

# EMC<sup>®</sup> NetWorker

Release 8.1

## Server Disaster Recovery and Availability Best Practices Guide

P/N 302-000-554  
REV 01

EMC<sup>2</sup>

Copyright © 1990 - 2013 EMC Corporation. All rights reserved. Published in the USA.

Published July, 2013

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

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC<sup>2</sup>, EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to the technical documentation and advisories section on the EMC online support website.

# CONTENTS

<b>Preface</b>	
<b>Chapter 1</b>	<b>Introduction</b>
	Revision history ..... 10
	NetWorker server disaster recovery roadmap ..... 10
<b>Chapter 2</b>	<b>Availability and Recovery Options for a NetWorker Server</b>
	Bootstrap and indexes ..... 14
	Bootstrap save set ..... 14
	Client file index save set..... 14
	Bootstrap recommendations and practices ..... 15
	How to obtain the bootstrap ..... 15
	Gathering the key information ..... 16
	Hardware information ..... 16
	Software information..... 16
	Disaster recovery scenario review ..... 17
	Basic disaster recovery (same host) ..... 17
	Advanced disaster recovery (different host)..... 18
	Ground level preparation for NetWorker server disaster recovery ..... 19
<b>Chapter 3</b>	<b>Data Storage and Devices</b>
	Capabilities and considerations ..... 22
	NetWorker metadata storage ..... 22
	Multi-path access and failover..... 23
	Storage devices and media ..... 23
	Method of connectivity..... 23
	Reliability and dependencies ..... 25
<b>Chapter 4</b>	<b>Disaster Recovery Use Cases</b>
	Basic disaster recovery scenario ..... 28
	Basic disaster recovery considerations ..... 30
	More advanced disaster recovery considerations..... 32
	Clustered solutions ..... 34
	Backup to disk ..... 35
	Index or configuration corruption ..... 36
	Corruption or loss of SAN storage ..... 36
	Loss of one server, Data Domain system, or site ..... 37
	Replication solutions ..... 37
	Replication of the NetWorker Server ..... 38



# PREFACE

*As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.*

*Contact your EMC representative if a product does not function properly or does not function as described in this document.*

---

**Note:** This document was accurate at publication time. New versions of this document might be released on the EMC online support website. Check the EMC online support website to ensure that you are using the latest version of this document.

---

## Purpose

This document describes how to design and plan for a NetWorker disaster recovery. However, it does not provide detailed disaster recovery instructions. The Disaster Recovery section of the EMC NetWorker SolVe Desktop (formerly known as the NetWorker Procedure Generator (NPG) provides step-by-step disaster recovery instructions that are tailored to your environment.

You can download the EMC NetWorker SolVe Desktop from the EMC Online Support Site at <https://support.emc.com/products/1095> under the **Tools and Utilities** section.

## Audience

This guide is part of the NetWorker documentation set, and is intended for use by system administrators who are responsible for setting up and maintaining backups on a network. Operators who monitor daily backups will also find this guide useful.

## Related documentation

The following EMC information products provide additional information:

- ◆ *EMC NetWorker Installation Guide*  
Provides instructions for installing or updating the NetWorker software for clients, console, and server on all supported platforms.
- ◆ *EMC NetWorker Cluster Integration Guide*  
Contains information related to installation of the NetWorker software on cluster server and clients.
- ◆ *EMC NetWorker Release Notes*  
Contain information on new features and changes, fixed problems, known limitations, environment and system requirements for the latest NetWorker software release.
- ◆ *EMC NetWorker Administration Guide*  
Describes how to configure and maintain the NetWorker software.
- ◆ *EMC NetWorker and EMC Data Domain Deduplication Devices Integration Guide*  
Provides planning and configuration information on the use of Data Domain devices for data deduplication backup and storage in a NetWorker environment.

- ◆ *EMC NetWorker and VMware Integration Guide*  
Provides planning and configuration information on the use of VMware in a NetWorker environment.
- ◆ *EMC NetWorker Snapshot Management Integration Guide*  
Provides the ability to catalog and manage snapshot copies of production data that are created by using mirror technologies on EMC storage arrays.
- ◆ *EMC NetWorker and EMC Avamar Integration Guide*  
Provides planning and configuration information on the use of Avamar in a NetWorker environment.
- ◆ *EMC NetWorker Error Message Guide*  
Provides information on common NetWorker error messages.
- ◆ *EMC NetWorker Performance Optimization Planning Guide*  
Contains basic performance tuning information for NetWorker.
- ◆ *EMC NetWorker Command Reference Guide*  
Provides reference information for NetWorker commands and options.
- ◆ *EMC NetWorker Licensing Guide*  
Provides information about licensing NetWorker products and features.
- ◆ *NetWorker License Manager 9th Edition Installation and Administration Guide*  
Provides installation, setup, and configuration information for the NetWorker License Manager product.
- ◆ *EMC NetWorker Software Compatibility Guide*  
Lists supported client, server, and storage node operating systems for NetWorker, NetWorker Modules, and options.
- ◆ NetWorker Management Console Online Help  
Describes the day-to-day administration tasks performed in the NetWorker Management Console and the NetWorker Administration window. To view Help, click Help in the main menu.
- ◆ NetWorker User Online Help  
Describes how to use the NetWorker User program, which is the Windows client interface connect to a NetWorker server to back up, recover, archive, and retrieve files over a network.
- ◆ EMC NetWorker SolVe Desktop (formerly known as the NetWorker Procedure Generator (NPG))  
The NetWorker Procedure Generator (NPG) is a stand-alone Windows application used to generate precise user driven steps for high demand tasks carried out by customers, support, and the field. With the NPG, each procedure is tailored and generated based on user-selectable prompts.  
  
To access the NetWorker Procedure Generator, log on to <https://support.emc.com/> and search for NetWorker Procedure Generator. You must have a service agreement to use this site.

- ◆ Technical Notes and White Papers  
Provides in-depth technical perspectives about product or products as applied to critical business issues or requirements. Technical notes and white paper types include technology and business considerations, applied technologies, detailed reviews, and best practices planning.

## Conventions used in this document

EMC uses the following conventions for special notices:

**NOTICE**

NOTICE presents information related to hazards.

**Note:** A note presents information that is important, but not hazard-related.

## Typographical conventions

EMC uses the following type style conventions in this document:

Normal	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> <li>• Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus</li> <li>• Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities</li> <li>• URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications</li> </ul>
<b>Bold</b>	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and man pages  Used in procedures for: <ul style="list-style-type: none"> <li>• Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus</li> <li>• What the user specifically selects, clicks, presses, or types</li> </ul>
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> <li>• Full titles of publications referenced in text</li> <li>• Emphasis, for example, a new term</li> <li>• Variables</li> </ul>
Courier	Used for: <ul style="list-style-type: none"> <li>• System output, such as an error message or script</li> <li>• URLs, complete paths, filenames, prompts, and syntax when shown outside of running text</li> </ul>
<b>Courier bold</b>	Used for specific user input, such as commands
<i>Courier italic</i>	Used in procedures for: <ul style="list-style-type: none"> <li>• Variables on the command line</li> <li>• User input variables</li> </ul>
< >	Angle brackets enclose parameter or variable values supplied by the user
[ ]	Square brackets enclose optional values
	Vertical bar indicates alternate selections — the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

## Where to get help

EMC support, product, and licensing information can be obtained as follows:

**Product information** — For documentation, release notes, software updates, or information about EMC products, licensing, and service, go to the EMC Online Support website (registration required) at:

<https://support.emc.com/>

**Technical support** — For technical support, go to EMC online support and select Support. On the Support page, you will see several options, including one to create a service request. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

**Online communities** — Visit EMC Community Network <https://community.EMC.com/> for peer contacts, conversations, and content on product support and solutions. Interactively engage online with customers, partners, and certified professionals for all EMC products.

## Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to:

[BSGdocumentation@emc.com](mailto:BSGdocumentation@emc.com)

# CHAPTER 1

## Introduction

This chapter includes the following sections:

- ◆ [Revision history](#) ..... 10
- ◆ [NetWorker server disaster recovery roadmap](#) ..... 10

## Revision history

E-mail your clarifications or suggestions for this document to:  
[BSGdocumentation@emc.com](mailto:BSGdocumentation@emc.com)

[Table 1 on page 10](#) lists the revision history of this document.

Table 1

### Revision History

Revision	Date	Description - added or changed sections
01	July 2013	Initial release

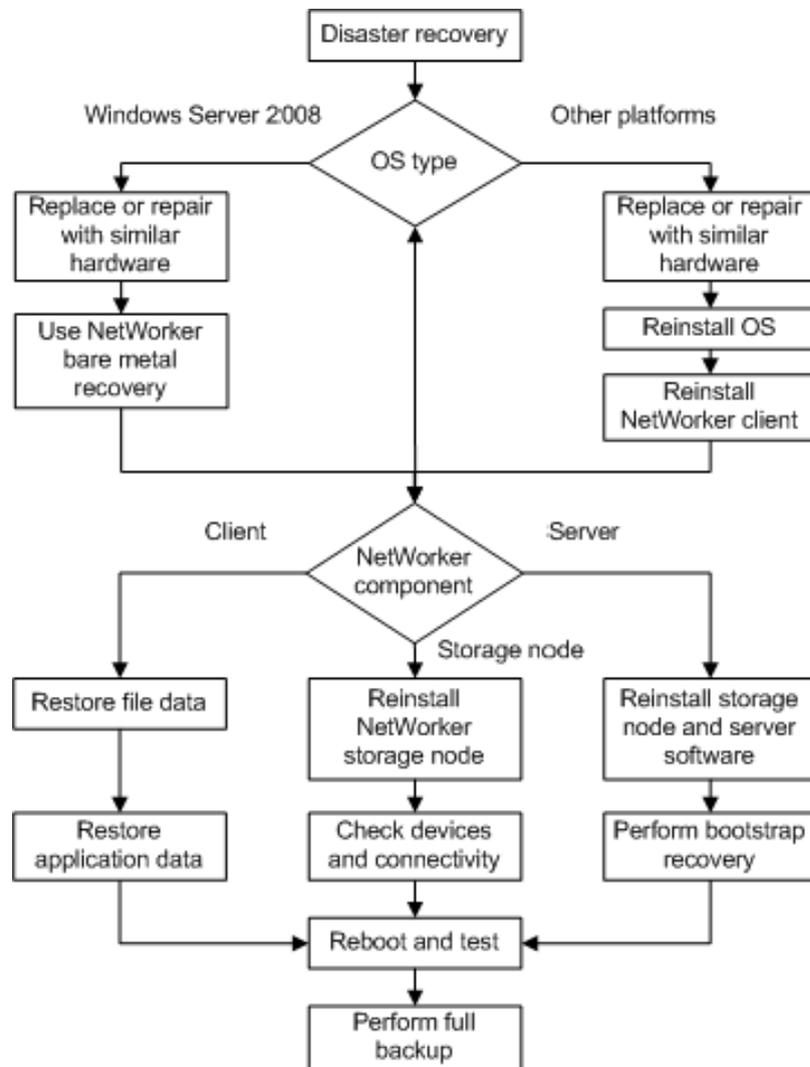
## NetWorker server disaster recovery roadmap

This guide provides an aid to disaster recovery planning. It does not provide detailed step-by-step disaster recovery instructions.

The Disaster Recovery section of the NetWorker Procedure Generator (NPG) provides step-by-step disaster recovery instructions that are tailored to your environment.

You can download the NPG from the EMC Online Support Site at <https://support.emc.com/products/1095> under the **Tools and Utilities** section.

[Figure 1 on page 11](#) lists the high-level steps to follow when performing a disaster recovery of the NetWorker server.



**Figure 1** Disaster recovery roadmap



# CHAPTER 2

## Availability and Recovery Options for a NetWorker Server

This section provides an overview of the various options that can be used to protect and recover a NetWorker server.

This chapter includes the following sections:

- ◆ [Bootstrap and indexes .....](#) 14
- ◆ [Gathering the key information .....](#) 16
- ◆ [Disaster recovery scenario review .....](#) 17

## Bootstrap and indexes

Backing up key configuration information is central to the recovery of a NetWorker server. This configuration information is stored in various locations on the NetWorker server and can change as different clients, devices, and volumes are used, updated, or changed.

The two main backup components that protect this stored data include the:

- ◆ [“Bootstrap save set” on page 14](#)
- ◆ [“Client file index save set” on page 14](#)

### Bootstrap save set

The bootstrap is a special save set that is generated by the backup server. The bootstrap backup contains key information about the current state and configuration of NetWorker clients, devices, volumes, and other important information for backup and recovery operations.

The bootstrap consists of three components that reside on the NetWorker server:

- ◆ The media database of the NetWorker server.
- ◆ The resource database of the NetWorker server that includes the jobs database.
- ◆ The server index.

The bootstrap backup typically occurs after each backup or savegroup completes and is generally small in size. Backing up this save set is the only guaranteed method to capture configuration information in a safe and consistent way. The availability of this save set is required to ensure a successful disaster recovery of the NetWorker server, regardless of any other protection methods that are used.

### Client file index save set

After all of the save sets in a scheduled backup for a client completes, the NetWorker software saves the client-specific backup information to the client file index. Each client has a client file index directory which is stored in the `nsr/index` directory on the NetWorker server. The client file index acts as a record of backup data and enables simple recovery and the ability to browse and restore the data. A client file index consists of many separate files and directories, and its size depends on the amount of client data backed up.

Each client file index contains the following information:

- ◆ Backups that have been performed for a client
- ◆ Backup level and type of backup
- ◆ File attributes

The client file index is not always required to recover data. You should back up the client file index and ensure that it is available for recovery by using the appropriate bootstrap information. The availability of the client file index greatly impacts the full restoration of backup and recovery services following a disaster recovery. You can also use the client file index to determine the time required to restore a NetWorker server to a fully functional state.

You can use the **nsrck** command to rebuild the client file index for a client from the index backup.

## Bootstrap recommendations and practices

By default, if the NetWorker server is a member of an active group, the bootstrap is backed up once all of the backups for a save group have completed. If the NetWorker server is not a member of an active group, the bootstrap backup is performed after all of the backups for every save group has completed.

To ensure that the latest NetWorker server configuration information is captured:

- ◆ Maintain a record of the bootstraps for reference. The record should be separate and independent from the backup server or any of its components. You can retain email or printed copies of the bootstrap record.
- ◆ Provide the following information in the bootstrap record:
  - The date and time of the bootstrap backup.
  - The volume and location that the bootstrap save set is stored on.
  - The save set ID of the bootstrap.
  - The starting file and record number of the bootstrap save set on the volume.
- ◆ Perform a bootstrap backup regularly, after all of the save sets in a save group have completed or at least once every 12 hours.
- ◆ Clone bootstrap volumes regularly to ensure that a single media failure or loss does not impact the recovery of the NetWorker server.
- ◆ Write the bootstrap save set to a device that is local to the NetWorker server.
- ◆ Write the bootstrap save set to separate, dedicated media.
- ◆ Do not mix the bootstrap save set with client backup data. This procedure speeds up the recovery process and ensures that the recovery of the NetWorker server is not dependant on client data volumes that might have inappropriate policies or protection.
- ◆ Ensure that the location of the media does not impact the access to the bootstrap data if a local disaster occurs such as a flood, fire, or loss of power. Although local copies of the bootstrap data are beneficial, they should maintain multiple copies of this information.

## How to obtain the bootstrap

You can obtain the bootstrap record in the following ways:

- ◆ Configure the bootstrap notification to email or print a copy of the bootstrap record.
- ◆ Use the **mminfo -B** command.
- ◆ Review the savegroup completion report. This report lists the bootstrap record when the save set is generated during a save-group backup.

## Gathering the key information

To aid in quick disaster recovery, maintain accurate records for each hardware, software, network, device, and media component.

### Hardware information

Maintain the following hardware information and ensure that is kept up to date:

- ◆ Volume or file-system configuration
- ◆ Fully qualified domain names, IP addresses, and host names
- ◆ References for Domain Name Servers (DNS) gateways, Active Directory, or domain servers
- ◆ Hard drive configuration
- ◆ Media device names and paths
- ◆ Hardware vendor contact information and contract numbers
- ◆ Configuration information for each piece of hardware, both active and inactive, for each system.

### Software information

Maintain the following software information and ensure that is kept up to date:

- ◆ Copies of the original operating system media and patches and where they are located
- ◆ Software enabler and authorization codes
- ◆ Software vendor contact information and contract numbers
- ◆ The operating system version and patches that were installed
- ◆ Operating system configuration
- ◆ Emergency media that can be used to recover a computer if a disaster occurs
- ◆ NetWorker bootstrap information for each NetWorker server
- ◆ Kernel configuration and location
- ◆ Device drivers
- ◆ A list of any Windows volume mount points and UNC paths

## Disaster recovery scenario review

The following disaster recovery scenarios might be encountered. Each scenario requires a different number of recovery steps and might be easier or more challenging to plan for or recover from. The NPG provides step-by-step guidelines on how to perform a disaster recovery by using the NetWorker software on different OS platforms.

In the simplest scenario, the same physical server remains in place with little or no changes to the original configuration or the surrounding environment. This is typical of a simple component failure such as a disk or power supply where the base operating system might have been removed or corrupted. In this scenario, a fresh install of the software is required.

In the more complex scenario, a major event has taken place such as a loss of an entire room or building due to flood or fire. Here, the same hardware might not be available and the surrounding environment might be disrupted or changed. The recovery process is more complex and some elements will need to be adapted or prioritized.

The following sections highlight the considerations to note when recovering the NetWorker server:

- ◆ [“Basic disaster recovery \(same host\)” on page 17](#)
- ◆ [“Advanced disaster recovery \(different host\)” on page 18](#)

### Basic disaster recovery (same host)

Recovering the NetWorker server to the same host is the simplest way to perform a disaster recovery. This base level of recovery should be planned for and in place for all NetWorker deployments.

In this disaster recovery scenario, the objective is to:

- ◆ Restore the NetWorker server as quickly as possible to the latest, last known good point before the server failed.
- ◆ Ensure that the original recovery media is available.
- ◆ Ensure that the original recovery devices are available.
- ◆ Ensure that the original environment such as SAN, IP, and storage units remain unchanged.

This is a simple recovery, if:

- ◆ An adequate bootstrap and index backups exists.
- ◆ The configuration details have not changed much and are well known or documented.
- ◆ You are able to access to the media and devices that are required for the recovery.
- ◆ The backup administrator has the appropriate skills and knowledge to perform the recovery task. NetWorker 8.1 and higher includes a command line wizard program named **nsrdr** that automates the recovery of the NetWorker server’s media database, resource files, and client file indexes. The NetWorker *Administration Guide* provides more details.

In some cases, the physical host might be subject to external issues that might prevent a disaster recovery to be performed or fully completed. This scenario might require manual adaptations to ensure that adequate or alternative connectivity is made available. This situation might not require restoring the bootstrap or client file index. To restore the server to the original state following a temporary change, you need to know the original configuration.

## Advanced disaster recovery (different host)

The recovery of a NetWorker server to a different host is more complex than performing a basic disaster recovery to the same host. The effort and skills required to recover to a different host is significantly greater than a basic disaster recover to the same host. Recovering to a different host will typically require additional information or resources coupled with the appropriate skills set to perform and complete the task.

While the loss of a building or site is less likely to occur, the effort and speed that is required to recover the NetWorker server has a direct impact on the time to restore or maintain critical business services following an incident. Business-critical services might also be affected and might require a disaster recovery or failover process that relies on the backup and recovery services that the NetWorker server provides. It is therefore essential that an advanced disaster recovery scenario is included in any disaster recovery or business continuity plan.

Although the objective is the same as for the basic disaster recover to the same host, in this situation:

- ◆ The NetWorker server hardware is likely to be different and its connectivity and configuration might be different from the original.
- ◆ Simply restoring the bootstrap and client indexes might not be as quick or as easy to perform.
- ◆ Additional changes to the configuration might also be required before the backup and recovery service is available.
- ◆ Immediate access to the original recovery media and the devices cannot be assumed.
- ◆ The environment is likely to be different so that the SAN, IP, and storage units might not match the original server.
- ◆ Additional steps might be required to make the NetWorker server available.
- ◆ The availability of adequate bootstrap and index backups is required, but these might be copies of the original save sets.
- ◆ Additional steps might be required to access the save set backups.

## Ground level preparation for NetWorker server disaster recovery

To optimize your chances for a successful disaster recovery of the NetWorker server, you must meet the following minimum requirements:

- ◆ Back up the bootstrap regularly, at least once per 12 hours.
- ◆ Back up the server OS configuration regularly.
- ◆ Back up the client file indexes for all clients. A separate, dedicated backup for all client indexes can be performed before or after a bootstrap. This step provides a comprehensive disaster recovery backup solution.
- ◆ Monitor, record, and store the status and contents of each bootstrap backup in a separate physical location from the NetWorker server.
- ◆ Use a dedicated pool for bootstrap backups.
- ◆ Clone the bootstrap backups.
- ◆ Record and maintain the connectivity and details of the SAN, IP, and all storage components.



# CHAPTER 3

## Data Storage and Devices

This chapter includes the following sections:

- ◆ Capabilities and considerations ..... 22
- ◆ NetWorker metadata storage ..... 22
- ◆ Multi-path access and failover..... 23
- ◆ Reliability and dependencies ..... 25

## Capabilities and considerations

Successful disaster planning and recovery relies on the availability of the media on which the data is stored and the availability of the devices to read that data. In some cases, the disaster might be localized and the devices and connectivity might be available. Other more serious or catastrophic incidents will impact the environment that the NetWorker server relies on. This scenario might render the devices inoperable or prevent access to devices or media.

A number of strategies can be used to cope with these scenarios and range from:

- ◆ Having multiple devices and copies of data.
- ◆ Ensuring that alternative devices, media, or paths are available within short time periods.

These recovery strategies will enable you to restore with minimal disruption, effort, and guess work.

## NetWorker metadata storage

Protecting the storage or data during its normal life can help to prevent disaster situations from occurring. These steps might also help to improve the speed or reliability of the disaster recovery.

To help to improve the speed, reliability, scalability, and performance of the backup server:

- ◆ Keep key configuration information and index data on separate LUNs to eliminate OS corruption issues and improve overall system performance.
- ◆ Host LUNs on RAID-protected or external storage systems to improve the performance, reliability, and resilience of this data.
- ◆ Ensure that you have the appropriate amount of storage.
- ◆ Ensure that the storage is protected and is performing at optimal levels.
- ◆ Consider using advanced protection technology such as replication or snapshots of this data since they offer additional protection.

## Multi-path access and failover

With any storage device that is used to store bootstrap information, consider the following:

- ◆ [“Storage devices and media” on page 23](#)
- ◆ [“Method of connectivity” on page 23](#)

### Storage devices and media

As the resilience and ease of deploying storage devices varies, so the disaster recovery strategies that are used should be changed to suit the circumstances. For example, the ability to obtain and move a single tape device is simpler than it would be for a virtual tape library (VTL), where the installation and configuration might take considerable effort and time to achieve.

For traditional tape, you can use a single tape deck that is manually loaded. It can be located next to or inside the same hardware as the physical server. In some cases, this might be an autoloader with multiple devices and an automated robotic arm that loads and unloads the media.

For other storage devices, such as VTL or disk systems, the device might be an appliance that includes CPU, memory, networking, and multiple disk units.

### Method of connectivity

The method of connectivity can vary from a simple cable for a standalone tape device, to multiple IP or SAN connections. Having the device or media available is of little use if the required connectivity is unavailable.

The availability of the following components are important aspects of disaster recovery planning:

- ◆ Spare cables
- ◆ Alternative ports and routes
- ◆ Resilient networks

### Configure devices with dual ports for multipath access

In some cases, devices can be configured with dual ports or multipath access which is transparent to the backup application and device. However, for other devices this configuration might be more difficult to configure. It is simpler to configure and make available spare ports or alternative host names and routes as a disaster recovery planning step than it is to create or configure them at the time of a disaster recovery.

Most manufacturers do not support dual path tape devices or library control ports, or they impose limitations that make these options impractical. However, you can reserve alternative ports and make available alternative or backup path connections.

## Make devices available in multiple locations

In some cases, you might be required to make devices available in multiple locations and then move the backup or direct the data to the appropriate devices. This scenario can provide a faster and more robust backup service. However, these configurations are often complex and might be difficult to configure, maintain, or troubleshoot. In these situations, it is often a choice between actively using and configuring the devices for normal use or having the devices in a standby state for only disaster recovery use.

Normal use is defined as actively using the devices in all locations during normal, non-disaster recovery operations. This can make the configuration more complex and presents operational and troubleshooting challenges. However, it does provide the benefit of being able to use the device and ensures that the device is operational at the time that it is required.

Standby use is defined as leaving the device in a standby state where it is not used in normal operations. This can simplify the configuration, but the device might be inoperative when it is required. This configuration is also a less efficient use of resources since the devices are not used during normal operations.

### Device failover

In both normal use and standby scenarios, device failover is an area that is often prone to error and can require some manual intervention. Although some of these issues are easy to resolve, they should be documented, understood, and practiced.

When planning for disaster recovery, consider the following:

- ◆ Device access paths might be different, might change, or might disappear. They can all impact the configuration and might require additional steps to correct.
- ◆ Device names should identify the location or use. This can facilitate easier troubleshooting and more reliable execution of disaster recovery procedures.
- ◆ Check the device status and availability. Devices that are not used regularly are more likely to exhibit issues at a time when they are most required.

Designing resilience into the backup service is a good practice and does not have to include idle devices. However, while this solution provides a better return of investment and increases the available capacity and performance for running backup operations, designing resilience also makes the solution more complex to configure and manage. Clustering and replication technologies are used to enhance resilience in the backup environment but also to reduce complexity.

Implementing clustering and replication technologies in your disaster recovery plan will:

- ◆ Help to manage and automate the different elements such as disk storage and network connections.
- ◆ Ensure that resources such as disk storage and network connections are available on the correct hardware.
- ◆ Ensure that the resources and configurations are appropriate for the software service that is running.

## Reliability and dependencies

The reliability of a backup and recovery service depends upon the reliability of the individual components, regardless of the chosen software, devices, and disaster recovery approach.

When designing a resilient backup and recovery service:

- ◆ Select devices that match the performance and operating expectations of the service.
  - You can use multiple devices and multiple paths to improve reliability and availability. Although this helps to eliminate single points of failure, it does not remove all of the single points of failure since no service can be completely reliable.
  - Careful design with consideration to the various disaster recovery scenarios will help to identify and eliminate the most common single points of failure.
- ◆ Consider the reliability of the expected duty cycle of the service and the components that are used.
  - Some devices cannot operate continuously, or may have limits on performance or functionality.
  - Using some devices, such as physical tape devices, for excessive periods might impact their reliability.
  - Appliances such as disk arrays, deduplication systems, or VTLs might also require maintenance periods in which backups cannot be performed or perform at a reduced speed or rate.
- ◆ Consider regular maintenance.
  - Issues might arise that will require some disruptive maintenance to the service.
  - The ability of a service and subcomponents of that service to be taken offline, failed over, or to recover automatically will also ensure that the maintenance is both performed and performed with minimal disruption to the service.
  - Software patches and updates will be required to ensure optimum performance, reliability, and support.



# CHAPTER 4

## Disaster Recovery Use Cases

This chapter includes the following sections:

- ◆ Basic disaster recovery scenario..... 28
- ◆ Basic disaster recovery considerations ..... 30
- ◆ More advanced disaster recovery considerations..... 32
- ◆ Clustered solutions ..... 34
- ◆ Backup to disk ..... 35
- ◆ Index or configuration corruption ..... 36
- ◆ Corruption or loss of SAN storage ..... 36
- ◆ Loss of one server, Data Domain system, or site ..... 37
- ◆ Replication solutions ..... 37

## Basic disaster recovery scenario

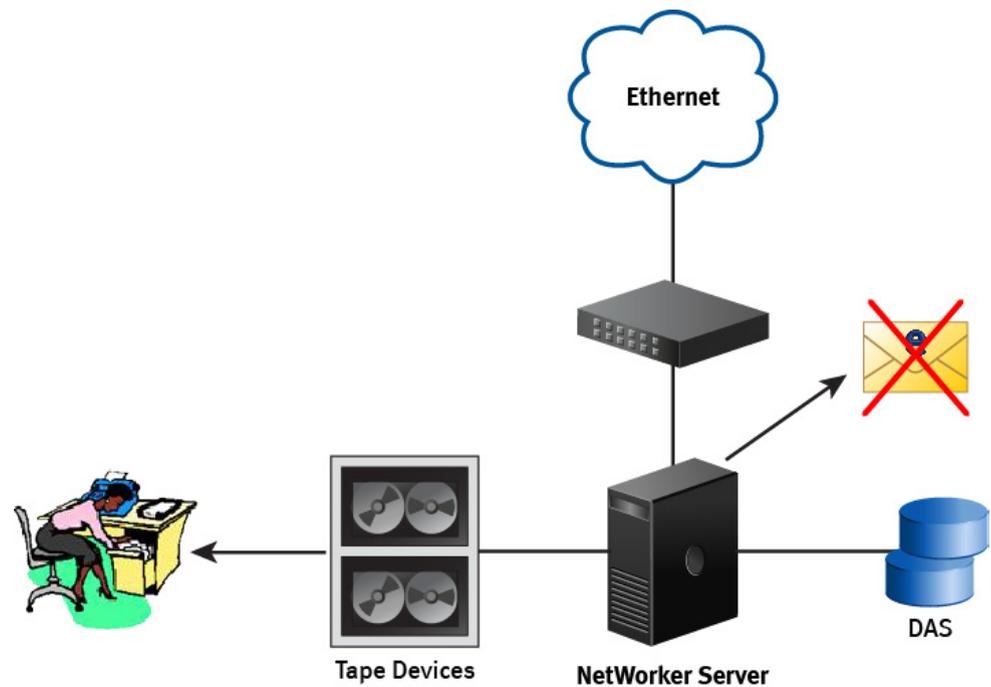
This section describes a basic NetWorker implementation to highlight important disaster recovery focus areas.

[Figure 2 on page 29](#) provides an example of a basic NetWorker solution that works well for a small office. If the server is powerful enough and the storage and connections are sized appropriately, it can protect 100 clients and a number of business systems.

In this example, the NetWorker server configuration offers very little resilience and highlights a number of disaster recovery issues that might make recovery difficult or even impossible:

- ◆ The NetWorker server:
  - Has a single ethernet connection and therefore is a single point of failure.
  - Is using internal disks and therefore is a single point of failure.
  - Has no mirroring or storage replication.
  - Is contained to a single space within a room or a data center and therefore is a single point of failure.
- ◆ The bootstrap email has not been configured and is not monitored, so the bootstrap backup emails are lost.
- ◆ The bootstrap and index backups are written to a single tape, which has three years of backups on it. The volume has not been changed or cloned and therefore is a single point of failure.
- ◆ The single copy of the bootstrap is created for disaster recovery purposes every three months and is stored in the office Administrator's desk in a different building. However, the secretary does not know the purpose of this tape and keeps it in a locked desk, in an office a few miles away from the main building and therefore is a single point of failure.

Unfortunately, in this example the management of the NetWorker server has been poor and little regard has been paid to the protection of the server.



**Figure 2** Basic NetWorker solution

In this example, the following issues might impede disaster recovery:

- ◆ Lack of resilience or redundancy in the backup environment. The NetWorker server is a single system and it uses RAID protected storage, but it is located locally through a direct attachment. This is the same for the tape devices that are located in a small autoloader near the system.
- ◆ A loss of the site might result in a loss of the tape devices, the server, and the storage. The customer in this situation only has one data room, so the use of a second site is not viable.
- ◆ The customer does not remove tapes from the site. The tapes are cycled on a monthly basis, but this is limited to a small number of monthly backups of key systems, with most tapes remaining on site.
- ◆ Bootstrap backups have been configured to run daily and are written to an index and bootstrap tape. This tape is changed, but with staff changes and an increasing workload, it is often left for several weeks. When it is changed, a new tape is labelled and the old tape is given to the office Administrator for storage. However, the office Administrator does not know the purpose of this tape and keeps it in a locked desk, in an office a few miles away from the main building.
- ◆ The bootstrap notifications have been configured to be sent by email. Unfortunately, no one monitors the email alias.
- ◆ The bootstrap notifications emails have failed for months and no one is aware of this situation.

In the event of a significant disaster, the company in this example will find it extremely difficult to recover its data and systems. Although some data is held offsite, the ability to recover it will rely on the NetWorker server and the infrastructure to be available.

While the hardware components may be quickly found, the ability to recover the NetWorker server to its previous state remains a challenge. The bootstrap tape from the office administrator's desk can be used and may only recently have been changed. The ability to use this tape depends on someone knowing where the tape is and who to ask and the office administrator being available to unlock the desk and deliver the tape. Unfortunately, without any records of the bootstraps, the entire tape will have to be scanned to rebuild the records on the new NetWorker server which is a time-consuming process. Since the tape was stored in an area that fluctuated in temperature, read errors might occur and the recovery might not be possible.

Although this situation may seem extreme, it highlights the ways in which, without careful consideration, a disaster recovery situation can have a major impact on the business.

If the following procedures were put in place, the recovery would have been much easier and faster to achieve:

- ◆ Regularly change the bootstrap tape
- ◆ Clone copies of the bootstrap and client file indexes
- ◆ Save the bootstrap notifications

Although some data is likely to have been lost forever, key data could have allowed the business to resume. Although it might not have been practical to have a second site with resilient links or remote storage, some simple measures with good management would have made the recovery situation far easier and faster.

Now that we have considered a poor example, the following examples provide information on improved levels of disaster recovery protection.

## Basic disaster recovery considerations

The following steps to improve the availability of a NetWorker server can be simple and cost effective:

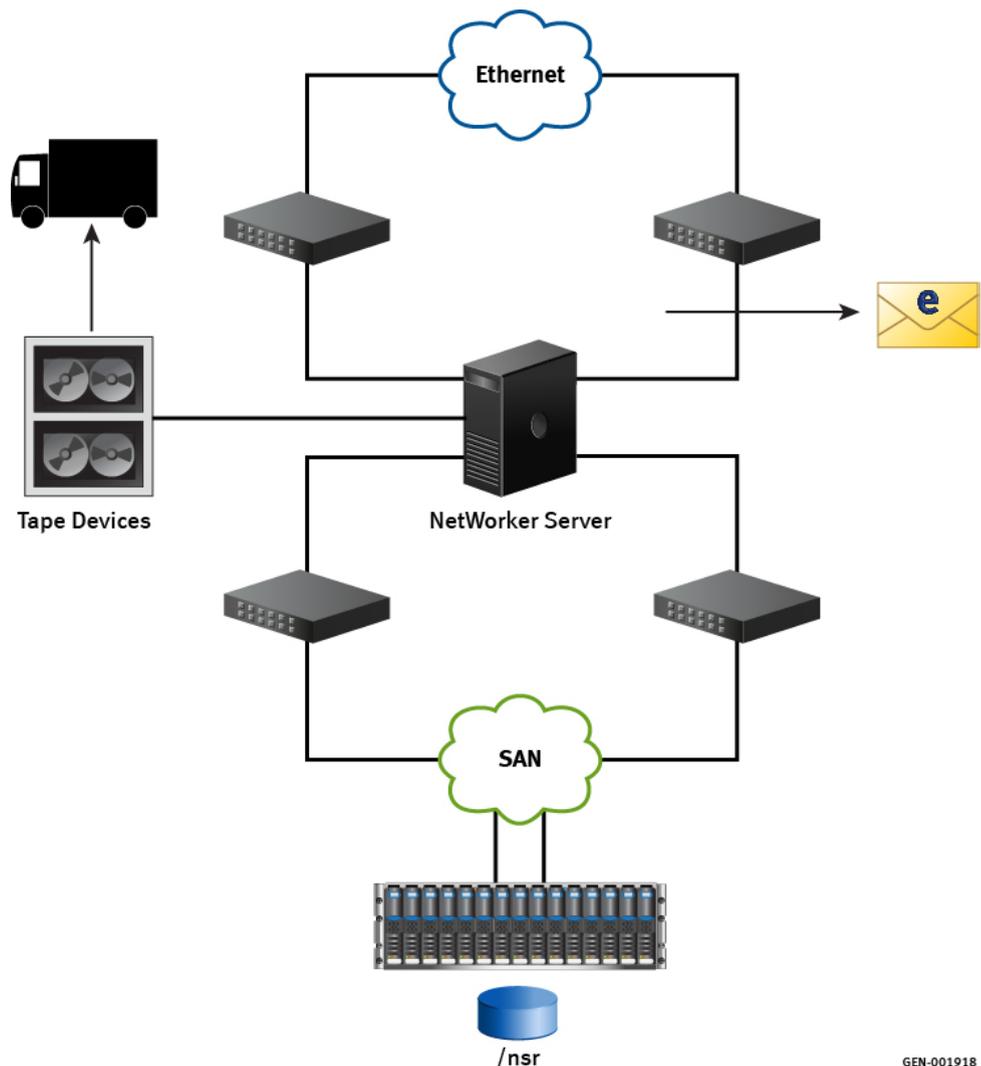
- ◆ Multiple paths for both network and storage connections are common and can help to reduce the likelihood of a failure that is due to a bad connection or failed NIC or HBA.
- ◆ Most storage systems use RAID to prevent one or more disk failures from impacting the system. These storage systems come in a range of sizes that suit any budget.

Implementing these procedures should be considered as a no-cost option, although the ongoing maintenance and management is likely to incur some expense. However, these options are very simple and cost effective and will have a big impact on the speed and ease that a disaster recovery demands.

[Figure 3 on page 31](#) highlights some of basic steps that can be used to improve the availability and disaster recovery capability of a NetWorker server. It shows a single site that is used for backup and recovery.

This example shows how a backup environment can be optimized to reduce single points of failure and improve the speed and ability of a recovery, should a disaster recovery be required:

- ◆ The bootstrap and index backups are cloned daily.
- ◆ Copies of the bootstrap and index backup clones are removed from the site and stored in a secure remote location.
- ◆ Dual Path Ethernet with automatic failover is configured and managed by a switch. This provides a single resilient IP connection.
- ◆ Email notifications are captured and stored in several locations and are available from an archive.
- ◆ The backup service and backup operations are monitored daily for nonfatal errors and warnings.
- ◆ A dual path SAN with a storage array that offers RAID protection, replication, and snapshot capabilities is used.



GEN-001918

**Figure 3** Standard disaster recovery deployment

In this example, the backup environment has been optimized to improve disaster recovery performance in the following ways:

- ◆ The same single NetWorker server is made to be more resilient and robust by adding some additional network and SAN links.
- ◆ The storage is RAID protected and has additional protection through snapshots, replication, and mirroring.
- ◆ Email notifications are sent to an alias that allows them to be accessed remotely. Email notifications are saved and monitored.
- ◆ Logs are monitored for errors so that issues can be detected early.
- ◆ Tapes are removed from site on a daily basis because there is only one site available.
- ◆ Tapes are stored in a secure and controlled location.
- ◆ Some data is cloned to ensure that multiple copies exist. This step aids in recovery and limits any exposure to media failure or loss.
- ◆ Bootstraps are cloned daily so that two copies always exist.

## More advanced disaster recovery considerations

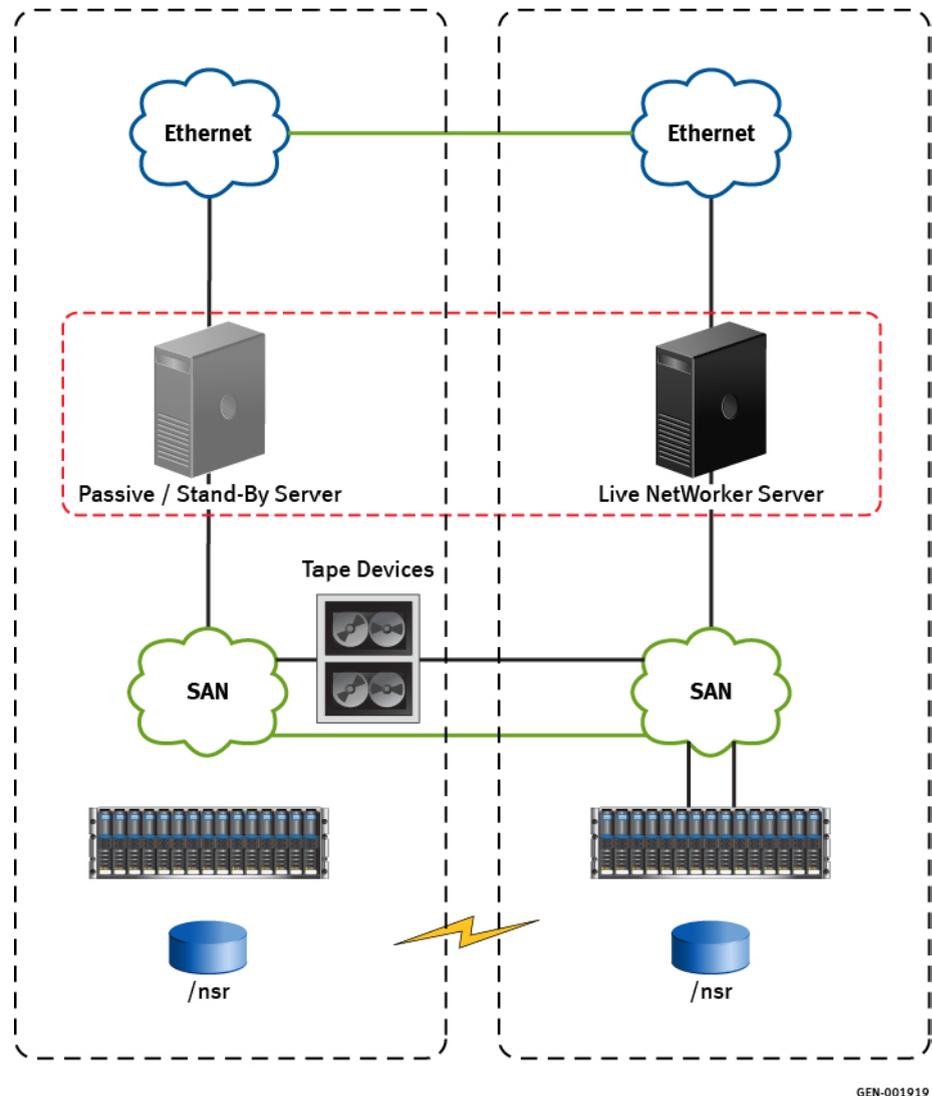
This section lists other options that build on resilience and offer higher levels of protection or recovery speed. In many cases, the recommendations from the previous section will provide adequate protection and allow the backup service to be recovered in a reliable manner and in a reasonable period of time. For others, this might not provide enough protection or might not deliver a solution that is as quick or as resilient as the business demands.

One of the best ways to improve recoverability and resilience is to introduce a second site. This practice allows the infrastructure and data to be present in two locations, which helps to mitigate the impact of an issue in a single site or with a single component within a site.

[Figure 4 on page 33](#) provides an example of a basic layout of a single NetWorker server that is configured to use two sites.

In this example:

- ◆ The same key infrastructure, such as SAN and network, is used.
- ◆ The infrastructure is configured with dual paths.
- ◆ The storage can be duplicated to provide the ability to replicate the NetWorker configuration on the second site.
- ◆ Tape devices are used to store the bootstrap and index backups. These devices are located in a different building.
- ◆ To reduce recovery time significantly, the index storage can be replicated or made available to the second site.
- ◆ To further reduce the unavailability of the backup and recovery service, add and cluster a second NetWorker server to make it highly available.



GEN-001919

**Figure 4** Single NetWorker server configured for two sites

In this example:

- ◆ One of the sites has a passive or stand-by server, which sits idle until it is required. However, a similar configuration can also allow a clustered solution where both sites have a node in which a cluster service is configured to run.
- ◆ The tape autoloader is the single point of failure in this example because it is located in one site. Although a second autoloader helps, it adds to the complexity of the configuration. Backup to disk solutions coupled with deduplication are better options in this environment.

One of the challenges with using this configuration, or any configuration in which a production backup server must be protected, is the ability to capture the system in a consistent manner. With backup and recovery operations taking place, the state of the server and the backup configuration files are in a constant state of change. The only way to reliably capture this information is to use the built-in bootstrap backup process.

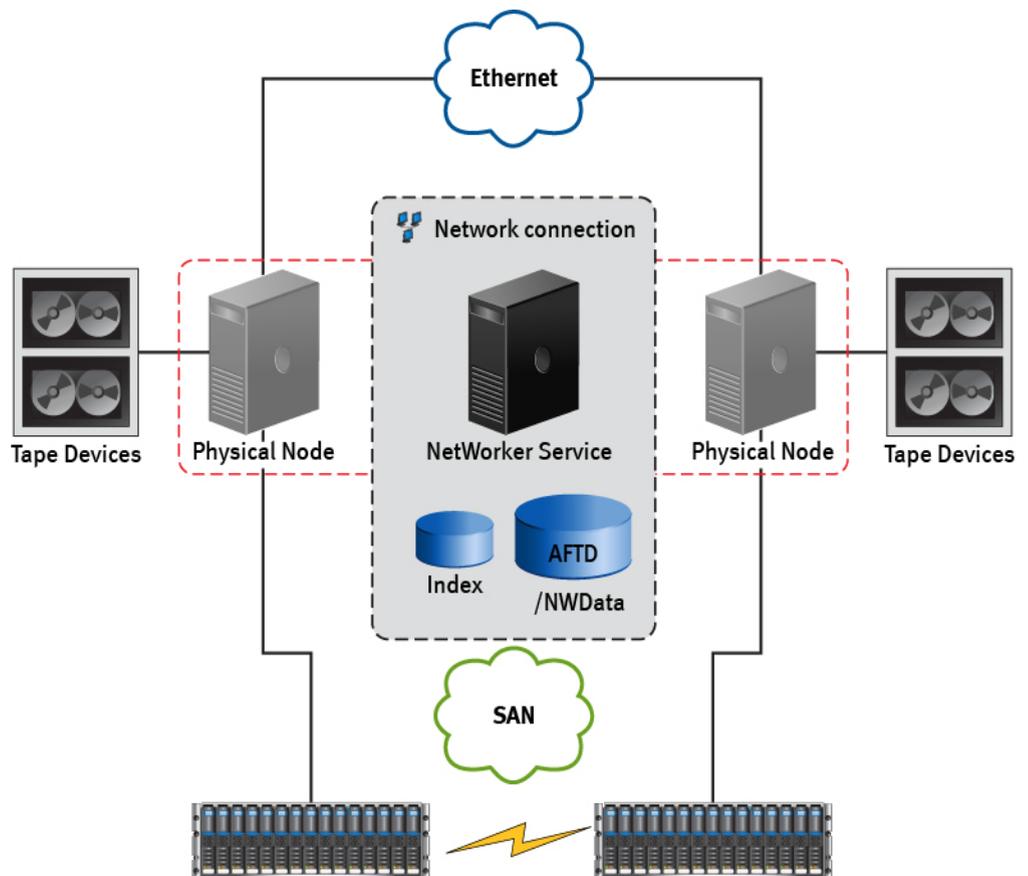
While replicating the configuration files is possible, the operation might result in a crash-consistent state. The bootstrap backup is the only method to ensure that the data is able to be recovered.

In this configuration, the SAN storage can be used to provide space for an AFTD device. These can be used for bootstrap backups and be cloned to the second site to ensure that a consistent copy is available.

## Clustered solutions

Figure 5 on page 34 illustrates how to maximize the benefit of having two sites where each site has a physical node with identical hardware.

In this example, the nodes are clustered together to provide a highly available service that is able to be active on either site. However, the tape device configuration is complex and it might be challenging to capture the system configuration in a consistent state.



GEN-001916

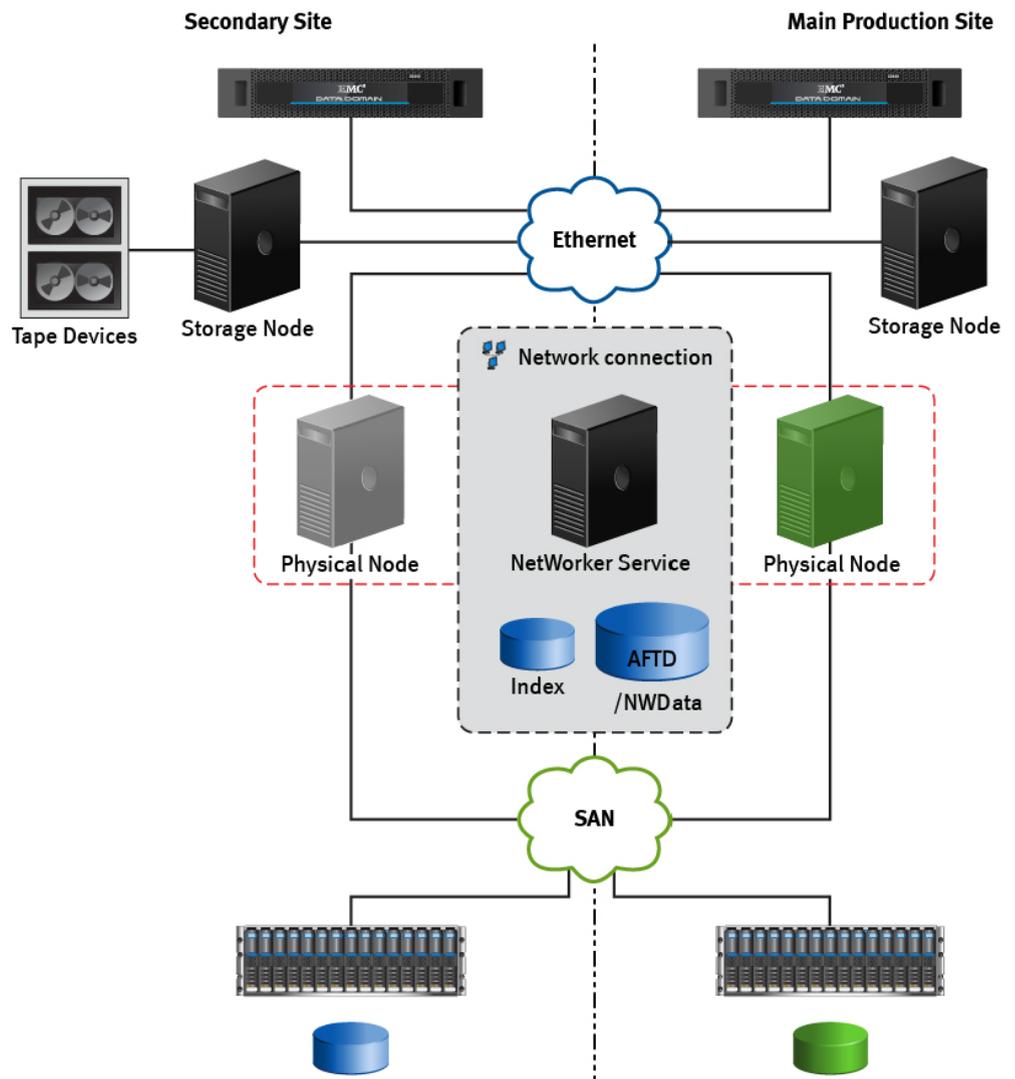
Figure 5 NetWorker servers in a clustered environment

# Backup to disk

Backup to disk based solutions simplify the configuration and help to improve the ability to capture the system configuration in a consistent state.

Although the diagram looks to be more complex, it provides a valid solution that helps to minimize the configuration complexity and helps to maintain a quick and easy failover in the event of a disaster in one of the sites.

Figure 6 on page 35 provides an example of a clustered solution that is similar to the previous example, where a two node cluster is configured to host a NetWorker service that can run on either site. Although this example might seem excessive, it meets the requirements of a number of disaster recovery scenarios.



GEN-001917

Figure 6 NetWorker server with backup to disk solution

In this example:

- ◆ The primary backup storage devices have been replaced with Data Domain systems that allow backup to disk functionality with AFTD or DD Boost devices.
- ◆ The client file index information, media database, and various configuration files are all located on SAN storage which is presented to the appropriate node through a SAN.
- ◆ The SAN storage is replicated between the sites. This step ensures that the storage is available even in the event of a site loss.
- ◆ There is still a requirement to store long term retention data on tape. This requirement is achieved by using the secondary site.
- ◆ A tape unit is available to the local storage node in this site for tape out purposes.
- ◆ A weekly copy of the bootstrap and index backups is cloned to tape and sent offsite.
- ◆ The bootstrap and index backups are cloned to both Data Domain systems to ensure that they are available in both sites.
- ◆ The bootstrap and index backups are performed regularly to the AFTD device. This step ensures that the environment can become consistent in the event of a failover and protects the backup service.

## Index or configuration corruption

Backing up the bootstrap and index backups on the AFTD will allow for rapid and immediate recovery, if the media database, or configuration areas could be corrupted because of a fault or due to human error. Backing up the bootstrap and index backups on the AFTD will allow for rapid and immediate recovery.

Consider that configuration corruption might make access to the DD Boost devices difficult, where an AFTD device is relatively easy to reconfigure.

## Corruption or loss of SAN storage

If SAN storage is lost or corrupt, you can:

- ◆ Reconfigure the DD Boost devices.
- ◆ Configure the tape device, since you will have bootstrap backups on both Data Domains systems as well as the autochanger.

## Loss of one server, Data Domain system, or site

If the server, Data Domain system, or site is lost, it will not result in the loss of backup and recovery service.

If the site or single server loss is the result of a network, power, or cooling event, then the other site should allow the backup service to remain functional after a short delay to allow for the failover to occur. The loss may be temporary, in which case additional recovery actions might not be necessary. You can restore the replication and fail over so that the main site is used once the problem resolved.

If the two sites are within a few miles of each other, you can use the tape out and offsite storage.

## Replication solutions

Replication is a term that is used differently by different vendors. Replication solutions across vendors are rarely the same and usually offer subtly different features or require different parameters to operate correctly. This section does not attempt to cover all technologies, vendors, or models, but aims to provide some background on which replication, mirroring, and snapshot features can be applied to NetWorker server disaster recovery.

When using any type of replication technology that the NetWorker server is constantly reading, changing, or updating, consider the following operations:

- ◆ Log files are updated with events and errors.
- ◆ Client file indexes are updated to reflect new backups or to remove backups that have reached their browse or retention policies.
- ◆ The media database is updated to reflect the location and state of each volume as it is being used.
- ◆ Save set information is created, deleted or changed.
- ◆ The general configuration is updated to reflect the current state of the NetWorker server, its storage nodes, devices, and clients.

All of these changes result in many IO operations on the disk. It is important to consider that anything that impacts the speed or reliability of the IO will impact the performance or reliability of the NetWorker server or the disaster recovery operation.

IO operations such as replication, mirroring, or snapshots all require some element of interception to capture the IO that has been requested. Once this IO is captured, it is simple to determine if it is a write, change, or read operation. Write operations provide the most concern as these require updates to disks as well as a confirmation that the write was successful.

If the disks are local, this activity might take very little additional time, especially with the more advanced array technologies. However, when the updates require changes on two different systems that are some distance apart, the time taken to send and acknowledge the change can be significant.

## Replication of the NetWorker Server

This section describes scenarios supported for the replication of storage used to host the NetWorker server. Replication is commonly used in conjunction with clustering or separately in scenarios where two separate hosts can act as the NetWorker server with one server being active and another ready to start NetWorker services in the event of a primary server failure. For specific performance requirements, refer to the *NetWorker Performance and Optimization Planning Guide*.

### Composite hostids for the NetWorker server

To avoid re-hosting NetWorker licenses in the event of a NetWorker server failover to a secondary host, use a composite hostid. A composite hostid is a way of creating a single ID for both the active and passive node so that single license can be used. A composite hostid is supported for any active-passive scenario, either in clustered or replicated NetWorker server environments. Setting up a composite hostid is covered in the *NetWorker Cluster Integration Guide*.

To avoid connection related issues on failover, ensure that all NetWorker clients in the datazone have both NetWorker server nodes in their servers file.

### Synchronous replication technologies

Any synchronous replication adds significant latency and due to the nature of IO produced by a NetWorker server (high amount of small random IO with 98% of write below 1Kb). Even a small increase in the disk request service time will have a significant performance impact on the NetWorker server and can even lead to server reliability issues. Therefore, only solutions that prove that synchronous replication does not introduce a significant increase in service time can be qualified by EMC.

This consideration applies to all array-based synchronous replication technologies such as EMC Symmetrix VMAX SRDF/S or EMC VNX MirrorView/S on local FC-based SAN clouds.

If IP-based links or SAN routing exist in the synchronous replication topology, an RPQ (Request for Product Qualification) is required to validate the performance impact.

Examples of configurations that require RPQs include:

- ◆ SRDF/S or MirrorView/S over IP-based replicas (FCoE, FCIP, Ethernet, and so on).
- ◆ SRDF/S or MirrorView/S over remote FC-based replicas (SAN routing, DWDM, and so on).

### Asynchronous replication technologies

Any hardware based asynchronous or nearly-synchronous replication is supported if the link that is used for replication has sufficient bandwidth so that replication is continuous and not restarted due to load. During periods when replication is being restarted, the target side is considered out-of-sync, therefore, storage failover during those times is unsupported.

This consideration applies to any asynchronous replication technologies such as SRDF/A or MirrorView/A over any type of link.

## Network attached storage

NetWorker server databases can be located on NFS (Unix/Linux) or CIFS (Windows) shares presented from a NAS filer if the connection to both the filer and NAS storage meets the performance requirements documented in *NetWorker Performance Optimization and Planning Guide*. Considerations have to be made with CIFS/NFS shares because they typically have higher latency compared to local or SAN storage.

If the CIFS/NFS share meets the minimum requirements for a NetWorker server, replication using the NAS filer's native replication technology is supported.

## Geo-replication technologies

Any array-based replication solution used for geo-clustering such as SRDF/CE should be qualified on a case-by-case basis as usage cases vary too much to provide a general qualification statement.

In general, their supportability is similar to asynchronous replication scenario, so if replication is continuous, without restarts due to link reliability or insufficient bandwidth, it is supported.

## Host-based replication technologies

Host based replication is also known as software based replication, that is, software running on a host. Any software-based remote mirroring or replication of the NetWorker server databases cannot be qualified due to the significant impact of IO latency and known incompatibilities with some filter-level drivers. This includes solutions such as Symantec Veritas VxVM remote replica or EMC RepliStor.

