



NETWORKER GEARS UP FOR CLOUD TRANSFORMATION AND MULTI-TENANCY



Anuj Sharma

EMC²

Table of Contents

Introduction	3
NetWorker 8 Architectural Enhancements	5
NetWorker Server Broker Code Redesign	5
JobsDB Redesign.....	6
Backup To Disk Enhancements.....	7
NDMP Backup Enhancements	9
New Attributes.....	10
NetWorker Firewall Requirements	11
NetWorker Server Service Port Requirements.....	11
NetWorker Storage Node	12
NetWorker Client.....	12
NetWorker Management Console.....	13
NetWorker Upgrade Methodology.....	15
Considerations for NetWorker Server Hosts File Changes	17
Adjust Antivirus Software Settings.....	18
Deciding on Backup Schedules and Policies	18
Deciding on Data Domain Configuration	19
Implementation Best Practices.....	24
NetWorker/Data Domain Best Practices	28
NetWorker Multi-Tenancy Best Practices	31
Troubleshooting Methodologies	33
References.....	36
Glossary.....	36

Disclaimer: The views, processes, or methodologies published in this article are those of the authors. They do not necessarily reflect EMC Corporation's views, processes, or methodologies.

Introduction

Enterprises using EMC NetWorker® with Data Domain® DD Boost have seen huge CapEX and OpEX benefits due to industry leading dedupe ratios and throughput. As enterprise interest grows toward public and private cloud offerings, NetWorker has evolved, recently introducing major underlying architectural changes. These changes have made the product more stable, flexible, efficient, robust, and interoperable with cloud infrastructures.

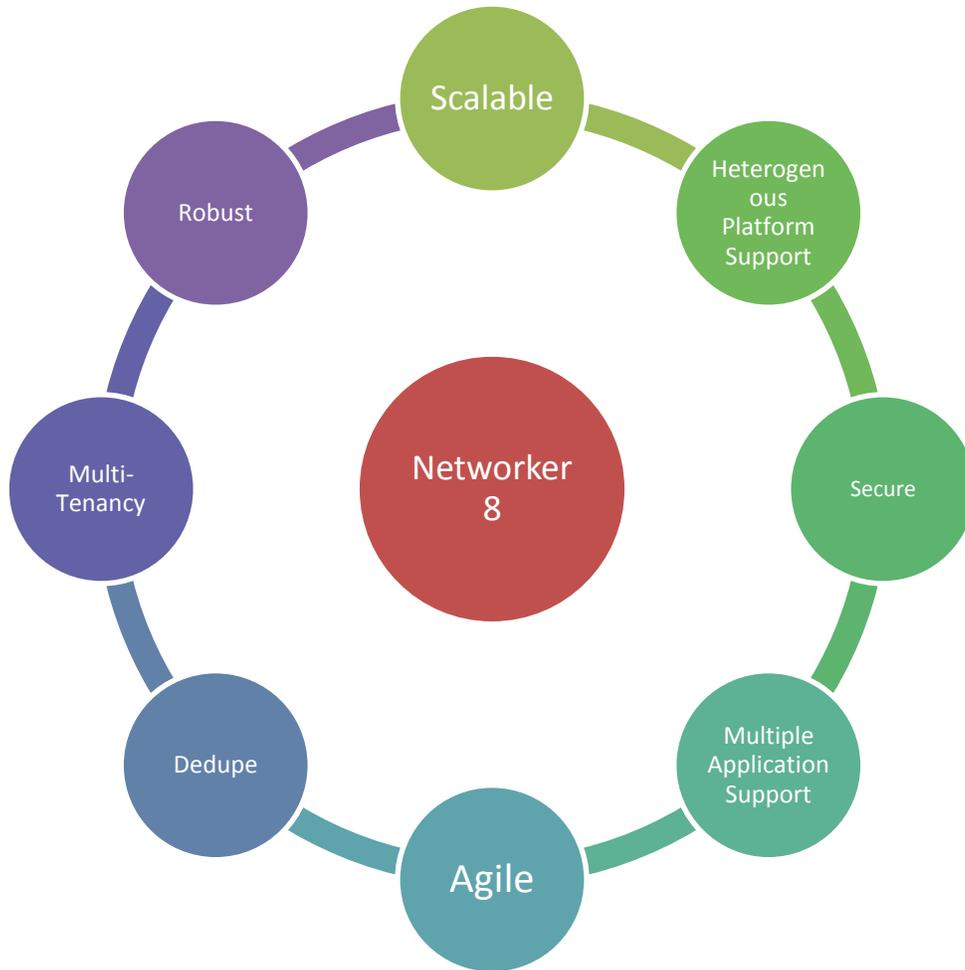


Figure 1: NetWorker 8 Features

This Knowledge Sharing article will examine the architectural changes of NetWorker 8, the benefits they provide, and best practices to make the most of the changes. Topics discussed in this article include:

- Backup to disk Enhancements and Best Practices.
- NSRMMD Enhancements and Benefits
- New NetWorker Daemons and Architecture
- NDMP Enhancements
- Security Enhancements
- JobsDB Enhancements
- Restricted Datazones for Multi-Tenancy Benefits and Implementation
- Introduction and Implementation of Synthetic Backups
- NetWorker Client Direct Architecture
- NetWorker Firewall Rules as per the new architecture changes in NetWorker 8
- Upgrade RoadMap from NetWorker 7.6 to NetWorker 8

We will also explore troubleshooting methodologies that will help an administrator to troubleshoot backup-related issues.

We will discuss how the introduction of multi-tenancy in NetWorker 8 will help organizations, the best practices to implement it, and how NetWorker 8 enables customers in their transformation to cloud.

NetWorker 8 Architectural Enhancements

NetWorker Server Broker Code Redesign

Challenge

It was often found in large environments that the NetWorker Server process (nsrd) was overloaded, leading to stability and performance issues.

Solution

NetWorker architecture has been changed from two-tier to three-tier by:

- Removing device monitoring from nsrd.
- Introducing one new master process nsrsnmd (storage node management daemon) to offload the device management workload from nsrd.
- Reducing design firewall port requirement between server and storage node which we will discuss later in detail.
- Eliminating the need for nsrd perform direct device management and regular polling of devices, due to the introduction of nsrsnmd.

Figure 2 shows that a new daemon NSRSNMD will be responsible for monitoring the storage node daemon NSRMMD which will offload the continuous monitoring/polling workload from NSRD daemon. NSRSNMD will be responsible for sending alerts to NSRD Daemon.

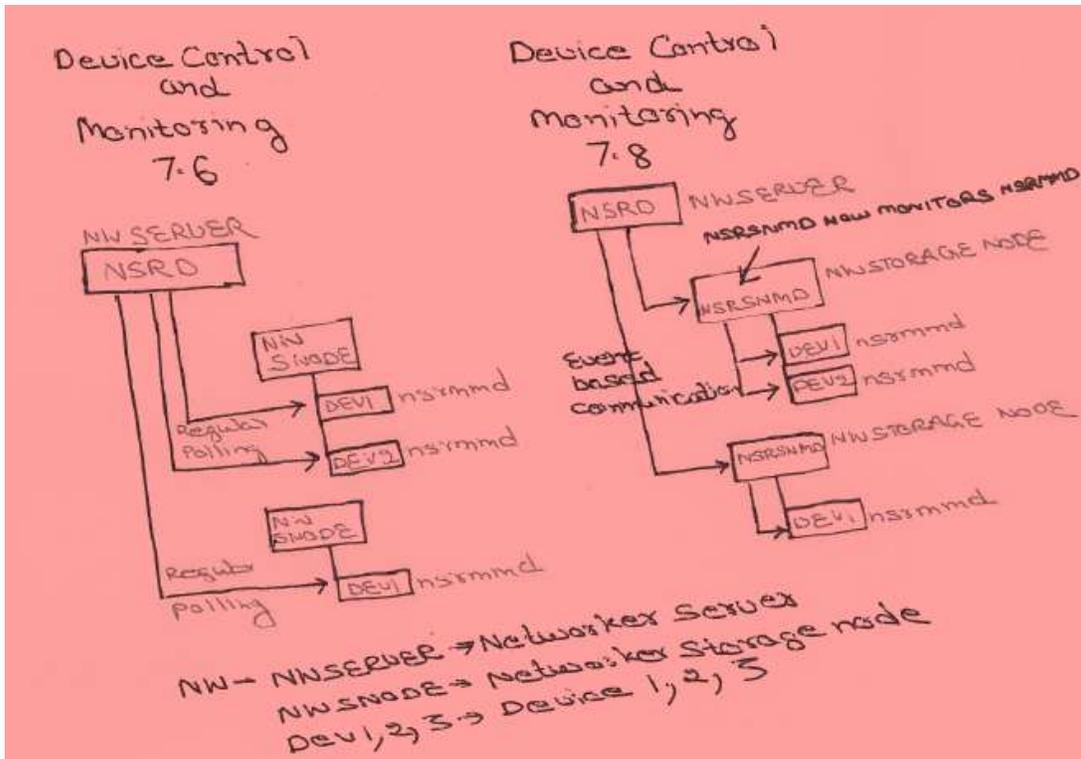


Figure 2: NetWorker Broker Redesign

JobsDB Redesign

Challenges

JobsDB database used to keep track of all backup jobs was a RAP Database, i.e. text-based database.

In NetWorker versions 7.3 to 7.6, if the environment was very busy and complex, the Jobs database could grow rapidly due to all the information being written to and stored in a text-based database. This would often cause performance issues—usually, high CPU usage as the database continued to grow in size. To streamline things and prevent performance issues, purging older data from the database was often necessary.

Since it requires regular purging, JobsDB caused performance issues in very large NetWorker environments.

Due to the RAP-based database, NetWorker Management Console typically experienced slow response in large environments. Also, sometimes it didn't display group details, etc.

Solution

In NetWorker 8.0, SQL has been introduced for JobsDB. The new SQLite database will allow all of the “jobs” information to be stored more efficiently in a new traditional database. This will result in numerous improvements in resource utilization, performance, and scalability.

However, the new database is no longer searchable as was the database used in previous versions of NetWorker.

The new database does not increase the hardware requirements needed for NetWorker. The size of the JobsDB is expected to remain below 1GB in size.

Backup to Disk Enhancements

Challenges

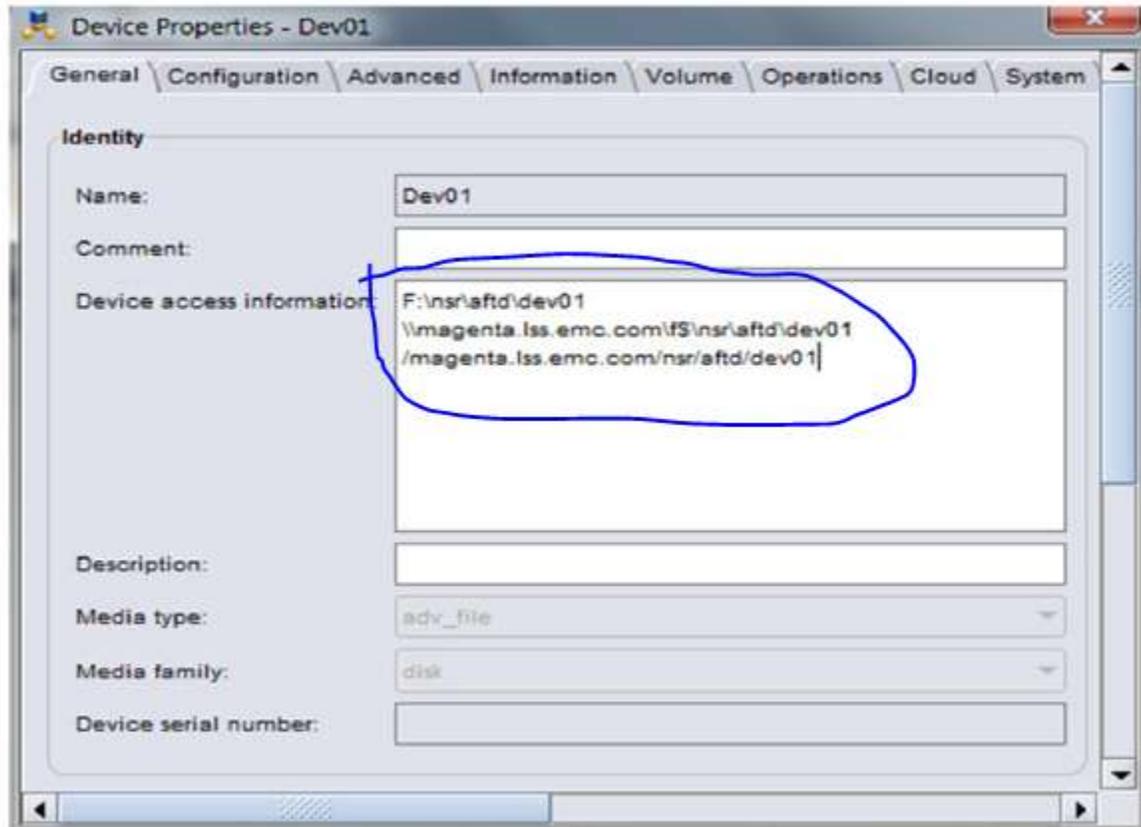
- Earlier NetWorker versions have two copies of the same device; one for read/write, another for Read Only.
- Concurrent recovery and cloning was not possible due to the type of tape simulation behavior.
- Single nsrmmd per device bottleneck. Each device was managed by single nsrmmd process, creating a daemon bottleneck .

Solution

- In NetWorker 8, there is no two instances of a single device.
- Multiple nsrmmd processes can work on the single device, eliminating the need for a dedicated .RO device.
- Nsrmmd process can be created dynamically or statically.
- Multiple nsrmmd process improve load balancing and performance.
- Each nsrmmd can handle multiple parallel requests of same type (save/recover/clone.)
- Earlier Device name was identified by the path of the device. Therefore, multiple paths for a single device can be listed in the new field introduced, i.e. Device Access Information.
- Same Volume can be mounted in multiple devices, enabling better use of the storage nodes. This means that a new session can be distributed to any other nsrmmd seeing the same volume, i.e. backup of Storage Node 1 and Clone on Storage Node 2.

Example

- Device on Storage node 1 : name dev01; path rd=sn1:/nfs_server/path_to_vol1
- Device on Storage node 2 : name dev02; path rd=sn2:/nfs_server/path_to_vol1



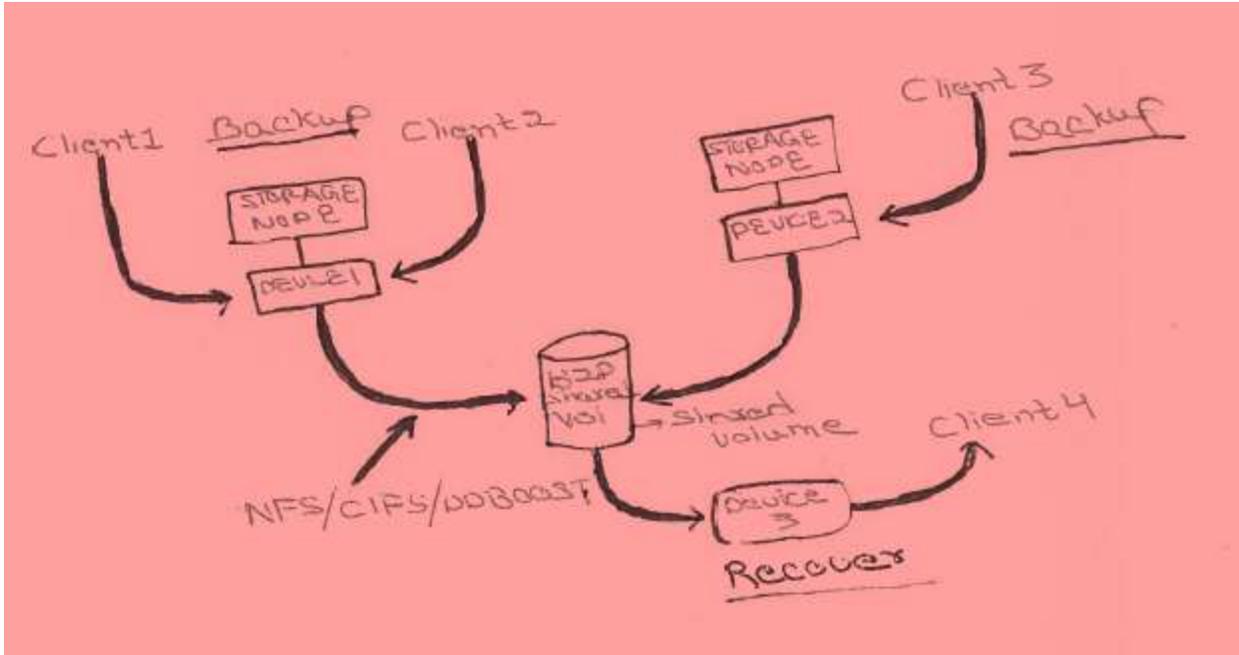


Figure 3: Backup to Disk Enhancements

NDMP Backup Enhancements

Challenges

In huge NAS environments with sizes of shares running in 100's of terabytes, a backup that fails to write most of the data has to be re-started from the beginning leading to the loss of backup media as well as backup time.

Solution

The re-startable backup feature has been extended to include NDMP in NetWorker 8.0. NDMP re-startable backup maintains the backup process if the backup job is uninterrupted. By default, a checkpoint where the backup can be restarted is written every 5 gigabytes. Backup is continued from the last checkpoint, saving significant time and allowing completion within the backup window.

New Attributes

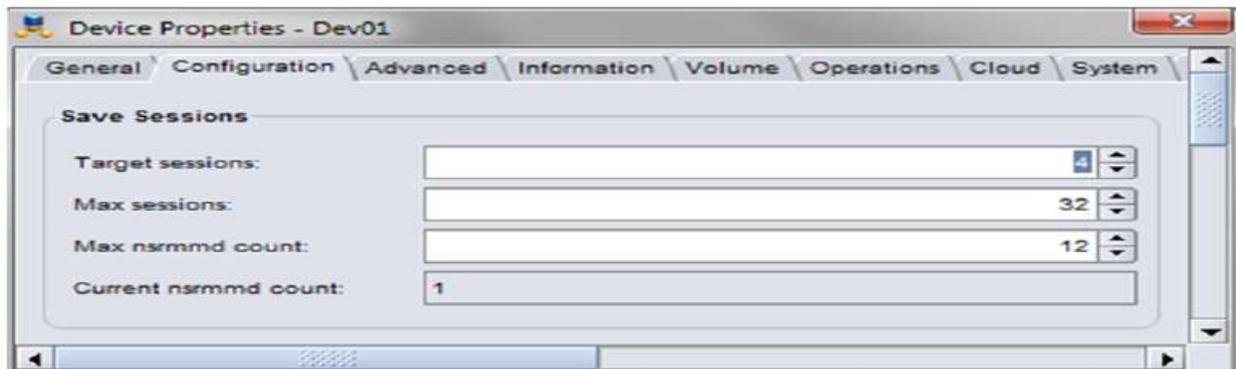
- The new max nsrmmnds attribute is used to define how many nsrmmnd processes can be started for a given device. the new default settings for target sessions and max sessions for each given device are:

For AFTD: TS = 4, MS = 32, MaxMMD = 12

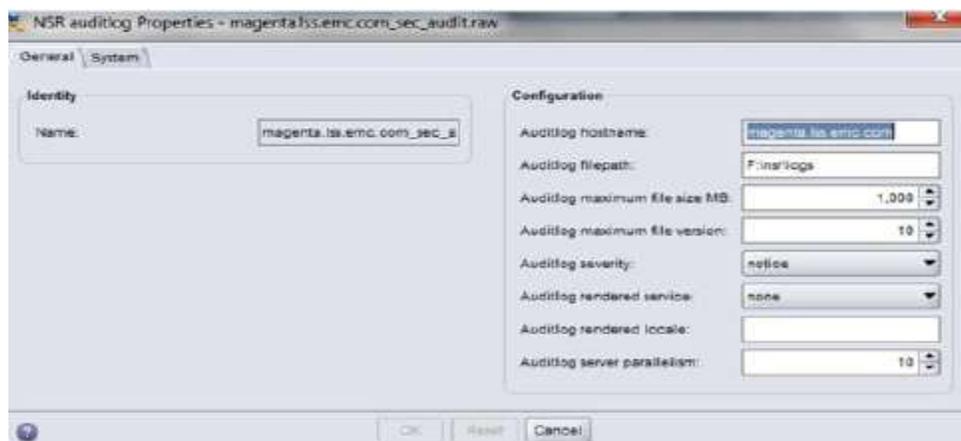
For DDBoost: TS = 6, MS = 60, MaxMMD = 14

To simulate NW7.6 behavior, set the MaxMMD setting to 2

Note that the settings for Max MMD will be automatically populated with the recommended settings for optimal performance but values can be overwritten if required.



- The audit log for NetWorker 8.0 is a property setting for every NetWorker client, storage node, and server and must be configured on a per client basis. It is modifiable only by a security or a full administrator. The audit log is not a replacement for the RAP log and does not have to necessarily be located on a NetWorker server. It can be on a separate secure system but the secure system must have the NetWorker client installed. Multiple systems can send their audit data to the same audit log server.



NetWorker Firewall Requirements

Due to the change in underlying architecture, there has been a change in the port requirements for NetWorker. Ports required between NetWorker Server and Storage Node has been reduced. Now nsrd communicates with nsrsnmd daemon instead of the nsrmmmd daemon's running on the Storage node.

NetWorker Server Service Port Requirements

nsrd	1
nsrexecd	4
nsrindexd	1
nsrjobd	1
nsrlcpd	One for each jukebox managed by the NetWorker server's storage node server's
nsrlogd (If the NetWorker server is the Audit Log server)	1
nsrmmmd (AFTDs or DD Boost devices only)	*Sum of the Max nsrmmmd Count settings of all these devices managed by the NetWorker server's storage node
nsrmmmd (devices that are not AFTD or DD Boost)	One for each device managed by the NetWorker server's storage node
nsrmmdbd	1
nsrmmgd	1
nsrpush	1
nsrsnmd	1

NetWorker Storage Node

nsrexecd	4
nsrjobd	1
nsrlcpd	One for each jukebox managed by the NetWorker Remote Storage node
nsrlogd (If the NetWorker Storage Node is the Audit Log Server)	1
nsrmmd (AFTDs or DD Boost devices only)	*Sum of the Max nsrmmd Count settings of all these devices managed by the NetWorker remote storage node
nsrmmd (devices that are not AFTD or DD Boost)	Two for each device managed by the NetWorker remote storage node

NetWorker Client

nsrexecd	4
nsrjobd	1
nsrlcpd	1 for each jukebox managed by the NetWorker Remote Storage node
nsrlogd (If the NetWorker Client is the Audit Log Server)	1

NetWorker Management Console

Console server requires the following ports:

- A HTTP port that is used for the Console-embedded web server, i.e. 9000, and is used to download the NMC user interface Java application.
- An RPC port 9001 is used for calls from the Console Java client to the Console server.
- The database queries port is 2638.
- An SNMP port 161 for getting information from the Data Domain system.
- An SNMPTRAPD port 162 for capturing Data Domain SNMP traps.

As per best practices, Connection Ports should be chosen as 10001-30000 using the nsrports command on the Backup Server and the Servers behind the firewall.

For Service, ports 7937-9936 should be opened as per best practices and nsrports should be specified on the Backup Server and the Servers behind the firewall to specify the port range.

Example

```
nsrports -s anuj1.anuj.com -S 7937-7938 10000-11000 11500
```

In this case, the service ports range for host: anuj1.anuj.com is set to use ports 7937-7938, 10000-11000, and 11500.

Example

The simplest configuration for the above firewalled environment will be a configuration that allows all the necessary communication through the firewall by opening 33 Service ports opened for incoming communication to the server. Among these 33 ports, 19 can also be used for outgoing communication from the server to the clients and storage nodes .

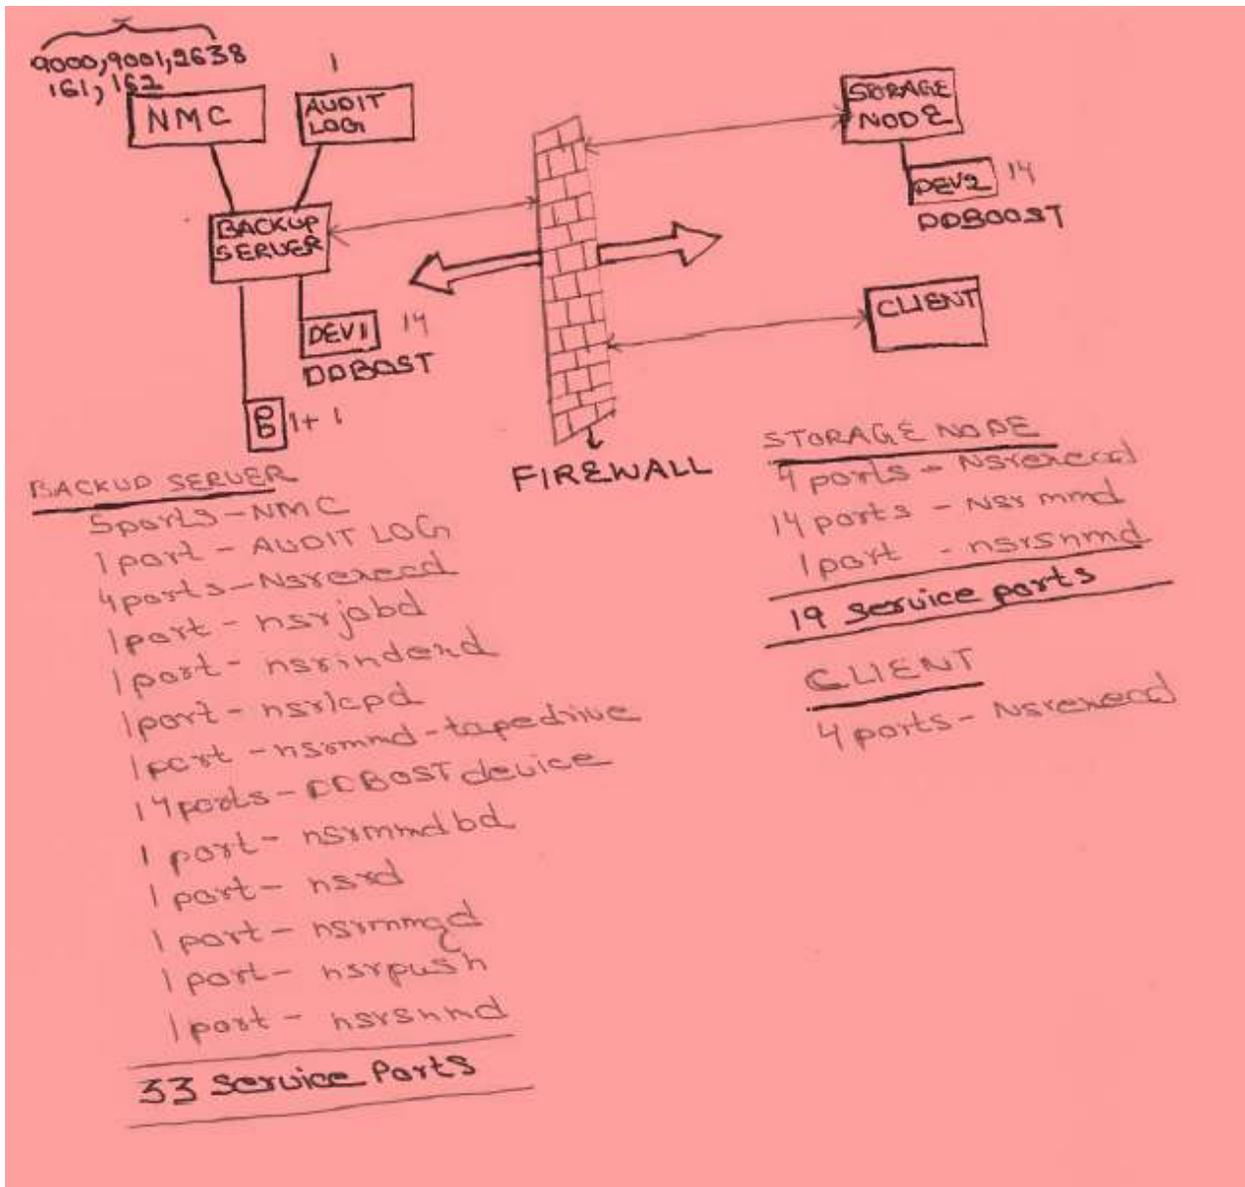


Figure 2: NetWorker Firewall Example

NetWorker Upgrade Methodology

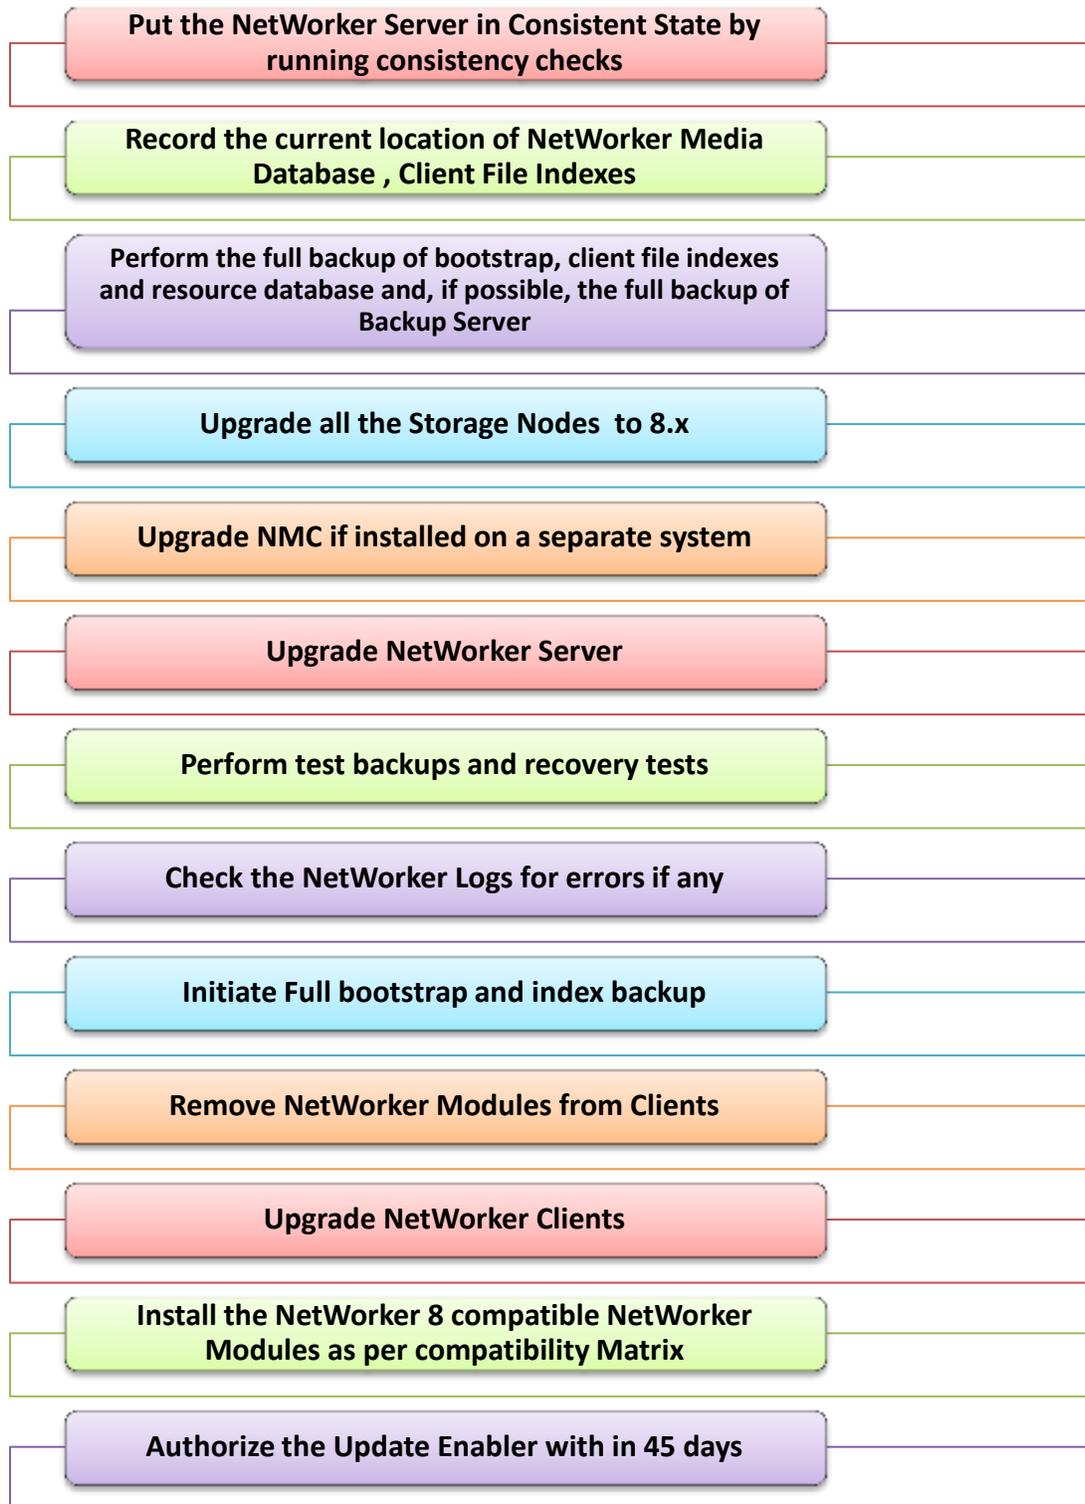


Figure 3: NetWorker 8 Upgrade Methodology

- Before updating the NetWorker server, ensure that:
 - The media database and client file indexes are in a consistent state.
 - You record the database locations.
 - You perform a backup of the NetWorker server databases.
- Put the NetWorker databases in a consistent state:
 - nsrim -X
 - nsrck -m
 - nsrck -L6
- Record the current location of the NetWorker media database:
 - nsrls -m
- Record the current location of the NetWorker client file indexes:
 - nsrls
- Perform a back up of the bootstrap, the client file indexes, and the resource database on the NetWorker server.

For example:

- savegrp -O group

The specified group must contain all of the NetWorker clients in the datazone. This ensures that all client file indexes are backed up. If a group that contains all of the clients does not exist, run the savegrp command more than once, specifying a different group each time, until all client indexes are backed up.

- Record the latest bootstrap save set ID (ssid) including the file number, the record number, and the associated volume label.

For example:

- mminfo -B

<i>date</i>	<i>time</i>	<i>level</i>	<i>ssid</i>	<i>file</i>	<i>record</i>	<i>volume</i>
10/11/11	16:29:40	full	4254377781	0	0	bootstrap_vol.001

In this example:

The save set ID (ssid) is 4254377781.

The file number is 0.

The record number is 0.

The label of the volume which contains the bootstrap save set is bootstrap_vol.001.

- If possible Backup the nsr directory of NetWorker Server.
- If possible, as a best practice Install the current version of NetWorker Server on another standby server and restore the nsr directory on that server so that in case during the upgrade if any issue occurs we can rename the standby server to the current backup server name and redirect the backups.
- Upgrade all Storage Nodes.
- Upgrade NMC if installed on a different server other than Backup Server.
- From the Control Panel, select the appropriate program to uninstall application software. When the machine is the Console server, uninstall the Console server software package before the NetWorker software package:
 - Select NetWorker Management Console and click Uninstall.

- Select NetWorker Management Console Server and click Uninstall.
- Now select NetWorker and click Uninstall.
- When the following window appears on a Windows systems with the vClient application running, click Ignore.

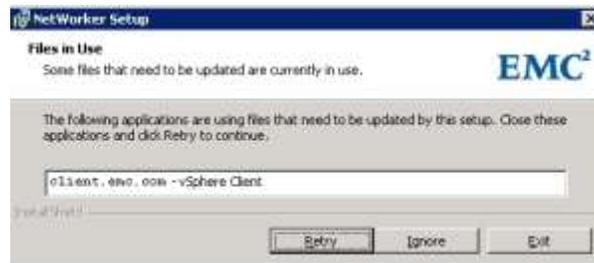


Figure 4: Upgrade Procedure

- Restart the NetWorker Daemons.
- Upgrade is successful now.
- Remove NetWorker Modules from the Clients.
- Upgrade the NetWorker Client Software.
- Install the NetWorker Module version compatible with NetWorker 8 as per compatibility matrix.

Considerations for NetWorker Server Hosts File Changes

NetWorker server always takes the first alias specified in the OS hosts file as the NetWorker server name. If the first alias in hosts file is different than the current NetWorker server name client resource, the next time the NetWorker daemons/services start, NetWorker will rename the NetWorker server to the first alias listed in the hosts file. This will assign a different clientid to the NetWorker server's client resource and will result in a failure to restore backups for both the NetWorker server and configured clients.

To avoid this issue, when modifying the hosts file ensure that the first alias for the IP address of the NetWorker server matches exactly the name defined for the client resource of the NetWorker server. Also, after the upgrade, all .RO devices will be removed as discussed earlier in the Backup to Disk Enhancement section.

Adjust Antivirus Software Settings

Undesirable behavior might occur if the antivirus software installed on a Windows machine is not tuned for backup environments.

Configure the antivirus software to:

- Avoid scanning files that are opened for backup:
 - Clear Opened for Backup in the Advanced Auto-Protect options for Norton Antivirus.
 - McAfee knowledge base article KB72334 describes how to disable scanning files that are opened for backup.
- Not monitor the following directories:
 - C:\Program Files\EMC or C:\Program Files\Legato
 - C:\Program Files\EMC NetWorker\nsr\res or C:\Program Files\Legato\nsr\res
 - C:\Program Files\EMC NetWorker\nsr\mm or C:\Program Files\Legato\nsr\mm
 - C:\Program Files\EMC NetWorker\nsr\index or C:\Program Files\Legato\nsr\index
 - AFTD directories

Deciding on Backup Schedules and Policies

Deciding on the backup policies revolves around:

- What needs to be backed up.
- When it needs to be backed up.
- How long it needs to be retained.
- The off-peak hours available for backups.

These pointers will help us to decide upon the backup schedules and policies.

- All stakeholders should be taken into confidence before deciding on the backup policies. By stakeholders, I mean Application Owners, Server Owners, and Backup Administrators.
- A Backup Form should be developed and circulated among all the stakeholders to get their input regarding their expected Backup Schedules, Recovery Point Objectives, Data Retention, etc. A Sample form is shown in Appendix A.
- Once we have input from all stakeholders, grouping of the servers should be done as per the requirements and the design should be shared with each stakeholder for agreement.

- Some pointers that can help you devise the Backup Schedules, Retention Policies, etc.
 - Categorize the Backup Clients as per application. Agree on a common Backup Schedule for the Application and Retention.
 - For Exchange, Weekly Full Backup should be taken followed by Incremental. But make sure that Exchange Admin has Circular Logging Disabled otherwise Incremental backups will be promoted to full by NetWorker.
 - To enable SQL Server to reduce the Recovery Time Objective and Backup Window during Weekdays, Weekly Full Backups should be taken along with Daily Differential Backup. Also, to reduce RPO we can take Incremental Backups for SQL Servers at 2-hour intervals since Incremental will back up only the transaction Logs.
 - Dedicated Media Pools should be created for Backup Groups.
 - For FileSystem Backups, daily Incremental and weekly Full Backups should be taken.
 - For File Servers with data in 100's of terabytes, to accommodate the Backup Window, schedule monthly Full Backup once followed by Incremental backups during the weekdays and then Synthetic Full Backups introduced in NetWorker 8.x during weekends.

Deciding on Data Domain Configuration

Data Domain can be configured in two ways for backups—as a VTL or a number of DDBoost devices. The choice depends on various factors. Figure 8 depicts the factors that decide the Data Domain design.

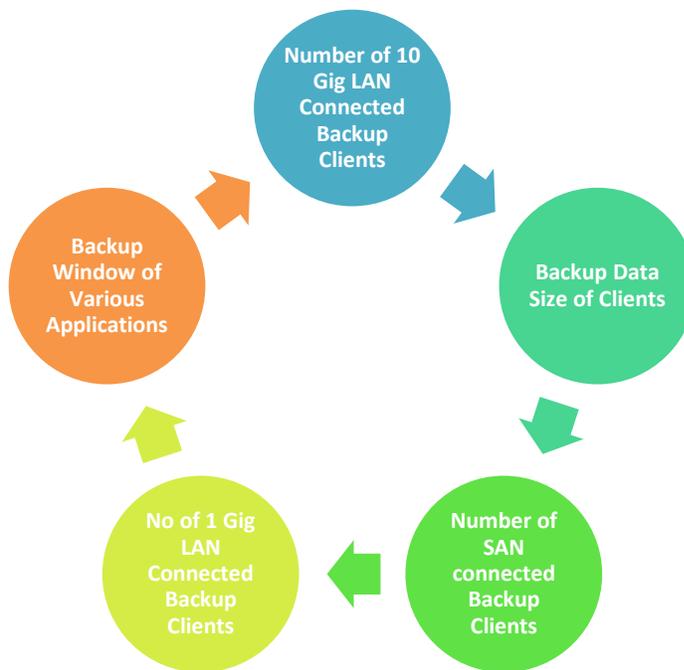


Figure 5: Data Domain Design

<p>SAN Connected Clients</p>	<ul style="list-style-type: none"> •VTL should be configured on Data Domain for SAN- connected clients and SAN-connected clients should be configured as Dedicated Storage Nodes with two Virtual Tape Library Tape Drives assigned.
<p>10 GigE Connected Clients</p>	<ul style="list-style-type: none"> •DD Boost devices should be created on Data Domain and assigned to 10 GigE-Connected Clients for Backups using Client Direct
<p>1 GigE Connected LAN Clients</p>	<ul style="list-style-type: none"> •For 1 GigE Connected Clients, their backups should be directed to Backup Server and Backup Server should be assigned VTL Tape Drives or DD Boost devices as per the connectivity of Backup Server. •For 1 GigE Connected Clients with stringent Backup Windows, we can assign the DDBoost device to the Client but it is preferable to have a dedicated Backup LAN Card for the Client .

Backup Window is the other aspect that has a major role in deciding the configuration irrespective of whether the Client is LAN/SAN Connected. The table below should be used as a reference to decide upon the configuration.

Client Name	Backup Size	SAN Connected	1GigE Connectivity	10GigE Connectivity	Backup Window	Device Configuration
Exchange 1	1.5 TB	Yes	Yes	No	6 Hours	Dedicated Storage Node 2 Tape Drives
SQL 1	1.5 TB	NO	NO	YES	6 Hours	Two DD Boost Devices
FileServer 1	1.5 TB	NO	Yes	NO	12 Hours	Two DD Boost Devices

*Assuming each Tape Drive and 10 GigE Connected DDBoost device can deliver average throughput of 40 MB/sec

*Assuming each 1 GigE Connected DDBoost device can give us a average throughput of 20 MB/sec.

Also note that the backup throughput also depends on:

- Size and number of files being backed up. If there are thousands of small files on a server, expect less backup throughput for that server as the Server will have to do many I/O operations and indexing of the files being backed up.
- Processing Power of the Server and resources available on the Server at the time of backup. If Server resources are scarce, backup throughput will be low depending upon the resources available.
- Backup Destination Device Throughput. If there are many servers backing up to the same set of Backup Devices, backup throughput will be shared across all the servers as in the case of LAN-based backup Servers where the Servers share the same backup Destination Devices.

Blocksize	Local backup performance	Remote backup performance
64 KB	173 MB/second	60 MB/second
128 KB	173 MB/second	95 MB/second
256 KB	173 MB/second	125 MB/second
512 KB	173 MB/second	130 MB/second
1024 KB	173 MB/second	130 MB/second

LTO-4 Tape Drive Performance as per the I/O Block Size

In essence, a NetWorker/DataDomain backup environment will consist of the following.

- SAN-connected Clients can be backed up to the VTL tape drives on the Data Domain. The number of tape drives assigned to the Client will depend on the data on the Client assuming 40 MB/sec throughput of the tape drive.
- 10 GigE Clients shall be DD Boost devices using the Client Direct feature of NetWorker 8.x. The number of devices will depend on the data on the Client assuming 40 MB/sec throughput of each DD Boost device.
- Please note that DD Boost devices take additional resources from the Server. Ensure that there is no resource contention on the Server.
- Allowing for other types of devices and services on a typical storage node, a storage node should have a minimum of 8 GB of RAM if hosting DD Boost devices.
- DD Boost clients require a minimum of 4 GB of RAM at the time of backup to ensure optimum performance for Client Direct backups.
- Each DD Boost device requires an initial 24 MB of RAM on the storage node and Client Direct client. Each DD Boost save session requires an additional 24 MB. To run 10 sessions requires 24 + 240 MB. The default max sessions of 60 sessions per DD Boost device requires 24 + 1440 MB.
- Ensure there is enough CPU power on the client to take advantage of DFA-DD increased performance capability. In most cases, Client Direct significantly improves backup performance. The DFA-DD backup requires approximately 2-10% more CPU load for each concurrent session.
- 1 GigE Clients shall be directed toward the devices configured on the Backup Server.

- There is aggregated throughput of each model of Data Domain and the accumulative throughput of all the devices configured should not exceed 80% of the Total Throughput Sustainable by the Data Domain box. Refer the datadomain support portal at my.datadomain.com for Maximum Throughput Sustainable by the different Data Domain boxes.

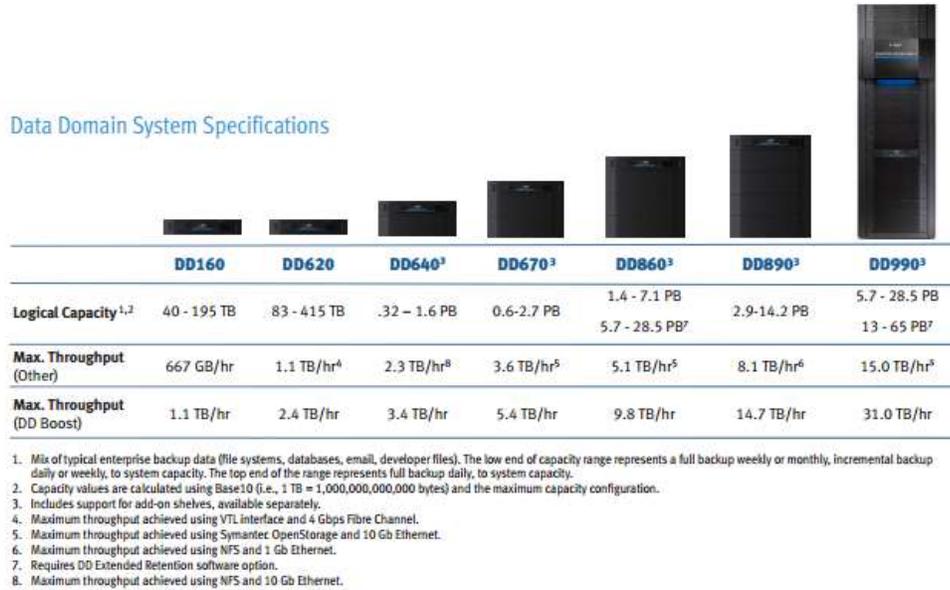


Figure 6: Data Domain Models

Implementation Best Practices

Once we have drafted the backup policies decided on the Data Domain configuration, we can proceed to the implementation phase. Below are some best practices that should be followed during the implementation phase to optimize the backup environment.

- Verify compatibility of the current Backup Server and Client OS to the NetWorker version. Be aware that NetWorker Management Console no longer supports any flavor of Windows 2003; NetWorker Backup Server is Supported Windows 2003 SP2 onwards.
- Ensure that the HBA firmware on the Servers connected with the SAN are updated.
- We should have dedicated FC Ports for Backup traffic to isolate the Production and Backup traffic.
- Single Initiator Zoning should be used and tape drives should be zoned to one fabric instead of multiple fabrics. This is because the Backup Server will see the same tape drive visible as two different drives, creating confusion for the backup server.
- Dedicated Backup LAN should be implemented to isolate the Production and Backup traffic.
- Ensure that the Client, Backup Server, and Ethernet Switches Network Setting are in sync, i.e. the network is configured as Full Duplex, Half Duplex, or Auto.
- Ensure that the drive on which the NetWorker is installed is a Mirrored or RAID 5 device to avoid the Backup Server as a performance bottleneck and achieve maximum performance. In large environments, use a dedicated RAID Disk for /nsr and dedicated disk for Indexes.
- NDMP backup method should be used for the NAS backup. In this method, the tape drives are zoned with the NAS device and NAS backup data travels over the SAN to the tape drive, ensuring optimum NAS backup performance.
- Target-based deduplication should be used for database files, as the deduplication process could eat up the host's CPU cycles, thus impacting application performance. DD Boost should not be used for Database clients.
- To optimize data deduplication, alter backup schedules and perform full backups more frequently in preference to incremental backups. This will increase the dedupe ratios and also result in speedy recoveries.

- Disable multiplexing, as it adversely affects dedupe ratios. Instead, create additional virtual tape drives for better throughput. Multiplexing interleaves backup streams—writing a little of save set 1, then a little of save set 2, and so on—so that none of the clients sending save sets need to wait for the other clients to finish. When using the Data Domain as a VTL with NetWorker, set both target sessions and max sessions on each virtual tape drive to a value of 1.
- Use of parallelism is not an issue when a Data Domain Storage System is used as Backup to Disk destination as each save set is stored as a separate file regardless of parallelism settings. In Data Domain systems that have been configured as both VTL and DDBoost devices, we can set both target sessions and max sessions on each virtual tape drive to a value of 1 and leave the target sessions and max sessions for DDBoost devices at default or tune them as necessary.
- Software compression should be disabled for data deduplication optimization. Since deduplication storage systems compress data after deduplication, doing so prior to deduplication results in the deduplication system having to perform additional tasks during the deduplication process.
- To achieve higher dedupe ratios and ensure data deduplication optimization, divide backup DDBoost clients in batches and start the backups one after another. This ensures that redundant data does not travel across the LAN, optimizing LAN/WAN link use. On the other hand, if backups are triggered simultaneously, data first travels across the LAN to the backup destination and then gets deduplicated.
- Update hosts file of the backup server with IPs, Hostnames, and Fully Qualified Domain Name (FQDN) of backup clients before starting the implementation.
- Disable the “Removable Storage Service” from the services panel if the backup server uses Windows as its OS. This is because this service is used by NTbackup program—the native Windows backup utility—which interferes with NetWorker services and can cause the tape to fill prematurely.
- Enable RAP by going to the backup servers’ properties after you install NMC and the NetWorker Server. After enabling RAP NetWorker, make a rap.log file that contains all the changes that an administrator makes on the NMC (i.e. client resources, attributes, media pools, backup devices). This will help administrators to track changes. One of our clients complained that their weekly full backup had not been done but after seeing the rap.log file we found that someone from the client side had modified the schedule resource and set it to skip on the weekend.

- Update the aliases of all the backup clients and the server interface attribute with the backup servers' IP to minimize connectivity issues. It is always better to divide a single saveset with data in terabytes into many to improve backup speeds, such as splitting a single saveset D:\ into various savesets like d:\data, d:\important.
- As per best practices, a separate media pool shall be created for Index and Bootstrap backups by specifying index: , bootstrap: in the selection criteria of the pool .
- In the case of SAN Dedicated storage nodes with tape drives assigned, ensure that only the drives are visible to the Dedicated Storage nodes. Autochanger should be visible only on the Backup Server.
- There should be different media pools for DDBoost devices and tape volumes.
- Some parameters that can be tuned to optimize backup performance, i.e. Server Parallelism, Client Parallelism , Group Parallelism, and Target Sessions.

Server Parallelism

Controls how many savestreams the server accepts at the same time. Server Parallelism at a minimum should be equal to the sum of all of the target sessions for each defined device on the NetWorker server (included devices in remote storage nodes.)

Target Sessions

Sets the target number of savestreams to write to a device at the same time. Because this value is not a limit, a device might receive more sessions than the target sessions attribute specifies (can happen when server parallelism exceeds sum of all target sessions). The larger the number of sessions you specify for target sessions, the more save sets are multiplexed, or interleaved, onto the same volume. However, be aware that the more the streams are interleaved the more time it will take to recover.

Client Parallelism

Controls how many savestreams a client can send at the same time. Its value should be equal to the sum of the savesets for each client (i.e. a Windows 2008 client with the c:\, d:\, SYSTEM STATE and SYSTEM DB defined would have a client parallelism of 4.)

Savegroup Parallelism

Number of server parallelisms that can be allocated for that savegroup. This is useful if you have more than one savegroup running at the same time and would like one group to take precedence over another. A value greater than 0 will override other parallelism settings the savegroup might use to avoid overutilizing system resources.

Thus, set the server Parallelism and Target Sessions attributes so that the total performance of the disk drives equals the total performance of the tape drives. There is no benefit to setting the Parallelism attribute to a higher value.

To select the correct values for the Parallelism and Target Sessions attributes, use the following equation:

$$\text{Parallelism} = \text{Number of Devices} * \text{Target Sessions}$$

For example, if you have three tape drives available for backup and you want each tape drive to accept two savestreams, set the value of server parallelism to 6 and the value of target sessions to 2.

When a NetWorker server is saving a large number of save sets, such as 500 or more, memory consumption and file descriptor consumption can reach values that are close to operating system limitations. In this event, the parallelism may need to be lowered. Decrease the server Parallelism and Target Sessions attributes to unload an overworked NetWorker server. With the correct settings, the normal operation of the computer should not be interrupted by backups or other NetWorker server activities. More memory may be needed to handle a higher parallelism setting if required. For the backup server, set the highest possible client parallelism to ensure that index backups are not delayed. This ensures that groups complete as they should.

- Ensure that scanning applications such as Antivirus are not running at the time of Backups.

The Client Server Disk can become a bottleneck if parallelism is defined for a client with non-RAID single physical disk. Take care while deciding on the Client Parallelism. If the Source has RAID Protected disk or has multiple non-RAID disks then we can have parallelism as per the number of disks or as tuned. For example, if there are four non-RAID disks, parallelism can be defined as 4; if there is one non-RAID disk, parallelism can be defined as 2 or 1.

- Avoid assigning more than 50 Backup Clients to a single backup group. Divide the number of clients into multiple groups and stagger the start times of the group such that all backups don't get triggered at the same time, leading to resource contention on the Backup Server.
- In real backup environments, manufacturer throughput and performance specifications based on theoretical environments are rarely, or never achieved. It is a best practice to never exceed 70 percent of the rated capacity of any component. Components include:
 - CPU
 - Disk
 - Network
 - Internal bus
 - Memory
 - Fibre Channel

Performance and response time significantly decreases when the 70 percent utilization threshold is exceeded.

Physical tape drives and solid state disks are the only exceptions to this rule, and should be used as close to 100 percent as possible. Neither suffers performance degradation during heavy use.

NetWorker/Data Domain Best Practices

- It is recommended not to send encrypted data to the Data Domain to achieve high dedupe ratios. If encryption is mandatory, avoid frequent changes to the data zone passphrase. This will make the current version of an unchanged file identical to the prior version.
- Disable multiplexing, as it adversely affects dedupe ratios. Instead, create additional virtual tape drives for better throughput. Multiplexing interleaves backup streams; writing a little of save set 1, then a little of save set 2, and so on, so that none of the clients sending save sets need to wait for the other clients to finish. This behavior has significant impact on deduplication efficiency when the Data Domain Storage System is used as a virtual tape library (VTL) because there are less chances of finding the redundant data because data from various clients are interleaved and it's difficult to find common data blocks because of the additional header information added to the data

with parallelism. When using Data Domain as a VTL with NetWorker by setting both target sessions and max sessions on each virtual tape drive to a value of 1.

- It is recommended to schedule full backups as soon as possible when Data Domain is introduced in an existing environment. If full backups cannot be scheduled out of the policy schedule, it's recommended that backups be directed to Data Domain on the day of Full Backup Schedule. This is due to the fact that if a recovery is required, data will need to be restored from both the previous and current Data Domain destination.
- As discussed earlier, since every device requires memory and CPU resources on the storage node and the NetWorker server, be vigilant while adding devices to the Backup Server or Storage Node. As a rule of thumb for a environment with more than 100 devices, after reaching 100 devices add 20 percent of the desired number at a time to see the impact of adding the devices and proceed in steps to avoid any adverse affect .
- When implementing a Data Domain VTL, consult the FC_Switch Compatibility List for the software release that is applicable to the specific Data Domain Storage System.
- In a SAN environment, a best practice is to limit FC extended fabric (ISL link) configurations to three hops between the backup server/storage node and the Data Domain Storage System.
- Ensure that persistent binding is enabled on the Operating System where VTL is assigned from the Data Domain .
- Use a dedicated network by configuring a separate network or use QoS features that guarantee network bandwidth. Alternatively, use virtual networks (VLANs) to segregate backup from production network traffic.
- As a best practice, do not use a single Data Domain box for more than 14 active NetWorker data zones.
- DD Boost devices use multiple concurrent nsrmmd (data mover) processes per device and multiple current save sessions (streams or threads) per nsrmmd process. Optimal device configuration for backup or clone operations reduces the number of active devices required and thereby reduces the impact on Data Domain system performance and maintenance. Balance session load among the available DD Boost devices so that new sessions attach to devices with the least load. For optimum performance, adjust the device Target sessions, Max sessions, and Max nsrmmd count attributes.
- Network Connectivity Best Practices for Data Domain System.
 - 10 GbE Ethernet connectivity is always recommended.

- Minimum of two 1GbE links each connected to a different switch and one 1GbE for administration.
- In environments where 10 GbE connectivity is not available or cost-prohibitive, two alternatives are available:
 - ✓ Use a dedicated 1 GbE connection from a storage node directly to the Data Domain system. This provides a private, high-bandwidth data connection and avoids the latency and complexity of a shared Ethernet connection. However, a separate traditional Ethernet connection is also required for administration and NMC access.
 - ✓ Use two or more NICs on the Data Domain system with 1 GbE connections aggregated by using the Data Domain ifgroup command. This will provide increased capacity and offers some resiliency. The Data Domain system provides automatic load balancing .
- Firewall port requirement between Data Domain, NetWorker, and NMC servers:
 - TCP 111 (NFS portmapper)
 - TCP 161 (for NMC server to query for alerts and statistics)
 - TCP 162 (SNMPTRAP for NMC server to monitor status and events)
 - TCP 2049 (NFS)
 - TCP 2051 (Replication, if clone-controlled replication is used, Data Domain to Data Domain systems)
 - TCP xxxx (select a random port for NFS mountd, 2052 is the default)
 - On the Data Domain system, type the following command from SE mode:


```
# nfs set mountd-port xxxx
```
- A Media pool in NetWorker that contains the DD Boost devices should not contain any other devices other than Data Domain devices. Also, devices should not be mixed from different Data Domain boxes in a single Medial Pool.
- In cases of Clone Controlled Replication, limit the number of target pools for simplification, i.e. if multiple pools have been created at the source site for different applications for weekly or daily backups, a single pool can be created at the destination for all weekly and all daily backups for simplification.

NetWorker Multi-Tenancy Best Practices

The architectural changes discussed earlier have given NetWorker 8.x a more robust underlying architecture to support cloud environments. These new architectural changes and features make NetWorker more agile, scalable, robust, and elastic, mandatory for any Cloud Backup Software.

NetWorker has added a new framework for multi-tenancy in cloud environments, enabling backup administrators to give control of Client Creation, Group Creation, Device Creation to a tenant, for example, an application owner. The application owner can create clients as per the quota assigned to him and corresponding devices. Tenants can view only the resources they are assigned without visibility of the other tenants in the environment. Figure 10 shows a Restricted DataZone created for a user rdzuser with the provision to create 4 clients, 2 devices, 3 storage nodes, and 3 juke boxes along with other privileges. The restricted data zone feature results in autonomy for tenants in a hosted (or service provider) environment and a simplified experience for NetWorker Administrators.

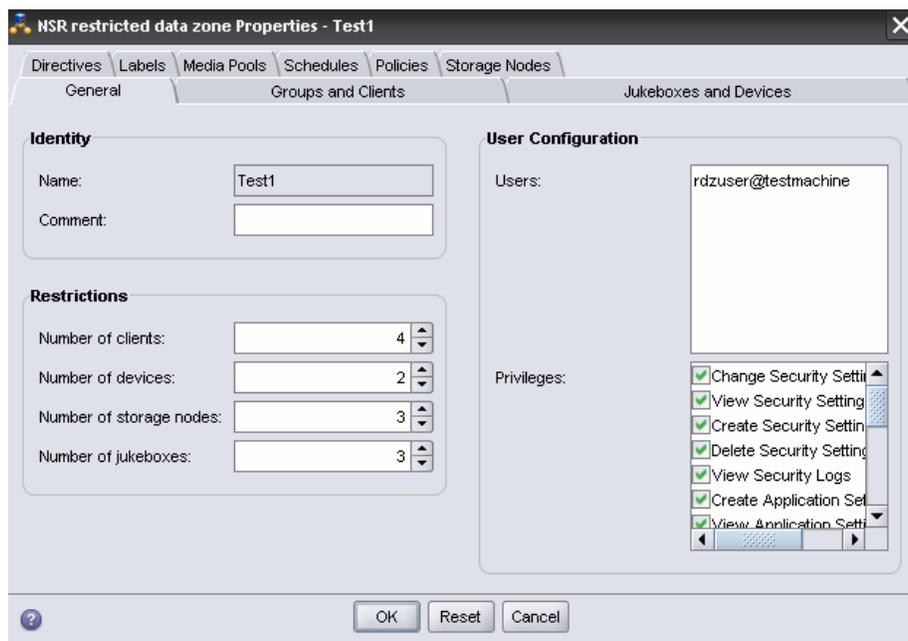


Figure 7: NetWorker Restricted DataZones

Some factors that should be considered while implementing the Multi-Tenancy with Restricted DataZones:

- With more flexibility will come more overheads for the administrator as the tenants will change the configuration on their own. A workflow should be drafted that whenever a tenant is planning a change, the Global Administrator is informed about it so that Global Administrator can provide recommendations.
- Multi-Tenancy should be done in a planned manner. The Global Administrator should decide the number of tenants that the current Global DataZone can support and should keep a continuous track of resource utilization so that resource over-utilization by one tenant doesn't affect another.
- Regular configuration audits should be performed by the Global DataZone Administrator for the different tenant configurations and environment recommendations should be provided if inconsistency is found or if improvements can be made .
- Regular Recovery Drills should be performed for each tenant at regular intervals; quarterly, etc.

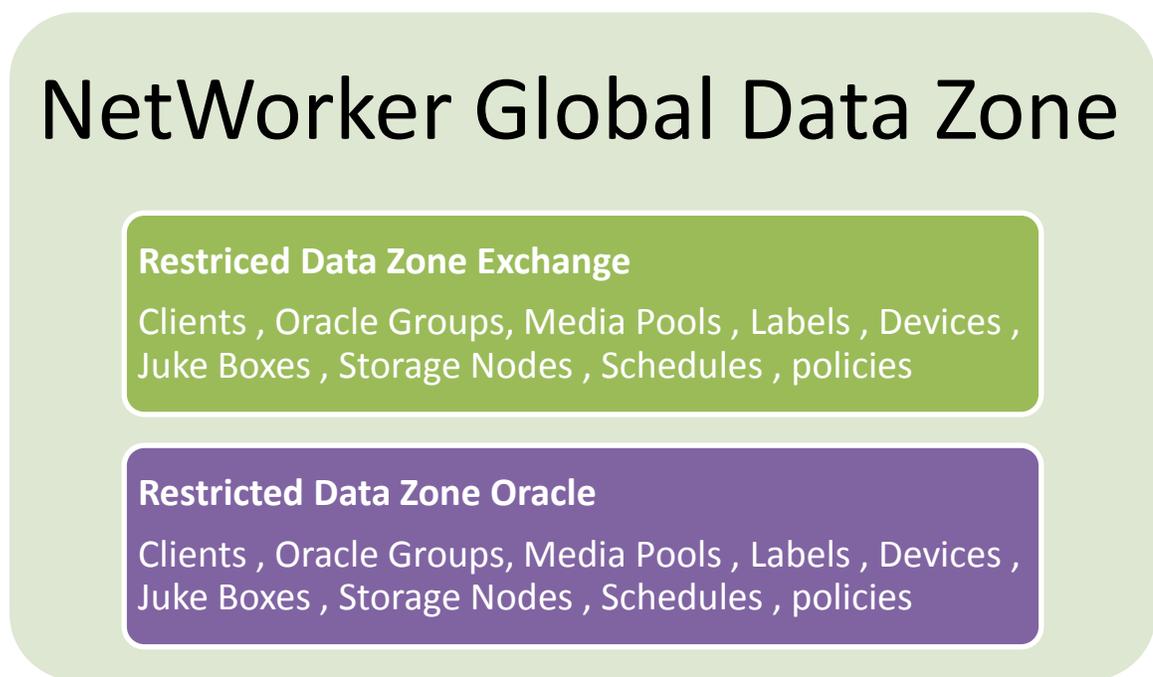


Figure 8: NetWorker Restricted Data Zone

Troubleshooting Methodologies

Once the environment is implemented, there may be days where you will encounter backup failures or backup throughput. This section discusses how to approach such issues and resolve them quickly and effectively.

Backup Failures

- If you are unable to determine the cause of failure from the SaveGroup Completion alert or GUI, first check the Daemon.raw file for errors when the backup failed.
- If still you are not able to find a valid reason, check the Daemon.raw located on the client machine.
- Check that NetWorker Services are running on the Client.
- If still you are not able to find the cause, check the connectivity between the backup server and client by using ping command , nslookup , rpcinfo .
- Check that the Backup Server is able to connect to the nsrexecd of the remote client by using the following command
 - **nsradmin -s <client_name> -p 390113**
 - Run the following command from the NetWorker client: **nsradmin -s**
 - **<server_name>, nsradmin -s <server_name> -p 390113** to test the NetWorker client connections to the nsrd and nsrexecd daemons on the backup server.
- If the command fails, there is likely a network issue; either a recent change in the firewall ports or Windows Firewall.
- If all the above commands are successful, run the following command for detailed debug logging
 - **savegrp -vvv -D9 -p -c source_client -G group_name > log.txt 2>&1**
- In my experience, most backup failures are because of changes in the environment, such as Firewall Changes, NIC Card Setting Changes on the Client, Low Disk Space on the Client, Services not running on the backup client, or Windows firewall being accidentally turned on if not turned on previously during the installation of NetWorker Client. It's obvious that if the backups were running correctly earlier, there might have been a change that has led to the failure and those must be found.
- If the client in question is hosting online Database Application modules, there might have been some change done by the administrator that the backups have started failing. For example, a password change, a setting being changed in the application, and so forth.

To troubleshoot application module backup failures, refer to the application log if we are not able to find the reason from daemon.raw. For example, nsrschsv.raw for exchange located in C:\Program File\legator\nsr\applogs\nsrxchsv.raw and nsrsqlsv.raw for sql located in C:\Program File\legator\nsr\applogs\nssqlsv.raw.

- With Microsoft SQL, I have seen that failures are mostly due to password changes by the administrator, or a database being offline.

Low Backup Throughput/Performance Issues

- Most Backup Throughput issues occur because of network changes, high resource utilization on the server, or inconsistency in the setting of NIC cards of Client and Backup Server. A best practice is to ensure that network settings are consistent throughout the network environment and complement each other, i.e. Full Duplex, Half Duplex, or Auto Negotiate.
- Another reason can be that the backup client has millions of small files causing low backup throughput compared to other clients as I/O's will increase and indexing of large number of files will add to the performance overhead.
- Some questions to ask to reach a solution or identify the cause of an issue:
 - Do backups perform better when started at a different time? Yes indicates that there is resource contention for the server. Backup should be scheduled for when the server is less utilized.
 - Is it consistent across all save sets for the clients? No indicates that the saveset experiencing low performance has millions of files or is being accessed by other applications at the same time of backups.
 - Is it consistent across all clients with similar system configuration using a specific storage node? Yes could indicate that the Storage Node is the bottleneck and Storage Node Performance should be monitored, as explained later.
 - Is it consistent across all clients with similar system configuration in the same subnet? Yes indicates that there a setting for this subnet differs from other subnets, affecting performance. Subnet settings should match. Also, we can match the operating system, LAN Card Drivers, and Patches Service pack of the client working as per expectation to the client with low performance to troubleshoot.

Monitoring Performance

Monitor I/O, CPU, Memory, Disk, and Network Performance using the native tools below at the time of backups to have an idea of what occurs during the backup process, i.e. resources consumed by each application before, during, and after the backups .

- Windows: Perfmon
- UNIX: iostat, vmstat, or netstat commands

If it is discovered that slow backups are due to excessive network use by other applications, this can be corrected by changing backup schedules. High CPU use is often the result of waiting for external I/O, not insufficient CPU power. This is indicated by high CPU use inside SYSTEM versus user space. On Windows, excessive time spent on Deferred Procedure Calls often indicates a problem with device drivers.

Determine Network Bottlenecks by using FTP

- Create a large data file on the NetWorker client and send it to the storage node via FTP.
- Make note of the time it takes for the file to transfer.
- Compare the time noted in previous with current backup performance:
 - If FTP performs much faster than the backups, the bottleneck might be with the tape devices.
 - If FTP performs at a similar rate, the bottleneck might be in the network.
 - If there is large difference in the transfer rate, or one type of FTP transfer has spikes, it might indicate the presence of network components that perform TCP packet re-assembly. This causes the link to perform in half-duplex mode, despite all physical parts that are in full-duplex mode.
- Test the device or drive performance by using Data Domain
 - Create a large data file on the storage node and use Data Domain to send it to the target device:

date; dd if=/tmp/5GBfile of=/dev/rmt/0cbn bs= 1MB; date
 - Note the time it takes for the file to transfer and compare it with the current tape performance.
 - If throughput is less when using NetWorker, device parameter tuning in NetWorker—Target Session, Read Block Size, etc.—may be required. Then test the backups.

References

<https://www.emc.com>

<https://powerlink.emc.com>

Glossary

Backup Server

NetWorker Server is the controlling backup entity that directs client backups and stores tracking and configuration information.

Client

NetWorker Client is the most fundamental host. NetWorker Client component is installed on all servers which needs to be backed up through NetWorker.

Cluster

A computer cluster is a group of linked computers, working together closely so that in many respects they form a single computer. The components of a cluster are commonly, but not always, connected to each other through fast local area networks. Clusters are usually deployed to improve performance and/or availability over that provided by a single computer, while typically being much more cost-effective than single computers of comparable speed or availability.

Deduplication

Data deduplication (often called "intelligent compression") is a method of reducing storage needs by eliminating redundant data. Redundant data is replaced with a pointer to the unique data copy. For example, a typical email system might contain 100 instances of the same one megabyte (MB) file attachment. If the email platform is backed up or archived, all 100 instances are saved, requiring 100 MB storage space. With data deduplication, only one instance of the attachment is actually stored; each subsequent instance is referenced back to the one saved copy. In this example, a 100 MB storage demand could be reduced to only one MB.

Disaster Recovery

Disaster recovery is the process, policies, and procedures related to preparing for recovery or continuation of technology infrastructure critical to an organization after a natural or human-induced disaster. Disaster recovery planning is a subset of a larger process known as business

continuity planning and should include planning for resumption of applications, data, hardware, communications (such as networking), and other IT infrastructure.

DMZ

In computer networks, a DMZ (demilitarized zone) is a computer host or small network inserted as a "neutral zone" between a company's private network and the outside public network. It prevents outside users from gaining direct access to a server that has company data. (The term comes from the geographic buffer zone that was set up between North Korea and South Korea following the UN "police action" in the early 1950s.) A DMZ is an optional and more secure approach to a firewall and effectively acts as a proxy server as well. In a typical DMZ configuration for a small company, a separate computer (or host, in network terms) receives requests from users within the private network for access to Web sites or other companies accessible on the public network. The DMZ host then initiates sessions for these requests on the public network. However, the DMZ host is not able to initiate a session back into the private network. It can only forward packets that have already been requested.

Firewall

A firewall is a set of related programs located at a network gateway server that protects the resources of a private network from users from other networks. (The term also implies the security policy that is used with the programs.) An enterprise with an intranet that allows its workers access to the wider Internet installs a firewall to prevent outsiders from accessing its own private data resources and for controlling what outside resources its own users have access to service and source ports.

LUN

In computer storage, a logical unit number (LUN) is simply the number assigned to a logical unit. A logical unit is a SCSI protocol entity, the only one which may be addressed by the actual input/output (I/O) operations. Each SCSI target provides one or more logical units, and does not perform I/O as itself, but only on behalf of a specific logical unit.

NAS

Network-attached storage (NAS) is file-level computer data storage connected to a computer network providing data access to heterogeneous network clients. A NAS unit is essentially a self-contained computer connected to a network, with the sole purpose of supplying file-based data storage services to other devices on the network. The operating system and other software

on the NAS unit provide the functionality of data storage, file systems, and access to files, as well as management of these functionalities.

NDMP

Network Data Management Protocol (NDMP) is meant to transport data between NAS devices, also known as filers, and backup devices. This removes the need for transporting the data through the backup server itself, thus enhancing speed and removing load from the backup server.

Recovery Point Objective

A Recovery Point Objective (RPO) is a point of consistency to which data must be restored. It is a measurement of time indicating how long a consistent point is expected to be compared to the time an incident occurred and can be from zero to minutes or hours. With synchronous data replication, RPO can be zero. For systems that don't need immediate recovery or where data can be rebuilt from other sources, RPO may be 24 hours or more.

Recovery Time Objective

Recovery Time Objective (RTO) is a measurement of the time permitted to recover an application to a consistent recovery point .This time can include some or all of the following;

- Time to bring up backup hardware
- Time to restore from backups
- Time to perform forward recovery on databases
- Time to provide data access

SAN

A storage area network (SAN) is an architecture designed to attach remote computer storage devices (such as disk arrays, tape libraries, and optical jukeboxes) to servers in such a way that the devices appear locally attached to the operating system. Although the cost and complexity of SANs are dropping, they are still uncommon outside larger enterprises.

Storage Node

The host that receives client-generated data, writes it on the backup device, generates the tracking information, and reads the data at the time of recovery. NetWorker Storage node component is installed on Backup Server itself.

Source :- www.Wikipedia.org , www.storagewiki.com

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.