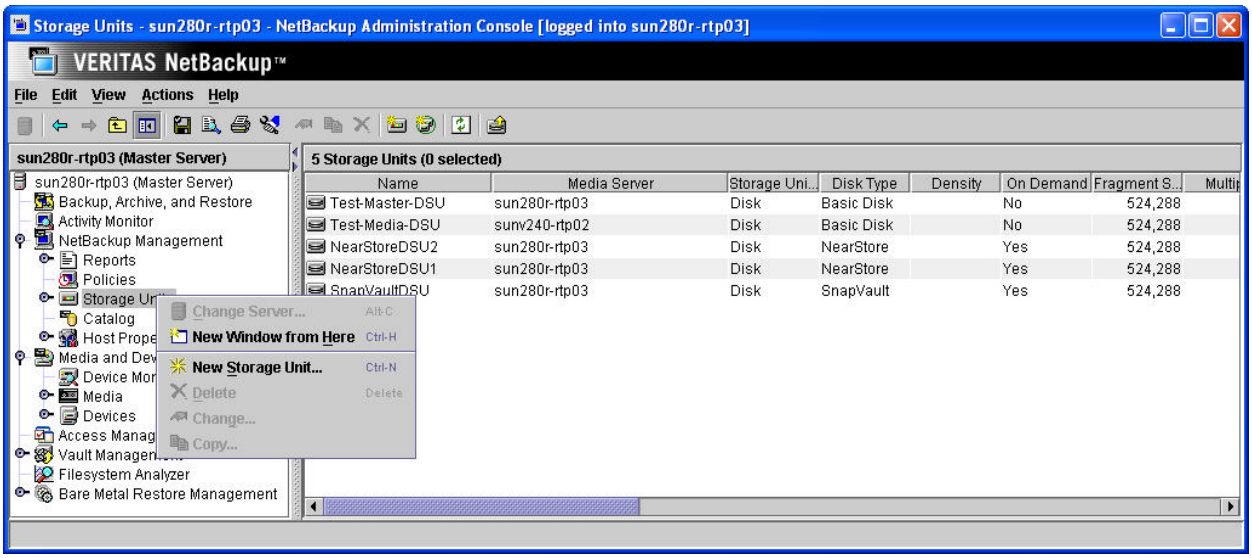

Figure 62. Creating a New Storage Unit.

The *Storage unit type* is set to *Disk*. *Disk type* is set to *SnapVault*. Select the appropriate *Media Server* (in this example, and often, this will be the Master Server); using the pull-down will provide a list of all the NetBackup Media Servers in the environment. The *SnapVault* server should be specified as the name of the IBM N series secondary storage system; using the pull-down will provide a list of all those devices which are NDMP authenticated with the previously selected *Media Server*. Finally, specify the *Absolute pathname to volume* as the volume which was created on the secondary; using the pull-down will provide a list of all those volumes which are on the *SnapVault server* specified in the previous step.

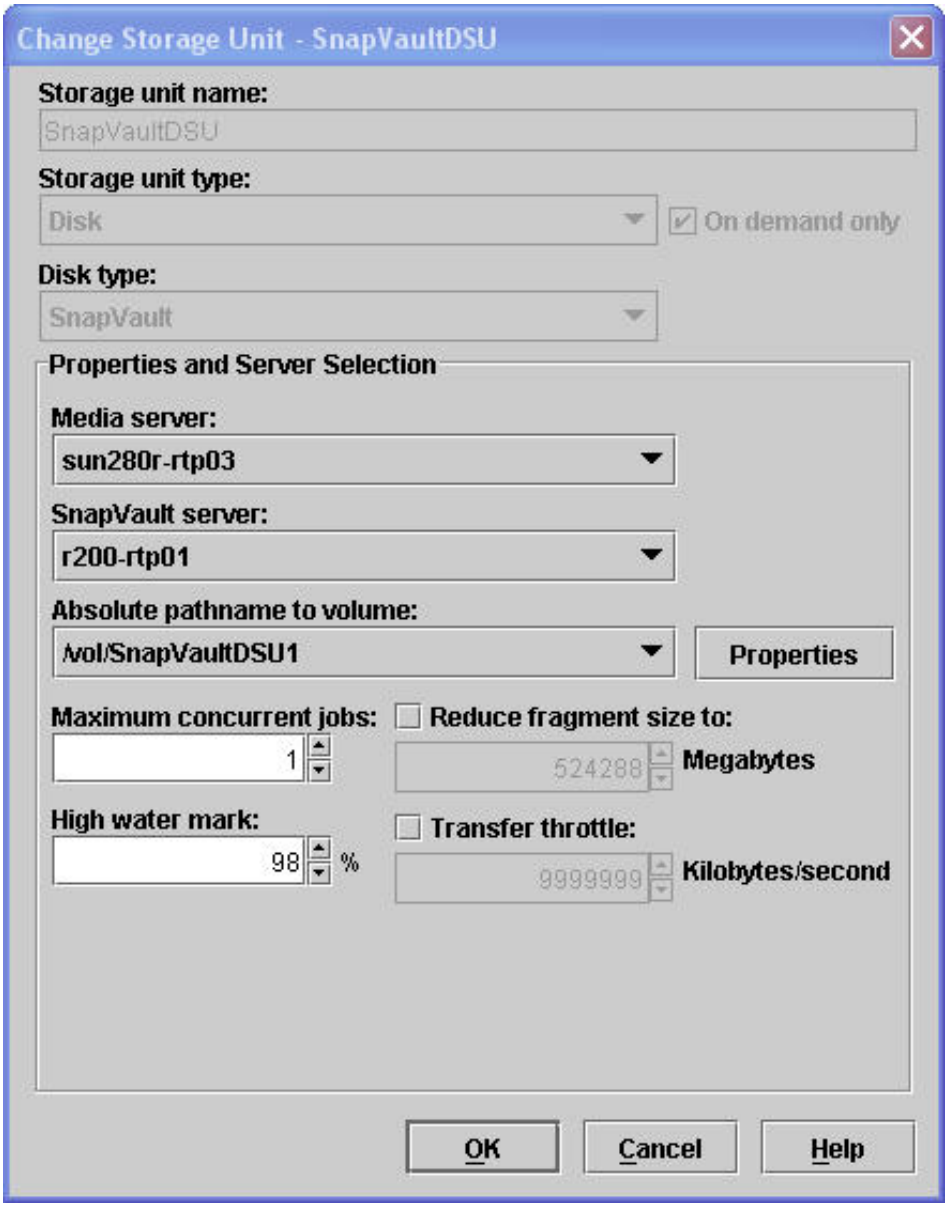

Figure 63. Configuring a SnapVault Storage Unit.

Note: When defining storage units, be aware that storage unit names are case-sensitive.

The next step is to configure a policy to use the newly created SnapVault storage unit. You can use the NetBackup Administration Console to configure a new policy by selecting *Policies* and then *Action* > *New Policy*. But in this example we'll simply copy the policy created in the NSM example.

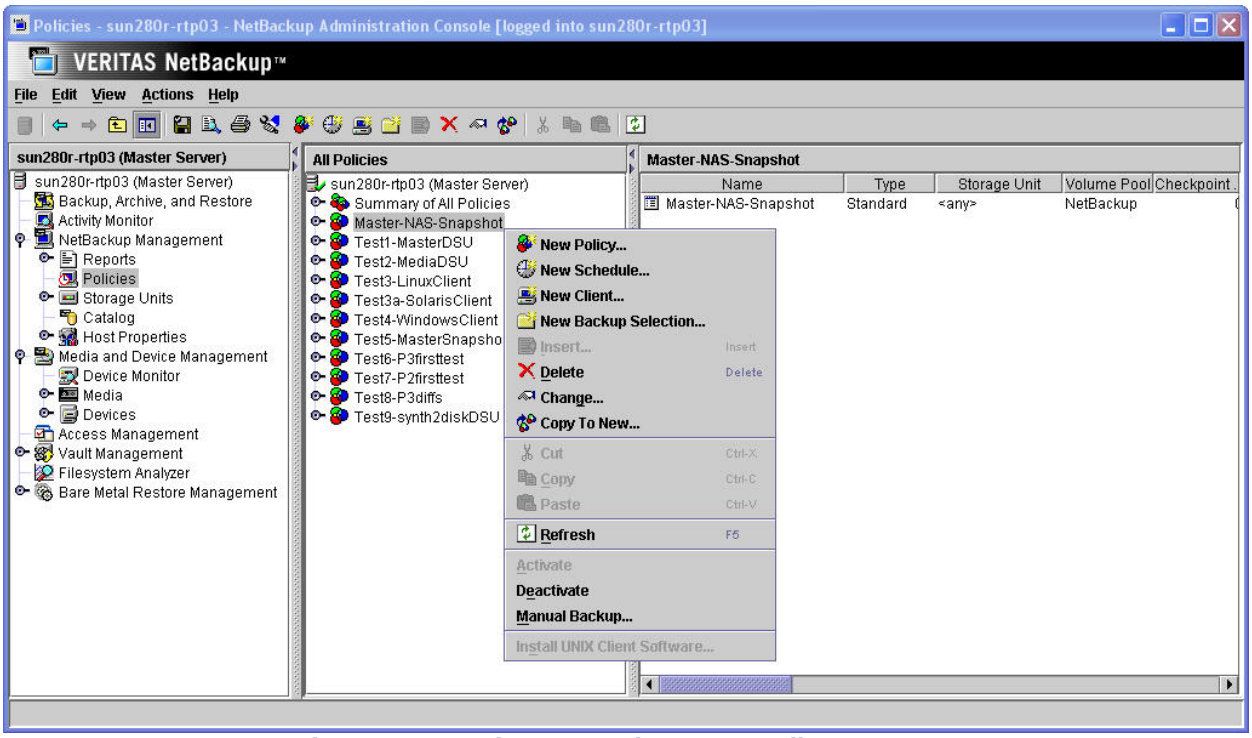

Figure 64. Copying NetBackup NSM Policy to NSVM.

First highlight the desired policy, and then right-click and select *Copy To New…*

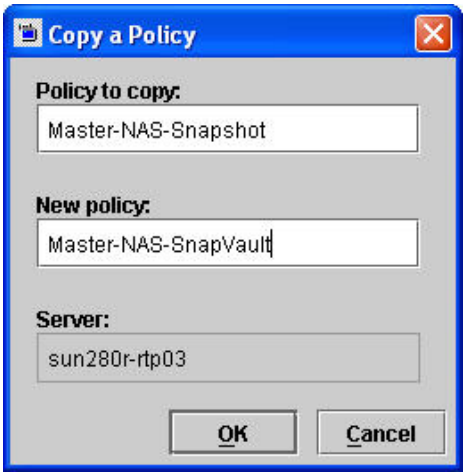

Figure 65. Naming the NetBackup NSVM Policy.

After specifying the new name, you'll need to modify the policy. The policy attributes are configured nearly identically to how they were for a NSM policy, with the following being selected:

Perform snapshot backups
Retain snapshots for Instant Recovery
Perform offhost backup
Use data mover (pull-down and specify Network Attach Storage).

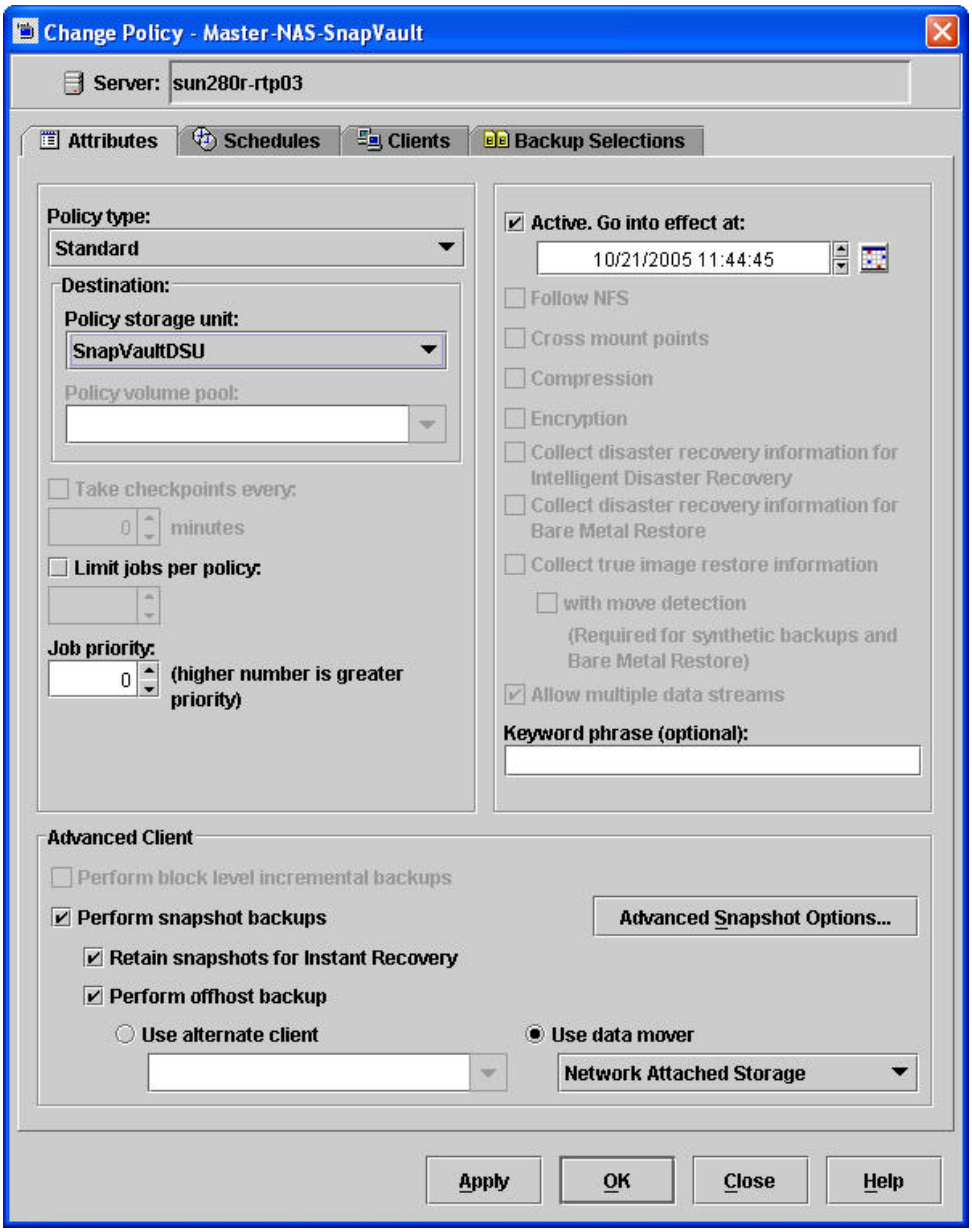

Figure 66. Configuring NetBackup NSVM Policy Attributes.

The key thing to configure differently from a NSM policy is to specify the *Policy storage unit* as the SnapVault storage unit which was created earlier.

It is also important to be aware that when selecting NAS_Snapshot for SnapVault backups, the *Maximum Snapshots (Instant Recovery only)* parameter specified under *Advanced Snapshot Options…* determines how many Snapshot copies can be kept on the SnapVault primary, not how many SnapVault copies are kept on the SnapVault secondary.

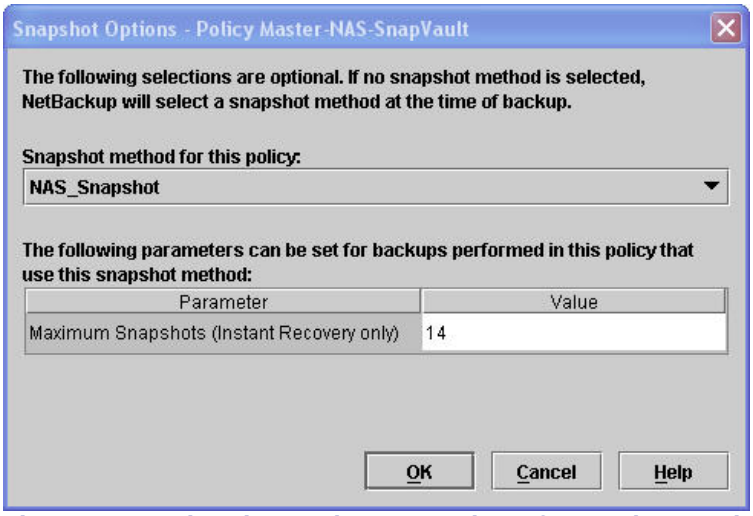

Figure 67. Setting the Maximum Number of Snapshot Copies to Keep on Primary.

On the primary, after n+1 snapshot copies are successfully created, then snapshot copies on the primary are managed by removing the oldest snapshot copy for the policy.

To set the retention for the NSVM backups (SnapVault copies) kept on the secondary, use the *Retention* field on the policy Schedule tab.

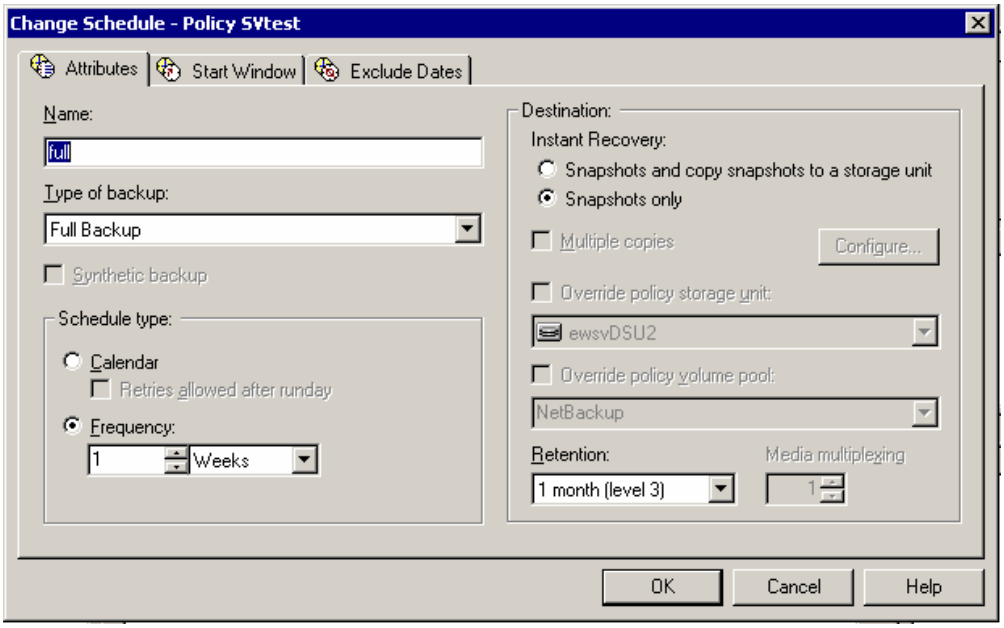

Figure 68. Setting the Retention for an NSVM Policy.

The *Clients* and *Backup Selections* tabs are configured identically to those defined for NSM backup policies.

The NSVM backup policy can be run manually or when the schedule is scheduled to execute. Use the Activity Monitor to watch the progress of the backup job – Job Id #65 shown below. Notice that for the previous NSM backup there was no *Storage Unit* listed, but for this NSVM backup there is.

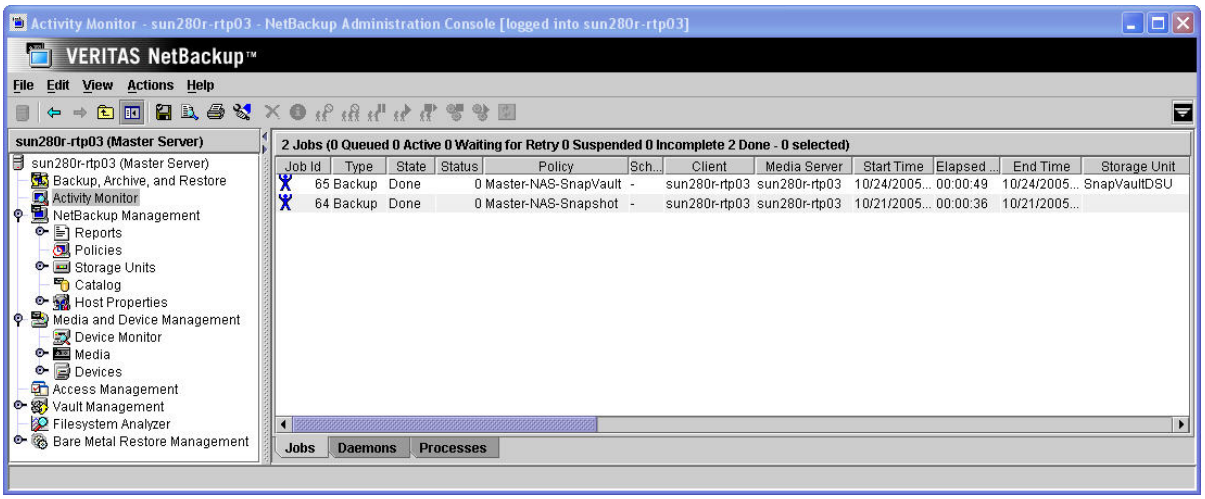

Figure 69. NetBackup Activity Monitor – NSVM Job.

Double-click on the job to see additional information on it, and select the *Detailed Status* tab to see all the details.

Figure 70. NetBackup Activity Monitor – NSVM Detailed Status.

NSVM restores

NSVM restores are accomplished the same way NSM restores are done (see “NSM Restores”), with the exception that *Point-In-Time Rollback* is not available since NSVM backs up a qtree and not an entire volume. After accessing the Backup *Archive Restore GUI*, you can click on the backup history icon -  - to select the image from which to restore.

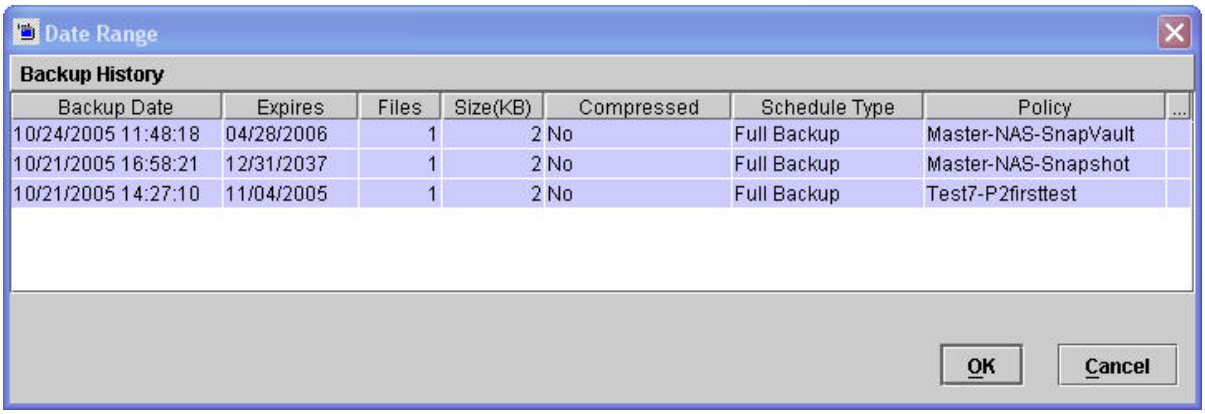

Figure 71. Backup Archive Restore Backup History.

Behind the scenes, restores from the secondary SnapVault storage system are accomplished via SnapVault Restore when appropriate.

It is important to remember that the NSVM backup caused a snapshot to be created both on the primary system (controlled by the “`Maximum Snapshots`” parameter of the policy’s Advanced Client settings) as well as the secondary (controlled by the “`Retention`” for the policy schedule.

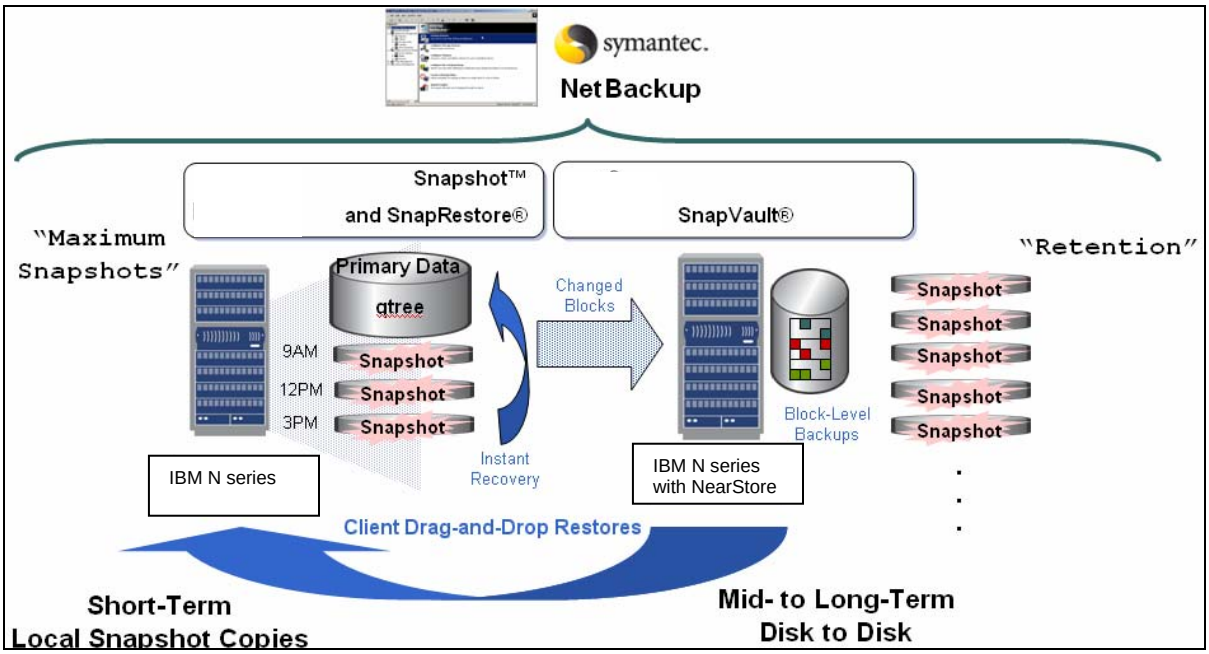

Figure 72. NSVM Optimized Restore.

When a restore is accomplished via NetBackup it does what is called an “Optimized Restore.” This means that it will try to do the restore from a snapshot on the primary first, and only if it has to will it restore the data from the SnapVault secondary.

SV-NBU configuration

This section contains specific steps about getting the SV-NBU solution up and running. It is assumed that in addition to the IBM N series storage array and the proper version of Data ONTAP being installed, the following items have also already been done:

Install / License NetBackup Enterprise Server
Install / Configure NDMP (on IBM N series 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 storage system since this enables the NearStore 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 ensuing example:

The master server is also the client.
Master Server = sun280r-rtp03
NearStore secondary = r200-rtp01
Client data is /usr/openv on the Master Server.

The remainder of this section steps through the following:

Configuring SnapVault on the secondary server
Making a volume on IBM N series secondary server
Configuring NETBACKUP SV-NBU Policy.

IBM N series SV-NBU configuration

Ensure NDMP authentication has been configured between the NetBackup master server (and any necessary media servers) and the NearStore storage system.

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, add, enable and configure SnapVault on the IBM N series secondary 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 NearStore system by entering the following command:

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

In this example the following is the set-up after the above is accomplished:

```
r200-rtp01> options snapvault
snapvault.access   host= fas3020-rtp01,sun280r-rtp03
snapvault.enable   on
```

First create a volume on the IBM N series storage system using FilerView. (Note that it must be a FlexVol™ volume.) When creating the volume, adhere to the following maximum usable volume size guidelines. (Refer to the size limitations in the Operating Characteristics section of this document for the technical reason for the restrictions.)

N5200	N5500 / R100	R200 / R1
1 TB	2 TB	4 TB

Table 2. Maximum SV-NBU Volume Sizes.

NearStoreDSU1 is the volume we've created for this example.

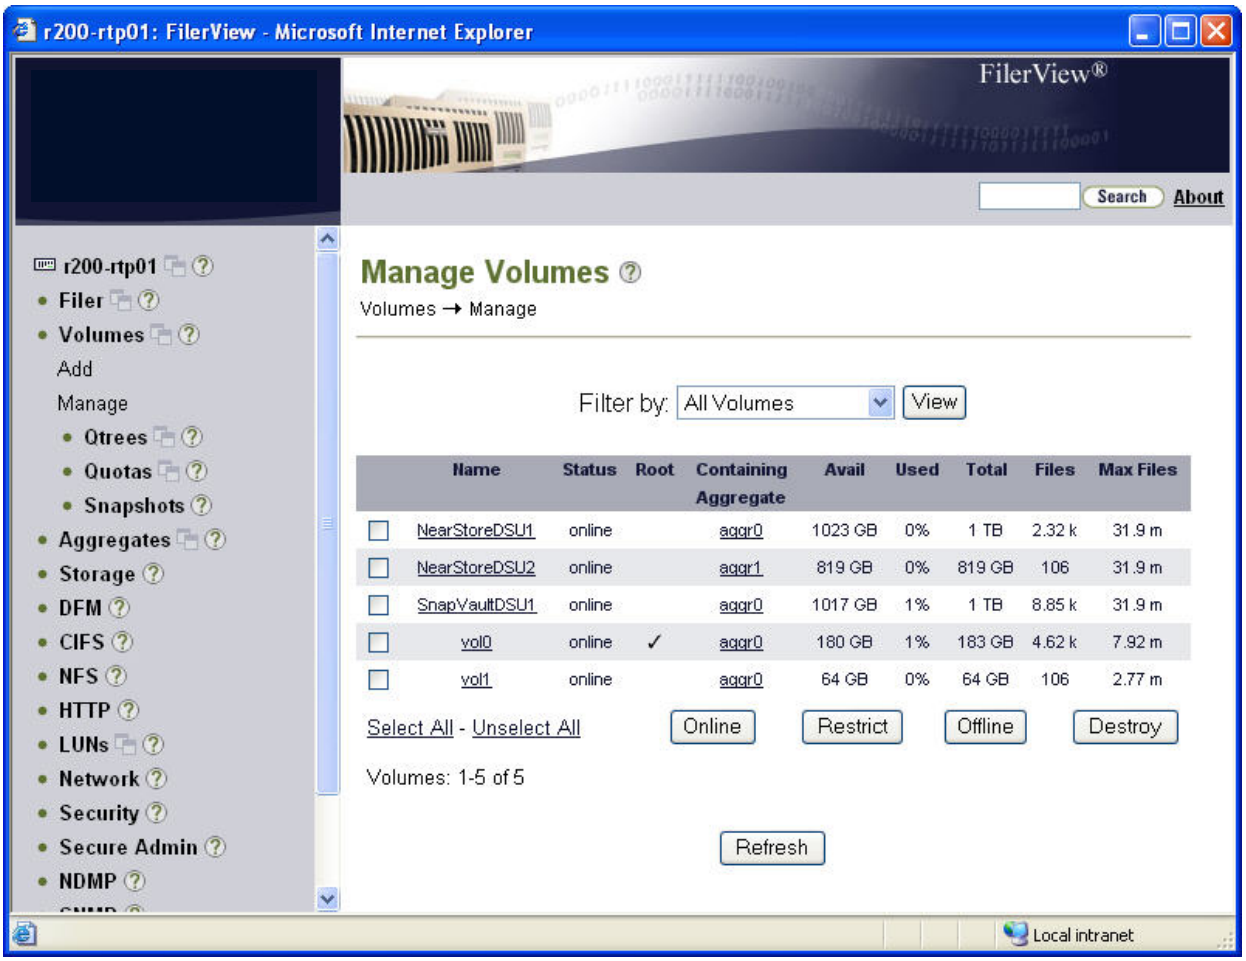

Figure 73. Creating a SV-NBU DSU Volume Via FilerView.

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

NetBackup SV-NBU configuration

The first thing to do is configure a storage unit which utilizes the volume which was created on the IBM N series secondary. Right-click on *Storage Units* and select *New Storage Unit*.

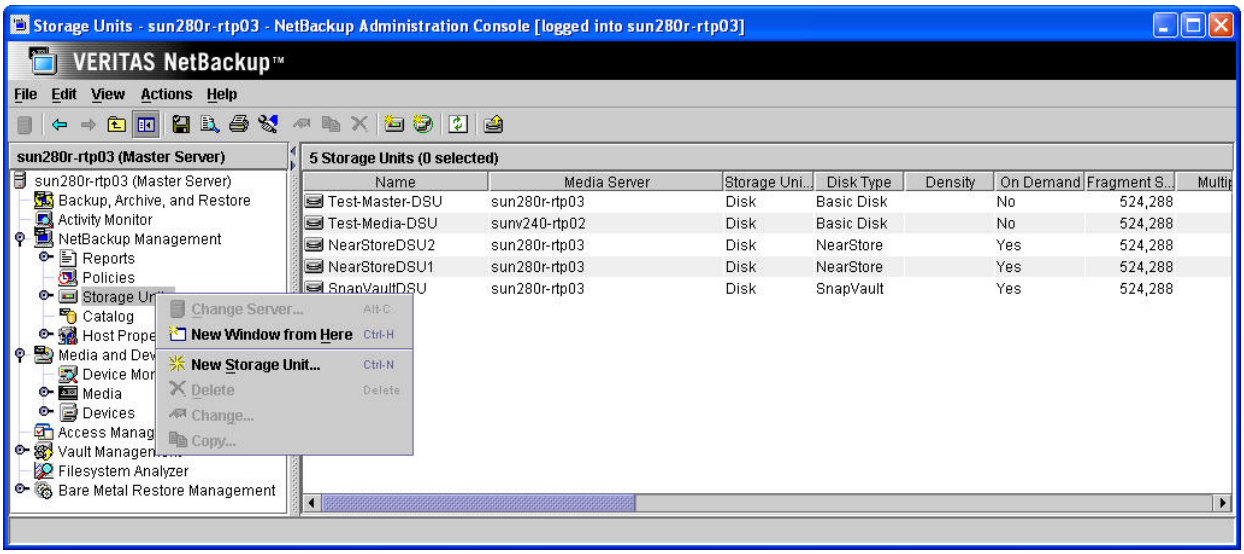

Figure 74. Creating a New Storage Unit.

The *Storage unit type* is set to *Disk*. *Disk type* is set to *NearStore*. Select the appropriate *Media Server* (in this example, and often, this will be the Master Server); using the pull-down will provide a list of all the NetBackup Media Servers in the environment. The *NearStore* server should be specified as the name of the IBM N series secondary storage system; using the pull-down will provide a list of all those devices which are NDMP authenticated with the previously selected *Media Server*. Finally, specify the *Absolute pathname to volume* as the volume which was created on the secondary; using the pull-down will provide a list of all those volumes which are on the *NearStore server* specified in the previous step.

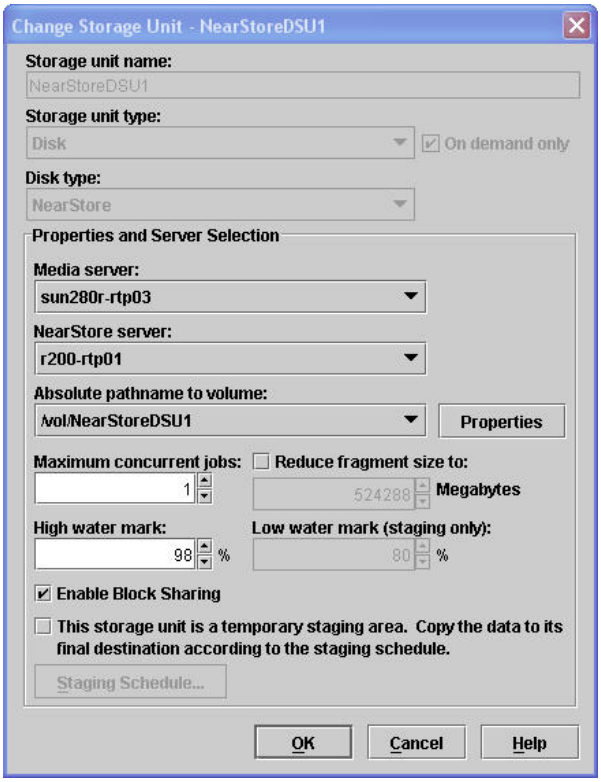

Figure 75. Configuring a NearStore Storage Unit.

Note: When defining storage units, be aware that storage unit names are case-sensitive.

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 *Action* > *New Policy*.

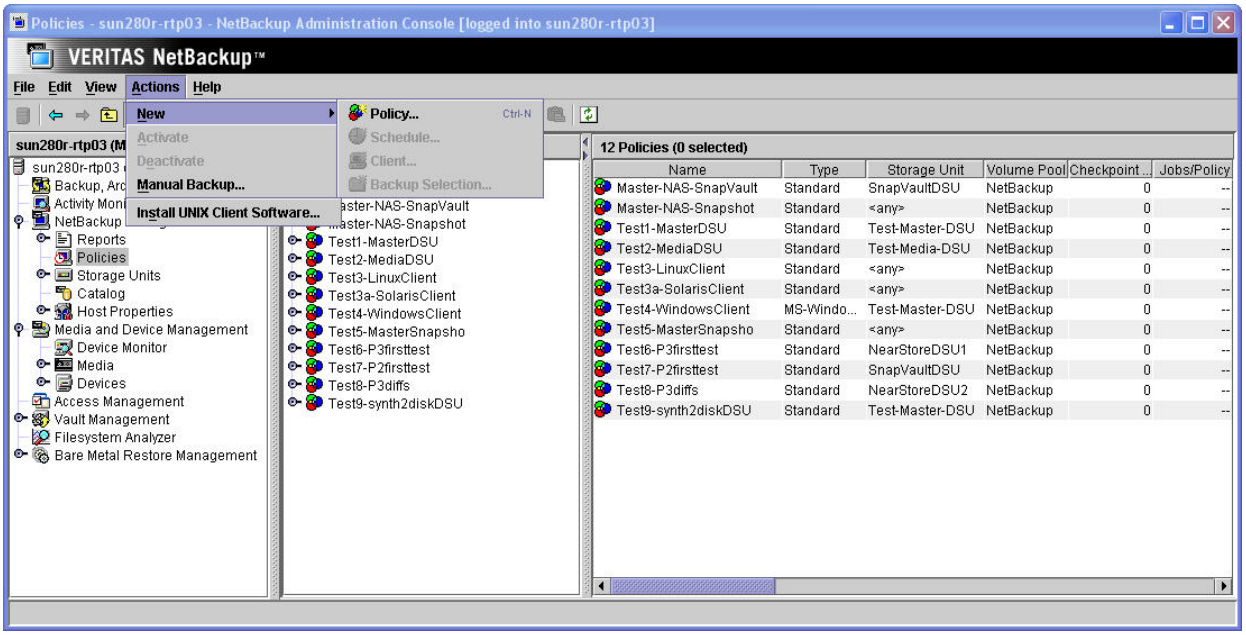

Figure 76. Creating NetBackup SV-NBU Policy.

Name the policy.

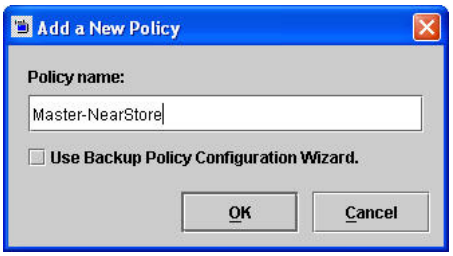

Figure 77. Naming the NetBackup SV-NBU Policy.

After specifying the new name, you'll need to modify the policy. The default policy attributes can all be utilized, except we need to specify the NearStore storage unit created earlier as the *Policy storage unit*.

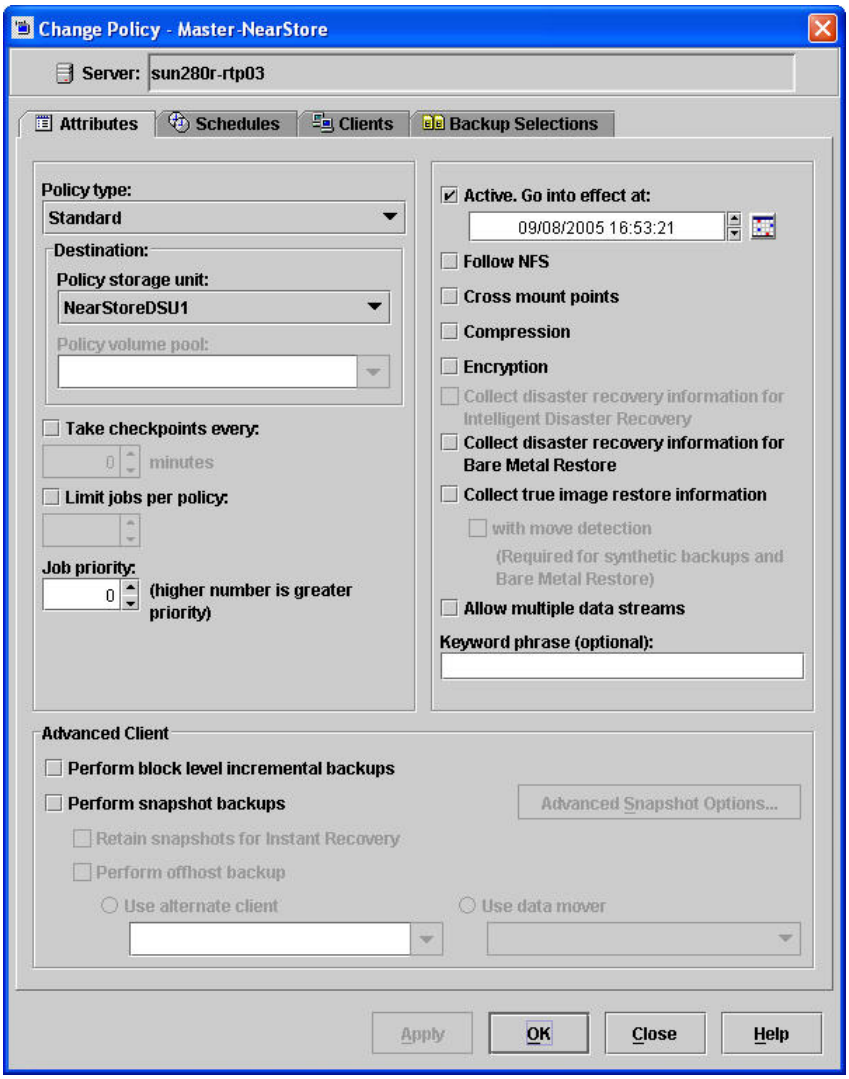

Figure 78. Configuring NetBackup SV-NBU Policy Attributes.

Now we need to create a schedule for the policy.

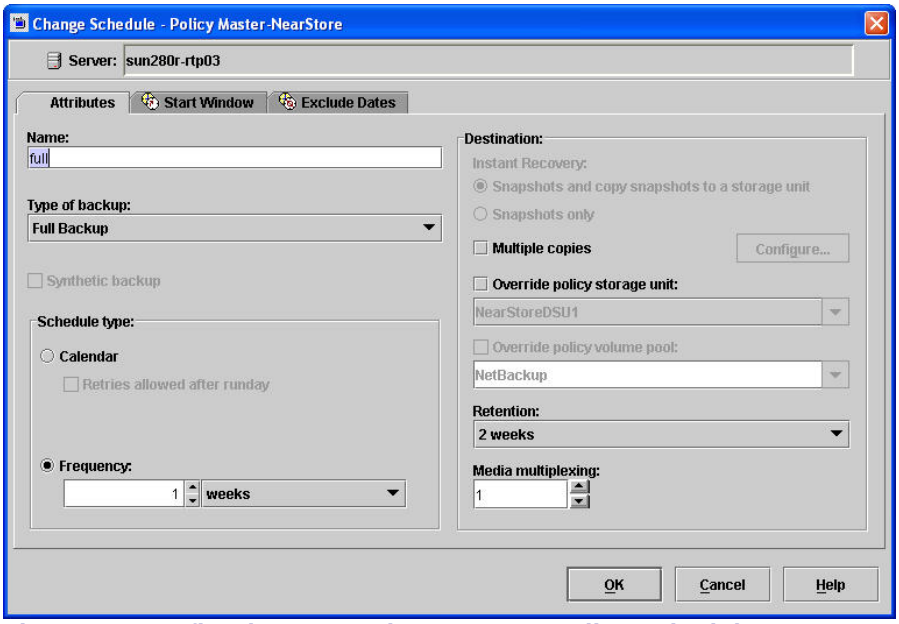

Figure 79. Configuring NetBackup SV-NBU Policy Schedule.

And then specify the client(s).

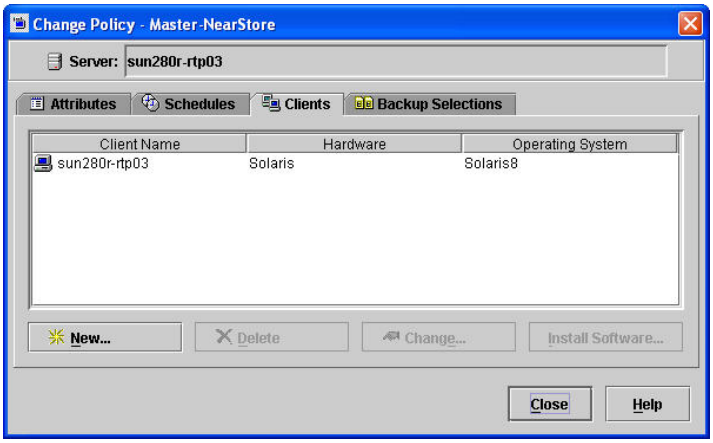

Figure 80. Configuring NetBackup SV-NBU Policy Client(s).

And finally the backup selections are specified.

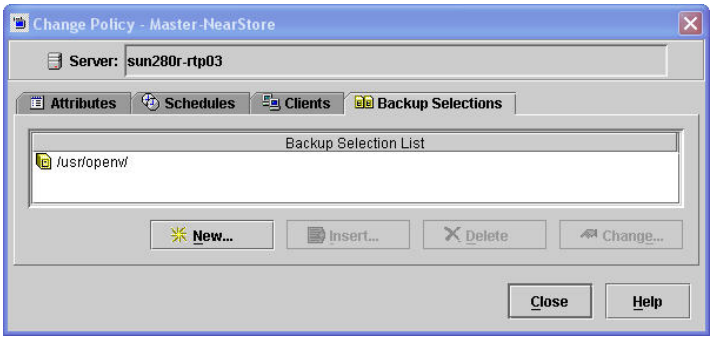

Figure 81. Configuring NetBackup SV-NBU Policy Backup Selections.

The SV-NBU backup policy can be run manually or when the schedule is scheduled to execute. Use the Activity Monitor to watch the progress of the backup job – Job Id #66 shown below.

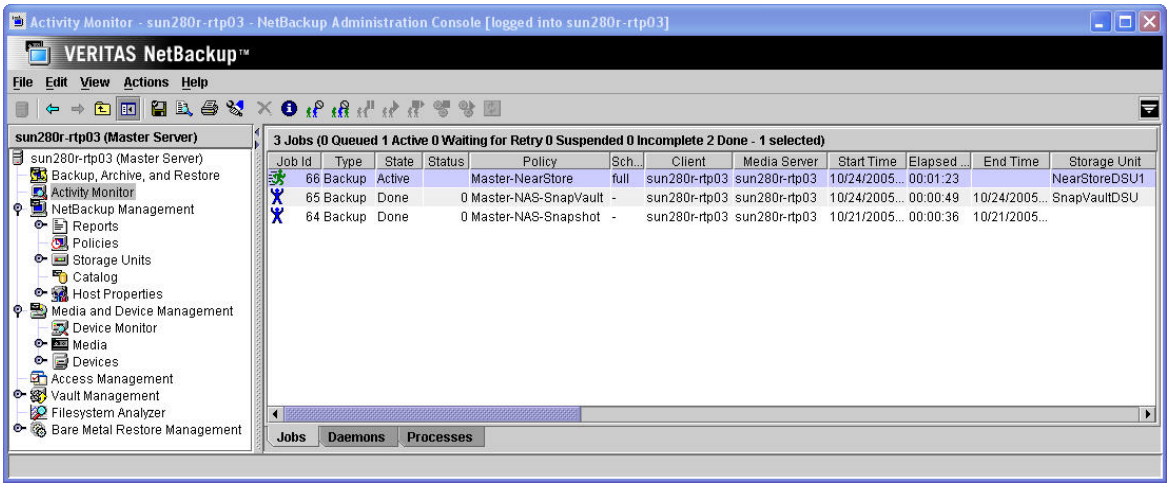

Figure 82. NetBackup Activity Monitor – SV-NBU Job.

Double-click the job to see more information on it, and select the *Detailed Status* tab to see all the details.

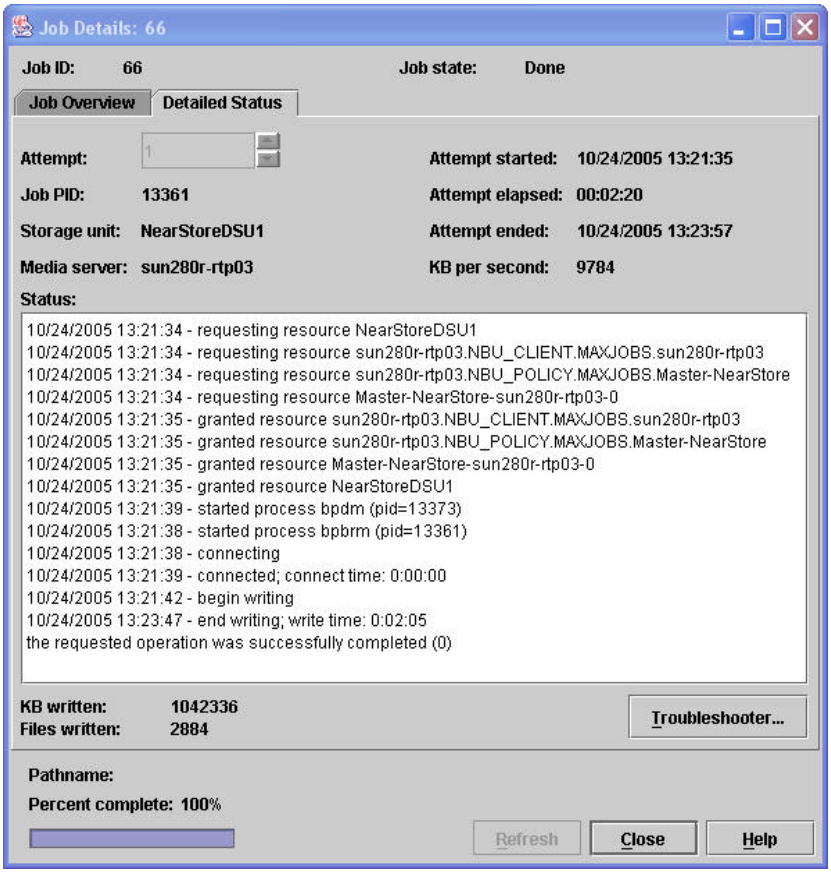

Figure 83. NetBackup Activity Monitor – SV-NBU Detailed Status.

After the first time a SV-NBU backup is accomplished to a volume on the NearStore system, the volume will show the "`sis`" – Single Instance Storage – attribute; see the `volume status` command to verify this.

```
r200-rtp01> vol status
         Volume State      Status            Options
  NearStoreDSU2 online     raid_dp, flex     create_ucode=on
                           sis
           vol0 online     raid_dp, flex     root
  SnapVaultDSU1 online     raid_dp, flex
  NearStoreDSU1 online     raid_dp, flex     create_ucode=on
                           sis
           vol1 online     raid_dp, flex     guarantee=none
```

After numerous SV-NBU backups are accomplished, you can use the "`df –s`" command to examine what sort of space savings are occurring on the NearStore system.

```
r200-rtp01> df –s
Filesystem               used     shared     saved       %saved
/vol/NearStoreDSU2/      3732          0         0           0%
/vol/vol0/            2768068          0         0           0%
/vol/SnapVaultDSU1/   6833656          0         0           0%
/vol/NearStoreDSU1/   2260928    1910232   2907752          56%
/vol/vol1/               3088          0         0           0%
```

The *shared* column shows how much of the used disk space is shared. The disk space savings generated by the *shared* space are shown in the *saved* column. The space *used* plus the space *saved* would be the total disk space usage, if no space was shared. The *%saved* is calculated as [*saved* / (*used* + *saved*)] * 100.

The space savings indicated by `df –s` can be a little deceptive after the first backup, since it will already show space savings (typically a few percent less than 50%). This is because `df -s` also takes into account the shared blocks in the expanded file system. Since each file is unpacked, and the unpacked files share the blocks with the tarball, they all have a block reference count of 2.

SV-NBU restores

SV-NBU restores are accomplished just like any other NetBackup restore, via the Backup Archive Restore GUI.

Operating characteristics

This section discusses the behavior one can expect of the solutions. Information included in this section comes from testing, observations, and knowledge of how the solutions are built.

The Performance sections will discuss the performance understandings of the various solutions. A lot of detailed performance-related testing has not occurred for any of the solutions, but continued testing and benchmarking will make this a much more robust and deterministic section going forward.

The Storage Overhead sections discuss storage savings the solutions can be expected to deliver. Again, more testing is on-going and results will continue to be delivered and enhanced in this document.

The Limitations sections discuss: what's not supported, what the reader 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 sections contain those items which might not have been covered elsewhere in this document, as well as (*in future revisions*) lessons learned via real-world implementations.

NSM

NSM performance

The performance characteristics are essentially those of Snapshot and SnapRestore.

Backup performance = Snapshot = matter of seconds!
Restore performance = Point-in-time rollback of entire file system = matter of seconds
= Other files and directories use intra-storage system ndmpcopy

NSM storage considerations

When it comes to storage overhead and NSM there are two areas of concern:

The first is the space allocation the snapshot copies will require. Even though snapshot space allocation is managed automatically by the system, it is good to understand the incremental disk space being consumed by the snapshot copies the NSM backups create, as this will impact the sizing of the storage.

NSM uses snapshots to maintain point-in-time or archival images of an active file system volume. An initial snapshot locks down the complete contents of a volume represented by the snapshot, but without any additional storage overhead beyond the 4k blocks used by the volume. Subsequent NSM archival snapshots require storage equal to the number of 4K blocks that changed between the previous and current snapshot. Therefore the total storage overhead required for NSM is the sum of 4k storage blocks locked down in snapshots that do not exist in the active file system volume.

There are 255 snapshot copies per volume allowed in Data ONTAP. Because of other services however, the recommended maximum number of snapshot copies available for backups is 250. Thus if you have a weekly full NSM backup for a particular volume, you could keep them around for about 5 years (5 x 52 = 260). For an increased Recovery Point Objective (RPO),

you could do a full every day, but then you could only keep about 8 months' worth of backups (250 / 30 = 8.3).

NSM limitations

We'll start by including a table of exactly what is supported for the NSM solution.

IBM N series Hardware	Any IBM N series, V-Series, or R-Series
Data ONTAP	Data ONTAP 7.1 or later
IBM N series Software	Snapshot, SnapRestore
NetBackup	NetBackup Enterprise 5.1MP2 or later NDMP Option Advanced Client Option[5]
Protocols	NFS, CIFS
Client	Solaris Windows NetBackup 5.1 or later
Media Server	N/A
Applications	File Sharing Oracle (8i or later) on Solaris

Table 3. NSM Solution Requirements Overview.

There is no integrated automation of moving/copying snapshot copies to tape, although a separate NetBackup NDMP policy could be configured and scheduled to run.

The number of snapshot copies limited to 255 per volume.

The following notes, taken from a NetBackup administration guide, apply to the NSM solution:

NetBackup will not restore to the root volume of a IBM N series filer by means of file promotion (SnapRestore), because SnapRestore causes the filer to reboot, thus disrupting service. Instead, you can "snaprestore" to a root volume using the SnapRestore command line tools.

Currently, the IBM N series with Data ONTAP operating system limits snapshots to 255 per volume. Note, however, that NetBackup Advanced Client controls the maximum number of NetBackup snapshots on a per client/per policy basis, using the policy's **Maximum Snapshots (Instant Recovery only)** parameter. When the configured maximum is reached, the oldest snapshot is deleted prior to creating the next snapshot.

If there are open references to a file (such as from snapshots or Oracle open file handles), a restore of the file cannot be done by file promotion "SnapRestore"). As a result, restoring the file may take longer.

Removing a file from the primary file system (such as with the UNIX rm command) will not increase disk free space if the file's blocks are referenced by one or more snapshots. The snapshot(s) referencing that file must be deleted before the file can actually be removed.

The NAS_Snapshot method is a copy-on-write type, which requires additional disk space for storing changes made to the client's data during the life of the snapshot. In Data ONTAP, this space is called snap reserve, and is configured on the NAS filer (not through NetBackup). The amount of space needed for snap reserve depends upon how much data is changed during the lifetime of the snapshot: the more data that changes, the more snap reserve space

[5] Advanced Client is not supported with NetBackup Server (need NetBackup Enterprise Server).

required. For Data ONTAP, the default snap reserve is 20% of the file system or volume. IBM recommends 10% for large ATA disk drives. NAS_Snapshot and write operations will fail if the snap reserve space is insufficient for the data change activity.

Unlike traditional NetBackup, NSM does not catalog the names of files contained in a particular backup. Therefore in NetBackup the administrator cannot determine which files changed between two backups. If there is a need to accomplish that functionality, refer to the “Finding the Snapshot you need” section in the IBM N series with Data ONTAP 7.1 Online Backup and Recovery Guide.

NSM best practices

For NSM you only need to put one Advanced Client license on one client, and many home directories (e.g.) can be protected (restored via Alternate Client Restore).

Even though for UNIX NSM only supports Solaris NFS clients, if other types of UNIX clients have their data mounted from the same location it will be backed up via the snapshot copy as well.

For NetBackup releases prior to 6.0: when configuring NAS volumes on the filer for NetBackup Windows clients, set the volume language type to en_US.UTF-8 (this is the UNICODE filer language). The default is POSIX, which is not appropriate for Windows. If the volumes are not configured to en_US.UTF-8, subdirectory and file names may not appear at all when browsing NetBackup snapshots for restore.

If the NAS volume was not configured with the correct language before the Windows client NAS_Snapshot was created, set the volume’s language to en_US.UTF-8 and then reboot the filer to make the change effective. When browsing for restore from the next NAS_Snapshot, directories and file names should display correctly.

Because the Data ONTAP operating system limits total snapshots to 255 per volume, consider disabling any unneeded scheduled Data ONTAP snapshots (configured with the `snap sched` command) when using the `NAS_Snapshot` feature on the same volume.

NSVM

NSVM performance

As with measuring performance for any product, there are may be external variables that could affect your specific performance throughput. In general, the clichés “it depends” and “your mileage may vary” could easily apply here, but some general rules of thumb are:

- It’s SnapVault, so only changed (4K) blocks are transferred, rather than the entire contents of changed files!
- If the data set being backed up is very dynamic (high data change rate), then backups will take longer because more 4K blocks will be transferred.
- Backup performance is also very dependent on the raw network throughput between the primary and secondary system. On a slow WAN or 10base-T network, backup times may not be acceptable.
- Data sets with a small number of large files will have better throughput than data sets with a large number of small files.
- Throughput is generally limited by CPU and disk I/O consumption on the destination.

In NSVM the SnapVault primary file system to be backed up is mounted and browsed by the NetBackup. Hence the overall NSVM backup performance also depends on the CPU of the NetBackup Media Server and the NetBackup Client, as well as the Network Bandwidth between the NetBackup Media Server/Client and the SnapVault primary.

The table below shows characteristics of various workloads when using SnapVault. Before a SnapVault transfer is performed, there is a scan that must take place to identify the files that have changed and the 4K blocks within those files that have changed. If there is a large amount of changed files, you will experience a higher scan time.

Workload	Time	Overall Throughput	Throughput after Scan Time	Primary CPU	Secondary CPU
Large Files Baseline Single Stream	0:02.38	63.43 MB/sec	95.39 MB/sec	54%	79%
Small Files Baseline Single Stream	0:07.43	17.02 MB/sec	25.82 MB/sec	21%	46%
Large Files Incremental (10% change) Single Stream	0:00.18	55.68 MB/sec	105.38 MB/sec	44%	64%
Small Files Incremental (10% change) Single Stream	0:02.08	24.78 MB/sec	31.33 MB/sec	35%	49%
Large Files Baseline Two Streams	0:05.02	65.72 MB/sec (combined)	77.16 MB/sec	27% (each)	79%
Small Files Baseline Six Streams	0:28.34	27.58 MB/sec (combined)	44.06 MB/sec	20% (each)	71%
Large Files Incremental (10%) Two Streams	0:01.22	67.05 MB/sec (combined)	85.07 MB/sec	26-47%	86%
Small Files Incremental (10%) Six Streams	0:08.06	32.57 MB/sec	39.04 MB/sec	39% (each)	70%
Large Files Restore One Stream	0:05.46	28.96 MB/sec	59.87 MB/sec	51%	33%
Large Files Restore Four Streams	0:09.36	69.60 MB/sec	91.55 MB/sec	82-88%	58%

Table 4. SnapVault Performance Summary.

In the table above, you will see the various workloads used in the test, along with their respective change rates. The time is the amount of total time required for the transfer. Overall throughput is equal to the amount of data transferred divided by transfer time. The throughput after scan time is the throughput of the data after the scan of the files has completed.

For more information on the space savings, network savings, and throughput characteristics of SnapVault, please refer to the SnapVault Performance Report.

NSVM storage considerations

NSVM storage overhead is identical to native SnapVault.

The same considerations that were discussed for NSM apply to NSVM for the primary IBM N series storage system: space consumed by snapshot copies and number of snapshot copies.

Unlike NSM, when using NSVM the primary storage system is not the final repository for the backups. Therefore, keeping numerous snapshot copies on the primary storage system might not be a requirement. Minimally one “reference snapshot” must be maintained on the primary storage system so that SnapVault can determine the changed 4K blocks between backups.

The number of snapshot copies maintained on the primary storage system (and the storage overhead associated with them) is thus a function of the customer’s restore requirements – as restoring directly from primary resident snapshots will be much quicker than restoring from the secondary storage system.

NSVM uses snapshots to maintain “point in time” or archival images of an active file system volume on the primary storage system in a manner similar to NSM. NSVM uses a similar archival snapshot technique on the secondary storage system. NSVM is different in that the secondary storage system volume can be the final backup repository for multiple qtrees (either from the same primary storage system or from multiple primary storage systems). Therefore there are additional NSVM storage considerations that affect both storage consumption and snapshot copies consumption:

- Number of primary storage system qtrees backed up to this secondary volume. (Affects the # of changed blocks….)
- Number of backups retained per primary storage system. (Affects the # of snapshot copies consumed…)
- The retention policy for each primary storage system qtree/policy. (Affects the # of snapshot copies…)

There are 255 snapshot copies per volume allowed in Data ONTAP. Because of other services however, the recommended maximum number of snapshot copies available for backups is 250.

NSVM limitations

We'll start by including a table of exactly what is supported for the NSVM solution.

IBM N series Hardware	Any platform combination SnapVault supports
Data ONTAP	Primary: Data ONTAP 7.1 or later Secondary: Data ONTAP 7.1 or later
IBM N series Software	Snapshot, SnapRestore[6] Primary: sv_ontap_pri Secondary: sv_ontap_sec
NetBackup	NetBackup Enterprise 6.0 or later NDMP Option Advanced Client Option IBM N series SnapVault Option
Protocols	NFS, CIFS
Client	Solaris Windows NetBackup 6.0 or later
Media Server	N/A (but Master must be NetBackup 6.0 or later)
Applications	File Sharing Oracle (8*i* or later) on Solaris

Table 5. NSVM Solution Requirements Overview.

For the most complete information on SnapVault limitations in general (e.g., concurrent streams, maxdirsize, etc.), see the SnapVault deployment/configuration guide.

There is no integrated automation of moving/copying snapshot copies to tape, although a separate NetBackup NDMP policy could be configured and scheduled to run.

NetBackup manages qtree-to-qtree only; it does not manage:

SnapVault to copy a whole volume to a qtree,
SnapVault with non-qtree data, or
Open Systems SnapVault.

Number of snapshot copies limited to 250 per volume on both primary and secondary.

There are no specific limitations of volume size with NSVM.

The Windows Server Appliance Kit (SAK) is not supported.

Also see the limitations discussed in the NSM Limitations section.

The following notes, taken from a NetBackup administration guide, apply to the NSVM solution:

Currently, the IBM N series with Data ONTAP operating system limits snapshots to 255 per volume. Note, however, that NetBackup Advanced Client controls the maximum number of NetBackup snapshots on a per client/per policy basis, using the policy's **Maximum Snapshots (Instant Recovery only)** parameter. When the configured maximum is reached, the oldest snapshot is deleted prior to creating the next snapshot.

[6] SnapRestore not needed, but recommended for NSM restores.

If there are open references to a file (such as from snapshots or Oracle open file handles), a restore of the file cannot be done by file promotion (SnapRestore). As a result, restoring the file may take longer.

Removing a file from the primary file system (such as with the UNIX rm command) will not increase disk free space if the file's blocks are referenced by one or more snapshots. The snapshot(s) referencing that file must be deleted before the file can actually be removed.

The NAS_Snapshot method is a copy-on-write type, which requires additional disk space for storing changes made to the client's data during the life of the snapshot. In Data ONTAP, this space is called snap reserve, and is configured on the NAS filer (not through NetBackup). The amount of space needed for snap reserve depends upon how much data is changed during the lifetime of the snapshot: the more data that changes, the more snap reserve space required. For Data ONTAP, the default snap reserve is 20% of the file system or volume. IBM recommends 10% for large ATA disk drives. NAS_Snapshot and write operations will fail if the snap reserve space is insufficient for the data change activity.

Differences between native SnapVault and NSVM:

All the qtrees in a NetBackup policy are backed-up sequentially, so if a volume has 1000 qtrees there is currently no way to back the whole thing in one shot via NSVM. This limitation will be addressed in the future such that the behavior of NSVM will be equivalent to native SnapVault.

Replication of NSVM secondary storage system volumes for disaster recovery purposes is not supported.

NSVM does not support use of pre-existing snapshot copies on the primary storage system. (I.e., native SnapVault's "`snapvault update -s`" option is not supported.

Unlike traditional NetBackup, NSVM does not catalog the names of files contained in a particular backup. Therefore in NetBackup the administrator cannot determine which files changed between two backups. If there is a need to accomplish that functionality, refer to the "Finding the Snapshot you need" section in the IBM N series with Data ONTAP 7.1 online backup and recovery guide.

Volumes containing both traditional SnapVault qtrees and NetBackup NSVM-controlled qtrees are not supported. (However, you can migrate native SnapVault qtree relationships to NSVM – see the Previous SnapVault Relationships section.)

NSVM storage units cannot be used for disk staging or as part of a NetBackup storage unit group.

NSVM best practices

Like NSM, for NSVM you only need to put one Advanced Client license on one client, and many home directories (e.g.) can be protected (restored via Alternate Client Restore).

And, similarly, even though NSVM for UNIX only supports Solaris NFS clients, if other types of UNIX clients have their data mounted from the same location it will be backed up via Snapshot and SnapVault as well.

For NetBackup releases prior to 6.0: when configuring NAS volumes on the filer for NetBackup Windows clients, set the volume language type to en_US.UTF-8 (this is the UNICODE filer language). The default is POSIX, which is not appropriate for Windows. If the volumes are not configured to

en_US.UTF-8, subdirectory and file names may not appear at all when browsing NetBackup snapshots for restore.

If the NAS volume was not configured with the correct language before the Windows client NAS_Snapshot was created, set the volume's language to en_US.UTF-8 and then reboot the filer to make the change effective. When browsing for restore from the next NAS_Snapshot, directories and file names should display correctly.

Due to Data ONTAP limiting total snapshots to 255 per volume, consider disabling any unneeded scheduled Data ONTAP snapshots (configured with the snap sched command) when using the NAS_Snapshot feature on the same volume. On secondary use one of the following:

```
snap sched MYVOL 0 0 0
vol options MYVOL nosnap true
```

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

Do not rename the volume.

Policies with similar backup schedules (frequency and retention) should go to the same NSVM secondary storage system volume. (This is to ensure the least number of snapshot copies are consumed.)

SV-NBU

SV-NBU target environment

SV-NBU in Data ONTAP 7.1 is ideally suited and supported for file services backups – e.g., home directories, file shares, etc.

SV-NBU in Data ONTAP 7.1 is not supported for the backup of application data – e.g., databases, email.

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

1. Customer wants to back up non-IBM N series primary storage and uses NetBackup 6.
2. Customer desires easy online access to a long history of backups - SV-NBU's space savings technology provides a cost-effective way of achieving this.
3. A majority of the data to be backed up remains unchanged between full backups – typical of home directories, file shares.

SV-NBU in Data ONTAP 7.1 is NOT recommended for environments where any of the following are true:

1. Raw backup throughput is the primary concern.
2. Data to be backed up consists mostly of large (> 500 MB) Windows files
3. Data to be backed up consists mostly of small (< 10 KB) files.

The remaining sections relate details which support the above-mentioned positioning.

SV-NBU performance

SV-NBU in Data ONTAP 7.1 has been optimized for storage efficiency of file systems backup, while providing backup throughput comparable to a NetBackup Basic DSU solution.

SV-NBU in Data ONTAP 7.1 can achieve the following levels of throughput when backing up a typical file system data set to an R200 with Gigabit Ethernet network connectivity:

35 MB/s single stream (limitation is ability to source the data).
60 MB/s for multiple concurrent streams (NearStore CPU is limit).
Restore performance is comparable.

When determining what a specific environment can expect for performance, consider these factors:

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 that consist primarily of large files (typically greater than 500 MB) due to the structure of the Windows “backup read” formatted files delivered by NetBackup. This limitation does not apply to UNIX data streams.
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 through the use of block-sharing and de-duplication 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 one could expect / 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 table below 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 tests only included performing full backups.
The home directory and web server data sets were static.
The Windows C$ data set was a live NetBackup Master Server.

Data Set	Size	# Files	# Full Backu (No incr)	Storage Overhead	Compres-si	Storage Savir
Home Dir (static)	1.1 TB	5.3 M	8	6.8%	5.9:1	83%
Big Files (static)	16 GB	920	12	4.7%	8.3:1	88%
Big Files (static)	16 GB	920	162	4.7%	20:1	95%
Web Server (static)	147 GB	277 K	12	6.1%	7.7:1	87%
Windows C$ (dynamic)	5.8 GB	22 K	12	9.6%	5.9:1	83%
Windows C$ (dynamic)	10 GB	23 K	111	-	11:1	91%

Table 6. SV-NBU Storage Savings.

One of the key methods to examine storage savings which are occurring is to utilize the "`df –s`" command on the IBM N series storage system. The space savings indicated by `df –s` can be a little deceptive after the first backup, however, since it will already show space savings (typically a few percent less than 50%). This is because `df -s` also takes into account the blocks shared between the tar image and the unpacked files in the expanded file system. Since each file is unpacked, and the unpacked files share the blocks with the tarball, they all have a block reference count of 2.

SV-NBU storage overhead

As previously stated, SV-NBU offers outstanding storage savings with block-sharing and de-duplication 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. NOTE: The overhead for typical home directory environment is less than 10% of the 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 4K 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 4K blocks and is variable as a function of number of files and directories. The overhead is typically 3.5K per directory, 3.5-7K per UNIX file, and 7-10K per Windows file. Storage efficiency will be negatively affected by small files. This effect is exacerbated if small files comprise a significant percentage of the total data.

Using the 1.1 TB home directory example in the previous Storage Savings section, only 250 GB of total overhead was required for the ~3.3 TB of stored backup images, roughly 7%.

SV-NBU limitations

We'll start by including a table of exactly what is supported for the SV-NBU solution.

IBM N series Hardware	NearStore R200, R150, R100 N5200, N5500
Data ONTAP	Data ONTAP 7.1 or later
IBM N series Software	sv_ontap_sec nearstore_option (for N5xxx)
NetBackup	NetBackup Enterprise 6.0 or later (NDMP Option – installed, but not licensed) Disk Optimization Option
Protocols	N/A
Client	(Whatever platform/OS NetBackup supports) NetBackup 5.0 or later
Media Server	(Whatever platform/OS NetBackup supports) NetBackup 6.0 or later
Applications	File Services

Table 7. 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. You will not be able to connect to the SV-NBU volume if the cluster license is present.

Backup of the SV-NBU volume to tape via NDMP or native dump is not supported. (Use NetBackup Inline Tape Copy or staging 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 volumes are supported. SV-NBU shouldn't share a volume with anything else, and cannot share one with native SnapVault.

There is no checkpoint restart. If a SV-NBU backup fails mid-transfer and a new transfer is started, it begins at the beginning of the transfer again. All data transferred prior to the failure is discarded.

FlexVol volumes 4 TB or smaller, depending on platform type. Maximum FlexVol size is function of platform memory:

Platform memory affects volume sizing because the `refcount` file for the SV-NBU volume must fit into memory to guarantee consistency when `WAFL_check` is run. (But since these processes run serially, you can have multiple volumes.) This translates into the following usable volume sizes for the associated IBM N series storage systems:

N5200	N5500 / R100	R200 / R150
1 TB	2 TB	4 TB

Table 8. Maximum SV-NBU Volume Sizes.

This could additionally be important to consider if the volumes would ever be moved to a different platform with a smaller maximum volume size.

NearStore storage units cannot be used as part of a NetBackup storage unit group.

SV-NBU best practices

The use of gigabit Ethernet between the Media Server and the IBM N series Storage system is strongly recommended in order to maximize backup performance. Use of 100base-T will very likely result in unacceptable performance. NetBackup policy considerations include the following:

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

For UNIX client backups the option DO_NOT_RESET_ATIME should be set on the NetBackup client. This will provide 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.
- After createinga storage unit with "Enable Block Sharing" and doing a backup, don't change it.
- An IBM N series storage system volume should not be configured as a SV-NBU storage unit on more than one 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.

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

```
snap sched 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 behind firewall.

Application-specific implementation

(As more sites are put into production, this section will be expanded with application-specific recommendations/limitations/etc.)

NSM

NSM supports NAS file system backup of Solaris NFS and Windows CIFS, configured as policy types *Standard* and *Windows NT* within NetBackup.

Even if an environment is predominantly a UNIX OS other than Solaris, the benefits of NSM can be seen.

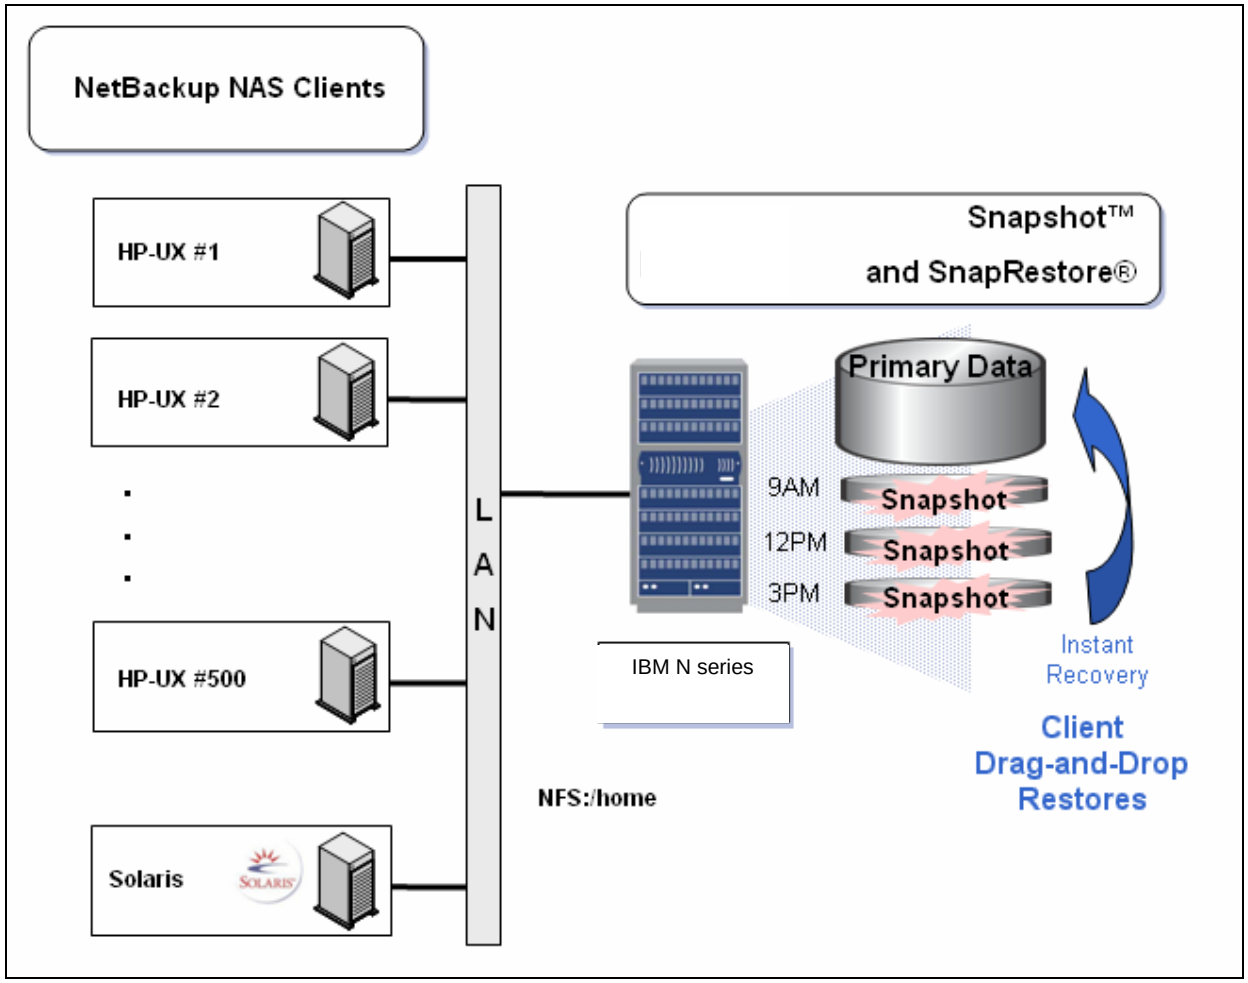

Figure 84. NSM in Predominantly non-Solaris Environments.

Oracle

The only integrated database application support is for Oracle8*i*™ or later running on Solaris. This is configured using the NetBackup Oracle Agent and an *Oracle* policy type.

RMAN needs to use the `proxy-copy` method.

With NSM integration, `.dbf` files (the vast majority of virtually all Oracle databases) are backed up with snapshot copies, but non-DBs (i.e., `.log` and `.ctl` files) are backed up to another storage unit.

For a (more) complete description of the functionality, see TR-3394, “Integrating Snapshot and SnapRestore with NetBackup in an Oracle Backup Environment.”

NSVM

NSVM has identical file systems and application support as NSM.

Even if an environment is predominantly a UNIX OS other than Solaris, the benefits of NSVM can be realized.

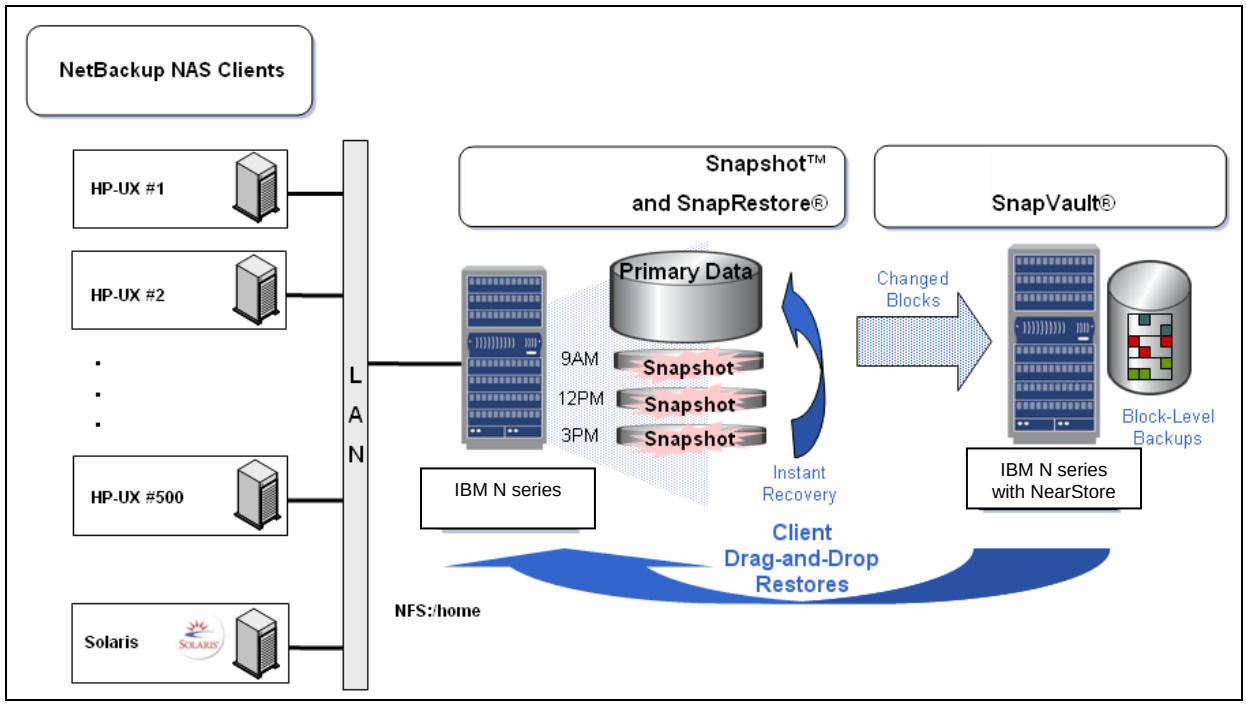

Figure 85. NSVM in Predominantly non-Solaris Environments.

SV-NBU

Unlike NSM and NSVM, SV-NBU is not limited to Solaris and Windows NAS clients only. Any NetBackup client data can be backed up regardless of where the data is stored.

File system

For heterogeneous file system backup, whatever clients NetBackup supports are also supported by SV-NBU. The exception is for file systems such as Novell and for Solaris ACLs, neither of which can be stored completely natively in WAFL. The security meta data is maintained (as is always the case) by NetBackup but not included in the unpacked WAFL structure as they can’t be stored natively.

Common problems and troubleshooting

This section covers those issues which have been known to come up when installing, configuring and deploying the solutions discussed in this document. This will most assuredly be a “living” section as some problems are fixed and different ones arise.

NSM

One of the most common problems is when NDMP authentication is configured incorrectly. When troubleshooting, begin by ensuring NDMP is installed, licensed and configured properly between the Master Server (and any media servers involved) and the IBM N series storage system.

tpautoconf –verify –nh <hostname>

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

Make sure NetBackup is at 5.1 MP2 or later and is the Enterprise Server version.

Make sure the volume that you have created is mounted (NFS) or mapped (CIFS) to the NetBackup client (the snapshot copy will be created in the root of this volume).

For NSM (and NSVM) in a Windows environment, the following items may cause problems:

- Be sure the pathname of the files/dirs in the backup list is correct (\\filer\vol\sharename).... else you’re likely to see error 156.
- When restoring if you run into permission problems make sure that you are running the NetBackup client service under an administrator account (not local system).

If the `~snapshot` directory is not visible when navigating to do restores, ensure the following:

- If using Windows, make sure you have folder options set-up to display hidden files.
- Although snapshots should be turned off on the primary IBM N series storage system (since NetBackup will do all the creation and management of snapshots), “Snapshot Directory Visible” for the volume should be selected.
- On the primary IBM N series storage system, be sure that `cifs.show_snapshot` is on:

```
fas3050-rtp01> options cifs.show_snapshot
cifs.show_snapshot          off
fas3050-rtp01> options cifs.show_snapshot on
```

NSVM

The same problems discussed for NSM apply for NSVM (except that you must be running NetBackup 6.0).

Make sure SnapVault is installed, licensed and configured properly on both the primary and the secondary.

Unlike NSM, NSVM requires that a Qtree (not a volume!) be mounted (NFS) or mapped (CIFS) on the NetBackup client.

Make sure that the Qtree is exported/shared – do NOT simply export/share and mount/map the volume and have the qtree below that.

If a problem is experienced, reproduce the problem while "`ndmpd debug 70`" and examine `/etc/log/ndmpdlog.*`

Due to the limit on the number of snapshot copies allowed per volume (255), killer snapshot schedules are easy to create; you'll see "`snapshot create failed`" on console and `/etc/log/snapmirror.`

Use "`snapvault destinations -s`" to see which qtrees hold oldest snapshot copies; delete older backups or move to tape using NetBackup GUI.

If the `~snapshot` directory is not visible when navigating to do restores, ensure the following:

If using Windows, make sure you have folder options set-up to display hidden files.

Although snapshots should be turned off on the secondary IBM N series storage system (since NetBackup will do all the creation and management of snapshots), "Snapshot Directory Visible" for the volume should be selected.

On the secondary IBM N series storage system, be sure that `cifs.show_snapshot` is on:

```
r200-rtp01> options cifs.show_snapshot
cifs.show_snapshot          off
r200-rtp01> options cifs.show_snapshot on
```

Snapshot sihedule Issue

Due to Data ONTAP limiting total snapshots to 255 per volume, consider disabling any unneeded scheduled Data ONTAP snapshots. Use one of the following commands:

```
snap sched MYVOL 0 0 0
vol options MYVOL nosnap true
```

In native SV, "`snapvault snap sched`" shows how many snapshot copies will be consumed, but for NetBackup NSVM backups, snapshot consumption must be calculated.

Good example:

```
/vol/280er/users (perf280er /users): 5:00, 13:00, 21:00
/vol/c38/home (fsr-c38 /home): 5:00, 12:00, 17:00, 23:00
/vol/c38/extra (fsr-c38 /extra): 13:00, 20:00
Needs 4 snapshot copies/day; 40-day retention OK.
```

Bad example (if one secondary volume is doing all backups):

```
/vol/280er/users (perf280er /users): 1:00, 3:00, 5:00
/vol/c38/home (fsr-c38 /home): 7:00, 9:00, 11:00, 13:00
/vol/c38/extra (fsr-c38 /extra): 15:00, 17:00, 19:00, 21:00
Needs 8 snapshot copies/day: 40-day retention IMPOSSIBLE.
```

Previous SnapVault relationships

If a SnapVault relationship exists for a primary qtree to a qtree in the secondary volume and if this existing SnapVault relationship needs to be brought into NSVM, simply create a policy with the SnapVault storage unit as the existing secondary volume, specify the source qtree and mount it on the client. Then schedule a backup. NetBackup will recognize the existing relationship and start updating the corresponding secondary qtree for the primary qtree.

Note that this done regardless of the user's preference. Hence the users should note that if a storage unit specified to NetBackup contains existing SnapVault qtrees, NetBackup will start updating them, if the corresponding primary qtrees are specified as sources in the policy.

To bring existing SnapVault relationships under NSVM control, perform the following steps:

Bring all the qtrees (i.e. SnapVault relationships) in the secondary volume under NetBackup. Before specifying the existing secondary volume as Storage Unit to NetBackup, disable SnapVault schedule for the volume.

It is not recommended/supported that for a given secondary volume, some qtrees are under traditional SnapVault control and some qtrees are under NSVM control.

If you don't want to bring an existing SnapVault relationship into NSVM but want to re-use the previously defined volume for something new and back it up with NSVM, a few things must be considered:

If you "`vol destroy VOL1`" which had native SnapVault qtrees on it and then "`vol create VOL1`" it will not work.
You must "`snapvault stop -f`" each old SnapVault qtree in order to use VOL1 for SV-NBU. "snapvault status -c" will help.)
"`snapvault snap unsched -f VOL1`" and "`snap sched VOL1 0 0 0`" would also be recommended configuration changes.

SV-NBU

The same NDMP authentication discussed for NSM and NSVM should be verified for SV-NBU.

If the drop down for the NearStore volumes doesn't appear make sure you have issued the "`options snapvault.access all`" (or the correct names of the Master and Media Servers are specified) command.

Check the version of Data ONTAP; it must be version 7.1 or above.

There *is* a new logfile: /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*` log file(s).

Note that this parameter is set to zero by default. Higher numbers provide increased debug information (in /etc/ log/nbu_snapvault), but backups will be slower. Therefore, when done 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 there are no WAFL snapshot copies 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.
```

Verify backup

The SV-NBU solution has a performance optimization based on the assumption that mtime will always change if a file changes. If the assumption may not be correct for a particular customer environment, this performance enhancement can be by-passed by turning on *Verify Backup* via the following command:

options snapvault.nbu.verify_backup

By turning on *Verify Backup*, a bcompare is performed of every file that's identified as not being changed.

Key things to understand about using this command are:

1. Off by default.
2. If set to on, will disregard the performance optimizations which use file's *mtime*.
3. If set to on, backups will be *significantly* slower.
4. Should be set to on, if *mtime* of the file does not change when the file's contents change.

If the customer is unsure about the `mtime` assumption made above, the can use this on a one-time basis to verify de-dup is working properly.

If they know they have apps that don't change the mtime but the data can change, this should be set to on.

Kernel settings

With NetBackup in general there can be issues if kernel settings aren't at least at some minimal settings. In existing NetBackup production environments they should already be fine, but if setting up a NetBackup environment from scratch, the below are reasonable ones 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
-------------------------------------
```

Designing solutions

This section documents how solutions can fit into customer environments. It is envisioned that it will be expanded more each time a new version of this document is published.

NSM

The NSM solution can easily and quickly be integrated into existing IBM N series and NetBackup environments.

The maximum number of snapshot copies allowed per volume (255) and amount of space that the backups are going to consume are the only things you really need to understand before deploying it.

NSVM

Sites which have multiple locations / data centers with IBM N series at each of them are especially well-suited for this solution. One example is having multiple data centers / remote offices backing up to a centralized "Backup Data Center."

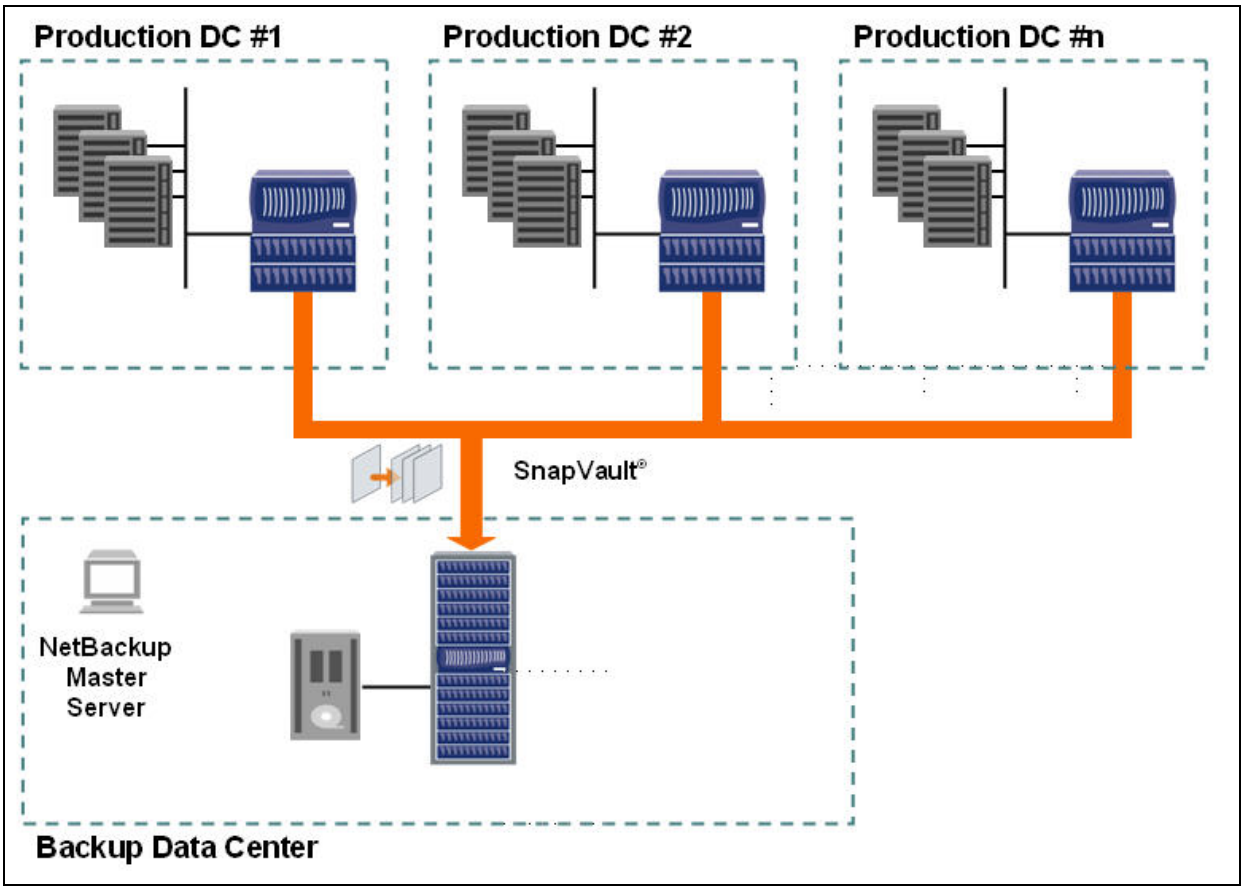

Figure 86. NSVM Solution for Remote Data Centers.

Another example is multiple production data centers, where SnapVault backups from each one are sent to the other. The example below is from a customer who has one data center north of and another data center south of a major city, with network connectivity between the two.

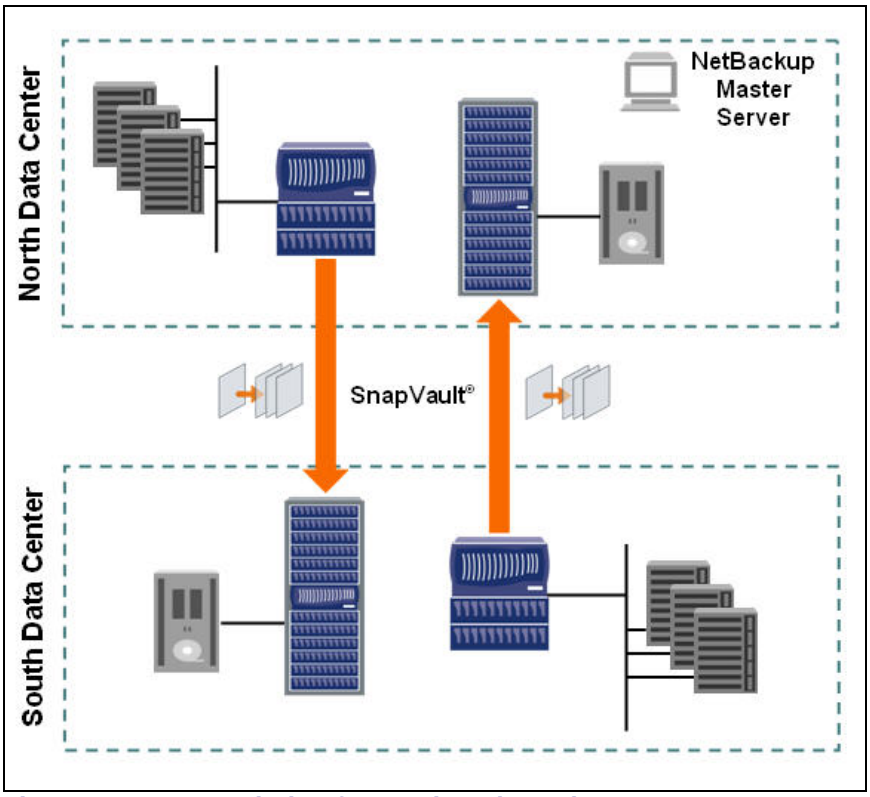

Figure 87. NSVM Solution for North and South Data Centers.

At each location are several tens of terabytes of IBM N series storage. The backups to local-attached tape drives were not completing in the weekend backup window and, like virtually all companies, data was growing fast.

Using the NSVM solution they can get the weekend “full” backups done very quickly because SnapVault is only sending the changed blocks. This meets their SLA to have backup completed in a specified period of time so production isn’t implemented, and for those backups to reside on other media (and, even better, in a different location).

At each destination site they can then do backups to tape and the “backup window” to accomplish this can be longer.

SV-NBU

Of the three solutions, SV-NBU has been the most talked about, probably because it’s brand new technology and solves problems differently than anybody else.

In general it fits where super-high performance isn’t the driving requirement for disk-based backups, but where storage savings are. The implication of which applications the environment has will be important to consider too as not all can take advantage of the de-duplication benefits.

File system environments especially well-suited are home directories, web server data, software development library directories, and Windows and UNIX operating system partitions.

Appendix 1 – Tape and disaster recovery scenarios

Although there are substantial benefits to be achieved with the joint disk-based backup solutions offered by Symantec and IBM N series, 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 disaster recovery purposes (and from there perhaps put it on tape). These two situations can be seen by looking again a picture presented earlier in this document.

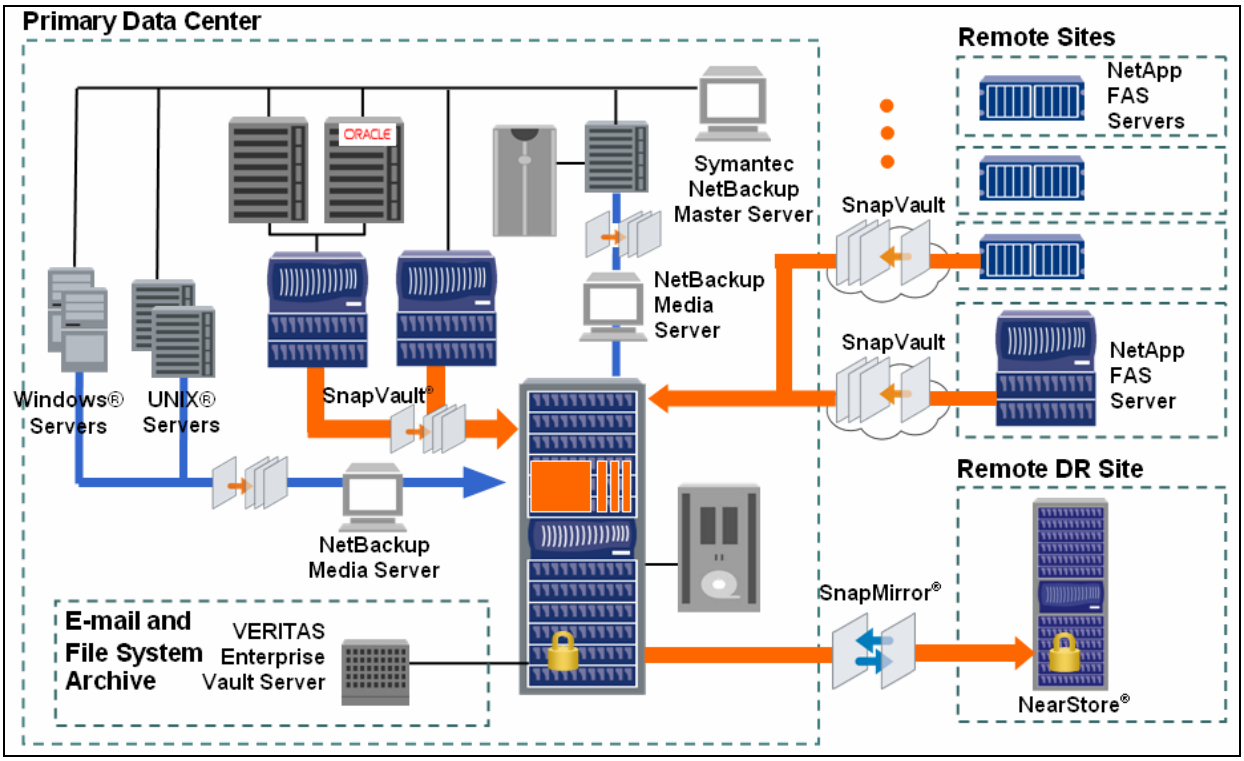

Figure 88. Joint Solutions and Disaster Recovery.

This section and those below will be enhanced to discuss the whole life-cycle management, addressing such things as:

- How tape integrates with each solution
- Automation of Snapshot copies to tape

Tape

Once backup images reside on the NearStore system, the customer can easily use other policies to put the images on tape for offsite storage. The sections below talk about specifics for each solution.

NSM

For best performance, use NDMP to backup the NearStore NSM backups to tape. Note that the NSM storage unit cannot be used with Inline Tape Copy (ITC), staging or vault.

The first step would be to attach tapes drives to the IBM N series storage system and configure them as NDMP storage units within NetBackup.

Then create an NDMP policy with NetBackup.

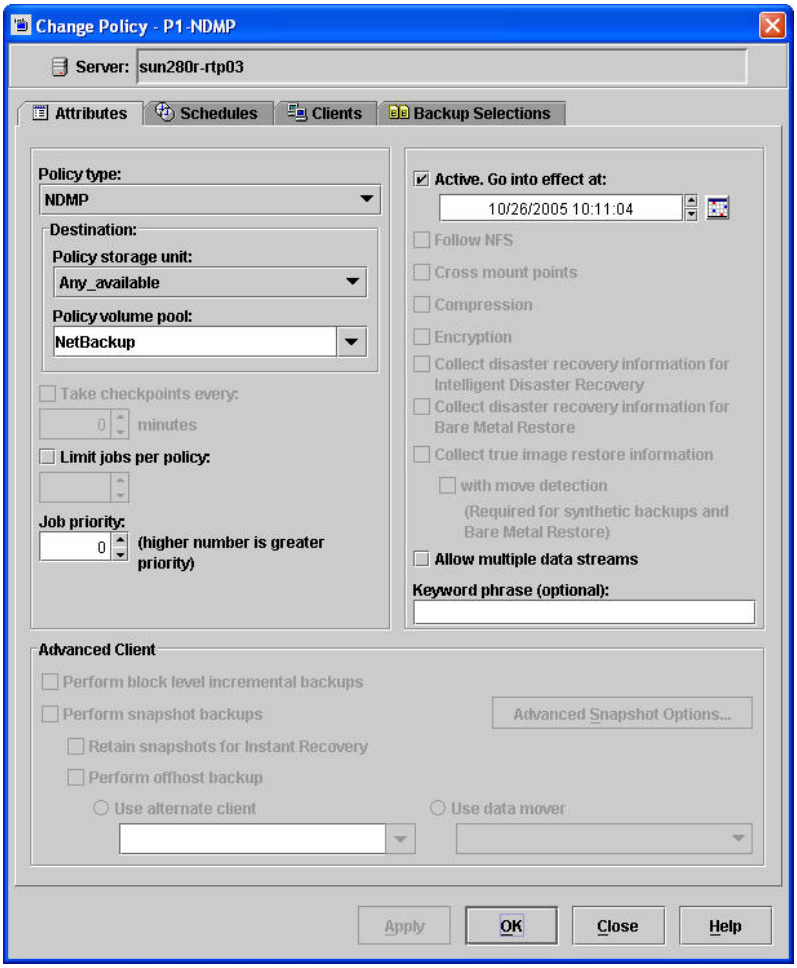

Figure 89. NetBackup NDMP Policy Attributes.

The schedules are defined like any other NetBackup schedule.

The client is specified as the NDMP host, which is the name of the IBM N series storage system.

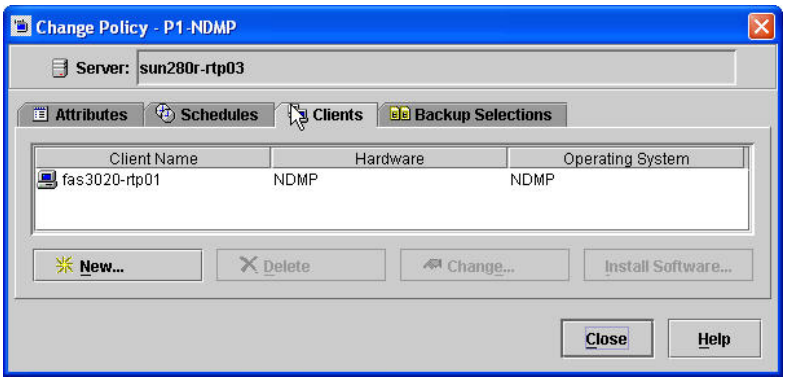

Figure 90. NDMP Backup of NSM – Clients,

The backup selections are specified as the volume or Qtree on the NDMP host.

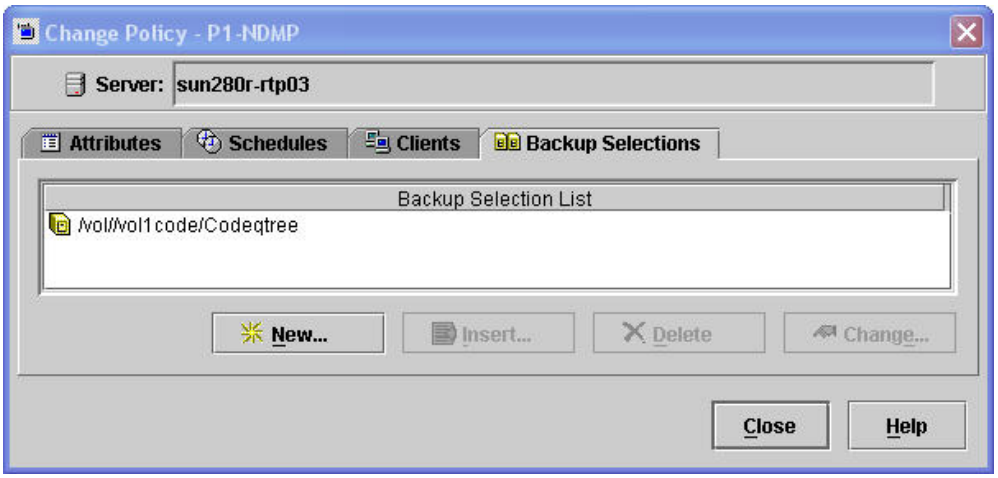

Figure 91. NDMP Backup of NSM – Backup Selections.

The resulting NDMP backup will not be automatically correlated to the original (Solaris or Windows) client backup. Although this means that the restore of the client requires a two-step restore, for point-in-time rollback the second step would take almost no time (individual file and directory restores would take longer).

NSVM

Like NSM, for best performance, use NDMP to backup the NearStore NSVM backups to tape. (Note that the NSVM storage unit cannot be used with Inline Tape Copy (ITC), staging or vault.)

The first step would be to attach tapes drives to the SnapVault secondary IBM N series storage system and configure them as NDMP storage units within NetBackup.

After creating the NDMP policy via NetBackup, the client is specified as the NDMP host, which is the name of the IBM N series SnapVault secondary.

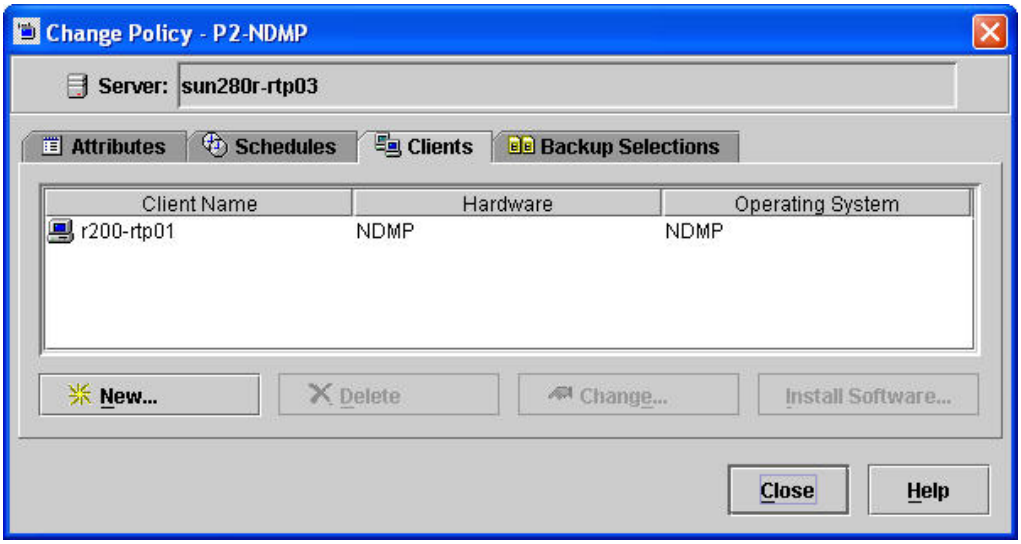

Figure 92. NDMP backup of NSVM Volume – Clients.

The backup selections are specified as the volume or Qtree on the NDMP host.

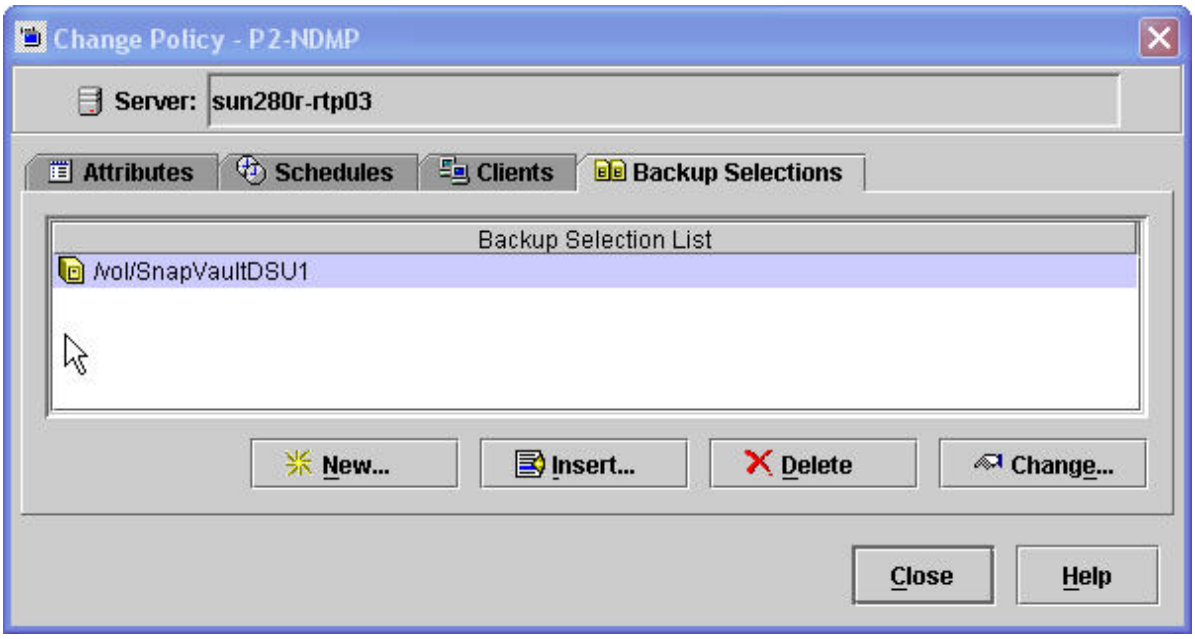

Figure 93. NDMP backup of NSVM Volume – Backup Selections.

The resulting NDMP backup will not be automatically correlated to the original (Solaris or Windows) client backup. This means that the restore of the client requires a two-step restore: (1) the NDMP restore to the SnapVault secondary, and (2) the SnapVault Restore from secondary to primary.

Important note: backup of the NSVM DSU volume to tape provides an additional backup of the qtree data, but does not provide disaster recovery protection for the overall NSVM DSU volume. In other words you can't recover the NSVM volume from tape and then resume NSVM backups to that volume.

This may be okay, however, as in an actual disaster scenario, the plan might be to make the secondary the production storage system, in which case the restore would only require the NDMP restore step (along with the appropriate actions to make the recovered volume/Qtree the active file system). Once the customer was back up and running in production, new backup policies would be configured to address protecting their new production data.

SV-NBU

For SV-NBU, NDMP is not supported. To move backup images to tape there are two choices:

1. One is to use Inline Tape Copy (ITC) for the policy schedule(s). This feature of NetBackup sends a copy of the backup stream to the tape and the NearStore disk storage unit simultaneously.

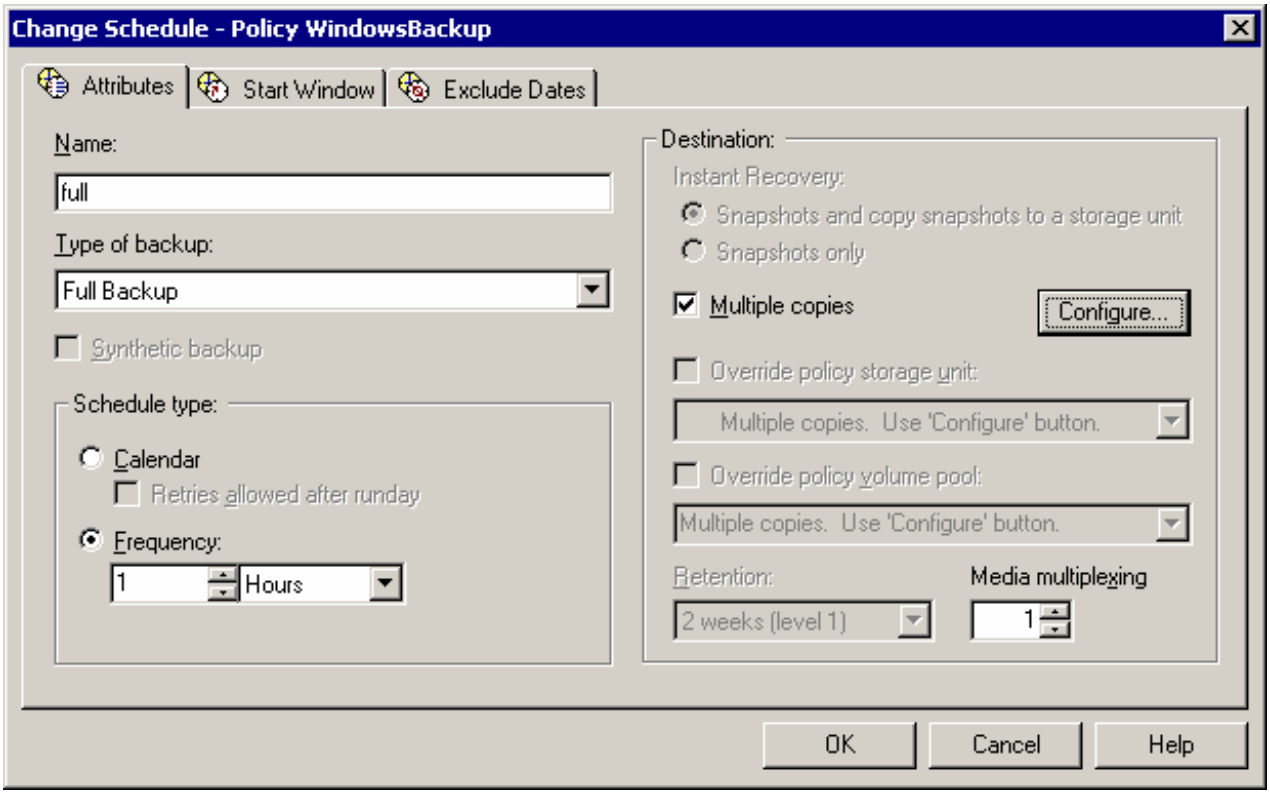

Figure 94. Enabling Inline Tape Copy (ITC).

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

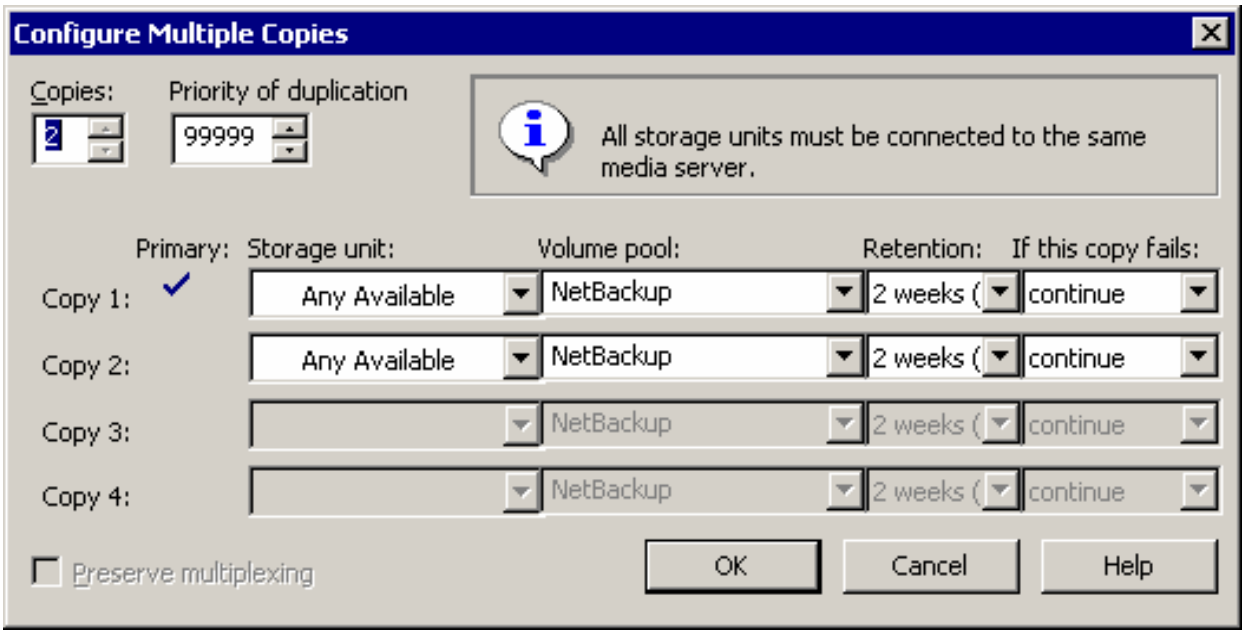

Figure 95. Specifying ITC Destinations.

2. The other is to enable the NearStore Disk Store unit as a temporary staging area, previously called a Disk Staging Storage Unit.

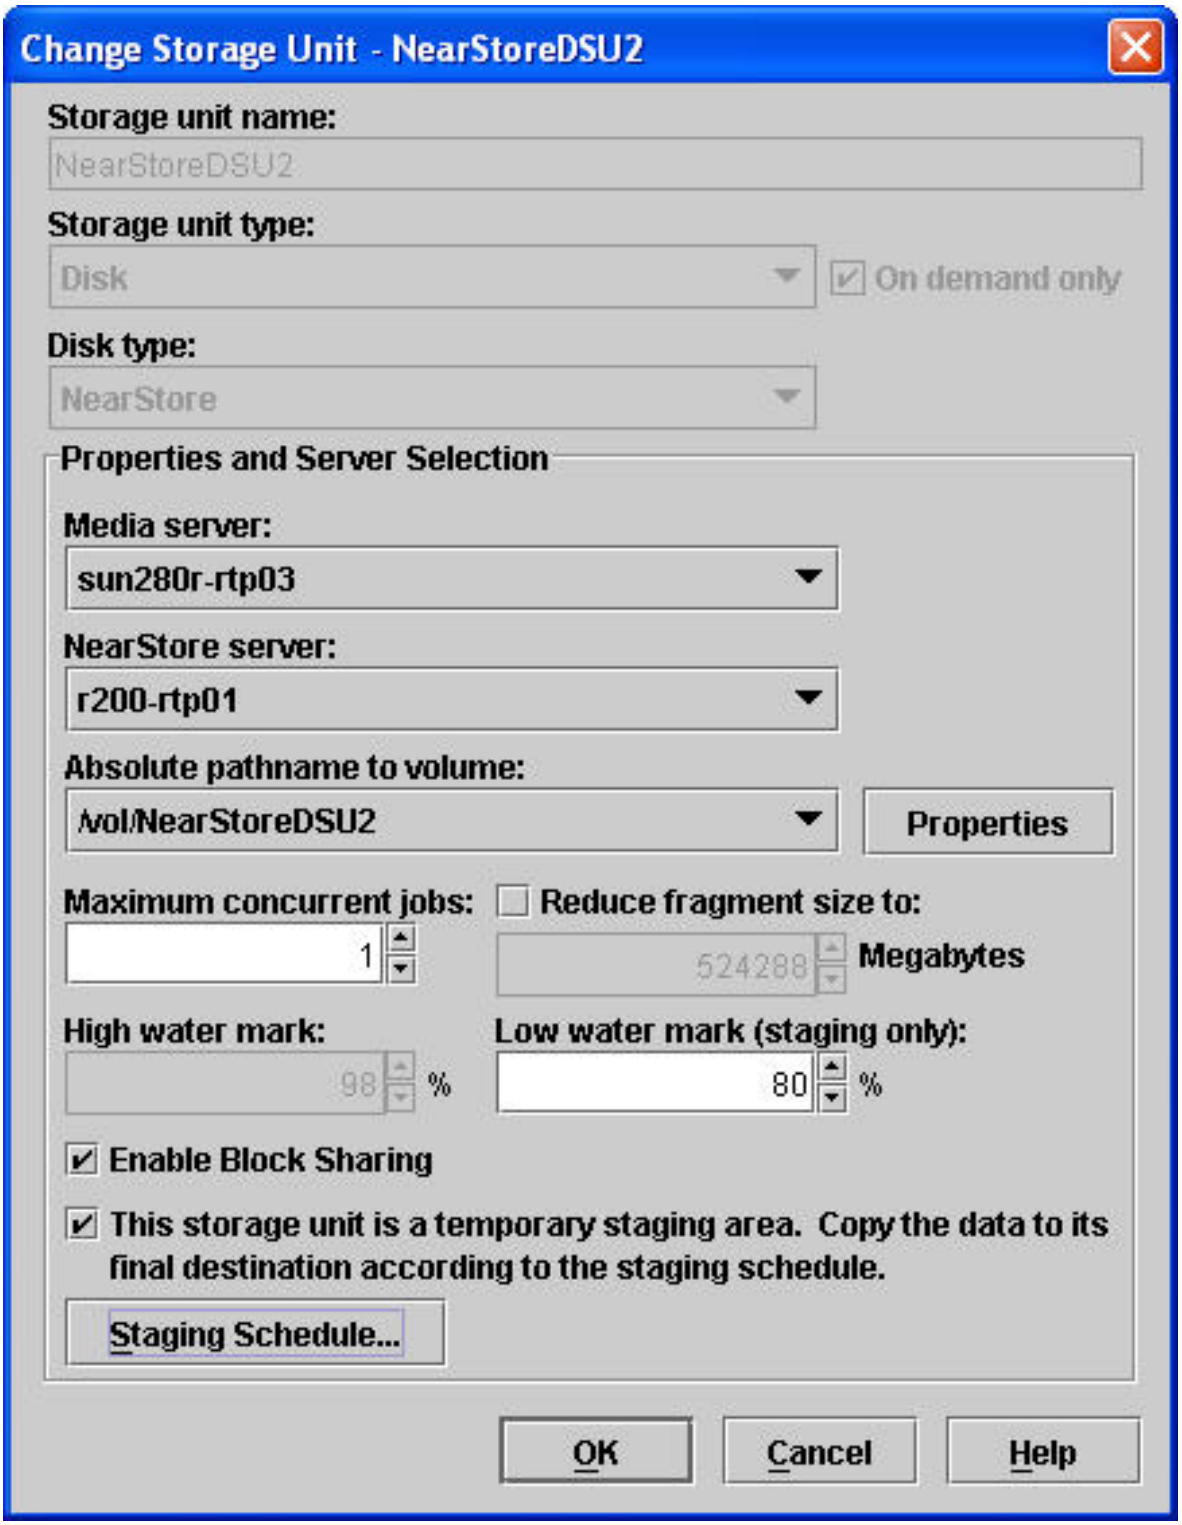

Figure 96. SV-NBU NearStore DSU – Enabling Staging.

After selecting the check-box for “temporary staging area,” you can configure a specific staging schedule for this storage unit.

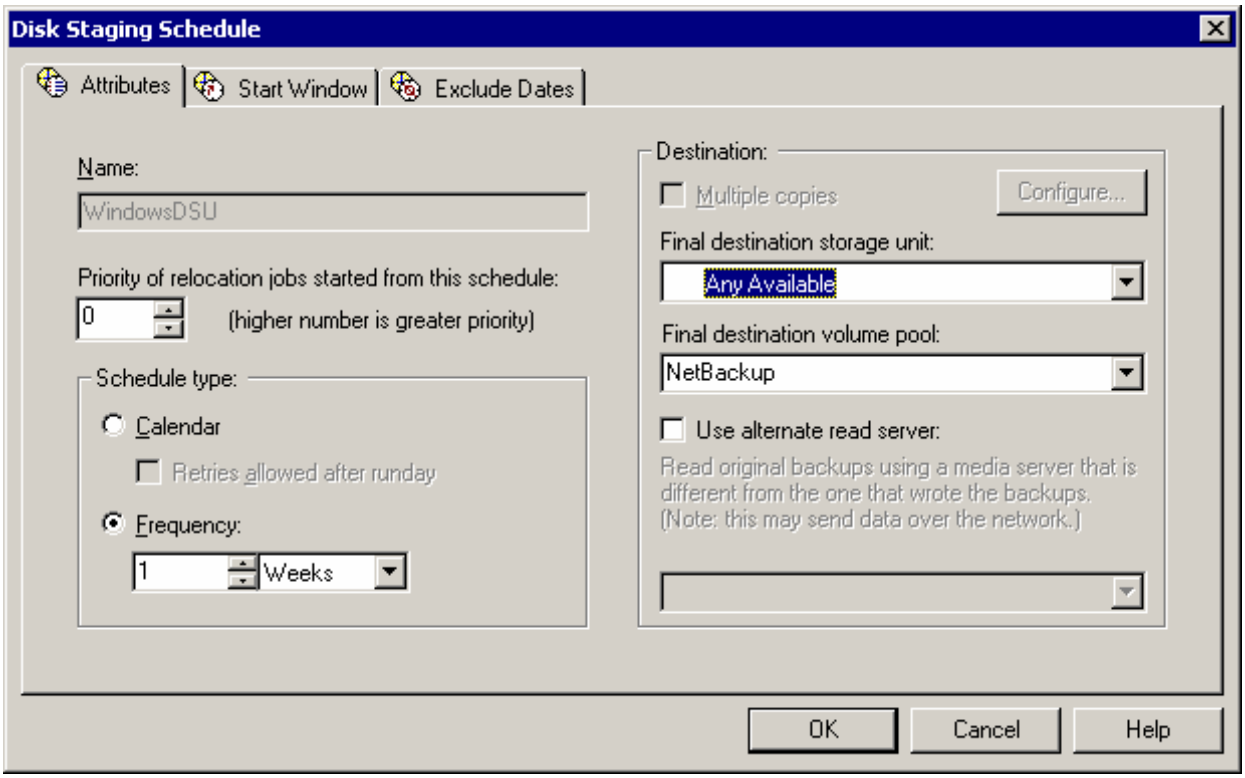

Figure 97. SV-NBU NearStore DSU – Configuring Staging Destination.

This means that NetBackup will automatically (based on schedules configured by the storage administrator) copy the data from the IBM N series with NearStore feature storage unit to a final storage unit (usually a tape drive).

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

In some customer environments a third method, the NetBackup Vault product, may be utilized to move SV-NBU images to tape.

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 storage system, IBM System Storage N series with SnapMirror® technology can be used to replicate the backup images to an online DR site, although NetBackup 6.0 and Data ONTAP don't provide an integrated automated solution currently.

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 appropriate sections in NetBackup documentation for bringing this NetBackup DR environment online.

For a complete discussion on SnapMirror, refer to theSnapMirror deployment/implementation guide.

NSM and SnapMirror

The NSM backup, essentially a snapshot on the primary IBM N series storage system, can be replicated to a remote site using SnapMirror.

NSVM and SnapMirror

SnapMirror with NSVM is not supported.

VSM would be able to provide an additional backup of the qtree data, but does not provide disaster recovery protection for the overall NSVM DSU volume. In other words you can't recover the NSVM volume from the SnapMirror site and then resume NSVM backups to that volume.

SV-NBU and SnapMirror

SnapMirror with SV-NBU is not supported.

When the gtar data stream from NetBackup is converted to 4K 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 storage system. Because of this, replication of the SV-NBU volume is currently not supported in any form.

Appendix 2 – Acronyms

ASIS	Advanced Single Instance Storage
CIFS	Common Internet File System
CLI	Command Line Interface
DB	Database
DR	Disaster Recovery
DSSU	Disk Staging Storage Unit
DSU	Disk Storage Unit
GB	Gordon Biersch
GUI	Graphical User Interface
ICS	Infrastructure Core Services
ITP	Inline Tape Copy
NAS	Network Attached Storage
NDMP	Network Data Management Protocol
NFS	Network File System
NSM	NetBackup Snapshot Management
NSVM	NetBackup SnapVault Management
PSE	Professional Services Engineer
QSM	Qtree SnapMirror
RPO	Recovery Point Objective
SE	Systems Engineer
SIS	Single Instance Storage
SLA	Service Level Agreement
SM	SnapMirror
SV	SnapVault
SV-NBU	SnapVault for NetBackup
TCE	Total Customer Experience
TR	Technical Report
VSM	Volume SnapMirror
VSS	Volume Shadow copy Service
VTL	Virtual Tape Library
WAFL	Write Anywhere File Layout
WORM	Write Once Read Many

Trademarks and special notices

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, FilerView, FlexVol, NearStore, SnapMirror, SnapVault, SnapRestore, Snapshot 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.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

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.