



EMC® NetWorker®
Module for Documentum®
Release 1.1
Multiplatform Version

BEST PRACTICES GUIDE
P/N 300-003-601
REV A03

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748-9103
1-508-435-1000
www.EMC.com



Copyright © 2004 - 2007 EMC Corporation. All rights reserved.

Published March, 2007

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

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

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

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Preface	vii
----------------------	-----

Chapter 1 Best Practices Guide

1. Objectives	1-2
2. Overview of NMD	1-2
2.1. Software features	1-3
2.2. Consistent repository backup	1-4
2.3. Backups with and without the SnapImage Module	1-7
3. Software components	1-11
3.1. Internal components of NMD	1-11
3.2. External modules	1-13
3.3. How NMD interfaces with external modules	1-14
3.4. Supported repository configurations	1-19
4. Prerequisites and checklist	1-20
4.1. SnapImage prerequisites	1-21
4.2. NMD prerequisites	1-22
4.3. NetWorker hardware and software prerequisites	1-22
4.4. NetWorker software prerequisites	1-23
4.5. NetWorker server sizing	1-24
5. Installation overview	1-27
5.1. NetWorker	1-27
5.2. Optional SnapImage 1.6.1 on UNIX	1-28
5.3. Optional SnapImage 2.0.2 and NDMP TapeServer 2.0.2 on Windows	1-28
5.4. NetWorker Module for Oracle, SQL Server, or Sybase (optional)	1-28
5.5. NMD 1.1	1-29
6. Configuration overview	1-29
6.1. NetWorker Client resource	1-29

6.2. NetWorker Device resource	1-30
6.3. NetWorker Group resource	1-30
6.4. NetWorker Server resource	1-31
6.5. Optional SnapImage 1.6.1 on UNIX	1-31
6.6. Optional SnapImage 2.0.2 and NDMP TapeServer 2.0.2 on Windows	1-32
6.7. NMD configuration file	1-32
6.8. Remote database server	1-40
6.9. Remote full-text index server	1-41
7. Overview of quick usage and verification	1-42
7.1. Ad hoc NMD backup	1-44
7.2. Scheduled NMD backup	1-46
7.3. Backup index query	1-46
7.4. Change of directory ownership and permissions during a restore	1-48
7.5. Content file restore	1-49
7.6 Database and log restore	1-50
7.7. Other component restore	1-51
8. Troubleshooting tips	1-53
8.1. How to verify the NMD release information	1-53
8.2. How to obtain debug information	1-54
8.3. NMD troubleshooting scenarios	1-55
8.4. SnapImage issues	1-59
8.5. Consideration for recovering the repository	1-62
8.6. Troubleshooting tools	1-63
9. Optimization recommendations	1-63
9.1. Parallelism for restores	1-63
9.2. SnapImage optimization	1-63
9.3. Other optimization	1-63
10. Special use case recommendations	1-64
10.1. Backups	1-64
10.2. Restores	1-67
10.3. UNIX symbolic links and mount points	1-67
11. Reconfiguration	1-68
11.1 Repository	1-68
11.2 SnapImage	1-68
12. Related documentation	1-69

Appendix A NMD Consulting Pre-site Qualifier (CPSQ)

Appendix B NetWorker Best Practices

B.1. Solution methodology: An industry best practice	B-2
--	-----

B.2. What is in a datazone?	B-2
B.2.1. Datazone requirements examples	B-2
B.2.2. Datazone conclusions	B-3
B.3. Single datazone	B-3
B.3.1. LAN clients, single datazone	B-3
B.3.2. LAN network, single datazone.....	B-4
B.3.3. Client optimization, single datazone	B-4
B.3.4. Host optimization, single datazone	B-4
B.3.5. NDMP optimization, single datazone	B-5
B.3.6. Server configuration, single datazone	B-6
B.3.7. Client options, single datazone	B-6
B.3.8. Storage nodes, single datazone	B-7
B.3.9. Storage node design, single datazone	B-7
B.3.10. Serverless design, single datazone.....	B-8
B.3.11. Single datazone review	B-8
B.4. Network best practices	B-8
B.4.1. Consideration.....	B-9
B.5. SAN best practices	B-9
B.5.1. Consideration: Recommended power up sequence.	B-9
B.6. Business best practices	B-10
B.6.1. Considerations	B-10
B.7. NetWorker best practices	B-10
B.7.1. Considerations	B-10
B.8. NetWorker 7.x best practices	B-11
B.8.1. File type device configuration on NetApp FAS960 (December 16, 2003).....	B-11

As part of an effort to improve and enhance the performance and capabilities of its product line, EMC from time to time releases revisions of its hardware and software. Therefore, some functions described in this EMC NetWorker Module for EMC Documentum Best Practices Guide may not be supported by all revisions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.

If this product does not function properly, or does not function as described in this guide, go to <http://Powerlink.EMC.com>.

Audience

This guide is part of the NetWorker Module for Documentum (NMD) documentation set, and is intended for use by professional services consultants and the Documentum administrators who are responsible for maintaining the Documentum Content Server and database server backup and restore system. Operators who monitor the Documentum data backups may also find this best practices guide useful.

Readers of this guide are expected to be familiar with the following:

- ◆ Documentum and relevant database terminology and concepts, especially those related to backups and restores.
- ◆ Backup and restore procedures on a Documentum Server.
- ◆ Disaster recovery procedures on a Documentum Server.

Organization This guide is organized as shown in the following table.

Chapter	Describes
Chapter 1, "Best Practices Guide"	Best practices for installing, configuring, and using NMD 1.1 on a Documentum Server host with a supported UNIX or Microsoft Windows operating system.
Appendix A, "NMD Consulting Pre-site Qualifier (CPSQ)"	NMD Consulting Pre-site Qualifier (CPSQ) form to be completed in order to assist EMC in planning for an NMD project initiation.
Appendix B, "NetWorker Best Practices"	Best practices for using the NetWorker software that works with the NMD software solution.

Related documentation

Documentation related to the use of this product can be found at the EMC website, <http://Powerlink.EMC.com>, including:

- ◆ The NetWorker Module for Documentum release 1.1 documentation set:
 - Administration guide
 - Installation guide
 - Release notes
 - Best practices guide
 - Command reference guide
- ◆ The NetWorker documentation set:
 - Administration guide
 - Installation guide
 - Release notes
 - Command reference guide
 - Disaster recovery guide
- ◆ Appropriate versions of the following documentation sets:
 - NetWorker Module for Microsoft SQL Server
 - NetWorker Module for Oracle
 - NetWorker Module for Sybase
 - NetWorker SnapImage Module
- ◆ Other EMC documentation:
 - Software compatibility guide
 - UNIX man pages

The following additional documentation may be useful:

- ◆ Documentum Content Server documentation
- ◆ Appropriate database (Microsoft SQL Server, Oracle, or Sybase) backup and recovery documentation

Conventions used in this guide

EMC uses the following conventions for notes and cautions.

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



CAUTION

A caution contains information essential to avoid data loss or damage to the system or equipment. The caution may apply to hardware or software.



IMPORTANT

An important contains information essential to operation of the software. The important may apply only to software.

Typographical conventions

EMC uses the following style conventions in this guide:

normal font

In running text:

- Interface elements (for example, button names, dialog box names) outside of procedures
- Items that user selects outside of procedures
- Java classes and interface names
- Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, filenames, functions, menu names, utilities
- Pathnames, URLs, filenames, directory names, computer names, links, groups, service keys, file systems, environment variables (for example, command-line and text), notifications

Bold

In procedures:

- Names of dialog boxes, buttons, icons, menus, fields
- Selections from the user interface, including menu items and field entries
- Key names
- Window names

In running text:

- Command names, daemons, options, programs, processes, notifications, system calls, man pages, services, applications, utilities, kernels

Italic

Used for:

- Full publications titles referenced in text
- Unique word usage in text

Bold Italic

Anything requiring extra emphasis

Courier	Used for: <ul style="list-style-type: none"> • System output • Filenames • Complete paths • Command-line entries • URLs
Courier, bold	Used for: <ul style="list-style-type: none"> • User entry • Options in command-line syntax
<i>Courier, italic</i>	Used for: <ul style="list-style-type: none"> • Arguments used in examples of command-line syntax • Variables in examples of screen or file output • Variables in path names
<i>Courier, bold, italic</i>	Variables used in a command-line sample
< >	Angle brackets enclose parameter or variable values supplied by the user
[]	Square brackets indicate optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces indicate content that you must specify (that is, x or y or z)
...	Ellipses indicate nonessential information omitted from the example

Where to get help

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

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

<http://Powerlink.EMC.com>

Technical support — For technical support, go to EMC Customer Service on Powerlink. To open a service request through Powerlink, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

Your comments Comments and suggestions about our product documentation are always welcome.

To provide feedback:

1. Go to:

<http://Powerlink.EMC.com>

2. Click the **Feedback** link.

This best practices guide describes the best practices for installing, configuring, and using the EMC® NetWorker® Module for EMC Documentum® (NMD) release 1.1 on a Documentum Server host with a supported UNIX or Microsoft Windows operating system.

Note: The *Software Compatibility Guide* on the EMC website at <http://Powerlink.EMC.com> provides details on the specific operating system versions that NMD supports.

The NMD software is distributed in a downloadable file, available from the EMC website.

This guide includes the following sections:

- ◆ 1. Objectives 1-2
- ◆ 2. Overview of NMD 1-2
- ◆ 3. Software components 1-11
- ◆ 4. Prerequisites and checklist..... 1-20
- ◆ 5. Installation overview 1-27
- ◆ 6. Configuration overview 1-29
- ◆ 7. Overview of quick usage and verification 1-42
- ◆ 8. Troubleshooting tips 1-53
- ◆ 9. Optimization recommendations 1-63
- ◆ 10. Special use case recommendations 1-64
- ◆ 11. Reconfiguration 1-68
- ◆ 12. Related documentation 1-69

1. Objectives

This best practices guide supplements the following documents:

- ◆ *NetWorker Module for Documentum, Release 1.1, Multiplatform Version, Installation Guide*
- ◆ *NetWorker Module for Documentum, Release 1.1, Multiplatform Version, Administration Guide*
- ◆ *NetWorker Module for Documentum, Release 1.1, Multiplatform Version, Release Notes*

This guide provides a quick summary of key installation and administration steps and basic NMD concepts. It also provides information on special configuration settings, optimization tuning, use case handling, and answers to general questions about the software.

2. Overview of NMD

NMD software works with the following software components to create a comprehensive Documentum data protection system:

- ◆ Documentum Content Server
- ◆ Documentum database server (Microsoft SQL Server, Oracle, or Sybase)
- ◆ (Optional) Documentum full-text index server (Documentum 5.3.x *only*)
- ◆ NetWorker server
- ◆ NetWorker client
- ◆ NetWorker storage node
- ◆ (Optional) NetWorker Module for the appropriate database:
 - NetWorker Module for Microsoft SQL Server
 - NetWorker Module for Oracle
 - NetWorker Module for Sybase
- ◆ (Optional) NetWorker SnapImage™ Module
- ◆ (Optional) NDMP TapeServer on Windows *only*

[“2.3. Backups with and without the SnapImage Module”](#) on page 1-7 provides information on using the NetWorker SnapImage Module for snapshot backups and restores of storage areas.

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides more information on the NetWorker software components.

2.1. Software features

NMD enables the Documentum repository administrator to perform streamlined backups and restores of one or more physical components of an online Documentum Content Server Document Base (docbase or repository). A *consistent* repository backup image is produced during the backups.

Note: The Documentum 5.3.x term *repository* is synonymous with the Documentum 5.2.x term *docbase*.

The physical components of a repository are the database, one or more file content storage areas, and zero or more full-text indexes per storage area.

- ◆ The database data can be under the control of an Oracle DBMS, Microsoft SQL DBMS, or Sybase DBMS.
- ◆ The storage area is typically a file system subdirectory hierarchy of possibly millions of files, which make up the actual contents of user documents.
- ◆ The full-text index is optional and consists of a much smaller number of files.

Note: NMD considers the installation and configuration files that make up a Documentum installation to be a pseudo-component of the repository.

NMD provides the following features:

- ◆ A consistent full repository backup, as defined in [“2.2. Consistent repository backup”](#) on page 1-4.
- ◆ Both ad hoc (manual or unscheduled) and automated scheduled backups of the following repository components:
 - Full-text indexes
 - Database datafiles for supported Oracle, Microsoft SQL Server, or Sybase databases
 - Database log files (if applicable)

- Storage areas
- Installation and configuration files
- ◆ Either regular (nonsnapshot) or snapshot backups of the storage areas.
- ◆ Browsing of the NMD backup records.

Note: Browsing of the backup records is recommended prior to restore.

- ◆ Restores of the following components from existing backups:
 - Storage areas
 - Content files and their associated renditions
 - Full-text indexes
 - Installation and configuration files

Note: Restores of database and database log backups must be performed by using the appropriate database NetWorker Module or user-created scripts, depending on how the backups were performed.

- ◆ Separate command-line interface programs for performing the supported backup, index browsing, and restore operations.

2.2. Consistent repository backup

A Documentum repository backup is consistent when the following conditions are met:

1. The backup image of the storage area is frozen within a few minutes at the end of the database backup. The backup image of the storage area will *not* contain any documents created after the image was frozen. This condition limits the number of orphaned documents that exist on restore of the database and storage area backup images; the only orphaned documents are those created within a few minutes.

Note: Orphaned documents can be purged following a restore. However, the purge requires clean-up processing time and can fragment the file system, causing subsequent problems with document access performance.

NMD achieves this condition by producing a frozen backup image of the storage area within a few minutes at the end of the database backup. When the database backup is restored, the restore must be to the point-in-time at the *end* of the database backup.



IMPORTANT

Do not back up the database by using logical exports/dumps or methods that place the sync point-in-time at the *start* of the database backup. NMD supports database backup methods that enable restores to the point-in-time at the *end* of the backup only. For example, NMD supports the appropriate NetWorker Module for database backups.

2. The backup image of the database refers only to content files and other objects that exist in the backup image of the storage area. This condition prevents errors and crashes when Documentum users access documents.

NMD achieves this condition by backing up the database component *before* the storage area component.

3. The backup image of the full-text index refers only to documents that exist in the backup image of the database. This condition prevents errors and crashes during Documentum user index searching.

NMD achieves this condition by backing up the full-text index component *before* the database component.

4. All objects recorded as indexed in the backup image of the database are those referred to by the backup image of the full-text index. This condition prevents documents from missing any required full-text index indexing after the restore of a repository backup.

Note: The database marks a document as having been indexed only *after* the document has been referenced by the full-text index as a result of a full-text index update job.

NMD achieves this condition by disabling all the system update processing of full-text indexes while the full-text index, database, and storage area components are backed up.



IMPORTANT

NMD temporarily disables the standard Documentum jobs for updating full-text indexes, as part of ensuring the consistency of NMD backups. Any other jobs for updating full-text indexes should *not* be active during NMD backups.

To achieve the preceding conditions for consistency, NMD performs the following during a repository backup:

1. Quiesces the full-text index.
2. Quiesces the storage area.
3. Backs up the full-text index.
4. Quiesces the database.
5. Backs up the database.
6. Unquiesces the database.
7. Backs up the storage area.
8. Unquiesces the storage area.
9. Unquiesces the full-text index.

Notes:

- ◆ Quiesce means to flush buffers and disable updates, or start the online write transaction logging.
- ◆ Unquiesce means to re-enable updates, or end the online write transaction logging.
- ◆ After the repository components are restored, run a job to update the full-text index, and a filesCAN job.

2.3. Backups with and without the SnapImage Module

NMD uses the SnapImage Module for snapshot backups of the storage area components. The following two limitations apply to the SnapImage Module:

- ◆ SnapImage can only back up an entire file system. If you are interested in only parts or directories under a file system, SnapImage cannot back up only those parts of the file system.
- ◆ SnapImage fast restore mode is fully destructive and is not used by NMD.

Note: If storage areas share a file system with other applications, consider the possible case where the restore of a SnapImage backup of the storage areas might require a destructive restore. [“8.4.5. Potential problem with the restore of storage areas that contain numerous files” on page 1-60](#) provides more information.

The following sections provide considerations for performing backups with and without the SnapImage Module:

- ◆ [“2.3.1. Recommendations for backups without SnapImage” on page 1-7](#)
- ◆ [“2.3.2. Recommendations for backups with SnapImage” on page 1-8](#)
- ◆ [“2.3.3. Caveats for the use of SnapImage” on page 1-10](#)

2.3.1. Recommendations for backups without SnapImage

Consider performing backups *without* the SnapImage Module if one or more of the following approximate conditions apply:

- ◆ The total number of content files for all the storage areas on each file system or drive is fewer than approximately 2 million, and the average content file size is equal to or greater than approximately 150 KB.
- ◆ The storage areas on each operating system file system occupy less than 75 percent of the file system’s total used space. That is, there is a significant amount of non-storage area data on the storage area file systems.

- ◆ The storage areas on each operating system file system occupy less than 75 percent of the overall number of files on the file system. That is, there is a significant amount of non-storage area data on the storage area file systems.
- ◆ The storage area file systems have significant disk write activity (for example, system or application logging) that cannot be redirected to a non-storage area file system.
- ◆ On Windows, the Oracle control and online redo logs are on storage area file systems and cannot be relocated for business reasons.
- ◆ On Windows, the Documentum Server's default port 10000 cannot be changed for business reasons.
- ◆ You do not have the extra memory and disk space for SnapImage file history work areas, caches, and general processing. Furthermore, for consistent repository backups, you do not have as many SnapImage caches as there are storage area file systems. (Storage area file system backups must all be in parallel to minimize the number of orphaned content files written to the storage area backup image.)
- ◆ Your system cannot risk SnapImage modifying the operating system file system stack to insert its disk write interceptor driver.
- ◆ You do not have expertise to maintain your SnapImage environment and/or you are uncomfortable with any one or more of the several caveats listed in the SnapImage guides.
- ◆ Your current Documentum platform (for example, operating system version, file system, or volume manager software, disk hardware) is not supported by SnapImage and you cannot change the platform.
- ◆ You intend to regularly back up a storage area by itself, and there is lots of other data on that storage area's file system.

2.3.2. Recommendations for backups with SnapImage

Consider performing backups *with* the SnapImage Module if the following approximate conditions apply:

- ◆ SnapImage supports the current Documentum Server platform (operating system version, file system or volume manager software, disk hardware).

- ◆ The total number of content files for all storage areas on each file system or drive is greater than approximately 2 million, and the average content file size is less than approximately 150 KB.

Note: The fewer the number of files and/or the greater the average file size, the less is the density (number of files per GB) of the file system. SnapImage is only recommended for high-density file systems.

- ◆ The storage areas on each file system make up more than 75 percent of the file system's total used space. There is little non-storage area data on the storage area file systems.
- ◆ The storage areas on each file system make up more than 75 percent of the overall number of files on the file system. There is little non-storage area data on the storage area file systems.
- ◆ Extra memory and disk space is available for SnapImage file history work areas, caches, and general processing.
- ◆ For consistent repository backups, there are as many SnapImage caches as storage area file systems.

Note: Storage area file system backups must all be performed in parallel to minimize the number of disconnected or orphaned content files written to the storage area backup image.

- ◆ The storage file systems have little disk write activity (for example, system or application logging) that cannot be redirected to a non-storage area file system.
- ◆ On Windows, Oracle control files and online redo logs on a storage area file system can be relocated.
- ◆ On Windows, the Documentum Server's default port 10000 can be changed.
- ◆ It is acceptable that SnapImage modifies the operating system file system stack to insert its disk write interceptor driver.
- ◆ You have expertise to maintain the SnapImage environment, and you accept the caveats listed in the SnapImage documentation.

The following EMC documentation provides more information on the supported releases of the SnapImage Module:

- ◆ *NetWorker Module for Documentum Release Notes*
- ◆ *Software Compatibility Guide*

2.3.3. Caveats for the use of SnapImage

Consider the following additional overhead that occurs with the use of the SnapImage Module for NMD backups:

- ◆ A SnapImage backup requires the entire file system's used blocks be backed up. If all the storage areas on a given file system form 50 percent of all files and use 50 percent of all used space, SnapImage must still back up the other 50 percent of disk space containing the unwanted files. Additionally, NMD requires that SnapImage writes file index entries for these unwanted files. Altogether, the unwanted overhead can be significant, and in some situations can lead to worse performance than for a backup without SnapImage.
- ◆ Although a SnapImage backup can achieve a synchronization point-in-time of the target file system within a few seconds or minutes of invocation, the snapshot image still must be synchronously written to a tape or disk target, which can take hours. The SnapImage backup session is still active during this time. The time to synchronize and create a snapshot image of the file system may depend on background disk write activity or other noise that occurs during the backup. Blocks that form part of the snap image but are modified and have not yet been copied to the backup device will have their pre-modification images cached for pending backup. This image caching overhead must be allowed for when specifying the size of the disk cache for SnapImage. The overhead is necessary to enable restores to the point in time at the start of the backup.
- ◆ SnapImage 2.0.2 and earlier can allow only as many concurrent file system backup sessions as there are SnapImage caches available. If there is only one cache available but multiple storage area file systems, then NMD with SnapImage sequentially snaps and backs up the storage areas in groups, one file system at a time. This violates the consistency between the database backup image and storage area backup image.
- ◆ NMD must use SnapImage file-by-file (FBF) restore for individual storage area content files. NMD must also use SnapImage FBF restore for an entire storage area directory, to prevent destroying other data on the storage area file system. NMD cannot use SnapImage non-FBF restore because it overwrites the entire file system. Generally, a SnapImage FBF restore is:
 - Slower than a SnapImage non-FBF restore (save set-based, destructive restore).
 - The same speed or slower than a restore without SnapImage.



IMPORTANT

“8.4.5. Potential problem with the restore of storage areas that contain numerous files” on page 1-60 provides information on how to use a SnapImage non-FBF restore in the specific case where one or more storage areas from the same file system were backed up by SnapImage and the storage areas contain a total of approximately 1 million or more files.

3. Software components

3.1. Internal components of NMD

NMD consists of four native platform binary executables, a single Java binary, and a configuration file:

- ◆ **nsrnmdiv**, a manual or automated user interface to back up one or more components of a repository.
- ◆ **nsrnmidx**, a manual user interface to list records of previous backup sessions and obtain details of each session.
- ◆ **nsrnmids**, a manual user interface to restore one or more components of a repository.
- ◆ **nsrnmidc**, a manual user interface to list candidate content file paths and to optionally restore those paths selected by the user.
- ◆ **nsrnmide**, a Java binary that sends requests to the Documentum Server for the following:
 - Component discovery and details.
 - Quiesce and unquiesce of components.
 - NMD customized file report dumps.
- ◆ NMD configuration file, an ASCII text file that defines mandatory and optional parameters for the software operations.

The binaries reside in the NetWorker bin directory, for example:

- ◆ /usr/sbin on Solaris
- ◆ C:\Program Files\Legato\nsr\bin on Windows

The configuration file typically resides in the NetWorker res directory, for example:

- ◆ /nsr/res on Solaris
- ◆ C:\Program Files\Legato\nsr\res on Windows

Figure 1-1 on page 1-12 illustrates the components involved in an NMD backup.

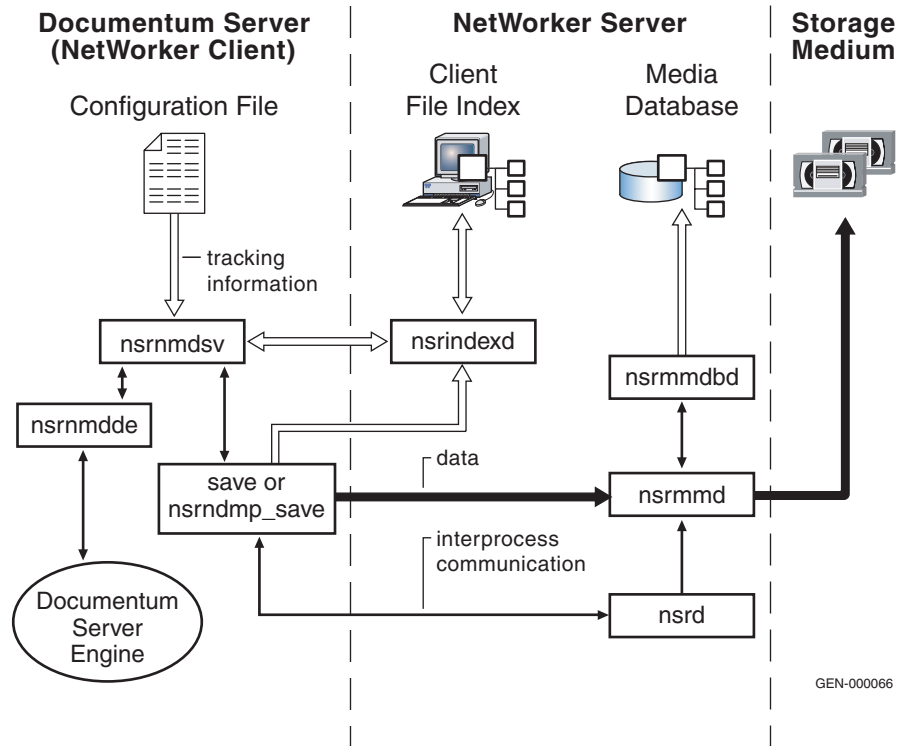


Figure 1-1 Components in an NMD backup

Figure 1-2 on page 1-13 illustrates the components involved in an NMD restore.

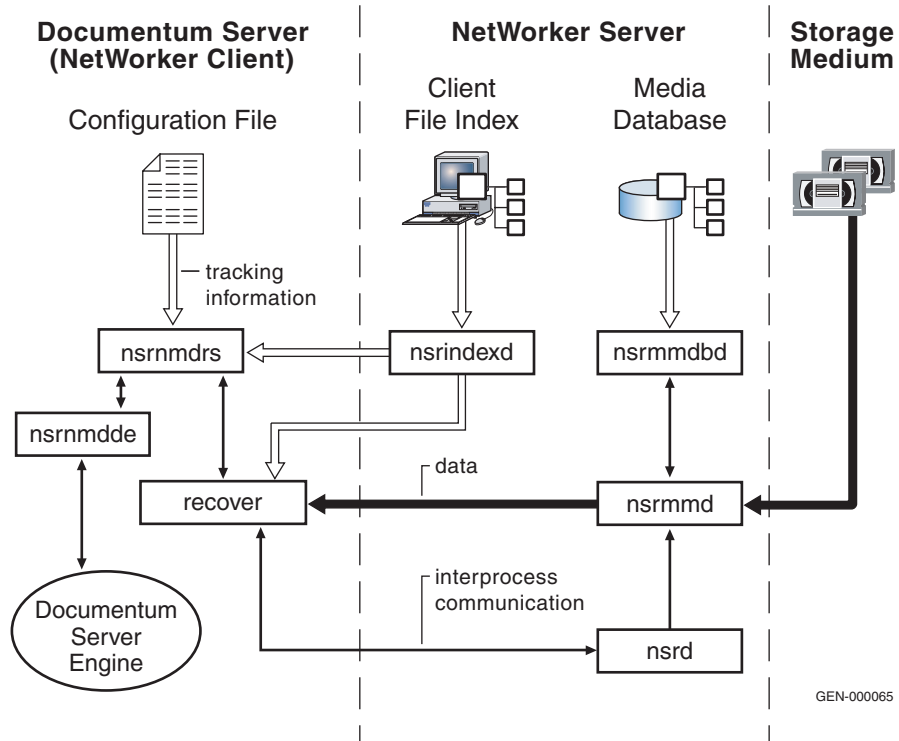


Figure 1-2 Components in an NMD restore

3.2. External modules

If the SnapImage Module will be used, the following external software must be installed:

- ◆ NetWorker 7.2.x client, storage node, and server (with new “DSA FBF/DAR” non-NDMP device support patch)
- ◆ NetWorker SnapImage Module 1.6.1 (or Celestra 1.6.1) on UNIX
- ◆ NetWorker SnapImage Module 2.0.2 on Windows
- ◆ NDMP TapeServer 2.0.2 on Windows
- ◆ EMC License Manager on Windows (for SnapImage 2.0.2 licensing)

The read1st.txt file that is included with the NMD software provides information on any required NetWorker and SnapImage software updates.

The NetWorker Module for the Oracle, SQL Server, or Sybase database is *optional*.

The Documentum Server has the following software installed:

- ◆ NetWorker 7.2.x client
- ◆ NMD 1.1
- ◆ (Optional) NetWorker SnapImage Module 1.6.1 (UNIX) or 2.0.2 (Windows)
- ◆ (Optional) NDMP TapeServer 2.0.2 (Windows)
- ◆ (Optional) NetWorker Module for Oracle, SQL Server, or Sybase

3.3. How NMD interfaces with external modules

NMD records the backup sessions by creating its own file index entries that are written to the backup index through the NetWorker **nsrindexd** service.

NMD interfaces with NetWorker to form legacy NetWorker save sets and client file index entries according to the type of component:

- ◆ Full-text index component — This component makes up a small fraction of the storage area and file system that contains the full-text index. Since NMD must be able to back up full-text index instances by themselves and a full-text index instance (excluding the log files) is typically far less than 4 GB, the instances are backed up through nonsnapshot means by using the NetWorker **save** program. All the files of a full-text index instance (excluding the log files) are packed into the same save set. Furthermore, the files of a full-text index instance should always be backed up as a unit; hence, a level full backup is always performed and a single save set is used.

Note: The backup index information for an incremental backup (that involves the full-text indexes) states that the full-text indexes were backed up at the incremental level, but they were actually backed up at the full level. The backup-level information in the backup index applies *only* to storage area backups.

- ◆ Install and configuration files component — The backup of an install and configuration files component is performed in the same way as the backup of a full-text index component. The backup is a full nonsnapshot backup through the NetWorker **save** program, with all the files packed into the same save set.
- ◆ Database component — NMD passes control of the database backup to an external module, such as NetWorker Module for Oracle or a user-defined script. NMD merely invokes the external database backup process at the appropriate step in the backup workflow sequence. There are hooks through the NMD configuration parameter definitions for invoking a full backup, incremental backup, and transaction logs backup where applicable.
- ◆ Storage area component — The backup of the storage area component is performed as either a snapshot or nonsnapshot backup, depending on the setting of NMD_USE_SNAPIMAGE in the NMD configuration file.
 - The NetWorker SnapImage Module and the NetWorker client **nsrndmp_save** program perform a *snapshot* backup of each file system that contains storage area instances. One save set is created for each file system. Both full (level=0) and differential (level=1) backups are supported.
 - The NetWorker client **save** program performs a *nonsnapshot* backup of all the storage area instances. For each storage area, one or more save sets are created, depending on the setting of NMD_SA_FULL_NUM_SAVESET for a full backup and NMD_SA_INCR_NUM_SAVESET for an incremental backup.

NMD uses the following programs to create the backup save sets:

- ◆ The **save** program for backups of full-text indexes, install and configuration files, and empty storage areas, and nonsnapshot backups of storage areas in use.
- ◆ The **nsrndmp_save** program for snapshot backups of storage areas in use.

The name of a save set is determined by the parameter settings in the NMD configuration file, as follows:

1. If `NMD_USE_DEFAULT_SAVESET_NAMES` is set to YES (strongly recommended), then the name is:
 - `DCB_<repository_name>:<component_type>_<component_name>`: for full-text index and install/configuration file components.
 - `DCB_<repository_name>:SI_SA_<component_name>:+`: for storage area components from a snapshot backup.
 - `DCB_<repository_name>:SA_<component_name>`: for storage area components from a nonsnapshot backup.
2. If `NMD_USE_DEFAULT_SAVESET_NAMES` is *not* set to YES (the default is NO):
 - If the `-N` option is specified with the `nsrnmadv` command, then the `-N` option value is used. NMD supports the `DCB_<repository_name>` format only.

Note: The save set attribute in the NetWorker Client resource (for scheduled backups) is used to specify the `-N` option value.

- If the `-N` option is *not* specified with the `nsrnmadv` command (for example, it is an ad hoc backup without the `-N` option), then the value of `NSR_SAVESET_NAME` is used. If `NSR_SAVESET_NAME` is not set, then the default names from [step 1](#) apply.

NMD creates NetWorker save sessions and save sets according to the command indicated:

- ◆ Full-text index — The data from each full-text index forms one save set:

```
save -l full [-f NMD_DIRECTIVES_FILE_to_exclude_logs]
           full-text_index_verity_subdir
```

- ◆ Install and configuration files — The install and configuration files form one save set:

```
save -l full
      [-f NMD_DIRECTIVES_FILE_to_exclude_unwanted_dirs]
      install_and_config_dirs_and_files
```

- ◆ Storage areas — The save sets created depend on whether a snapshot or nonsnapshot backup is performed.
 - During a snapshot backup, all the storage areas that share the same file system form one save set:

```
nsrndmp_save -l full | 1 -M -T image | celestra  
file_system_mount_point_or_drive
```

Note: By default, empty storage areas are backed up without SnapImage. Empty storage areas are unused (never had any content files), and only have an empty root directory structure that is backed up.

- During a nonsnapshot backup, each storage area is formed from one or more save sets, depending on the setting of `NMD_SA_FULL_NUM_SAVESET` (for a full backup) and `NMD_SA_INCR_NUM_SAVESET` (for an incremental backup):

```
save -l full -I listfile_of_paths
```

where *listfile_of_paths* is a temporary list file returned by the **nsrnmde** process, which contains the paths of the storage area's leaf subdirectories.

Notes:

- The **-l full** option is passed to **save** even for an incremental backup. It is the listfiles that determine the storage area's incremental changes (new documents added).
- An incremental backup without SnapImage is truly incremental, including only changes from the last backup (whether that last backup was full or incremental). A SnapImage backup always backs up all the changes since the last full backup.
- The listfiles of a storage area altogether specify a frozen part of the storage area, within a minute or so at the end of the database backup.
- A storage area's documents are never modified; they are only added or removed. This allows the backup without SnapImage to readily obtain a frozen specification of a storage area.

NMD creates NetWorker recover sessions according to whether the components were originally backed up with or without the SnapImage Module:

◆ Full-text index:

- For the restore of a backup created with NMD_USE_SNAPIMAGE set to YES, each full-text index has a **recover -S** *save_set_ID* session, created with the following command:

```
recover -S save_set_ID [-d relocation_dir]
[-i {Y|R|N}]
```

- For the restore of a backup created with NMD_USE_SNAPIMAGE set to NO, each full-text index has a **recover -t** session, created with the following command:

```
recover -t as_of_timestamp [-d relocation_dir]
[-i {Y|R|N}] -a FTI_root_subdirectory
```

◆ Install and configuration files:

- For the restore of a backup created with NMD_USE_SNAPIMAGE set to YES, the install and configuration files have a **recover -S** *save_set_ID* session, created with the following command:

```
recover -S save_set_ID [-d relocation_dir]
[-i {Y|R|N}]
```

- For the restore of a backup created with NMD_USE_SNAPIMAGE set to NO, the install and configuration files have a **recover -t** session, created with the following command:

```
recover -t as_of_timestamp [-d relocation_dir]
[-i {Y|R|N}] -a list_of_ICF_files_and_directories
```

◆ Storage areas:

- For the restore of a snapshot backup, all the storage areas that share the same file system are restored from one save set:

```
recover [-t as_of_timestamp] [-d relocation_dir]
-a storage_area_subdirs_exclusive_of
full-text_index_subdirs
```

Note: The **recover -t** command must be used to avoid a full file system restore by SnapImage, which would corrupt non-storage area data. By default, empty storage areas are restored without SnapImage.

- For the restore of a nonsnapshot backup, each storage area has a **recover -t** session, created with the following command:

```
recover [-t as_of_timestamp] [-d relocation_dir]
          -a storage_area_root_subdir_exclusive_of_
          full-text_index_subdir
```

Note: These recover sessions are created when a storage area backup is restored with the **nsrnmdrs** command. If one or more storage areas from the same file system were backed up by the SnapImage Module and the storage areas contain a total of approximately 1 million or more files, manually restore the storage areas with the **nsrndmp_recover** command to a staging file system, as described in [“8.4.5. Potential problem with the restore of storage areas that contain numerous files”](#) on page 1-60.

3.4. Supported repository configurations

NMD currently supports only the following:

- ◆ Documentum Server releases 5.2.x and 5.3.x
- ◆ Both single server and separate database server configurations.
 - The Content Server and database server can be on either the same host or separate hosts.
 - NMD must be installed on *both* the Content Server host and a separate database server.
 - If a database NetWorker Module is used, it must be installed on the database server host.

Note: The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides backup script examples that apply to *both* the single server and separate database server configurations.

- ◆ A separate full-text index server with Documentum 5.3.x *only*.
 - The Content Server and full-text index server can be on either the same host or separate hosts.
 - NMD must be installed on *both* the Content Server host and a separate full-text index server.

Note: The use of a full-text index server for multiple repositories is *not* recommended.

- ◆ Only for backups *without* the SnapImage Module, configurations where the storage areas are on different hosts but all storage areas are directly accessible from at least one host through local disks, SAN, shared SCSI, NAS, or NFS.

Note: The SnapImage Module does *not* support remote disks mounted through NAS or NFS.

- ◆ The following types of storage areas: file store, blob, and turbo.
- ◆ If the optional SnapImage Module is used, UNIX or Windows and Volume Manager versions as limited by the SnapImage Module support restrictions. The following sources provide more information on the restrictions:
 - *NetWorker Module for Documentum, Release 1.1, Multiplatform Version, Installation Guide*
 - *NetWorker SnapImage Module, Release 1.0, UNIX Version, Installation and Administrator's Guide*
 - *NetWorker SnapImage Module, Release 1.6.1, UNIX Version, Release Supplement*
 - *NetWorker SnapImage Module, Release 2.0, Microsoft Windows Version, Installation and Administrator's Guide*
 - *NetWorker SnapImage Module, Release 2.0.2, Microsoft Windows Version, Release Notes*

4. Prerequisites and checklist

Best practices require that the customer environment has the recommended hardware and the software is properly configured prior to initiating implementation processes. The required external modules listed in [“3. Software components” on page 1-11](#) must be installed.

Ensure that the prerequisites in the following sections are met:

- ◆ [“4.1. SnapImage prerequisites” on page 1-21](#)
- ◆ [“4.2. NMD prerequisites” on page 1-22](#)
- ◆ [“4.3. NetWorker hardware and software prerequisites” on page 1-22](#)
- ◆ [“4.4. NetWorker software prerequisites” on page 1-23](#)
- ◆ [“4.5. NetWorker server sizing” on page 1-24](#)

4.1. SnapImage prerequisites

Ensure that the following SnapImage prerequisites are met:

- ◆ One raw unmanaged disk partition cache is allocated for each storage area file system or drive. The size of the cache should be approximately 5 percent to 15 percent of the used area of the largest file system to be backed up.

Note: A smaller cache size is required if there is low file system write activity during the snapshot backup. SnapImage on UNIX can use a normal file system for the cache.

- ◆ Extra file system space is reserved for the NDMPHOME directory. The extra space should be approximately 5 percent to 15 percent of the cumulative maximum data to be backed up, which is the aggregate of the used parts of all storage area file systems. The NDMPHOME directory is typically as follows:
 - /usr/ndmphone on UNIX
 - C:\Program Files\Legato\NDMPSvc on Windows
- ◆ Swap area or page file extra space is allocated. The extra space should be approximately as follows:
 - 100 MB for UNIX and 160 MB for Windows for each 1 million storage area files to be backed up.
 - 350 MB for UNIX for each 1 million storage area files to be restored.

The *NetWorker SnapImage Module Installation Guide* provides more information.

- ◆ On Solaris, at least 1 GB of memory is available for a storage area backup.
- ◆ An evaluation license is obtained by calling the EMC Sales/Licensing group. The license is installed as follows:
 - On UNIX, use the NDMPHOME **install_license** utility.
 - On Windows, enter the following command:


```
lgto lic -s ""::1" -c enabler
```
- ◆ An NDMP client license is installed on the NetWorker server or License Manager.

4.2. NMD prerequisites

Ensure that the following NMD prerequisite is met:

An extra file system of sufficient space is available for the optional file reports. The amount of space required depends on the number of repository documents and the frequency of the full or incremental dumps. Typically, at least 2 GB of extra space should be reserved. The Documentum administrator needs to handle the periodic purging of old documents.

4.3. NetWorker hardware and software prerequisites

Ensure that the following NetWorker hardware prerequisites are met:

- ◆ Legacy disk or tape devices can be used with the NetWorker software. The devices must be non-NDMP devices, with the NDMP flag set to NO.
- ◆ For sufficient parallelism, multiple devices are used. In the NetWorker Device resource for each device, the Target Sessions attribute is set to 4 or greater. In the NetWorker Client and Server resources, the Parallelism attribute value is increased accordingly.

Note: “[6.7. NMD configuration file](#)” on page 1-32 provides information on parallelism parameter settings in the NMD configuration file.

- ◆ All hosts are installed, configured, and operating on the network and SAN, if applicable.
- ◆ All hosts can **ping** the servers which will be the NetWorker server and storage nodes.
- ◆ If using DNS for name resolution, all hosts can perform reverse DNS lookups of hostnames by using IP addresses. For example, if there is a client named web1.legato.com with the IP address 137.69.200.8, enter the following **nslookup** command:

```
nslookup 137.69.200.8
Server: ns.legato.com
Address: 137.69.200.200
Name: web1.legato.com
Address: 137.69.200.8
```

In this example, **nslookup** returned the proper hostname for the address, web1.legato.com. If **nslookup** does not report the proper hostname or any hostname, then the DNS configuration must be corrected.

- ◆ There is enough space allocated to the NetWorker server for the client indexes. Typically, this space should be on its own file system. Additionally, the EMC best practice is that this space is *not* on a RAID 5 type disk device. A RAID 1+0 disk device is the most preferable disk configuration for the indexes, although it is not required. NetWorker documentation provides more information.
- ◆ All tape libraries and tape devices are installed, configured, and working properly.
- ◆ All tape devices can be seen and accessed by native operating system utilities on systems that are connected directly to the devices or to a SAN. The utilities include Windows Backup on Windows 2000 systems, and **mt** or **tar** on UNIX systems.
- ◆ All SAN components such as host bus adapters (HBAs), bridges, switches, hubs, routers, MUXs, and so on are installed and working properly.
- ◆ The latest firmware is installed on all SAN bridges, switches, hubs, routers, MUXs, and so on. New devices do not always have the latest firmware. Contact the vendor representative or manufacturer website for verification of the latest firmware.
- ◆ Ensure that the SAN switches are zoned properly so that the NetWorker server and any storage nodes can access the appropriate tape drives and/or disk arrays.
- ◆ If restore testing will be performed, ensure that an appropriate environment is available. This may be a separate directory to restore files, a separate database instance on a production machine, or possibly a completely separate server, depending on the restores to be tested.

4.4. NetWorker software prerequisites

Ensure that the following software prerequisites are met:

- ◆ All EMC software products have been purchased and delivered. EMC Consulting does not deliver any software media or manuals. If any products have not been received, contact your EMC Sales Representative or Reseller.
- ◆ All database applications which will be backed up through NetWorker modules have been installed, configured, and are running.

- ◆ Required operating system patches are installed. Information on these can be found in the *NetWorker Installation Guide* and/or *NetWorker Release Notes*.

4.5. NetWorker server sizing

This section addresses the higher level server sizing for data transfer to a tape or disk target. This is not a complete server architecture sizing guide. For instance, it discusses the amount of memory and memory transfer speed. It does not address types and speeds of the memory chips themselves, nor does it discuss the details of how many HBAs a computer can support. The local engineers and vendors decide this level of detail after the data transfer requirements are determined.

There are basically two avenues that can be followed to determine the number of target drives and supporting server specifications:

1. Determine server specifications based on the appropriate number of target drives.
2. Determine the number of target drives that can be supported based on a current server's specifications.

4.5.1. Server sizing factors

No matter which avenue is followed, the following information is critical to determine the "right fit." Obtain the information from the appropriate hardware vendor:

- ◆ Target drive I/O rate in MB/s
- ◆ Expected ratio of compression
- ◆ CPU speed in MHz
- ◆ Memory bus speed in MB/s = (memory bus MHz * bus width) / 8
- ◆ Speed of peripheral buses
- ◆ HBA or SCSI card speed
- ◆ Number of cards that can be supported
- ◆ Method of delivery for data being protected (network versus direct)
- ◆ Capabilities and bandwidth of network interface card (NIC)

For example, does a GBE NIC have a TCP Offload Engine (TOE) to offset the main CPU resource use?

Kernel and resource requirements (file descriptors, memory heap) must also be properly configured in the operating system. The

operating system documentation provides information on kernel tuning.

The modifiers to server sizing are the following:

- ◆ Data change rate
- ◆ Frequency of backup
- ◆ Frequency and type of target drive operations (clone, archive, stage, recover, use of B2D)

These modifiers can increase or reduce the resources needed for server sizing, but are outside the scope of this discussion.

4.5.2. Server sizing example

In the following tape-specific example, the calculations are approximations. Some of the formulas are based on lab testing of certain devices and extrapolated for others. For example, if it takes 200 MHz to support a DLT7000 tape drive writing at 18 MB/s, then by extrapolation 400 MHz of CPU is required to drive a device at 36 MB/s.

4.5.2.1. Sizing example assumptions

The following assumptions are made:

- ◆ The 9840Bs used in the “CustomerX datazone” have a native speed of 19 MB/s, so we could use 200 Mhz of CPU for each drive, but that doesn’t allow for compression.
- ◆ There are four 600 Mhz CPUs.
- ◆ There is 4 GB of memory.
- ◆ STK has recommended a limit of two 9840Bs per HBA.
- ◆ There are two HBAs, each having a 2 GB/s transfer rate.
- ◆ During the time frame of the health check, the number of 9840B drives per HBA was reduced from five to four and finally to three, for a total of six 9840B tape drives on the NetWorker server.
- ◆ One GBE NIC is being used for network backups.
- ◆ The memory bus transfer speed is unknown.
- ◆ The Fibre ready 9840Bs are behind a new 2 GB/s IBM fibre switch.

To perform a thorough analysis, we still need to know: the memory bus transfer rate, the transfer rate of the fibre switch, and whether the GBE NIC is TOE enabled.

4.5.2.2. Sizing example logic flow

The responses to the following questions are based on the known NetWorker server specifications:

1. Is one GBE NIC sufficient to keep six 9840B tape drives spinning?

Answer: No

Assumption: The GBE NIC operates at 60 percent of the theoretical speed (60 MB/s).

With the native rate of six 9840B drives being a total of 114 MB/s, an input rate provided by a single GBE NIC is insufficient to keep six tape drives writing at capacity even with 1:1 compression, especially with higher compression rates.

2. Are there enough CPUs to support six 9840B tape drives?

Answer: Possibly

Using the model of # CPUs = (n+1) and considering that the 9840B drives require 200 MHz of CPU to write at 1:1 compression and 300 MHz for 1:1.5 compression, three 600 MHz/s CPU drives should be sufficient for six tape drives, with one left to be dedicated to processing. However, performance may decrease if the GBE NIC is not TOE enabled and requires a portion of processing power expected to be used for I/O to the 9840B drives.

3. Are two HBAs sufficient to support six 9840B drives?

Answer: Yes

Assumption: The HBAs in use have a true transfer rate of 2 GB/s.

At 1:5 compression and a 2 GB/s HBA transfer rate (2000 MB/s) and with three 9840B drives being written to at an I/O rate of approximately 90 MB/s, the tape drive-to-HBA ratio can support the six 9840B drives. However, keep in mind that STK has recommended only two 9840B drives per HBA.

4. Are the internal bus transfer rates sufficient?

Answer: Unknown

As previously stated, it is very important to know the memory bus transfer rate to determine if the data can be transferred from the GBE NIC to the CPU to the HBA to the 9840B at the necessary transfer rate.

4.5.2.3. Conclusion

The current configuration of the NetWorker server is insufficient to support six 9840B tape drives because only one GBE NIC is used to bring in the data. More information is needed about the specifications of the GBE NICs in use before it can be determined whether using another NIC can be of any benefit.

5. Installation overview

The following sections provide instructions for installing the required software on the Documentum Server, also known as the NMD NSR client.

5.1. NetWorker

To install the NetWorker client software on the Documentum Server, perform the following:

- ◆ On UNIX, enter the appropriate installation command for the platform, as described in the *NetWorker Installation Guide*.

For example, on Solaris, select the LGTOclnt package for installation. If the Documentum Server is also to be a NetWorker storage node, then also select the LGTOdrv and LGTONode packages.

- ◆ On Windows, run **setup.exe** or the equivalent self-extracting binary.

The NetWorker server host can be any platform supported by NetWorker release. On the NetWorker server, install the NetWorker packages for the server, client, and storage node, and if applicable, drivers.

If SnapImage 2.0.2 is used on Windows, the EMC License Manager (LLM) must also be installed and functional for SnapImage 2.0.2 licensing. LLM is installed on the NetWorker server host.

The *NetWorker Installation Guide*, *Administration Guide*, and *Release Notes* provide more information. The *NetWorker Administration Guide* contains useful information about SnapImage, which is one type of NDMP Data Server host.

The read1st.txt file (included with the NMD software) provides information on any required NetWorker and SnapImage software updates.

5.2. Optional SnapImage 1.6.1 on UNIX

To install the optional SnapImage 1.6.1 on UNIX, use the information in the following sources:

- ◆ *NetWorker SnapImage Module, Release 1.0, UNIX Version, Installation and Administrator's Guide*
- ◆ *NetWorker SnapImage Module, Release 1.6.1, UNIX Version, Release Supplement*

The read1st.txt file that is included with the NMD software provides information on any required SnapImage software updates.

5.3. Optional SnapImage 2.0.2 and NDMP TapeServer 2.0.2 on Windows

To install the optional SnapImage 2.0.2 and NDMP TapeServer 2.0.2 on Windows, use the information in the following sources:

- ◆ *NetWorker SnapImage Module, Release 2.0, Microsoft Windows Version, Installation and Administrator's Guide*
- ◆ *NetWorker SnapImage Module, Release 2.0.2, Microsoft Windows Version, Release Notes*

Note: The NDMP TapeServer from the NDMP Services package must be installed *before* the SnapImage software.

The read1st.txt file that is included with the NMD software provides information on any required SnapImage software updates.

5.4. NetWorker Module for Oracle, SQL Server, or Sybase (optional)

To install the required NetWorker Module, follow the instructions in the corresponding NetWorker Module Installation Guide.

For example, to install the NetWorker Module for Oracle 4.1 on Solaris, enter the following command:

```
pkgadd -d dir_pathname LGTONmo
```

where *dir_pathname* is the complete pathname of the directory containing the **LGTONmo** package.

The appropriate NetWorker Module documentation provides more information on installation.

5.5. NMD 1.1

To install NMD 1.1, follow the instructions in the *NetWorker Module for Documentum, Release 1.1, Installation Guide*.

6. Configuration overview

The following sections describe the requirements for configuring resources for the NMD backups.

6.1. NetWorker Client resource

To configure the NetWorker Client resource, use the NetWorker Administrator program to specify the following values for the required resource attributes:

- ◆ Application Information:
 - For snapshot backups:
 - DDIMAGE=n
 - DIRECT=y
 - HIST=y
 - REMOTE=y
 - For nonsnapshot backups — Leave blank
- ◆ Backup Command — **nsrnmdiv -I *client_name* -f *NMD_config_filepath*** (for example, **-f /nsr/res/nmd.cfg**)
- ◆ Group — Name of the NetWorker Group resource, required for *both* ad hoc and scheduled backups
- ◆ NDMP:
 - For snapshot backups — Yes
 - For nonsnapshot backups — No
- ◆ Parallelism — For nonsnapshot backups, 16 or 32 (recommended, but superceded by the setting of NSR_PARALLELISM)
- ◆ Password — *Password_of_Remote_User*; mandatory for snapshot backups
- ◆ Remote User — Name of the root user (UNIX) or a member of the administrators group (Windows); mandatory for snapshot backups; optional for Windows
- ◆ Save Set — *DCB_repository-name*

6.2. NetWorker Device resource

To configure the NetWorker Device resource, use the NetWorker Administrator program to specify the following values for the required resource attributes:

- ◆ NDMP — No
- ◆ Target Sessions — For nonsnapshot backups, 4 or 8

Note: Is it recommended to keep Target Sessions at the default value of 4 and add more devices.

6.3. NetWorker Group resource

To configure the NetWorker Group resource, use the NetWorker Administrator program to specify the following values for the required resource attributes:

- ◆ Autostart — Enabled
- ◆ Inactivity Timeout (hidden attribute) — 0, to disable the inactivity timeout

Note: In the NMD configuration file, set the NMD_TASK_TIMEOUT_MINS parameter to a suitable value such as 120, as the time beyond which a single active or inactive task is forcefully terminated.

- ◆ Snapshot — False, even if the SnapImage Module is being used
- ◆ Start Time — Backup start time

The NetWorker Group resource is required for *both* ad hoc and scheduled NMD backups. Before a backup, a Group resource must be configured for the backup group that is specified by one of the following:

- ◆ The Group attribute in the NetWorker Client resource.
- ◆ The **-g** option of the **nsrnmdsv** backup command.
- ◆ The NSR_GROUP parameter in the NMD configuration file.

The **nsrnmdsv** command option takes *precedence* over the configuration file parameter.

Enter the name of the NetWorker group in the Group attribute of the Client resource.

Note: To have a scheduled backup automatically use a volume pool associated with the backup group, specify the group name in the Pool resource for the volume pool.

To use the Default group for testing scheduled NMD backups, change its Autostart attribute to Enabled.

6.4. NetWorker Server resource

To configure the NetWorker Server resource for *nonsnapshot* backups, use the NetWorker Administrator program to specify the following value for the required resource attribute:

- ◆ Parallelism — 32 or 64 (recommended)

6.5. Optional SnapImage 1.6.1 on UNIX

To configure the optional SnapImage 1.6.1 on UNIX, perform the following:

1. Configure the cache by entering the following command:

```
/usr/ndmphone/bin/cacache  
-c raw_unmanaged_disk_partition
```

2. Install the required license by entering the following command:

```
/usr/ndmphone/bin/install_license
```

Follow the onscreen prompts.

The following sources provide more information:

- ◆ *NetWorker SnapImage Module, Release 1.0, UNIX Version, Installation and Administrator's Guide*
- ◆ *NetWorker SnapImage Module, Release 1.6.1, UNIX Version, Release Supplement*

6.6. Optional SnapImage 2.0.2 and NDMP TapeServer 2.0.2 on Windows

To configure the optional SnapImage 2.0.2 and NDMP TapeServer 2.0.2 on Windows, perform the following:

1. Configure the cache by performing one of the following:
 - Enter the following command:


```
SIConfigCache /c raw_device_path
```

 For example, *raw_device_path* is X:\.
 - Use the SnapImage Administrator program.
2. Install the required license by entering the following command:

```
lgtolic -s "::<1" -c enabler
```

The following sources provide more information:

- ◆ *NetWorker SnapImage Module, Release 2.0, Microsoft Windows Version, Installation and Administrator's Guide*
- ◆ *NetWorker SnapImage Module, Release 2.0.2, Microsoft Windows Version, Release Notes*

6.7. NMD configuration file

In the NMD configuration file:

- ◆ Each parameter setting must occur on a separate line.
- ◆ Each line containing a parameter setting must *not* include any line breaks except at the end of the line.
- ◆ Each parameter must be specified in the following format:


```
parameter=value
```
- ◆ Parameter names and values are case-insensitive *except* for UNIX pathnames, which are case-sensitive.
- ◆ Comment lines must start with the # symbol.

A working sample of the configuration file (that must be customized) is provided by the NMD installation process in the following location:

- ◆ On UNIX:


```
installation_path/nsrnmd_unix.cfg
```

 or (if NMD is relocated)


```
relocation_path/nsrnmd_unix.cfg
```
- ◆ On Windows:


```
%SystemDrive%\Program Files\Legato\nsr\bin\nsrnmd_win.cfg
```

You can copy this sample file to any location and customize the parameter settings as required for your specific NMD configuration. “6.7.1 Recommended parameter settings” on page 1-38 provides recommended settings for the NMD configuration parameters.



IMPORTANT

Before you start any type of NMD backup, ensure that the mandatory parameters are set. Specific parameters are mandatory, depending on how NMD will be used. The *NetWorker Module for Documentum Administration Guide* provides complete details on the configuration file parameters.

Example 1-1 Sample configuration file on UNIX

The sample NMD configuration file on UNIX is named /nsr/res/nsrnmnd_unix.cfg, and contains the following parameter settings.

Note: In the configuration file, each line that contains a parameter setting (*without* the # symbol at the start) must *not* include any line breaks except at the end of the line.

```
#####
#NetWorker Module for Documentum v1.1 nsrnmnd_unix.cfg
#D5.2/D5.3 example with typical set of working
#parameters. THIS FILE MUST BE SITE-CUSTOMISED.
#
#Parameters not shown can be set in this file (as per
#site customisation) or from the command-line.
#
#Please refer to the user guides for details on all
#parameters, including those not listed below.
#####
#Can get values for most of below from doing as root:
#su - dctmadmin -c env
DOCUMENTUM=/db/dctm
DOCUMENTUM_SHARED=/db/dctm/shared/5.3SP1
DM_HOME=/db/dctm/product/5.3SP1
JAVA_HOME=/db/dctm/shared/5.3SP1/java/1.4.2_04
JAVA_PATH=/db/dctm/shared/5.3SP1/java/1.4.2_04/bin
ORACLE_HOME=/db/oracle/product/9.2.0
PATH=.: /db/dctm/shared/5.3SP1/java/1.4.2_04/bin:
/db/dctm/product/5.3SP1/bin:/usr/bin:/usr/sbin:/usr/ucb:
/etc:/db/oracle/product/9.2.0/bin:/db/dctm/dba
#AIX uses LIBPATH and HP-UX uses SHLIB
```

```

LD_LIBRARY_PATH=:/db/oracle/product/9.2.0/lib:/db/dctm/
shared/5.3SP1/dfc:/db/dctm/product/5.3SP1/bin:/db/dctm/
shared/5.3SP1/javad1.4.2_04/lib:/db/dctm/share/clients/
unix/solaris:/db/dctm/product/5.3SP1/fusion
#See also manifest dctm.jar file for class path locations
CLASSPATH=./usr/sbin/nsrnmde.jar:/db/dctm/product/
5.3SP1/dctm-server.jar:/db/dctm/product/5.3SP1/bin:
/db/dctm/shared/5.3SP1/dctm.jar:/db/dctm/shared/5.3SP1/
config:/db/dctm/shared/5.3SP1/share/clients/classes
#####
#If not using nsrnmdev -m
#ALL|DB|DB_LOG|FTI|FTI_ALL|ICF|SA|SA_ALL, set for backup
#NMD_SCOPE=ALL
#####
NMDDE_DM_DOCBASE=dctmrep01
NMDDE_DM_USER=dctmadmin
#NMDDE_DM_PASSWD must be set via running:
#nsrnmdev -f <nmdcfg> -P <pwd>.
#####
#DB related hooks to invoke arbitrary scripts:
#Set if DB is on a remote host
NMD_DB_HOST=dbhost
#Pure basename implies remote host execution; absolute
#path ... local execution as in NMD v1.0.
#
#Remote execution requires script be put in remote
#nsrexecd bin directory and D5.3 host be added to remote
#/nsr/res/servers file w/ nsrexecd recycled.
#
#Refer to user guides for sample script code.
NMD_DB_FULL_BACKUP_CMD=nsrnmddb.sh
NMD_DB_LOG_BACKUP_CMD=nsrnmddb.sh
NMD_DB_INCR_BACKUP_CMD=nsrnmddb.sh
#####
#For D5.3 only: NMD_FTI_HOST, NMD_FTI_USER,
##NMD_FTI_PASSWD & NMD_FTI_SUBDIRS.
#Mandatory for remote D5.3 FTI server
#NMD_FTI_HOST=ftihost
#Recommended for D5.3 FTI server quiesce/unquiesce
#NMD_FTI_USER=dctmadmin
#Recommended for D5.3 FTI server quiesce/unquiesce
#NMDDE_FTI_PASSWD must be set via running:
#nsrnmdev -f <nmdcfg> -P <pwd> -M FTI.
#Pure basename implies remote host execution; absolute
#path ... local execution as in NMD v1.0.

#Remote execution requires script be put in remote
#nsrexecd bin directory and D5.3 host be added to remote
#/nsr/res/servers file w/ nsrexecd recycled.
#
#See example nsrnmdf*ti*.sh examples.
#

```

```

#Mandatory for D5.3
#NMD_BACKUP_FTI_QUIESCE=nsrnmfdftiq.sh
#NMD_BACKUP_FTI_UNQUIESCE=nsrnmfdftiu.sh
#Mandatory for D5.3. No spaces in paths or around comma
#separators.
#NMD_FTI_SUBDIRS=/db/dctm_idx/data/fulltext/fixml,
#/db/dctm_idx/data/fulltext/index
#####
#Mandatory for D5.3, optional for D5.2. No spaces in
#paths or before comma separators in NMD_ICF_SUBDIRS_xxx:
#NMD_ICF_INCLUDED=yes
#NMD_ICF_SUBDIRS_AUGMENT=/db/other_misc
#NMD_ICF_SUBDIRS=/db/dctm/dba,/db/dctm/product/5.3SP1,
#/db/dctm/dmcl.ini
#####
NMD_FILEREPORT_INCLUDED=no
#NMDDE_METADATA_OUTPUT_DEST=/filereports/fp
#####
#Other misc recommended NMD_xxx parameters
#Recommended to get more meaningful saveset names
NMD_USE_DEFAULT_SAVESET_NAMES=yes
#Use following to skip unwanted ICF, FTI and
#non-SnapImage SA dirs/files.
#For example, "<</>> +skip: dm_ftwork_dir" line will skip
#non-data files.
#
#The path will be the same on D5 and remote FTI hosts.
#NMD_DIRECTIVES_FILE=/nsr/res/nsrnmddirectives.txt
#For non-SnapImage SA backup
#NMD_SA_FULL_NUM_SAVESET=16
#NMD_SA_INCR_NUM_SAVESET=4
#NMD_USE_SNAPIMAGE=yes
#####
NSR_SERVER=nsvrhost
NSR_CLIENT=d5svrhost
NSR_GROUP=Default
NSR_DATA_VOLUME_POOL=Default
#NSR_SNAPIMAGE_DATA_VOLUME_POOL=Default
#Relocation dir will be the same on D5 and remote FTI
#hosts.
NSR_RELOCATION=/reloc
#NSR_PARALLELISM=16
#NSR_DEBUG_FILE=/nsr/applogs/nmd.log
#NSR_DEBUG_LEVEL=9
#####

```

Example 1-2 Sample configuration file on Windows

The sample NMD configuration file on Windows is named `%SystemDrive%\Program Files\Legato\nsr\res\nsrnm_win.cfg`, and contains the following parameter settings.

Note: In the configuration file, each line that contains a parameter setting (*without* the # symbol at the start) must *not* include any line breaks except at the end of the line.

```
#####
#NetWorker Module for Documentum v1.1 nsrnm_win.cfg
#D5.2/D5.3 example with typical set of working
#parameters. THIS FILE MUST BE SITE-CUSTOMISED.
#
#Parameters not shown can be set in this file (as per
#site customisation) or from the command-line.
#
#Please refer to the user guides for details on all
#parameters, including those not listed below.
#####
#Can get values for most of below from doing as the dctm
#inst owner: cmd> set
DOCUMENTUM=F:\Documentum
DM_HOME=F:\Documentum\product\5.2
JAVA_HOME=F:\Documentum\RTE\jdk\131_04
JAVA_PATH=F:\Documentum\RTE\jdk\131_04\bin
ORACLE_HOME=e:\db\oracle\ora920
PATH=E:\WINNT\system32;E:\WINNT;E:\WINNT\system32\WBEM;
E:\Program Files\Legato\nsr\bin;F:\Documentum\RTE\
Shared;F:\Documentum\product\5.2\bin;E:\db\oracle\
ora920\bin;E:\Program Files\Oracle\jre\1.1.8\bin;
E:\Program Files\Oracle\jre\1.3.1\bin
#See also manifest dctm.jar file for class path locations
CLASSPATH=.;E:\Program Files\Legato\nsr\bin\
nsrnmde.jar;F:\Documentum\RTE\jdk\131_04\lib\tools.jar;
F:\Documentum\RTE\Shared\dfc.jar;F:\Documentum\RTE\
Shared\dfcbase.jar;F:\Documentum\RTE\Shared\dctm.jar;
F:\Documentum\RTE\Shared\workflow.jar;F:\Documentum\RTE\
Shared\log4j.jar;F:\Documentum\RTE\jdk\131_04\lib\
dt.jar;F:\Documentum\RTE\jdk\131_04\lib\dt.jar;
#####
#If not using nsrnmdev -m
#ALL|DB|DB_LOG|FTI|FTI_ALL|ICF|SA|SA_ALL, set for backup
#NMD_SCOPE=ALL
#####
NMDDE_DM_DOCBASE=dctmrep01
NMDDE_DM_USER=dctmadmin
#NMDDE_DM_PASSWD must be set via running:
#nsrnmdev -f <nmdcfg> -P <pwd>.
```



```
#####
#DB related hooks to invoke arbitrary scripts:
#Set if DB is on a remote host
#NMD_DB_HOST=dbhost
#Pure basename implies remote host execution; absolute
#path ... local execution as in NMD v1.0.
#
#Remote execution requires script be put in remote
#nsrexecd bin directory and D5.3 host be added to remote
#nsr\res\servers file w/ nsrexecd recycled
#
#Refer to user guides for sample script code.
NMD_DB_FULL_BACKUP_CMD=nsrnmddbfbat
NMD_DB_LOG_BACKUP_CMD=nsrnmddbblbat
NMD_DB_INCR_BACKUP_CMD=nsrnmddbibat
#####
#For D5.3 only: NMD_FTI_HOST, NMD_FTI_USER,
#NMD_FTI_PASSWD & NMD_FTI_SUBDIRS.
#Mandatory for remote D5.3 FTI server
#NMD_FTI_HOST=ftihost
#Recommended for D5.3 FTI server quiesce/unquiesce
#NMD_FTI_USER=dctmadmin
#Recommended for D5.3 FTI server quiesce/unquiesce
#NMDDE_FTI_PASSWD must be set via running:
#nsrnmddsv -f <nmdcfg> -P <pwd> -M FTI.
#Pure basename implies remote host execution; absolute
#path ... local execution as in NMD v1.0.
#
#Remote execution requires script be put in remote
#nsrexecd bin directory and D5.3 host be added to remote
#nsr\res\servers file w/ nsrexecd recycled.
#
#See example nsrnmddfti*.sh examples.
#
#Mandatory for D5.3
#NMD_BACKUP_FTI_QUIESCE=nsrnmddftiq.sh
#NMD_BACKUP_FTI_UNQUIESCE=nsrnmddftiu.sh
#Mandatory for D5.3. No spaces in paths or around comma
#separators.
#NMD_FTI_SUBDIRS=F:\Documentum_idx\data\fulltext\fixml,
#F:\Documentum_idx\data\fulltext\index
#####
#Mandatory for D5.3, optional for D5.2. No spaces in
#paths or before comma separators in NMD_ICF_SUBDIRS_xxx:
#NMD_ICF_INCLUDED=yes
#NMD_ICF_SUBDIRS_AUGMENT=F:\Documentum\other_misc
#NMD_ICF_SUBDIRS=F:\Documentum\dba,F:\Documentum\
#product\5.2
#####
NMD_FILEREPORT_INCLUDED=no
#NMDDE_METADATA_OUTPUT_DEST=F:\filereports\fp
#####
```

```

#Other misc recommended NMD_xxx parameters
#Recommended to get more meaningful saveset names
NMD_USE_DEFAULT_SAVESET_NAMES=yes
#Use following to skip unwanted ICF, FTI and
#non-SnapImage SA dirs/files.
#For example, "<</>> +skip: dm_ftwork_dir" line will skip
#non-data files.
#
#The path will be the same on D5 and remote FTI hosts.
#NMD_DIRECTIVES_FILE=E:\Program Files\Legato\nsr\
#res\nsrnmddirectives.txt
#For non-SnapImage SA backup
#NMD_SA_FULL_NUM_SAVESET=16
#NMD_SA_INCR_NUM_SAVESET=4
#NMD_USE_SNAPIMAGE=yes
#####
NSR_SERVER=nsvrhost
NSR_CLIENT=d5svrhost
NSR_GROUP=Default
NSR_DATA_VOLUME_POOL=Default
#NSR_SNAPIMAGE_DATA_VOLUME_POOL=Default
#Relocation dir will be the same on D5 and remote FTI
#hosts.
NSR_RELOCATION=F:\reloc
#NSR_PARALLELISM=16
#NSR_DEBUG_FILE=E:\Program Files\Legato\nsr\aplogs\
#nmd.log
#NSR_DEBUG_LEVEL=9
#####

```

6.7.1 Recommended parameter settings

The following are recommended parameter settings in the NMD configuration file:

- ◆ NMD_APPEND_TASK_LOGS=yes
- ◆ NMD_DIRECTIVES_FILE=*NetWorker_directives_filepath*
 With Documentum 5.2.x *only*, to improve the performance of a full-text index backup by excluding the log files, specify the pathname of a NetWorker directives file, such as /nsr/res/nmd_directives.txt, where the file contains the following line:


```
<< / >> +skip: dm_ftwork_dir
```
- ◆ NMD_SA_FULL_NUM_SAVESET=16 (for nonsnapshot backups of storage areas; tune accordingly)

This is the approximate number of save sets for each full nonsnapshot backup of a storage area. For example, if you have only two large storage areas and each is 200 GB, a value of 25 to

50 might be better since the ideal upper limit for NetWorker save set sizes is approximately 4 to 8 GB (200 GB/50). The default value is 10.

- ◆ NMD_SA_INCR_NUM_SAVESET=4 (for nonsnapshot backups of storage areas; tune accordingly)

This is the approximate number of save sets for each incremental nonsnapshot backup of a storage area. The default value is 2.

- ◆ NMD_USE_DEFAULT_SAVESET_NAMES=yes
- ◆ NMD_USE_SNAPIMAGE=yes or no
 - For snapshot backups of storage areas, set to yes (mandatory).
 - For nonsnapshot backups, set to no or comment out the parameter.

- ◆ NSR_DEBUG_FILE=*debug_logfile_pathname*

The value “default” causes the log file /nsr/applogs/nsrnmdsv.pid to be used.

- ◆ NSR_DEBUG_LEVEL=*debug_level*

Debug levels from 0 to 9 are supported.

- A level of 9 produces the maximum trace logging.
- A level of 0 (default) produces no logging, which is suitable for normal production use.

To reduce the size of the NMD debug log file, set NMD_DRIVER_DELAY_SECS to 10 or greater.

- ◆ NSR_PARALLELISM=8 or 16 (recommended for nonsnapshot backups)

Alternatively, comment out NSR_PARALLELISM and set NMD_SAVE_PARALLELISM and NMD_RESTORE_PARALLELISM to separate values.

- ◆ NSR_QUIET=yes

When set to no or commented out, and NSR_DEBUG_LEVEL is 1 or greater, NMD backup and restore tasks generate detailed debug information.

- ◆ NSR_SNAPIMAGE_DATA_VOLUME_POOL=*NetWorker_pool_name*

This volume pool is used exclusively for save sets from SnapImage backups of the storage areas.

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides more information on the NMD configuration file and its parameters.

6.8. Remote database server

NMD software supports a remote database server configuration, whereby the database server is located on a separate host from the Content Server.

To set up backups of a remote database server:

1. Ensure that the NetWorker and NMD software is installed on the remote database server. The *NetWorker Module for Documentum, Release 1.1, Installation Guide* provides more information.
2. Ensure that the required parameters are set in the NMD configuration file on the Content Server. Pay particular attention to the following parameters:
 - NMD_DB_HOST is *mandatory* for a remote database server.
 - One or more of the following must be set properly to specify the backup scripts on the remote database server:
 - NMD_DB_FULL_BACKUP_CMD
 - NMD_DB_INCR_BACKUP_CMD
 - NMD_DB_LOG_BACKUP_CMD
 - If the backup includes preprocessing or postprocessing, the following must be set properly for a remote database server:
 - NMD_DB_POSTCMD
 - NMD_DB_PRECMD

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides complete details on the configuration file parameters.

3. Ensure that the backup script complies with the guidelines in the *NetWorker Module for Documentum, Release 1.1, Administration Guide*. The script must repeat (display as output) the following to indicate success of the backup:

```
##completed
```

4. Add the hostname of the Content Server to the `/nsr/res/servers` or `NetWorker_install_path\res\servers` file on the remote database server.

5. Shut down and restart the **nsrexecd** service on the remote database server. The *NetWorker Administration Guide* provides more information.

6.9. Remote full-text index server

With Documentum 5.3.x, NMD software supports a remote full-text index server, located on a separate host from the Content Server.

To set up backups of a remote full-text index server:

1. Ensure that the NetWorker and NMD software is installed on the remote full-text index server. The *NetWorker Module for Documentum, Release 1.1, Installation Guide* provides more information.
2. Ensure that the required parameters are set in the NMD configuration file on the Content Server. Pay particular attention to the following *mandatory* parameters:
 - NMD_BACKUP_FTI_QUIESCE
 - NMD_BACKUP_FTI_UNQUIESCE
 - NMD_FTI_HOST
 - NMD_FTI_PASSWD
 - NMD_FTI_SUBDIRS
 - NMD_FTI_USER

Note: Set NMD_FTI_PASSWD *only* by entering the **nsrnmdev** command with the **-f, -P,** and **-M FTI** options at the operating system command line. The parameter must appear *only once* in the configuration file. Do *not* set the parameter by entering the value as text in the configuration file.

NMD_FTI_HOST, NMD_FTI_PASSWD, and NMD_FTI_USER are mandatory with a remote full-text index server *only*.

NMD_BACKUP_FTI_QUIESCE, NMD_BACKUP_FTI_UNQUIESCE, and NMD_FTI_SUBDIRS are mandatory with Documentum 5.3.x *only*.

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides complete details on the configuration file parameters.

3. Ensure that scripts specified by NMD_BACKUP_FTI_QUIESCE and NMD_BACKUP_FTI_UNQUIESCE repeat (display as output) the following to indicate the success of the operation:

```
##completed
```

Refer to the following sample scripts for full-text index quiesce and unquiesce operations, respectively, which are installed with the NMD software:

- nsrnmdftiq.[bat | sh]
- nsrnmdftiu.[bat | sh]

4. Add the hostname of the Content Server to the `/nsr/res/servers` or `NetWorker_install_path\res\servers` file on the remote full-text index server.
5. Shut down and restart the **nsrexecd** service on the remote full-text index server. The *NetWorker Administration Guide* provides more information.

7. Overview of quick usage and verification

Verify the NMD configuration by performing small-sized ad hoc backups and then isolated restores to a relocation directory. Monitor any error or warning messages sent to the screen or designated command session log file. [“8.3. NMD troubleshooting scenarios” on page 1-55](#) provides more information on quick verification tests and troubleshooting.

To verify that all backup and restore operations are functioning properly, perform the following:

1. Without using NMD, test a SnapImage backup and restore by entering the following commands:

```
nsrndmp_save -c NMD_client -M -T type
             small_file_system_or_drive_fullpath
```

where *type* is **image** for Windows and **celestra** for UNIX.

```
recover -d relocation_dir [-a
                          file_or_dir_fullpath_from_small_file_system_or_drive]
```

2. Back up the installation and configuration files by entering the **nsrnmdsv -M ICF** command.
3. Back up all the full-text indexes by entering the **nsrnmdsv -M FTI_ALL** command.
4. Back up one small, in-use storage area by entering the **nsrnmdsv -M SA -a small_storage_area_name** command. (The **-a** option is *not* supported with **-M FTI** and Documentum 5.3.x *only*.)

5. Back up the database by entering the **nsrnmdiv -M DB** command.
6. Query the previous backups by entering the **nsrnmdiv** command and inspect the command output.
7. Restore the installation and configuration files to a relocation directory on a file system with adequate disk space by entering the **nsrnmdiv -d relocation_dir -C ICF** command. Quickly validate the restored data.
8. Restore all the full-text indexes to the same relocation directory by entering the **nsrnmdiv -d relocation_dir -C FTI** command.
9. Restore a single content file by entering, for example, the **nsrnmdiv -B repository_name -A filestore_name -P filereport_path -T timebound -d doc_name -D relocation_dir** command.
10. Restore the small storage area to the same relocation directory by entering the **nsrnmdiv -d relocation_dir -C SA -a small_SA_name** command. (The **-a** option is *not* supported with **-C FTI** and Documentum 5.3.x *only*.)
11. Run a scheduled backup of the installation and configuration files:
 - a. Perform one of the following:
 - Set **NMD_SCOPE=ICF** (temporarily) in the NMD configuration file.
 - Specify the **nsrnmdiv -M ICF** command in the Backup Command attribute of NetWorker Client resource.
 - b. Set **NMD_FILEREPORT_INCLUDED=no** (temporarily) in the NMD configuration file.
 - c. Perform one of the following:
 - Schedule the backup through the NetWorker Administrator program.
 - Start the backup by entering the **savegrp -c NMD_client_name** command on the NetWorker server.
 - d. When the backup finishes, clear the relocation direction.

- e. Query the backup by entering the **nsrnmidx** command and inspect the command output.
 - f. Remove the temporary settings (NMD_FILEREREPORT_INCLUDED=no and possibly NMD_SCOPE=ICF) from the NMD configuration file.
12. Run a scheduled backup of the entire Documentum system. This might take a long time, depending on the repository size and whether a file report dump is requested (the default).

7.1. Ad hoc NMD backup

To perform an ad hoc (manual or unscheduled) NMD backup, enter the **nsrnmdiv** command at the operating system command line, as either the root user on Solaris or a member of the Microsoft Windows Administrators group.

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides more information on the **nsrnmdiv** command and its options.

Example 1-3 Sample nsrnmdiv commands for different types of backups

The following **nsrnmdiv** command performs a Documentum system backup of the repository docora92:

```
nsrnmdiv -B docora92 -f /space10/home/dmuser/nmd.cfg
-M ALL -l full
```

The following **nsrnmdiv** command performs an incremental system backup of the repository docora92:

```
nsrnmdiv -B docora92 -f /space10/home/dmuser/nmd.cfg
-M ALL -l incr
```

The following **nsrnmdiv** command performs an incremental backup of the storage area named filestore_01:

```
nsrnmdiv -B docora92 -f /space10/home/dmuser/nmd.cfg
-M SA -a filestore_01 -l incr
```

Example 1-4 Backup information displayed by the nsrnmdiv command

The following **nsrnmdiv** command performs an incremental backup of the storage area named filestore_01 for the repository babilon by using the NetWorker server birdman:

```
nsrnmdiv -f e:\nmd\config\nmd.cfg -M SA -a filestore_01
-s birdman -l incr
```


The **nsrnmdev** command displays the following information during the backup process:

```

12/20/04 16:47:02 nsrnmdev: Backup parameters validation:
    Option NMD_USE_SNAPIMAGE is either not set, or set to
    empty string. Using default value no.
12/20/04 16:47:02 nsrnmdev: Backup parameters validation:
    Option NSR_CLIENT is either not set, or set to empty
    string. Using default client kitchener.
12/20/04 16:47:09 nsrnmdev: SA quiesce: Succeeded
12/20/04 16:47:22 nsrnmdev: SA backup: Successfully
    obtained the subdirs/files of SA(s) for the backup
    CLASSPATH: \
"E:\Program Files\Legato\lsr\bin\lsrnmde.jar;
F:\Documentum\RTE\jdk\131_04\lib\tools.jar;
F:\Documentum\RTE\Shared\dfc.jar;F:\Documentum\R
\
TE\Shared\dfcbase.jar;F:\Documentum\RTE\Shared\
dctm.jar;F:\Documentum\RTE\Shared\dfcbase.jar;
F:\Documentum\RTE\Shared\workflow.jar;
F:\Documentum\RTE\Shared\log4j.jar;F:\Documentum\
\
RTE\jdk\131_04\lib\dt.jar;F:\Documentum\RTE\
jdk\131_04\lib\dt.jar;";
    DM_HOME:
"F:\Documentum\product\5.2";
    DOCBASE_NAME: babilon;
    DOCUMENTUM: "F:\Documentum";
    JAVA_PATH:
"F:\Documentum\RTE\jdk\131_04\bin";
    NMD_DB_FULL_BACKUP_CMD:
"E:\nmd\config\lsrnm_doc.bat";
    NMD_DB_INCR_BACKUP_CMD:
"E:\nmd\config\lsrnm_log.bat";
    NMD_DB_LOG_BACKUP_CMD:
"E:\nmd\config\lsrnmlog.bat";
    NMD_OBJECT_NAME: filestore_01;
    NMD_SCOPE: SA;
    NMD_USE_SNAPIMAGE: no;
    NMDDE_DM_DOCBASE: babilon;
    NMDDE_DM_PASSWD: <T5dr`nl;
    NMDDE_DM_USER: Administrator;
NMDDE_METADATA_OUTPUT_DEST:
"E:\nmd\config\filereports\DCB_";
    NSR_BACKUP_LEVEL: INCR;
    NSR_CLIENT: kitchener;
    NSR_DATA_VOLUME_POOL: Default;
    NSR_GROUP: Default;
    NSR_QUIET: YES;
    NSR_SERVER: birdman;
    OBJECT_NAME: filestore_01;
    SCOPE: SA;

```

```

12/20/04 16:47:34 nsrnmdsv: SA backup: TM SA save job
succeeded
12/20/04 16:47:38 nsrnmdsv: SA unquiesce: Succeeded
12/20/04 16:47:38 nsrnmdsv: Physical backup is completed.
Storing backup entries.
12/20/04 16:47:44 nsrnmdsv: NMD backup is fully
successful.
completed savetime=1103579264
12/20/04 16:47:44 nsrnmdsv: Generating filereports...
12/20/04 16:47:44 nsrnmdsv: Started filereport using
"F:\Documentum\RTE\jdk\131_
04\bin\java.exe" nsrnmdde -o query -M MD -t
"filestore_01:-2147467043:-2147467035"
12/20/04 16:47:53 nsrnmdsv: Finished filereport.
12/20/04 16:47:53 nsrnmdsv: Filereports are successful.

```

7.2. Scheduled NMD backup

To perform a scheduled NMD backup, specify the **nsrnmdsv** command in the Backup Command attribute of the Client resource on the NetWorker server. [“7.1. Ad hoc NMD backup” on page 1-44](#) provides more information on the **nsrnmdsv** command.

Ensure that the other NetWorker resources required for a scheduled backup are configured, as described in the *NetWorker Module for Documentum, Release 1.1, Administration Guide*.

7.3. Backup index query

To browse and display the NMD backup index information, enter the **nsrnmdnx** command. For example, if only the mandatory **-B** option is specified for the repository name, the **nsrnmdnx** command displays the records for every backup of the repository, starting with the most recent record and ending with the oldest.

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides more information on the **nsrnmdnx** command and its options.

Example 1-5 Sample nsrnmdnx commands for index searches

The following **nsrnmdnx** command searches on the NetWorker server jupiter for all the backup records of the repository docora92 that belong to the client mars:

```
nsrnmdnx -B docora92 -c mars -s jupiter
```

The following **nsrnmidx** command searches for all the incremental backups of scope SA for the storage area named filestore_01 in the repository docora92 that occurred after the last week:

```
nsrnmidx -B docora92 -M SA -a filestore_01 -L "last week"
-l incr
```

The following **nsrnmidx** command searches for all the backups of scope FTI_ALL containing the full-text indexes of the repository docora92 that occurred between last week and yesterday:

```
nsrnmidx -B docora92 -M FTI_ALL -L "last week"
-U yesterday
```

Example 1-6 Backup information displayed by the nsrnmidx command

The following **nsrnmidx** command searches the backup index on the client named almaza for backups of the repository dctm_syb:

```
nsrnmidx -B dctm_syb -c almaza -s aesir
```

The **nsrnmidx** command displays the following backup index information and shows the records for two backup of the specific storage area:

```
===== Record Head #1 =====
database_conn: almaza
database_name: db_dctm_syb
database_type: Sybase
DM_DOCBASE: dctm_syb
DM_HOME: /space2/dctm/product/5.2.5
DOCBASE_NAME: dctm_syb
DOCUMENTUM: /space2/dctm
FTI_COUNT: 2
FTI_END: 1101732286 (Mon Nov 29 07:44:46 2004)
FTI_START: 1101732269 (Mon Nov 29 07:44:29 2004)
JAVA_PATH: /space2/dctm/dfc/jdk/bin
NMD_USE_SNAPIMAGE: no
NSR_BACKUP_LEVEL: FULL
NSR_CLIENT: abc.xyz.com
NSR_DATA_VOLUME_POOL: Default
NSR_GROUP: Default
NSR_SERVER: aesir
SA_COUNT: 0
SCOPE: FTI_ALL
TAG: 1101732296
type: DCTM_PHY
===== Record Head #2 =====
database_conn: almaza
database_name: db_dctm_syb
database_type: Sybase
DM_DOCBASE: dctm_syb
```

```
DM_HOME: /space2/dctm/product/5.2.5
DOCBASE_NAME: dctm_syb
DOCUMENTUM: /space2/dctm
FTI_COUNT: 0
JAVA_PATH: /space2/dctm/dfc/jdk/bin
NMD_USE_SNAPIMAGE: no
NSR_BACKUP_LEVEL: FULL
NSR_CLIENT: abc.xyz.com
NSR_DATA_VOLUME_POOL: Default
NSR_GROUP: Default
NSR_SERVER: aesir
SA_COUNT: 5
SA_END: 1101393077 (Thu Nov 25 09:31:17 2004)
SA_START: 1101393056 (Thu Nov 25 09:30:56 2004)
SCOPE: SA_ALL
TAG: 1101393085
type: DCTM_PHY
```

7.4. Change of directory ownership and permissions during a restore

You must run an NMD restore as the root user on UNIX or Linux, or a member of the Administrators group on Windows. During a restore or disaster recovery of a full-text index, storage area, or installation and configuration file component, if an upper-level parent directory of the component is missing, the directory is automatically re-created and the ownership and permissions of the directory are changed. For example, the ownership of a parent directory is changed to the user that restores the backup, such as the root user on UNIX.

If `/usr/abc/file1` is being restored and the `/usr/abc` directory is missing on the Documentum Server, `/usr/abc` is re-created during the restore and the `/usr` directory ownership is changed to the user that performs the restore.

The user that performed the restore must manually change the ownership and permissions of the directory back to those of the original user (dminstall user) by using the appropriate operating system commands.

Change the directory ownership by one of the following methods:

- ◆ On UNIX or Linux, enter the following command:

```
chown dminstall_userid dir_pathname
```

where *dir_pathname* is the pathname of the parent directory.

- ◆ On Windows, perform the following:
 1. In Windows Explorer, right-click the parent directory and select **Properties**.
 2. On the **Properties** screen, click **Security**.
 3. Select the correct name under **Group or user names**, and click **Advanced**.
 4. On the **Advanced Security Settings** screen, click **Owner**.
 5. Select the correct name under **Change owner to**, and click **OK**.

7.5. Content file restore

To browse Documentum file reports (optionally generated at the end of storage area backups) and display a list of content files from the file reports that match the specified search criteria, enter the **nsrnmocr** command. The command can optionally restore the content files found during the browse operation.



IMPORTANT

Restore the content files to a relocation area, and hand them to the user. The user then re-adds the corresponding document to the repository by using the restored files.

If the `NMD_FILEREPORT_INCLUDED` and `NMDDE_METADATA_OUTPUT_DEST` parameters are set properly in the configuration file before a storage area backup, Documentum file reports containing content file metadata are generated at the end of the backup. `NMD_FILEREPORT_INCLUDED` must *not* be set to `NO`. The backup must be of scope `ALL`, `SA`, or `SA_ALL`. A separate file report is generated for *each* storage area backed up.

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides more information on the **nsrnmocr** command and its options.

Example 1-7 Sample nsrnmocr command for file report searches

The following **nsrnmocr** command searches file reports for the storage area `filestore_01` of the repository `babilon`. The file reports searched are those created on May 22, 2006. The search looks for

document names containing the (case-insensitive) string "mobile", with the folder name mobilitydocs and version number 1.0. The documents were checked in from the client waterloo.

```
nsrnmocr -B babilon -A filestore_01
-P "F:\Documentum\stores\Filereports\current"
-T "20060522" -V -c "waterloo" -d "mobile"
-f "mobilitydocs" -v "1\1.0"
```

Note: In the regular expression "v1\1.0", the escape character \ is included before the dot.

7.6 Database and log restore

The **nsrnmocr** command can be used to search for the required database or database log backup in the backup index before starting a restore. For example, the index might be searched for a backup of scope ALL, DB, or DB_LOG that contains the required database or logs. The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides more information on the **nsrnmocr** command.

To restore database files and database logs from an NMD backup, use the appropriate method, depending on how the files were originally backed up. For example, if the database files were backed up by using a NetWorker Module or user-created scripts, the corresponding method must be used for the restore.

Example 1-8 Oracle database restore

When an Oracle database is backed up during an NMD backup by using the NetWorker Module for Oracle, you can restore the Oracle database to a point in time as follows:

1. Enter the **nsrnmocr** command to browse the backup index and select the correct backup that includes the Oracle database. From the **nsrnmocr** command output, record the SA_START value that specifies the start time of the backup record of type SA.

Note: If a remote database host is used and its system time is out-of-sync with the Documentum host system time, then the times reported by **nsrnmocr** must be adjusted accordingly. The reported times for database and storage area backups are based on the Documentum host system time.

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides more information on the **nsrnmocr** command.

2. Restore the database files and (if required) logs from the backup by using an RMAN restore script that includes the appropriate Oracle restore commands. The following sources provide more information on RMAN restore scripts:
 - *NetWorker Module for Oracle Administrator's Guide*
 - Oracle backup and recovery documentation
3. Recover the database to the point in time specified by the SA_START value in the **nsrnmndx** command output from [step 1](#). For the recovery, use the appropriate Oracle commands to apply the archived redo logs and online redo logs. There are two ways to use the Oracle recovery commands:
 - Include the Oracle commands in the RMAN restore script in [step 2](#).
 - After the RMAN restore script of [step 2](#) has completed successfully, enter the Oracle commands at the operating system command line.

The Oracle backup and recovery documentation provides more information on the Oracle commands.



IMPORTANT

The database must be recovered to the point in time specified by SA_START in [step 1](#). Do *not* recover the database beyond the SA_START time of the backup by applying the log files.

7.7. Other component restore

All defective full-text index and storage area components being restored should be offline and quiesced before the restore. The components can be quiesced manually outside of NMD, for example, through the Documentum Administrator GUI. Only as a last restore for Documentum 5.2.x, set NMD_RECOVER_QUIESCE=YES to force **nsrnmdrs** to attempt the quiesce operation.

Note: On Documentum 5.3.x systems, NMD_RECOVER_QUIESCE does *not* apply, and the **nsrnmdrs** command issues screen output stating that quiescing is not performed even if NMD_RECOVER_QUIESCE is set to YES.

To browse the NMD backup index and restore selected components from a specified backup, enter the **nsrnmdrs** command, as either the root user on UNIX or a member of the Microsoft Windows

Administrators group. The command restores Documentum Server storage areas, full-text indexes, and installation and configuration files.



IMPORTANT

It is strongly recommended that you do *not* restore a backup over the original, existent source data. Use a relocation directory.

Alternatively, perform the following:

1. Place the corresponding repository component in offline mode.
2. Rename the root directory of the component.
3. Re-create the root directory.
4. Perform the restore.

When restoring the database component, restore the backup to the point in time at the end of the database backup. To determine this point in time, use the `nsrnmndx` command. Do *not* roll forward the database by applying logs.

When restoring a storage area, randomly open some of the files to verify that there is no corruption.

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides more information on the `nsrnmdrs` command and its options.

Example 1-9 Sample nsrnmdrs commands for different types of restores

The following `nsrnmdrs` command restores the latest backup that contains all the storage areas of the repository docora92:

```
nsrnmdrs -B docora92 -f /space10/home/dmuser/nmd.cfg
-C SA
```

The following `nsrnmdrs` command restores the storage area named filestore_01 from the latest backup that occurred between 12/12/04 and 12/13/04 and relocates the storage area to the directory /space2/sa_dec12:

```
nsrnmdrs -B docora92 -f /space10/home/dmuser/nmd.cfg -L
"12/12/04" -C SA -a filestore_01 -d /space2/sa_dec12
```

The following `nsrnmdrs` command restores all the full-text indexes from the latest backup that occurred before today:

```
nsrnmdrs -B docora92 -f /space10/home/dmuser/nmd.cfg
-C FTI -U "today"
```


Example 1-10 Restore information displayed by the nsrnmldr command

The following **nsrnmldr** command searches the backup index and restores the latest backup of the storage area named **dm_notes_file_store3** of the repository **dm_notes**, relocating it to the **/export/home/dm20/logs/temp_sa** directory:

```
nsrnmldr -B dm_notes -f /nsr/nmd/nmd_config
-C SA -a dm_notes_file_store3
-d /export/home/dm20/logs/temp_sa
```

The **nsrnmldr** command displays the following information during the restore process:

```
12/17/04 14:50:59 nsrnmldr: Note: This session will
      restore the backup with the latest timestamp
12/17/04 14:50:59 nsrnmldr: --- Component(s) selected for
      restore ---
nsrnmldr: --- Storage Area Start ---
nsrnmldr: =====
nsrnmldr: --end Save Time:Fri Dec 17 13:30:56 2004--
nsrnmldr: --name:dm_notes_file_store3--
nsrnmldr:
      --path:/export/nfs2-3/dm20/data/dm_notes/storage_03/00
      000001--
nsrnmldr: --start Save Time:Fri Dec 17 13:27:08 2004--
nsrnmldr: --ticket:-2147045069--
nsrnmldr: --type:SA--
nsrnmldr: =====
12/17/04 14:50:59 nsrnmldr: Would you like to proceed
      with the restore ? [y/n]:y
12/17/04 14:59:60 nsrnmldr: Restore successful
```

8. Troubleshooting tips

This section contains general troubleshooting tips and specific recommendations to address NMD backup and restore issues.

8.1. How to verify the NMD release information

To verify the NMD release information, use the following procedure to display the release of the **nsrnmldrsv** program or NMD package:

- ◆ On UNIX or Linux, enter one of the following commands, where *filenm* is the full pathname of the **nsrnmldrsv** program file:

```
what filenm
strings filenm | grep "@(#)"
pkginfo -l LGTONmd (Solaris only)
```

- ◆ On Windows:
 - a. In Windows Explorer, locate the **nsrnmadv.exe** program file.
 - b. Right-click the file icon, and select **Properties**.
 - c. In the **Properties** window, click **Version** and view the release information.

8.2. How to obtain debug information

Use the information in this section to help diagnose and resolve issues.

To obtain debug information:

1. On the Documentum Server, set the `NSR_DEBUG_LEVEL` and the `NSR_DEBUG_FILE` parameters in the NMD configuration file, as follows:

```
NSR_DEBUG_LEVEL=9
```

```
NSR_DEBUG_FILE=path_to_logfile or default
```

If `NSR_DEBUG_FILE` is set to default, the following log files are generated:

- `/nsr/applogs/nsrnmadv.pid.log` or `NetWorker_install_path\nsr\applogs\nsrnmadv.pid.log` for a backup operation
 - `/nsr/applogs/nsrnmadvr.pid.log` or `NetWorker_install_path\nsr\applogs\nsrnmadvr.pid.log` for a restore operation
2. To have the problematic NMD backup or restore operation stop when it first encounters an error, set the following parameter in the NMD configuration file:

```
NSR_STOP_ON_ERROR=YES
```

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides information about this parameter.

3. Run the problematic NMD backup or restore operation.

Debug entries are created in the log file, which is located in the path that was specified in the [step 1](#).

Note: When a **save**, **nsrndmp_save**, or **recover** task produces an error or when `NSR_DEBUG_LEVEL=n` (where *n* is greater than 0), the task's log file remains in the `/nsr/applogs` or `NetWorker_install_path\nsr\applogs` directory.

4. From the end of the log file, search the log file entries backwards (most recent to oldest) for the following types of text strings: error, warn, fail, abort, bypass.

- On UNIX, you can search the entries with the following type of command:

```
egrep -i "error|warn|fail|abort|bypass"
      log_filepath
```

- On Windows, you can search the entries with the following types of commands at the DOS prompt:

```
find "error" log_filepath
find "warn" log_filepath
find "fail" log_filepath
find "abort" log_filepath
find "bypass" log_filepath
```

Note: The NetWorker server also writes diagnostic information from an ad hoc or scheduled backup to specific log files on the NetWorker server. The *NetWorker Administration Guide* provides more information about these log files.

5. Search for the error message and refer to its corresponding resolution. The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides information about error messages.

8.3. NMD troubleshooting scenarios

This section contains solutions to some common issues that are related to an incorrect NMD configuration.

8.3.1. DOCUMENTUM_SHARED is required on UNIX and Linux

For troubleshooting the **nsrnmdd** process at the command line on UNIX and Linux, ensure that the environment variable `DOCUMENTUM_SHARED` is set to the pathname of the directory where Documentum Foundation Classes (DFC) is installed.

8.3.2. Cannot connect to the database

- ◆ Ensure that the path to `nsrnmdd.jar` is specified in the `CLASSPATH` parameter.
- ◆ Ensure that the Java path specified in the `JAVA_PATH` parameter is the same Java path used by the Documentum Content Server.
- ◆ Ensure that the username specified in the `NMDDE_DM_USER` parameter is a valid Documentum Content Server user with administration-level privileges.
- ◆ Ensure that the password in the `NMDDE_DM_PASSWD` parameter was set with the **-P** option of the `nsrnmddsv` command at the operating system command line. Additionally, ensure that the password is correct for the repository user specified by the `NMDDE_DM_USER` parameter. `NMDDE_DM_PASSWD` must *not* appear more than once in the configuration file.
- ◆ With Documentum 5.3.x, if the full-text index server is on a *separate* host from the Content Server, ensure that the password in the `NMD_FTI_PASSWD` parameter was set with the **-P** and **-M FTI** options of the `nsrnmddsv` command at the operating system command line. Additionally, ensure that the password is correct for the repository user specified by the `NMD_FTI_USER` parameter. `NMD_FTI_PASSWD` must *not* appear more than once in the configuration file.
- ◆ Ensure that the following parameters are set correctly:
 - `DOCUMENTUM`
 - `DM_HOME`
 - `NMDDE_DM_DOCBASE`
- ◆ Test the connection by using either of the following:
 - The Documentum `idql` or `iapi` utility (`idql32` and `iapi32` on Windows). For example, enter the following command:

```
idql -Uusername -Ppassword repository_name
```

where:
 - `username` is the username from the parameter `NMDDE_DM_USER`.
 - `password` is the user password from the parameter `NMDDE_DM_PASSWD`, but in *unencrypted* form.
 - `repository_name` is the repository name from the parameter `NMDDE_DM_DOCBASE`.

- The NMD `nsrnmddde.jar` binary:
 - a. Change to a temporary test directory. For example, enter the following command on Solaris:


```
cd /tmp/nmd_test/
```
 - b. Copy the `nsrnmddde.jar` and `log4j.properties` files from the NetWorker bin directory to the temporary test directory. For example, enter the following command on Solaris:


```
cp /usr/sbin/nsrnmddde.jar
   /usr/sbin/log4j.properties .
```
 - c. Edit the `log4j.properties` file and change INFO to DEBUG.
 - d. Export the following NMD parameters as environment variables by using the values from the NMD configuration file *except* where noted:
 - DOCUMENTUM
 - DM_HOME
 - CLASSPATH (add the `nsrnmddde.jar` pathname as the first path; in the Solaris example, add `/tmp/nmd_test/nsrnmddde.jar` as the first path)
 - JAVA_PATH
 - PATH
 - NMDDE_DM_DOCBASE
 - NMDDE_DM_USER
 - NMDDE_DM_PASSWD (set to the *unencrypted* password)
 - e. Export the parameter `NMDDE_OUTPUT_FILE`, for example, set to `/tmp/nmd_test/nsrnmddde.out`.
 - f. Enter the following commands, where the **-verbose** option is optional and the screen output is captured to a log file by appending `>>logfile 2>&1` to the first command:

```
java -verbose nsrnmddde -h >>logfile 2>&1
java nsrnmddde -o query -C
java nsrnmddde -o query -D
java nsrnmddde -o query -s ALL
java nsrnmddde -o query -f ALL
```

- g. Check the screen output (or optional log file) for error messages.
- h. Check the output file specified by NMDDE_OUTPUT_FILE for valid output that describes the database.

8.3.3. Full-text index backs up slowly

With Documentum 5.2.x *only*, the full-text index (FTI) may have many large log files. This can increase the backup time. To decrease the number of log files, set the NMD_DIRECTIVES_FILE parameter to skip FTI log files during the backup. [“10.1.1. Excluding the full-text index log files with Documentum 5.2.x” on page 1-64](#) provides more information.

8.3.4. Installation and configuration files back up slowly

The installation and configuration files backup may include many unnecessary directories, such as older Documentum Content Server software versions. Specify which subdirectories to back up by using the NMD_ICF_SUBDIRS parameter.

8.3.5. Backup or recovery is slow

Insufficient data throughput can cause poor backup performance. To improve throughput, ensure that the device parallelism/sessions values are sufficiently high for the following parameters:

- ◆ NSR_PARALLELISM
- ◆ NMD_SAVE_PARALLELISM
- ◆ NMD_RESTORE_PARALLELISM

If NSR_PARALLELISM is set, its value overrides NMD_SAVE_PARALLELISM and NMD_RESTORE_PARALLELISM. The default value of each parameter is 8. [“6.7. NMD configuration file” on page 1-32](#) provides other recommended values of these parameters.

The *NetWorker Administration Guide* provides more information about device parallelism and sessions attributes.

8.3.6. Backup components were not fully saved

When NMD_PARTIAL_SUCCESS=YES is set, backups in which some of the target components were not fully saved can still be browsed through the **nsrnmidx** command and/or used to restore data through the **nsrnmdrs** or **nsrnmldr** command.

To determine which components were not fully backed up:

- ◆ Browse the backup logs and/or screen messages for “error: task failed for component(s): <list of one or more component names>.” The messages altogether give the names of all failed or partially successful components.
- ◆ Browse the backup index through the **nsrnmidx -v** command.
 - A full-text index or install/configuration file backup without a save time value has totally failed.
 - A storage area backup (with SnapImage) without a save time value has totally failed.
 - A storage area backup (without SnapImage) without as many save times values as the listed (positive) number of subdirectories has totally or partially failed.
 - All other components have been fully backed up or bypassed. For example, a non-existent root directory or a component with no incremental changes is bypassed.

8.4. SnapImage issues

This section includes troubleshooting tips that are specific to the SnapImage Module.

8.4.1. Identifying SnapImage issues

SnapImage issues can be identified by doing the recommended **nsrndmp_save** and **recover** tests. You may need to deploy versions of **nsrndmp_save** and **recover** that have been provided specifically for NMD by an EMC Customer Support Representative.

8.4.2. Backups are slow

Put the storage areas on their own dedicated file system to avoid the backup of unrelated disk write activity.

8.4.3. Snapshot backup of storage areas is suspended when memory is low

On a Documentum Server host running Solaris, if the available memory is low, a snapshot backup of storage areas might become suspended before it completes. A corresponding error message is not generated unless the debug level for the backup is set to a high value.

If an error message for the backup suspension is generated in the backup logs, the error incorrectly states that the backup is "out of space." The error should state that the backup is "out of memory."

To avoid problems during a storage area backup, ensure that the Solaris system has at least 1 GB of memory.

8.4.4. Oracle Server crashes or hangs on Windows

Move the Oracle control file and redo logs to a file system that does *not* host any repository storage areas.

8.4.5. Potential problem with the restore of storage areas that contain numerous files

On Solaris (*without* the required SnapImage hotfix) and HP-UX, if you attempt to use the **nsrnmdrs** command to restore a snapshot backup (performed by the SnapImage Module) of one or more storage areas from the same file system where the backup contains approximately 1 million or more files, the restore becomes suspended. The EMC bugs LGTpa73035 and LGTpa79120 contain information on this issue.

Note: For all other storage area backups (whether or not the SnapImage Module performed the backups), restore the backups by entering the **nsrnmdrs** command with the appropriate options. The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides more information.

If one or more storage areas from the same file system were backed up by the SnapImage Module on either Solaris (without the SnapImage hotfix) or HP-UX and the storage areas contain a total of approximately 1 million or more files, restore the storage area backup as follows:

1. Prepare a separate unmounted file system or raw partition to be used as a staging file system for the restore. The file system must contain at least the same amount of disk space as the backed-up storage area file system.
2. Determine the save set ID of the storage area backup by using one of the following methods:

- To use the NetWorker Administrator program on UNIX or NetWorker User program on Windows, follow the instructions in the *NetWorker Administration Guide*.

From the program display, select the correct save set name and save set version with the required save time, and determine the save set ID value (to be used in [step 3](#)).

- To use the **mminfo** command, enter the following:

```
mminfo -s NetWorker_server -c NetWorker_client
-N "save_set_name"
-r ssid,name,savetime,nsavetime,totalsize,nfiles
```


From the **mminfo** command output, select the row that contains the required save time, and determine the save set ID value (to be used in [step 3](#)).

The *NetWorker Command Reference Guide* or the **mminfo** man page provides more information on the **mminfo** command.

- To use the **nsrnmidx** and **mminfo** commands, perform the following:
 - a. Enter the required **nsrnmidx** command:

```
nsrnmidx -v {-M ALL|SA} [-L lower_bound_time]
          [-U upper_bound_time]
```

From the **nsrnmidx** command output, determine the save time of the storage area backup.

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* or the **nsrnmidx** man page provides more information on the **nsrnmidx** command.

- b. Enter the required **mminfo** command:

```
mminfo -qsavetime=savetime
        -r ssid,savetime,nsavetime
```

where *savetime* is the save time determined in [step a](#). From the **mminfo** command output, determine the save set ID value (to be used in [step 3](#)).

The *NetWorker Command Reference Guide* or the **mminfo** man page provide more information on the **mminfo** command.

3. Perform the restore to the staging file system by entering the following command:

```
nsrndmp_recover -s NetWorker_server
                -c NetWorker_client -r raw_device_pathname
                -S save_set_ID
```

where:

- *raw_device_pathname* is the raw device pathname of the staging file system, prepared in [step 1](#).
- *save_set_ID* is the save set ID of the storage area backup, obtained in [step 2](#).

For example, enter the following `nsrndmp_recover` command to perform the restore:

```
nsrndmp_recover -s nwsvr -c nmsapclnt01 -r
/dev/rdisk/c1t3d0s5 -S 2747749942
```

4. Mount and validate the restored staging file system.
5. Disable the storage areas on the original file system.
6. Copy the required storage area directories from the staging file system to the original file system by using a standard operating system utility, such as the `cp -pr` command.

For example, enter the following `cp` command to copy the directories:

```
cp -pr /stagingfileSYS/dir_SA1 /srcfileSYS/dir_SA1
```

8.4.6. SnapImage does not support the IBM RDAC driver on UNIX or Windows 2003

Support for the IBM RDAC driver is provided by the latest SnapImage release on Windows 2000 only.

SnapImage does *not* support the IBM RDAC driver on UNIX or Windows 2003.

8.4.7. Other SnapImage limitations

The following documents provide information on additional limitations:

- ◆ *NetWorker SnapImage Module, Release 1.0, UNIX Version, Installation and Administrator's Guide*
- ◆ *NetWorker SnapImage Module, Release 1.6.1, UNIX Version, Release Supplement*
- ◆ *NetWorker SnapImage Module, Release 2.0, Microsoft Windows Version, Installation and Administrator's Guide*
- ◆ *NetWorker SnapImage Module, Release 2.0.2, Microsoft Windows Version, Release Notes*

8.5. Consideration for recovering the repository

If the directories specified in the `DOCUMENTUM` or `DM_HOME` parameter were re-created during a repository recovery, reset the ownership of these directories to the user who installed Documentum.

8.6. Troubleshooting tools

To avoid negatively affecting other applications, an EMC Customer Support Representative may ask that you set one or more of the following NMD troubleshooting parameters:

- ◆ `NMD_SI_SAVE_CMD=fullpath_of_patched_binary`
- ◆ `NMDTM_SI_SAVE_CMD=...`
- ◆ `NMDTM_RESTORE_CMD=...`

9. Optimization recommendations

9.1. Parallelism for restores

The following types of restores use the **recover -t** command:

- ◆ Restores of storage areas with SnapImage
- ◆ Restores of storage areas, full-text indexes, and installation and configuration files without SnapImage

Note: Restores of full-text indexes and installation and configuration files with SnapImage use the **recover -S** command.

Currently on UNIX, the use of a single **recover -t** task provides inherent parallelism during a restore, both with a single tape and across multiple concurrent tapes. As a result, it may be more beneficial to set the restore parallelism to a value of one or two, instead of a higher value that can lead to tape head thrashing.

9.2. SnapImage optimization

SnapImage backups are most beneficial under the following conditions:

- ◆ When there is sufficient memory, swap space, and disk space, as described in [“4.1. SnapImage prerequisites” on page 1-21](#).
- ◆ When there is no or minimal disk write activity on the storage area file systems.

9.3. Other optimization

The following provide more information on optimization settings:

- ◆ [“6. Configuration overview” on page 1-29](#)
- ◆ [Appendix A, “NMD Consulting Pre-site Qualifier \(CPSQ\)”](#)

10. Special use case recommendations

10.1. Backups

10.1.1. Excluding the full-text index log files with Documentum 5.2.x

With Documentum 5.2.x *only*, to speed up a full-text index backup by excluding the logs, you can perform the following:

1. In a convenient location on the Documentum Server, create a NetWorker directives file that contains only the following line:

```
<< / >> +skip: dm_ftwork_dir
```

For example, create a file C:\temp\nmd_directive.txt that contains the directive line.

The *NetWorker Administration Guide* and the **nsr_directive(5)** man page provide more information on NetWorker directives.

2. In the NMD configuration file, set the parameter NMD_DIRECTIVES_FILE to the complete pathname of the NetWorker directives file created in [step 1](#).

For example, set the parameter in the configuration file as follows:

```
NMD_DIRECTIVES_FILE=C:\temp\nmd_directive.txt
```

[“6.7. NMD configuration file” on page 1-32](#) provides more information on the configuration file parameters.

10.1.2. Scheduling repository backups

The recommended NetWorker schedule for repository backups is as follows:

- ◆ Level full on late Saturday or early Sunday, with a default file report dump.
- ◆ Level incr (also known as level 1) on all other days, with a default incremental file report dump.

10.1.3. Scheduling separate backups of installation and configuration files

NMD can back up the installation and configuration files, either as part of a repository system backup or as a separate backup, depending on the setting of the parameter NMD_ICF_INCLUDED.

Unless your environment requires the backup of the installation and configuration files to be part of the system backup, configure the backup of those files to occur separately.

By default, the backup includes the installation and configuration files discovered automatically by the software. To specify nondefault directories and files to be backed up, you can set the additional parameter `NMD_ICF_SUBDIRS` or `NMD_ICF_SUBDIRS_AUGMENT`.

Note: With Documentum 5.2.x, backup of the installation and configuration files might not automatically discover all the required directories and files in certain situations, for example, in certain customized installations. With Documentum 5.3.x, NMD does *not* autodiscover any part of the installation and configuration files during a backup, and the parameter `NMD_ICF_SUBDIRS` is *mandatory*.

The schedule of installation and configuration file backups should be infrequent because the files do not change as often as other repository components. An installation and configuration file backup should be in its own savegroup session, and should *not* overlap with a regular NMD backup.

A level full save set is always generated for an installation and configuration file backup.

10.1.4. File report generation and purging

If the `NMD_FILEREPORT_INCLUDED` and `NMDDE_METADATA_OUTPUT_DEST` parameters are set properly in the configuration file before a storage area backup, Documentum file reports containing content file metadata are generated at the end of storage area backup. `NMD_FILEREPORT_INCLUDED` must *not* be set to `NO`. The backup must be of scope `ALL`, `SA`, or `SA_ALL`. A separate file report is generated for *each* storage area backed up.

An extra file system of sufficient disk space should be used for the file reports. The amount of space required depends on the number of repository documents and the frequency of full and incremental dumps. Typically, at least 2 GB of extra space should be reserved. The Documentum Server administrator needs to handle the periodic purging of old documents.

To browse the content file metadata in the file reports, enter the `nsrnmocr` command at the operating system command line. “7.5. Content file restore” on page 1-49 provides more information.

A file report filename has the following format:

`DCB_repository-name_storage-area-name_level_data-ticket_timestamp.txt`

Example 1-11 A file report filename

A sample file report filename is as follows:

`DCB_dmnotes_filestore_01_FULL_-2147046105_20041214194643.txt`

where:

- ◆ `dmnotes` is the repository name.
- ◆ `filestore_01` is the storage area name.
- ◆ `FULL` is the backup level.
- ◆ `-2147046105` is the data ticket.
- ◆ `20041214194643` is the timestamp.

Example 1-12 Contents of a file report

A sample of the contents of a file report is as follows:

```
U 0900000180481992 /lsultan/WP Architecture.ppt
Linda Sultan 06/23/2004 15:05:02 1.0 2jpeg_th
C:\Program Files\Documentum\MediaServer\cache\0000000001
.file-933889016.jpg abc.xyz.com
/export/nfs2-2/dm20/data/dm_notes/thumbnail_storage_01/
00000001/80/00/0d/e6.jpg
U 0900000180481992 /lsultan/WP Architecture.ppt
Linda Sultan 06/23/2004 15:05:02 1.0 2
jpeg_story C:\WINNT\TEMP\1088029318656batchFile.tar
abc.xyz.com /export/nfs2-2/dm20/data/dm_notes/
thumbnail_storage_01/00000001/80/00/0d/e8.jpg
```

10.1.5. LAN, SAN, and local backup devices

For the best performance in backing up and restoring a repository, use a local NetWorker device or a SAN NetWorker device.

Backups to a remote NetWorker device (on a remote NetWorker server or storage node) are limited by the network bandwidth, for example, approximately 8 MB/s for a 100 Mb/s LAN.

10.2. Restores

10.2.1. Repository and disaster recovery

The *NetWorker Module for Documentum, Release 1.1, Administration Guide* provides information on repository and disaster recovery.

10.2.2. Intentionally inconsistent restore

An inconsistent repository restore occurs when only part of a repository is restored. This type of restore introduces inconsistencies in the repository, as explained in [“2.2. Consistent repository backup” on page 1-4](#).

For instance, to avoid the need to restore the entire repository, you decide to restore a single storage area. The restored storage area does not contain documents that the database refers to.

Documentum application errors should be expected and manually resolved. For example, run a DQL script to remove all bad references to nonexistent documents. The relevant users should also be notified.

10.2.3. Using NetWorker directly

For diagnostic purposes or workarounds to problems encountered, an EMC Customer Support Representative may ask you to use the NetWorker **recover** program to restore repository data, although it is not supported by NMD.

[“3.3. How NMD interfaces with external modules” on page 1-14](#) provides information on save set formations and the **recover** commands to use.

Alternatively, you can use the **recover** command to browse for backup history and perform actual restores. This type of operation is *not* supported by NMD.

Note: If one or more storage areas from the same file system were backed up by the SnapImage Module and the storage areas contain a total of approximately 1 million or more files, manually restore the storage areas with the **nsrndmp_recover** command to a staging file system, as described in [“8.4.5. Potential problem with the restore of storage areas that contain numerous files” on page 1-60](#).

10.3. UNIX symbolic links and mount points

NMD automatically expands and follows symbolic links that are in a root directory path of the storage area, full-text index, or installation and configuration file component. (This is any path listed in the NMD log files when tracing is turned on.)

NMD does *not* follow symbolic links that are *under* the root directory path of a storage area, full-text index, or installation and configuration file component. These types of “inner” symbolic links should never occur for a storage area or full-text index component, unless a user manually creates them. NMD can indirectly handle “inner” symbolic links for the installation and configuration file component when the fully expanded directory and file paths are listed in one of the following parameters:

- ◆ NMD_ICF_SUBDIRS
- ◆ NMD_ICF_SUBDIRS_AUGMENT

Similar to symbolic links, NMD does not follow mount points during a backup if they are *under* the root directory path of a storage area, full-text index, or installation and configuration file component. Again, these types of mount points should never occur, unless a user manually creates them.

11. Reconfiguration

11.1 Repository

To add or remove a storage area or full-text index, no additional configuration is required. NMD automatically detects the new repository configuration.

11.2 SnapImage

The following sources provide information on adding or removing disks or an extra cache:

- ◆ *NetWorker SnapImage Module, Release 1.0, UNIX Version, Installation and Administrator’s Guide*
- ◆ *NetWorker SnapImage Module, Release 1.6.1, UNIX Version, Release Supplement*
- ◆ *NetWorker SnapImage Module, Release 2.0, Microsoft Windows Version, Installation and Administrator’s Guide*
- ◆ *NetWorker SnapImage Module, Release 2.0.2, Microsoft Windows Version, Release Notes*

12. Related documentation

The information in this guide may be supplemented with the following resources:

- ◆ Appropriate versions of the documentation for the following EMC products:
 - NetWorker
 - NetWorker Module for Documentum
 - NetWorker Module for Microsoft SQL Server
 - NetWorker Module for Oracle
 - NetWorker Module for Sybase
 - NetWorker SnapImage Module
- ◆ UNIX man pages (for NetWorker and NetWorker Module commands)
- ◆ Documentum Content Server documentation
- ◆ Appropriate database (Oracle, Microsoft SQL Server, or Sybase) backup and recovery documentation

The complete set of EMC documentation is provided in PDF form on the Documentation Suite CD-ROM shipped with the EMC software. The *most* up-to-date EMC documentation is available on the website at <http://Powerlink.EMC.com>.

NMD Consulting Pre-site Qualifier (CPSQ)

The information provided on this form will assist EMC in planning for project initiation. It is important that all questions are answered completely and accurately to ensure site readiness before consultants arrive on site. Please complete this form in its entirety and submit electronically as a Microsoft Word file to the required address.

Business Assessment		
Installation environment	Production ____	Test ____
What is the target installation date?		

Account Information			
Date:			
Company Name:			
City, State:			
Country, if not U.S.:			
Site ID:			
	Name	Telephone	Email
Customer Project Leader:			
ESG Account Executive:			
ESG SE:			

Account Information			
ESG Consultant:			
EMC Account Executive:			
EMC TC/CSL:			

Existing NetWorker Backup Solution	
General	Completed?
Create or designate a Group to use for NMD backup	
Create or designate a Pool to use for NMD backup	
Create or designate a Schedule to use for NMD backup	

NetWorker Module for Documentum (Sheet 1 of 2)	
General	Completed?
Is all the required licensing available and ready for installation?	
What is the operating system version?	
What is the database version?	
Is there any clustering or failover involved?	
Database layout: One repository per server? Or more than one repository per server?	
Is the database on the same server as Documentum?	
Will NetWorker SnapImage be used for backups? (If yes, complete the SnapImage section)	
Verify hardware compatibility in the EMC Hardware Compatibility Guide and Release Supplements for all modules	
Upgrade and patch all existing software to supported versions	
Upgrade and test any required hardware patches	
Configure and test all new backup media, for example, jukeboxes, arrays, NAS devices	
Set network interfaces on Documentum and NetWorker servers to Full Duplex	
Documentum server name:	
NetWorker server name (if different):	

NetWorker Module for Documentum (Sheet 2 of 2)	
NetWorker storage node name (if used):	
NetWorker installation path:	
Documentum installation path:	
Database installation path:	
NetWorker server Administrator name:	
Documentum installation owner user name:	
Database user with SA rights name:	
NetWorker Group:	
NetWorker Pool:	
NetWorker Schedule:	

NetWorker SnapImage Module	
General	Completed?
What is the operating system version where the Documentum storage area resides?	
Total number of files in the storage area?	
What is the average file size in the storage area?	
Does the total file system size for the storage area exceed 2 terabytes?	

NetWorker Module for Oracle ^a	
General	Completed?
Register the database name in the Recovery Catalog	
Verify that the database is in archivelog mode	
Create or designate a user with SA rights to use with RMAN	

a. Please note that all operating system, database, and application versions that exist in the customer environment must meet the minimum required rev levels in order for this solution to be supported by EMC. To ensure that all EMC software components in this solution meet the necessary compatibility requirements click on the following link: <http://portal2.legato.com/resources/compatibility/>.

This appendix describes the best practices for using the NetWorker software that works with the NMD software solution.

Note: The *Software Compatibility Guide* on the EMC website at <http://Powerlink.EMC.com> provides details on the NetWorker software releases that NMD supports.

This appendix includes the following sections:

- ◆ B.1. Solution methodology: An industry best practice B-2
- ◆ B.2. What is in a datazone? B-2
- ◆ B.3. Single datazone B-3
- ◆ B.4. Network best practices B-8
- ◆ B.5. SAN best practices B-9
- ◆ B.6. Business best practices B-10
- ◆ B.7. NetWorker best practices B-10
- ◆ B.8. NetWorker 7.x best practices B-11

B.1. Solution methodology: An industry best practice

The primary components of a best practices-based solution methodology are as follows:

- ◆ Define requirements — Determine the Service Level Objectives (SLOs) for backup and restore and the roles and responsibilities of personnel, as well as DR/BC processes.
- ◆ Design and architecture — Based on SLOs. Design a solution in the environment to meet defined requirements (template design for repeatable deployment).
- ◆ Implement — Execute a pilot in order to demonstrate that the design meets requirements.
- ◆ Roll out — Roll pilot into production, and train personnel to support the environment.
- ◆ Document — Create runbook documentation detailing operations and restore scenarios so that the backup/restore solution is supportable and meets SLOs.
- ◆ Knowledge transfer — Perform hands-on transfer of knowledge with the customer and designated personnel for the critical administrative tasks and required functions.
- ◆ Continuous improvement — Engage EMC Professional Services (health checks) at regular intervals to validate that the deployment:
 - Continues to meet requirements.
 - Scales.
 - Integrates new environments.

B.2. What is in a datazone?

B.2.1. Datazone requirements examples

Application and technology specific requirements are as follows:

- ◆ SAP
- ◆ Oracle
- ◆ EMC Timefinder[®]/SRDF

- ◆ Library sharing and tape rotation
- ◆ Availability
- ◆ Budget and available technology
- ◆ Operations staff and skill set (roles and responsibilities)

B.2.2. Datazone conclusions

The four most often neglected keys to success are as follows:

- ◆ Perform detailed implementation planning.
- ◆ Define use requirements first, then market, then build.
- ◆ Run a pilot.
- ◆ Document the design and operation.

B.3. Single datazone

Consider the following data protection solutions for a single datazone:

- ◆ LAN-based — Distributed application servers require LAN backups.
- ◆ LAN-free — High-capacity data application servers require local backups that would overwhelm LAN-based backups. Application servers have CPU and I/O resource capacity to handle data protection overhead.
- ◆ Serverless — Business critical application servers have little CPU and I/O capacity to handle data protection overhead that compromises business application availability.

B.3.1. LAN clients, single datazone

LAN-based data protection issues are as follows:

- ◆ Client data structure limitations
- ◆ High density file system
- ◆ Large amounts of data
- ◆ Network attached storage
- ◆ NetApp
- ◆ EMC Celerra®

B.3.2. LAN network, single datazone

LAN-based data protection issues are as follows:

- ◆ Network
- ◆ High-speed switched network is required to maximize throughput
- ◆ 100 base -T
- ◆ Gigabit
- ◆ Backup server or storage node
- ◆ NICs limited by server backplane capacity and CPU resources or network bandwidth
- ◆ Devices limited by CPU and I/O resources (backplane bandwidth and number of buses)

B.3.3. Client optimization, single datazone

A performance improvement strategy is as follows:

- ◆ Use application modules to improve data protection efficiency
- ◆ Oracle RMAN support
- ◆ SAP
- ◆ Informix
- ◆ Microsoft SQL Server
- ◆ Microsoft Exchange Server
- ◆ Use NDMP-based backup for NAS clients
- ◆ NetWorker SnapImage tape server

B.3.4. Host optimization, single datazone

A performance improvement strategy (non-NDMP) is as follows:

- ◆ Optimize the server or storage node configuration
- ◆ Immediate save versus TCP/IP (CPU is 400 MHz Sparc, tape is DLT7K)
- ◆ $n + 1$ CPUs ($n = (\text{number of tape devices})/2$)
- ◆ 50 percent CPU utilization (immediate save/recover, one stream)
- ◆ Add CPUs for NICs and HBAs as needed

B.3.4.1. Considerations

- ◆ 256 MB per CPU (scale up for recovery operations, if needed)
- ◆ (Memory bandwidth) * (.70) * (.25) = maximum backplane I/O
- ◆ Two tape devices per CPU
- ◆ One tape per SCSI bus
- ◆ RAID 1 + 0 (or RAID 10) index and media database disk for maximum scalability
- ◆ Local host tables before DNS name resolution
- ◆ Kernel and resource requirements (file descriptors, memory heap)
- ◆ Treat the CPU speeds and tape speeds as base lines and extrapolate out for your environment.
- ◆ For memory bandwidth, multiply the speed of the memory times the width.
- ◆ The reference SCSI bus here was 20 MB/s. In other words, 1XDLT7K per 20 MB/s bus. To extrapolate this, identify what the maximum compressed throughput is for a tape drive and use that as the capacity requirement. For example, the DLT7K internal bus is about 18 MB/s, which is the maximum throughput you get with highly compressible data and hardware compression enabled, thus the recommendation for a minimum dedicated 20 MB/s bus.
- ◆ The host files recommendation is a troubleshooting step, not necessarily (unless DNS performance issues cannot be resolved, which is frequently the case) the production deployment requirement.

B.3.5. NDMP optimization, single datazone

Note: These considerations are for NetWorker pre-7.2. For NetWorker 7.2, the memory requirement is less by at least a factor of 10.

A performance improvement strategy (NDMP) is as follows:

- ◆ Optimize the server or storage node configuration
- ◆ n +1 CPUs (n = CPU for NDMP indexing)
- ◆ RAID 1 + 0 (or RAID 10) for the index and media database
- ◆ Local host tables before DNS name resolution

- ◆ 256 MB per CPU
- ◆ NDMP Backup — 150 MB memory/wwap per 1 million files
- ◆ NDMP Recover — 300 MB memory/wwap per 1 million files
- ◆ DAR for efficient tape recovery operations

B.3.6. Server configuration, single datazone

A performance improvement strategy is as follows:

- ◆ Use multiplexing (non-NDMP).
- ◆ Maximize the data protection device usage for disk or network source limited clients.
- ◆ Balance the total data. Stagger the backup schedule so that a nightly backup reduces the total amount of data per backup window.
- ◆ Optimize **savegrp** usage. Minimize the number of groups and group overlap while keeping the number of clients per group consistent with the capacity of the datazone. (server parallelism and savegrp parallelism)
- ◆ Use storage nodes to improve NetWorker server capacity and data protection performance.
- ◆ Increase the number of datazone devices by using local application server data protection devices.

B.3.7. Client options, single datazone

A performance improvement strategy is as follows:

- ◆ NAS
- ◆ NetApp
- ◆ Celerra
- ◆ NetWorker SnapImage
- ◆ Snapshot — Local block image backup
- ◆ Three party copy
- ◆ Application modules
 - Oracle
 - SAP
 - IBM Informix, DB2

- Microsoft SQL Server, Microsoft Exchange Server
- NMO
- NMES

B.3.8. Storage nodes, single datazone

Consider the following LAN-free data protection implementation issues related to storage nodes.

When are storage nodes used?

- ◆ When protecting large amounts of application server data (immediate save and recover).
- ◆ When application servers are isolated by slow or multiple protocol network configurations.
- ◆ When poor LAN performance requires local tape support.
- ◆ When LAN backup CPU load must be reduced.

B.3.9. Storage node design, single datazone

B.3.9.1. Single or multiple datazones

Consider the capacity of the datazone design:

- ◆ Are there datazone split/consolidation questions?
- ◆ Does the number of clients exceed the NetWorker server capacity (*nsr.res* and *nsrd* performance issues)?
- ◆ Does the number of LAN storage nodes exceed the server capacity (number of active sessions or remote devices)?
- ◆ Is the NetWorker server index and media database within the design capacity?
- ◆ Allow for growth over the life of the design.

B.3.9.2. NDMP, single datazone

Consider the following items in a single datazone with NDMP:

- ◆ When vendor application specific
- ◆ NAS
- ◆ NetApp — NW 6.0.1-001 supported
- ◆ Celerra — NW 6.0.1-001 supported
- ◆ NDMP devices required

- ◆ NetWorker SnapImage
- ◆ Local block image backups for NDMP and non-NDMP based application servers

B.3.10. Serverless design, single datazone

Consider the following in a serverless backup implementation:

- ◆ Secondary host or data mover workstation
- ◆ Celestra — Data mover workstation or CESD
- ◆ EMC TimeFinder Modules (SAP and Oracle) — Secondary host
- ◆ More coming with NetWorker 7.x Advanced file type device
- ◆ Block-level backup performance
- ◆ Can be connected via FC or SCSI attached devices

B.3.11. Single datazone review

- ◆ Use a LAN-based, LAN-free, and serverless solution where each fits the requirements.
- ◆ A one-size solution does not solve all problems.
- ◆ If the NetWorker server is already at capacity for LAN-based backups, do not add storage nodes or consolidate datazones. Use multiple datazones.
- ◆ A blend of each type of solution is usually required to provide the best solution.
- ◆ Using this approach will achieve the big picture.

B.4. Network best practices

- ◆ All NICs and switches should be hard-set to 100 Mb/s Full-duplex where appropriate.
- ◆ Local host tables for NetWorker servers and storage nodes (minimum)
- ◆ Connect lower performance NICs (100 Base-T) to a switch and high-performance NICs (Gb) between a switch and the NetWorker server/storage node (minimum).

- ◆ Configure Windows servers larger than 70-100 GB and UNIX servers larger than 100-150 GB as storage nodes on a 100 Mb/s LAN or upgraded to Gb Ethernet.

B.4.1. Consideration

The issue of autonegotiation and autosensing is a controversial one. For more information, go to the Network World website, www.networkworld.com, and search on those keywords.

B.5. SAN best practices

- ◆ Always check for compatibility between SAN components
- ◆ Multi-vendor can be acceptable but dependencies become complex quickly
- ◆ With less than the current JNI/EMULEX/QLogic, any two of the three work fine together, but all three may need to be zoned to isolate QLogic. The latest firmware/drivers may not have this problem.
- ◆ Always observe the SAN power up sequence.
- ◆ Always run a pilot before deploying.
- ◆ Separate the SAN for disk and tape.
- ◆ Set the data router (Crossroads, ADIC, and so on) to block SCSI resets.
- ◆ Use persistent binding for tape devices.

B.5.1. Consideration: Recommended power up sequence

- ◆ SCSI devices — Tape and disk devices that reside behind a FC-to-SCSI bridge/router
- ◆ FC switch and/or hub devices
- ◆ FC-to-SCSI bridge or routing devices
- ◆ FC RAID or other disk devices
- ◆ Host systems

B.6. Business best practices

- ◆ Test business continuity
- ◆ Validate backups
- ◆ Cross-check compatibility
- ◆ Use the NetWorker disaster recovery guide
- ◆ Use the NetWorker performance tuning guide
- ◆ Engage Professional Services
- ◆ Train personnel at all sites

B.6.1. Considerations

- ◆ Test once a year and understand your business continuity plan.
- ◆ Validate backups by testing recovers.
- ◆ Cross-check hardware and software with the current Compatibility Guides on the EMC website.
- ◆ Use the NetWorker disaster recovery guide and performance tuning guide practices where they apply.
- ◆ Engage Professional Services where appropriate.
- ◆ Train personnel at all sites in business continuance processes.

B.7. NetWorker best practices

Follow these best practices in addition to configuring and running the NetWorker software:

- ◆ Use automatic cloning for light cloning requirements only.
- ◆ Keep a printed copy of the daily bootstrap.
- ◆ Keep an offsite bootstrap and bootstrap printout.
- ◆ Schedule indexing outside the backup window (**savegrp -O -G group_name**).
- ◆ Keep the configuration simple.

B.7.1. Considerations

- ◆ Do not use automatic cloning if your groups overlap (drive and tape contention, etc.).
- ◆ Keep a printed copy of bootstrap daily.

- ◆ Send a copy of the bootstrap tape and printout offsite for disaster recovery along with your offsite data shipment.
- ◆ Save the bootstrap and index to the same pool and make a copy (clone) to simplify DR.
- ◆ You can turn indexing off to increase backup performance (eliminates indexing overhead on backup server).
- ◆ Create a superset Group (for example, NeverRun) that you never run, to consolidate the index backups (`savegrp -O -G group_name`). In high utilization environments, this needs to be run separately (just as with cloning).
- ◆ Keep your configuration simple for ease of management and less administrative overhead.

B.8. NetWorker 7.x best practices

- ◆ Advanced file type devices should be dedicated and partitioned or zoned to the desired size.
- ◆ Advanced file type device auto staging should be enabled.
- ◆ Turn on the resource logging and ensure that users login by using their own ID, *not* the root ID.
- ◆ The number of save sets on a volume directly impacts bulk volume operations.
 - Purge (save set operations performed in series).
 - Label volumes.
 - Media database migration.
- ◆ Follow the NetWorker 7.x disaster recovery guide for the new `mmrecov` operation.
- ◆ All NDMP devices must be identified as remote devices

B.8.1. File type device configuration on NetApp FAS960 (December 16, 2003)

The best practices are simple:

- ◆ Using one file device per LUN or file system makes it more easy to calculate throughput.
- ◆ Do not mix file devices in one file system, to prevent one from filling up the file system.

- ◆ Use as many file devices as concurrent recovery sessions. For example, if you need to recover three servers at the same time, use at least three file devices and back up each server to a separate file device.
- ◆ Use a journaling file system. If a crash did not occur, a reboot can take a lot of time with its checking.
- ◆ One pool per file device is a general rule.
- ◆ If you perform cloning or staging, consider the read throughput of each file device. It should be similar to the tape drive throughput.
- ◆ Use a large file system (1 to 3 TB) for a file device, especially for an advanced file device that should never become full.
- ◆ An advanced file device requires two **nsrmmmd** processes and two devices for NetWorker licencing and sizing.
- ◆ Work closely with NetApp when configuring the array.