

EMC NetWorker Foundations

Upon completion of this module, you will be able to:

- Describe the EMC NetWorker solution
- Explain NetWorker advantages
- Explain NetWorker hosts and their roles
- Describe NetWorker control data and its use
- Explain NetWorker resources and their purpose
- Explain the NetWorker data-tracking process
- List NetWorker-supported device types and devices
- Describe the NetWorker backup process, types and levels
- Describe the NetWorker recovery process and types

The objectives for this module are shown here. Please take a moment to read them.


Lesson 1: EMC NetWorker Overview

Upon completion of this lesson, you will be able to:


- Describe the EMC NetWorker solution
- Describe the NetWorker host roles
- List the advantages of NetWorker
- Explain how multiplexing and Open Tape Format work
- Describe the NetWorker software editions and add-on modules

In this lesson you will gain an understanding of the EMC NetWorker solution.

The objectives for this lesson are shown here. Please take a moment to read them.




There are many pieces in an overall solution for your business computing needs. EMC NetWorker works within the existing framework of hardware, operating system software and network communication protocols to provide protection for the critical application data that the framework supports. NetWorker protects critical business data by centralizing, automating and accelerating backup and recovery operations across an enterprise. NetWorker provides support for the latest disk backup and snapshot technologies.



The three primary types of NetWorker hosts in a NetWorker Data Zone are:

- NetWorker Server
- NetWorker Storage Node(s)
- NetWorker Client(s)


Together, a single NetWorker Server, with its Storage Nodes and Clients, forms a *NetWorker Data Zone* within which data is protected. An enterprise may have more than one NetWorker data zone. NetWorker servers and storage nodes may belong to only one data zone. NetWorker clients may be backed up by multiple NetWorker servers and, therefore, may belong to multiple data zones.



The *NetWorker Client* is the largest NetWorker software component and the fundamental host. The client's most important functions are to generate backups called *save sets*, push them to a NetWorker storage node and retrieve them during a recovery. NetWorker clients are usually the data servers in an IT environment. The types of data that are typically backed up as save sets include file system data and applications.

While performing a backup, the client also generates *tracking information*, including the file and directory names in the backup and the time of the backup, and sends it to the server to facilitate point in time recoveries.

The client software includes graphical user interfaces (GUIs) and command line utilities that allow users to manually perform backup and recovery operations. NetWorker client software is installed on all participating hosts in the data zone, including hosts that also play the roles of server and storage node. Every host in a NetWorker data zone is a NetWorker client.




NetWorker Storage Nodes are hosts with directly-attached or SAN-accessible devices that provide the NetWorker interface for the backup devices and volumes. Storage node hosts have both the NetWorker client and storage node software installed.

A backup device attached to the NetWorker server is referred to as a *local device*. If a device is controlled by a host other than the NetWorker server, it is considered a *remote device* and the storage node controlling the device is referred to as a *remote storage node*. The NetWorker server is always a storage node and is the default storage node for backups. A NetWorker server can manage many storage nodes, but a NetWorker storage node can be managed by only one NetWorker server.

During a backup, a NetWorker client sends backup data to a particular storage node based on that client's configuration. The storage node organizes the client's data and writes the client's data to one of its devices. Storage nodes also send tracking information about the save sets written to the volume during the backup to the NetWorker server. This information is used for future backups as well as for recoveries.

During a recovery, the client reads from the storage node. The storage node provides the device that contains the necessary volume.



Starting with NetWorker 7.3, administration of a NetWorker server is performed using the *NetWorker Management Console Server* (NMC) GUI, commonly referred to as the *console server*, a Java-based interface accessible from any web browser.

The console server provides a global view of your NetWorker environment, allowing you to centrally configure and manage one or more NetWorker data zones.

The console server must have the NetWorker 7.3 or later NetWorker client software already installed. The console server will generally be backed up as a normal client by at least one NetWorker server.

The console server can prepare a number of preconfigured reports generated using information gathered from any or all of the NetWorker servers. Detailed customized reports can also be created.


The NetWorker Advantage



- Network-based backup protection
- Ease of administration
- Efficient backups
- Strengthened security for communications between NetWorker hosts
- A full range of data protection functions

EMC NetWorker is an award-winning, extremely efficient, highly configurable, distributed storage management solution. NetWorker offers the following advantages:


- Network-based backup protection
- Ease of administration
- Efficient backups
- Strengthened security for communications between NetWorker hosts
- A full range of data protection functions




NetWorker provides network-based backup protection that functions in LAN, WAN and SAN environments on open system protocols, including TCP/IP and Network Data Management Protocol, or NDMP.

NetWorker uses the client/server model, which distributes the workload and improves performance. It is supported on both Windows and UNIX operating systems. The supported platforms seamlessly co-exist within the data zone. For example, a Microsoft NetWorker client can backup to a UNIX NetWorker storage node and vice versa.

Consult the *EMC Information Protection Software Compatibility Guide* for a current list of NetWorker supported operating systems.



NetWorker can be used in a SAN to back up data to disk or library devices on the SAN. The configuration shown in the diagram is called LAN-free backup because no backup data travels over the LAN.



NDMP, Network Data Management Protocol, is a TCP/IP-based protocol that specifies how network components talk to each other for the purpose of moving data contained on each component across the network for backup and recovery.

NDMP is able to communicate with several different interface types for data transfer during backup and recovery. NDMP enables disparate vendors to use a common protocol for the backup architecture. Data is copied from disk to tape using NDMP, regardless of the operating system or platform. Tracking information is also passed to and from backup software using the same protocol, regardless of software type.

NDMP works with NetWorker for backups, cloning, recoveries, and updates or queries to resource files. Any communication to and from the NetWorker Server to any NDMP device including servers, disk and storage is done via NDMP.

By using NDMP with NetWorker, the NDMP-enabled NAS device sends data to the NDMP-enabled tape device where NetWorker keeps track of the backed up files. The NDMP-controlled tape device attaches directly to the NAS device or another location on the network.

With NetWorker 7.2 and later, NDMP hosts have the option to send their data to a storage node.

EMC Education Services

EMC²
where information lives[®]

Ease of Administration

- GUI (web-based browser using NMC)
- Command line interface (CLI)
- Highly configurable

The image shows two overlapping windows. The background window is 'leg-win1 - NetWorker Administration', which has a menu bar (File, Edit, View, Configuration, Window, Help) and a toolbar with icons for Monitoring, Configuration, Devices, and Media. A tree view on the left lists various configuration items like Archive Requests, Clients, Directives, Groups, Local Hosts, Notifications, Policies, Registrations, Schedules, Snapshot Policies, Staging, and User Groups. The foreground window is 'NetWorker Management Console - http://leg-win1:9000', which has a menu bar (File, Edit, View, Enterprise, Window, Help) and a toolbar with icons for Events, Enterprise, Libraries, Reports, and Setup. It displays a table with columns Type, Name, and Application Features, showing a single entry for 'NetWorker' with features 'Capture Events, Gather Reporting Data'. A red arrow points from the 'leg-win1' entry in the tree view of the background window to the foreground window. A yellow box labeled 'NetWorker Management Console Window' points to the foreground window, and another yellow box labeled 'NetWorker Administration Configuration window' points to the background window. A green box contains the text 'NetWorker Management Console provides centralized access to each NetWorker server'.

NetWorker Administration Configuration window

NetWorker Management Console Window

NetWorker Management Console provides centralized access to each NetWorker server

© 2006 EMC Corporation. All rights reserved.

EMC NetWorker Foundations - 13 Mercer Road
EMC Education Services

With NetWorker Management Console, you can easily administer NetWorker from any host having a supported web browser. For example, you can administer a UNIX NetWorker server from a Windows machine and vice versa. NetWorker also provides many command line utilities.

NetWorker is highly configurable. With its many configuration points, NetWorker can be tailored to meet the unique backup requirements of your specific site.

EMC Education Services

EMC²
where information lives[®]

NetWorker Management Console Reporting

NetWorker Management Console - http://USENLGILLES1E.eng.emc.com:9000

Events Enterprise Libraries Reports Setup

File Edit View Reports Window Help

Reports

- NetWorker Backup Statistics
- NetWorker Backup Status
 - Server Summary
 - Group Summary
 - Monthly Summary
 - Daily Summary
 - Group Summary by Server
 - Group Details
 - Group Summary by Client and Server
 - Daily Save Set Summary
 - Save Set Details
 - Save Set Details by Client
 - Save Set Output
 - Daily Backup Status
 - Group Status over Time
 - Group Status by Server
 - Client Status over Time
- Events
- Hosts
- Users
- Devices

Report: Daily Summary

Configure View Report

Server Name: all
Group Name: all
Group Start Time: from earliest available to 1/13/06 4:44:12 PM
Status: all

Daily Summary

Date	Total Duration (S...	Total Group Runs	Successful	Failed	Interrupted	Success Ratio
12/29/05	537	2	2	0	0	100%
1/13/06	506	2	2	0	0	100%

© 2006 EMC Corporation. All rights reserved.

EMC NetWorker Foundations - 14 Mercer Road
EMC Education Services


To facilitate NetWorker administration, a number of reports are available through the NetWorker Management Console Reports feature. Data maintained in the NMC server database, gathered from any or all of the NetWorker servers, is used to prepare reports on backup statistics and status, events, hosts, users and devices. Customized reports can also be created. Here you see an example of the NetWorker Backup Status Daily Summary report.



Efficient Backups

NetWorker supports:

- Multiple concurrent backups
- Backups from different operating systems
- Backups to multiple devices




© 2006 EMC Corporation. All rights reserved.

EMC NetWorker Foundations - 15 

NetWorker allows you to perform more efficient backups than most native operating system backup software programs.


Many native solutions limit you to one backup at a time, one operating system and only local devices. NetWorker is not bound by these restrictions and is capable of backing up multiple save sets from clients running different operating systems to any NetWorker-configured device. Features that support these capabilities include:

- Multiplexing
- Open Tape Format
- Remote device support




One of the ways that NetWorker achieves backup efficiency is by interleaving, or *multiplexing*, multiple backups onto a backup device at the same time.

In an environment without multiplexing, as illustrated here, only one stream of data is written to the device or volume at a time. This situation is not ideal because as more clients perform simultaneous backups, the tape drive's throughput is not optimized.




Multiplexing enables the device to simultaneously write more than one save set to a storage volume. This allows the device to write to the volume at the collective data rate of the save streams, up to the maximum data rate of the device.




Open Tape Format, or OTF, is a data format that allows multiplexed, heterogeneous (UNIX, Windows, NetWare, etc.) data to reside on the same tape.

As shown in the diagram, clients wrap their backup data into packages containing platform-independent data, like boxes, which are called save set chunks, before sending them to the storage node. The storage node receives the packages and arranges them on a pallet, or media records and files.


These pallets are then stored in the warehouse, or tape volume. The way the storage node organizes the pallets is also platform independent allowing any NetWorker storage node to read the data. Because of Open Tape Format, a NetWorker storage node can be migrated to a host running a different operating system.



Another feature of NetWorker that promotes efficient backups is the support of remote devices. This means that data can be backed up to devices controlled by remote storage nodes. Recall that a *remote storage node* is any storage node other than the NetWorker server. The remote device support feature allows you to distribute the backup workload throughout the data zone and can reduce network traffic.



With version 7.3, NetWorker authentication provides increased security for communications between NetWorker hosts. Users are authenticated using file system privileges. Host machines are authenticated using the Secure Sockets Layer (SSL) protocol and self-signed certificates.



NetWorker performs the data protection functions listed in the diagram. Following is a description of each function.

A *backup* is a copy of data you want to protect by storing it in another location. NetWorker provides a scalable solution to manage backups for a small network or an entire enterprise, and enables you to automate and configure this process for speed and efficiency.

Tracking and Reporting is the process of storing and accessing information about the backup. NetWorker saves and tracks information such as the location and volumes of the backup, the client that generated the backup, the backup creation date and time, and the backup type.

Recovery is the process of restoring data to a given point in time. NetWorker makes this process possible from anywhere within the NetWorker data zone. This feature enables users to retrieve single files or restore an entire file system.

Aging determines the length of time that the backup data is available for recovery. NetWorker allows you to specify how long to keep the tracking information viewable and how long to protect each volume from overwriting.

Staging is the optional process of moving a save set from one storage volume, the physical medium where the data is stored, to another.

Cloning is the optional process of making a copy of the backup save set or volume.

These data protection functions are presented in more detail throughout the module.

NetWorker Software Editions

NetWorker Edition	Max Clients	Max Jukeboxes /Slots	Remote Storage Node	Add-on Modules	Max Parallel-ism	Max Devices	Immediate Save and Recover
Workgroup /Business	8	1-32/ 1-26	No	Limited	12	4	No
Network	10 - unlimited	Unlimited/ unlimited	Yes	Yes	32-512	16-256	Yes
Power	10 - unlimited	Unlimited/ unlimited	Yes	Yes	64-512	32-512	Yes

Four NetWorker software editions are offered to meet your data protection needs:

- Workgroup Edition
- Business Edition
- Network Edition
- Power Edition

The table shows the different combinations of client connections, add-on modules, and architectural and performance improvements that are available with each edition.

NetWorker Add-on Modules

Add-on Module	Description	License Required?
Client Connection	A connected client is a computer that is configured on the NetWorker Server and has contact with the server to attempt a backup	Yes
ClientPak	Required for any NetWorker client that has a different operating system from the NetWorker server	Yes
Storage Node	Allow the NetWorker server to use remote devices	Yes
NetWorker Module	Allows online application and database backups	Yes

EMC offers numerous add-on licenses and modules that work in conjunction with your base edition of NetWorker to provide additional features and benefits.

Client Connection licenses are required for any computer that is configured on the NetWorker server and has contact with the server to attempt a backup.

ClientPak licenses are required for any NetWorker client that has a different operating system from the NetWorker server.

Storage Node licenses allow the NetWorker server to use remote devices.

NetWorker Module licenses allow online applications and databases to be backed up.

Each add-on module requires a separate license and some add-on modules require additional software to be installed.

Lesson Summary

Key points covered in this lesson include:

- NetWorker Clients generate backups and tracking information, perform recoveries and provide administrative interfaces.
- NetWorker Storage Nodes read and write backup data to and from devices.
- The NetWorker Server supports backups and recoveries, and stores tracking and configuration information.
- NetWorker advantages include:
 - Network-based backup protection
 - Ease of administration
 - Efficient backups through the use of multiplexing, open tape format and remote device support
 - Strengthened security through host and user authentication
 - Full range of data protection functions including backup, tracking and reporting, recovery, aging, cloning and staging
- NetWorker software editions and add-on modules are offered to meet your unique data protection needs.


These are the key points covered in this lesson. Please take a moment to review them.

Lesson 2: NetWorker Architecture

Upon completion of this lesson, you will be able to:

- Describe the types of NetWorker control data
- Describe NetWorker resources and their purpose
- Explain the various types of NetWorker tracking information
- List NetWorker administrative interfaces
- Describe the types of NetWorker-supported devices

In this lesson you will explore the architecture and product features of NetWorker in greater detail. The objectives for this lesson are shown here. Please take a moment to read them.




NetWorker *control data* collectively represents the NetWorker configuration information and the backup tracking information stored on the NetWorker Server.

The *Resource Directory* is the directory that contains the configuration resource files. The resource directory is also called the resource database.


The *Media Database* is the NetWorker database used to track the backups and the volumes that store the backups.

Client File Indexes, or CFIs, are the NetWorker databases that track each file or pathname in a client's backup, allowing clients to browse their backups for files from a particular point in time. The NetWorker Server creates and maintains one client file index per client.



Resources are used to configure a NetWorker environment. Resources are managed as configurable objects by the NetWorker administrator. NetWorker supports a resource type for each configurable component and there can be multiple configurations for each type. Examples of resource types include clients, devices, tape libraries, backup start times and policies.


Nearly all NetWorker resources are stored on and managed by the NetWorker server in the resource database. A few resources are managed on NetWorker clients.



The NetWorker Server maintains tracking information for save sets and volumes in a client file index for each client and in the media database.

The *Client File Index* (CFI) stores the records for each backed up file for each NetWorker Client. There is one CFI per client host. The information answers the question “what files were backed up and when?” for each client. This information is used to support browsable recoveries which allow clients to more easily recover to a point in time. As the save sets age, the CFI records are deleted automatically to save space. The amount of time the records are kept is configurable through the browse policy.

The *Media Database* tracks the NetWorker labeled volumes and the save sets on the volumes. The information answers the question, “on which volumes are the save sets located?” All volumes and save sets are listed in the Media Database. This information is needed to support recoveries.



Save sets and volume records in the media database are checked for aging based on browse and retention policies. The browse policy determines how long file entries remain in the CFI. The retention policy determines how long a save set is protected from being automatically overwritten through the recycling (relabeling) of the volume on which it resides.

Status changes are made to the CFI when a save set changes from browsable to recoverable to recyclable to appendable.

NetWorker Interfaces


	UNIX	Microsoft Windows
Administer NetWorker Server	NetWorker Management Console nsradmin (character-based)	NetWorker Management Console nsradmin (character-based)
Backup	nwbackup NetWorker Management Console	NetWorker User NetWorker Management Console
Recover	nwrecover	NetWorker User

NetWorker provides administrative interfaces, including GUI and command line interfaces, to manage, configure and monitor the NetWorker Server.


The interfaces for administering the NetWorker server are the NetWorker Management Console GUI and **nsradmin**. **nsradmin** is a command line utility used to view, create, delete and modify resources.

There are also user interfaces for initiating NetWorker backups and recoveries. These interfaces include the NetWorker Management Console GUI for both UNIX and Windows, NetWorker User on Windows, and **nwbackup** and **nwrecover** on UNIX.

All user interfaces are discussed in more detail later in this module.



The NetWorker Management Console provides administrative, configuration and monitoring interfaces for both UNIX and Windows environments. The Console window provides a global view of all NetWorker servers in your environment. From this interface, you can open a NetWorker Administration window for each NetWorker server.



In the NetWorker environment, three types of devices can be used to back up data: tape, optical and file type.

Tape devices have several associated media types including 8mm, AIT, DLT and LTO.

Optical devices include magneto optical devices such as a Hewlett-Packard 5200ex.


A *File type* device is an existing file system directory configured in NetWorker as a device resource with media type file or advanced file. Once configured, NetWorker uses the directory as a backup volume to read or write save set data. Each save set directed to the device is written to a separate file within the directory. Storing data to a file device often reduces the time it takes to both save and recover data.

A device can be either a standalone device or an autochanger device.

A *stand-alone device* is any type of device that does not have a robotic arm for loading volumes. Media are manually loaded or preloaded before a backup runs. NetWorker does not have to load the media as part of the backup procedure.

An *autochanger*, sometimes called a library, silo, medium changer, TLU, or jukebox, is a multiple-volume device that uses a robotic arm.


Refer to the *EMC NetWorker 7.x Hardware Compatibility Guide* for a current list of supported NetWorker devices.




Autochangers can be either dedicated to a single storage node or several systems in the network can share autochangers if the autochanger has more than one drive.

In a *dedicated autochanger*, all devices are controlled by a single storage node. Backup data from clients other than the NetWorker server must be sent to the storage node using the TCP/IP network.

A *shared autochanger* is one in which each tape drive is statically assigned to a specific storage node and two or more storage nodes are assigned drives. The robotic arm is controlled by a single storage node, typically a storage node controlling one of the drives. Autochangers cannot be shared across data zones.




With *Dynamic Drive Sharing*, or DDS, one or more tape drives are accessed and used by two or more storage nodes within the same data zone. The devices can be stand-alone or in an autochanger. Dynamic Drive Sharing is useful in reducing hardware needs and costs by allowing multiple storage nodes to use the same device. Once configured, the administration of shared drives is the same as for non-shared drives.



File type devices are generally faster than tape devices for both backup and recovery. There are two types of file type devices, regular and advanced. The advanced file type provides several benefits over the regular file type:

- Advanced file type devices support multiple backups and one read operation simultaneously. This allows you to recover, clone, or stage data from the device while backups are in progress.
- With NetWorker 7.3 or newer, multiple save sets on a single advanced file type volume can be recovered simultaneously.
- An advanced file type device can reside on a Windows file share, CIFS.

Because of these benefits, advanced file type devices are well-suited for use with NetWorker's staging feature.




Advanced file type devices respond differently than file type devices to the disk full condition. File type devices behave much like tape devices. When there is no more room on the volume, NetWorker marks the volume full and continues the backup on another volume. This may be another disk volume or a tape volume.

A save set being written to an advanced file type device will never continue on another volume. When the disk becomes full, NetWorker suspends all saves being directed to that device until more space is made available on the volume.


NetWorker Pools

- Represent a collection of NetWorker volumes
- Used to group backup data onto specific media



A pool is a NetWorker resource that represents a collection of NetWorker-labeled volumes. A volume is associated with a pool when the volume is labeled. NetWorker determines which pool of volumes to use for a backup by comparing the characteristics of the save set to corresponding attributes of configured backup pools. If a save set does not explicitly match any other pool, NetWorker sends the save stream to a volume in the default pool.

Pools allow you to group backup data onto specific media. For example, you can use pools to retain certain save sets longer than others or to write sensitive data from a particular client to its own set of volumes.



NetWorker provides the ability to further manage and protect save sets and volumes through cloning and staging.

Cloning copies save sets or a volume of save sets to another volume. With cloning, you can create identical copies of save sets in case of damage to original media or for offsite storage. No volume contains more than one instance of a save set. Cloning can be done either manually or automatically.

Staging moves save sets to another volume. Staging is often used to move save sets from file type and advanced file type devices to long term media such as tape. This allows the most recent backups to be written to and recovered from disk and then moved to tape to free space for subsequent backups.

Lesson Summary

Key points covered in this lesson include:

- NetWorker control data includes configuration and backup tracking information.
- NetWorker provides administrative interfaces, including both GUI and command line interfaces, to configure, manage and monitor NetWorker servers.
- NetWorker can use tape, optical and file type devices.
- Autochangers can be dedicated to a specific storage node or shared among multiple storage nodes.
- Tape drives on a SAN may be dynamically shared.
- With NetWorker 7.3 or newer, multiple save sets on a single advanced file type volume can be recovered simultaneously.
- Cloning copies save sets to another volume; staging moves save sets to another volume.


These are the key points covered in this lesson. Please take a moment to review them.

Lesson 3: NetWorker Backups

Upon completion of this lesson, you will be able to:

- Describe NetWorker backup types
- Describe NetWorker backup levels
- Identify NetWorker backup GUIs and commands
- Explain the data flow in a client-initiated backup
- Explain the data flow in a server-initiated backup

In this lesson you will learn about the two different types of NetWorker backups.
The objectives for this lesson are shown here. Please take a moment to read them.



A backup is the process of copying data to another location. The destination can be another directory on the same system, another server, or some type of storage media.

NetWorker manages this backup data, which includes client and tracking information, and directs it to the proper destination. NetWorker also enables you to automate and schedule the backup process.


In a NetWorker backup, the NetWorker client pushes the data to the destination storage node where it formats the data and writes it to a volume in a backup device which it manages. The NetWorker server stores the tracking information.

NetWorker Backup Types

Client-initiated Backup	Server-initiated Backup
Manual	Scheduled
User on a NetWorker client initiates the backup	NetWorker server requests a backup of a client
Performed using a GUI or command line	Performed using a GUI or command line

NetWorker allows you to perform two types of backups:

- A *client-initiated backup* is initiated from a NetWorker client. This is a manual process that the user must perform.
- A *server-initiated backup* is started from the NetWorker server. The NetWorker server sends a backup request to one or more NetWorker clients. A server-initiated backup is usually configured to start automatically, but may also be performed manually, either from the NetWorker Administration window or the command line.




A *client-initiated backup* is a manual process performed on a NetWorker client from either a GUI or the command line. This type of backup is useful any time you want to perform a backup outside the regular backup schedule. The user specifies which files, directories and file systems to save.

Although the NetWorker server does not initiate a client-initiated backup, it manages the backup after the client makes a request. This management includes authorizing the backup and determining to which storage node and backup device the client should send its save stream.

For a client to execute any type of backup, it must first be configured as a client resource on the NetWorker server. This can be as simple as specifying the client host name in a client resource and using default values for all attributes.

When the client performs a **save**, it generates a save stream, sends it to the assigned storage node and sends tracking information to the NetWorker server. The storage node also generates tracking information which it sends to the server.




NetWorker provides UNIX and Windows GUIs for performing client-initiated backups.

On UNIX, the nwbackup command is the GUI front end to the save command. The nwbackup interface provides a File Manager look and feel from which you select the files and directories to backup.

The NetWorker User GUI is used to perform a client-initiated backup on Windows. NetWorker User provides a Windows Explorer look and feel from which you select the files and directories to backup.

Instead of using a GUI, you can execute the save command directly from the command line on any NetWorker client.




A server-initiated backup is a backup request from the NetWorker Server to a configured group of clients. Although a server-initiated backup can be run manually from the NetWorker Administration window or the **savegrp** command, it is normally scheduled to start automatically.

The remote execution request causes the clients to run their backup commands for their configured save sets. The client and other configured resources determine what, when and how to back up. Tracking information is sent to the NetWorker server and is also stored on the backup volumes.

Server-initiated backups also cause the NetWorker server to back up the *client file indexes* and the *bootstrap save set* to protect the NetWorker configuration and tracking data. The bootstrap is a special save set containing all the information in the media and resource databases. NetWorker generates a *savegroup completion report* and either mails it to root on the NetWorker server on UNIX or appends it to the savegroup log on the NetWorker server on Windows. The report contains information about the backup status, backup contents, data size, and backup creation date and time.


The NetWorker server generates and prints information about the bootstrap save set such as the volume used and the bootstrap's save set ID. This information is essential for NetWorker server disaster recovery.

The NetWorker server also does a consistency check on the client file indexes and media database and performs the aging of save sets.



The diagram shows the steps in the backup process flow that are performed during both server-initiated and client-initiated backups.

1. The client sends the save stream and tracking information.
2. The server writes tracking information to the client file index.
3. The storage node receives the save stream and sends it to the backup device.
4. The storage node sends save set and volume information to the NetWorker server.



The diagram shows the steps in the backup process flow that are performed in server-initiated backups only. These steps are done by the NetWorker Server:

1. The NetWorker server requests the backup from the NetWorker client.
2. NetWorker backs up metadata: The NetWorker server backs up client file indexes for all clients in the group; the NetWorker Server backs up the Media Database and Resource files as the bootstrap save set.
3. The NetWorker Server triggers preconfigured notifications.
4. The NetWorker Server enforces data life-cycling policies through database management.

EMC Education Services

EMC²
where information lives[®]

Customizing the Backup Environment

Client Resource

Properties

General \ Apps & Modules \ Globals (1 of 2) \ Globals (2 of 2) \ Info & Licensing \ Index \ System

Identity

Name: leg-win1

Comment:

Index Management

Browse policy: Month

Retention policy: Year

Backup

Scheduled backup: ☒

Directive: All

Save set: All

Group:

☐ astronomy

☒ Default

☐ test

Schedule: Default

Policies

Name	Number	Period
Month	1	Months
Decade	10	Years
Week	1	Weeks
Year	1	Years
Quarter	3	Months
Day	1	Days

Schedules

Name	Start time
Consolidate on 1st of Month	
Consolidate 1st Friday of Month	
Consolidate Every Friday	
Full on 1st Friday of Month	
Default	
Full on 1st of Month	
Full Every Friday	
Consolidate Quarterly	
Quarterly	

Groups

Name	Start time
IT	2:30
astronomy	23:00
Default	3:33

© 2006 EMC Corporation. All rights reserved.

EMC NetWorker Foundations - 48

Mercer Road
EMC Education Services

NetWorker backups can be customized by configuring client resources that determine the specifics of how a client is backed up. You define the save sets that will be backed up by the client. *Policies* establish the aging parameters to be used for a client's save sets. *Schedules* define the backup level used for a client backup on any given day. The *Group* resource specifies the start time for server-initiated backups.

EMC Education Services

EMC²
where information lives™

NetWorker Backup Levels

A recover here requires only the full or consolidated backup

Full backs up all files

Consolidated performs a level 1 and merges it with most recent full to create a new "consolidated full"

Incremental backs up all files changed since last backup

Skip the backup this day

Level 1-9 backs up all files changed since most recent lower level backup

Create Schedule wizard

© 2006 EMC Corporation. All rights reserved. EMC NetWorker Foundations - 49 Mercer Road EMC Education Services

NetWorker supports several backup levels.

Full backups back up all files and directories in a save set and are the lowest level backup. This type of backup requires the most storage space and takes the longest time to perform.

A *level backup* backs up only files that have changed since the most recent lower level backup. Valid levels include level 1 through level 9. For example, level 5 backs up any files and directories that changed since a previous level 4 or lower level backup. A full backup is equivalent to a level 0.

Incremental backups back up files and directories that have changed since the last backup of any level.

Using level [1–9] and incremental backups takes less time than full backups and uses less volume space. However, using them can also slow file recovery if several save sets are required to recover to a particular point in time.

Consolidated backups perform a level 1 backup which is then merged with the save set's last full backup creating a new full backup. The difference between a full and a consolidated backup is that consolidated backups are typically less taxing on the network and client. As a result, they increase backup performance.

Skip backups do not back up the client's save sets on a specific day. The `savegrp` program generates a savegroup completion report that shows that the backup for the clients using this schedule were skipped and that no data was backed up.

Lesson Summary

Key points covered in this lesson include:

- A backup is the process of copying data to another location.
- A client-initiated backup is a manual process that the user starts from a NetWorker client.
 - Useful any time you want to perform a backup outside of the regular backup schedule
 - User specifies which files to save
- A server-initiated backup is a backup request from the NetWorker server to a configured group of clients.
 - Usually configured to start automatically
 - Server-initiated backups cause the server to back up the client file indexes and the bootstrap save set
- Backup levels include full, levels 1-9, incremental and consolidated.

These are the key points covered in this lesson. Please take a moment to review them.


Lesson 4: NetWorker Recoveries

Upon completion of this lesson, you will be able to:

- Describe client roles in a recovery
- Describe the NetWorker recovery types
- Identify NetWorker recovery GUIs and commands


In this lesson you will learn about the three types of client recoveries that are supported by NetWorker.

The objectives for this lesson are shown here. Please take a moment to read them.



A recovery restores data to its original state at a specific point in time. NetWorker is flexible in how recoveries are performed while at the same time maintaining necessary security to avoid recovery of data by non-authorized persons. Recoverable data can include files, directories, file systems, or application data. NetWorker detects and can be configured to automatically resolve naming conflicts. Files can be recovered to a directory other than the directory from which they were backed up.

There are three types of recoveries: browsable, save set and directed. All three types are manual processes initiated from a NetWorker client. NetWorker does not schedule recoveries or perform them automatically.




In any recovery, there are three client roles performed by one or more NetWorker hosts as illustrated in the diagram.

The *source client* is the NetWorker client whose data is being recovered. It is the client that created the save set.

The *destination client* is the NetWorker client that receives the recovered data. It is the client where the data will be recovered.

The *administering client* is the client who initiates the recovery by running the NetWorker recover program.


Often, a single host performs all three roles in a recovery.



A *browsable recovery* allows you to recover files and directories from browsable save sets. Browseable recoveries are performed by NetWorker users and administrators using interfaces that require information from the client file index. Only those files for which the user has read permission can be recovered. In a browsable recovery, the administering and destination clients are the same host.

During the recovery, the user selects the set of files and directories to be recovered. When recovering an entire directory or file system, a point-in-time recovery is automatically performed. This restores the directory or file system to the way it looked as of the most recent backup. You can also recover a version of a file other than the most recent version.


A browsable recovery is useful when many files have been deleted or renamed since the last full backup.



A *save set recovery* allows you to recover data from a save set rather than browsing and marking data for recovery. System administrator privileges are required to perform a save set recovery. Data that you can recover includes individual files, directories and the entire save set itself. With the save set recovery, the administering and destination clients are the same host.


Since a save set recovery does not utilize client file index information, it does not perform a point-in-time recovery. Recovering to a specific point in time using save set recoveries may require multiple recoveries followed by manual deletion and renaming of files.

A save set recovery is useful when recovering an entire save set or when the save set being recovered is no longer in the client file index.



A *directed recovery* is defined as a browsable recovery in which the administering and destination clients are different NetWorker hosts. The source client may be the administering client, the destination client or another host. However, because backed up data has a specific data format, the source and destination clients must have similar operating systems and file system formats. Directed recoveries are performed only by a NetWorker Administrator. In the example in the diagram, a NetWorker administrator on C3 initiates a directed recovery causing C2 to recover data from C1's save sets.

Directed recovery is useful when you cannot log on to the destination client to perform the recovery. This remote administration capability gives you a high degree of flexibility and convenience in recovering and redirecting data.



NetWorker provides command line and GUIs for performing recoveries.

To perform browsable, save set and directed recoveries, you can execute the `recover` command directly from the command line on the administering client on both UNIX and Windows hosts. Illustrated here is an example of using `recover` in interactive mode to perform a browsable recovery.

On UNIX, the `nwrecover` command is the GUI front end to perform browsable and save set recoveries. This `nwrecover` example shows the selection of a specific file, `/etc/default/inetinit`, in a browsable recovery.

On Windows, NetWorker User provides the GUI front end for browsable, save set and directed recoveries. As with `nwrecover`, NetWorker User enables users to mark one or more items for recovery. The example shows the use of NetWorker User to perform a save set recovery.

Recovered files retain the modification time and permissions they had when backed up. The modification time does not change when the file is recovered.

Lesson Summary

Key points covered in this lesson include:

- A recovery restores data to its original state at a specific point in time.
- The source client is the NetWorker Client whose data is being recovered. It is the client that created the save set.
- The destination client is the NetWorker Client that receives the recovered data.
- The administering client is the client who initiates the recovery by running the NetWorker recover program.
- Browsable recoveries allow you to recover files and directories from browsable save sets.
- Save set recoveries allow you to recover files and save sets.
- Directed recoveries allow you to centrally administer the recovery of files, directories and save sets.

These are the key points covered in this lesson. Please take a moment to review them.

Module Summary

Key points covered in this module include:

- EMC NetWorker protects critical business data by centralizing, automating and accelerating backup and recovery operations across an enterprise.
- NetWorker advantages include network-based backup protection, ease of administration, efficient backups, strengthened security and a full-range of backup protection functions.
- Host roles include NetWorker server, NetWorker storage nodes, NetWorker clients and the NetWorker Management Console Server.
- There are two types of backups: client-initiated and server-initiated. Backup levels include full, levels 1-9, incremental and consolidated backups.
- Types of NetWorker recoveries include browsable, save set and directed recoveries.

These are the key points covered in this module. Please take a moment to review them.