WHITE PAPER: DATA PROTECTION

Confidence in a connected world.

Best Practice for NDMP Backup Veritas NetBackup™

Paul Cummings

January 2009

Best Practice for NDMP Backup Veritas NetBackup™

Contents

1.0 Introduction and overview

Many customers are confronting the confusing number of options and configurations when it comes to backing up and restoring their network attached storage (NAS) servers. Most NAS vendors and data protection providers (such as Symantec) have multiple options for protecting your NAS environment. This paper will help you make sense of the numerous options available to you for NAS backup and restore. You may be asking yourself questions such as:

- What is the best way to protect my particular NAS environment?
- What is NDMP and how does it help protect my NAS environment
- What are the options for improving NDMP backup speeds?
- I currently back up my NAS over CIFS or NFS. Is this safe?
- What is needed to recover individual files and directories at a granular level with my NAS backups?
- What about backups to tape, VTL or disk?
- Can NetBackup™ client and NAS backups share the same devices?

This paper addresses these questions by describing the best practices for protecting and recovering NAS devices utilizing the NDMP (Network Data Management Protocol) from vendors such as NetApp, EMC, and IBM as well as discussing the challenges associated with protecting NAS appliances and multiple strategies that can be used to address them.

This paper does not discuss other NAS protection options from NetBackup such as Snapshot Client and SnapVault™, to name a few. These are covered in the NetBackup Snapshot Client Administration Guide.

1.1 NetBackup and NDMP

NetBackup provides a comprehensive data protection solution that supports a wide range of platforms and applications found in today's data centers. It includes centralized administration and reporting, media management, automated policy-based backups, and restores.

The NetBackup NDMP option extends the capabilities of NetBackup to include native backup and restore of NAS appliances. Supported versions of software for these vendors' NAS appliances are listed on the NetBackup NDMP Hardware Compatibility List (HCL). This allows you to create backups of data on an NAS without interrupting client access to the data. NetBackup incorporates the protection of NDMP-enabled NAS into a single solution by enabling tape/VTL library sharing, drive sharing, direct access recovery, and auto configuration.

The following is an overview of the feature set provided by the NetBackup for NDMP option:

- Full, differential incremental, cumulative incremental, and snapshot backups (Enterprise Client is required for snapshot backups of application data)
- Alternate NDMP client and path restore
- NDMP direct copy
- Auto configuration
- Direct Access Recovery (DAR), which provides individual file level restore
- Advanced database integration
- Broad platform and protocol support
- Dynamic tape drive sharing (Shared Storage Option)

1.2 Glossary

The following terms are used throughout this document:

- Network Data Management Protocol or NDMP—An industry standard protocol created to ease the process of integrating backup and restore for network attached storage (NAS) with data protection software such as NetBackup. NDMP is the data transfer protocol and one piece of the overall data protection solution. For more information about the NDMP protocol and the NDMP standards group visit http://www.ndmp.org/info/faq.shtml.

- Network Attached Storage or NAS—A self-contained computer commonly referred to as an appliance that is connected to a network or SAN, with the sole purpose of supplying file-based data storage services to other devices on the network.

- Storage Area Network or SAN—A network created for the main purpose of sharing disk and/or tape storage to other devices. SAN is most commonly a Fibre Channel (FC) topology with a switch to manage the connections between devices. One benefit of SAN is the ability to make disk and/or tape look as locally attached devices to the server operating systems.

- Direct Access Restores or DAR—A feature of NDMP that Provides faster recovery of directories, an individual file, or selection of files. When files are backed up, the location is recorded. At restore time, this information can be used to position the exact location of the file on the media rather than reading sequentially through the whole backup. This feature is inherent to all NDMP backup methods, but may not be supported by all vendors.

- Storage Unit—A logical target to which NetBackup writes backup data. Storage Units may map to either disk or tape storage. The precise nature of the mapping depends on the type of storage unit.

- Disk Storage Unit or DSU—Disk storage utilized by NetBackup to store backup data. NetBackup supports several different types of DSU, all of which can be used as targets for NDMP backup.

1.3 Additional resources

The following documents provide more background on the subjects discussed in this paper:

- A NDMP Hardware Compatibility List (HCL) indicating which products work with NetBackup is available here: http://support.veritas.com/docs/251713

- The Veritas NetBackup Backup Planning and Performance Tuning Guide is available at: http://seer.entsupport.symantec.com/docs/307083.htm

- The NetBackup Server Hardware Compatibility list is available at: http://support.veritas.com/docs/284599

- NetBackup administrator's guides

- The NetBackup NDMP Administrator's Guide for UNIX, Linux and Windows is available at: http://seer.entsupport.symantec.com/docs/290205.htm

- The NetBackup Snapshot Client Administrator's Guide is available at http://seer.entsupport.symantec.com/docs/290224.htm

2.0 Architecture options and selections

There are several options for architecting NAS NDMP protection. The most common and effective are:

- Local NDMP—One of the more common options for architecting NAS NDMP protection is local NDMP backup. This is accomplished by attaching a single tape drive or tape library directly to the NAS and sending backups directly across SCSI or SAN. This tape device can be a standalone drive, library, VTL, or any number NDMP dedicated drives in a library.

- Direct NDMP—Direct NDMP backup is identical in practice to local NDMP backup, but differs in the implementation with sharing of SAN tape drives in a library with a NetBackup media server. With the Shared Storage Option (SSO) NetBackup can share tape resources between the media servers and NAS. This requires the NDMP host to be SAN-attached and zoned to see the library or drives. The master server controls access to the tape device.

- Remote NDMP—Remote NDMP backup incorporates the same tape device support as direct backup, but sends the data stream over the network and through an NetBackup media server. This can provide a few advantages; one such notable feature is support for writing backup data to disk with the introduction of NetBackup 6.5.2. This can also result in some disadvantages, such as slower network transfer speeds.

- 3-way NDMP—In a three-way backup or restore, data is sent from an NDMP host over a LAN to a storage device that is attached to another NDMP host. This backup contrasts with local NDMP backup or restore where the data is sent directly to a storage device attached to the NDMP host.

An additional option for protecting NAS devices is to utilize file sharing protocols such as CIFS or NFS and “walking” the file system to back it up. While this may be an effective architecture for smaller NAS environments, it typically is not appropriate for most enterprise-class NAS devices and therefore is not discussed in this document.

	NDMP Backup Architecture		
	Remote	Local	Direct
Recommended for	NAS environments looking at LAN backup to NetBackup media server tape or disk	NAS environments with SAN, SCSI or FC attached tape or VTL devices	NAS environments with SAN and shared tape/VTL devices (Shared Storage Option)
Basic NDMP backup and recovery	✓	✓	✓
Centralized management	✓	✓	✓
NDMP backup to NetBackup Disk Storage Units	✓	✗	✗
DAR file level granular recovery	✓	✓	✓
"LAN Free" direct to tape/VTL backup	✗	✗	✗
Direct "on host" copy of tape and VTL media with NDMP Direct Copy	✗	✓	✓
Use SSO to share tape resources with NetBackup media servers and NDMP devices	✓	✗	✓

Table 1: Solution Comparison

You would typically use remote if: The backups are smaller in size, no SAN has been implemented, or you do not require backup to NetBackup disk solutions.

You would typically use local or direct if: Backup directly to tape or VTL is a top priority.

For detailed NetBackup and NDMP tuning recommendations, refer to the Veritas NetBackup Backup Planning and Performance Tuning Guide. This guide discusses tuning options such as shared memory (number and size of data buffers) that if configured correctly can increase remote NDMP backup transfer speeds.

2.1 Remote NDMP backup

With remote NDMP, architecture data is sent from the NAS device via the LAN through a NetBackup media server, which then writes the data to either disk or tape.

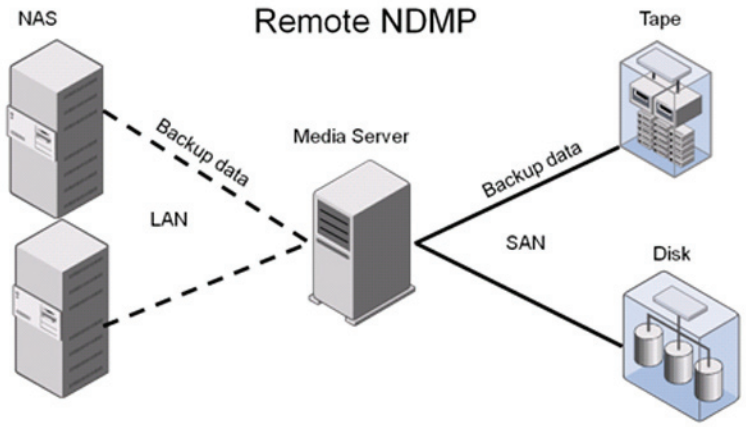

Figure 1: Remote NDMP Architecture

Because the backup traffic traverses the LAN, it is important to ensure there is adequate bandwidth to support the backup operation in the time window provided without unduly disturbing other network traffic. In case of heavy backup loads on the LAN, it may be ideal to create a separate LAN or vLAN dedicated exclusively to carrying the backup traffic to avoid saturating the production LAN.

Remote NDMP provides significant flexibility and functionality over other methods of NAS protection because the data flows through a NetBackup media server, including:

- Backup to Disk—Backup data can be written directly to disk using any NetBackup DSU type, including Basic Disk, AdvancedDisk, OpenStorage, and the PureDisk™ deduplication option (the PureDisk deduplication option will provide little to no deduplication with NDMP backups) with NetBackup 6.5.2.
- Encryption—Backup data can be encrypted using either the NetBackup media server encryption option or tape encryption in combination with NetBackup Key Management Services.
- Storage Lifecycle Policies—With Storage Lifecycle Policies (SLPs) data can be automatically duplicated to different storage media with different retention policies; for example, backup to fast disk and retain 24 hours, duplicate to inexpensive disk and retain for 30 days, duplicate to a remote site over the WAN and retain for 30 days, and duplicate to tape and retain for three years.

Advantages:

- Simple to implement
- Support for remote NDMP backup to disk storage units (requires 6.5.2 or greater)
- Take advantage of media server load balancing for increased performance, efficiency, and high availability
- Media server encryption option (MSEO) and tape (LTO4) encryption supported
- Automate data lifecycles with SLP to migrate NDMP data from disk to disk to tape as desired

Disadvantages:

- Potentially slower backups due to LAN vs. SAN speeds
- Typical 1 GigE LAN will limit backups to a real-world 60-80 MB/sec
- Additional LAN overhead if not using a dedicated backup LAN
- No support for NDMP direct copy
- Remote NDMP backup to disk supports PureDisk disk option with limited deduplication results

Recommendations:

- When building large NAS configurations, consider keeping the volumes in a manageable size to assist with backup performance. When starting a NDMP backup, the NAS device must walk the inodes for the entire volume before transferring data. By keeping the number of files and the size of the volumes in check you can help with backup times.
- Depending on the size of your NDMP backups, you might dedicate a network for only NDMP backups, commonly referred to as a "backup LAN." With NDMP backups larger than 1 TB, 1 GigE links should be the minimum, with many customers looking at 10 GigE for even larger NDMP backups in the 10 TB range.
- Look at backup to disk for the added flexibility and performance. When waiting for NDMP backups to supply data, tape drives can sit waiting. Backups writing to disk would not have the same limitations and would keep your tape drives free for other backups.
- With NetApp filers, backing up sets of qtrees instead of volumes in a backup is NOT recommended as NetApp walks the nodes of the entire volume for each backup.

2.2 Local and direct NDMP backup

A local or direct NDMP backup architecture requires that a tape drive be either directly attached to a NAS device (for example, local NDMP backup) or connected via a SAN (for example, local, direct or SAN (SSO) NDMP backup). While the devices are considered "directly connected," all scheduling and management is still handled by the NetBackup master server. Also known as "LAN free," these architectures take advantage of the typically faster speed of the SAN (2 Gbps, 4 Gbps, or 8 Gbps SANs that are very common today compared with 1 GigE LANs, though increasingly 10 GigE LANs are being deployed). With these architectures, data is sent directly from the NAS device through the SAN to the target tape drive (or VTL). Catalog data is sent from the NAS device over the LAN to the NetBackup master server.

If a SAN is not an option, you can always attach a single drive or more from your tape library directly to the NAS and run them as a local NDMP backup. NetBackup will see these drives as part of the library; however, they are only available to perform NDMP backups from the NAS device.

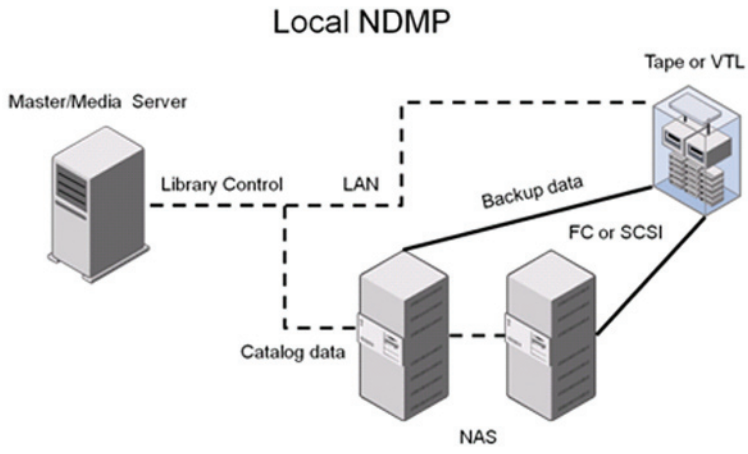

Figure 2: Direct NDMP Architecture

This architecture avoids sending data over the LAN and through a NetBackup media server. There is no need to create a separate backup LAN, and the additional load on the media server is avoided. When using this architecture with the SSO, tape drives do not need to be dedicated to a NAS device for NDMP-only backups, but can be shared among NAS devices and NetBackup media servers for true distributed NDMP protection.

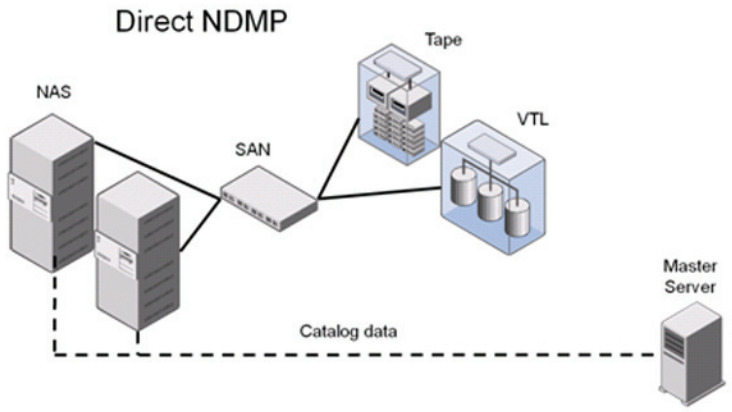

Figure 3: Shared NDMP Architecture

Advantages:

- Backups avoid the LAN (no additional LAN traffic) and media server (no additional load) and take advantage of direct SCSI or SAN attached tape or VTL for potentially faster backup and recovery.
- Typical 2 Gbps SAN backups can send data at 175 Mbs compared to 80 MB/Sec on a 1 GigE LAN. The actual speed at which data is sent depends on the speed at which the NAS reads from its file system and provides the data for backup. This figure will vary from one NAS device to another depending on the vendor and underlying technology used.

Disadvantages:

- Requires SAN for shared NDMP backup and recovery
- Local only (non-SAN) attached NDMP drives are only available for NDMP backups
- Not all NAS vendors support FC or SCSI HBAs and they lack the ability to support either local or direct NDMP backups
- Local or direct backup to FC/SCSI attached tape or VTL only; disk backup is not supported
- Encryption options limited to tape drive, vendor key management systems (KMS), or inline appliance only

Recommendations:

- When building the NetApp NAS configuration consider keeping the volumes in a manageable size to assist with backup performance. When NetApp NAS starts an NDMP backup it must walk the inodes for the entire volume. By keeping the number of files on the volumes in check you can help with backup times. In summary, if the options are creating one volume with 100 million files or 10 volumes with 10 million files, the latter would be ideal if applicable.

- If protecting volumes with hundreds of millions of files or more, look at backup to disk for the added flexibility and performance. When waiting for NDMP backups to supply data, tape drives can sit idle. Backups writing to disk would not have the same limitations and would keep the physical tape drive free for other backups; however, there is a downside. LAN-based backups can be significantly slower than SAN-based backups to fast tape or VTL. This option should only be considered when NDMP backups wait for long periods of time before transferring data. This is common with volumes housing hundreds of millions to billions of files.

3.0 Duplicating NDMP backup images

NDMP direct copy is another advantage to local or direct NDMP backup architecture. Typically, backup administrators want to make two backup copies of the data; one for on-site storage and another for off-site storage (perhaps for DR or legal reasons). Direct copy allows copying backup images from a supported VTL or NDMP attached physical tape drive to a another physical tape drive attached to either the VTL or NDMP NAS without using media server I/O or network bandwidth. Both source and destination drives used for NDMP direct copy must have NDMP device paths. Refer to the NetBackup Hardware Compatibility Lists for details of VTLs that support direct copy.

NetBackup can also copy NDMP images between a NAS-attached tape or VTL device and media server managed devices such as tape, VTL, or DSU. NetBackup can directly restore the NDMP image from either the original or duplicate back to the NDMP NAS.

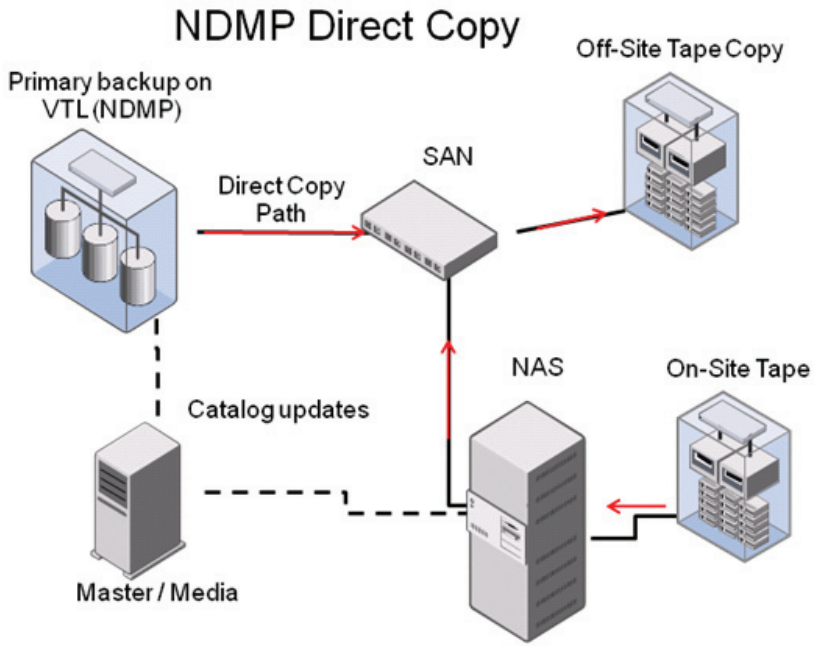

Figure 4: NDMP Direct Copy Architecture

4.0 Policy and device configuration

NDMP protection supports a number of NetBackup policies. This can range from the normal NDMP full or incremental backup to the more advanced snapshot, SnapVault or database integration. Some of the more popular policies and configurations are highlighted below. The following is a high level overview of configurations. For a more detailed explanation refer to the NetBackup, NDMP or Snapshot Client administrator's guides.

4.1 Full and incremental configurations

NDMP policies support the normal full, differential incremental and cumulative incremental backup types that you would select with NetBackup clients. When selecting one of these options you are actually manipulating the NDMP dump levels on the NAS. Full is the equivalent of dump level 0; differential incremental is dump level 1; and cumulative incremental is dump level 1–9. When performing backups, NetBackup will automatically increment the numbers from 0–9 as needed when selecting full, differential incremental, or cumulative incremental backups.

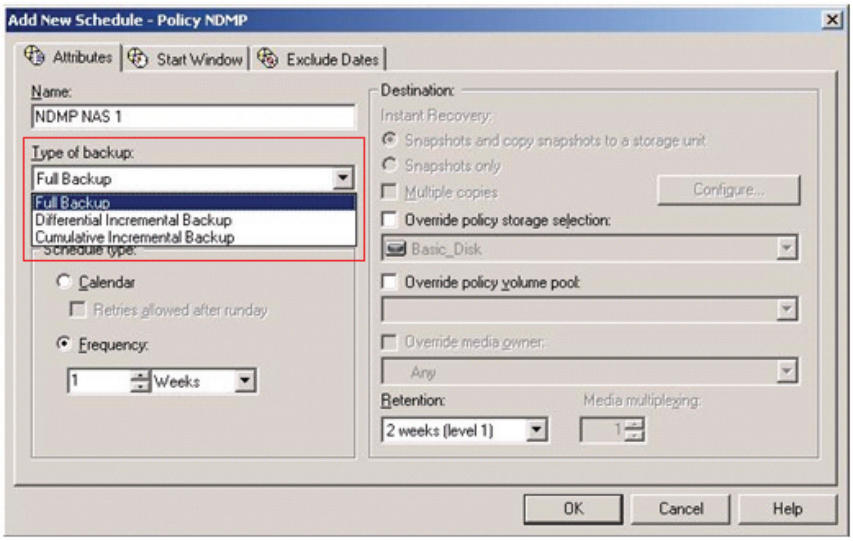

Figure 5: NetBackup NDMP Policy Attributes

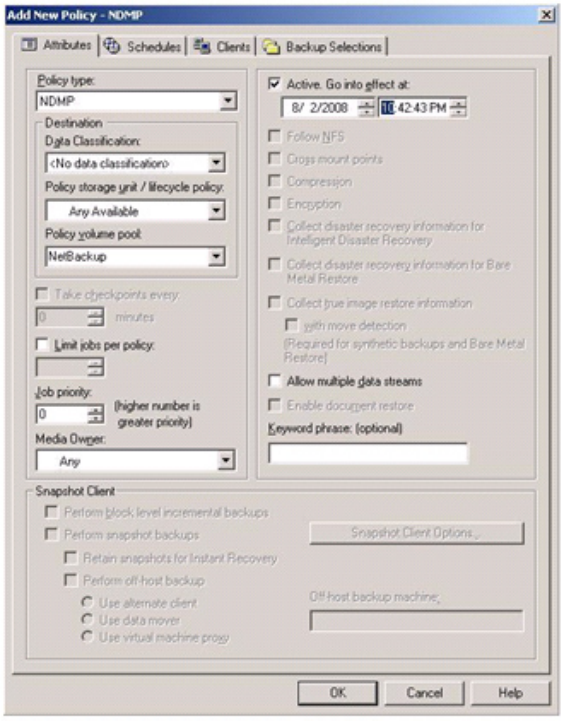

Figure 6: NetBackup NDMP Policy

The NetBackup policy for this configuration is straightforward. Select a policy type equal to “NDMP”, and then select a “Policy storage unit / lifecycle policy”. This method of creating a policy is similar to all methods of NDMP backup. The major difference is which storage unit is targeted. For local, direct, and 3-way NDMP backups, the storage unit must be an NDMP storage unit. For remote NDMP backups, a Media Manager or Disk storage unit may be used.

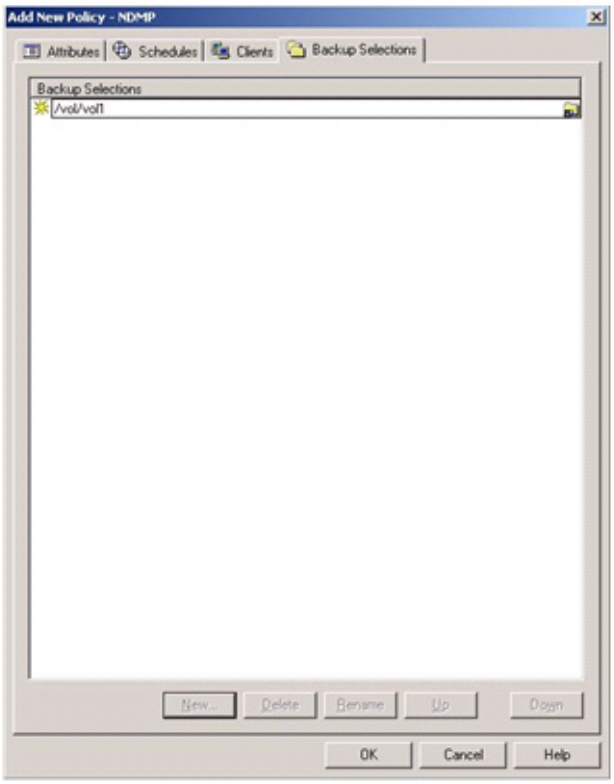

Figure 7: NetBackup Policy Backup Selection

With NDMP dump levels NetBackup schedules have the options of full, differential incremental, or cumulative incremental NDMP backups. Type of backup and retention periods for each backup type can also be selected on the Schedule tab.

The NetBackup policy "Backup Selections" can be populated by providing the volume and/or directory path. In this example, a path has been used to populate the backup selections list with "/vol/vol1".

4.1.1 NDMP backup to disk storage unit (DSU) configuration

With the release of NetBackup 6.5.2 support for remote NDMP backup to disk storage units such as Basic Disk, AdvancedDisk, SharedDisk, OpenStorage Disk and PureDisk has been added. For more information about creating a DSU, refer to the NetBackup administrator's guide.

NDMP backup to disk is as easy as selecting the DSU from the Policy Storage Unit in the Add New Policy screen. Another benefit of the NetBackup NDMP option is support for Storage Lifecycle Policies (SLP). With SLP, the NDMP backup can be written to disk or tape, and then automatically duplicated to other NetBackup managed storage devices, such as tape or DSU. Each copy created by an SLP can have a different retention period allowing copies on higher cost storage to expire earlier than those on lower cost (slower access) storage.

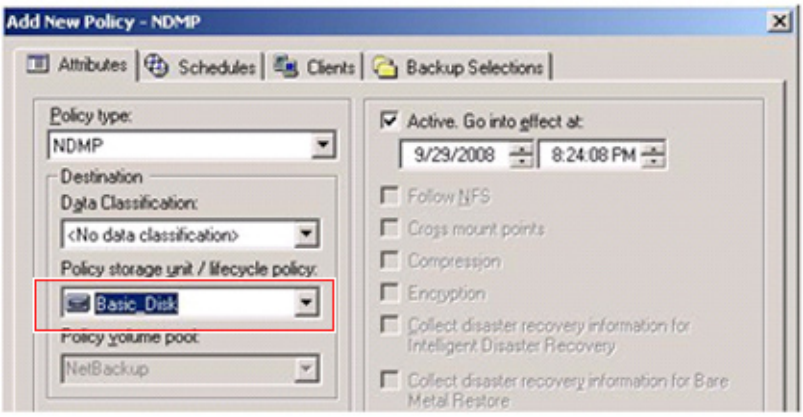

Figure 8: Disk Storage Unit

4.2 Shared storage option configuration

When configuring NDMP backups into a new or existing SAN environment, many customers want to share physical for virtual (in a VTL) tape drives within a library among NetBackup media servers and NDMP hosts. NetBackup supports sharing all supported SAN-attached physical/virtual tape drives between media servers and NAS, as long as the NAS vendor also supports the attached tape drive and their software supports SCSI Reserve/Release or SCSI Persistent Reservation. This support can be determined from the NetBackup NDMP HCL listing. This is easily configured with the NetBackup Device Wizard. For a more information about configuration, refer to the NDMP and Shared Storage Options Administration guides.

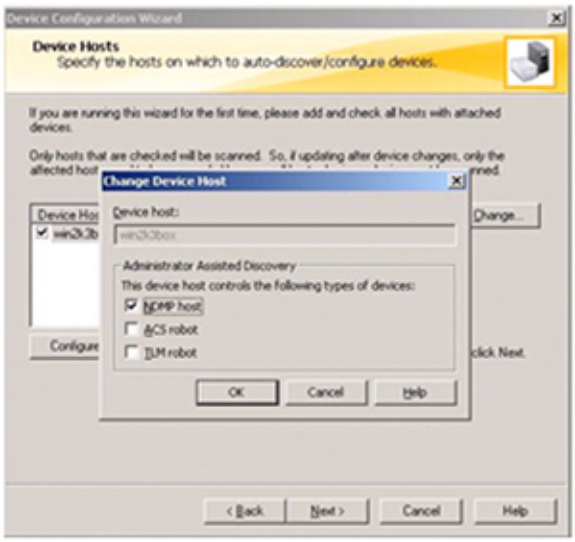

Figure 9: Shared Storage Option for NDMP Host

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world. The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

02/09 20016956