EMC²
where information lives®

Best Practices for Implementing and Administering EMC NetWorker®

Anuj Sharma

EMC Proven Professional Knowledge Sharing 2009

EMC² PROVEN PROFESSIONAL

Anuj Sharma
Implementation Engineer
Ace Data Devices Pvt. Ltd.
anuj.sharma@ace-data.com

Table of Contents

Executive Summary 5

Abstract 8

Introduction 9

Essentials 11

- Backup Server 11
- Client 11
- Storage Node 11
- Cluster 11
- Disaster Recovery 11
- Deduplication 12
- DMZ 12
- Firewall 12
- LDAP 13
- LUN 13
- NAS 13
- NDMP 13
- Recovery Point Objective 14
- Recovery Time Objective 14
- SAN 14
- SLA 14

Section 1 : Pre Implementation Phase 15

Analyzing the Backup Infrastructure 16

Categorizing Data 24

RPO and RTO Requirements 24

Backup Schedules and Policies 25

Section II: Implementation Phase 26

Data Deduplication Using EMC Avamar 27

Persistent Binding 30

NDMP Backups 34

Cluster Client Backups 35

Probe Based Backups 36

Email Alerts 37

LDAP Integration 38

Section III: Post Implementation / Administration Phases 40

Troubleshooting 43

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

LIST OF FIGURES

S.NO	Figure	Page
1.	Data Loss Causes	4
2.	Backup Solution Characteristics	5
3.	Deployment Phases	9
4.	Pre Implementation Phases	15
5.	Backups across bi-directional firewall	19
6.	Avamar Integration With NetWorker	27
7.	SAN Backup Environment	30
8.	Backup Server Goes Offline	38

Executive Summary

Pick up any newspaper and you will see an article on information loss. The June 12, 2008 USA Today headline read: "Lost digital data costs businesses billions." This is an example of that reinforces the problem and the implications. Data has become a critical asset; organizations must be able to deal with natural disasters, government compliance, database corruption, component failure, human error etc.

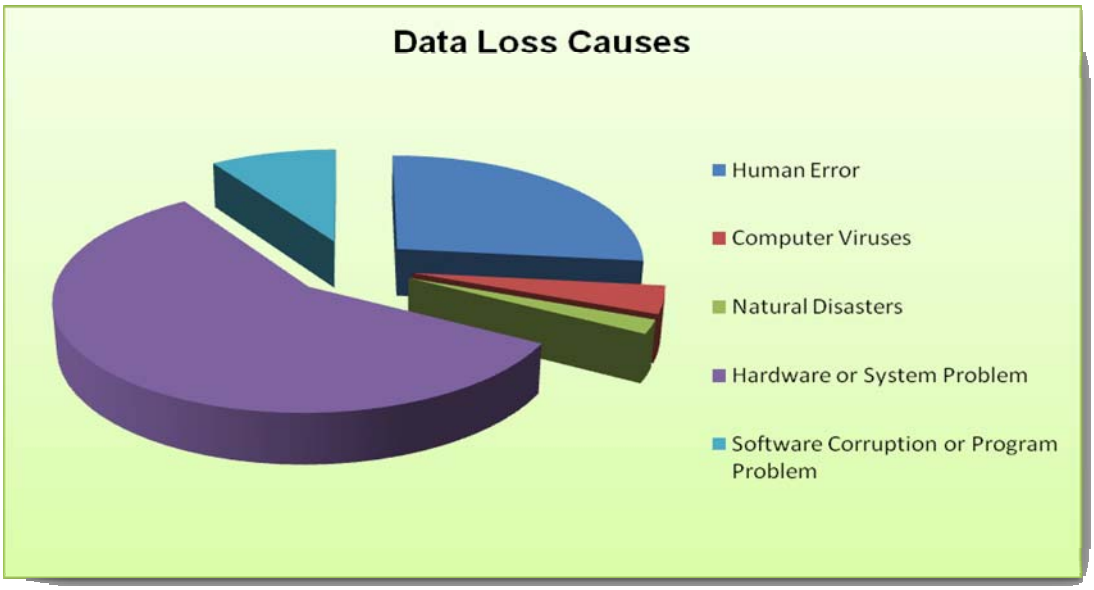

Figure 1 Causes of Data Loss

Elevated demand for increased application availability confirms the need to ensure business continuity practices are consistent with business needs. Interruptions are classified as either planned or unplanned. Failure to address these specific outage categories compromises a company's ability to meet business goals. Fortunately, the most devastating events rarely happen, but when they do the potential outage could put any business in financial jeopardy.

Businesses generate and maintain vast amounts of data that may include names of customers, partners, inventory, and pricing of products and services. For example, a bank has to accurately and securely maintain account information about several million customers. Businesses create information from the data they collect. Traditionally, businesses stored data because they had to. We all expect our banks to accurately reflect our current balances. What will happen if all that data is lost? Banks won't be able to justify our statements or balance

sheets. In short, there will be total chaos. Let's consider New York's World Trade Center on September 11, 2001. Many companies did not back up their data to a remote site. As a result, they never resumed operation. Data Backup and Recovery Solutions is an enterprises' most critical asset. Enterprises should consider the following questions before choosing a Backup and Recovery Solution →

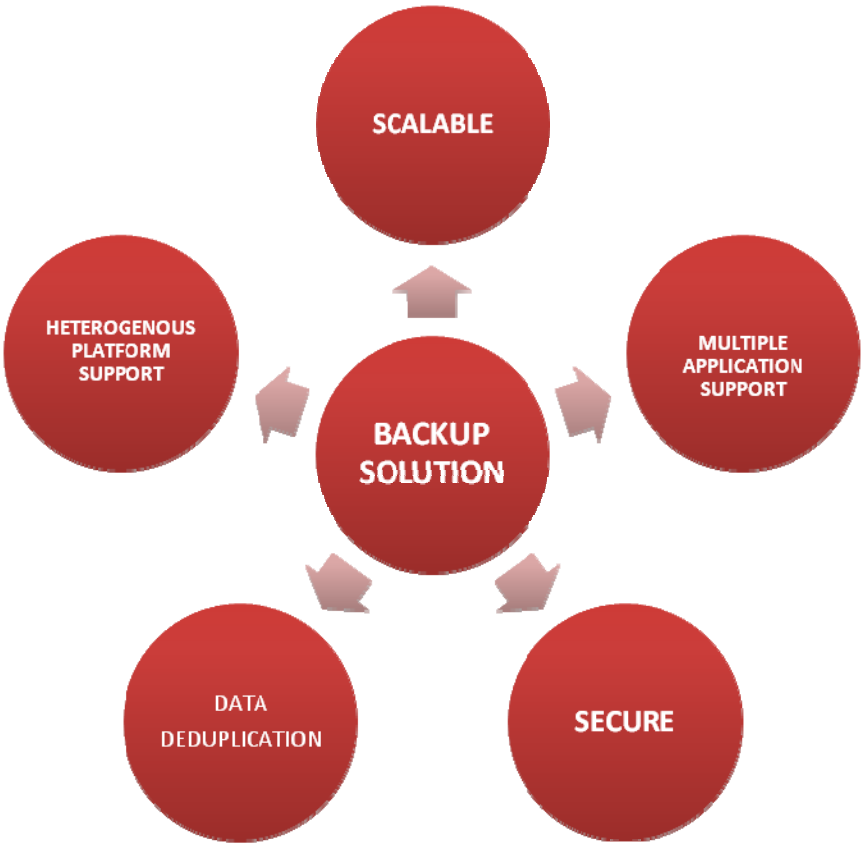

Figure 2 Backup Solution Characteristics

- **Is the solution Secure, Scalable, Efficient and Reliable?**
- **Does the solution permit Granular Recovery?**
- **Does the solution support heterogeneous platforms?**
- **What applications does the solution support?**
- **Does the solution have de duplication?**
- **Does the solution support an open architecture?**

EMC NetWorker is the answer to all these questions. However, only having EMC NetWorker as the Backup and Recovery Solution doesn't guarantee that an enterprise is immune from data loss. Feedback from EMC's customers suggests that they can recover data but still experience data and time loss.

Recovery time objectives (RTOs) are difficult to meet when IT managers encounter unreadable tapes. IT executives report that they have no formal remote-office backup procedures in place. Respondents have reported a direct impact to their business from unrecoverable data and SMB's are pressed when it comes to recovery times. Recognizing the increasing value of information, 76% of those surveyed reported that unrecoverable data had a direct impact on their business.

We need to give special attention to the implementation and administration of EMC NetWorker to ensure that backed up data can be recovered without difficulty. This article covers the post-purchase implementation and administration practices that make enterprises more secure in the event of data loss.

Abstract

EMC NetWorker is the fastest performing backup application in the market. Integration with replication and snapshot technologies helps you meet the most aggressive RTO and RPO requirements and transform backup to disk or backup to tape in an off-host, off-hours process. It supports a broad set of OS, databases, applications and topologies.

EMC NetWorker's compatibility with various operating systems, applications, and databases is the root cause of its success. However, it should be implemented and administered properly to ensure the ease of backup and recovery. There are some practices that the implementation specialist can keep in mind while implementing the product, and some key points that the NetWorker administrator should keep in mind while administering the product. These practices will help to meet the RTO and RPOs set by the enterprise.

This article details practices that I performed when implementing and administering EMC NetWorker.

- **To be fool proof in case of EMC NetWorker Server Disaster**
- **Implementing persistent binding through EMC NetWorker**
- **Integrating EMC Avamar for deduplication**
- **Implementing EMC NetWorker in case of a bidirectional as well as unidirectional hardware firewall including various scenarios for e.g. when some of the clients are in DMZ**
- **Working with the EMC NetWorker ports**
- **Implementing EMC NetWorker in a cluster**
- **Integrating modules with EMC NetWorker**
- **Integrating Email alerts with EMC NetWorker**
- **Practices that should be performed after the Implementation of EMC NetWorker i.e. while administering EMC NetWorker**
- **Implementing EMC NetWorker on heterogeneous platforms**
- **Probe based backups**
- **Backup and Recovery Drills**
- **Backup Infrastructure Audits**

Introduction

This article describes the various practices for deploying a smooth backup infrastructure using EMC NetWorker software (i.e top of the line Software for deployment, monitoring, administering Backup Infrastructure in an enterprise). I have divided the deployment of a smooth, reliable and efficient backup infrastructure to meet the SLA's, RPO's and RTO's set by an enterprise into phases:

- Pre Implementation Phase
- Implementation Phase
- Post Implementation Phase

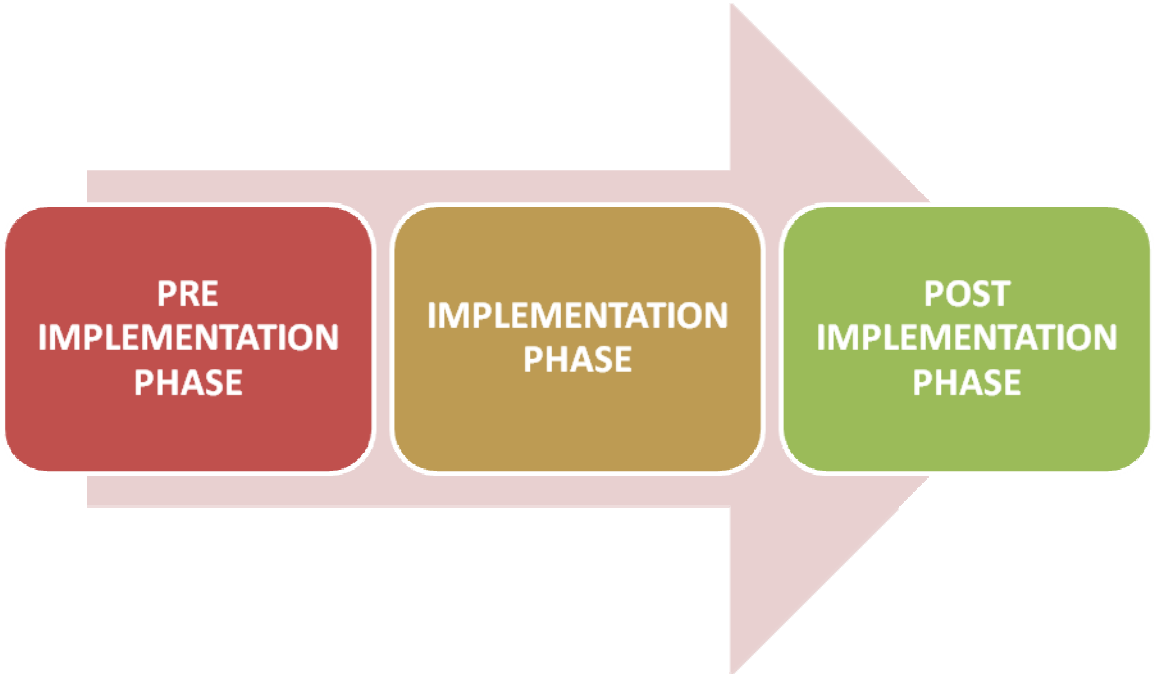

Figure 4 Backup Solution Deployment Phases

These phases will be further divided into sub phases. Carry out all the phases in systematically to ensure smooth functioning of the backup infrastructure and minimal backup infrastructure downtime in the event of a Backup Server Disaster.

This article will also cover the new features included in 7.4 onwards version of EMC NetWorker that include enhanced reporting, LDAP integration, probe based backups, persistent binding etc . It will review the concept of Deduplication and integrating EMC Avamar with EMC NetWorker, and will describe each phase along with the sub- phases. I've included basic terminology required to better understand the article in the Essentials sub topic.

Essentials

Backup Server

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

Client

NetWorker Client is the most fundamental host. The NetWorker Client component is installed on all servers that need to be backed up through EMC NetWorker.

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. The NetWorker Storage node component is installed on backup server itself.

Cluster

A computer cluster is a group of linked computers that work 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 being more cost-effective than single computers of comparable speed or availability.

Disaster Recovery

Disaster recovery encompasses the process, policies and procedures related to preparing for recovery or continuing a 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.

Deduplication

Data deduplication (often called "intelligent compression") is a method to reduce storage 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 of storage space. With data deduplication, only one instance of the attachment is 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.

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. A DMZ is an optional, more secure approach to a firewall and effectively acts as a proxy server. 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 other networks' users. (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 can access.

LDAP

The Lightweight Directory Access Protocol (LDAP) is an application protocol for querying and modifying directory services running over TCP/IP. A directory is a set of objects with similar attributes organised in a logical and hierarchical manner. The most common example is the telephone directory that consists of a series of names (either of persons or organizations) organized alphabetically, with each name having an address and phone number. An LDAP directory tree often reflects various political, geographic, and/or organizational boundaries, depending on the model chosen. LDAP deployments today tend to use Domain name system (DNS) names for structuring the topmost levels of the hierarchy. Deeper inside the directory entries representing people, organizational units, printers, documents, groups of people or anything else that represents a given tree entry (or multiple entries) might appear.

LUN

In computer storage, a logical unit number (LUN) is the number assigned to a logical unit. A logical unit is a SCSI protocol entity, the only one that 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, access to files, and the management of these functionalities.

NDMP

Network Data Management Protocol (NDMP) is a protocol that transports data between NAS devices, also known as filers and backup devices. This removes the need to transport the data through the backup server itself, 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. It can range 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 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 to attach remote computer storage devices (such as disk arrays, tape libraries, and optical jukeboxes) to servers so that the devices appear to be locally attached to the operating system. Although the cost and complexity of SANs are dropping, they are still uncommon outside larger enterprises.

SLA

A service level agreement (SLA) formally defines a service contract. In practice, the term *SLA* is sometimes used to refer to the contracted delivery time (of the service) or performance.

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

Section 1 : Pre Implementation Phase

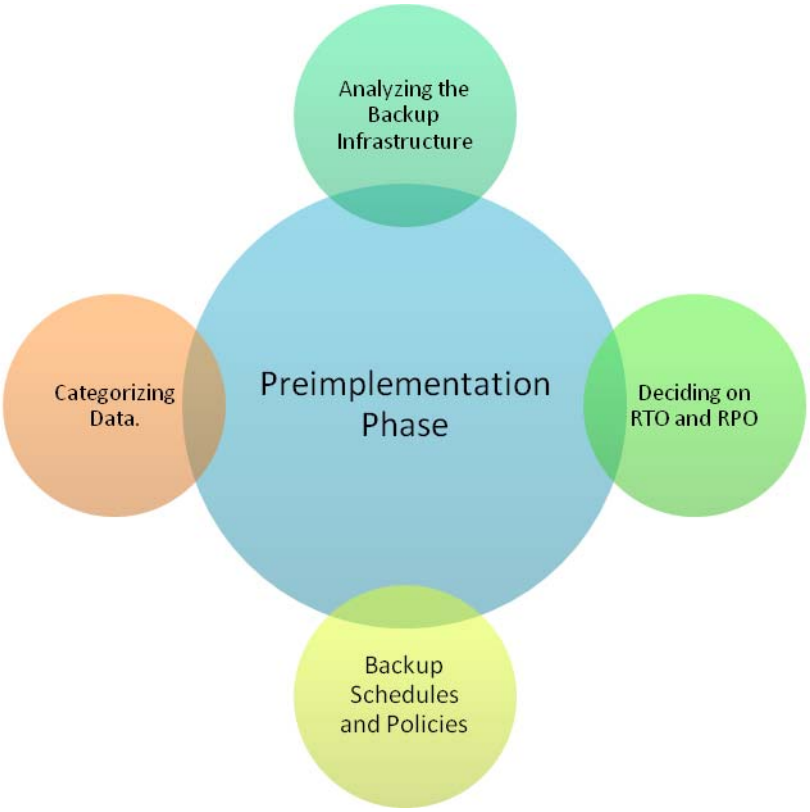

Figure 4 Pre Implementation Phase

Analyzing the Backup Infrastructure

1(a) Connectivity, Bandwidth , Hardware

Backup speed depends on:

- **Connectivity Between the backup server and backup clients**
- **Connectivity between the backup server and the backup device i.e. FC , ISCSI**
- **Network Traffic**
- **Existence of SAN between the backup server, clients and backup device**
- **Backup Device whether Disk Based backup device, tape based backup device ie library , LTO 3 , LTO 4 tape technology , etc.**
- **Backup Server and Clients Hardware Configuration**
- **Server Parallelism attribute on backup server; optimize Client Parallelism attribute on backup clients for better performance**

In simple language, we need to calibrate all the factors above to achieve required backup performance. For example, suppose we have an LTO 4 tape drive that can write data at the speed of 120 MB/s, but we don't have the network to provide the LTO 4 tape drive with the data to write at that speed. The resources are not optimally utilized; in this case the LTO 4 tape drive will not serve the purpose. In addition, we should consider the network.

Consider another case of Physical Tape Library (PTL) and Virtual Tape Library (VTL). Having a VTL in place of PTL won't make backups faster. The network should also be there to provide the data at faster transfer rates to the VTL.

We should collect the information about the backup server, backup clients, applications running on the backup clients, and data size in the pre implementation phase. The existing network infrastructure needs to be carefully examined (i.e. information about firewalls and connectivity between the backup server and clients). It is very important to have the accurate knowledge of the backup infrastructure. At one of our clients, the NDMP backups were failing after the configuration. In the end, we found that we were using a Control Station whereas the data mover IP is required for NDMP backups.

We should make the request for the necessary ports to be opened on the firewall to the concerned team for smooth implementation of the backup solution. If possible, we should use a dedicated backup network so that the existing network infrastructure is not impacted. Special attention should be given to the available bandwidth available for backups to take place in case of a shared network. The deployment of the backup solution on the existing network infrastructure should not choke the network creating chaos. Any implications to the customer should be communicated at this phase to avoid future conflicts.

With one client, we had to deploy EMC NetWorker as the backup solution. When we started implementing it, the clients were unable to communicate with the backup server. We later found later a firewall exists and the backup server is in DMZ. This was never mentioned earlier in the design document. This is another reason you should have a clear image of the organizations' networking infrastructure. When you are working on a network consisting of a firewall between the backup server and clients, consider the following aspects to determine the ports that need to be opened on the firewall for barrier free communication between the backup clients and backup server, or when the backup server is in DMZ:

- Backup Server performs backups and recoveries using a number of TCP ports (service and source ports). Two of the TCP ports are reserved by the NetWorker host and are used as follows:

  ➔Port 7937 as a service port for the nsrexecd daemon

  ➔Port 7938 as a connection port for the EMC NetWorker portmapper

- A NetWorker 7.3 or later client uses nsrexecd that requires four service ports: the reserved ports 7937 and 7938 and two user-configurable ports from the service port range.

- A NetWorker storage node (SN) is also a NetWorker client; it uses all of the ports for a client. In addition to the four ports for a client, a storage node requires ports for nsrmmd and nsrlcpd daemons. There is one nsrmmd per tape or file device on the machine to handle backup and recover data. An advanced file type device counts as two devices since it creates a read-only device for simultaneous restores, and thus has two nsrmmd. When spanning from one device to another, a helper nsrmmd is launched to mount the new tape. The helper nsrmmd also requires a port. There can be up to two

nsrmmd per device on a system. There is one nsrlcpd per robot in an autochanger. A storage node requires a minimum of: **4 + (2 * #devices) + (#jukeboxes)** service ports.

- A NetWorker server is also a NetWorker storage node; and uses all of the ports for a storage node. In addition to the ports for a storage node, a server requires ports for nsrd, nsmmdbd, nsrindexd, nsrmmgd, and nsrjobd daemons. A NetWorker 7.3.x server requires a minimum of: **11 + (2 * #devices) + (#jukeboxes)** service ports.

- NetWorker 7.4 introduces a new daemon, the client push daemon, which also consumes a TCP service port. As a result, a NetWorker 7.4 server requires a minimum of: **12 + (2 * #devices) + (#jukeboxes)** service ports

- The Console server component of NMC uses 3 ports:

  - ➔ One port (9000 by default) is used for the web server. It provides a way to download the java application code that acts as the Console front end. This port is selected during the installation process.
  - ➔ The second port (9001 by default) is used for RPC calls from the Console Java client to the Console server. This port is selected during the installation process.
  - ➔ The last port (2638 by default) is used for database queries.

- We always recommend that you open all the ports from range 10001 to 30000 from the firewall for fast and smooth backups

- The EMC Avamar integration with NetWorker uses port 27000 (or 29000 if secure sockets layer – ssl is used)

- NDMP uses port 10000

- EMC AlphaStor uses ports 44444, 41025, 41114, 44460, and 44455

Let's consider a case for **calculating service ports on a bidirectional firewall**.

This example shows how to apply the basic rules for a sample network with NetWorker clients Alpha, Charlie, Bravo, NetWorker storage nodes A and B, and a NetWorker server with a single firewall that blocks both ways. Each storage node and the NetWorker server have 2 tape libraries and 4 drives, and there are no pre-NW 7.3 clients. The hosts table:

192.167.10.101 Alpha
192.167.10.102 Bravo
192.167.10.103 Charlie
...
196.167.10.124 SN A
192.167.10.125 SN B
192.167.10.126 NetWorker Server

The firewall in the figure below is bidirectional; it blocks traffic both ways. Suppose the NetWorker server has eight devices in two libraries. It needs **11 + 2 * (num devices) + (num libraries) = 11+ 2* (8) + 2= 29** service ports.

Two ports must be 7937 and 7938 and preferably select ports 7937–7966. A NetWorker 7.4 server requires one additional port to accommodate the client push daemon. The firewall must allow traffic to the NetWorker server's IP address on all the service ports configured. The firewall rule for the service ports:

TCP, Service, src 192.167.10., dest 192.167.10.126, ports 7937-7966, action accept*

There are NetWorker storage nodes on the right of the firewall. Storage node A has 8 devices and two libraries. It needs **4 + 2 * (num devices) + (num libraries) = 22** service ports.
Storage node B is identical, and needs the same number of ports. It can use the same port range as well, 7937–7959. Each NetWorker SN must be configured to use 22 service ports 7939–7959, and the firewall must allow to each SN's IP address on all the service ports.

Figure 5 Backups Across Bi Directional Firewall

The Firewall Rule will be
TCP, Service, src 192.167.10.*, dest 192.167.10.124, ports 7937-7959, action accept
TCP, Service, src 192.167.10.*, dest 192.167.10.125, ports 7937-7959, action accept

Clients Alpha, Bravo, Charlie will have the requirements of 4 service ports. Configure each client to use at least four service ports, 7937–7940, and configure the firewall to allow traffic rightward to each client's IP address on all of the service ports configured. The firewall rule for the service ports would be:

TCP, Service, src 192.167.10., dest 192.167.10.101, ports 7937-7940, action accept*
TCP, Service, src 192.167.10.*, dest 192.167.10.102, ports 7937-7940, action accept
TCP, Service, src 192.167.10.*, dest 192.167.10.103, ports 7937-7940, action accept

In the previous pseudo syntax, the firewall is configured to allow incoming service connections to the NetWorker server's IP address on ports 7937–7966, from the IP addresses of each of the storage nodes or client machines (as well as any other machines on that subnet). The firewall is also configured to allow connections to the IP addresses for each storage node on ports 7937–7959, and to each client IP address on ports 7937–7940. Similarly, the calculations can be carried out for a unidirectional firewall and you can set rules accordingly.

You should also request the networking team to configure NICs for the backup network on Full Duplex or Half Duplex, and 100 or 1000 MBps depending on the NIC used. It should not be set to auto negotiate. Do not use Full/Half Duplex and 10/100/1000 Mbps type options. Many network issues are known to be resolved by this.

In case of a Software firewall, make sure that the NetWorker daemon nsrexecd is in the exceptions list of the firewall; and in case of a hardware firewall NAT should be disabled.

It is very important to check the consistency of Tape Library Drivers on all the dedicated storage nodes and backup server In a Mixed Environment, i.e. LAN and SAN based backup. You must check that consistency is maintained and the drivers are current. At one of our clients, we faced some issues with data recovery. We were not able to recover a client's data and we checked a directory recovery from a different backup client and it was successful. We later found that the tape library drivers were not consistent and caused the unsuccessful recovery.

1(b) Zoning

Due to complexities in multi-hosting tape devices on SANs, use zoning tools to help keep the backup/restore environment simple and less susceptible to the effects of changing or problematic SANs. Zoning provides a way for servers, disk arrays, and tape controllers to see only the hosts and targets they need to see and use.

Among the benefits of zoning:

- The potential to greatly reduce target and LUN shifting
- Limiting unnecessary discoveries on the FC interface controllers
- Reducing stress on backup devices by polling agents
- Reducing the time it takes to debug and resolve anomalies in the backup/restore environment
- Reducing the potential for conflict with untested third-party products

Zoning may not always be required for small or simple configurations. The bigger the SAN, the more zoning is needed.

• Small fabric (16 ports or less)—may not need zoning.

If no zoning is used, tape controllers should reside in the lowest ports of the switch.

• Small to medium fabric (16 - 128 ports)—use host-centric zoning.

Host-centric zoning is implemented by creating a specific zone for each server or host, and adding only those storage elements that the host will utilize. Host-centric zoning prevents a server from detecting any other devices on the SAN including other servers; it simplifies the device discovery process.

Large fabric (128 ports or more)—use host-centric zoning and split disk and tape targets.

Splitting disk and tape targets from the same zone will free the tape controllers from discovering disk controllers that it doesn't need to see, unless extended copy is required. Dedicate HBAs for disk and tape, where practical, for optimal performance.

You can use this checklist with the SAN Administrator to verify that zoning is appropriate:

Checklist

To ensure that all components on the SAN are logged in and configured properly, you must be able to answer YES to each of the following questions so that the library doesn't misbehave after the implementation:

☑ Are all hardware components at the minimum supported firmware revision (HBA, Fibre Channel switch, interconnects, tape library drives, and tape library robot)?

☑ Is the minimum patch level support for the operating system installed?

☑ Are the minimum supported drivers installed (HBA, tape drives)?

☑ Is the tape library online?

☑ Are the Fibre Channel ports of the tape library correctly logged into the Fibre Channel switch?

☑ Is the host server correctly logged into the Fibre Channel switch?

☑ If the Fibre Channel switches are cascaded or meshed, are all Interswitch Link (ISL) ports correctly logged in?

☑ Are the Fibre Channel ports of the tape library and the host server HBA in the same switch zone (by World Wide Name (WWN) or port)?

☑ Does the host server detect all of the tape and robotic devices intended to be used?

Categorizing Data

Categorize backup clients based on the criticality of data residing on the clients. There could be data that is not useful for the organization on some servers, for example mp3, avi, and jpg files. These files should be excluded from the data that requires backup. This reduces the backup window. We should categorize data into what needs to be backed up and what does not need to be backed up. The data that needs to be backed should be further divided into categories. In some enterprises, application database data is more important than file system data. The backup policies and schedules should be decided accordingly.

How much critical data does the enterprise generate every day? If you know the size and expected performance, it will help you when deciding on the schedules.

Maintain a list of backup clients according to the criticality of data residing on the servers and the approximate size of the data that requires backup. This will help to decide on schedules, browse, and retention policies. Consider the future data growth trends to optimally utilize resources.

RPO and RTO Requirements

You must work with the appropriate person in the enterprise to identify RPO and RTO requirements. You should ask questions like, "What frequency of recovery requests does the enterprise expect? And how long do you think they can afford to wait for a recovery request to happen?" Document the customers' expectations regarding RPO and RTO to ensure a mutual understanding. It is also very important to include the time for shipment of tapes when they are shipped offsite.

This will help you to decide the frequency of backups of clients, schedules, backup levels, and browse and retention policies. You can decide the backup device and its connectivity including storage and dedicated storage nodes.

Backup Schedules and Policies

Work done in the previous two steps will help to decide backup schedules and policies. The frequency of backup cycles for critical servers can be set higher than other servers. You should take full backups of more critical servers more frequently so that the recovery is faster. In case of critical database backups, such as in an SAP Production Environment, you can schedule archive backups more than once a day so that there can be point in time restores. At one of our clients, we backup archive logs every six hours so there will be minimal data loss on the event of a disaster. Normally, you can schedule weekly full and daily incremental schedules.

You can also use the data growth trend to choose the schedules. You can schedule frequent full backups as the client with frequently changing data, and incremental backups on other clients.

Section II: Implementation Phase

Develop an implementation plan before you begin to document EMC NetWorker's implementation. Make sure that the environment meets the pre-requisites for EMC NetWorker (Hardware Configuration, Network Configuration etc).

Follow these general practices:

- Update hosts file of the backup server with IPs, Hostnames, and FQDN's of backup clients before starting the implementation.

- Update the servers file (in the /nsr/res folder of backup clients) with the backup servers' IP and Hostname. If the file is left blank, every backup server on the network will get authorized to take a backup of the client by default. For security purposes, and to prevent unauthorized access, it should be regular practice to enter the name of authorized backup server in the servers file at the backup clients /nsr/res folder.

- You should 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, and it 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.

- Check the connectivity between the backup client and the backup server by using save command; for example, Save -s <backup server> <save set>

- When configuring the Exchange backup, specifying different Information Stores in the save set will result in faster backup and will produce more save streams. This practice has helped us to achieve better backup throughput for exchange backups.

- It is always better to divide a single saveset into many to improve backup speeds: such as splitting a single saveset D:\ into various savesets like d:\data, d:\important.

Data Deduplication Using EMC Avamar

This method uses a new backup agent at the host. Post NetWorker 7.4.x version has deduplication with the integration of EMC Avamar. EMC's Avamar is an example of a source-side deduplication product. The agent on each host connects to a centralized consolidated store of previously backed-up data. The agent makes comparisons of what's on the local server and what's on the backed-up store. It then sends only the unique blocks across the network. While this works well in remote office backup situations, it does not scale well in the data center, especially if there is a moderate-sized server count of more than 20. Having that many servers doing redundant data analysis consumes time and server resources. In some cases, if the change rate is high -- such as in a database -- the amount of additional resources required by the deduplication process can render the application itself unusable while it is being scanned for duplicate data.

In addition to factors such as data type and source, deduplication rates will vary depending on change rate, length and retention period. From a data perspective, there are two processes common to deduplication: backup and archiving. The first full backup will generate some level of deduplication as redundancy is identified across files, volumes and servers within the enterprise. Subsequent incremental backup jobs will typically capture efficiencies of 6X or 7X. Most of the data in an incremental backup consists of either new or modified documents, or updated database or email stores. Even if the documents are new, a comparison can be made to similar files for redundant patterns.

The data segments that represent the modified files will be compared to the original copy's data segments, and only the changed segments need to be stored. Because they tend to be very large files, databases will gain the most. For example, a 200 GB Oracle database that only had a 1% change during the course of the day will require storing 2 GB of new data rather than the entire 200 GB that would be stored without deduplication. Subsequent full backups will see a 50X to 60X reduction in stored data. This is because, as a percentage, there is not much changed data between two full backups.

In the case of deduplication, a high percentage of changes were captured during the incremental jobs. From a storage perspective, subsequent full backups require no more space than the prior incremental. A NetWorker de-duplication node is an EMC Avamar server that stores de-duplicated backup data. The initial backup to a de-duplication node should be a full backup. During subsequent backups, the Avamar infrastructure identifies redundant data segments at the source and backs up only unique segments and not entire files. This reduces the time required to perform backups, as well as both the network bandwidth and storage space used for backups. The Avamar server must be configured as a NetWorker de-duplication node, and the Avamar server must be available when a de-duplication client resource is created.

The Avamar server receives backup data from NetWorker de-duplication clients. You can configure NetWorker clients to take advantage of EMC Avamar de-duplication technology to dramatically decrease the amount of time, network bandwidth, and disk capacity required to back up client data.

.

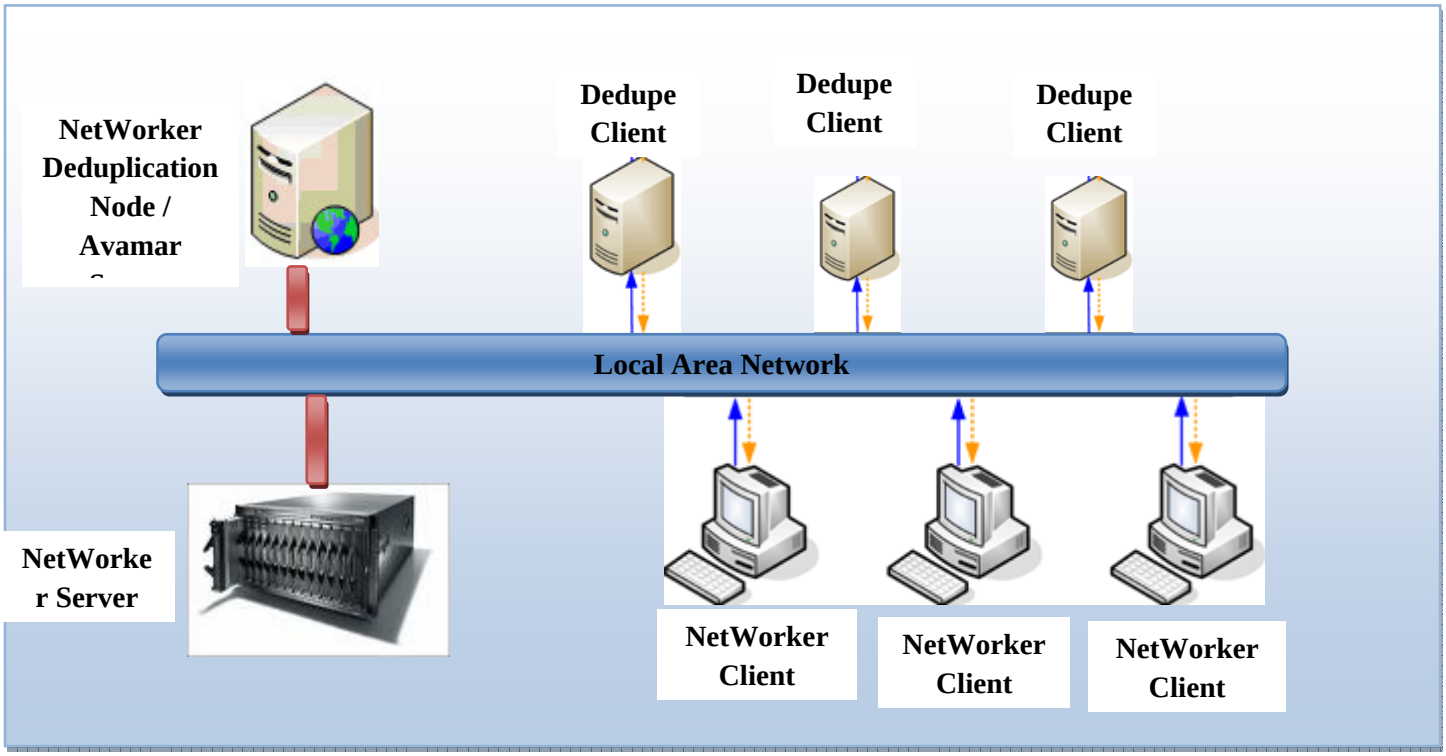

FIGURE 6 Avamar Integration with NetWorker

All typical client operations such as adding, deleting, and drag-and-drop, work the same for de-duplication clients as they do for other types of clients

A de-duplication save set is treated as a regular NetWorker save set. It has two parts: metadata (hash info) and the backed-up client data. Only the metadata is stored on a NetWorker storage node.

Creating a de-duplication node

Deduplication nodes exist on Avamar servers. You create access to them from the NetWorker side once EMC Avamar's implementation. EMC Avamar reduces the Backup window significantly since data volume is reduced. EMC NetWorker integration with EMC Avamar plays a vital role in organizations with large amounts of redundant data. The customer will see the effect of deduplication on network bandwidth and required disk capacity.

To create a NetWorker de-duplication node:

- Click the Devices button in the NetWorker server's Administration interface.
- Right-click De-duplication Nodes in the navigation tree, and select New.
- Type the fully qualified domain name or short name of the de-duplication node (an Avamar server) in the Name field.
- Enter the credentials for the Avamar server that is the de-duplication node.

Configuring a De-Duplication Client

After creating a de-duplication node, configure a NetWorker de-duplication client as follows:

1. Create a new client as a normal client.
2. In the Backup area:

   – Ensure that you have assigned the de-duplication client to a group that contains only de-duplication clients.

   -You backup schedule for de-duplication clients depends on whether your priority is faster backups or faster, less complicated recoveries.

3. In Apps & Modules tab:

a. Select the De-duplication backup attribute in to enable this for de-duplication backups.

b. Select the name of the de-duplication node to which this client's backup data will be sent. If the de-duplication node for this client's backup data changes, the next backup done must be at level 0 (full) backup.

4. Complete the configuration as for a normal NetWorker client.

Persistent Binding

Persistent binding is a mechanism to create a continuous route from a storage device object in a host to a volume storage array across the fabric within a SAN environment. Without persistent binding, the host cannot maintain persistent logical routing of the communication from a storage device object across the fabric to a storage array volume. If the physical configuration of the switch is changed, ie. cables are being swapped or the host reboots, the logical route becomes inconsistent. Solaris Example →

The following example shows binding by WWPN.

fcp-bind-WWPN="500e800300000000:lpfc1t10"

This example binds the storage port defined by "500e800300000000" to Emulex HBA lpfc1 and assigns all the discovered LUNs on that storage port to target 10. Devices will appear to the host as: cxt10dz for example c5t10d12, where x represents the controller number assigned to the HBA card and z is the LUN number assigned to the host on the storage port.

At one of our customer sites, we faced many issues because persistent binding was not done for the library. Many times, the drive symbolic names changed for the corresponding element address on dedicated storage nodes in a SAN environment. The library was continuously misbehaving when the backups started. We decided to do persistent binding using EMC NetWorker and it worked wonders. Persistent binding can be achieved by binding to WWPN (world wide port name), WWNN (world wide node name), or DID (device ID). NetWorker 7.4.2 introduced a feature for configuring persistent binding using the NSR RAP resource, (NSR Jukebox, which is another name for tape library).

NetWorker supports persistent binding for two major platforms – Windows Server 2003 and Linux. Persistent binding can be done by using Command Line Interface and Graphical User Interface. On Windows 2003, persistence of symbolic names assigned to tape LUN's can be enabled manually by editing a Windows Registry. Symbolic name persistence means that tape devices will be assigned the same symbolic name across reboot cycles irrespective of the order in which the operating system actually discovers the device .The persistence registry key is :H_KEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Tape\Persistence Persistence =1 symbolic tape names are persistent Persistence=0 non-persistence .

After enabling persistence \\.\ Tape0 will become \\.\Tape2828929982 where 2828929982 can be any OS assigned number. On Linux /dev/nst1 becomes /dev/tape/by-id/HP_LTO4_4k82h97_-nst.

For persistent binding we use inquire with p switch i.e. inquire –p and jbconfig with p switch i.e. jbconfig –p for configuring library through CLI using persistent binding. To configure library using GUI perform the following steps:

- **Launch NMC and then right click libraries and choose Configure All Libraries or Scan for devices.**
- **Select the library type SCSI/NDMP.**
- **Select the Target Storage Node on which you are configuring the library.**
- **Select use persistent names as yes.**
- **Click finish.**

Figure 7 SAN BACKUP ENVIRONMENT

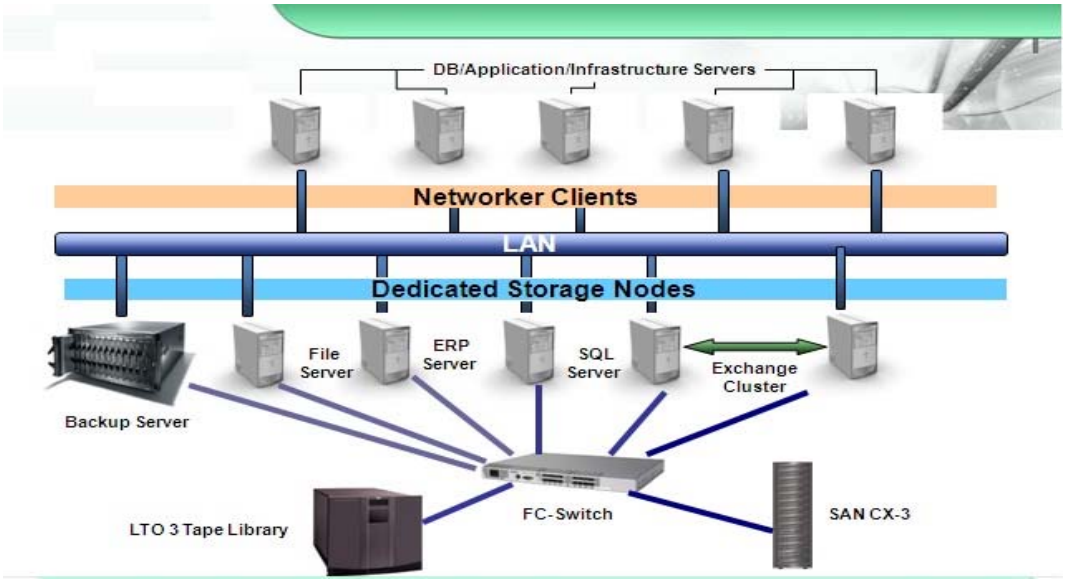

Please see next page.

```
C:\Documents and Settings\ceilbkpadm>jbconfig -p

Jbconfig is running on host cairn-44.cairnenergy.com (Windows Server 2003 5.2),
  and is using ace1.ace-data.com as the Networker server.

        1) Configure an AlphaStor Library.
        2) Configure an Autodetected SCSI Jukebox.
        3) Configure an Autodetected NDMP SCSI Jukebox.
        4) Configure an SJI Jukebox.
        5) Configure an STL Silo.
        6) Configure a Microsoft Removable Storage Jukebox.

What kind of Jukebox are you configuring? [1] 2
14484:jbconfig: Scanning SCSI buses; this may take a while ...
Incorrect function.

Installing 'Standard SCSI Jukebox' jukebox - scsidev03.5.1.

What name do you want to assign to this jukebox device? HP-MSL8096
Turn Networker auto-cleaning on (yes / no) [yes]? no
The drives in this jukebox cannot be auto-configured with the available
 information. You will need to provide the path for the drives.
Is (any path of) any drive intended for NDMP use? (yes / no) [no] no
Is any drive going to have more than one path defined? (yes / no) [no] yes

You will be prompted for multiple paths for each drive.
Pressing <Enter> on a null default advances to the next drive.


Please enter the device path information in one of the following formats:

\\.\Tape0 --for local path or
host:device-path --for remote node or NDMP device(s) or
host:drive-letter:directory path --for Windows disk file

Drive  1, element 1
Device path 1 ? \\.\Tape2147483646
Device path 2 ? ggndb1.ace-data.com:\\.\Tape5
Device path 3 ? ggnfs1.ace-data.com:\\.\Tape2147483641
Device path 4 ? ggnfs2.ace-data.com:\\.\Tape2147483643
Device path 5 ? ggnex1.ace-data.com:\\.\Tape2147483643
Device path 6 ? ggnex2.ace-data.com:\\.\Tape2147483641
Device path 7 ?

Drive  2, element 2
Device path 1 ? \\.\Tape2147483643
Device path 2 ? ggndb1.ace-data.com:\\.\Tape4
Device path 3 ? ggnfs1.ace-data.com:\\.\Tape2147483642
Device path 4 ? ggnfs2.ace-data.com:\\.\Tape2147483644
Device path 5 ? ggnex1.ace-data.com:\\.\Tape2147483644
Device path 6 ? ggnex2.ace-data.com:\\.\Tape2147483642
Drive  3, element 3
Device path 1 ? \\.\Tape2147483644
Device path 2 ? ggndb1.cairnenergy.com:\\.\Tape3
Device path 3 ? ggnfs1.cairnenergy.com:\\.\Tape2147483643
Device path 4 ? ggnfs2.cairnenergy.com:\\.\Tape2147483645
Device path 5 ? ggnex1.cairnenergy.com:\\.\Tape2147483645
Device path 6 ? ggnex2.cairnenergy.com:\\.\Tape2147483643
Device path 7 ?

Drive  4, element 4
Device path 1 ? \\.\Tape2147483645
Device path 2 ? ggndb1.cairnenergy.com:\\.\Tape2
Device path 3 ? ggnfs1.cairnenergy.com:\\.\Tape2147483644
Device path 4 ? ggnfs2.cairnenergy.com:\\.\Tape2147483646
Device path 5 ? ggnex1.cairnenergy.com:\\.\Tape2147483646
Device path 6 ? ggnex2.cairnenergy.com:\\.\Tape2147483644
Device path 7 ?


Please select the appropriate drive type number:
  1) 3480                   25) 9840C                 48) SAIT-1
  2) 3570                   26) 9940                  49) SAIT-2
  3) 3590                   27) 9940B                 50) SD3
  4) 3592                   28) adv_file              51) sdlt
  5) 4890                   29) dlt                   52) sdlt320
  6) 4mm                    30) dlt vs160             53) sdlt600
  7) 4mm 12GB               31) dlt-s4                54) SLR
  8) 4mm 20GB               32) dlt-v4                55) T10000
  9) 4mm 4GB                33) dlt1                  56) tkz90
 10) 4mm 8GB                34) dlt7000               57) travan10
 11) 4mm DAT160             35) dlt8000               58) TS1120
 12) 4mm DAT72              36) dst (NT)              59) TS1130
 13) 8mm                    37) dtf                   60) tz85
 14) 8mm 20GB               38) dtf2                  61) tz86
 15) 8mm 5GB                39) file                  62) tz87
 16) 8mm AIT                40) himt                  63) tz88
 17) 8mm AIT-2              41) logical               64) tz89
 18) 8mm AIT-3              42) LTO Ultrium           65) tz90
 19) 8mm AIT-4              43) LTO Ultrium-2         66) tzs20
 20) 8mm AIT-5              44) LTO Ultrium-3         67) VXA
 21) 8mm Mammoth-2          45) LTO Ultrium-4         68) VXA-172
 22) 9490                   46) optical               69) VXA-2
 23) 9840                   47) qic                   70) VXA-320

Enter the drive type of drive 1? 44
Are all the drives the same model? (yes / no) [yes] yes
14421:jbconfig:
A Dedicated Storage Node can backup only local data to its devices.
Should ggnex2.cairnenergy.com be configured as a Dedicated Storage Node? (yes /
no) [no] yes
14421:jbconfig:
A Dedicated Storage Node can backup only local data to its devices.
Should ggnex1.cairnenergy.com be configured as a Dedicated Storage Node? (yes /
no) [no] yes
14421:jbconfig:
A Dedicated Storage Node can backup only local data to its devices.
Should ggnfs2.cairnenergy.com be configured as a Dedicated Storage Node? (yes /
no) [no] yes
14421:jbconfig:
A Dedicated Storage Node can backup only local data to its devices.
Should ggnfs1.cairnenergy.com be configured as a Dedicated Storage Node? (yes /
no) [no] yes
14421:jbconfig:
A Dedicated Storage Node can backup only local data to its devices.
Should ggndb1.cairnenergy.com be configured as a Dedicated Storage Node? (yes /
no) [no] yes


Jukebox has been added successfully
```

The above pictogram is a jbconfig –p output for configuring persistently bound jukebox for corresponding element numbers, and configuring dedicated storage nodes at the same time.

To add a new persistent binding configured tape drive , that is , \\.\Tape21577383

Jbedit –s server –j ATL@2.0.0 –a –f \\.\Tape21577383

To delete a persistent binding configured tape drive, that is, \\.\Tape21577383

Jbedit –s server –j ATL@2.0.0 –d –f \\.\Tape21577383

Note: It is very important to set the block sizes to the vendor recommended specification as this will make the backups run faster.

NDMP Backups

Using the NDMP protocol to backup a NAS appliance can achieve improved backup performance. Some people use NDMP protocol and backup to their device connected directly to the backup server. This sounds good for NAS data, but better NDMP performance can be achieved if the backup device is directly connected to the NAS box. A NDMP client license is required for NDMP Based backups. For NDMP Based backups →

1. Create a Client Resource with a hostname of Data Mover. (We should not use hostname and IP of Celerra control station in case of Celerra).
2. Log into Celerra to find the Data Mover's Hostname and IP. Note the Data Mover hostname. Also note the IP of theData Mover, which can be found by clicking Network in the Data Mover section.
3. Update etc host file on Backup server.
4. Assign savesets. Eg. /vol, /vol0 etc.
5. Put backup command as : nsrndmp_save –M –T dump
6. Enable NDMP in Client Resource.
7. Enter a valid NDMP username and password (account must be valid NDMP user account on Celerra) in the Remote username and password section
8. In Application Attributes, enter DIRECT=y HIST=y UPDATE=Y
9. Set the parallelism to 1.
10. Create a separate pool for NDMP backup; label a tape for this pool. This is required because NDMP and non-NDMP data cannot reside on same tape.

Cluster Client Backups

NetWorker software backs up only the shared disk data that exists on the virtual server. Virtual servers in the cluster become NetWorker virtual clients.

To back up virtual client data:

- Create a client.
- Configure a Client resource for each cluster client in the cluster.
- For the **Save Set** attribute for the virtual client:
  • Specify **All** to back up all of the shared and active physical drives on a client.
  • Specify the drive volume letter (for example:**G:**) to back up only the drive volumes of shared disks that a virtual client owns.
- Add each NetWorker client that corresponds to a physical node for the **Remote Access** attribute on the **Globals 2 of 2** tab. For example:
  SYSTEM@physical-client1 (root@physical-client1 for UNIX)
  SYSTEM@physical-client2(root@physical-client2 for UNIX)
- In case of any application specific module you should append <–c virtual client name> to the backup command.
- For example in case of SAP Cluster the backup command is
  nsrsapsv –f /nsr/resnsrsapsv.cfg –c <virtual client name>

  Enter <curphyhost> if there is a dedicated storage node.

If the virtual NetWorker server is listed in the \nsr\res\servers file, the physical nodes must also be listed there. A backup fails if a virtual NetWorker server is listed in the servers file and you create a savegroup and add a physical node that does not own the NetWorker Server resource. To avoid this problem, do one of the following:

Leave the servers file blank.

Note: If the servers file is blank, any NetWorker server can back up the client.

Ensure that if the virtual NetWorker server is added to the servers file, all physical nodes are also added to the list.

Probe Based Backups

Probe based backups are introduced in version 7.5. NetWorker server schedules probe-based backups are based on user-defined events for clients and NetWorker modules, not just on time. We have experienced issues where resource utilization of the backup client is very high; initiating backups can cause the server application to crash. In those cases, we have to choose the backup time very carefully. With the help of probe based backups, we can schedule event based backups; we can write a script that passes the return code 0 when the CPU utilization is below 50 % and then the backup starts automatically. To configure a probe based backup →

1. Click Configuration from the NetWorker Administration window.
2. Right-click probes, and select New. The Create NSR probe window opens.
3. Type the name of the probe in the Name: field.
4. Type the name and path of the probe script in the Command: field. Place the probe resource script in the same directory as the nsr binaries for each client referencing the probe. For example, /usr/sbin on Solaris and Linux. A user defined probe is any program that passes a return code. Return codes as interpreted by NetWorker:

   • Return code 0: backup is required
   • Return code 1: backup is not required

   Note: The probe script name must begin with save, or nsr.

To associate a probe resource with a Client resource:

1. Click Clients, and right-click the client in the Configuration screen of the NetWorker Administration window.
2. Select Properties, and the Properties window opens.
3. In the Apps & Modules tab, select the probe from the Probe resource name: list.

Email Alerts

You can configure Email alerts so that the various inbuilt reports are automatically generated and mailed to the administrator email accounts. This helps in better administration. Bootstrap report email should be configured so that the bootstrap saveset ID's are well known in case of a NetWorker Server Disaster.

To configure a custom notification for NetWorker 7.3 and later; and create an SMTP email notification:

1. Click Configuration from the Administration window.
2. Click Configuration in the server's Administration interface.
3. Right-click Notifications, select New.
4. For Name, type a unique name for this second custom notification, such as Second adv_full Notice.
5. For Event, clear or unselect all choices except adv_file.
6. For Priority, clear or unselect all choices except Critical, Emergency, and Alert.
7. In the Action attribute, for a single host, enter: smtpmail -s "subject" -h hostname user
8. In the Action attribute, if multiple hosts are required, enter each additional address with a space: smtpmail -s "subject" -h hostname user1 user2 user3...

   where -h hostname is the hostname of the SMTP server

   where -s subject is the subject text enclosed in double quotes

   where user is the email address of the user to whom the notification will be sent

   Note: If the email address contains special characters, be sure to place the entire address in double quotes.

9. Click OK.

LDAP Integration

LDAP integration has been introduced in NetWorker Version 7.5. LDAP authentication enables you to log in to the Console server with user names and passwords that are maintained on a centralized LDAP v3 compliant server such as a Microsoft Active Directory server. Console user privileges are controlled by mapping LDAP user roles or LDAP user names to Console user roles. There is no need to add user names and passwords on the Console server. However, one must still add LDAP user names to user groups on the NetWorker server to enable privileges on the NetWorker server.

To enable LDAP authentication:

1. Log in to the Console server as a user, such as the default administrator, who belongs to the Console Security Administrator role.

2. On each NetWorker server, add an external LDAP user to the NetWorker server administrator's user group. This step ensures that once LDAP is enabled, at least one user will be able manage the NetWorker server and to add additional NetWorker users as required. The LDAP user that you add should also belong to the LDAP user roles or LDAP user names.

a. Click the Enterprise button on the taskbar.

b. Highlight a host in the navigation tree, right-click NetWorker, and select Launch Application.

c. Click the Configuration button on the taskbar.

d. In the navigation tree, select User Groups.

e. In the User Groups list, right-click Administrators and select Properties.

f. Add the LDAP user to the User attribute. Use the following format to add the user:
User=LDAP_username, host=console_host

g. Click OK.

3. Select Configure Login Authentication to launch the Configure Login Authentication wizard from the Console Setup menu.

4. Select the External Repository radio button and click Next. The Manage Authentication Authorities panel appears.

5. Click Add; provide information about your authentication authority in the remaining fields.

7. Enter the LDAP user roles or LDAP user names that will be mapped to the Console Security Administrator role and click Finish.

8. Restart the Console. application.

You can now log in to the Console server using an LDAP user name and password that belongs to the LDAP role that was mapped earlier.

Note: The user name Administrator is not allowed even if it is defined In LDAP mode.

Section III: Post Implementation / Administration Phases

Follow these practices for smooth functioning of the EMC NetWorker backup solution and accelerate recovery in case of disaster →

1. Configure a separate media pool for Index & bootstrap backups. This makes disaster recovery quicker as you need to look for only one media to locate the latest bootstrap. Clone these tapes regularly for additional protection.

2. For faster disaster recovery, make a note of the bootstrap Saveset ID on a daily basis. You don't need to spend time on a scanner to locate this.

3. Maintain similar nomenclature for a set of resources. If you want to name a Group "SQL", name its media pool "SQL". It is easier to correlate while troubleshooting.

4. Create separate media pools for critical databases.

5. Run the nsrck –L6 command at regular intervals to check the consistency of all the indexes.

6. Conduct backup and recovery drills after a fixed interval of time so that any inconsistency in the setup is known and addressed well in advance.

7. Carry out audits so that any enhancements required for smooth functioning of the Backup Solution can be carried out before a problem arises.

8. Sample Audit Reports are enclosed in Appendix A for your reference.

9. Many believe that once the backup has been triggered, even if the backup server link is lost the backups take place on the configured dedicated storage node. This is not the case; the backup server must remain online for successful backups. The reason is explained below.

Even though file system data is backed up onto a storage node, its tracking information is sent to the server by both the storage node and the client where server databases will be updated accordingly. If connectivity fails in between, it will fail to contact server and update the databases, hence the backup will be failed.

Will the backup start automatically from the point of failure when the link is up after some time? It depends on the " Client Retries " & "Timeout " parameters that are configured under group properties. When connectivity breaks, the client will retry as many times as the value set. If connectivity is restored within the retrial period, backup will be restarted; if not it will fail. You can manually restart the backupso that it will backup only failed savesets.

For example, a client's retry value is 2. When connectivity fails, it will try for the 1st time. When the timeout value is over it will retry for 2nd time. Backup will fail if connectivity does not restore and times out.

Figure 8 Backup Server Goes Offline

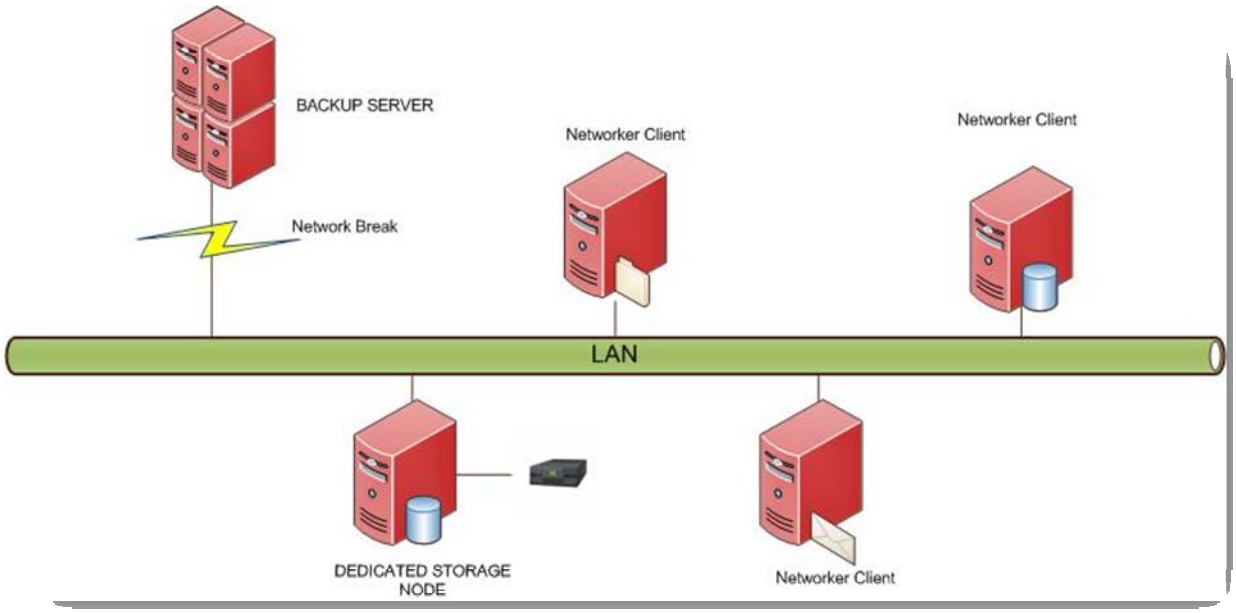

10. Never disable bootstrap/index savesets that are a part of every savegroup.

11. Never disable the index backup of NDMP savesets in a savegroup. This is the only way to retrieve the index of NDMP client in case of a NetWorker Server Crash.

12. If a firewall arbitrarily cuts connections that seem idle, there will be a network error (transport layer problem) to NetWorker, exactly as if a network cable had been unplugged. To prevent this, configure the firewall not to close idle connections, or have an operating system timeout long enough to accommodate the backup window. The status connection to nsrjobs is frequently idle for most of the backup: if there are no error messages to report, the connection will not have traffic until the success message when the backup is done. As a result, this connection is the most common connection to be cut by an aggressive firewall. If it is impossible to eliminate a firewall timeout, change the operating system's TCP keepalive interval on each machine if connections to other daemons time out. The following examples set value of the OS TCP Keep Alive to 57 minutes to be below default 60 minute timeout on most firewalls:

    Windows

    REG_DWORD: KeepAliveTime

    Key: Tcpip\Parameters

    Value Type: REG_DWORD—time in milliseconds

    Value: 3420000

    Solaris

    ndd -set /dev/tcp tcp_keepalive_interval 3420000

    Linux

    net.ipv4.tcp_keepalive_intvl = 30

    net.ipv4.tcp_keepalive_probes = 8

    1. net.ipv4.tcp_keepalive_time = 3420

Troubleshooting

1. The basic troubleshooting starts from studying the daemon.log file in /nsr/logs folder and other log files are rap.log , savegrp.log. You can view the logs to troubleshoot the cause of error if any. After 7.4 version daemon.raw file is created You can use nsr_render_log –mepathy daemon.raw >> daemon.log to covert .raw file into .log file

2. Block size error while recovering data of a client.
   The block size error may occur if we try to recover data of a particular client using another client particularly when the backup is performed using a dedicated storage node. Different drivers of HBA & tape drivers in different dedicated storage node clients are the root cause of the issue. Always maintain consistency between the HBAs and drivers loaded in all the dedicated storage node clients sharing the drives.

3. I/O reset marking the Tapes prematurely full.
   The I/O resets are caused primarily by hardware issues in drives. Cleaning drives may solve the issue. Removable Storage Service of the OS should also be checked & disabled to solve the issue.

4. Tape drive stops responding giving serial number mismatch error.
   This issue arises particularly when the persistent binding is not done on the drives. When the library reboots, the OS gives the new symbolic name to the drives and NetWorker stops recognizing them. When persistent binding is properly done, each time the library reboots the OS will match the WWNN # & will give it the same symbolic name as was previously configured.

5. Expected volume `xxxxx' in slot `<slot_number>'. The actual volume is `<NULL>'.
   The inventory is out of sync. Rerun the inventory to solve the issue.

6. Drives go into service mode.
   The drive goes into service mode if consecutive 20 errors are encountered in the drive operation. Change the maximum consecutive errors attribute from the properties of the device to 200.

7. Run the following command from the NetWorker server to test the NetWorker server connection to the nsrexecd daemon running on the client:
   nsradmin –s <client_name> -p 390113

8. 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.

9. Use the following command for detailed debugged log file when you are unable to figure out the error:
   savegrp -vvv -D9 -p -c source_client -G group_name > log.txt 2>&1

10. Regularly clean tape drives with the cleaning cartridge to avoid the tapes getting full prematurely or failed backups. Cleaning drives should be a regular practice.

Appendix A

Setup Details						
Backup Clients/Storage Nodes						
General Details						
Client Host Name						
Client IP Address						
Client Operating System						
Operating System Service Pack						
Any Other Application in use	MS SQL	MS Exchange	Lotus			
	Oracle	SAP	DB2			
Application Version						
Application Service Pack						
Legato Networker Details					**Network Details**	
Legato Networker Version in use					Connectivity Type	
Networker Installation Folder					NIC Card	
Application Backup	Yes/No				NIC Speed	
Application Module Version					NIC Setting	Full Duplex/Half Duplex/Auto
Client ID					Speed Setting	10/100/1000/Auto
Storage Node	Yes/No				Backup to	
Backup Device						
Media Type						
Connectivity						
Save Sets Details						
Save Set	Size	Browse	Retention	Time	Command	Level/Group
Schedule Details						
Name	Description					

BIOGRAPHY

I was born in the holy city of temples Jammu, the winter capital of Jammu and Kashmir (India), on 5th December 1986. I designed my own website and a website for a social organization when I was 13 years old. I am a Computer Engineer by profession and have completed my BE from the Model Institute of Engineering and Technology Jammu (Jammu and Kashmir).

I achieved my EMC Proven Professional Information Storage and Management certification on December 22, 2007. I love listening to music, reading technology-oriented books and shopping. I joined Ace Data Devices on September 4th 2008 after EMC referred me to ACE DATA DEVICES Pvt Ltd. Gurgaon, Haryana (India). Since then I have worked on products like EMC NetWorker®, EMC RepliStor®, EMC VisualSRM™, EMC Celerra®, EMC CLARiiON®, and EMC AutoStart™.

.