

Dell EMC NetWorker

Version 19.1

Server Disaster Recovery and Availability Best Practices Guide

302-005-697

REV 01

Copyright © 1990-2019 Dell Inc. or its subsidiaries. All rights reserved.

Published May 2019

Dell 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." DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Preface		5
Chapter 1	Introduction	11
	NetWorker Server disaster recovery roadmap.....	12
	Bootstrap and indexes.....	12
	Bootstrap save set.....	13
	Client file index save set.....	13
Chapter 2	Disaster Recovery Use Cases	15
	Basic disaster recovery scenario.....	16
	Basic disaster recovery considerations.....	18
	More advanced disaster recovery considerations.....	19
	Index or configuration corruption.....	21
	Corruption or loss of SAN storage.....	21
	Loss of one server, Data Domain system, or site.....	21
	Replication solutions.....	21
	Configuring RecoverPoint for virtual machines replication	22
Chapter 3	Planning and Preparing for a Disaster Recovery	23
	Bootstrap recommendations and practices.....	24
	Backing up the NetWorker Server.....	24
	Creating a Server Backup action.....	24
	Performing a manual backup of the NetWorker server.....	28
	Gathering the key information.....	28
	How to obtain the bootstrap.....	28
	Hardware information.....	32
	Software information.....	32
	Disaster recovery scenario review.....	32
	Basic disaster recovery (same host).....	33
	Advanced disaster recovery (different host).....	33
	Ground level preparation for NetWorker Server disaster recovery....	34
	Data storage devices.....	34
	Capabilities and considerations.....	34
	NetWorker metadata storage.....	35
	Multi-path access and failover.....	35
	Basic disaster recovery scenario.....	37
	Basic disaster recovery considerations.....	39
	More advanced disaster recovery considerations.....	40
	Index or configuration corruption.....	43
	Corruption or loss of SAN storage.....	43
	Loss of one server, Data Domain system, or site.....	43
	Replication solutions.....	43
	Configuring RecoverPoint for virtual machines replication	44
Chapter 4	NetWorker Server Disaster Recovery Procedures	45

	Downloading the NetWorker software and documentation.....	46
	Information required before recovering a NetWorker Server.....	46
	Replacing the hardware, then reinstalling and upgrading the operating system.....	48
	Reinstalling the NetWorker Server software.....	48
	Opening NMC and connecting to the NetWorker Server.....	49
	Configuring NetWorker device and client resource.....	51
	Recovering critical NetWorker Server databases.....	52
	Consider the recovery options.....	53
	Recovering critical NetWorker server databases.....	54
	Recovering the NetWorker server application and user data.....	64
	Avoiding slow startups due to hostname resolution issues.....	65
	Workaround for hostname resolutions issues.....	65
	Making changes to the etc/nsswitch.conf file.....	66
Chapter 5	NMC Server Disaster Recovery Procedures	67
	Recover the NMC Server database.....	68
	Prepare for an NMC Server recovery.....	68
	Recovering the NMC Server.....	68

Preface

As part of an effort to improve product lines, periodic revisions of software and hardware are released. Therefore, all versions of the software or hardware currently in use might not support some functions that are described in this document. The product release notes provide the most up-to-date information on product features.

If a product does not function correctly or does not function as described in this document, contact a technical support professional.

Note

This document was accurate at publication time. To ensure that you are using the latest version of this document, go to the Support website <https://www.dell.com/support>.

Purpose

This document describes how to design and plan for a NetWorker disaster recovery. However, it does not provide detailed disaster recovery instructions.

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.

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
01	May 20, 2019	First release of this document for NetWorker 19.1

Related documentation

The NetWorker documentation set includes the following publications, available on the Support website:

- *NetWorker E-LAB Navigator*
Provides compatibility information, including specific software and hardware configurations that NetWorker supports. To access E-LAB Navigator, go to <https://elabnavigator.emc.com/eln/elhome>.
- *NetWorker Administration Guide*
Describes how to configure and maintain the NetWorker software.
- *NetWorker Network Data Management Protocol (NDMP) User Guide*
Describes how to use the NetWorker software to provide data protection for NDMP filers.
- *NetWorker Cluster Integration Guide*
Contains information related to configuring NetWorker software on cluster servers and clients.
- *NetWorker Installation Guide*

Provides information on how to install, uninstall, and update the NetWorker software for clients, storage nodes, and servers on all supported operating systems.

- *NetWorker Updating from a Previous Release Guide*
Describes how to update the NetWorker software from a previously installed release.
- *NetWorker Release Notes*
Contains information on new features and changes, fixed problems, known limitations, environment and system requirements for the latest NetWorker software release.
- *NetWorker Command Reference Guide*
Provides reference information for NetWorker commands and options.
- *NetWorker Data Domain Boost Integration Guide*
Provides planning and configuration information on the use of Data Domain devices for data deduplication backup and storage in a NetWorker environment.
- *NetWorker Performance Optimization Planning Guide*
Contains basic performance tuning information for NetWorker.
- *NetWorker Server Disaster Recovery and Availability Best Practices Guide*
Describes how to design, plan for, and perform a step-by-step NetWorker disaster recovery.
- *NetWorker Snapshot Management Integration Guide*
Describes the ability to catalog and manage snapshot copies of production data that are created by using mirror technologies on storage arrays.
- *NetWorker Snapshot Management for NAS Devices Integration Guide*
Describes how to catalog and manage snapshot copies of production data that are created by using replication technologies on NAS devices.
- *NetWorker Security Configuration Guide*
Provides an overview of security configuration settings available in NetWorker, secure deployment, and physical security controls needed to ensure the secure operation of the product.
- *NetWorker VMware Integration Guide*
Provides planning and configuration information on the use of VMware in a NetWorker environment.
- *NetWorker Error Message Guide*
Provides information on common NetWorker error messages.
- *NetWorker Licensing Guide*
Provides information about licensing NetWorker products and features.
- *NetWorker REST API Getting Started Guide*
Describes how to configure and use the NetWorker REST API to create programmatic interfaces to the NetWorker server.
- *NetWorker REST API Reference Guide*
Provides the NetWorker REST API specification used to create programmatic interfaces to the NetWorker server.
- *NetWorker 19.1 with CloudBoost 19.1 Integration Guide*
Describes the integration of NetWorker with CloudBoost.
- *NetWorker 19.1 with CloudBoost 19.1 Security Configuration Guide*
Provides an overview of security configuration settings available in NetWorker and Cloud Boost, secure deployment, and physical security controls needed to ensure the secure operation of the product.

- **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 the online 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, to connect to a NetWorker server to back up, recover, archive, and retrieve files over a network.

Special notice conventions that are used in this document

The following conventions are used for special notices:

NOTICE

Identifies content that warns of potential business or data loss.

Note

Contains information that is incidental, but not essential, to the topic.

Typographical conventions

The following type style conventions are used in this document:

Table 2 Style conventions

Bold	Used for interface elements that a user specifically selects or clicks, for example, names of buttons, fields, tab names, and menu paths. Also used for the name of a dialog box, page, pane, screen area with title, table label, and window.
<i>Italic</i>	Used for full titles of publications that are referenced in text.
Monospace	Used for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, file names, file name extensions, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Used for variables.
Monospace bold	Used for user input.
[]	Square brackets enclose optional values.
	Vertical line indicates alternate selections. The vertical line means or for the alternate selections.
{ }	Braces enclose content that the user must specify, such as x, y, or z.
...	Ellipses indicate non-essential information that is omitted from the example.

You can use the following resources to find more information about this product, obtain support, and provide feedback.

Where to find product documentation

- <https://www.dell.com/support>

- <https://community.emc.com>

Where to get support

The Support website <https://www.dell.com/support> provides access to product licensing, documentation, advisories, downloads, and how-to and troubleshooting information. The information can enable you to resolve a product issue before you contact Support.

To access a product-specific page:

1. Go to <https://www.dell.com/support>.
2. In the search box, type a product name, and then from the list that appears, select the product.

Knowledgebase

The Knowledgebase contains applicable solutions that you can search for either by solution number (for example, KB000xxxxxx) or by keyword.

To search the Knowledgebase:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Knowledge Base**.
3. In the search box, type either the solution number or keywords. Optionally, you can limit the search to specific products by typing a product name in the search box, and then selecting the product from the list that appears.

Live chat

To participate in a live interactive chat with a support agent:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Contact Support**.
3. On the **Contact Information** page, click the relevant support, and then proceed.

Service requests

To obtain in-depth help from Licensing, submit a service request. To submit a service request:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Service Requests**.

Note

To create a service request, you must have a valid support agreement. For details about either an account or obtaining a valid support agreement, contact a sales representative. To get the details of a service request, in the *Service Request Number* field, type the service request number, and then click the right arrow.

To review an open service request:

1. Go to <https://www.dell.com/support>.
2. On the **Support** tab, click **Service Requests**.
3. On the **Service Requests** page, under **Manage Your Service Requests**, click **View All Dell Service Requests**.

Online communities

For peer contacts, conversations, and content on product support and solutions, go to the Community Network <https://community.emc.com>. Interactively engage with customers, partners, and certified professionals online.

How to provide feedback

Feedback helps to improve the accuracy, organization, and overall quality of publications. You can send feedback to DPAD.Doc.Feedback@emc.com.

CHAPTER 1

Introduction

This chapter includes the following sections:


- [NetWorker Server disaster recovery roadmap](#)..... 12
- [Bootstrap and indexes](#)..... 12

NetWorker Server disaster recovery roadmap

This guide provides an aid to disaster recovery planning and detailed step-by-step disaster recovery instructions.

The following figure shows the high-level steps to follow when performing a disaster recovery of the NetWorker Server.

Figure 1 Disaster recovery roadmap



Bootstrap and indexes

The successful recovery of a NetWorker Server requires a current backup of key configuration information. NetWorker stores this configuration information in various directory locations on the host and this information changes when you modify the clients, devices, and volumes in the datazone.

The main backup components that protect the configuration information are the bootstrap save set and the client file index save sets.

Bootstrap save set

The NetWorker server generates a special save set called the bootstrap. 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 five components that reside on the NetWorker server:

- Media database
- Resource files
- License server files (`dpa.lic` and `licspec.properties`)
- NetWorker Authentication Service database
- Lockboxes

The Server Backup action, which is part of the Server Protection policy performs the bootstrap backup. By default, the Server Backup action also backs up all the client file indexes. The only guaranteed method to safely and consistently capture the NetWorker server configuration information is to perform a bootstrap backup. The bootstrap 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 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 that 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 that is backed up.

The client file index contains the following information for each backup save set:

- Name of each file
- File type of each file
- Save time
- Size of each file (UNIX only)
- File attributes

The client file index is not always required to recover data, but it is recommended that you back up the client file indexes, and ensure that it is available for a disaster recovery. The availability of the client file index greatly impacts the full restoration of backup and recovery services following a disaster recovery. The client file index helps you determine the time that is required to restore a NetWorker Server to a fully functional state.

Use the `nsrck` command to rebuild the client file index for a client from the index backup.

CHAPTER 2

Disaster Recovery Use Cases

This chapter includes the following sections:


- [Basic disaster recovery scenario](#)..... 16
- [Basic disaster recovery considerations](#)..... 18
- [More advanced disaster recovery considerations](#)..... 19
- [Index or configuration corruption](#)..... 21
- [Corruption or loss of SAN storage](#)..... 21
- [Loss of one server, Data Domain system, or site](#)..... 21
- [Replication solutions](#)..... 21

Basic disaster recovery scenario

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

The following figure 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 up to 100 clients and a number of business systems.

Figure 2 Example of insufficient NetWorker Server management



GEN-001915

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 email messages 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 Administrators 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.

Basic NetWorker solution

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.

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 situation 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 process 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 labeled 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 email messages have failed for months and no one is aware of this situation.

If a significant disaster occurs, the company in this example might find it extremely difficult to recover its data and systems. Although some data is held offsite, the ability to recover it relies 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 must 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.

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. The following figure highlights some of basic steps that can be used to improve the availability and disaster recovery capability of a NetWorker Server and 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.

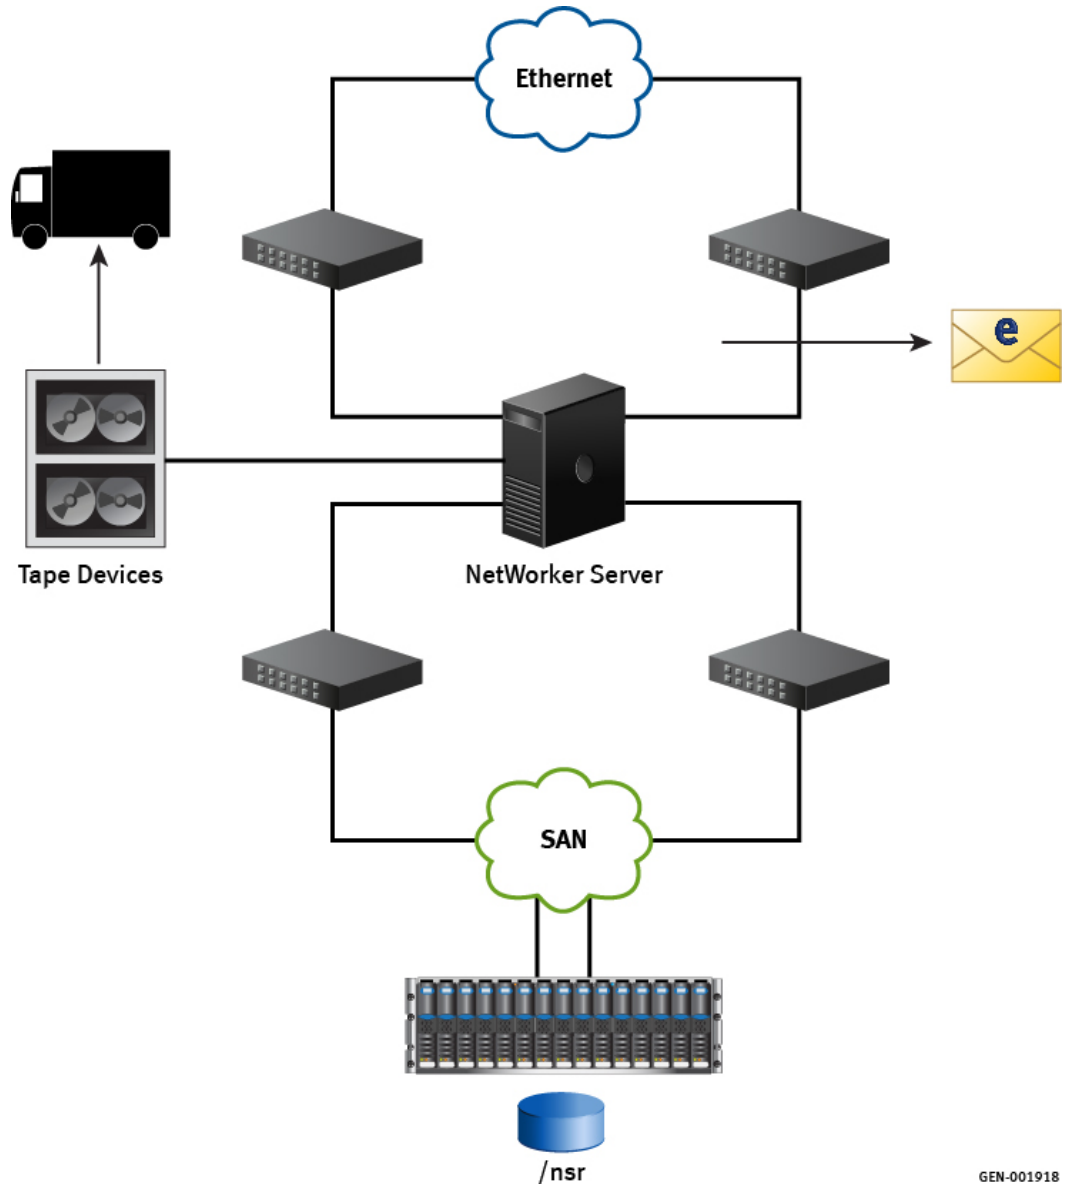
Standard disaster recovery deployment

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 steps aids in recovery and limits any exposure to media failure or loss.

- Bootstraps are cloned daily so that two copies always exist.

Figure 3 Standard disaster recovery deployment



GEN-001918

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.

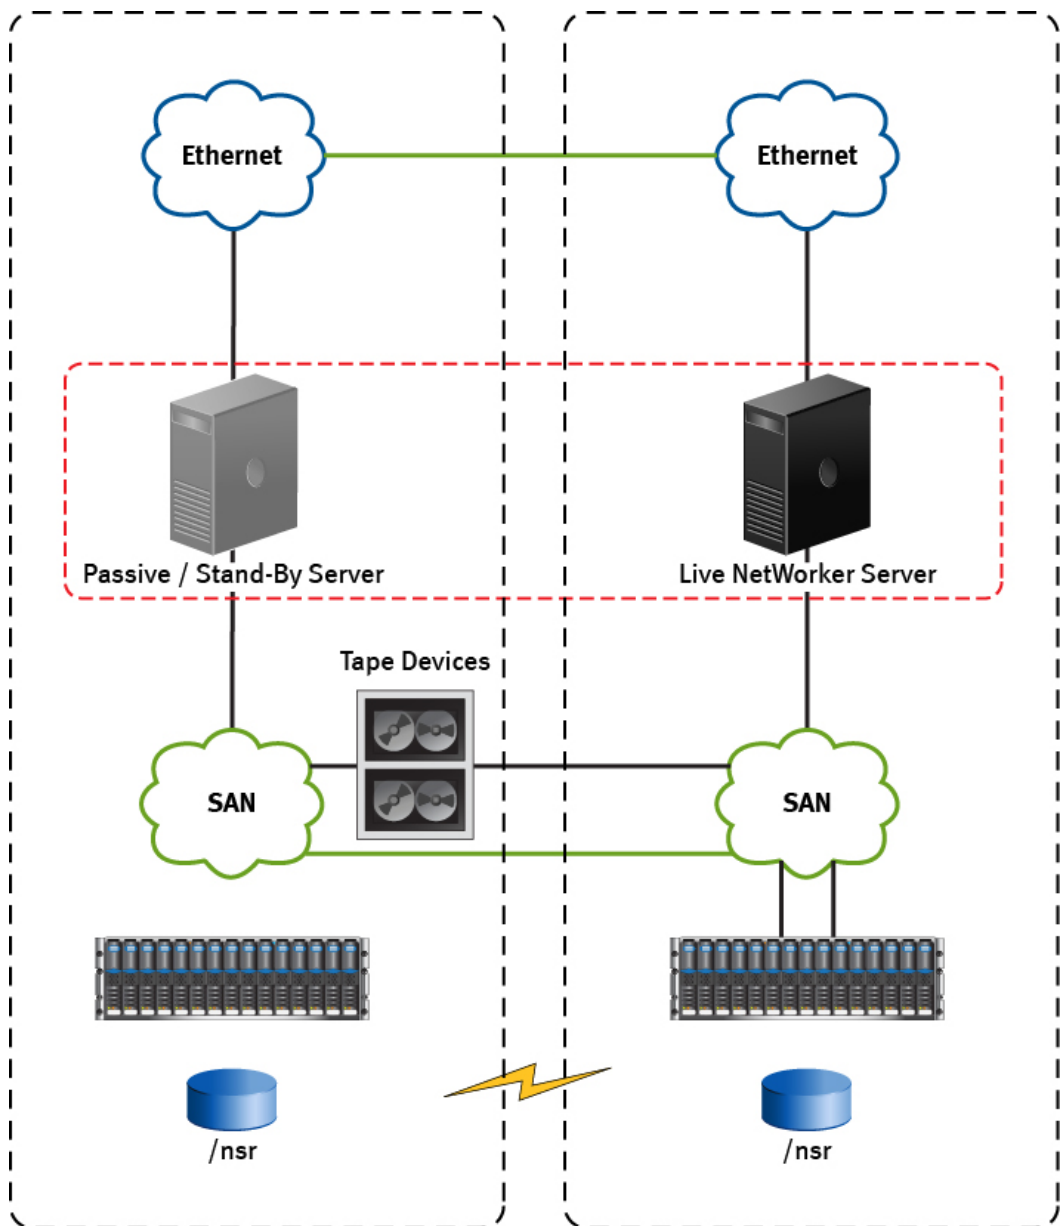
One of the best ways to improve recoverability and resilience is to introduce a second site. This practice allows the infrastructure and data to be present in two locations, which helps to mitigate the impact of an issue in a single site or with a single component within a site.

Single NetWorker Server configured for two sites

This figure provides an example of a basic layout of a single NetWorker Server that is configured to use two sites, where:

- 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.

Figure 4 Single NetWorker Server configured for two sites



GEN-001919

In this example:

- One of the sites has a passive or stand-by server, which sits idle until it is required.
- 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. 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.
- 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.

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.

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. 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 and replication solutions are rarely the same. The features that are offered can be subtly different and require different parameters to operate. This section provides some basic background

on the support and qualification of the various replication, mirroring, and snapshot features you need to consider for disaster recovery of the NetWorker Server.

When planning a replication solution, consider the potential impacts on the NetWorker Server, which is constantly at work reading, changing, and updating information:

- 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 expiration polices.
- The media database is updated to reflect the location and state of each volume used.
- Save set information is created, deleted, or changed.
- The general configuration is updated to reflect the current state of the NetWorker Server with all its storage nodes, devices, and clients.

These activities require many I/O operations on the server's disk. Any impacts on the speed and reliability of the I/O operations will impact the performance and reliability of the NetWorker Server and the disaster recovery.

Replication, mirroring, and snapshot operations all require interception and capture of any requested read, write, and change IOs that occur during the operation. Write I/Os require extra processing not only for the disk updates but to confirm that the updates are successful.

If the replication disks are local, the I/O activity might take very little time, especially with advanced array technologies. However, if the replication requires operations on systems that are separated by distance, the time required to perform and confirm the operations can have a significant impact.

The *NetWorker Performance Optimization Planning Guide* provides details on specific performance requirements.

You can validate the performance impacts and support of replication solutions by a Request for Product Qualification (RPQ), which you can submit through Professional Services.

Configuring RecoverPoint for virtual machines replication

NetWorker 18.1 and later supports replication through RecoverPoint for virtual machines.

RecoverPoint for virtual machines replicates virtual machines and their datastores to a secondary site for high availability and faster disaster recovery. Replication is managed within the vSphere Web Client using the RecoverPoint plug-in.

In order use RecoverPoint for virtual machines with NetWorker, you must enable MAC address replication for each virtual machine that will be replicated. When you are configuring a consistency group to protect the virtual machines with the Protect Volumes Wizard, select **Advanced Settings > MAC Address Replication** and deselect **Disable for local copy**.

The *RecoverPoint Administrator's Guide* provides information about configuring replication with the Protect Volumes Wizard.

CHAPTER 3

Planning and Preparing for a Disaster Recovery

This chapter includes the following sections:

- [Bootstrap recommendations and practices](#)..... 24
- [Backing up the NetWorker Server](#).....24
- [Gathering the key information](#).....28
- [Disaster recovery scenario review](#).....32
- [Data storage devices](#).....34
- [Basic disaster recovery scenario](#)..... 37
- [Basic disaster recovery considerations](#)..... 39
- [More advanced disaster recovery considerations](#).....40
- [Index or configuration corruption](#)..... 43
- [Corruption or loss of SAN storage](#).....43
- [Loss of one server, Data Domain system, or site](#)..... 43
- [Replication solutions](#)..... 43

Bootstrap recommendations and practices

To ensure that you have access to the latest bootstrap backup, perform the following tasks:

- Maintain a record of the bootstrap save set information. The NetWorker Server provides you with the ability to configure policy completion and failure notifications in the policy, workflow, and action resources. Ensure that you retain email or printed copies of the notification for the Server Backup action, which contains the following information about the bootstrap backup:
 - Backup date and time
 - Volume name and location
 - Save set ID (SSID)
 - Starting file and record number on the volume.

Note

The *NetWorker Administration Guide* provides more information about how to configure policy completion and failure notifications.

- Perform a bootstrap backup regularly, at least once every 24 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 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 dependent on client data volumes that might have inappropriate policies or protection.
- Ensure that the physical location of the backup media does not impact 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, you should maintain multiple copies of this information in other locations.

Backing up the NetWorker Server

When you install or upgrade the NetWorker Server, the installation or upgrade process creates a default Server Protection policy that backs up the NetWorker Server and the NMC Server database.

The Server Protection policy includes a server backup workflow. The server backup workflow performs a bootstrap backup of the NetWorker Server for disaster recovery purposes. The workflow is scheduled to start at 10:00 a.m. A full backup occurs on the first day of the month, and incremental backups occur the remaining days of the month. The workflow is assigned to the default Server Protection group, which contains a dynamically generated list of the Client resources for the NetWorker Server and the NMC Server.

Creating a Server Backup action

A Server Backup action performs a bootstrap backup of the NetWorker media and resource databases, and can also include the client file indexes. By default, the NetWorker server configuration contains a Server Protection policy that contains NMC server backup and Server db backup workflows. The Server db backup workflow

contain a server backup action. This section describes how to create a new server db backup action, if required.

Before you begin

Create the policy and workflow that contain the action. The Server Backup action should be the first action in the workflow.

Procedure

1. In the expanded left pane, select the policy's workflow, and then perform one of the following tasks in the right pane to start the **Policy Action** wizard:

- If the action is the first action in the workflow, select **Create a new action**.
- If the workflow has other actions, right-click an empty area of the **Actions** pane, and then select **New**.

The **Policy Action** wizard opens on the **Specify the Action Information** page.

2. From the **Action Type** list, select **Server Backup**.
3. If you create the action as part of the workflow configuration, the workflow appears automatically in the **Workflow** box and the box is dimmed.
4. Specify the order of the action in relation to other actions in the workflow:
 - If the action is part of a sequence of actions in a workflow path, in the **Previous** box, select the action that should precede this action.
 - If the action should run concurrently with an action, in the **Previous** box, select the concurrent action, and then select the **Concurrent** checkbox.
5. Specify a weekly or monthly schedule for the action:
 - To specify a schedule for each day of the week, select **Weekly by day**.
 - To specify a schedule for each day of the month, select **Monthly by day**.
6. Click the icon on each day to specify the type of backup to perform.

To perform the same type of backup on each day, select the backup type from the list and click **Make All**.

7. Click **Next**.

The **Server Backup Options** page appears.

8. From the **Destination Storage Node** list, select the storage node with the devices on which to store the backup data.
9. From the **Destination Pool** list, select the media pool in which to store the backup data.
10. From the **Retention** lists, specify the amount of time to retain the backup data.

After the retention period expires, the save set is marked as recyclable during an expiration server maintenance task.

11. Specify whether to include the client file indexes in the server backup by selecting or clearing the **Perform CFI** checkbox.

When you clear this option, the action will only backup the bootstrap.

12. Specify whether to include a bootstrap backup in the server backup by selecting or clearing the **Perform Bootstrap** checkbox.

When you clear this option, the action will only backup the client file indexes.

NOTICE

You must select either the **Perform CFI** checkbox, the **Perform Bootstrap** checkbox, or both checkboxes. Otherwise, the server backup action does not back up any data.

13. Click **Next**.

The **Specify the Advanced Options** page appears.

14. In the **Retries** field, specify the number of times that NetWorker should retry a failed probe or backup action, before NetWorker considers the action as failed. When the **Retries** value is 0, NetWorker does not retry a failed probe or backup action.
-

Note

The **Retries** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option for other actions, NetWorker ignores the values.

15. In the **Retry Delay** field, specify a delay in seconds to wait before retrying a failed probe or backup action. When the **Retry Delay** value is 0, NetWorker retries the failed probe or backup action immediately.
-

Note

The **Retry Delay** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. When you specify a value for this option in other actions, NetWorker ignores the values.

16. In the **Inactivity Timeout** field, specify the maximum number of minutes that a job run by an action can try to respond to the server.

If the job does not respond within the specified time, the server considers the job a failure and NetWorker retries the job immediately to ensure that no time is lost due to failures.

Increase the timeout value if a backup consistently stops due to inactivity. Inactivity might occur for backups of large save sets, backups of save sets with large sparse files, and incremental backups of many small static files.

Note

The **Inactivity Timeout** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types. If you specify a value for this option in other actions, NetWorker ignores the value.

17. In the **Parallelism** field, specify the maximum number of concurrent operations for the action. This is applicable if multiple rollover is implemented at an action level.

For Direct-NDMP backups, set the parallelism value to the number of available NDMP drives.

If you set the parallelism attribute to a higher value, there will not be enough drives to support all the queued backup save sets. Large save sets might fail due to the inactivity timeout limit.

When NDMP groups back up simultaneously, divide the number of drives by the number of groups. Use this value for each of the parallelism attributes.

Setting the parallelism value for the group overrides the parallelism value that is defined for the NDMP clients.

18. From the **Failure Impact** list, specify what to do when a job fails:
- To continue the workflow when there are job failures, select **Continue**.
 - To abort the current action if there is a failure with one of the jobs, but continue with subsequent actions in the workflow, select **Abort action**.

Note

The **Abort action** option applies to probe actions, and the backup actions for the Traditional and Snapshot action types.

- To abort the entire workflow if there is a failure with one of the jobs in the action, select **Abort workflow**.

Note

If any of the actions fail in the workflow, the workflow status does not appear as interrupted or cancelled. NetWorker reports the workflow status as failed.

19. From the **Send Notifications** list box, select whether to send notifications for the action:
- To use the notification configuration that is defined in the Policy resource to send the notification, select **Set at policy level**.
 - To send a notification on completion of the action, select **On Completion**.
 - To send a notification only if the action fails to complete, select **On Failure**.
20. From the **Soft Limit** list, select the amount of time after the action starts to stop the initiation of new activities. The default value of 0 (zero) indicates no amount of time.
21. From the **Hard Limit** list, select the amount of time after the action starts to begin terminating activities. The default value of 0 (zero) indicates no amount of time.
22. (Optional) Configure overrides for the task that is scheduled on a specific day.
- To specify the month, use the navigation buttons and the month list box. To specify the year, use the spin boxes. You can set an override in the following ways:
- Select the day in the calendar, which changes the action task for the specific day.
 - Use the action task list to select the task, and then perform one of the following steps:
 - To define an override that occurs on a specific day of the week, every week, select **Specified day**, and then use the lists. Click **Add Rules based override**.
 - To define an override that occurs on the last day of the calendar month, select **Last day of the month**. Click **Add Rules based override**.

Note

- You can edit or add the rules in the **Override** field.
 - To remove an override, delete the entry from the **Override** field.
-

23. Click **Next**.

The **Action Configuration Summary** page appears.

24. Review the settings that you specified for the action, and then click **Configure**.

After you finish

(Optional) Create a clone action to automatically clone the bootstrap backup when the backup completes or create an expire action.

Note

NetWorker only supports one action after the server backup action.

Performing a manual backup of the NetWorker server

Perform the following steps to manually back up the NetWorker server including the bootstrap and client indexes, from a command prompt.

Procedure

1. To perform a full backup of the backup server, type following command:

```
nsrpolicy start -p <server_protection> -w <server_backup>
```

2. To view the status of the policy, type the following command:

```
nsrpolicy monitor -p <server_protection> -w <server_backup>
```

3. To obtain the latest bootstrap information, type the following command:

```
mminfo -B
```

Keep the latest bootstrap information in a safe place for future reference in a disaster recovery.

Gathering the key information

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

How to obtain the bootstrap

Use one of the following methods to obtain information about the bootstrap:

- Review the `policy_notifications.log` file at `nsr/logs/location`, or the target destination that you configured for the policy resource notification. The "Server backup Action report" section contains information about the bootstrap and client file index backups. For example:

```
---Server backup Action report---
Policy name:Server Protection
Workflow name:Server backup
Action name:Server db backup
```

```

Action status:succeeded
Action start time:10/27/15 07:52:34
Action duration:0 hours 0 minutes 34 seconds
--- Successful Server backup Save Sets ---
4079980473/1445957561 bu-iddnserver: index:edward-soll10x64
level=1, 1 KB, 0 files
3979317182/1445957566 bu-iddnserver: index:edward-w2k12r2
level=1, 1 KB, 2 files
3945762750/1445957566 bu-iddnserver: index:bu-iddnserver
level=1, 35 MB, 43 files
3777990608/1445957584 bu-iddnserver: bootstrap level=full,
752 KB, 224 files
--- Bootstrap backup report ---
date time level ssid file record volume
10/22/15 08:43:06 full 3609789450 0 0 DDclone.001
10/22/15 08:43:06 full 3609789450 0 0 bu-iddnserver.001
10/22/15 10:00:30 full 3156809262 0 0 DDclone.001
10/22/15 10:00:30 full 3156809262 0 0 bu-iddnserver.001
10/23/15 10:00:31 full 2351589295 0 0 DDclone.001
10/23/15 10:00:31 full 2351589295 0 0 bu-iddnserver.002
10/24/15 10:00:30 full 1630255406 0 0 DDclone.001
10/24/15 10:00:30 full 1630255406 0 0 bu-iddnserver.001
10/25/15 10:00:28 full 908921516 0 0 DDclone.001
10/25/15 10:00:28 full 908921516 0 0 bu-iddnserver.002
10/26/15 10:00:30 full 204364846 0 0 DDclone.001
10/26/15 10:00:30 full 204364846 0 0 bu-iddnserver.001
10/27/15 07:53:04 full 3777990608 0 0 bu-iddnserver.002

```

In this example, the SSID/CloneID for the latest bootstrap backup is 3777990608/1445957584 on volume bu-iddnserver.002.

- If the media database is not lost and the volume list is available, use the `mminfo` command to obtain bootstrap information. For example, `mminfo -av -B -s server_name`. Where *server_name* is the hostname of the NetWorker server.
- Run the `nsrdr` command, which scans the device for the bootstrap information. For existing devices, the `nsrdr` command detects the latest bootstrap on a volume that contains the bootstrap information.

Gathering NetWorker bootstrap information

This section outlines the NetWorker bootstrap information requirements.

Is a local device available?

The NetWorker Server requires a local device resource to recover data from a bootstrap backup. In a disaster recovery situation, the resource database is lost, and you must recreate the local device to recover from the bootstrap save set.

When you recreate the device, keep the following considerations in mind:

- Do not relabel the volume when you create the device. Relabeling a volume with bootstrap backups, or any other backups, renders the data unrecoverable.
- Additional requirements for disk based devices such as AFTD.
 - Do not allow the device wizard to label the disk volume. The **Label and Mount** option on the wizard's **Device Label and Mount** window has this option selected by default. Uncheck the **Label and Mount** option.

- Specify the local path to the AFTD or DD or Tape volume in the device wizard **Select Storage Node** window. Ensure that this is the same path on which the bootstrap data is stored.

Note

If the AFTD device is created in the local server itself, data stored in the local device might be lost, if server crashes. It is recommended either to use DD/Tape device or AFTD, which is not created in the local disk. For example: AFTD created in NFS share

Is the bootstrap on a remote device, including a CloudBoost device?

NetWorker supports cloning the bootstrap backup to a local or remote device. NetWorker does not support bootstrap recoveries from a remote device. To recover the bootstrap from a cloned save set on a remote device, you must clone the save set from the remote device to a device that is local to the NetWorker Server.

To recover from a clone copy of a bootstrap backup that resides on a remote device, including a CloudBoost device, perform the following steps:

Procedure

1. Re-create the device that contains the cloned bootstrap save set on the NetWorker Server.

Note

To ensure that the wizard does not re-label the cloned device and result in data loss, in the Device Configuration Wizard, on the **Pools Configuration** page, clear the **Configure Media Pools for devices** checkbox.

2. Create a new local device on the NetWorker Server.

Note

To prevent data loss, it is recommended that you create a new AFTD device on the NetWorker Server, to which you can recover the bootstrap data.

3. Re-populate the media database with information about the cloned bootstrap save set by performing the following steps:

- a. Determine the SSID of the save set by using the `scanner` command.

For example,
`scanner -B device_name`

where *device_name* is the name of the Cloud Boost device, for example,
`rd=bu-idd-cloudboost.iddlab.local:base/bkup`

- b. Re-populate the media database with information about the cloned save set, by using the `scanner` command.

For example, `scanner -m -S SSID/CloneID device_name`

4. Mount the DD Cloud Tier device.

5. Determine the SSID/CloneID of the bootstrap backup on the DD Cloud Tier device, by using the `mminfo -B` command.
6. Clone the bootstrap save set from the DD Cloud Tier device to the Data Domain device, by using the `nsrclone` command or create a save set group.
7. Determine the SSID/CloneID of the bootstrap backup on the Data Domain device, by using the `mminfo -B` command.
8. Recover the bootstrap backup from the local device, by using the `nsrdr` command.

Is the bootstrap on a Cloud Tier device?

NetWorker supports cloning the bootstrap backup to a Cloud Tier device. NetWorker does not support bootstrap recoveries from a Cloud Tier device. To recover the bootstrap from a Cloud Tier device, you must clone the save set from the Cloud Tier device to a Data Domain device, and then recover the bootstrap backup from the Data Domain device.

To recover from a bootstrap backup that resides on a Cloud Tier device, perform the following steps:

Procedure

1. Create a new Data Domain device on the same Data Domain system and storage unit as the DD Cloud Tier device that contains the bootstrap.
2. Label and mount the new Data Domain device.
3. Re-create the DD Cloud Tier device on the NetWorker Server.

Note

Do not label the DD Cloud Tier device.

4. Re-populate the media database of the NetWorker Server with information about the save set on the DD Cloud Tier device, by performing the following steps:
 - a. Determine the SSID/CloneID of the save set by typing the `scanner -B device_name` command, where `device_name` is the name of the Cloud Boost device.

For example, `scanner -B rd=bu-idd-cloudboost.iddlab.local:base/bkup`
 - b. Re-populate the media database with information about the cloned save set, by using the `scanner` command.

For example:


```
scanner -s networker_server -m ddct_device
```

 where:
 - `networker_server` is the hostname of the NetWorker Server.
 - `ddct_device` is the name of the DD Cloud Tier device.
5. Mount the DD Cloud Tier device.
6. Determine the SSID/CloneID of the bootstrap backup on the DD Cloud Tier device, by using the `mminfo -B` command.

7. Clone the bootstrap save set from the DD Cloud Tier device to the Data Domain device, by using the `nsrclone` command or create a save set group.
8. Determine the SSID/CloneID of the bootstrap backup on the Data Domain device, by using the `mminfo -B` command.
9. Recover the bootstrap backup from the Data Domain device, by using the `nsrdr` command.

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 it 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.
- List of any Windows volume mount points and UNC paths.

Disaster recovery scenario review

You might encounter the following disaster recovery scenarios. Each scenario requires a different number of recovery steps and might be easier or more challenging to plan for or recover from.

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 you might need to adapt or prioritize some elements.

The following sections highlight the considerations to note when recovering the NetWorker Server.

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 license server files 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 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. For more information on `nsrdr`, see the *NetWorker Administration Guide*.

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.

Note

The bootstrap also contains the license file `dpa.lic`, which resides in `/nrs/lic` directory.

-
- 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, and client file indexes for all clients regularly, at least every 24 hours.
- Back up the server OS configuration regularly.
- 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.

Data storage devices

Review the following sections for information about the data storage devices that you use in for a disaster recovery.

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 estimation.

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 information outlined in the following sections.

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 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 and reduce complexity.

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

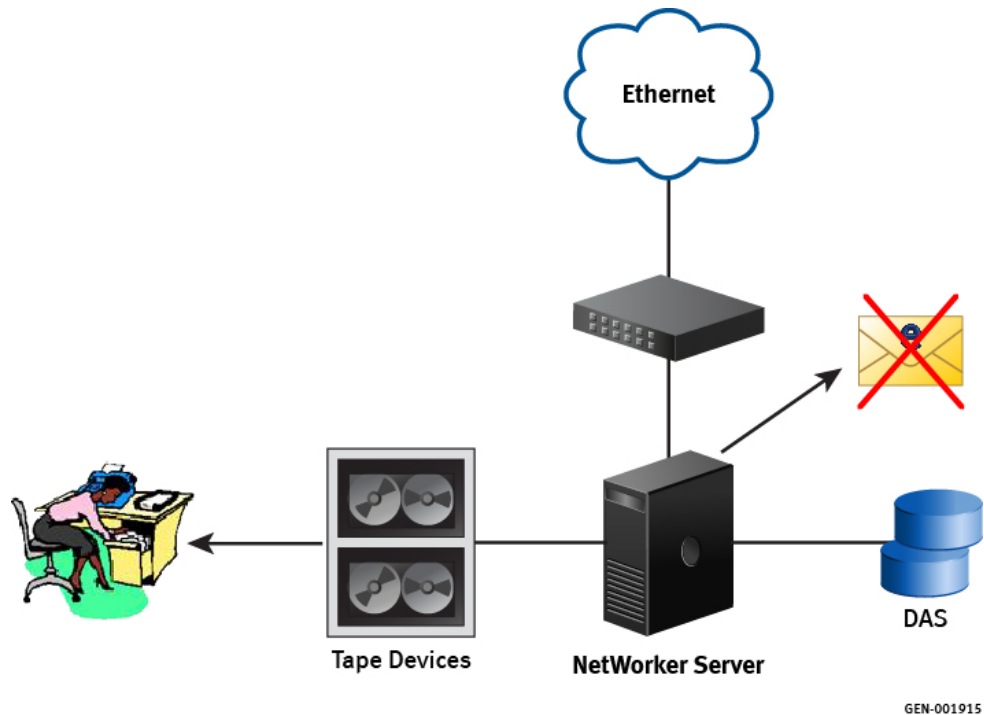
- Help to manage and automate the different elements such as disk storage and network connections.
- Ensure that resources such as disk storage and network connections are available on the correct hardware.
- Ensure that the resources and configurations are appropriate for the software service that is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

Basic disaster recovery scenario

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

The following figure 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 up to 100 clients and a number of business systems.

Figure 5 Example of insufficient NetWorker Server management



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 email messages 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.

Basic NetWorker solution

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.

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 situation 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 process 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 labeled 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 email messages have failed for months and no one is aware of this situation.

If a significant disaster occurs, the company in this example might find it extremely difficult to recover its data and systems. Although some data is held offsite, the ability to recover it relies 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 must 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.

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.

The following figure highlights some of basic steps that can be used to improve the availability and disaster recovery capability of a NetWorker Server and 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.

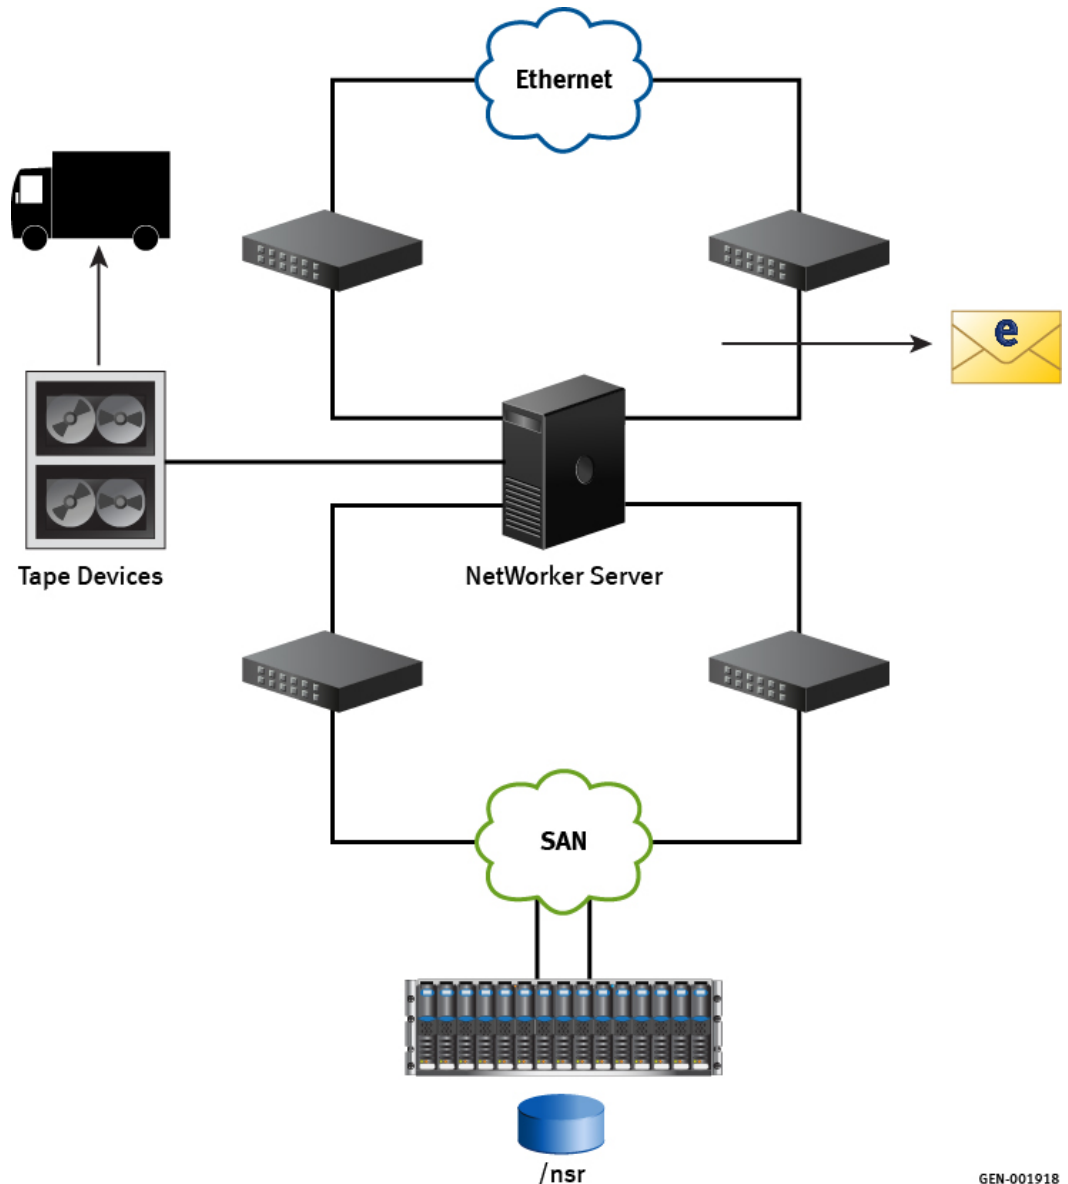
Standard disaster recovery deployment

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 steps aids in recovery and limits any exposure to media failure or loss.
- Bootstraps are cloned daily so that two copies always exist.

Figure 6 Standard disaster recovery deployment



GEN-001918

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.

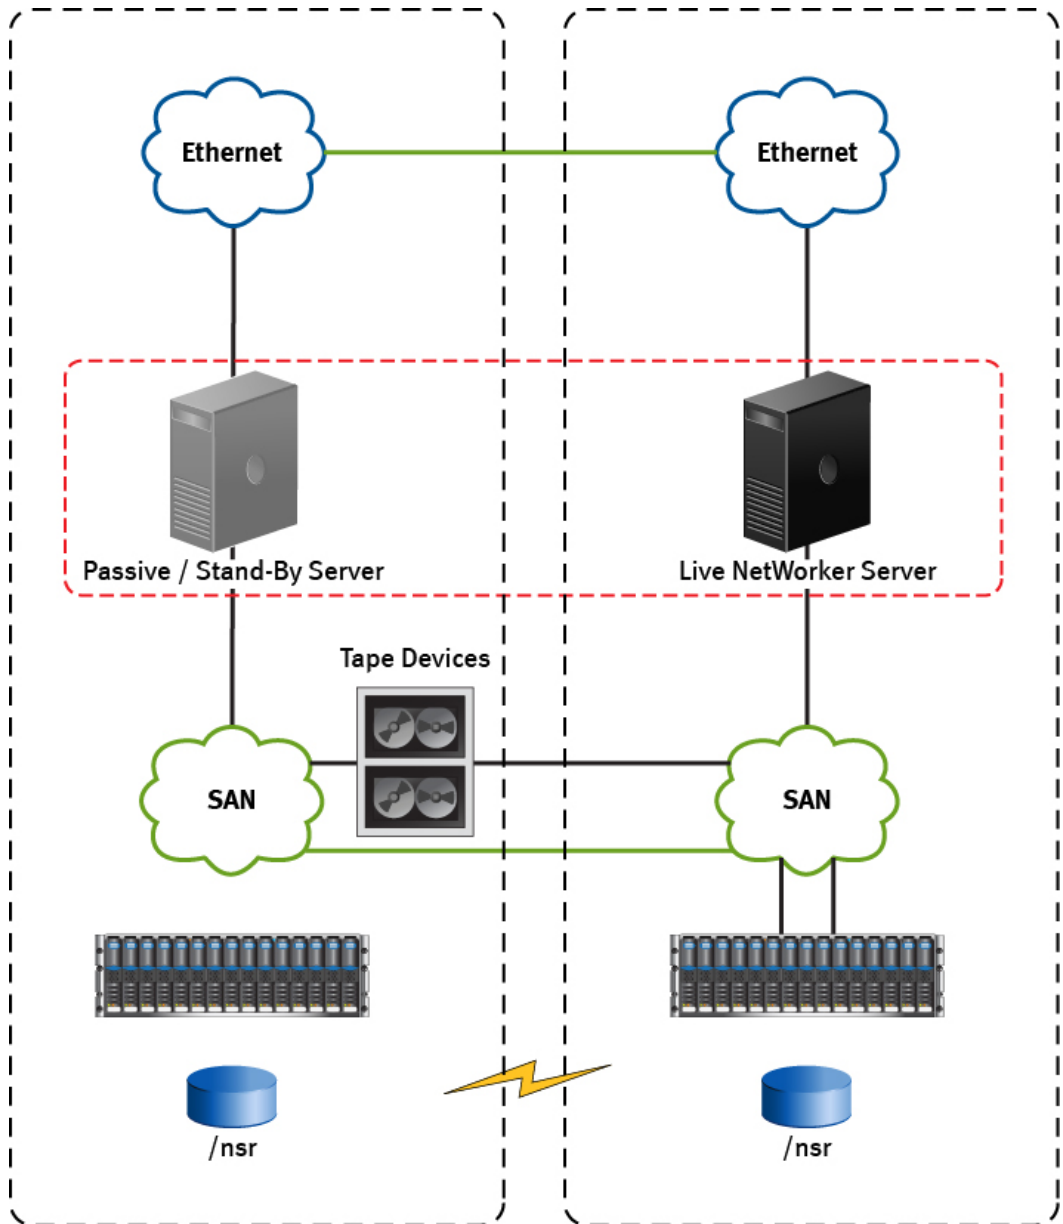
One of the best ways to improve recoverability and resilience is to introduce a second site. This practice allows the infrastructure and data to be present in two locations, which helps to mitigate the impact of an issue in a single site or with a single component within a site.

Single NetWorker Server configured for two sites

This figure provides an example of a basic layout of a single NetWorker Server that is configured to use two sites, where:

- 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.

Figure 7 Single NetWorker Server configured for two sites



GEN-001919

In this example:

- One of the sites has a passive or stand-by server, which sits idle until it is required.
- 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. 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.

- 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.

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.

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. 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 and replication solutions are rarely the same. The features that are offered can be subtly different and require different parameters to operate. This section provides some basic background on the support and qualification of the various replication, mirroring, and snapshot features you need to consider for disaster recovery of the NetWorker Server.

When planning a replication solution, consider the potential impacts on the NetWorker Server, which is constantly at work reading, changing, and updating information:

- 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 expiration polices.
- The media database is updated to reflect the location and state of each volume used.
- Save set information is created, deleted, or changed.

- The general configuration is updated to reflect the current state of the NetWorker Server with all its storage nodes, devices, and clients.

These activities require many I/O operations on the server's disk. Any impacts on the speed and reliability of the I/O operations will impact the performance and reliability of the NetWorker Server and the disaster recovery.

Replication, mirroring, and snapshot operations all require interception and capture of any requested read, write, and change IOs that occur during the operation. Write I/Os require extra processing not only for the disk updates but to confirm that the updates are successful.

If the replication disks are local, the I/O activity might take very little time, especially with advanced array technologies. However, if the replication requires operations on systems that are separated by distance, the time required to perform and confirm the operations can have a significant impact.

The *NetWorker Performance Optimization Planning Guide* provides details on specific performance requirements.

You can validate the performance impacts and support of replication solutions by a Request for Product Qualification (RPQ), which you can submit through Professional Services.

Configuring RecoverPoint for virtual machines replication

NetWorker 18.1 and later supports replication through RecoverPoint for virtual machines.

RecoverPoint for virtual machines replicates virtual machines and their datastores to a secondary site for high availability and faster disaster recovery. Replication is managed within the vSphere Web Client using the RecoverPoint plug-in.

In order use RecoverPoint for virtual machines with NetWorker, you must enable MAC address replication for each virtual machine that will be replicated. When you are configuring a consistency group to protect the virtual machines with the Protect Volumes Wizard, select **Advanced Settings > MAC Address Replication** and deselect **Disable for local copy**.

The *RecoverPoint Administrator's Guide* provides information about configuring replication with the Protect Volumes Wizard.

CHAPTER 4

NetWorker Server Disaster Recovery Procedures

This chapter contains the following sections:

- [Downloading the NetWorker software and documentation](#)..... 46
- [Information required before recovering a NetWorker Server](#)..... 46
- [Replacing the hardware, then reinstalling and upgrading the operating system](#)
..... 48
- [Reinstalling the NetWorker Server software](#)..... 48
- [Opening NMC and connecting to the NetWorker Server](#)..... 49
- [Configuring NetWorker device and client resource](#)..... 51
- [Recovering critical NetWorker Server databases](#)..... 52
- [Recovering the NetWorker server application and user data](#)..... 64
- [Avoiding slow startups due to hostname resolution issues](#)..... 65

Downloading the NetWorker software and documentation

To obtain the latest NetWorker software and documentation, perform the following steps.

Procedure

1. Review the online NetWorker documentation, such as the *NetWorker Administration Guide*, *NetWorker Installation Guide*, and *NetWorker Release Notes*, for the latest information.
2. Obtain the required NetWorker cumulative hotfix media kits that provide customers with the opportunity to install the latest version of NetWorker including important hotfixes. Cumulative builds are released approximately once a month and each build contain a rollup of the fixes in each previous build.

If additional hotfixes are required in an environment where a cumulative build is installed, hotfixes can be generated for use with the latest cumulative version. The cumulative releases for specific NetWorker versions are available at the Online Support website.

3. Open the NetWorker Cumulative Hotfix document for details regarding fixes that are in each build, knowledge base articles that are related to the fixes in each build, and download instructions.

Information required before recovering a NetWorker Server

The following information is required before recovering a NetWorker Server:

- NetWorker version and patch level
- NetWorker installation path
- NetWorker bootstrap information
- Software enabler and authorization codes
- Operating system version and patches installed
- TCP/IP properties:
 - Adapter type
 - IP address
 - Default gateway
 - Subnet mask
 - DNS server hostnames and IP addresses
- Computer properties:
 - Hostname
 - DNS domain name
 - Superuser password
- Backup or clone volume that contains the NetWorker Server's most recent:
 - Bootstrap
 - Client file indexes
 - Device and SCSI drivers

- Media device names
- Kernel configuration and location.
- For Linux, the following boot files required for starting the kernel:
 - /unix
 - /boot
 - /etc/default/boot
 - /stand/vmunix
- If you routinely move NetWorker backup media offsite for safekeeping, ensure that all necessary volumes are available so that you can avoid delays during a recovery.
- To help ensure that you are prepared to replace and reconfigure a hard drive, maintain a current record of the system information. If one or more hard drives fail, refer to the operating system documentation and hard drive vendor documentation for detailed instructions on how to replace the hard drive.
- Obtain the following information by using the appropriate operating system commands:
 - Size of the drive
 - File system volume information
 - Volume label assigned to each disk partition
 - How the disk is partitioned
 - How the disk is loaded
 - The size of the disk
 - Each logical volume (size and label)
 - Each file system

Note

- Although it does not affect NetWorker operation, It is recommended that you note any use of mirroring, Redundant Array of Independent Disk (RAID), striping, compression, or volume sets. To ensure that you can recover all the drive's data, install a new drive that is the same size or larger than the original drive.
 - NMC has separate db, which is isolated from serverdb. If you want NMC to display all the resources after `nsrdr`, you need to recover nmc db as well. Perform `recoverpsm` to recover nmc db. For example: If you add a Data Domain before performing Server DR, to get the Data Domain listed after performing Server DR, you need to perform `recoverpsm`. Refer *NMC Server Disaster Recovery Procedures* for further details.
 - The server protection policy running on NetWorker does not backup jobsdb. Post Disaster recovery, the `jobsdb` does not contain any information. All the previous status of workflows and actions are lost and the status will change to `Never Run`.
-

Replacing the hardware, then reinstalling and upgrading the operating system

To replace the hardware, then reinstall and upgrade the operating system, perform the following steps.

Procedure

1. Identify the defective or suspect hardware, and then replace the hardware as required.
2. Reinstall the same operating system version that was installed on the original host.
3. Configure the operating system by using the same IP address, hostname (shortname and FQDN) that the original host used.
4. Install any patches or upgrade the operating system to the same level as before the disaster.

Note

When using the LVM and the file systems of the root VolumeGroups are mounted, check which additional file systems outside of the root VolumeGroups are part of the backup. If VolumeGroups are imported, use the same logical volumes that were used in the backup. On Linux, ensure that all LVM device files in `/dev` are unique after the operating system installation.

5. Verify connectivity with the library and tape drives on the operating system.
6. To ensure that the library and tape drive devices are correct, run the following commands:

```
inquire
sjirdtag devname
```

where *devname* is the control port of the jukebox.

For example:

```
inquire
sjirdtag scsidev0.2.4
```

Reinstalling the NetWorker Server software

Before you can recover the NetWorker databases, reinstall the NetWorker Server software.

Perform the following steps to reinstall the NetWorker software.

Procedure

1. Reinstall the same version of the NetWorker Server software in its original location.

Ensure that you install the NetWorker client, storage node, and Authentication service packages. Installation instructions are provided in the *NetWorker Installation Guide*.

2. Run the `/opt/nsr/authc-server/scripts/authc_configure.sh` configuration script.
3. To upgrade the NetWorker Server, first recover the server to its original state, and then perform the upgrade.
In a Linux environment, you are not required to reload the license enablers if the NetWorker configuration files exist. By default, the configuration files are located in the `/nsr/res/nsrdb` directory.
4. Reinstall any NetWorker patches that were installed before the disaster.
5. Name the NetWorker Server with the same name that was used before the changes. For example, ensure that the new installation or the new server is configured with the same fully qualified name.
6. Name the devices shortname the same as it was before the changes.

Opening NMC and connecting to the NetWorker Server

Perform the following steps to open the NetWorker Management Console (NMC) and connect to the NetWorker Server.

Note

If the NMC Server and the NetWorker Server are installed on different hosts, the process owner of the NMC daemon (`gstd`) and the Administrator account for NMC must be added to the Administrators List of the NetWorker Server. In this scenario, perform all steps, otherwise skip to step 3.

Procedure

1. On the NetWorker Server, open a command prompt, and then add the NMC Server administrator account:


```
nsraddadmin -u "user=administrator, host=shortname_console_host"
nsraddadmin -u "user=administrator, host=longname_console_host"
```

 where:
 - `shortname_console_host` is the short name of the NMC Server hostname.
 - `longname_console_host` is the long name of the NMC Server hostname.
2. Add the process owner of the `gstd` daemon:


```
nsraddadmin -u "user=username, host=shortname_console_host"
nsraddadmin -u "user=SYSTEM, host=longname_console_host"
```

 where:
 - `shortname_console_host` is the short name of the NMC Server hostname.
 - `longname_console_host` is the long name of the NMC Server hostname.
 - `username` is SYSTEM on Windows hosts and root on Linux hosts.
3. Reset or delete the peer information from NMC.
4. Open a browser, and then type the following in the address bar:


```
http://<gst_server_name>://9000
```

 where `gst_server_name` is replaced with the hostname of the server where the NMC Server was installed.

Note

The NMC Server is also known as the GST server.

The **NetWorker** page displays.

5. Right-click on the **Server** in the **Enterprise** tab, and then select **Launch Application**.

The `Launching NetWorker Java` message appears and then the **NMC-Create shortcut** and **NMC Login** window appears.

6. On the **NMC Login** window:
 - a. For the username, type `administrator`.
 - b. For the password, type the password that was used for the Administrator user of the Authentication service.
 - c. Click **OK**.
7. On the **License Agreement** window, review the license text, and then click **Accept**.
8. In the **Set NetWorker License Manager Server Name** window:
 - a. Type the hostname of the License Manager server. If you did not install the License Manager software, leave this field blank.
 - b. Click **Next**.
9. In the **Set Database Backup Server** window:
 - a. In the **NetWorker Server** field, type the hostname of the NetWorker Server.
 - b. In the **Client name** field, type the hostname of the NMC Server.
 - c. Check the checkbox for **Create client resource on this server**.

This step ensures that a client resource is created to backup the Management Console database.
 - d. Click **Next**.
10. In the **Add NetWorker Servers** window:
 - a. In the **NetWorker Servers** field, type the hostname of each NetWorker Server in the environment that is required to be managed. Add one hostname per line.

The host appears in the **Enterprise** section of the Console.
 - b. Click **Finish**.

The NMC GUI opens.
11. Close the **Getting Started** window.
12. To start using the NetWorker software, click **Enterprise** from the top level toolbar.

In the left navigation pane, the top level object **Enterprise** displays. Underneath the navigation tree, each NetWorker Server hostname which was typed earlier in the window appears.
13. Click the hostname of the NetWorker Server.

On the right side pane, **Host: Server_Name** and **Managed Applications** appear. The NetWorker Administration GUI launches.

Configuring NetWorker device and client resource

Perform the following steps to configure NetWorker device and Client resources.

Procedure

1. In the NetWorker **Administration** window, on the toolbar, click **Devices**. Create, and then configure the device resources for the NetWorker Server.

If you are recovering data using an autochanger, perform the following steps:

- a. Click on the **Device** button, and in the left navigation pane, expand **Libraries**. Ensure that an autochanger resource exists.

If the autochanger resource does not exist, create the autochanger resource. The *NetWorker Administration Guide* describes how to configure an autochanger.

- b. Reset the autochanger by using the `nsrjb -vHE` command.

This command resets the autochanger, ejects backup volumes, reinitializes the element status, and checks each slot for a volume.

- If the autochanger does not support the `-E` option, initialize the element status by using the `ielem` command.
- Inventory the autochanger by using the `nsrjb -I` command. This helps to determine whether the volumes required to recover the bootstrap are located inside the autochanger.

Note

None of these volumes are in the media database. The contents of the tape cannot be viewed through NMC and the volume name appears as `-*`.

2. On the toolbar, click **Protection**, and then perform the following steps:
 - a. In the left navigation pane, click **Clients**.
 - b. Select **Diagnostic mode** under **View** tab to get the `Retention policy` attribute, in the **Client properties** window.
 - c. In the **Clients** window, right-click the client resource for the NetWorker Server, and then select **Modify Client Properties**.
 - d. On the **General** tab, in the **Retention policy** attribute, select **Decade**.

The retention policy is one month by default. This enables the recovery of all the records in the database files.

Note

If the retention policy set for the client instance of the NetWorker Server is long enough to cover all of the save sets, all of the NetWorker Server's records are recovered. However, if the retention policy set for the client is not changed and save sets for the NetWorker Server that have a retention policy greater than one month, are discarded as the default browse policy is one month.

- e. On the **Globals (1 of 2)** tab, verify that the **Aliases** attribute contains the correct host names for the NetWorker Server.

For example, for a NetWorker Server named Kingdome, the following aliases appear:

```
kingdome
kingdome.seattle.washington.com
```

Recovering critical NetWorker Server databases

Protecting a NetWorker Server including its critical databases requires careful planning and preparation. The recovery methods that are described in this section may not work if the NetWorker Server is not adequately protected.

Note

- Symbolic links to bootstrap save sets are not recovered to the default directory. Instead, they are recovered to the symbolic link's target directory. For example, If `/nsr/res` bootstrap save set is linked to `/bigres/res` directory, then the resource database is recovered to `/bigres/res` directory. Ensure that there is sufficient free space in the target directory to recover the resource database.
 - Use the `nsrdr` command to recover NetWorker 9.x databases only. To perform a roll back of the NetWorker server to an earlier version of the NetWorker software, contact Customer Service.
-

The databases that are critical to the recovery of a NetWorker Server include the bootstrap and the client file indexes.

A bootstrap includes the following components:

- Media database—Which contains the volume location of each save set.
 - Resource files—Which contains all the resources, such as NetWorker Clients and backup groups, that are defined on the NetWorker Server.
 - The NetWorker Authentication Service database.
 - Lockboxes.
-

Note

The lockbox folder in the resource directory stores confidential information, for example, Oracle client passwords and the DD Boost password, in an encrypted format. NetWorker uses this information to perform backup and recovery operations.

The client file indexes include tracking information for each file that belongs to a client's save sets. There is one client file index for each NetWorker Client.

The `nsrdr` command line program simplifies the recovery of the media bootstrap, and optionally the client file indexes for a NetWorker server. Previous releases of NetWorker required the `mmrecov` command to recover the media database and resource files, and the `nsrck` command to recover client file indexes. UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about the `nsrdr` command.

Note

The `mmrecov` command is deprecated in NetWorker 9.0 and later and replaced by the `nsrdr` command. It is recommended that you perform disaster recovery by using the `nsrdr` command.

Use the procedures in this section to recover lost or corrupted bootstrap or client file indexes (CFIs). If the server databases are not corrupted and you only want to restore expired save set entries into the client file index or the media database, use the procedures in the Recovering expired save sets topic of the *NetWorker Administration Guide*. Save sets are removed from the client file index when their browse policy time has expired. Save set entries are removed from the media database when their retention policy time expires.

The `nsrdr` command is flexible. You can run the `nsrdr` program in fully interactive mode and respond to questions or you can run the program silently with command line options. You can recover the media database, resource files, and all CFIs in one operation, or recover just one item by itself. You can recover individual CFIs or all CFIs in one operation.

To help troubleshoot issues with the wizard, the `nsrdr` command logs messages to the following location:

- On Linux, `/nsr/logs/nsrdr.log`
- On Windows, `NetWorker_install_path\nsr\logs\nsrdr.log`

Consider the recovery options

The `nsrdr` command is flexible and can be run in a variety of ways. However, the major options to consider before running the `nsrdr` command are outlined in this section.

Do you need to recover all client file indexes?

Recovering all client file indexes can take a long time. If you only need to recover the client file indexes for a limited set of clients, use the `nsrdr -I` option, for example:

```
nsrdr -c -I clientA clientB clientD
```

[Options for running the nsrdr command](#) on page 62 provides more options for recovering specific client file indexes with the `nsrdr` command.

Were save sets backed up after the last bootstrap backup?

If save sets were backed up after the last bootstrap backup, then these backup records might be overwritten after the bootstrap is recovered. This situation can only occur when a manual backup is taken. A manual backup does not trigger a bootstrap backup immediately, therefore the manual backup are not recorded in the bootstrap

until the next scheduled backup. To protect against losing save sets that were backed up after the last bootstrap backup, use the `nsrdr -N` or `nsrdr -N -F` options.

For example:

- Use the `nsrdr -N -F` command in a NetWorker datazone that contains tape devices, file type devices, and AFTDs when you only want to protect file type devices and AFTDs against loss of save sets.
- Use the `nsrdr -N` command in a NetWorker datazone that contains tape devices, file type devices, and AFTDs when you want to protect tape devices, file type devices, and AFTDs against loss of save sets.

If you know that manual backups were not taken after the last bootstrap backup or you are not concerned about losing these backups, do not use the `-N` or `-N -F` options. These options can increase the time and complexity of the recovery considerably.

Recovering critical NetWorker server databases

Use the `nsrdr` command to recover the NetWorker server databases from a command prompt.

The `nsrdr` command line options that you use to recover the database depends on the type of devices that are used in the datazone, and how you want to perform the recovery.

Setting nsrdr tuning parameters

You can specify the following tuning parameters for `nsrdr`, the NetWorker server disaster recovery command.

- You can specify the path to the NetWorker services, such as `nsrdr`, if the default path was not used during the installation.
- The number of parallel threads that can be spawned when recovering client file indexes (CFIs) for multiple NetWorker clients. The default value is 5, which means that up to five parallel threads are spawned to recover CFIs. If you are recovering many client CFIs, increasing this value can shorten the disaster recovery time.

If you do specify any of these parameters, they must be set up before running the command. You can set up these parameters by creating an ASCII plain text file, naming it `nsrdr.conf`, typing the parameter values in the file, and placing the file under the debug folder of the NetWorker installation path. Use the following procedure to set the tuning parameters:

Procedure

1. Create a text file, and then give it the name `nsrdr.conf`.

Note

Some text editors append `.txt` to the end of the file name. If this occurs, remove the `.txt` extension so that the file name is `nsrdr.conf`.

2. To specify a non-default path to the NetWorker services, add the following entry:
 - On Linux:

```
NSRDR_SERVICES_PATH = /non_default_path/nsr
```

- On Windows:

```
NSRDR_SERVICES_PATH = drive:\non_default_path\EMC
NetWorker\nsr\bin
```

where *non_default_path* is the path to the NetWorker services.

3. To specify the number of parallel threads that can be spawned when recovering CFIs for multiple clients, add the following entry:

```
NSRDR_NUM_THREADS = number
```

where *number* is a value that is greater than 1.

Note

If a value of zero (0) or a negative value is typed, a default value of 5 is automatically assigned instead.

Ensure that a space is added before and after the equals (=) sign. If you specify both tuning parameters, ensure that each value is typed on a separate line.

4. Save the `nsrdr.conf` file as a plain text file, and then place it in the following directory:
 - On Linux: `/nsr/debug/`
 - On Windows: `NW_install_path\nsr\debug`

The tuning parameters take effect the next time the `nsrdr` command is run.

Using nsrdr to perform a disaster recovery

Before you begin

Before you perform a disaster recovery of the NetWorker server databases, ensure that the authentication database directory does not contain a recovered database file that is more recent than the bootstrap that you want to recover. The name of the recovered database file is in the following format: `authcdb.h2.db.timestamp`.

The steps in this section assume that you are running the NetWorker server disaster recovery command, `nsrdr`, in fully interactive mode. Dell EMC recommends that you use the `nsrdr` command to perform a disaster recovery of the NetWorker server. To avoid data loss, Dell EMC recommends using the `-N` option. [Options for running the nsrdr command](#) provides information on additional command line options that are available for use with the `nsrdr` command.

Procedure

1. To connect to the NetWorker server and unmount all the volumes including tape, file type, advanced file type devices, and cloud volumes, use NMC.
 - a. In the **NetWorker Administration** window, click **Devices**.
 - b. Select **Devices** in the navigation tree.

The **Devices** detail table appears.
 - c. Right-click a device, and then select **Unmount**.
2. Enable the common device interface (CDI) attribute.

Note

NDMP and optical devices do not support CDI.

- a. From the **View** menu, select **Diagnostic Mode**.
 - b. Select **Devices** in the left navigation pane.
The **Devices** detail table appears.
 - c. In the **Devices** table, double-click a device.
 - d. Select the **Advanced** tab.
 - e. In the **Device Configuration** area, locate the CDI settings and select **SCSI commands**.
 - f. Stop and restart the NetWorker server services/daemons.
3. Log in to the NetWorker server as root for a Linux host, or Administrator on a Windows host.
 4. To prevent the possibility of overwriting manual backups that were taken after the last bootstrap backup, type:

```
nsrdr -N
```

When you use `-N` option, consider the following:

- For AFTD devices, you can still write to the disk, however, recover space operations are suspended until the **Scan Needed** flag is removed. A recover space operation clears the disk device of any save sets that do not have a corresponding entry in the media database.
- For tape devices, when you try to write data to a tape-based device that has newer save sets than what is recorded in the media database, a message displays that explains how to update the media database to avoid the possibility of overwriting the newer data.

If you are sure that backups were not done after the last bootstrap backup or you do not need to recover that data, omit the options.

5. At the `Do you want to continue?` prompt, type `y` for yes.
6. (Optional) If you have more than one configured device, the `Configured device output` appears with a list of configured devices. At the `What is the name of the device that you plan to use?` prompt, specify the number that is assigned to the device that contains the NetWorker server bootstrap save set.
7. At the `Enter the latest bootstrap save set id` prompt, type the save set ID of the latest bootstrap.

If you do not know the save set ID of the latest bootstrap, leave this entry blank, and then press **Enter**, and perform the following steps:

- a. At the `Do you want to scan for bootstrap save set ID on the device?` prompt, type `y` for Yes.

Note

The option to scan for a bootstrap save set ID is not supported for non-English locales. In this case, use the `scanner` command to find the bootstrap ID.

- b. **At the** `Do you want to recover the bootstrap save set with the selected ID?` prompt, type `y` for yes, to recover the bootstrap save set.
-

Note

If you are recovering from a cloud device, you are prompted to type the name of the cloud volume that contains the bootstrap save set. If you are recovering from a cloud device, you are prompted to type the datazone ID of the NetWorker server. Ensure that the datazone ID is for the NetWorker server datazone used to back up the bootstrap.

The `scanner` program is run and the bootstrap save set is recovered. Data from the bootstrap save set replaces the media database.

8. **At the** `Do you want to replace the existing NetWorker resource configuration database folder, res, with the folder being recovered?`, type `y` for yes.

The recover process performs the following tasks:

- The recovered resource database is saved to a temporary folder named `res.R`.
- The NetWorker server services are shut down because `nsrdr` cannot overwrite the resource database while these services are running.
- The recovery process replaces the existing resource database folder with the recovered resource database. The replaced folder is renamed to `res.timestamp`.

9. **At the** `Do you want to replace the existing NetWorker Authentication Service database file, authcdb.h2.db, with the recovered database file?` prompt, type `y` for yes.

10. When prompted to continue, type `y` for yes.

The NetWorker server services are restarted after the authentication database is replaced with the recovered authentication database. The replaced file is renamed to `authcdb.h2.db.timestamp`.

11. **At the** `Do you want to recover the client file indexes?`, perform one of the following tasks:

- To recover all the client file indexes:
 - a. Type `y` for yes.
 - b. Type `y` for yes again when asked to confirm the choice.

The disaster recovery operation recovers a client file index for each NetWorker client that was backed up including the client file index for the NetWorker server. The disaster recovery operation completes after all the client file indexes are recovered.

- To recover the client file index for selected clients only:
 - a. Type `n` for no.
The disaster recovery operation completes.
 - b. Re-type the `nsrdr` command with the `-c -I` options.
 - c. Provide a list of client names with each name separated by a space.
For example: `nsrdr -c -I clientA clientB clientD`
The `nsrdr` command skips the bootstrap recovery and you are prompted to complete the recovery of the specified client file indexes.
The disaster recovery operation completes after all the client file indexes that you specified are recovered.
12. Run the `/opt/nsr/authc-server/scripts/authc_configure.sh` configuration script.
 13. Open the **Administration** window in NMC, and then check that all the NetWorker Server resources appear:
 - a. Click the **Protection** icon, and then check that all resources appear as they were before recovery.
 - b. Click the **Devices** icon, and then check that all resources appear as they were before recovery.
 - c. Click the **Media** icon, and then check that all resources appear as they were before recovery.
 - d. Select **Tape Volumes** or **Disk Volumes** from the **Media** screen.
 - e. Check the mode status of the volume, **Tape Volumes**, which appears in the window on the right:
 - All volumes should have the same mode that existed before the recovery.
 - All devices that are written to should be in the appendable mode.

Remove the Scan Needed flag from volumes

Review the following sections for instructions on how to remove the Scan Needed flag from AFTD, Cloud, and Tape devices.

Removing the Scan Needed flag from AFTDs and Data Domain Devices

If you used the `nsrdr` command to set the scan needed option, all the volumes that are appendable (non read-only) and are in the recovered media database are set to Scan Needed. If you suspect that the volumes have save sets that were saved after the last bootstrap backup, you can run the `scanner -i` command to populate the recovered media database and the client file indexes with the missing save set information.

A manual save operation is the only way a save set can get backed up without triggering a save of the CFI data. If a manual backup was performed before the next scheduled backup, which always backs up the bootstrap and client file indexes, then the last CFI will not have a record of the save sets that were backed up manually.

NOTICE

The `scanner -i` command can take a very long time to complete, especially on a large disk volume. For volumes that you do not suspect have save sets that were backed up after the last bootstrap backup or for volumes where you do not need to keep these manual backups, you can skip this step and remove the Scan Needed flag from the volume.

For AFTD and DD volumes that you suspect may have save sets that were saved after the last bootstrap backup, perform the following steps:

Procedure

1. If you know the device name then skip to step 2. If you do not know the device name that corresponds to the AFTD or DD volume, use the `nsrmm` command with the `-C` option:

```
nsrmm -C
```

Output similar to the following is displayed:

```
32916:nsrmm: file disk volume_name mounted on device_name,
write enabled
```

where *device_name* is the device that corresponds to the AFTD or DD volume_name.

2. Use the `scanner` command to repopulate the CFI and media database with the save set information:

```
scanner -i device_name
```

where *device_name* is the AFTD or DD device name and not the volume name.

Note

The `scanner -i` command can take a very long time to complete, especially on a large disk volume.

3. Unmount the AFTD or DD Device.
 - a. Use NMC to connect to the NetWorker server. On the **Administration** window, select **Devices**, and then click **Devices** in the left panel.
 - b. Use NMC to connect to the NetWorker server. On the **Administration** window, select **Devices**, and then click **Devices** in the left panel.
 - c. Identify the device in the right panel that you want to unmount. Note the volume that is associated with the device.
 - d. Right-click the device, and then select **Unmount**.
 - e. Repeat for all devices that require the Scan Needed status to be removed.
4. Remove the Scan Needed status.
 - a. On the **Administration** window, select **Media**, and then click **Disk Volumes** in the left panel.
 - b. Identify the volume in the right panel that is associated with the device in the previous step.
 - c. Right-click the volume, and then select **Mark Scan Needed**.

- d. Select **Scan is NOT needed**, and then click **OK**.
 - e. Repeat for all volumes that require the Scan Needed status to be removed.
5. Mount the AFTD or DD volume.
 - a. On the **Administration** window, select the **Devices**, and then click **Devices** in the left panel.
 - b. Identify the device in the right panel that you want to mount.
 - c. Right-click the device, and then select **Mount**.
 - d. Repeat for all devices that were unmounted.
 - e. Ensure that all devices are mounted and that the Scan Needed status has been removed for the associated volumes.

Results

You can now use normal recovery procedures to recover application and user data on the NetWorker server.

NOTICE

If the recovered NetWorker server was protecting virtual cluster clients or an NMM protected virtual DAG Exchange server, the `nsrdr.log` file contains false error messages that are related to the CFI recovery of the underlying physical hosts. Using an NMM protected virtual DAG Exchange server as an example, a messages similar to the following is displayed:

```
9348:nsrck: The index recovery for 'EXCH2010-2.vll1.local'
failed.9431:nsrck: can't find index backups for
'EXCH2010-2.vll1.local' on server 'sa-wq.vll1.local'
```

You can ignore error messages that are related to the physical hosts, because NetWorker does not backup the underlying physical host in a virtual environment.

Removing the Scan Needed flag from Cloud devices

If you used the `nsrdr` command to set the scan needed option, all the volumes that are appendable (non read-only) and are in the recovered media database are set to Scan Needed. If you suspect that the volumes have save sets that were saved after the last bootstrap backup, you can run the `scanner -i` command to populate the recovered media database and the client file indexes with the missing save set information.

A manual save operation is the only way a save set can get backed up without triggering a save of the CFI data. If a manual backup was performed before the next scheduled backup, which always backs up the bootstrap and client file indexes, then the last CFI will not have a record of the save sets that were backed up manually.

NOTICE

The `scanner -i` command can take a very long time to complete, especially on a large disk volume. For volumes that you do not suspect have save sets that were backed up after the last bootstrap backup or for volumes where you do not need to keep these manual backups, you can skip this step and remove the Scan Needed flag from the volume.

For Cloud volumes that you suspect may have save sets that were saved after the last bootstrap backup, perform the following steps:

Procedure

1. If you do not know the Cloud device name that corresponds to the Cloud volume, use the `nsrmm` command with the `-C` option:

```
nsrmm -C
```

Output similar to the following is displayed:

```
32916:nsrmm: file disk volume_name mounted on device_name,
write enabled
```

where *device_name* is the device that corresponds to the Cloud *volume_name*.

2. To repopulate the CFI and media database with the save set information, use the `scanner` command:

```
scanner -i -V cloud_volume -Z datazone_ID cloud_device
```

where *datazone_ID* is the NetWorker server datazone ID if it is in a different datazone than the cloud device.

Results

You can now use normal recovery procedures to recover application and user data on the NetWorker server.

NOTICE

If the recovered NetWorker server was protecting virtual cluster clients or an NMM protected virtual DAG Exchange server, the `nsrdr.log` file contains false error messages that are related to the CFI recovery of the underlying physical hosts. Using an NMM protected virtual DAG Exchange server as an example, a messages similar to the following is displayed:

```
9348:nsrck: The index recovery for 'EXCH2010-2.v111.local'
failed.9431:nsrck: can't find index backups for
'EXCH2010-2.v111.local' on server 'sa-wq.v111.local'
```

You can ignore error messages that are related to the physical hosts, because NetWorker does not backup the underlying physical host in a virtual environment.

Removing the Scan Needed flag from tape devices

If you used the `nsrdr` command to set the scan needed option, all the volumes that are appendable (non read-only) and are in the recovered media database are set to Scan Needed. If you suspect that the volumes have save sets that were saved after the last bootstrap backup, you can run the `scanner -i` command to populate the recovered media database and the client file indexes with the missing save set information.

Note

Use `nsrdr -F` to set scan needed for all the devices.

A manual save operation is the only way a save set can get backed up without triggering a save of the CFI data. If a manual backup was performed before the next scheduled backup, which always backs up the bootstrap and client file indexes, then the last CFI will not have a record of the save sets that were backed up manually.

NOTICE

The `scanner -i` command can take a very long time to complete, especially on a large disk volume. For volumes that you do not suspect have save sets that were backed up after the last bootstrap backup or for volumes where you do not need to keep these manual backups, you can skip this step and remove the Scan Needed flag from the volume.

If you used the `-N` option with the `nsrdr` command and you try to mount a tape volume that has save sets that are newer than what is recorded in the media database, a message similar to the following appears:

```
nw_server nsrd media info: Volume volume_name has save sets unknown
to media database. Last known file number in media database is ###
and last known record number is ###. Volume volume_name must be
scanned; consider scanning from last known file and record numbers.
```

For tape volumes that you suspect may have save sets that were saved after the last bootstrap backup, perform the following steps:

Procedure

1. Make a note of the file number and record number that is displayed in the message.
2. To repopulate the CFI and media database with the save set information, use the `scanner` command:

```
scanner -f file -r record -i device
```

3. To remove the Scan Needed flag from the tape volume, use the `nsrmm` command:

```
nsrmm -o notscan volume_name
```

Results

You can now use normal recovery procedures to recover application and user data on the NetWorker server.

NOTICE

If the recovered NetWorker server was protecting virtual cluster clients or an NMM protected virtual DAG Exchange server, the `nsrdr.log` file contains false error messages that are related to the CFI recovery of the underlying physical hosts. Using an NMM protected virtual DAG Exchange server as an example, a messages similar to the following is displayed:

```
9348:nsrck: The index recovery for 'EXCH2010-2.v111.local'
failed.9431:nsrck: can't find index backups for
'EXCH2010-2.v111.local' on server 'sa-wq.v111.local'
```

You can ignore error messages that are related to the physical hosts, because NetWorker does not backup the underlying physical host in a virtual environment.

Options for running the nsrdr command

You can run the NetWorker server disaster recovery wizard command (`nsrdr`) with various command line options instead of running the wizard in fully interactive mode. The following table includes a brief description of the `nsrdr` command line options.

For a complete description of the `nsrdr` command and its options, refer to the *NetWorker Command Reference Guide* or the UNIX man pages.

Table 3 Command line options for the `nsrdr` command

Option	Description
<code>-a</code>	Runs the command line wizard in non-interactive mode. At a minimum, the <code>-B</code> and <code>-d</code> options must be specified with this command. You must specify a valid bootstrap ID with the <code>-B</code> option when running this command in non-interactive mode. Otherwise, the wizard exits as though it was canceled without providing a descriptive error message.
<code>-B bootstrap_ID</code>	The save set ID of the bootstrap to be recovered.
<code>-d device_name</code>	The device from which to recover the bootstrap.
<code>-K</code>	Use the original resource files instead of the recovered resource files.
<code>-v</code>	Verbose mode. Generates troubleshoot information.
<code>-q</code>	Quiet mode. Display only error messages.
<code>-c</code>	Recover client file indexes only. If specified with the <code>-a</code> option, you must also specify the <code>-I</code> option.
<code>-I</code> <code>-I client1 client2...</code>	Specify which CFIs (client file indexes) to recover: <ul style="list-style-type: none"> Each client name must be typed at the command prompt and separated with a space. If no client names are specified, all client file indexes are recovered. When the <code>-I</code> option is specified, ensure that it is the last option in the command string because any entries after the <code>-I</code> option are interpreted as client names.
<code>-f path/file_name</code>	Specify which CFIs to recover by using an ASCII text file. <ul style="list-style-type: none"> Place each client name on a separate line in the file. Must be used with the <code>-I</code> option. Ensure that each client name is typed correctly because there is no validation of client names.
<code>-t date/time</code>	Recover CFIs from the specified date or date and time. <ul style="list-style-type: none"> You must type a date and optionally, a time, format that is accepted by the <code>nsr_getdate</code> program. The <i>NetWorker Command Reference Guide</i> or the UNIX man pages provide more information about <code>nsr_getdate</code>.
<code>-N</code>	If tape volumes have save sets that are newer than what is recorded in the recovered bootstrap backup, they are marked as Scan Needed, to prevent the possibility of losing backed up data. For AFTD devices, this option prevents NetWorker from running recover space operations until you remove the Scan Needed flag. A recover space operation clears the disk device of any save sets that do not have a corresponding entry in the media database.

Table 3 Command line options for the nsrdr command (continued)

Option	Description
-F	This option sets the Scan Needed flag on File type devices, AFTD devices, and Cloud devices only. The <code>nsrdr</code> command will not mark tape volumes as Scan Needed. This option requires the <code>-N</code> option.

Examples

The following examples depict some common `nsrdr` commands.

- To recover the bootstrap data and selected client file indexes only, type:

```
nsrdr -I client1 client2 client3
```

 where each client name is separated with a space.
- To recover the bootstrap data and selected client file indexes by using an input file, type:

```
nsrdr -f path\file_name -I
```

 where *file_name* is an ASCII text file with one client name on each line.
- To skip the bootstrap recovery and recover selected client file indexes by using an input file, type:

```
nsrdr -c -f path\file_name -I
```

 where *file_name* is an ASCII text file with one client name on each line.
- To skip the recovery of bootstrap data and recover all client file indexes, type:

```
nsrdr -c -I
```
- To skip the recovery of bootstrap data and recover selected client file indexes, type:

```
nsrdr -c -I client1 client2
```
- To skip the recovery of bootstrap data and recover selected client file indexes from a specified date, type:

```
nsrdr -c -t date/time -I client1 client2
```

 where the *date/time* is the date and/or time from which the client file indexes are recovered. The date and time format is specified in *MM/DD/YYYY* and *hh[:mm[:ss]]* format respectively, or any date and time accepted by the `nsr_getdate` command. The *NetWorker Command Reference Guide* or the UNIX man pages provide more information about the `nsr_getdate` command.
- To run `nsrdr` in non-interactive mode and to recover the bootstrap data and all client file indexes, type:

```
nsrdr -a -B bootstrap_ID -d device -I
```

Recovering the NetWorker server application and user data

Perform the following steps to recover the application and user data that was on NetWorker server.

Procedure

1. Log in as root.
2. Load and inventory the devices.

This ensures that the NetWorker server can recognize the location of each volume.

Note

If you load a clone volume, either delete the original volume from the media database or mark the desired save set as suspect in the media database. If you are using the clone volume, it is used for the remainder of the recovery process.

3. From a command prompt, type the `recover` command.
 4. Mark all the directories or files to be recovered. Use option `a` to select or add all the files for recovery.
-

Note

Overwriting operating system files may cause unpredictable results.

5. Type `recover` to begin the recovery.

The `recover` UNIX man page and the *NetWorker Command Reference Guide* provides detailed information about how to recover the data.

Avoiding slow startups due to hostname resolution issues

Hostname resolution issues can cause the NetWorker server to become unresponsive or start very slowly. If either of the following situations apply to the environment, consider the workaround included in this section.

- The NetWorker server uses DNS, but the DNS server is not available.
- The NetWorker server cannot resolve all of its NetWorker client hosts.

Workaround for hostname resolutions issues

Perform the following steps if the NetWorker server starts slowly or becomes unresponsive due to hostname resolution issues.

Procedure

1. Disable DNS lookup for the host being recovered, and then use the local `hosts` file on the NetWorker server for the hostname resolution.
 - On Windows, the `hosts` file is located in `C:\Windows\System32\Drivers\etc\`
 - On Linux, the file is located in `/etc`

Modify the `/etc/nsswitch.conf` file to look up the `hosts` files, before DNS, as a hostname resolution method.
 2. Ensure that the `hosts` file is set up so that the full qualified domain name (FQDN) is first and is followed by the corresponding shortname.
-

Note

If the NetWorker server's client name was originally the shortname, then the shortname must be first, followed by the longname. For example, `10.10.2.5 host.emc.com`.

3. When the DNS server is available, re-enable DNS lookup.

Making changes to the `etc/nsswitch.conf` file

Perform the following steps to edit the `etc/nsswitch.conf` file:

Procedure

1. If the `hosts` line in the `/etc/nsswitch.conf` file contains the following:

```
hosts: dns files
```

This indicates that the DNS feature is used first. If the DNS servers are not available or cannot resolve the address, then modify the `hosts` line to use the `/etc/hosts` file for hostname resolution. Change the `hosts` line to the following:

```
hosts: files
```

2. On the NetWorker server, populate the local `hosts` file with the known client's valid IP address. For those clients whose IP address is unknown, use `127.0.0.1`. `127.0.0.1` is the standard IP address used for a loopback network connection. When the NetWorker server comes up, a DNS check is performed for each client. For clients that are offline or not available, the server connects to `127.0.0.1` which immediately loops back to the same host. This approach helps the NetWorker server to become available faster instead of waiting to resolve all the clients DNS lookup.
3. When the DNS server is available, re-enable DNS lookup.

Note

This method for resolving the hostnames is preferred when the NetWorker server is recovered after a DNS server failure.

When the NetWorker server is available, a valid client IP address must be updated in the local `hosts` file to perform a backup or recovery for the critical clients.

When the DNS server is available and running, remove the client details from the local `hosts` file.

CHAPTER 5

NMC Server Disaster Recovery Procedures

This chapter contains the following section:

- [Recover the NMC Server database](#).....68

Recover the NMC Server database

The NMC Server database contains management data such as report information. You can recover the NMC Server databases to the original NMC host or to a new NMC host.

Before you can perform a NMC Server database recovery, you must have an NMC Server database backup.

An NMC backup contains the following components:

- NMC database files
- NMC database credential file (`gstd_db.conf`)
- NMC lockbox files
- Legacy authentication configuration files

The "NMC Server management" chapter provides more information about NMC Server database backups.

Prepare for an NMC Server recovery

Before you recover an NMC Server, review the following information.

- If required, install the operating system on the target NMC Server.

Note

To recover an NMC Server from one host to another, both hosts must run on the same operating system.

- If required, install the NetWorker and NMC Server software on the target host. When you are prompted to specify the NetWorker Authentication Service host, specify the NetWorker Authentication Service host that was used by the source NMC Server.
- If you use a License Manager server, install and configure the License Manager software first. If you use the License Manager software and the License Manager server moves to a new host, specify the new License Manager hostname in the **Console** window.
- By default, the recover process overwrites existing NMC files. To recover to the original location, stop the NMC services by typing the following command from a prompt:

On Windows:

```
net stop gstd
```

On Linux:

```
/etc/init.d gst stop
```

Recovering the NMC Server

Perform the following steps to recover the NMC Server to the original host or a different host, from a point-in-time backup or the last backup time.

Procedure

1. Optional, to recover from an earlier backup, determine the *nsavetime* of the save set.

For example, on the NetWorker Server, type the following command:

```
mminfo -avot -q client=NMC_Server,level=full -r
client,name,savetime,nsavetime
```

where *NMC_Server* is the hostname of the NMC Server.

Output similar to the following appears:

On Windows:

```
client name date save time
bu-iddnwserver C:\Program Files\EMC NetWorker\Management
\nmcd_b_stage\pgdata 13/03/2017 1489431765
```

On Linux:

```
client name date save time
bu-iddnwserver /nsr/nmc/nmcd_b_stage 13/03/2017 1489431765
```

The *nsavetime* value appears in the last column.

2. On a target Linux NetWorker Authentication Service, set the *LD_LIBRARY_PATH* environment variable to include the postgres library path.

For example:

```
export LD_LIBRARY_PATH=NMC_Installation_dir/postgres/lib
```

where *NMC_installation_path* is */opt/lgtnm* by default.

3. Change the directory to the NMC bin directory:

On Windows the bin directory is :

```
C:\Program Files\EMC NetWorker\Management\GST\bin
```

On Linux the bin directory is:

```
/opt/lgtnm
```

4. Follow step 1-9 of *Moving the NMC Server* topic of *NetWorker Administration Guide*.
5. On the target NetWorker Authentication Service, restore the NetWorker Authentication Service backup by typing the following command:

```
recoverpsm -s NetWorker_server -c source_NMC_server /nsr/nmc/
nmcd_b_stage
or
recoverpsm -f -s NetWorker_server -c source_NMC_server -p
AES_Passphrase staging_dir -d dir_name
```

Note

If you had set datazone pass phrase during backup, then *-p AES_Passphrase* is required.

where:

- *-f* instructs the recovery operation to delete the database files that currently exist in the database directory. Do not use this option if you want to restore the database files to a different location.

- *NetWorker_server* specifies the name of the NetWorker Server.
- *source_NMC_server* specifies the name of the source NetWorker Authentication Service, when you recover the database to a different NetWorker Authentication Service host.
- *AES_Passphrase* specifies the passphrase that was used during the NMC database backup.
- *staging_dir* specifies the staging directory that was used during the backup of the database on the source NetWorker Authentication Service.
- *dir_name* specifies the directory to relocate the recovered database files. When you use this option, you must manually copy the database files from the destination directory to the database directory defined for the NetWorker Authentication Service. Ensure that you retain the same ownership and permissions on the database files and the credential files after the copy completes.

During a recovery of the NetWorker Authentication Service database, the console GUI is unavailable. Consequently, messages such as mount requests cannot be addressed from the console. Consider the following during a recovery of the NetWorker Authentication Service database:

- Monitor the daemon log files for messages. The use of the NetWorker `nsr_render_log` command can make the `daemon.raw` file more user friendly for interpretation.
 - Use the `nsrwatch` command to view messages and use commands such as `nsrjb` to address those messages. The *NetWorker Command Reference Guide* provides more information about `nsr_render_log`, `nsrwatch`, `nsrjb` and other NetWorker commands.
6. After the recovery completes, if you stopped the NMC services, start the NMC services, by typing the following command from a prompt:

On Windows:

```
net start gstd
```

On Linux:

```
/etc/init.d gst start
```