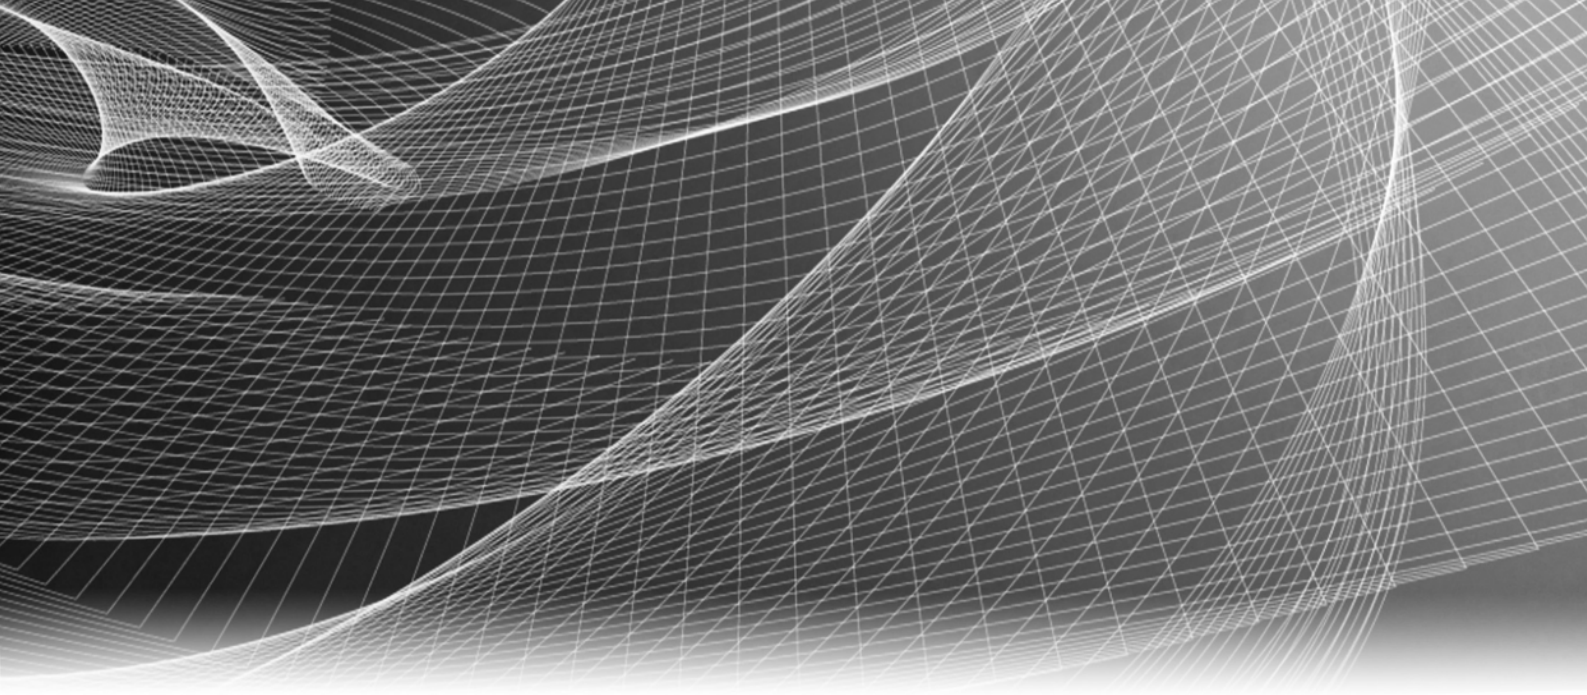

EMC® NetWorker®
Module for Databases and Applications

Release 1.6

Administration Guide

P/N 300-015-313
REV 02

EMC²

Published January 2014

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 regulatory document for your product line, go to EMC Online Support (https://support.emc.com).

CONTENTS

Preface

Revision History

Chapter 1 Overview of NMDA Features

Road map for NMDA operations 24
Terminology used in this guide 25
Importance of backups 26
NMDA features for all supported applications 26
Basic backup and recovery features 27
Scheduled backups versus manual backups 28
Client Direct feature 29
Deduplication backups and restores with EMC Avamar 29
Deduplication backups and restores with EMC Data Domain 30
Snapshot backups and restores 31
Virtualization support 32
Cluster backups and restores 33
Configuration wizards 33
Configuration conversion 34
Internationalization (I18N) 34
IPv6 support for backups and restores 34
Multiple database or application installations on the same host 34
Simplified scheduled backup configuration of multiple databases or applications 35
NMDA features specific to DB2 35
DB2 full, incremental, and delta backups 35
DB2 transaction log backups 35
DB2 history pruning 36
DB2 10.5 features 36
Restores of Fujitsu NetWorker Module for DB2 backups 37
NMDA features specific to Informix 37
Informix terminology 37
Informix full and incremental backups 37
Logical log backups 37
Restores of Fujitsu NetWorker Module for Informix backups 38
Informix 12.10 features 38
NMDA features specific to Lotus 38
Files backed up during Lotus backups 38
Lotus full and incremental backups 40
Lotus transaction log backups 40
Lotus DAOS backups and restores 40
Restartable scheduled backups 41
Types of Lotus restores 42
NetWorker User for Lotus program 42
Partitioned Domino servers 42
NMDA features specific to MySQL 43
MySQL online and offline backups 43
MySQL full and incremental backups 43

MySQL binary log backups 43
MySQL backup granularity 44
MySQL restore and recovery operations 45
MySQL validate operation 46
MySQL 5.6 features 46
NMDA features specific to Oracle 46
Oracle full and incremental backups 47
Backups of archived redo logs 47
Oracle DBA and NetWorker backup administrator collaboration 48
Control file and server parameter file autobackups 48
Automatic channel allocation 48
Backup and restore optimization 49
Backup copies 49
Backup and deletion of backup sets 49
Oracle Data Guard support 50
Oracle Exadata support 50
Restartable backups 51
Retention policies 51
Restores of Fujitsu NetWorker Module for Oracle backups 52
Save set bundling 52
Policy uniformity 55
Other Oracle features 56
NMDA features specific to Sybase 59
Sybase full and incremental backups 59
Sybase up-to-the-minute recovery 60
Password-protected database backups and restores 61
Database consistency check 61
Database backup and restore verification 62
Exclusion of multiple user-defined temporary databases from backup 62
Sybase in-memory and relaxed durability databases 62
Sybase ASE 15.7 features 62
NetWorker User for Sybase 63
NMDA components 63
NMDA backup and restore processes 65
Regular scheduled backup processes 65
Regular restore processes 67

Chapter 2 Backup Configuration

Configuring NMDA backups 70
Configuring internationalization (I18N) support 71
I18N support requirements 73
Configuring I18N support for Informix operations 73
Configuring I18N support for Oracle operations 73
Configuring I18N support for Sybase operations 74
Verifying the database or application server configuration 75
Verifying the basic NetWorker resource configurations 75
NetWorker Server resource 75
NetWorker User Group resource 75
NetWorker Client resource 78
NetWorker Device resource 78
NetWorker Pool and Label Template resources 78
Firewall support 79
Converting client-side to server-side configurations 79
NetWorker resources modified by the conversion 80

Using the nsrdaadmin conversion command 81
Configuring scheduled backups with the wizard 82
Requirements for using the backup configuration wizard 83
Configuring a scheduled backup with the wizard 83
Configuring scheduled backups without the wizard 85
Configuring the NMDA parameters without the wizard 85
Configuring the NetWorker resources with NMC 86
Configuring the Group resource 87
Configuring the Schedule resource 87
Configuring the Client resource 89
Configuring manual backups 94
Configuring deduplication backups with Avamar 95
Requirements for Avamar deduplication backups 95
Best practices for Avamar deduplication backups 95
Configuring scheduled Avamar deduplication backups 98
Configuring manual Avamar deduplication backups 99
Configuring Data Domain backups 99
Best practices for Data Domain backups to any device 100
Configuring DD Boost backups to use Client Direct 101
Configuring DD Boost backups to use a storage node 103
Configuring transaction log backups to traditional NetWorker devices. 104
Configuring backups to AFTD storage 104
Configuring Client Direct backups to AFTD 104
Configuring storage node backups to AFTD 105
Configuring probe-based backups 105
NMDA probe 106
User-defined probe 107
Configuring a probe-based backup 107
Configuring parallel backups 112
Configuring DB2 parallel backups 114
Configuring Informix parallel backups 114
Configuring Lotus parallel backups 115
Configuring Oracle parallel backups 115
Configuring Sybase parallel (multistripe) backups 115
DB2 backup considerations 116
Configuring automatic deletion of DB2 recovery objects 116
Configuring automatic backups of DB2 transaction logs 116
Informix backup considerations 117
Configuring automatic (continuous) backups of Informix logical logs... 117
Lotus backup considerations 118
Setting the environment for Lotus manual backups on Solaris or Linux 118
Considerations for Lotus database or directory link backups 118
Considerations for Lotus transaction log backups 118
Lotus incremental backups with the comfort span option 120
Configuring partitioned Domino server backups 120
Configuring Lotus DAOS backups 121
Configuring Lotus restartable scheduled backups 124
MySQL backup considerations 125
Configuration parameter settings 125
Considerations for MySQL whole and partial backups 126
Considerations for MySQL incremental backups 126
Considerations for MySQL binary log backups 127
MySQL backup privileges 127
Configuring MySQL 5.6 features 128
Oracle backup considerations 128

RMAN scripts for manual backups 129
RMAN scripts for scheduled backups 130
Configuring automatic channel allocation 132
Creating Oracle backup copies 133
Configuring save set bundling for scheduled Oracle backups 135
Configuring policy uniformity for scheduled Oracle backups 136
Configuring operations in an Oracle Data Guard environment 137
Sybase backup considerations 137
Setting up Sybase roles and permissions 137
Specifying verification of Sybase database backups 138
Specifying volume pools for Sybase incremental backups 138
Sybase transaction log backups 139
Configuring a Sybase threshold procedure 141
Setting the environment for Sybase backups on HP-UX 143
Configuring Sybase ASE 15.7 backup features 143

Chapter 3 Backup Procedures

Performing scheduled backups 146
Testing scheduled backups 146
Canceling scheduled backups 146
Restarting failed scheduled backups 147
Monitoring scheduled backups 147
Performing manual backups 148
Performing DB2 manual backups with the db2 backup command 148
Performing DB2 manual backups with the DB2 GUI 149
Performing Informix manual backups with the onbar command 150
Performing Lotus manual backups with the nsrdasv command 150
Performing Lotus manual backups with NetWorker User for Lotus 151
Performing MySQL manual backups with the nsrdasv command 153
Performing Oracle manual backups with the rman command 153
Performing Oracle manual backups with Oracle Enterprise Manager 154
Performing Sybase database consistency checks before backups 154
Performing Sybase manual backups with the nsrdasv command 155
Performing Sybase manual backups with NetWorker User for Sybase 156
Canceling manual backups 158
Monitoring manual backups 158
Verifying backup information in NetWorker indexes 159
Verifying Avamar backup information in NetWorker indexes 159
Verifying Data Domain backup information in NetWorker indexes 161
Lotus restartable backup information in NetWorker indexes 161
MySQL backup information in NetWorker indexes 162
Synchronizing backup catalogs and deleting backups 164
Deleting DB2 backups 164
Deleting Informix backups 166
Cross-checking and deleting Oracle backups 166

Chapter 4 Data Restore and Recovery

Data restore and recovery terminology 168
NetWorker indexes and policies used for restores 168
Performing NMDA data restore and recovery 169
Performing DB2 data restore and recovery 171
Determining how many restore devices and sessions to use 172
Performing DB2 data restores with the db2 restore command 173

Performing DB2 data recovery 176
Performing DB2 data restore and recovery with the DB2 GUI..... 178
Performing Informix data restore and recovery..... 178
Determining the Informix restore mode 179
Performing Informix data restores with the onbar command 180
Performing Lotus data restore and recovery..... 181
Performing Lotus database recovery with the nsrnotesrc command..... 182
Performing Lotus database recovery with NetWorker User for Lotus 185
Performing recovery of partitioned Domino servers..... 190
Performing Lotus DAOS data recovery..... 191
Performing Lotus document-level recovery with nsrdocrc command 192
Performing Lotus document-level recovery with the Notes client GUI ... 194
Performing MySQL data restore and recovery..... 196
Determining the binary logs for MySQL recovery 197
Performing MySQL recovery of whole instance backups..... 198
Performing MySQL restores of partial backups 199
Performing MySQL restores of binary log backups 201
Performing MySQL redirected restores..... 201
Performing MySQL list image operations 202
Performing MySQL extract operations..... 203
Performing MySQL extract and prepare operations 204
Performing MySQL copy back operations..... 205
Performing MySQL restores of InnoDB tables outside data directory.... 206
Performing MySQL validate operations..... 206
Performing Oracle data restore and recovery 207
Determining the volumes for Oracle restores 207
Preventing restore performance degradation with Oracle 10.2 or later. 211
RMAN scripts for restore and recovery 211
Using the Oracle recovery configuration wizard 212
Performing Oracle restores with the rman command..... 215
Performing Oracle restores with the Oracle Enterprise Manager..... 215
Performing Oracle data recovery..... 215
Performing Sybase data restore and recovery 216
Specifying verification of Sybase database restores 216
Performing Sybase data restores with the nsrsybrc command 216
Performing Sybase data restores with NetWorker User for Sybase 220
Canceling a data restore or recovery..... 223

Chapter 5 Disaster Recovery

Preparing for disaster recovery..... 226
Preparing for DB2 disaster recovery..... 227
Preparing for Informix disaster recovery..... 228
Performing an Informix imported restore 229
Recovering from an Informix server disk crash..... 231
Preparing for Lotus disaster recovery..... 231
Recovering a nonlogged Lotus environment 232
Recovering a logged Lotus environment 233
Preparing for MySQL disaster recovery..... 235
Performing a MySQL disaster recovery..... 236
Preparing for Oracle disaster recovery 236
Creating the DBID text file 238
Setting up a postcommand script for backup of Oracle files 238
Setting up RMAN backups of the database and related files..... 239
Setting up RMAN backups of Recovery Catalog..... 240

Performing an Oracle disaster recovery 240
Preparing for Sybase disaster recovery 240
Recovering the master database 241
Recovering user databases after database device failure..... 242

Chapter 6 Cluster Systems

Active-passive clusters..... 246
Backup failover 247
Cluster configuration requirements 248
Converting client-side to server-side configurations in a cluster..... 249
Configuring scheduled backups in a cluster with the wizard 249
Configuring scheduled backups in a cluster without the wizard 250
Configuring manual backups in a cluster..... 252
Configuring Avamar deduplication backups in a cluster..... 252
Configuring Data Domain deduplication backups in a cluster 253
Configuring probe-based backups in a cluster..... 253
Setting up nodes to back up to a local storage node..... 253
Configuring recoveries in a cluster..... 254
Active-active application clusters (overview) 254
Active-active cluster configuration requirements 255
Restarting backups after a system failure 256
DB2 DPF systems 256
Configuring backups and restores in a DPF environment 257
Configuring manual DPF backups 258
Performing manual DPF backups 258
Configuring scheduled DPF backups..... 259
Restoring and recovering DPF backups 260
DB2 pureScale systems..... 260
Configuring backups and restores in a DB2 pureScale system 261
Configuring manual backups in a DB2 pureScale system..... 262
Performing manual backups in a DB2 pureScale system..... 262
Configuring scheduled backups in a DB2 pureScale system 262
Configuring automatic backups of transaction logs in a DB2 pureScale system 263
Restoring and recovering backups in a DB2 pureScale system..... 263
Informix MACH systems..... 264
Using Continuous Log Restore for backup and restore 265
Backups of primary and HDR or RSS or SDS servers 266
Restores of primary and HDR or RSS or SDS servers 266
Oracle RAC systems 267
Configuring backups and restores in a RAC environment 267
Connect-time failover 268
Parallel RAC backups and restores 269
Creating RMAN backup and restore scripts 269
Archived redo logs 271
Sybase ASE Cluster Edition systems 272
Configuring backups and restores for ASE Cluster Edition..... 273
Skipping the backup of temporary databases..... 274
Restoring an ASE Cluster Edition backup 274
Performing backup and recovery of the quorum device..... 274
Recovering the master database on clusters..... 275

Chapter 7 Multiple Installations on a Single Host

Multiple database or application installations 278
Multiple databases or applications with the same bitness 278
Multiple databases or applications with 32-bit and 64-bit coexistence 278

Chapter 8 Oracle DBA and NetWorker Backup Administrator Collaboration

DBA disk backup issues in Oracle environments 282
NMDA solution for DBA disk backup issues 282
Configuring backups of DBA disk backups 284
Advanced NMDA parameters 286
Configuring optional probe-based backups of DBA disk backups 287
Reporting on backups of DBA disk backups 288
Purging DBA disk backups 290
Restore and recovery of DBA disk backups 291

Chapter 9 Snapshot Backups and Restores

Snapshot operations with NetWorker Snapshot Management 294
Software requirements for NSM snapshot operations 294
Types of NSM backups 295
Types of NSM restores 296
NSM backup processes 297
Configuring NSM snapshot backups 299
Configuring Snapshot Pool resources for NSM backups 300
Configuring internationalization (I18N) support 300
Configuring Snapshot Policy resources for NSM backups 301
Configuring Group resources for NSM backups 301
Configuring Client resources for NSM backups 301
Configuring deduplication backups with NSM 302
Testing an NSM snapshot backup 302
Canceling an NSM snapshot backup 303
Configuring NSM snapshot restore and recovery 303
Setting the NSM parameters 304
DB2 considerations for NSM snapshot operations 305
Restoring a DB2 NSM backup 306
Managing and deleting DB2 NSM backups 306
Oracle considerations for NSM snapshot operations 308
Configuring the required Oracle settings 308
Configuring the NWORA resource file 310
Creating RMAN scripts for NSM snapshot backups 310
Checking configuration consistency 313
Performing Oracle NSM backups 315
Verifying Oracle NSM backup information in NetWorker indexes 317
Restoring an Oracle NSM backup 319
Catalog synchronization for Oracle NSM backups 322
NSM backups and restores on cluster systems 335
Replication Manager snapshot operations with Oracle ASM 338
NMDA integration with Replication Manager 339
Requirements for backups of Oracle database components 340
Requirements for snapshot backups of Oracle ASM data 340
Configuring snapshot backups of Oracle ASM data 341
Performing backups and restores of Oracle ASM data 348
Sample preprocessing and postprocessing scripts 349

Appendix A NMDA Parameters and Configuration File

Setting the NMDA parameters 354
NMDA configuration file 354
Syntax rules for the NMDA configuration file 355
Common NMDA parameters 358
NMDA DB2 parameters 366
NMDA Informix parameters 370
NMDA Lotus parameters 371
Wildcard support for Lotus operations 378
Setting the NSR_RECOVER_TIME parameter 379
NMDA MySQL parameters 380
NMDA Oracle parameters 387
NMDA Sybase parameters 392

Appendix B Oracle RMAN Commands

The pool option of the backup command 398
The send command 398
Syntax rules 398
Two ways to run the send command 400
Precedence rules 402

Appendix C Troubleshooting and Error Messages

General troubleshooting tips 406
Debug log files 407
Wizard backup configuration fails, authentication denied 408
Backup becomes suspended 408
NMDA error messages 408
NetWorker XBSA error messages 409
DB2 troubleshooting tips and error messages 412
Delta or incremental backup fails in a DB2 pureScale system 412
Removal of failed backup 412
The load libnsrdb2 command 413
DB2 error messages 413
Informix troubleshooting tips and error messages 415
No dbspaces/blobspaces found to back up or restore 415
Unable to open connection to server 416
Default value assigned to LTAPEDEV causes failure 416
ON-Bar returned codes 2 through 34 416
Lotus troubleshooting tips and error messages 417
"Invalid Time Specified" error 417
Backup failure due to time conversion 417
Notes API 525 error 417
NMC Group Details window displays “command not found” 418
Restore problem due to NSR_BACKUP_PATHS=NOTES: 418
Failure of a document-level recovery on a remote Domino server 418
NMDA Lotus error messages 419
MySQL troubleshooting tips and error messages 421
NMDA MySQL backup issues 421
NMDA MySQL restore issues 423
NMDA MySQL error messages 423
Oracle troubleshooting tips and error messages 425
Diagnosing NMDA Oracle problems 425
RMAN error messages 426

The trace option of the backup command 426
Canceling unresponsive Oracle backups on UNIX 427
Canceling unresponsive Oracle backups on Windows 427
NMDA Oracle error messages 428
Sybase troubleshooting tips and error messages 440
NMDA Sybase error messages 440

Glossary

Index

FIGURES

	Title	Page
1	Regular scheduled NMDA backup	66
2	Backup window in NetWorker User for Lotus	151
3	Backup Options dialog box in NetWorker User for Lotus	152
4	Backup Options dialog box in NetWorker User for Sybase	157
5	Lotus restartable backup information in NMC Console	162
6	Recover window in NetWorker User for Lotus	186
7	Recover Options dialog box in NetWorker User for Lotus	187
8	Recover to Sybase Server dialog box in NetWorker User for Sybase	222
9	Single database node with failover capability	246
10	Multiple database nodes with failover capability	247
11	Multiple database nodes with mutual failover capability	247
12	Single DPF host with shared memory	256
13	Two DPF hosts with multiple nodes	257
14	Software used in the NMDA environment for NSM snapshot operations	295
15	NSM snapshot backup and restore data flow	298
16	Application Credentials window in the Replication Manager console	342
17	Advanced Replication Settings in the Replication Manager console	343
18	Mount options in the Replication Manager console	344
19	Starting the Job window in the job wizard	345

TABLES

	Title	Page
1	NMDA components	63
2	NMDA parameters that support non-ASCII values	72
3	NetWorker Server resource attributes	75
4	User group privileges for NMDA operations	76
5	Database-specific user for each database or application	77
6	Options of the **nsrdaadmin** command for configuration conversion	81
7	Mandatory parameters for scheduled NMDA backups	85
8	Backup levels specified in NetWorker Schedule resource	87
9	NetWorker Client resource attributes	89
10	Mandatory parameters for manual NMDA backups	94
11	Configuring a Data Domain backup	102
12	NetWorker Probe resource attributes	107
13	Configuration parameters for InnoDB partial backups	126
14	Sybase roles and permissions	138
15	Full backups with NSR_DUMP_LOG_OPT	139
16	Incremental backups with NSR_DUMP_LOG_OPT	140
17	Threshold procedure location	142
18	Manual backup procedures	148
19	Database consistency check options of the **nsrsybcc** command	155
20	Data restore procedures	171
21	Configuration file parameters corresponding to Recover Options fields	188
22	Commands or library names for 32-bit and 64-bit coexistence	279
23	Advanced NMDA parameters for backups of DBA disk backups	286
24	NetWorker Probe resource attributes	287
25	Typical configuration for NSM snapshot operations with NMDA	294
26	Supported types of NSM backups	295
27	Supported types of NSM restores	296
28	Sample NetWorker Group resource attributes for a snapshot backup	301
29	NSM parameters	304
30	DB2 resource file parameters	307
31	NWORA parameter resources	325
32	NWORA SID resource components	327
33	Common NMDA parameters	358
34	NMDA DB2 parameters	366
35	NMDA Informix parameters	370
36	NMDA Lotus parameters	371
37	NMDA MySQL parameters	380
38	NMDA Oracle parameters	388
39	NMDA Sybase parameters	392
40	Option values in the **send** command	400
41	NetWorker XBSA error messages	409
42	Path and suffix for the load libnsrdb2 command	413
43	DB2 SQL error messages	413
44	Error messages from Lotus backups and restores	419
45	Error messages from MySQL backups and restores	424
46	Trace option values and conditions traced	427
47	Error messages from Oracle backups and restores	428
48	Error messages from the nsroraadmin program	437
49	Error messages from the nsrorainfo program	439
50	Error messages from Sybase consistency checks, backups, and restores	441

51 Sybase backup server and libnsrsyb error messages 444

PREFACE

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

Contact your EMC technical support professional if a product does not function properly or does not function as described in this document.

Note: This document was accurate at publication time. Go to EMC Online Support (https://support.emc.com) to ensure that you are using the latest version of this document.

Purpose

This document describes how to configure and use the EMC NetWorker Module for Databases and Applications (NMDA) release 1.6.

Audience

This document is intended for system administrators or database administrators (DBAs) who are responsible for installing software and maintaining backup and recovery systems for databases or applications.

Users of this guide must be familiar with these topics:

- Backup, recovery, database, applications, and network terminology
- Backup and recovery procedures
- Disaster recovery procedures

Related documentation

You can find additional EMC publications for this product release and related NetWorker products at EMC Online Support.

The *EMC NetWorker Software Compatibility Guide* at EMC Online Support provides a full list of supported environments and platforms.

The following additional documentation might be useful:

- Database or application server documentation
- Database or application backup and recovery documentation

Conventions used in this document

EMC uses the following conventions for special notices:

NOTICE

NOTICE is used to address practices not related to personal injury.

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

Typographical conventions

EMC uses the following type style conventions in this document:

Normal	Used in running (nonprocedural) text for: • Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus • Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities • URLs, pathnames, file names, directory names, computer names, links, groups, service keys, file systems, and notifications
Bold	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and man pages Used in procedures for: • Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus • What the user specifically selects, clicks, presses, or types
Italic	Used in all text (including procedures) for: • Full titles of publications referenced in text • Emphasis, for example, a new term • Variables
`Courier`	Used for: • System output, such as an error message or script • URLs, complete paths, file names, prompts, and syntax when shown outside of running text
`Courier bold`	Used for specific user input, such as commands
`Courier italic`	Used in procedures for: • Variables on the command line • User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user
[ ]	Square brackets enclose optional values
|	Vertical bar indicates alternate selections—the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

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

Product information — For documentation, release notes, software updates, or information about EMC products, go to EMC Online Support at:

https://support.emc.com

Technical support — Go to EMC Online Support and click Service Center. You will see several options for contacting EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

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

Your comments

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

BRSDocumentation@emc.com

REVISION HISTORY

The following table presents the revision history of this document.

Revision	Date	Description
02	January 6, 2014	Updated the following information: • Table 1 on page 63 - Deleted two obsolete references in the NMDA components table. • "Simplified scheduled backup configuration of multiple databases or applications" on page 35 - Deleted the references to "instances" in the topic. • Change libnsrdb2.*xx* to libnsrdb2.so in the following information: - Table 1 on page 63 - "Configuring automatic backups of DB2 transaction logs" on page 116 - "Performing DB2 manual backups with the db2 backup command" on page 148 • "NetWorker User Group resource" on page 75 - Added details about ensuring that the user has the required privileges for NMDA operations. • "Configuring MySQL 5.6 features" on page 128 - Deleted information about setting **innodb_undo_directory** or **innodb_page_size** in the MYSQL_MEB_OPTIONS parameter. • "Configuring automatic backups of transaction logs in a DB2 pureScale system" on page 263 - Replaced information with a cross-reference to a previous topic. • "Restoring and recovering backups in a DB2 pureScale system" on page 263 - Updated the **nsrdb2rlog** examples and added information about the **-c** *client_name* option. • Figure 14 on page 295 - Modified the NetWorker storage node information in the figure. • Example 30 on page 305 - Corrected the NSR_SNAP_TYPE parameter value in the VMAX example. • Table 30 on page 307 - Updated the DB2PATH and NSR_DB2CAT_MODE parameter information. • "NSR_ORACLE_NLS_LANG" on page 325 - Updated the default value of the parameter. • "NSR_NWPATH" on page 364 - Added a recommendation for Oracle operations on Windows and deleted references to Sun branded software.
01	November 14, 2013	Initial release of NMDA 1.6.

CHAPTER 1
Overview of NMDA Features

This chapter includes the following topics:

- Road map for NMDA operations 24
- Terminology used in this guide 25
- Importance of backups 26
- NMDA features for all supported applications 26
- NMDA features specific to DB2 35
- NMDA features specific to Informix 37
- NMDA features specific to Lotus 38
- NMDA features specific to MySQL 43
- NMDA features specific to Oracle 46
- NMDA features specific to Sybase 59
- NMDA components 63
- NMDA backup and restore processes 65

Road map for NMDA operations

Follow this road map to configure and perform backup and restore operations with EMC® NetWorker® Module for Databases and Applications (NMDA) for the supported databases and applications.

Before you perform NMDA operations, review the *EMC NetWorker Module for Databases and Applications Release Notes* for the most up-to-date information about NMDA.

Ensure that you have installed or updated NMDA according to the instructions in the *EMC NetWorker Module for Databases and Applications Installation Guide*.

NOTICE

Before you click any cross-reference in this guide, determine if you will need to return to the section that contains the cross-reference. Ensure that you know how to return to the section, if required.

Complete the required steps to configure and perform the NMDA operations.

1. Review the following information about this guide and the major features of NMDA:

   - "Terminology used in this guide" on page 25
   - "Importance of backups" on page 26
   - "NMDA features for all supported applications" on page 26
   - Information about the NMDA features for a specific database or application:
     - "NMDA features specific to DB2" on page 35
     - "NMDA features specific to Informix" on page 37
     - "NMDA features specific to Lotus" on page 38
     - "NMDA features specific to MySQL" on page 43
     - "NMDA features specific to Oracle" on page 46
     - "NMDA features specific to Sybase" on page 59

   If you have multiple database installations on the NMDA host, review Chapter 7, "Multiple Installations on a Single Host."

   The following topics provide additional general information about NMDA components and the backup and restore workflows:

   - "NMDA components" on page 63
   - "NMDA backup and restore processes" on page 65

NOTICE

In an Oracle environment where the Oracle DBA and NetWorker backup administrator collaborate to enable simplified backups and restores of Oracle disk backups, ignore step 2 to step 6 in this road map. Instead, follow the NMDA configuration, backup, and restore instructions in Chapter 8, "Oracle DBA and NetWorker Backup Administrator Collaboration."

2. Configure the NMDA backups:
   - In a stand-alone environment, configure regular (not snapshot) backups according to "Configuring NMDA backups" on page 70.
   - In a cluster environment, configure regular backups according to the information in Chapter 6, "Cluster Systems."
   - Configure DB2 or Oracle snapshot backups according to the information in Chapter 9, "Snapshot Backups and Restores."

   NOTICE

   To enable NMDA operations on UNIX systems, ensure that the /nsr/apps and /nsr/apps/tmp directories have the drwxrwxrwt access permissions.

3. Perform the NMDA backups:
   - In a stand-alone environment, perform regular backups according to "Performing scheduled backups" on page 146 or "Performing manual backups" on page 148.
   - In a cluster environment, perform regular backups according to the information in Chapter 6, "Cluster Systems."
   - Perform DB2 or Oracle snapshot backups according to the information in Chapter 9, "Snapshot Backups and Restores."

   You can verify information about NMDA backups according to "Verifying backup information in NetWorker indexes" on page 159.

   You can synchronize information in the application backup catalog and NetWorker indexes according to "Synchronizing backup catalogs and deleting backups" on page 164.

4. Use the following information to perform the restore and recovery of NMDA backups:
   - In a stand-alone environment, restore regular backups according to "Performing NMDA data restore and recovery" on page 169.
   - In a cluster environment, restore regular backups according to the information in Chapter 6, "Cluster Systems."
   - Restore DB2 or Oracle snapshot backups according to the information in Chapter 9, "Snapshot Backups and Restores."
5. Use the information in Chapter 5, "Disaster Recovery," to perform disaster recovery procedures.
6. Use the following information to troubleshoot any issues with NMDA operations:
   - Appendix C, "Troubleshooting and Error Messages"
   - *EMC NetWorker Module for Databases and Applications Release Notes*

Terminology used in this guide

The generic sections of this guide use the term "transaction logs" for the logs required to recover data backed up by NMDA. The different applications that NMDA supports use application-specific terms for the logs, such as archived logs, binary logs, logical logs, and so on.

In this guide, the UNIX references apply to both UNIX and Linux operating systems, unless specified otherwise. The Windows references apply to all the supported Microsoft Windows operating systems, unless specified otherwise.

Unlike NetWorker processes that use the term recovery for all backup retrieval activities, NMDA processes distinguish between the restore and recovery of a database:

- Restore means to retrieve individual data from backup and store the data on disk.
- Recover means to apply the transaction logs to make the database consistent.

A legacy NetWorker Module is one of the NetWorker Modules that NMDA replaces:

- NetWorker Module for DB2 (NMDB2)
- NetWorker Module for Informix (NMI)
- NetWorker Module for Lotus (NML)
- NetWorker Module for Oracle (NMO)
- NetWorker Module for Sybase (NMS)

"Glossary" on page 447 provides details about the terms used in this guide.

Importance of backups

For the complete protection of a database system or application system, a viable backup strategy must include regular backups of the following data:

- Database or application data
- Transaction logs
- Configuration and control files for a database or application

These backups are important for the following reasons:

- Without backups, you cannot restore a database at all.
- Without transaction logs, you can restore a database only to the time of its last consistent backup. Transaction logs enable you to roll forward the database restore to the current time or a point-in-time.
- Without the configuration files and control files, you cannot start up the database and application even though the data is intact on the primary storage. These files are important, especially in disaster recovery scenarios.

NMDA features for all supported applications

NMDA provides a data protection solution for DB2, Informix, Lotus Domino/Notes, MySQL, Oracle, and Sybase data. NMDA operates with the NetWorker software and the supported database software or application software.

NMDA integrates database backups and file system backups to a centralized backup server, which relieves the burden of backup from the database administrator while enabling the database administrator to retain control of the restore process.

NMDA provides high performance backups and restores of databases and applications to and from different NetWorker backup devices, including Data Domain devices, Avamar devices, and tape devices.

The following sections provide details about each different feature that NMDA supports.

Basic backup and recovery features

NMDA supports the following basic backup and recovery features:

- Online backup—NMDA can back up an online database or online application without requiring any downtime.
- Full and incremental backup—NMDA can back up all the data or only the data that has changed. Application-specific topics in this guide describe the supported types of backups.
- Backup of transaction logs and other files required for recovery—Application-specific topics in this guide provide details about the backup of transaction logs and other files.
- Automatic recovery of a database or application to the current time or to an arbitrary point-in-time.
- Recovery to the original location or to an alternate location.
- Granular backup and granular recovery—NMDA supports the following backup and recovery granularities:
  - DB2 database and tablespace backup and recovery
  - Informix instance and dbspace backup and recovery
  - Lotus database backup and Lotus database and document-level recovery
  - MySQL instance, database, and table backup and recovery
  - Oracle database, tablespace, and datafile backup and recovery

    Note: In addition, Oracle software can perform block-level recovery.

  - Sybase ASE server and database backup and recovery
- Parallelism—NMDA supports concurrent data streams for a backup or restore as described in "Configuring parallel backups" on page 112.

Scheduled backups versus manual backups

An NMDA backup can be either a scheduled backup or a manual backup:

- The NetWorker server initiates a scheduled backup through the **savegrp** program interface. The scheduled backup start time depends on the settings in the NetWorker resources. The scheduled backup automatically backs up the NetWorker bootstrap and client index, which are required for NetWorker server disaster recovery:
  - A regular (time-based) scheduled backup starts at a time specified in the NetWorker Group resource.
  - A probe-based (event-based) scheduled backup starts when specified conditions are met, as determined by the probe program. NMDA supports two types of probes, an NMDA probe and user-defined probes:
    - An NMDA probe is implemented through the **nsrdaprobe** program, which determines the amount of logs generated since the previous backup. The probe-based backup runs when that amount is equal to or exceeds the user-configured log threshold.
    - A user-defined probe checks if any user-defined condition has been met since the previous backup.

    You must configure a probe-based backup according to "Configuring probe-based backups" on page 105.

    Note: For MySQL probe-based backups, NMDA supports only user-defined probes, not the NMDA **nsrdaprobe** program.

    A snapshot backup with EMC NetWorker Snapshot Management is supported only through a scheduled backup.

- A manual backup is initiated by a user on the application host through a database or application-specific command or procedure that runs the appropriate backup utility. The manual backup does not back up the NetWorker bootstrap and client index, which are required for NetWorker server disaster recovery. You can perform the bootstrap and client index backup separately. The NetWorker disaster recovery documentation provides details.

Chapter 2, "Backup Configuration," describes the configuration of both scheduled backups and manual backups.

Chapter 3, "Backup Procedures," describes the procedures for both scheduled backups and manual backups.

The NetWorker server maintains an online client file index and online media database, which together comprise the online indexes. During an NMDA backup, the NetWorker server makes an entry in the online client file index. The server also records the location of the data in the online media database. These entries provide recovery information required for all backed-up data.

During an NMDA scheduled backup, the NetWorker server backs up the client index and other server configuration information (called bootstrap), such as the resource database. The client index and bootstrap information is mandatory for performing disaster recovery of the NetWorker server.

If you do not run scheduled backups with NMDA, you must back up the bootstrap information manually to prepare for a NetWorker server disaster. The NetWorker disaster recovery documentation describes how to back up the bootstrap information and how to prepare for NetWorker server disaster recovery.

Client Direct feature

NMDA supports the Client Direct feature, which enables client backups to bypass the NetWorker storage node and perform one of the following operations:

- Send deduplicated backup data directly to DD Boost devices.
- Send backup data directly to advanced file type device (AFTD).

The Client Direct feature, also known as direct file access or DFA, reduces bandwidth usage and bottlenecks at the storage node, and provides highly efficient transmission of backup data.

If an NMDA backup cannot use the Client Direct feature due to the permissions or accessibility to the target device, the backup uses the storage node instead.

Deduplication backups and restores with EMC Avamar

NMDA supports deduplication backups and restores with an EMC Avamar® server.

An EMC Avamar server interacts with the NetWorker server and NMDA during deduplication backups and restores. After you configure the Avamar server as a NetWorker deduplication node, the Avamar server deduplicates the backup data from various clients, including the NMDA clients.

The initial backup to a deduplication node (Avamar server) backs up all the specified data. During subsequent deduplication backups, the Avamar software identifies the redundant data blocks on the NMDA client host and backs up only the unique blocks that contain changes.

Because the software performs data deduplication on the client host, deduplication backups typically require less time, less network bandwidth, and less storage space than nondeduplication backups.

Note: You can offload deduplication to a proxy host for the application that supports snapshot-based backups. The NetWorker Snapshot Management integration guide describes deduplication support.

During a restore, the Avamar server converts the stored data to its original nondeduplicated format before sending the data over the network.

An Avamar deduplication backup can be a manual backup or scheduled backup, including a probe-based backup.

You must configure a deduplication backup with an Avamar server according to "Configuring deduplication backups with Avamar" on page 95.

Deduplication backups and restores with EMC Data Domain

NMDA supports deduplication backups and restores with an EMC Data Domain® system. You can configure the Data Domain system as NetWorker AFTDs, virtual tape library (VTL) devices, or DD Boost devices.

A Data Domain deduplication backup to an AFTD, VTL, or DD Boost device can be a manual backup or a scheduled backup, including a probe-based backup. NMDA supports backups to DD Boost devices over IP or Fibre Channel (FC).

The first Data Domain backup backs up all the specified data and achieves the least amount of data deduplication, referred to as compression or global compression in Data Domain documentation. Subsequent Data Domain backups achieve improved deduplication rates as the backups identify more and more redundant data blocks.

Backups to AFTD or VTL devices on Data Domain

An NMDA backup to a Data Domain system configured as a NetWorker AFTD or VTL device sends all the data to the Data Domain system where the deduplication occurs. During a restore, the Data Domain system converts the stored data to its original nondeduplicated format before sending the data over the network.

Backups to DD Boost devices

An NMDA backup to a Data Domain system configured as a DD Boost device can take advantage of the DD Boost feature by using the following two components:

- The DD Boost library API enables the NetWorker software to communicate with the Data Domain system.
- The distributed segment processing (DSP) component reviews the data that is already stored on the Data Domain system and sends only unique data for storage.

NMDA can perform DD Boost backups through either the Client Direct workflow or the NetWorker storage node workflow:

1. Deduplication using Client Direct—With the Client Direct workflow, NMDA uses the DD Boost components to deduplicate the data locally and send only the unique blocks directly to the target device, bypassing the storage node. This configuration eliminates the sending of data over the network between the NMDA client and the storage node. NMDA sends the deduplicated data directly from the NMDA client to the DD Boost device.
2. Deduplication on a remote storage node—With the NetWorker storage node workflow, NMDA sends all the data to a remote storage node, which deduplicates the data and sends only the unique blocks to the target device. NMDA uses this workflow if the NMDA client does not have direct access to the Data Domain system or if the Client Direct workflow is not possible for other reasons.
3. Deduplication on a local storage node (located on the NMDA client)—This configuration has similar performance to deduplication using Client Direct. However, this configuration requires a storage node on the NMDA client and additional memory resources.

By default, NMDA tries to perform a Client Direct backup. If the NMDA client does not have direct access to the DD Boost device on the Data Domain system or if the Client Direct workflow is not possible for other reasons, then the backup is automatically routed through the storage node.

During a restore, if the NMDA client has direct access to the DD Boost device, then the client restores data directly from the device, regardless of whether the backup was done directly from the NMDA client or through the storage node. If the NMDA client cannot access the data directly, then the restore process reverts to the traditional method that uses the storage node. Regardless of the restore method used, the Data Domain system converts the stored data to its original nondeduplicated state before sending the data over the network.

You must configure a deduplication backup with a Data Domain system according to "Configuring Data Domain backups" on page 99.

Snapshot backups and restores

Snapshot technology provides enhanced protection and improved availability of data and greatly reduces the use of resources on the production host to perform backups.

NMDA supports snapshot backups and restores of DB2 and Oracle data through the NetWorker Snapshot Management (NSM) feature.

Note: With NetWorker release 8.1 or later, the NSM feature integrates and replaces the previous stand-alone EMC NetWorker PowerSnap™ Modules. NSM is available as part of the NetWorker client software.

NMDA works with NSM to perform snapshot backups and restores of DB2 or Oracle server data that resides on specific types of primary storage:

- Snapshot backups create point-in-time copies or snapshots of DB2 data or Oracle data on primary storage devices that NSM supports, such as EMC VMAX® (EMC Symmetrix®) and EMC VNX® Block (EMC CLARiiON®). Snapshot backups optionally perform a snapshot rollover by mounting the snapshot to a mount host and backing up the data to secondary storage, such as a Data Domain device or tape, from the point-in-time copies.
- Snapshot restores restore the DB2 or Oracle data from the snapshots or secondary storage.

Oracle software does not support NSM backups of Oracle data that resides on the Oracle Automatic Storage Management (ASM). However, NMDA uses EMC Replication Manager to perform snapshot (split-mirror-based) backups and restores of Oracle ASM data through disk replication technology.

You can optionally configure a DB2 or Oracle snapshot backup with NSM as a Cllient Direct backup to an AFTD or DD Boost device. The Client direct feature applies only to a snapshot rollover operation and sends the backup of the snapshot from the mount host to the device directly, bypassing the storage node.

Chapter 9, "Snapshot Backups and Restores," provides more details about snapshot backups and restores.

Virtualization support

NMDA supports several types of virtualization software, such as VMware, Solaris zones, Microsoft Hyper-V, and so on. The *EMC NetWorker Software Compatibility Guide* provides a full list of supported environments and platforms.

Note: You must install NMDA in a guest operating system.

VMware support

NMDA supports regular backups and restores of the database or application on a VMware virtual machine (VM) on an ESX server.

NMDA also supports the following advanced features of a VMware ESX server:

- VMotion—The VMotion feature enables migration of virtual machines from one ESX server to another while the servers are on. The migration is seamless to the applications that run on the virtual machines. Users do not experience disconnections during the migration. A migration that occurs during an NMDA backup or restore does not interrupt the backup or restore. VMware documentation describes the virtual machine requirements for VMotion.
- Distributed Resource Scheduler (DRS)—The DRS feature enables dynamic balancing and allocation of resources across multiple ESX servers. Depending on the DRS policies set by the user, the DRS can migrate a virtual machine or recommend that users migrate a virtual machine to a different ESX server by using VMotion. DRS can also start (at boot-up) a virtual machine on a different ESX server. As this feature uses VMotion, if a migration occurs during an NMDA backup or restore, the migration does not interrupt the backup or restore.
- High Availability (HA)—The HA feature enables the restart of VMware on the same ESX server, on a different ESX server, or on a physical machine, depending on the type of VMware cluster configured. During a restart, users are disconnected and must reconnect.

  If a restart occurs during an NMDA backup or restore, the backup or restore fails.

  - You must restart a manual backup or restore manually when the guest operating system restarts.
  - For a scheduled backup, the NetWorker server retries the backup if the Client Retries attribute in the Group resource has a nonzero value.

Solaris zones support

NMDA supports Solaris global zones, whole root zones, and sparse root zones. The *EMC NetWorker Module for Databases and Applications Installation Guide* provides installation details for sparse root zone environments.

Cluster backups and restores

NMDA supports backups and restores of active-passive clusters and active-active application clusters:

- An active-passive cluster includes multiple hosts (nodes) connected by a shared SCSI bus with common storage attached. You can define cluster services, such as disk services, and assign the services their own IP addresses and names (virtual cluster hosts). The services and their associated storage can migrate for failover between the hosts in the cluster.
- An active-active application cluster includes multiple database nodes on the same host or on different hosts that each contain a data server and can continue to process data if other nodes in the cluster fail. DB2 Database Partitioning Feature (DPF), DB2 pureScale, Informix Multi-node Active Clusters for High Availability or High Availability Clusters, Oracle Real Application Cluster (RAC), and Sybase ASE Cluster Edition systems are examples of active-active application clusters that NMDA supports.

  Note: NMDA documentation uses the MACH acronym for both Informix technologies: Informix Multi-node Active Clusters for High Availability and High Availability Clusters.

Chapter 6, "Cluster Systems," describes cluster systems and their configuration for NMDA backups and restores.

Configuration wizards

NMDA supports configuration wizards for configuring backup and recovery. The wizards are integrated with the NetWorker Management Console (NMC).

You can run the NMDA wizards from the NMC Administration window. You can start the NMC window on any supported host by using a web browser session and specifying the Console server URL.

NMDA supports the following configuration wizards:

- Backup configuration wizard that configures scheduled backups (either typical or custom) of DB2, Informix, Lotus Domino/Notes, MySQL, Oracle, and Sybase data.

  "Configuring scheduled backups with the wizard" on page 82 describes how to use the backup configuration wizard.
- Recovery configuration wizard that generates Oracle Recovery Manager (RMAN) scripts. The scripts perform the following operations:
  - Restore and recover Oracle data to the original host.
  - Create a duplicate Oracle database on a local or remote host.

  "Using the Oracle recovery configuration wizard" on page 212 describes how to use the recovery configuration wizard.

The wizards provide security and ease of management for the configurations of NMDA backups and recoveries. For example, the wizards use NetWorker lockbox services to store sensitive data.

Configuration conversion

NMDA supports two types of conversions:

- Conversion of a scheduled backup configuration when you update from a legacy NetWorker module to NMDA.

  You can use the **nsrdaadmin -M** command to perform either of these operations:

  - Convert the client-side configuration of a legacy NetWorker module to an NMDA client-side configuration.
  - Convert the server-side configuration of a legacy NMDB2 or NMO module to an NMDA server-side configuration.

  The *EMC NetWorker Module for Databases and Applications Installation Guide* describes how to convert the backup configuration of a legacy NetWorker module.

- Conversion of an NMDA client-side configuration to an NMDA server-side configuration.

  Except for a MySQL client-side configuration, you can use the **nsrdaadmin -W** command to convert an NMDA client-side configuration to an NMDA server-side configuration. The server-side configuration uses the configuration storage framework that the wizard supports.

  "Converting client-side to server-side configurations" on page 79 provides details.

Internationalization (I18N)

NMDA I18N is the capability of nonlocalized NMDA to operate in a non-English environment or locale. After you configure NMDA I18N as described in "Configuring internationalization (I18N) support" on page 71, NMDA can process and display non-ASCII data that the operating system, the NetWorker software, and the database or application software pass to NMDA. The non-ASCII data can include text messages, dates, times, numbers, and so on.

IPv6 support for backups and restores

NMDA supports the Internet Protocol version 6 (IPv6) for both nonsnapshot and snapshot backups and restores.

The *EMC NetWorker Installation Guide* describes the use of NetWorker software in an IPv6 environment.

Multiple database or application installations on the same host

NMDA supports multiple installations of the same application or different applications on the same host.

Chapter 7, "Multiple Installations on a Single Host," provides information about how to configure manual backups, scheduled backups, and recoveries with multiple applications on the same system, and details about coexistence of 32-bit and 64-bit NMDA applications.

Simplified scheduled backup configuration of multiple databases or applications

NMDA supports the scheduled backup of multiple databases or applications on the same client host by using a single NetWorker Client resource.

When you configure a scheduled backup without the wizard, you can specify multiple databases or applications in the Save Set attribute of a single Client resource as described in "Save Set" on page 91. You do not need to configure a separate Client resource for each database or application.

When you use a single Client resource for multiple databases or applications, you also use a single NMDA configuration file and specify unique parameter settings for each database or application.

The Parallelism attribute in the NetWorker Client resource or Group resource determines whether NMDA backs up the multiple databases or applications simultaneously or sequentially. The *EMC NetWorker Administration Guide* provides details about the NetWorker client and savegrp parallelism.

NMDA features specific to DB2

NMDA supports specific features of DB2 backup and restore operations and DB2 history pruning. In addition to restores of NMDA DB2 backups, NMDA supports restores of Fujitsu NetWorker Module for DB2 backups.

DB2 full, incremental, and delta backups

NMDA supports the following types of DB2 backups:

- Full backup—A DB2 full backup is the building block of any recovery strategy. You can start a restore only with a full backup. If you perform an online backup, you need the logs of all the transactions that transpired during the backup, unless you specify the **include logs** option during the online backup.

  Note: You can send the DB2 data and logs to separate backup pools during a single NMDA backup.

- Incremental backup—A DB2 incremental backup includes all the changes since the last full backup.
- Delta backup—A DB2 delta backup includes all the changes since the last backup of any type.

DB2 transaction log backups

NMDA supports DB2 backups of transaction logs and rollforward recovery with the DB2 transaction logs. These transaction logs keep records of changes made to DB2 databases.

DB2 software provides two types of transaction logging:

- Circular logging is the default behavior when you create a DB2 database. With this type of logging, each full backup deletes the transaction logs. You can restore only full backups. Circular logging supports only full offline backups of databases.

- Archive logging supports online backups and rollforward recovery. With this type of logging, the transaction logs are retained as archive logs and you can back up the logs. You can recover a database or tablespace to a specific point-in-time by "rolling forward" through the transaction logs in sequence until the specified point.

Ensure that the DB2 archive transaction logs are backed up, for example, by configuring the automatic backup of DB2 transaction logs as described in "Configuring automatic backups of DB2 transaction logs" on page 116.

DB2 history pruning

NMDA supports the synchronous removal of snapshot backup entries from the DB2 history file when the corresponding entries expire and are removed in the NetWorker indexes. "Pruning DB2 snapshots" on page 307 includes configuration details for this synchronous removal ("pruning") of snapshot entries from the DB2 history file.

For regular backups, you can delete the backup entries for databases or tablespaces from the DB2 server and NetWorker server with the **db2 prune** command. Deletion of backup entries might be necessary if the NetWorker index and DB2 history files become excessively large and the retention period is high. "Deleting DB2 backups" on page 164 provides details.

Note: You cannot delete snapshot backups with the **db2 prune** command.

DB2 10.5 features

NMDA supports the new features of DB2 release 10.5 except for DB2 high-availability disaster recovery (HADR) in a DB2 pureScale environment.

NMDA supports the following new features of DB2 release 10.5:

- Increased availability support for topology changes in a DB2 pureScale environment.

  For example, after you add a new pureScale node, DB2 software no longer requires an offline immediate database backup. You can restore backup images created before the topology change, and roll forward. To enable the automatic backup of the archived logs on the new node, ensure that NMDA is installed and configured on the new node as described in "Configuring backups and restores in a DB2 pureScale system" on page 261.

- Easier recovery in a DB2 pureScale environment.

  For example, you can restore a database or tablespace backup image to a pureScale instance with a different topology.

- Restores of backup images between a DB2 pureScale instance and a DB2 Enterprise Server Edition instance.

- DB2 Advanced Workgroup Server Edition.

The IBM DB2 documentation provides details about the DB2 10.5 features.

Restores of Fujitsu NetWorker Module for DB2 backups

NMDA can restore backups performed with Fujitsu NetWorker Module for DB2 version 2.1 or 4.0 on an operating system that this NMDA release supports. The *EMC NetWorker Software Compatibility Guide* at EMC Online Support provides details about the operating systems that NMDA supports.

If the backup with Fujitsu NetWorker Module for DB2 used advanced backup compression with GZIP or BZIP2, NMDA will automatically uncompress the backup during the restore.

If you convert from using Fujitsu NetWorker Module for DB2 to using NMDA, you must configure backups of DB2 data according to the NMDA documentation. For example, NMDA operations ignore the Fujitsu-specific compression options GZIP and BZIP2 if the options are set.

NMDA features specific to Informix

NMDA supports Informix full and incremental backups and logical log backups. In addition to restores of NMDA Informix backups, NMDA supports restores of Fujitsu NetWorker Module for Informix backups.

Informix terminology

In this guide, the term IDS refers to the following software:

- For Informix releases earlier than11.70, Informix Dynamic Server.
- For Informix release 11.70 or later, Informix Database Server.

For Informix release 11.70 or later, replace all instances of “Informix Dynamic Server” in this guide with “Informix Database Server.”

Informix full and incremental backups

NMDA supports all the Informix backup levels that the IDS ON-Bar utility supports:

- Level 0 backup—An Informix level 0 backup is a full backup, which includes all the (full) records in the selected dbspaces. This is the only type of Informix backup that enables a complete restore without recovery steps, such as applying logs and incremental backups.
- Level 1 backup—An Informix level 1 backup is an incremental backup, which backs up the (incremental) records that have changed since the last level 0 backup in the selected dbspaces.
- Level 2 backup—An Informix level 2 backup backs up the records that have changed since the last level 1 backup in the selected dbspaces.

Logical log backups

NMDA supports the following types of IDS logical log backups:

- Automatic, continuous logical log backups (recommended)
- Manual logical log backups

"Configuring automatic (continuous) backups of Informix logical logs" on page 117 describes the configuration of the IDS logical log backups to ensure that ON-Bar automatically backs up the logical logs as they become full.

Note: You can send the Informix data and logs to separate backup pools during a single NMDA backup.

Restores of Fujitsu NetWorker Module for Informix backups

NMDA can restore backups performed with Fujitsu NetWorker Module for Informix version 2.0 on an operating system that this NMDA release supports. The *EMC NetWorker Software Compatibility Guide* at EMC Online Support provides details about the operating systems that NMDA supports.

If the backup with Fujitsu NetWorker Module for Informix used advanced backup compression with GZIP or BZIP2, NMDA will automatically uncompress the backup during the restore.

If you convert from using Fujitsu NetWorker Module for Informix to using NMDA, you must configure backups of Informix data according to the NMDA documentation. For example, NMDA operations ignore the Fujitsu-specific compression options GZIP and BZIP2 if the options are set.

Informix 12.10 features

NMDA supports Informix 12.10, including new features such as **onmode** utility enhancements. However, the following limitations apply with Informix 12.10:

- You must not adjust ONCONFIG parameters including BAR_* dynamically by using **-wf** or **-wm** during backup or recovery operations as doing so might have a negative impact on the backup or recovery.
- NMDA does not support backup and recovery filters.

NMDA features specific to Lotus

NMDA supports specific features of Lotus Domino and Notes backup and restore operations.

Files backed up during Lotus backups

NMDA backs up the following Lotus files by default:

- .nsf, .ntf, and .box files

  NMDA considers files with these file name extensions to be database files, and backs up the files by using the Lotus application programming interface (API) for databases.

- .ncf, .njf, .nrf, .dic, .dsk, .id, and notes.ini files

  NMDA considers these files to be regular operating system (nondatabase) files and backs up the files at the file system level.

- Domino transaction logs in archive mode

Note: NMDA does not back up transaction logs in circular or linear mode.

You can also use NMDA to back up the following files when you set the NSR_BACKUP_ALL_EXTENSIONS parameter to TRUE:

- DAOS files (with .nlo file name extension)

  "Configuring Lotus DAOS backups" on page 121 describes Lotus DAOS backups.

- Any other files that exist in a Lotus directory, such as .gif, .html, and .doc files

"Lotus full and incremental backups" on page 40 describes which files are backed up and at which backup level.

By default during Lotus backups, NMDA follows directory links (with a .dir extension) and database links (with a .nsf extension) when forming the file list for backup. NMDA backs up both the Lotus link files and the data files or data directories that the link files point to. "Considerations for Lotus database or directory link backups" on page 118 describes how to disable the default behavior during database link backups or directory link backups.

After backing up the Domino data, NMDA creates a browselist file and backs up the file. The browselist is required for restores because the browselist contains browsing information for the backed-up Lotus files.

If a backup of Lotus data succeeds but the browselist backup fails, you can still recover the data by using the **nsrnotesrc** command. However, you must specify each file to be restored in the NSR_BACKUP_PATHS parameter setting or in a file pointed to by the NSR_RECOV_LIST_FILE parameter. "Performing Lotus database recovery with the nsrnotesrc command" on page 182 describes how to use the **nsrnotesrc** command.

Note: If a Lotus data backup succeeds but the browselist backup fails, you cannot use the NetWorker User for Lotus GUI to recover the data.

During a Lotus backup, the first 10 MB of generated browselist data is stored in a memory cache, which is typically enough for most Lotus backups. For larger backups, additional browselist data is stored in a temporary file at a location that you can specify with the NSR_BROWSELIST_CACHE_DEST parameter. "NSR_BROWSELIST_CACHE_DEST" on page 373 provides details.

Lotus full and incremental backups

NMDA supports full and incremental backups of Lotus data:

- Full backup—Backs up the specified files, regardless of whether or not the files have changed since the last backup operation. If you enable Lotus transactional logging and set the logging to archive mode, NMDA does not back up the logs unless you request the log backup through the NSR_BACKUP_LOGS_MODE parameter setting or corresponding wizard option.
- Incremental backup—The behavior of an incremental backup depends on the settings for Lotus transactional logging on the Domino server:
  - If you enable Lotus transactional logging and set the logging to archive mode, then an incremental backup backs up specific files:
    - Backs up the transaction logs unless the NSR_INCR_BACKUP_LOGS_MODE setting or corresponding wizard option specifies to not back up the logs.
    - Backs up any database files that are not logged and have changed since the last backup.
    - Backs up the database files that are logged and have had their database instance ID (DBIID) property changed since the last backup of the database.
  - If you disable Lotus transactional logging (or enable the logging but do not set it to archive mode), then an incremental backup backs up only the database files that have changed since the last backup.

  The incremental backup always backs up the nondatabase files that have changed since the last backup, regardless of the transactional logging mode.

Lotus transaction log backups

NMDA supports the backup of Lotus transaction logs only, without the backup of other Lotus files. If you enable Lotus transactional logging and set the logging to archive mode, then you can configure an NMDA manual or scheduled backup of the transaction logs only.

Transaction log only backups cannot replace full backups or incremental backups. You must still perform full and incremental backups to completely protect the Domino server. A transaction log backup does not back up any Lotus (database or nondatabase) data files. You might want to perform a full or incremental backup every 24 hours at a minimum and schedule more frequent backups of transaction logs only.

You cannot run multiple simultaneous transaction logs only backups for the same Domino server.

Lotus DAOS backups and restores

NMDA supports backups and restores of the attachments that the Domino Attachment Object Service (DAOS) manages. Domino 8.5 and later supports the DAOS feature.

You can configure a manual or scheduled backup to back up the DAOS base directory, separate DAOS subdirectories, or individual DAOS objects, and to run as either of the following backups:

- A stand-alone backup
- An integrated backup that backs up the DAOS files after the Domino database data

The appropriate IBM documentation describes the features and setup of DAOS directories and the DAOS files that IBM calls NLO files.

"Configuring Lotus DAOS backups" on page 121 describes how to configure backups of Domino with DAOS enabled.

"Performing Lotus data restore and recovery" on page 181 describes how to restore the databases that contain links to missing DAOS files.

Restartable scheduled backups

NMDA supports Lotus restartable scheduled backups through the NetWorker checkpoint restart feature, which enables you to restart an NMDA Lotus backup from a good point called a checkpoint. If NMDA successfully backs up some files during a failed backup, NMDA will not back up the files again during a restarted backup. NMDA restarts the backup only within the restart window of the backup group.

The *EMC NetWorker Administration Guide* describes restarted backups for a backup group.

For Lotus restartable backups, the checkpoint granularity is always at the file level. The backups ignore the setting of the Checkpoint granularity attribute in the NetWorker Client resource.

Note: Configuring a scheduled backup as checkpoint-enabled might impact the backup speed due to the synchronous communication between the NMDA Lotus client and the NetWorker server and storage node.

Depending on when a backup fails, the restarted backup might not back up certain files:

- If the backup fails while backing up data files, the restarted backup:
  1. Rescans all the data files.
  2. Backs up any files that were not backed up previously or that have changed since the previous backup.
- If the backup fails while saving transaction logs, NMDA does not rescan the data files. NMDA restarts the backup at the transaction log level.
- If the integrated DAOS backup fails while saving the NLO files, NMDA does not rescan the database data files, but it does rescan the DAOS files.

You must configure a restartable backup for an NMDA Lotus client according to "Configuring Lotus restartable scheduled backups" on page 124.

Types of Lotus restores

NMDA supports the following types of restore methods for Lotus data:

- Database-level (file-level) restore—Restores the databases of a Domino server.
- Document-level restore—Restores modified or deleted Notes documents in a single database, whether the database is logged or not:
  - You can perform document-level recovery of deleted Notes documents in the local database through the **nsrdocrc** command line program. "Performing Lotus document-level recovery with nsrdocrc command" on page 192 provides details.
  - On Windows systems only, you can use the Lotus Notes client GUI to perform document-level recovery of modified and deleted Notes documents either in the local Notes or Domino database or in a remote Domino database. "Performing Lotus document-level recovery with the Notes client GUI" on page 194 provides details.

NetWorker User for Lotus program

On Windows systems only, you install the NetWorker User for Lotus GUI program (**nwbml.exe**) with NMDA.

The NetWorker User for Lotus program provides a graphical interface for performing manual backups and recovery operations.

You cannot use the NetWorker User for Lotus program to perform a transaction logs only backup.

The following topics describe how to use the NetWorker User for Lotus program for backups and restores:

- "Performing Lotus manual backups with NetWorker User for Lotus" on page 151
- "Performing Lotus database recovery with NetWorker User for Lotus" on page 185

Partitioned Domino servers

NMDA supports the backup and recovery of partitioned Domino servers.

"Configuring partitioned Domino server backups" on page 120 describes how to configure the backups of partitioned Domino servers.

NMDA features specific to MySQL

NMDA supports specific features of MySQL backup and restore operations. NMDA interacts with the MySQL Enterprise Backup (MEB) software through the Oracle SBT interface to perform MySQL backups and restores.

MySQL online and offline backups

A MySQL server instance can be either online or offline during an NMDA MySQL backup:

- The online backup mode depends on the type of MySQL storage engine:
  - With an InnoDB storage engine, NMDA performs a hot online backup with the database in a read/write state during the backup.
  - With a MyISAM storage engine, NMDA performs a warm online backup. The MEB software places the databases or tables in a read-only state during the backup.
- An offline backup requires the user or administrator to shut down the instance before running the backup.

MySQL full and incremental backups

NMDA supports full and incremental level backups of MySQL data:

- A full backup backs up all the data of a specified MySQL instance or all the data of specified databases or tables.
- An incremental backup backs up only the data that has changed since the last backup. An incremental backup can be cumulative or differential:
  - A cumulative incremental backup includes only the data that has changed since the last full backup.
  - A differential incremental backup includes only the data that has changed since the last full or incremental (cumulative or differential) backup.

An incremental backup includes changed blocks or files, depending on the type of MySQL storage engine:

- With an InnoDB storage engine, an incremental backup includes the data blocks that have changed since the last backup.

  Note: Starting with MEB 3.7, you can also perform a redo log only incremental backup for an InnoDB storage engine, which is a differential incremental backup that includes only the redo log changes since the last full or incremental backup.

- With a MyISAM storage engine, an incremental backup includes the table datafiles that have changed since the last backup.

MySQL binary log backups

NMDA supports backups of the MySQL binary logs. When you enable binary logging for a MySQL instance, the binary logs record all the data changes for the MySQL instance. The MySQL documentation describes binary logs and how to enable binary logging.

The MySQL instance must be online to perform binary log backups.

NMDA can restore the MySQL binary log backups for a server instance recovery to the current time or a point-in-time.

NMDA supports both of the following types of binary log backups:

- Backup of the binary logs immediately after a full backup of a whole instance.
- Backup of only the binary logs for an instance.

You might want to perform a full or incremental backup every 24 hours at a minimum, and schedule more frequent backups of only the binary logs.

Note: You can send the MySQL data and binary logs to separate backup pools during a single NMDA backup.

You can optionally specify that NMDA deletes the binary logs from the disk after completing a binary log backup.

During a binary log backup, NMDA uses the NetWorker **save** program to back up all the binary logs created since the last successful log backup.

Note: NMDA does not back up any binary logs during an incremental backup or a partial backup.

MySQL backup granularity

When NMDA backs up a MySQL instance or MySQL databases or tables, NMDA performs either a whole instance backup or a partial backup:

- NMDA performs a whole instance backup by default. A whole instance backup backs up an entire MySQL instance, including all the databases in the instance and all the tables in each database. A whole instance backup can be a full or incremental backup.
- NMDA performs a partial backup to back up any combination of specified databases and tables for an InnoDB or MyISAM storage engine. A partial backup can be a full or incremental backup.

  Note: A partial backup for an InnoDB storage engine always includes the system tablespace and all the tables within it.

NMDA supports all the following types of partial backups:

- With a MyISAM storage engine, the backup of any combination of specified databases and specified tables within the same database or different databases.
- With an InnoDB storage engine:
  - The backup of specified databases.
  - The backup of specified tables when the file-per-table option is enabled.

    Note: If the file-per-table option is disabled, then the backup includes all the tables in the InnoDB database.

- With MEB 3.7 or later, the backup of all the tables and their associated .frm files in a specified instance.
- With MEB 3.7 or later, the backup of specified tables (when the file-per-table option is enabled) and their associated .frm files.

MySQL restore and recovery operations

NMDA supports the restore and recovery of MySQL backups performed with NMDA. You can use the NMDA program **nsrmysqlrc** to perform the following restore and recovery operations:

- Restore and recovery of a whole instance backup to the current time or a point-in-time.
- Restore and recovery of a partial backup to the current time or a point-in-time.
- Restore of one or more binary log backups.
- Redirected restore and recovery.
- Disaster recovery.
- Additional advanced restore-related operations:
  - List image operation—Lists backup files from a backup image.
  - Extract operation—Extracts backup files from a backup image.
  - Extract and prepare operation—Extracts backup files from a backup image and processes the files to produce a prepared backup that is ready for restore.
  - Copy back operation—Copies a prepared backup to a specified directory, for example, to complete a restore to the database server.

The NMDA program **nsrmysqlrc** provides enhanced functionality and usability over the MEB utilities **mysqlbackup** and **mysqlbinlog**. NMDA tracks all the required details about NMDA MySQL backups that enable you to perform any of the supported restore and recovery operations through a single **nsrmysqlrc** command.

The restore and recovery of a whole instance backup and the restore of a partial backup both include an “extract and prepare” operation and a “copy back” operation.

Any restore or recovery that includes a copy back operation to the MySQL data directory requires a shutdown of the database server. Unless you disable the prompt, the following operations prompt you to shut down the server before the prepared backup is copied back to the data directory:

- Restore and recovery of a whole instance backup to the current time or a point-in-time
- Restore of a partial backup to the current time or a point-in-time
- Copy back operation

Note: The **nsrmysqlrc** program restores all the data of a whole instance backup or partial backup. For example, when you restore a partial backup that includes three databases db1, db2, and db3, the **nsrmysqlrc** program restores all three databases. You cannot use **nsrmysqlrc** to restore only part of a backup, such as only one of the three databases included in the partial backup.

MySQL validate operation

With MEB 3.7 or later, you can use NMDA to validate the integrity of the backup image produced by an NMDA MySQL backup. The validate operation does not alter the backup image.

You do not need to extract any files from the backup image prior to the validate operation. You must only ensure that the device containing the backup image is mounted.

MySQL 5.6 features

NMDA supports the new features of MySQL release 5.6 except for the use of an encrypted login file, mylogin.cnf, for setting the backup credentials. You must set the NMDA MySQL backup credentials in the MySQL configuration file or in the NMDA configuration file.

NMDA supports the following new features in MySQL release 5.6 with MEB 3.8.*x*. The *EMC NetWorker Module for Databases and Applications Release Notes* describes the supported MEB versions:

- You can store InnoDB tables in a specified directory outside the MySQL data directory. MySQL 5.6 uses the InnoDB file-per-table mode as the default tablespace management mode, storing an InnoDB table and its indexes in a separate .ibd file.
- You can store the InnoDB undo logs or rollback segments in one or more separate tablespaces outside of the system tablespace. You can also optionally store the InnoDB undo logs outside of the MySQL data directory. The tablespace files for the undo logs are named undo001, undo002, and so on.
- You can use the **innodb_page_size** parameter to specify a page size for InnoDB tablespaces in a MySQL instance.
- You can use the **innodb_checksum_algorithm** parameter to specify how to generate and verify the checksum stored in each disk block of each InnoDB tablespace.

The MySQL documentation provides details about the MySQL 5.6 features.

"Configuring MySQL 5.6 features" on page 128 describes how to configure the MySQL 5.6 features for NMDA operations.

"Performing MySQL restores of InnoDB tables outside data directory" on page 206 describes the MySQL 5.6 restore procedure to a new directory structure.

NMDA features specific to Oracle

NMDA supports specific features of Oracle Server backup and restore operations. In addition to restores of NMDA Oracle backups, NMDA supports restores of Fujitsu NetWorker Module for Oracle backups.

NMDA interacts with RMAN software through the Oracle SBT interface to perform Oracle backups and restores.

Note: Oracle software does not support RMAN backup encryption to NetWorker or any backup vendor software except Oracle Secure Backup. For Oracle backup encryption, you can use NetWorker AES encryption by setting the NSR_AES_ENCRYPTION parameter. "NSR_AES_ENCRYPTION" on page 358 provides details.

Oracle documentation describes the Oracle RMAN backup and restore operations, including the RMAN commands that specify the operations.

Oracle full and incremental backups

NMDA supports full and incremental backups of Oracle data:

- A full (or stand-alone full) backup includes every used block of the database objects specified in the RMAN backup script. The backup skips never-used blocks. This type of backup occurs when you do not specify a backup level with the RMAN **backup** command.

  Note: A full backup cannot be the parent of a subsequent incremental backup. Incremental backups cannot be dependent on a stand-alone full backup.

- An incremental backup includes blocks that have changed since the previous specified backup. Incremental backups occur when you specify either **incremental level=0** or **incremental level=1** with the RMAN **backup** command. Incremental backups are dependent on preceding incremental backups in the same scheduled backup cycle:

  - A level 0 incremental is physically identical to a full backup, but is recorded as incremental in the RMAN repository.

    Note: A level 0 backup is also referred to as "full" in other topics in this guide.

  - A level 1 incremental can be either of the following backups:
    - A differential backup, which contains only the data blocks that changed since the most recent incremental backup, whether level 0 or 1. The differential backup is dependent on the preceding level 0 or 1 backup. Incremental backups are differential by default.
    - A cumulative backup, which contains only the data blocks changed since the most recent level 0 incremental backup. The cumulative backup is dependent on the preceding level 0 backup.

Backups of archived redo logs

Backups of archived redo logs enable recovery of the database to its predisaster state. Without these backups, you can recover the database only to the time of the last consistent Oracle backup. In this case, you will lose the transactions that occurred between the time of the last consistent backup and the time of the database corruption.

You might want to perform a full or incremental backup every 24 hours at a minimum, and schedule more frequent backups of only the archived redo logs.

You can back up archived redo logs by using the appropriate option of the RMAN **backup** command.

You can configure scheduled backups of archived redo logs by using the Oracle configuration wizard.

"Backing up all archived logs from each node" on page 272 provides a sample script to back up the archived redo log files in an Oracle Real Application Cluster (RAC) system.

Oracle DBA and NetWorker backup administrator collaboration

NMDA supports the collaboration of the Oracle DBA and NetWorker backup administrator for simplified backups and restores of Oracle disk backups. Oracle disk backups are RMAN disk type backups performed without the use of any backup software.

The Oracle DBA and NetWorker backup administrator have separate roles. NMDA empowers both the DBA and backup administrator to perform their separate tasks by enabling specific functionality:

- Without requiring NMDA or NetWorker knowledge, the DBA performs RMAN disk backups to an FRA (Fast Recovery Area or Flash Recovery Area). The DBA performs a one-step recovery, without knowing whether the backups are on the FRA or on a NetWorker device.
- Without requiring Oracle knowledge, the backup administrator moves the disk backups to the NetWorker server, catalogs the backups, and reports on what was backed up.

Chapter 8, "Oracle DBA and NetWorker Backup Administrator Collaboration," provides details about the Oracle DBA and NetWorker backup administrator collaboration and the NMDA support of Oracle DBA disk backups.

Control file and server parameter file autobackups

In addition to regular control file backups, Oracle performs a control file autobackup after each RMAN **backup** command if you have enabled the control file autobackup.

With a control file autobackup, RMAN can recover the database even if the current control file, the recovery catalog, and the server parameter file are inaccessible. The path used to store the autobackup follows a well-known format.

You can specify persistent settings for the control file autobackups with the **configure controlfile autobackup** command. For example, you can enable control file autobackup and specify the persistent setting for the format of the control file autobackup name with the following commands:

```
configure controlfile autobackup on
configure controlfile autobackup format for device type 'sbt_tape' to
   '/NMDA_%F/'
```

If you enable the control file autobackup, NMDA backs up the control file.

With the control file autobackup, Oracle software backs up the current server parameter file.

Automatic channel allocation

Oracle RMAN supports automatic channel allocation. This feature enables the configuration of persistent settings for automatic channels, for use in all RMAN sessions.

"Configuring automatic channel allocation" on page 132 provides configuration details.

Backup and restore optimization

If you enable Oracle backup optimization with the **configure backup optimization on** command, RMAN skips selected files during a backup, based on several criteria. The Oracle backup and recovery documentation describes these criteria.

To force a backup that would otherwise be skipped due to backup optimization, you can use the **force** option in the **backup** command.

The restore optimization function prevents RMAN from restoring a file if the original file is already in the correct location and contains the expected information.

To force a restore that would otherwise be skipped due to restore optimization, you can use the **force** option in the **restore** command.

NOTICE

If you use Oracle backup optimization with NMDA backups and restores, run the **crosscheck** command regularly to synchronize the Recovery Catalog and NetWorker indexes. This synchronization ensures that backups expired by the NetWorker server are also marked as expired in the Recovery Catalog and RMAN does not skip a backup when a referenced backup has already expired in NetWorker.

Backup copies

RMAN with NMDA can create copies of the backup, also known as duplexing the backup. RMAN can produce up to four identical copies of each backup piece on different NetWorker volumes with one **backup** command.

NMDA supports backup copies with manual Oracle backups only. NMDA does not support the use of the RMAN backup copies commands during scheduled Oracle backups.

"Creating Oracle backup copies" on page 133 describes the requirements for generating backup copies during manual Oracle backups.

Backup and deletion of backup sets

Oracle RMAN supports the backup of backup sets. If you use the **device type disk** option to back up Oracle data, you can use NMDA to back up these backup sets from a disk to NetWorker volumes.

For example, to back up all backup sets from a disk to NetWorker volumes, you can use the following command:

```
backup device type sbt backupset all
```

You can also delete the backup set on a disk with the **delete input** option in the **backup device type sbt backupset** command. For example, to back up the backup sets created on a disk more than a week ago and then remove the backup sets from the disk, you can use the following command:

```
backup device type sbt backupset completed before sysdate-7 delete
   input
```

Oracle Data Guard support

NMDA supports Oracle Data Guard, which is an Oracle data availability and protection solution for a primary database and one or more standby databases over an IP network. As transactions occur in the primary database and as Oracle writes redo data to the local redo logs, Data Guard automatically performs the following operations:

- Transfers this redo data to the standby sites.
- Applies the redo data to the standby databases, which synchronizes the standby databases with the primary database.

You can offload RMAN backups of datafiles, archived redo logs, and possibly other files to a physical standby database. You can then use the backups to recover the primary or standby database. RMAN and Data Guard documentation describes how to configure and back up a physical standby database, and use the backups to recover the primary or standby database.

"Configuring operations in an Oracle Data Guard environment" on page 137 provides configuration details for NMDA operations.

Oracle Exadata support

NMDA supports Oracle Exadata in both of the following configurations:

- Exadata Database Machine
- Exadata Storage Server, attached to an external database server

NMDA supports the same environment for Oracle Exadata (including the Oracle database versions, operating system versions, and NetWorker versions) that NMDA supports for Oracle RAC, as described in the *EMC NetWorker Software Compatibility Guide*.

You must install and configure NMDA on each Exadata database server in the same way that you install and configure NMDA for Oracle RAC. The installation and configuration procedures are described in the *EMC NetWorker Module for Databases and Applications Installation Guide* and in this administration guide.

To perform backups and restores of Exadata, you must follow the instructions for Oracle RAC environments in this administration guide.

If the Oracle Exadata database servers are connected to a NetWorker remote storage node, then the following requirements apply to the connection:

- The only supported protocol is IP, over InfiniBand or other networks like Ethernet. Native InfiniBand or any protocol other than IP is not supported.
- There could be multiple network interfaces on the Oracle Exadata database servers to the storage node, such as one for an InfiniBand connection and one for an Ethernet connection. In such a case, you must configure the network settings and the NetWorker Client resource to send the backup data from the Oracle Exadata server to the remote storage node over the preferred network, such as InfiniBand. The configuration details are available in the NetWorker documentation.

Restartable backups

Oracle RMAN can back up files that the software has not backed up since a specified time. For example, to continue the backup of an Oracle database that you cancelled two days ago, you can use the following command:

```
backup device type sbt database not backed up since time 'sysdate-2'
```

RMAN compares the time specified in this command with the completion time of the most recent backup of any datafile to determine if it requires backup. The appropriate Oracle backup and recovery documentation provides details.

Retention policies

Oracle RMAN provides an Oracle retention policy for backups. An Oracle retention policy is based on the recovery window or on redundancy. The retention policy is not based on a static time period, such as a year. Oracle RMAN considers a backup to be obsolete when the backup is no longer required according to the Oracle retention policy. Oracle RMAN checks the retention policy of a backup when you run the **report obsolete** or **delete obsolete** command.

NMDA supports the Oracle retention policy with some restrictions. The NetWorker server has its own browse policy and retention policy to specify how long data is available for recovery. The NetWorker policies are based on a user-defined time period. As the Oracle retention policy is independent from that of the NetWorker server, the NetWorker and Oracle policies could conflict.

To prevent conflicts, do not use both the NetWorker and Oracle policies. Instead, perform either of the following actions:

- If you want to use only the NetWorker server policy, disable the Oracle retention policy with the following command:

  ```
  configure retention policy to none
  ```

- If you want to use only the Oracle retention policy, set the NSR_ORACLE_RETENTION parameter to TRUE in the NMDA configuration file. NMDA supports this parameter only for scheduled backups, not manual backups.

  With NSR_ORACLE_RETENTION set to TRUE, NMDA performs the following actions:

  1. NMDA retrieves the Oracle retention policy from RMAN and sets the NetWorker retention policy accordingly.
  2. NMDA automatically enables policy uniformity so that backup pieces do not expire unless all the dependent backups expire. For example, a full backup does not expire unless all the dependent incremental backups expire.

  Note: NMDA supports the NSR_ORACLE_RETENTION parameter only with a backup window-based Oracle retention policy. Do not use a redundancy-based policy with the NetWorker server.

Because the NetWorker server automatically expires the backups based on the Oracle retention policy, you do not need to run the **delete obsolete** command. The **delete obsolete** operation is an expensive operation on the NetWorker server. The operation also requires special privileges on the NetWorker server as described in Table 4 on page 76.

If you use the Oracle retention policy, the DBA should run the RMAN **crosscheck** command regularly to keep the RMAN catalog in sync with the NetWorker index. The **crosscheck** command guarantees that all the expired backups will not be flagged as expired in the RMAN catalog and will not be used by RMAN. Due to the overhead on the NetWorker server, target the **crosscheck** command for a specific time period, rather than to crosscheck everything in the catalog. For example, run the following command:

```
crosscheck backup device type sbt completed after 'sysdate-45' before
   'sysdate-30';
```

Restores of Fujitsu NetWorker Module for Oracle backups

NMDA can restore backups performed with Fujitsu NetWorker Module for Oracle version 5.0 on an operating system that this NMDA release supports. The *EMC NetWorker Software Compatibility Guide* at EMC Online Support provides details about the operating systems that NMDA supports.

If the backup with Fujitsu NetWorker Module for Oracle used advanced backup compression with GZIP or BZIP2, NMDA will automatically uncompress the backup during the restore.

If you convert from using Fujitsu NetWorker Module for Oracle to using NMDA, you must configure backups of Oracle data according to the NMDA documentation. For example, NMDA operations ignore the Fujitsu-specific compression options GZIP and BZIP2 if the options are set.

Save set bundling

If you configure NMDA save set bundling, NMDA automatically creates a save set bundle to group all dependent save sets from the same backup cycle. Save sets are dependent when two or more save sets are required to restore a database object.

A backup cycle includes the following backups:

- A level 0 incremental backup of the database object.
- All subsequent level 1 incremental backups that are dependent on the level 0 backup.

NOTICE

NMDA does not support save set bundling for nonsnapshot manual backups or snapshot backups. NMDA performs save set bundling for nonsnapshot scheduled Oracle backups only.

Save set bundling automatically enables the following features for Oracle operations:

- Improved staging

  Oracle-aware staging causes NMDA Oracle save sets that have a dependency on each other to be staged together:

  - During automatic staging, the NetWorker server stages an entire save set bundle if the staging criteria determine both of the following conditions are true:
    - A particular NMDA save set must be staged.
    - The save set is part of the save set bundle.
  - During manual staging with the **nsrstage** command, if one or more save sets being staged are from a save set bundle, the command stages all the save sets in the bundle.

- Policy uniformity

  Policy uniformity is enabled automatically whenever you enable save set bundling. If you do not want to use save set bundling, you can enable policy uniformity separately. "Policy uniformity" on page 55 provides details.

Note: After a staging operation that stages all the save sets in a bundle, the resulting available space on the staging device might exceed the lower-water mark specified in the staging policy.

The *EMC NetWorker Administration Guide* describes how to use staging policies and perform automatic and manual staging operations through the NetWorker server.

"Configuring save set bundling for scheduled Oracle backups" on page 135 describes how to configure save set bundling for NMDA scheduled backups.

If an error occurs during save set bundling, the bundling operation fails but the scheduled backup can finish successfully. Information about the bundling failure appears in the **savegrp** output and the NMDA debugging file.

NetWorker staging restrictions

When planning the strategy for NMDA save set bundling, consider the following NetWorker staging restrictions:

- The NetWorker server cannot simultaneously stage all the save sets from a save set bundle if some of the save sets were backed up to separate volumes. The server simultaneously stages save sets only if they reside on the same staging volume. Example 3 on page 55 provides details.

  To ensure the proper staging of all the save sets from a save set bundle, do not split the backup between different staging volumes. If required, split the backup into different backup cycles, with each cycle going to a separate volume.

- NetWorker staging policies must not cause staging of the save sets of an NMDA backup cycle before the cycle is complete. For example, if a one-week NMDA cycle starts on Sunday, the staging policy must not cause staging of the partially complete save set bundle before the final backup of the cycle occurs on Saturday.

To prevent a staging operation from splitting an NMDA backup cycle, adjust the NetWorker staging policy. For example, adjust the policy to stage older save sets before new ones, or adjust the high-water and low-water marks.

The *EMC NetWorker Administration Guide* describes how to work with staging policies and perform automatic and manual staging operations through the NetWorker server.

Save set bundling processes

The NMDA program **nsrdasv** automatically places save sets into a save set bundle at the end of a scheduled backup.

Note: The save set bundling operation requires specific privileges as described in Table 4 on page 76.

To perform save set bundling, the **nsrdasv** program connects to the Oracle database by attempting to use the login and password from the RMAN script. If the script does not include a login and password, the program performs the following actions:

- Uses the ORACLE_SID value from the NMDA configuration file to search the nwora.res file for the NSR_ORACLE_CONNECT_FILE parameter.
- Uses the connection strings from the specified connection file.

After connecting to the Oracle database, the **nsrdasv** program obtains all the required information about the backups by using the V$ views. "Configuring save set bundling for scheduled Oracle backups" on page 135 describes the nwora.res file and the save set bundling requirements.

The **nsrdasv** program creates a save set bundle for each incremental level 0 backup. The program adds the save sets from subsequent incremental backups to the bundles of the level 0 backups that they are dependent on. Example 1 on page 54 and Example 2 on page 55 illustrate different scenarios for forming a save set bundle.

The name that the **nsrdasv** program assigns to a save set bundle is the save time of the oldest save set in the bundle.

After a scheduled backup, the NetWorker server stores the save set bundle name and the list of save sets it contains in the media database. You can view the bundle information by using the **mminfo** command, as described in "Save set bundling information in the media database" on page 55.

Example 1 Save set bundling for a one-week scheduled backup cycle of a tablespace

A one-week scheduled backup cycle of a tablespace includes a level 0 backup of the tablespace on Sunday and a level 1 backup every other day of the week. The Sunday backup creates the save set bundle for the cycle. Each level 1 backup adds save sets into the same bundle. The complete bundle contains the save sets from the seven daily backups of the tablespace. The next backup cycle creates a new bundle during the following week.

Example 2 Save set bundle join

This example illustrates a scenario where NMDA combines existing bundles into a new save set bundle.

Separate level 0 backups of files A and B create two save set bundles. Then a level 1 backup of both files A and B creates a backup set. As the new backup set depends on both of the preceding level 0 backups, NMDA combines all three backups into the same save set bundle.

Example 3 Splitting a save set bundle across volumes

In this example, a save set bundle is split across multiple volumes. A level 0 backup of file A is performed to volume A. An incremental backup of file A is then performed to volume B. Although both backups are recorded as belonging to the same save set bundle, the save set bundle is split across volumes. During staging, only the save sets on the same volume can be staged together.

Save set bundling information in the media database

NMDA stores information about each save set bundle in the NetWorker media database.

Query the media database by using the NetWorker command, **mminfo**, with the appropriate options:

- The **mminfo -r** command can display the name of the bundle associated with a save set. For example, the following command displays a list of all save sets and their bundles:

  ```
  mminfo -a -r "ssid,ssbundle"
  ```

- The **mminfo -q** command can display all the save sets in a specific bundle. For example, the following command displays all the save sets in the bundle named 12983479182:

  ```
  mminfo -a -q "ssbundle=12983479182"
  ```

The *EMC NetWorker Command Reference Guide* and the UNIX man pages describe the **mminfo** command and its available options.

Policy uniformity

If you enable policy uniformity, NMDA automatically enforces the uniformity of browse and retention policies between all the dependent save sets in a scheduled backup cycle, whether or not you have enabled save set bundling.

After NMDA performs an incremental Oracle scheduled backup, if the policies of save sets in the backup are longer than the policies of preceding dependent save sets in the same backup cycle, NMDA changes the policies of all save sets in the cycle to match the longest policy of the new incremental save sets. NMDA modifies the policies that are recorded in the NetWorker media database. As a result, backups cannot expire and become recyclable before other dependent backups from the same backup cycle.

Note: NMDA does not support policy uniformity for nonsnapshot manual backups and snapshot backups. NMDA supports policy uniformity for nonsnapshot scheduled backups only.

Policy uniformity does not depend on whether save sets are on separate volumes. For example, if parts of a save set bundle are on separate volumes, all the save sets in the bundle still receive the same browse and retention policies.

"Configuring policy uniformity for scheduled Oracle backups" on page 136 describes how to configure policy uniformity for NMDA backups.

Other Oracle features

NMDA supports other Oracle RMAN features, for example:

- Fast incremental backups that use change tracking files
- Management of backup duration and throttling
- Backups and restores of data that resides on Oracle ASM
- Backup of the fast recovery area
- Data Recovery Advisor as described in "Data Recovery Advisor" on page 58
- Archival backup through the RMAN **backup...keep** command as described in "Archival backup feature" on page 59
- Improved archived redo log management through the **configure archivelog deletion policy** command
- Improved block media recovery when the **recover...block** command replaces the **blockrecover** command
- Improved integration and block change tracking support in Data Guard
- Backup of read-only transportable tablespaces
- Oracle Enterprise Manager enhancements with new interfaces for the Data Recovery Advisor

The Oracle Recovery Manager documentation describes the RMAN features.

When you use Oracle RMAN features with NMDA, keep in mind the following considerations:

- A fast recovery area stores and manages files related to the recovery of a particular database. You must complete the following steps to back up the RMAN disk backups, the control file autobackups, and the archived redo logs from the fast recovery area to NetWorker volumes:

  a. Allocate or configure one or more channels with the **sbt_tape** device type.

  b. Back up the files with one of the following RMAN commands:

  ```
  backup recovery area
  backup recovery files
  ```

Note: Whether or not you enable a fast recovery area, you can use the **backup recovery files** command to perform the backup.

For example, you can use the following sequence of RMAN commands to configure an automatic channel for NMDA and back up the files from the fast recovery area:

```
configure default device type to 'sbt_tape';
configure channel device type 'sbt_tape' send
'NSR_ENV=(NSR_SERVER=server1)';
backup recovery files;
```

- NMDA supports channel backup failover and backup piece restore failover. If you use multiple channels for an RMAN **backup** command and one of the channels fails, Oracle fails over to another channel to continue the backup job.

  For example, if you configure two channels with different NetWorker volume pools and one of the channels fails over to the other channel during a backup, the entire backup goes to the volumes in the pool of that remaining channel.

- Before you use the **backup** command with the **duration...minimize load** option, consider the following possible results:

  - The **minimize load** option might impact the tape streaming because the transfer rate of data that RMAN sends might be slow with this option, depending on the **duration** value.

    Note: The **minimize load** option does not impact the NetWorker backup-to-disk feature.

  - The **minimize load** option might cause timeout of a scheduled backup if RMAN does not send data to NMDA within the time specified in the Inactivity Timeout field of the NetWorker Group resource.

Oracle 11gR2 features

The following are examples of Oracle features that NMDA supports with Oracle 11gR2 and later. NMDA supports the new features to the extent that RMAN supports the features:

- Enhanced **duplicate** command that can duplicate a database without connecting to a target database by using the NMDA backups of the target database (connections to a catalog and an auxiliary database are required)
- Tablespace Point-in-Time Recovery (TSPITR) enhancement that recovers a dropped tablespace and recovers to a point-in-time before you bring the tablespace online
- Advanced Compression Option

  Note: Do not enable both Oracle and NetWorker compression for NMDA Oracle backups.

- Oracle Grid infrastructure for either a stand-alone database or a RAC

- Oracle ASM Dynamic Volume Manager (Oracle ADVM), a new feature of Oracle ASM that provides volume management services and a standard disk device driver interface to clients
- Policy-managed RAC databases

The appropriate Oracle documentation describes the Oracle 11g features.

Oracle 12cR1 features

The following are examples of Oracle features that NMDA supports with Oracle 12cR1. NMDA supports the new features to the extent that RMAN supports the features:

- Backup and recovery of container databases (CDBs) and pluggable databases (PDBs)
- Table recovery or table partition recovery from NMDA backups
- SYSBACKUP privilege
- Flex ASM
- Oracle Automatic Storage Management Cluster File System (ACFS) extension, such as database files on ACFS

The appropriate Oracle documentation describes the Oracle 12c features.

The *EMC NetWorker Module for Databases and Applications Release Notes* provides details about any limitations in the NMDA support of Oracle 12c features.

Data Recovery Advisor

The Oracle Data Recovery Advisor is a new tool in Oracle 11g. Integrated with RMAN and Oracle Enterprise Manager (OEM), the tool enables a database administrator (DBA) to diagnose and repair database failures.

Before you can use the Data Recovery Advisor to run an RMAN restore script that causes NMDA to repair a database failure, you must configure automatic channels to specify parameters such as NSR_SERVER and NSR_CLIENT and the restore parallelism.

Complete the required configurations to enable the use of Data Recovery Advisor with NMDA:

- If you have not configured automatic channels for NMDA, use the following commands to ensure the basic automatic channel configuration:
  - With Oracle version 11gR1:

    ```
    configure channel device type 'sbt_tape' parms
    'ENV=(NSR_SERVER=NetWorker_server_name,
    NSR_CLIENT=NMDA_client_name)';
    ```

  - With Oracle version 11gR2 or later:

    ```
    configure channel device type 'sbt_tape' parms
    'SBT_PARMS=(NSR_SERVER=NetWorker_server_name,
    NSR_CLIENT=NMDA_client_name)';

    configure channel device type 'sbt_tape' parallelism
    number_of_restore_channels;
    ```

- If you have already configured automatic channels for NMDA, no additional configuration steps are required.

Archival backup feature

The RMAN **backup...keep forever** command enables the creation of an archival backup that is exempt from Oracle backup retention policies, but not automatically exempt from NetWorker retention policies. The archival backup is all-inclusive because it backs up every file required to restore a database to a single disk or to a tape location.

To enable the use of the RMAN **backup...keep forever** command with NMDA:

1. Configure an Archive Pool resource through the NetWorker server.
2. Specify that the backup data must go to the Archive pool by performing one of the following actions:
   - Set the pool selection criteria on the NetWorker server.
   - Set the NSR_DATA_VOLUME_POOL parameter in the RMAN backup script.
3. Set the parameter value NSR_SAVESET_RETENTION=forever in the RMAN backup script.

Note: You must not set the NSR_ORACLE_RETENTION parameter in the RMAN backup script used with NMDA.

The *EMC NetWorker Administration Guide* describes how to configure resources and specify pool selection criteria in the NetWorker server.

NMDA features specific to Sybase

NMDA supports specific features of Sybase ASE backup and restore operations.

Sybase full and incremental backups

NMDA supports full and incremental backups of Sybase data:

- A full Sybase backup sends the **dump database** command to Sybase, which backs up the entire database, including both the data and the transaction log. If incremental backups are not supported for a database, Sybase truncates the inactive portion of the transaction log.

  Note: You can send the Sybase data and logs to separate backup pools during a single NMDA backup.

- An incremental Sybase backup sends the **dump transaction** command to Sybase, which backs up the transaction log and truncates the inactive portion of the transaction log.

Under certain conditions, the Sybase software does not support an incremental backup, such as the following conditions:

- The database and the transaction logs are on the same device.
- The select into/bulk copy option is selected and the database contains unlogged data.

- The truncate log on checkpoint option is selected.
- The database is an in-memory or relaxed durability database.

The Sybase documentation describes all the conditions that do not support an incremental backup.

When an incremental Sybase backup is not supported, the backup result depends on whether you perform a whole instance backup:

- If you perform a whole instance incremental backup, then NMDA skips any database that does not support incremental backups (for example, when the database data and logs are on the same device), regardless of the NSR_PROMOTE_FULL setting.
- If you perform an incremental backup of one or more databases (but not the whole instance), then the Sybase software promotes the backup of the databases that do not support incremental backups (for example, when the database data and logs are on the same device) to a full backup unless the NSR_PROMOTE_FULL parameter value is FALSE. After the backup, the Sybase software truncates the inactive portion of the database log.

Note: If you have any database that does not support incremental backups, perform full backups of the database separately.

Sybase up-to-the-minute recovery

After a database failure, you typically want to recover your database to the latest time possible before the failure. Sybase refers to this type of recovery as up-to-the-minute recovery.

Up-to-the-minute recovery retrieves the intact logs from a partially corrupted or lost database when the failure happens only on the data device, and uses the logs to recover the database up to the time of the failure.

Note: Other sections of this guide refer to this type of recovery as a recovery to the current time.

You can use up-to-the-minute recovery only if you can run the **dump transaction** command on the database and the failure happened on the data device only.

During an up-to-the-minute recovery, NMDA performs a logtail backup of the logs still on the application host by using the **dump tran** command. NMDA then restores the full backup and all the incremental backups, including the logtail backup, and performs the recovery to the current time.

Because this option uses information in the master database to determine the location of the log, the master database must still be accessible.

If the transaction log is not on a separate device, you cannot perform up-to-the-minute recovery because you cannot perform incremental backups. You can only recover to the most recent full backup. This is a Sybase limitation.

Note: Sybase recommends storing the data and transaction logs on separate devices.

If the transaction logs are on separate devices from the data, without using an up-to-the-minute recovery, you can only recover the database to the state corresponding to the latest backup, regardless of whether that backup is a full or incremental backup.

You must meet the following Sybase requirements before you can perform an up-to-the-minute recovery:

- The data and the log are on separate devices.
- The failure is not a failure on the log device.
- A dump of the database (full backup) is already done.
- The master database is still available.
- A minimally logged transaction has not been executed in the database since the last database or log dump. (A minimally logged transaction is a specific Sybase transaction type that does not generate a transaction log record for each affected row.)

The Sybase documentation provides details about Sybase logtail backups and up-to-minute recovery.

Password-protected database backups and restores

NMDA supports password protection of Sybase database backups and restores.

If you specified password protection of the Sybase data at the backup time, you must specify the same password to restore the backed-up data. If the password does not match the one used for the backup, then the restore fails and an error message appears in the log file and on the output terminal.

In a client-side backup configuration, specify the password for the backup by using the NSR_ASE_PASSWORD parameter in the NMDA configuration file, as described in "NSR_ASE_PASSWORD" on page 393.

In a server-side backup configuration, specify the password for the backup on the appropriate screen of the configuration wizard. The wizard stores the password in the NetWorker Lockbox resource.

You can specify the password for the restore by using the **-r** option with the **nsrsybrc** restore command.

Database consistency check

NMDA supports different types of database consistency checks in preparation for a Sybase manual backup.

You must use the **nsrsybcc** command to perform a database consistency check.

"Performing Sybase database consistency checks before backups" on page 154 describes the procedures for Sybase database consistency checks.

Database backup and restore verification

NMDA supports Sybase database backup and restore verification.

NMDA supports verification at the header or a full verification.

"Specifying verification of Sybase database backups" on page 138 and "Specifying verification of Sybase database restores" on page 216 describe how to specify the verification of Sybase backups and restores.

Exclusion of multiple user-defined temporary databases from backup

The Sybase server supports multiple user-defined temporary databases.

In addition to the system-defined temporary database, tempdb, you can create user-defined temporary databases. These temporary databases are attached with a specific user login or a database. Creating temporary databases enhances the performance of a database where many transactions take place at a time. Also, creating temporary databases prevents the critical applications from failing when the system-defined database fails.

NMDA excludes both the user-defined temporary databases and system-defined temporary database during a backup. During the backup, NMDA communicates with the Sybase server to determine whether a database is a user-defined temporary database (to exclude from the backup) or a normal database.

Sybase in-memory and relaxed durability databases

Sybase introduced two new types of databases in ASE 15.5: in-memory databases and relaxed durability databases. NMDA performs full backups and restores of these types of databases in the same way as backups and restores of the regular type of database. NMDA does not support incremental backups of these new types of databases.

NMDA supports the restore of any supported type of database to any other supported type of database as described in "Different database types" on page 219.

Sybase documentation describes the supported types of databases.

Sybase ASE 15.7 features

NMDA supports the features of Sybase ASE release 15.7 except for concurrent Sybase database backups and log backups.

Note: With Sybase 15.7, if you try to run concurrent NMDA Sybase backups, the first backup runs successfully and any other concurrent backups terminate with an appropriate error message.

NMDA supports the other new Sybase 15.7 features, including Sybase database compression, Sybase ASE Cluster Edition systems, and the shrink log feature. The Sybase documentation provides details about the Sybase 15.7 features.

"Configuring Sybase ASE 15.7 backup features" on page 143 describes how to configure Sybase 15.7 features for NMDA operations.

NetWorker User for Sybase

On Windows systems only, you install the NetWorker User for Sybase GUI program (**nwbms.exe**) with NMDA.

The NetWorker User for Sybase program provides a graphical interface for configuring and performing Sybase manual backups and recovery operations.

The following topics describe how to use the NetWorker User for Sybase for manual backups and restores:

- "Performing Sybase manual backups with NetWorker User for Sybase" on page 156
- "Performing Sybase data restores with NetWorker User for Sybase" on page 220

NMDA components

The following table lists the major software components that you install on the NMDA client host during an NMDA installation. Unless you specify otherwise, the files are installed in the same directory as the NetWorker client software.

Table 1 NMDA components (page 1 of 3)

Component name	Description
nmda_*application*.cfg	Located in the directory /nsr/apps/config (UNIX) or *NetWorker_install_path*\apps\config (Windows). NMDA configuration file template for DB2, Informix, Lotus, MySQL, Oracle, or Sybase parameters that apply to client-side configurations only. For MySQL, NMDA provides both a backup configuration file template, nmda_mysql_backup.cfg, and a restore configuration file template, nmda_mysql_restore.cfg. For other applications, NMDA provides a single configuration file template.
nsrdaadmin(.**exe**)	Program that performs one of the following types of conversion: • Conversion of the client-side configuration of a legacy NetWorker module to an NMDA client-side configuration. You create a client-side configuration manually with NMC, without the backup configuration wizard. • Conversion of the server-side configuration of a legacy NetWorker module to an NMDA server-side configuration. You create a server-side configuration with the backup configuration wizard. • Conversion of an NMDA client-side configuration to an NMDA server-side configuration.
nsrdaprobe(.**exe**)	Program that probes for the number or size of generated logs as a condition that triggers a probe-based backup.
nsrdasv(.**exe**)	Main NMDA program that performs one of the following backups: • A scheduled backup of a database server. • A manual backup of a Domino server or Notes client. • A manual backup of a MySQL server. • A manual backup of a Sybase server.
nsr*application***ra**(.**exe**)	Program that performs operations for the NMDA configuration wizard on a DB2, Informix, Lotus, MySQL, Oracle, or Sybase host.
NMDA DB2 components:	
libnsrdb2.so	Located in the directory /usr/lib on UNIX only. NMDA library that interacts with the DB2 backup and restore utilities and the NetWorker server to perform backup, inquiry, and restore processes for DB2 data.
nsrdb2cat(.**exe**)	Catalog synchronization program that prunes (removes) snapshot entries from the DB2 advanced copy services backup history as NetWorker removes expired snapshot entries from its index.

Table 1 NMDA components (page 2 of 3)

Component name	Description
nsrdb2rlog(.**exe**)	Restore utility that copies DB2 transaction logs stored on the NetWorker server to a local file system so that the logs may be used to perform rollforward recovery.
NMDA Informix components:	
libnsrifmx.*xx* libxbsa.dll	Located in the directory /usr/lib on UNIX only. NMDA libraries that interact with the Informix backup and restore utilities and the NetWorker server to perform backup, inquiry, and restore processes for Informix data.
NMDA Lotus components:	
nsrdoclb.dll	On Windows systems only. Library for document-level recovery of Lotus data. The library is only in the 32-bit NMDA Windows install package because the Lotus Notes client or admin client is 32-bit only.
nsrdocrc(.**exe**)	Program for document-level recovery of Lotus data.
nsrlotus_remrecov(.**bat**)	Script that enables NMDA to perform a directed recovery from the NetWorker User for Lotus program on Windows.
nsrnotesrc(.**exe**)	Program for recovery of Lotus database files.
nwbml.dll **nwbml.exe**	On Windows systems only. Library and program for the NetWorker User for Lotus.
NMDA MySQL components:	
libmysqlapi.so.*version*	Located in the directory /usr/lib on Linux only. NMDA API library that interacts with the MySQL library libmysqlclient.so for communicating with the MySQL server.
libmysqlapiwrap.so	Located in the directory /usr/lib on Linux only. NetWorker API library that interacts with the NMDA API library libmysqlapi.so.
libnsrmysql.so **nsrsbtmysql**	The libnsrmysql.so file, located in the directory /usr/lib on Linux only, is the main NMDA library loaded by the MYSQL MEB backup or restore thread. The libnsrmysql.so library starts and uses **nsrsbtmysql** to perform any corresponding NetWorker operations.
nsrmysqlrc	Program for recovery of MySQL data.
NMDA Oracle components:	
libnsrora.*xx*	Located in the directory /usr/lib on UNIX only. Library (known as Media Management Library in Oracle documentation) that is loaded by an Oracle backup or restore process.
nsrnmodrpostcmd(.**bat**)	Sample Oracle postcommand script that can be customized to back up specific files at the end of a scheduled Oracle backup, in preparation for disaster recovery.
nsroraadmin(.**exe**)	Program used to create resource settings in the NWORA resource file for an Oracle backup configuration.
nsroraclecat(.**exe**)	Only available on platforms that support snapshot backups. Program used to remove RMAN catalog entries during automatic catalog synchronization for snapshot backups of Oracle data.
nsrorainfo(.**exe**)	Program that determines the NetWorker volumes required to restore specified Oracle backup pieces from NMDA backups.
nsrsbtcn.exe orasbt.dll	On Windows systems only. The orasbt.dll file is the main NMDA library (known as Media Management Library in Oracle documentation) loaded by the Oracle backup or restore thread. The orasbt.dll library uses **nsrsbtcn.exe** to perform any corresponding NetWorker operations.
NMDA Sybase components:	
libnsrsyb.*xx*	Located in the directory /usr/lib on UNIX only. NMDA library that interacts with the Sybase backup and restore utilities and the NetWorker server to perform backup, inquiry, and restore processes for Sybase data.

Table 1 NMDA components (page 3 of 3)

Component name	Description
nsrsybcc(.exe)	Program for database consistency check of Sybase databases.
nsrsybrc(.exe)	Program for recovery of Sybase databases.
nwbms.dll **nwbms.exe**	On Windows systems only. Library and program for the NetWorker User for Sybase.

NMDA backup and restore processes

Regular (not snapshot) NMDA backups and restores involve the specific process interactions described in the following topics.

Chapter 9, "Snapshot Backups and Restores," describes the processes involved in snapshot backups and restores.

Regular scheduled backup processes

A regular scheduled backup includes the following interactions:

1. At the scheduled backup start time, the main NetWorker server-side service, **nsrd**, starts the configured group's backup by invoking the **savegrp** program.
2. The **savegrp** program requests that the NetWorker client-side service, **nsrexecd**, run the **savefs** program, which sends information back to the **savegrp** program.
3. For each NMDA client in the backup group and each of the client's save sets, the **savegrp** program contacts the **nsrexecd** service to start the **nsrdasv** program. The following steps occur, depending on the type of database backup or application backup:
   - For a DB2, Informix, MySQL, Oracle, or Sybase backup:
     a. The **nsrdasv** program communicates with the appropriate database server or backup server to start a backup session through one of the following methods:
        – DB2 API
        – Informix **onbar**
        – MySQL **mysqlbackup** executable
        – Oracle RMAN executable
        – Sybase Open Client/Server API
     b. Each backup session created by the database server or backup server loads the NMDA shared library for the specific database to perform the backup.
   - For a Lotus backup:
     a. The **nsrdasv** program invokes another **nsrdasv** process called the parent process.
     b. The parent **nsrdasv** process determines the Lotus Domino or Notes files that require backup and starts the child **nsrdasv** processes to back up the files. The number of processes spawned depends on the number of files for backup and the specified parallelism.

4. The NMDA shared library (for DB2, Informix, MySQL, Oracle, or Sybase) or the child **nsrdasv** process (for Lotus) performs the following actions:
   - Contacts the NetWorker server service, **nsrd**, to obtain the required authorization.
   - Sends the backup data to the NetWorker media service, **nsrmmd**, to store on the appropriate backup volumes.
5. The NetWorker online indexes store the tracking information:
   - The **nsrmmd** service records tracking information in the NetWorker media database by using the **nsrmmdbd** service.
   - The backup session sends tracking information to the NetWorker client file index by using the **nsrindexd** service.
6. At the end of a scheduled backup, the **savegrp** program also automatically backs up the NetWorker server bootstrap and the corresponding client file indexes. A manual backup does not automatically back up the bootstrap and client indexes.

Note: A manual backup includes steps 3b to 5 only.

The following figure shows how the database server or application server, NetWorker server, and NMDA processes interact during a regular scheduled backup.

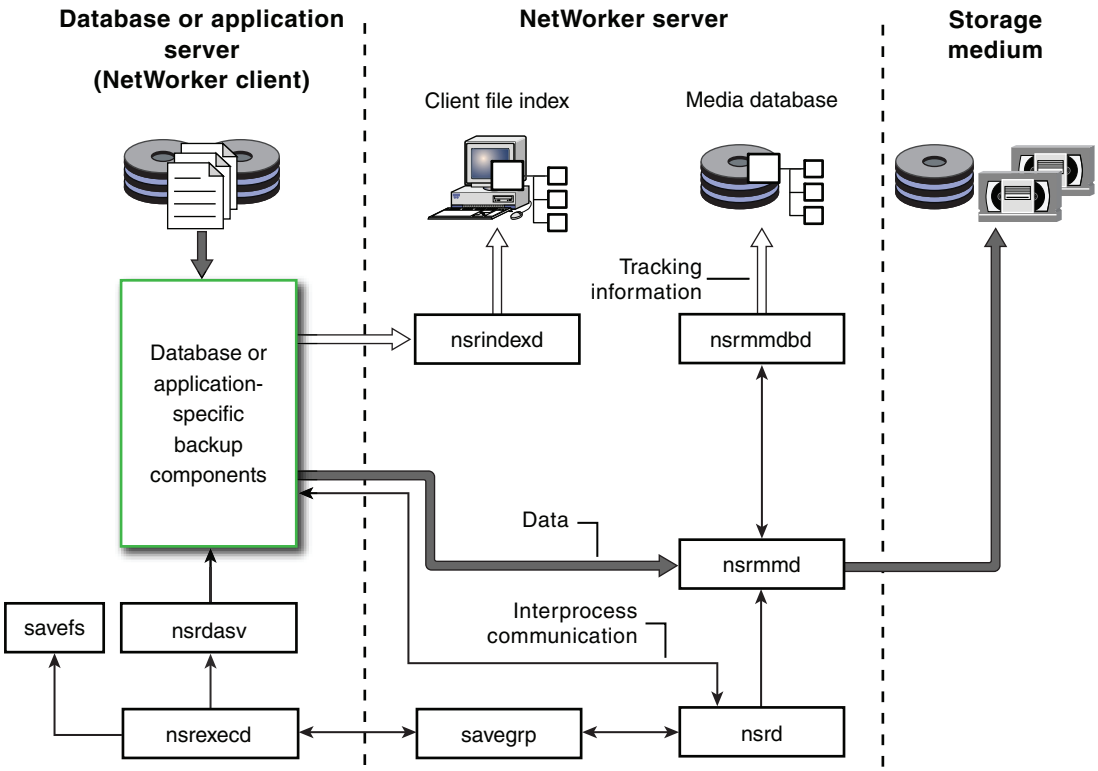

Figure 1 Regular scheduled NMDA backup

Regular restore processes

A regular NMDA restore includes specific process interactions, depending on the type of database restore or application restore.

For a DB2, Informix, MySQL, Oracle, or Sybase restore

1. A user starts the restore by running the proper database or NMDA restore utility through one of the following methods:
   - DB2 command interface
   - Informix **onbar** command
   - NMDA **nsrmysqlrc** command
   - Oracle RMAN interface
   - NMDA **nsrsybrc** command
2. The restore session loads the NMDA shared library for the proper database, which performs the following actions:
   a. Translates the object names requested for restore into a format that the NetWorker server understands.
   b. Forwards the restore object names to the NetWorker service, **nsrindexd**, which verifies that the objects exist in the client file index.
3. The restore session works with the NetWorker server services to mount the volumes required for the restore and read the data from the volumes.
4. The restore session passes the data to the database server or backup server, which writes the data to the disk.

For a Lotus restore

1. A user starts the restore by running the NMDA **nsrnotesrc** command.
2. The **nsrnotesrc** process performs the following actions:
   a. Queries the NetWorker server to obtain a list of Domino files to recover, based on the specified options.
   b. Spawns child **nsrnotesrc** processes to recover Domino data. The number of processes spawned depends on the number of Domino files for recovery and the specified parallelism.
   c. When the child processes finish the recovery successfully, restores the Domino logs, if requested by the Domino server.
3. Each child **nsrnotesrc** process performs the following actions for the subset of files that it restores:
   a. Works with the NetWorker server services to mount the volumes required for the restore and read the data from the volumes.
   b. Writes the data to the disk.

CHAPTER 2
Backup Configuration

This chapter includes the following topics:

- Configuring NMDA backups 70
- Configuring internationalization (I18N) support 71
- Verifying the database or application server configuration 75
- Verifying the basic NetWorker resource configurations 75
- Converting client-side to server-side configurations 79
- Configuring scheduled backups with the wizard 82
- Configuring scheduled backups without the wizard 85
- Configuring manual backups 94
- Configuring deduplication backups with Avamar 95
- Configuring Data Domain backups 99
- Configuring backups to AFTD storage 104
- Configuring probe-based backups 105
- Configuring parallel backups 112
- DB2 backup considerations 116
- Informix backup considerations 117
- Lotus backup considerations 118
- MySQL backup considerations 125
- Oracle backup considerations 128
- Sybase backup considerations 137

Configuring NMDA backups

NOTICE

Before configuring NMDA backups, determine the files to back up to prepare the system for disaster recovery. Chapter 5, "Disaster Recovery," provides details.

To enable NMDA operations on UNIX systems, ensure that the /nsr/apps and /nsr/apps/tmp directories have the drwxrwxrwt access permissions.

You must use the following steps to configure manual backups, scheduled backups, and deduplication backups. Subsequent chapters provide the additional configuration steps for specific environments and technologies.

Before you perform any configuration, backup, or restore operation, ensure that the NetWorker client service, **nsrexecd**, is running on the application host.

To configure backups of Oracle disk backups, follow the instructions in "Configuring backups of DBA disk backups" on page 284.

1. If you have an internationalization (I18N) environment, configure I18N support according to "Configuring internationalization (I18N) support" on page 71.
2. Verify the configuration of the database server or application server according to "Verifying the database or application server configuration" on page 75.
3. Verify the basic NetWorker resource configurations according to "Verifying the basic NetWorker resource configurations" on page 75.
4. If you have upgraded from a legacy NetWorker module to NMDA and want to use the existing scheduled backup configuration, convert the configuration to the NMDA configuration. You can convert the configuration by using either the **nsrdaadmin** command or the manual conversion methods described in the *EMC NetWorker Module for Databases and Applications Installation Guide*.
5. Configure the backup by using the appropriate method:
   - "Configuring scheduled backups with the wizard" on page 82 (server-side configuration)
   - "Configuring scheduled backups without the wizard" on page 85 (client-side configuration)
   - "Configuring manual backups" on page 94

   Note: A backup configuration created with the wizard is a *server-side* configuration, which stores specific configuration parameters in an attribute in the NetWorker Client resource on the NetWorker server host.
   A backup configuration created without the wizard is a *client-side* configuration, which stores the configuration parameters in a configuration file on the client host.
   With all configuration methods, the NetWorker server always stores certain configuration attributes, including the client, group, device, and other attributes.

   If you want to use the wizard to modify a client-side configuration, then you must convert the configuration to a server-side configuration according to the instructions in "Converting client-side to server-side configurations" on page 79.

6. Complete any additional procedures for a specific type of configuration:
   - "Configuring deduplication backups with Avamar" on page 95
   - "Configuring Data Domain backups" on page 99
   - "Configuring backups to AFTD storage" on page 104
   - "Configuring probe-based backups" on page 105
   - "Configuring parallel backups" on page 112
7. Ensure that you meet any application-specific requirements:
   - "DB2 backup considerations" on page 116
   - "Informix backup considerations" on page 117
   - "Lotus backup considerations" on page 118
   - "MySQL backup considerations" on page 125
   - "Oracle backup considerations" on page 128
   - "Sybase backup considerations" on page 137

Configuring internationalization (I18N) support

NMDA documentation refers to internationalization as the capability of NMDA to process non-ASCII data in a non-English locale.

The extent of the NMDA I18N support depends on the following support:

- I18N support that the operating system provides on the NMDA client host.
- I18N support that the NetWorker client and server provide.
- National Language Support (NLS) or globalization support that the database or application provides.

For example, if NetWorker software does not support non-ASCII data in a specific NetWorker resource attribute, NMDA cannot support non-ASCII data in that attribute or the corresponding configuration parameter. NetWorker documentation provides more details about I18N support.

NMDA supports non-ASCII data in the following items:

- Database or application object names and parameters to the extent supported by the database vendor or application vendor, for example, database names, tablespace names, and datafile paths.
- Command line options to database commands, such as the **db2** command, Informix **onbar** command, and other commands, to the extent supported by the database vendor.
- Command line options to NMDA commands, such as **nsrdaadmin**, **nsrdasv**, **nsrdb2rlog**, **nsrmysqlrc**, **nsrnotesrc**, **nsroraadmin**, **nsrorainfo**, **nsroraclecat**, **nsrsybcc**, and **nsrsybrc**.

Note: When you use the **nsrdaadmin** command to encrypt a password for a scheduled backup for DB2 or Sybase, you must use ASCII characters only in the password or the operation fails.

- REMOTE_RECOVCMD environment variable used for Lotus directed recovery, as described in "Configuring Lotus directed recovery" on page 189.
- Values for the parameters listed in the following table.

  Appendix A, "NMDA Parameters and Configuration File," describes the parameters.

Table 2 NMDA parameters that support non-ASCII values

Database or application	Parameters that support non-ASCII values
All databases and applications	• NSR_DIAGNOSTIC_DEST • PRECMD • POSTCMD
DB2	• NSR_DR_FILE_LIST
Informix	• INFORMIXSQLHOSTS (Windows only) • NSR_DR_FILE_LIST • ONCONFIG (Windows only)
Lotus	• Notes_ExecDirectory • NSR_BACKUP_PATHS • NSR_CATALOGFILE • NSR_EXCLUDE_FILE • NSR_LOG_DIR • NSR_LOTUS_DATA_DIR • NSR_RELOCATION_DEST • NSR_RESOURCE_DIR • PATH
MySQL	• MYSQL_BACKUP_DIR • MYSQL_BACKUP_NAME • MYSQL_CFG_FILE • MYSQL_DATABASES • MYSQL_DATADIR • MYSQL_INCLUDE • MYSQL_INCR_DIR • MYSQL_MEB_PATH • MYSQL_SBT_LIB_PATH • NSR_DR_FILE_LIST
Oracle	• NSR_ORACLECAT_LOG_FILE • NSR_RMAN_ARGUMENTS
Sybase	• NSR_ASE_PASSWORD • NSR_BACKUP_PATHS • NSR_EXCLUDE_FILE

NMDA generates error messages in the nmda_*app*.messages.raw file in a language-independent form, readable by the **nsr_render_log** program only. The log file does not contain database server errors or application server errors.

The *EMC NetWorker Administration Guide* describes how to use the **nsr_render_log** program to read any language-independent file, such as nmda_*app*.messages.raw.

I18N support requirements

To configure I18N support, you must meet the following I18N requirements:

- The NMDA client host includes a supported internationalized version of the operating system, configured to operate in the non-English locale.
- The database or application provides the required National Language Support (NLS) or globalization support. You must configure the database with the required non-ASCII character set.
- The appropriate host contains the required NetWorker software:
  - The NMDA client or a remote host contains internationalized NetWorker server software.
  - If a remote host contains the NetWorker server, then the NMDA client contains internationalized NetWorker client software or storage node software.

  The *EMC NetWorker Software Compatibility Guide* identifies the different supported languages, the operating system requirements, and the NetWorker version requirements for I18N support.
- For I18N support during snapshot operations, a supported release of the NetWorker client (including the NSM feature) is installed and configured.

Configuring I18N support for Informix operations

The default locale for Informix IDS is English (en_us).

To configure I18N support for Informix operations in a non-English locale, set the following parameters in the NMDA configuration file for scheduled backups and as environment variables for manual backups and restores:

- DB_LOCALE
- DBLANG
- SERVER_LOCALE
- CLIENT_LOCALE

The Informix documentation describes these parameters.

Configuring I18N support for Oracle operations

For Oracle backups, NMDA supports non-ASCII data in the following items:

- The **format** string of an RMAN **backup** command unless you use the **nsrdaadmin** command to convert a legacy backup configuration on Windows
- The **tag** string of an RMAN **backup** command
- Usernames in the connection strings to the target database and the recovery catalog

Note: Due to Oracle limitations, you must use ASCII text for the following items:
- ORACLE_HOME path
- Net service name of the Oracle target database, recovery catalog, or duplicate database
- The password of the database user with DBA privileges

Do not use non-ASCII text in the Oracle database usernames. Support of non-ASCII values for ORACLE_SID and TNS_ADMIN depends on the Oracle software.

To configure I18N support for Oracle operations, complete these steps:

1. Set the environment variable NLS_LANG to the character set supported by the operating system and Oracle database, and then restart the Oracle Server.

   The Oracle Globalization Support documentation describes the NLS_LANG variable.

   For example, to ensure that Oracle operations return Japanese text in a Japanese locale, set NLS_LANG as follows:

   ```
   export NLS_LANG=JAPANESE_JAPAN.JA16EUC
   % lsnrctl stop
   % lsnrctl start
   % sqlplus /nolog
   SQL*Plus: Release 10.1.0.2.0 - Production on Thu Apr 26 15:12:03
   Copyright (c) 1982, 2004, Oracle.  All rights reserved.
   SQL> connect sys/oracle as sysdba;
   SQL> shutdown;
   SQL> startup;
   SQL> quit;
   ```

2. For scheduled backups configured without the wizard, set the NLS_LANG parameter in the configuration file to the same value as the environment variable NLS_LANG.

   For example, in a Japanese locale, set NLS_LANG in the configuration file as follows:

   ```
   NLS_LANG=JAPANESE_JAPAN.JA16EUC
   ```

   Note: If you configure the scheduled backup with the configuration wizard, you can set NLS_LANG on a wizard screen. If you set NLS_LANG in the NWORA resource file, the wizard autopopulates the NLS_LANG field.

Configuring I18N support for Sybase operations

To configure I18N support for Sybase operations, follow the appropriate procedure for the particular type of Sybase backup:

- For an NMDA Sybase client-side scheduled backup, if the locale in which the Sybase ASE server runs is different than the locale in which the NetWorker client (**nsrexecd**) runs, then set the NSR_LOCALE value in the NMDA Sybase configuration file to the locale of the Sybase ASE server.
- For an NMDA Sybase manual backup, if the execution environment of the **nsrdasv** program has a different locale than that of the Sybase ASE server, then set the NSR_LOCALE value in the NMDA Sybase configuration file to the locale of the Sybase ASE server.

Verifying the database or application server configuration

Ensure that the following components are installed and configured according to the DB2, Informix, Lotus, MySQL, Oracle, or Sybase server documentation:

- Database server or application server software
- Required target databases
- Networking software
- Any other required components

Verifying the basic NetWorker resource configurations

Ensure that the basic NetWorker resources are configured on the NetWorker server to enable backup and restore operations. The basic resources include the NetWorker Server resource, User Group resource, Client resource, Device resource, and the Pool and Label Template resources.

"Configuring the NetWorker resources with NMC" on page 86 describes the NetWorker resources required for scheduled backups.

The *EMC NetWorker Administration Guide* and the NetWorker server online help describe how to configure the NetWorker resources and attributes.

NetWorker Server resource

After you install the NetWorker server, the NetWorker configuration includes a preconfigured Server resource with attribute settings that influence the performance and security of backups.

The following table describes the main NetWorker Server resource attributes. Ensure that the attribute settings in the Server resource are valid for the NMDA backup environment. Modify the settings as required.

Table 3 NetWorker Server resource attributes

Attribute	Description
Name	Specifies the hostname of the NetWorker server.
Parallelism	Specifies the maximum number of backup save streams that the NetWorker software allows to arrive concurrently at the server. The NetWorker server edition determines the maximum parallelism value. The simultaneous backup of multiple data streams increases the efficiency of the storage devices.
Datazone pass phrase	Specifies an optional key or optional pass phrase, used only if you specify AES encryption for a backup. You need the pass phrase to restore encrypted data from the backup. "NSR_AES_ENCRYPTION" on page 358 provides details.

NetWorker User Group resource

The NetWorker server includes an access control feature that enables NetWorker administrators to assign users to NetWorker user groups. Each user group has a specific set of privileges, defined in the Privileges attribute of the User Group resource.

The NetWorker server installation includes preconfigured user groups. Certain groups are for administrative purposes; other groups are for users. You can create additional user groups, if required.

By default, the NetWorker server assigns the following privileges to a user in the preconfigured Users group:

- Backup Local Data
- Monitor NetWorker
- Recover Local Data

These default user group configurations are sufficient for most NMDA operations. If you do not use the default user group configurations, or you use restricted data zones, you must ensure that the user has the required privileges. When you use restricted data zones, you can also grant the privileges by using the Users and Privileges fields of the Restricted Data Zone resource. The *EMC NetWorker Administration Guide* provides details.

Different operations require specific privileges, as described in the following table.

Note: In a cluster system, grant privileges to the users as listed in the following table on all the physical hosts in the cluster.
NSM snapshot backups and restores of DB2 data or Oracle data require the same privileges as nonsnapshot backups and restores, plus the privileges that NSM requires. The NSM documentation describes the required privileges.

Table 4 User group privileges for NMDA operations (page 1 of 2)

Operation	Operating system user that performs operation	Required user group privileges
Manual backup	Database-specific user on the database server or application server as described in "Database-specific user definition" on page 77	Backup Local Data Monitor NetWorker
Scheduled backup	Database-specific user on the database server or application server as described in "Database-specific user definition" on page 77	Backup Local Data Monitor NetWorker
	Root user (UNIX) or System user (Windows) on the database server or application server	Backup Local Data Monitor NetWorker
Restore	Database-specific user on the database server or application server as described in "Database-specific user definition" on page 77	Recover Local Data Monitor NetWorker
Backup deletion	Database-specific user on the database server or application server as described in "Database-specific user definition" on page 77	**With NetWorker 7.*x*:** Operate NetWorker and any prerequisite privileges **With NetWorker 8.0 or later:** Change Application Settings and any prerequisite privileges
Configuration of NetWorker resources for backup or restore	NMC user	Configure NetWorker and any prerequisite privileges

Table 4 User group privileges for NMDA operations (page 2 of 2)

Operation	Operating system user that performs operation	Required user group privileges
Conversion of scheduled backup configuration with the **nsrdaadmin** command	Root user or a member of the Microsoft Windows Administrators group that runs the **nsrdaadmin** command on the database server or application server	Configure NetWorker and any prerequisite privileges
Oracle operations only:		
RMAN crosscheck	Database user on the Oracle Server	Recover Local Data Monitor NetWorker
Save set bundling or policy uniformity	Root user (UNIX) or System user (Windows) on the Oracle Server	**With NetWorker 7.*x*:** Operate NetWorker and any prerequisite privileges **With NetWorker 8.0 or later:** Change Application Settings and any prerequisite privileges

Database-specific user definition

Table 4 on page 76 lists the user group privileges required for specific NMDA operations and refers to the database-specific user on the database server or application server.

The following table provides the definition of the database-specific user for each supported database or application.

Table 5 Database-specific user for each database or application

Database or application	Definition of database-specific user
DB2	• On UNIX: OS user that runs the DB2 instance. • On Windows: User that runs the DB2 Windows services, which is typically the System user by default.
Informix	OS user that runs the **onbar** or **onsmsync** command.
Lotus	Domino or Notes OS user.
MySQL	MySQL OS user.
Oracle	• On UNIX: - If you use the Net service: OS user that starts the Net service. - If you do not use the Net service: OS user that runs the Oracle instance. **Note:** If Oracle ASM is used, the database-specific user also includes the OS user that runs the ASM instance. • On Windows: User that runs the Oracle Windows services.
Sybase	• On UNIX: OS user that runs the Sybase server. • On Windows: User that runs the Sybase Windows services, which is typically the System user by default.

Requirements for backup deletion operations

NMDA tries to delete a backup entry from the NetWorker index in the following cases:

- DB2—When you use the **db2 prune** command or **db2acsutil delete** command
- Informix—When you use the **onsmsync** command
- Oracle—When you use the RMAN **delete** command or you cancel a running Oracle backup

To enable NMDA to delete a backup from the NetWorker index, ensure that the database-specific user has the required NetWorker privileges for backup deletion as listed in Table 4 on page 76.

Note: If the Oracle user does not have the required NetWorker privileges, NMDA fails to remove the backup save set entries from the NetWorker index. However, RMAN might remove the corresponding entries from the RMAN catalog, which would leave the NetWorker index and RMAN catalog unsynchronized. To resynchronize the NetWorker index and RMAN catalog, run the appropriate NetWorker media management command to manually remove the inconsistent save set entries from the NetWorker index.

NetWorker Client resource

If you install the NetWorker server on the NMDA client host, the installation automatically creates a basic Client resource for the client.

Before you can use NMDA for backups or restores, a NetWorker Client resource must exist for the client host:

- For manual backups, a basic Client resource must exist.
- For scheduled backups, a customized Client resource must exist.

For both manual backups and scheduled backups, set the Parallelism attribute in the Client resource to the maximum number of data streams that the NMDA client sends in parallel to the NetWorker server.

NetWorker Device resource

The NetWorker server or NetWorker storage node uses a supported tape or a disk storage device to write data or read data during a backup or restore.

The NetWorker server configuration must include a Device resource for each physical storage device used for backups and restores. In addition, each storage device must contain a labelled and mounted volume.

The *EMC NetWorker Hardware Compatibility Guide* provides a complete list of the storage devices that the NetWorker server supports.

NetWorker Pool and Label Template resources

A pool is a specific collection of backup volumes that the NetWorker server uses to store, sort, and organize backup data. Each NetWorker volume belongs to either a preconfigured pool or a user-created pool.

Leave the Save Sets attribute blank in the Pool resource. NMDA sets the Save Sets during each backup session.

You can specify the backup group in the Pool resource so that a scheduled backup automatically uses a volume pool. The scheduled backup uses that pool unless you set the NSR_DATA_VOLUME_POOL parameter. The parameter setting takes precedence over any pool associated with the scheduled backup group.

With NetWorker server 7.6.1 or later, you can also associate the Client resource with a specific pool for a scheduled backup. This selected pool takes precedence over a pool associated with a group or the NSR_DATA_VOLUME_POOL setting.

Each pool has a specific label template associated with it. The label provides an automated method to identify the media assigned to a pool. NetWorker software uses the volume pools and the label templates to track the data location on each volume.

Firewall support

NMDA supports a firewall. The ports that NMDA uses for the firewall depend on the corresponding ports configured for the NetWorker server.

To configure the firewall that NMDA uses, follow the firewall configuration instructions in the *EMC NetWorker Administration Guide* for the particular NetWorker server system.

Converting client-side to server-side configurations

Before you can use the wizard to modify a scheduled backup configuration created without the wizard (client-side configuration), you must convert the configuration by using the **nsrdaadmin -W** command.

The **nsrdaadmin -W** command converts the scheduled backup configuration to the configuration storage framework that the wizard supports, which is a server-side configuration. The conversion does not create a Client resource, but modifies an existing Client resource so that you can then use the wizard to modify the resource. After the conversion, you must use the wizard only to modify the configuration.

Note: NMDA does not support the conversion of a MySQL client-side configuration to a server-side configuration.

To convert the client-side configuration of a legacy NetWorker module to an NMDA server-side configuration, you must first use the **nsrdaadmin -M** command. This command converts the legacy configuration to an NMDA client-side configuration, as described in the *EMC NetWorker Module for Databases and Applications Installation Guide*. Then you can use the **nsrdaadmin -W** command.

An administrative user must run the **nsrdaadmin** command on the NMDA host. The user can be the root user on UNIX or a member of the Windows Administrators group. The administrative user must have the Configure NetWorker user group privilege, as described Table 4 on page 76.

The **nsrdaadmin -W** command can convert only the following configurations:

- A correct (working) NMDA client-side configuration.

  If the configuration is not correct, then the conversion might fail or might generate an invalid wizard configuration.

- A configuration that includes configuration files or configuration scripts on the host where the **nsrdaadmin** command runs.

  To convert the Client resources for different physical hosts, you must run the **nsrdaadmin** command on each physical host, or create a script to automate the process.

Note: If an existing Oracle RMAN script contains non-ASCII characters on Windows, you cannot use the **nsrdaadmin -W** command to convert the NMDA Oracle client-side to server-side configuration. For this conversion, you must perform one of these tasks:
- Remove the non-ASCII characters from the RMAN script, if possible.
- Create a configuration with the wizard by using the same configuration options and then delete the old configuration.

NetWorker resources modified by the conversion

The configuration conversion with the **nsrdaadmin -W** command updates the NetWorker resources as follows:

- In the Client resource:
  - The conversion changes the Backup Command attribute from:

    nsrdasv [**-T** *app*] [**-c** *client*] **-z** *NMDA_configuration_file*

    to:

    nsrdasv -T *app* [**-c** *client*]

    where:

    – *app* is one of the following values, depending on the database or application: **db2**, **informix**, **lotus**, **oracle**, **sybase**.

    – *client* is the client name. This name is set if the client is a virtual client on a cluster.

  - The conversion creates the Backup Config attribute (hidden attribute) to contain:

    – Information from the NMDA configuration file.

    – For Oracle only, information from the RMAN script specified in the Save Set attribute.

  - The conversion creates the Backup Type attribute to contain the value DB2, Informix, Lotus, Oracle, or Sybase.

- If needed, the conversion updates the NetWorker Lockbox resource with sensitive data from the NMDA configuration file or Oracle RMAN script, such as the database username or database password. The conversion adds the following entries to the ACL of the Lockbox resource:
  - system@*client_hostname* (Windows) or root@*client_hostname* (UNIX)
  - Other ACL entries that the user provides during the conversion

Using the nsrdaadmin conversion command

As an administrative user on the NMDA host, type the following **nsrdaadmin** command to convert a client-side configuration to a server-side configuration:

```
nsrdaadmin -W -s server_name [-c client_name] [-g group_name]
   [-N save_set_name] [-Y]
```

Command options and settings in brackets ([ ]) are optional. Do not include the brackets when typing the command.

The following table describes the **nsrdaadmin** command options. Example 4 on page 81 provides an example of the command usage.

Table 6 Options of the nsrdaadmin command for configuration conversion

Option	Description
-W	Mandatory. Specifies the conversion from an NMDA client-side to wizard (server-side) scheduled backup configuration.
-c *client_name*	Optional. Specifies the hostname of the NetWorker client to be converted. The default value is the hostname of the local physical client. In a cluster, you typically set this option to a virtual cluster client. **Note:** Do not specify a hostname that is different from the name of the host where you run the **nsrdaadmin** command.
-g *group_name*	Optional. Specifies the conversion of only the configurations of the client (either the default or set through the **-c** option) that is in the *group_name* group.
-N *save_set_name*	Optional. Specifies the conversion of only the configurations of the client (either the default or set through the **-c** option) that has the *save_set_name* save set.
-s *server_name*	Mandatory. Specifies the hostname of the NetWorker server that backs up the client.
-Y	Optional. Specifies that the conversion proceeds without prompting for confirmation. After the conversion, you must add the usernames to the NetWorker Lockbox resource for all the users that will use the wizard to modify the configuration. The *EMC NetWorker Administration Guide* describes the Lockbox resource. If you do not specify the **-Y** option, the **nsrdaadmin** program asks for user confirmation before proceeding with the conversion.

Example 4 Using the nsrdaadmin command for configuration conversion

You run the following command on the NMDA host, hostA, to convert all the Client resources for hostA to a wizard configuration. The conversion stores the Client resources on the NetWorker server host, hostNWS:

```
nsrdaadmin -W -s hostNWS
```

Configuring scheduled backups with the wizard

NMDA supports the scheduled backup configuration wizard, also known as the Client Backup Configuration wizard in NMC, which is integrated with a supported NMC release.

The backup configuration wizard can perform the following actions:

- Configure a new NetWorker Client resource for a scheduled backup.
- Configure a new NetWorker Group resource, or use an existing resource, for the scheduled backup.
- Configure a new browse policy and new retention policy for the scheduled backup, or use existing policies.

  Note: You must configure other NetWorker resources through NMC without the wizard, as described in "Configuring the NetWorker resources with NMC" on page 86.

- Save a copy of the configuration settings to a specified file on the NMDA host for reference purposes.
- Modify a backup configuration created with the configuration wizard.

The wizard stores the configuration information, except for sensitive data such as passwords, in a hidden attribute named Backup Config in the Client resource.

Note: Do not manually modify the Backup Config attribute in the Client resource. If you use the wizard to create a backup configuration, you must use the wizard to modify the configuration.

The wizard stores sensitive data securely by using NetWorker lockbox services.

The wizard provides options for typical and custom configurations:

- The "typical" workflow enables you to configure a backup with fewer clicks by using default backup settings. The wizard help describes the default settings for a typical scheduled backup.
- The "custom" option provides a more detailed workflow that enables you to:
  - Use an existing NMDA configuration file as initial input to the wizard configuration.
  - Customize backup options, for example, to configure:
    - Number of backup sessions to use
    - Specific data objects to back up
    - Additional backup parameters
  - Save a copy of the configuration settings to a specified file for reference purposes.
- For Oracle backup configurations, the wizard also provides an option to configure a backup of Oracle DBA disk backups. Chapter 8, "Oracle DBA and NetWorker Backup Administrator Collaboration," provides details.

To use the wizard to modify an NMDA configuration not created with the wizard (client-side configuration), you must first convert that configuration according to "Converting client-side to server-side configurations" on page 79.

The *EMC NetWorker Module for Databases and Applications Release Notes* describes the limitations and NetWorker requirements for the wizard. Descriptive inline text and online help in the wizard describe how to use the wizard.

Requirements for using the backup configuration wizard

You must meet specific requirements before you use the backup configuration wizard.

- The NMC user that starts the wizard (the wizard user) has the Configure NetWorker privilege and any prerequisite privileges on the NetWorker server.
- Communication between the NMC server, the NetWorker server, and the NMDA client uses NetWorker **nsrauth** authentication, which is the default NetWorker setting. The NetWorker documentation describes the **nsrauth** authentication requirements.
- You have installed the required NetWorker releases on the NMC server host, the NetWorker server host, and the NMDA client host, as described in the *EMC NetWorker Module for Databases and Applications Release Notes*.

You can run the wizard on a host with no NetWorker software installed and with no direct communication with the NMDA client. You do not need administrator privileges or root user privileges on that system.

Configuring a scheduled backup with the wizard

You must run the NetWorker Management Console (NMC) GUI and select the appropriate options to use the configuration wizard for configuring an NMDA scheduled backup.

Ensure that you have configured the NetWorker Schedule, Device, and Pool resources according to "Configuring the NetWorker resources with NMC" on page 86. You cannot configure the Schedule, Device, and Pool resources with the wizard.

1. In the NMC **Enterprise** window, select the NetWorker server name and double-click the **NetWorker** application to launch the application.
2. In the **Configuration** window under the NetWorker server name, click **Clients**.
3. Start the wizard by using the appropriate method:
   - To create a backup configuration, right-click **Clients** and select **Client Backup Configuration** › **New**.
   - To modify a backup configuration created with the wizard, right-click the NMDA client in the right pane and select **Client Backup Configuration** › **Modify**.
4. On each wizard screen that appears, specify the required options and values for the backup configuration. Click **Next** to go to the next screen.

   Each wizard screen includes an online help button that you can click to access descriptions of the fields and options on the screen.

   You can click a link in the wizard steps panel to go directly to the screen you want to modify.

5. DB2, MySQL, Oracle, or Sybase only—If you use the wizard to modify a configuration and you select to change a database user password used for backup, the wizard enables you to apply that password change to other NMDA backup configurations with the same credentials. The feature facilitates easy password changes across different server-side configurations.

   NMDA applies the password changes according to the type of database:

   - DB2, MySQL only—If you select to change the password for a specific OS username on the client host, NMDA applies the password change to all the backup configurations that use the same OS username on the host.
   - Oracle only—If you use database authentication on the client host and you select to change the password for a specific database (Oracle Net service name) and database username, NMDA applies the password change to all the backup configurations that use the same database and username on the host.

     Note: If you use the same username for different databases on the same host, you can select an option to apply the same password change to all the backup configurations of all the databases with that username. This feature does not apply for OS authentication on the client host, which does not require a password to be managed.

     If you use an Oracle recovery catalog on the client host and you select to change the password for a specific recovery catalog database and username, NMDA applies the password change to all the backup configurations that use the same recovery catalog database and username on the host.

     Note: If you use the same recovery catalog database across different hosts, you can select an option to apply the same password change to all the backup configurations with the recovery catalog database and username across the different hosts.

   - Sybase only—If you select to change the password for a specific Sybase server and username on the client host, NMDA applies the password change to all the backup configurations that use the same Sybase server and username on the host.

     Note: If you use the same username for different Sybase servers on the same host, you can select an option to apply the same password change to all the backup configurations of all the Sybase servers with that username.

6. On the last screen titled **Review and Accept the Client Configuration**, click **Create** to create the configuration or click **Modify** to modify the configuration, and click **Finish** to exit the wizard.

To configure NMDA scheduled backups of Oracle disk backups with the wizard, you must follow the additional instructions in "Configuring backups of DBA disk backups" on page 284.

Configuring scheduled backups without the wizard

As an alternative to configuring a scheduled backup with the wizard (server-side configuration), you can create and modify a scheduled backup configuration without the wizard (client-side configuration).

1. For Oracle backups only, create an RMAN script for the Oracle backups, and set the required parameters in the script. "Oracle backup considerations" on page 128 provides details.
2. Create an NMDA configuration file, including the required parameter settings. "Configuring the NMDA parameters without the wizard" on page 85 provides details.
3. Set the required NetWorker resources directly in the NMC program. "Configuring the NetWorker resources with NMC" on page 86 provides details.

Note: The configuration file pathname in the Backup Command attribute of the Client resource is the pathname of the configuration file from step 2.

Configuring the NMDA parameters without the wizard

For the client-side configuration of a scheduled backup, set the required parameters in the NMDA configuration file. Appendix A, "NMDA Parameters and Configuration File," describes the NMDA configuration file, including syntax, templates, and all the common parameters and application-specific parameters.

For Oracle backups, set only certain parameters, such as ORACLE_HOME, in the configuration file. Set the other parameters in the RMAN script. Store the configuration file and RMAN script on the NMDA client host.

Set the mandatory parameters for a scheduled backup of the database or application, as shown in the following table.

Note: Do not set the parameters NSR_SERVER and NSR_GROUP for a scheduled backup because NMDA automatically passes the server information and group information from the NetWorker server that started the backup to the backup processes. However, you can set NSR_SERVER and NSR_GROUP for a manual backup.

Table 7 Mandatory parameters for scheduled NMDA backups (page 1 of 2)

Database or application	Mandatory parameters for scheduled NMDA backups
DB2	• DB2INSTANCE (UNIX only) • DB2_NODE_NAME • DB2PATH (Windows only) • DB2_TBS_LIST (tablespace backup only) • DB2_USER • INSTHOME (UNIX only) • USER_PSWD **Note:** Set the encrypted DB2 user password in the USER_PSWD parameter with the **nsrdaadmin -P** command as described in "USER_PSWD" on page 370.

Table 7 Mandatory parameters for scheduled NMDA backups (page 2 of 2)

Database or application	Mandatory parameters for scheduled NMDA backups
Informix	• INFORMIXDIR • INFORMIXSQLHOSTS (UNIX) • INFORMIXSQLHOSTS (Windows, with Informix 12.10 or later) • ONCONFIG
Lotus	• LOTUS_USER (UNIX only) • Notes_ExecDirectory • NSR_RESOURCE_DIR (UNIX only) • PATH (UNIX only)
MySQL	• MYSQL_CFG_FILE • MYSQL_DATADIR
Oracle	• ORACLE_HOME • ORACLE_SID (specific cases only) • TNS_ADMIN (specific case only)
Sybase	• LD_LIBRARY_PATH or LD_LIBRARY_PATH_64 or LIBPATH (for specific operating systems) • PATH (specific cases only) • SYBASE • SYBASE_USER • USER_PSWD (when Sybase server has a password) **Note:** Set the encrypted Sybase user password in the USER_PSWD parameter with the **nsrdaadmin -P** command as described in "USER_PSWD" on page 396.

Configuring the NetWorker resources with NMC

1. In the NMC **Enterprise** window, select the NetWorker server name, and double-click the **NetWorker** application to launch it.
2. In the **Configuration** window, under the NetWorker server name, select the appropriate backup resources to configure, such as **Groups**, **Schedules**, **Clients**, and so on. For example:
   - To create a client, right-click **Clients** and select **New**.
   - To modify a client, double-click the client name to open its **Properties** window.

   On each of the properties tabs, specify the appropriate options and values for the backup configuration.

The following topics highlight the resource configuration details for NMDA. The *EMC NetWorker Administration Guide* and NMC online help describe the NetWorker resources.

Configuring the Group resource

Configure the NetWorker Group resource to specify the attributes of the scheduled backup group. The Group resource specifies the NetWorker Client resources that all start to back up data at a specified time after you set the following attributes in the Group resource:

- The Autostart attribute is Enabled.
- The Start Time for the backup is set.

You can create and modify all the NetWorker backup groups. You can delete all the backup groups except the Default group.

1. For a regular scheduled backup, set the Snapshot attribute in the Group resource to False.
2. Set the values of these backup level attributes, if required:
   - Level attribute
   - Force Incremental attribute

   The values of the Level attribute and Force Incremental attribute determine the scheduled backup level for the group.
3. Specify the name of the NetWorker group in the Group attribute of the Client resource.
4. To have a regular scheduled backup automatically use a volume pool associated with the backup group, specify the group name in the Pool resource for the volume pool.

Configuring the Schedule resource

1. Create or select the NetWorker Schedule resource to specify the days of the week when the scheduled NMDA backup runs.

   The NetWorker server provides preconfigured schedules that you can modify, but you cannot delete the schedules. If the preconfigured schedules meet the backup requirements, use the schedules as is. Otherwise, create schedules to accommodate the site needs. You can modify and delete any schedules that you create.
2. Configure the NetWorker Schedule resource for the scheduled backup by selecting the backup level for each day of the week as shown in the following table.
3. Specify the name of the NetWorker schedule in the Schedule attribute of the Client resource.

Table 8 Backup levels specified in NetWorker Schedule resource (page 1 of 2)

Database or application	Supported backup levels
DB2	• Level full—DB2 full backup of all the database data. • Level incr—DB2 incremental backup. • Level 1 to 9—DB2 delta backup. • Level skip—NetWorker server skips the backup on that day. To support the NetWorker backup levels from the Schedule resource: • Set the DB2_APPLY_NW_LEVELS parameter to TRUE in the NMDA configuration file, as described in "DB2_APPLY_NW_LEVELS" on page 367. • Set the TRACKMOD parameter to ON with a DB2 command at the operating system command line: **db2 update db cfg for *sample* using TRACKMOD ON** where *sample* is the name of the database to be backed up.

Table 8 Backup levels specified in NetWorker Schedule resource (page 2 of 2)

Database or application	Supported backup levels
Informix	• Level full—ON-Bar level 0 backup of all the selected dbspaces of the database instance. • Level 1—ON-Bar level 1 backup of the data that has changed since the last level 0 (full) backup. • Level 2 to 9—ON-Bar level 2 backup of the data that has changed since the last level 1 backup. • Level skip—NetWorker server skips the backup on that day.
Lotus	• Level full—Lotus full backup of all the data. • Level incr—Lotus incremental backup of only data that has changed since the last backup. • Level 1—Lotus transaction logs only backup. • Level skip—NetWorker server skips the backup on that day. **Note:** NMDA does not support the NetWorker backup levels 2 to 9 for a Lotus backup. If you specify one of these levels, the backup fails.
MySQL	• Level full—Full backup of all the data. • Level incr—Differential incremental backup of only the data that has changed since the last full or incremental backup. • Level 1—Cumulative incremental backup of only the data that has changed since the last full backup. **Note:** NMDA does not support the NetWorker backup levels 2 to 9 for a MySQL backup. If you specify one of these levels, NMDA uses level 1 instead for the backup.
Oracle	• Level full, incr, or 1 to 9—NetWorker server runs the backup script on that day. The backup level set in the RMAN backup script determines the Oracle backup level. **Note:** Always specify the full backup level on the days when you want to run a backup. This setting prevents unnecessary processing on a NetWorker server that processes other level backups. • Level skip—NetWorker server skips the backup on that day.
Sybase	• Level full—Sybase full backup of all the data. • Level incr—Sybase incremental backup of only the transaction logs generated since the last full or incremental backup. • Level skip—NetWorker server skips the backup on that day. **Note:** NMDA does not support the NetWorker backup levels 1 to 9 for a Sybase backup. If you specify one of these levels, the backup fails. A whole instance incremental backup skips the backup of any database that does not support incremental backups (for example, when the database data and transaction logs are on the same device).

Configuring the Client resource

To customize the Client resource for a scheduled backup, specify the attribute values as shown in the following table.

For an NMDA application that supports multiple entries in the Save Set attribute of the Client resource, consider the following facts:

- NMDA performs the backups for the entries in arbitrary order, possibly in parallel.
- If the NMDA configuration file also contains a setting for PRECMD or POSTCMD, the precommand and postcommand files will be as follows:
  - Common for all the backups
  - Executed once for each backup

If NMDA retries a scheduled backup, NMDA executes the specified precommand and postcommand again for the backup. To include separate preprocessing and postprocessing for each backup, define a separate Client resource for each backup.

Note: In the following table, the information in the NMC tab column is based on NMC 7.6 SP2 and might change in other NMC releases.

Table 9 NetWorker Client resource attributes (page 1 of 5)

Attribute name	NMC tab	Attribute setting
Aliases	Globals	Specify all known aliases for the system that contains the NMDA software.
Backup Command	Apps & Modules	Specify the command to be used for the backup: **nsrdasv**(**.exe**) **-z** *configuration_file* where *configuration_file* is the pathname of the configuration file that contains the NMDA parameter settings for the scheduled backup. On Windows systems only, if the configuration file pathname includes any spaces, perform one of the following actions: • Change the configuration file pathname so it does not include spaces. • Use the Windows short pathname format. For example: **nsrdasv**(**.exe**) **-z C:\Progra~1\Legato\nsr\apps\config\config.txt** • Use quotes and double backslashes in the pathname. For example: **nsrdasv**(**.exe**) **-z "C:\\Program Files\\Legato\\nsr\\apps\\config\\config.txt"**
Browse Policy	General	Specify the length of time that the NetWorker server retains an entry for the backup in the online client file index. The value of the Browse Policy attribute cannot exceed the value of the Retention Policy attribute. **Note:** If you set the parameter NSR_SAVESET_BROWSE as described in Appendix A, "NMDA Parameters and Configuration File," the parameter value overrides the Browse Policy attribute setting in the Client resource.
Checkpoint enabled	General	Select this attribute to enable a checkpoint restart (CPR) backup for Lotus only. NMDA backups ignore this attribute setting for other applications.
Checkpoint granularity	General	Do not set this attribute. All NMDA operations ignore this attribute setting. This setting has no impact on a Lotus restartable scheduled backup, described in "Restartable scheduled backups" on page 41.

Table 9 NetWorker Client resource attributes (page 2 of 5)

Attribute name	NMC tab	Attribute setting
Client Direct	General	Select this attribute to enable the Client Direct feature for a scheduled backup to a DD Boost device or AFTD. If a Client Direct backup is not possible, the backup routes the data through the storage node to the NetWorker device. **Note:** This attribute is only available with NetWorker server release 8.0 or later.
(Avamar) Deduplication Backup	Apps & Modules	Select this attribute to enable a deduplication backup to Avamar. **Note:** You cannot select this attribute together with the Data Domain Backup attribute.
(Avamar) Deduplication Node	Apps & Modules	Specify the hostname of the Avamar server that will store the deduplicated backup.
Data Domain Backup	Apps & Modules	Optionally select this attribute to ensure that NMDA always stores the deduplication backup on a DD Boost device when the backup pool contains a mixture of DD Boost devices and other types of devices. **Note:** You cannot select this attribute together with the (Avamar) Deduplication Backup attribute.
Data Domain Interface	Apps & Modules	Specify the interface over which the backup will occur to a Data Domain device. This option is used only when you enable Data Domain Backup for the client. Select the appropriate setting for the attribute: • Fibre Channel—The backup will use a Fibre Channel (FC) interface. • IP—The backup will use an IP interface. • Any—The backup will use either an FC or IP interface, depending on the available Data Domain device. **Note:** This attribute is only available with NetWorker server release 8.1 or later. To enable a restore from a Data Domain device over an FC interface, you must set this attribute to Fibre Channel.
Group	General	Specify the NetWorker backup group to use for the scheduled backup. "Configuring the Group resource" on page 87 describes backup groups.
Name	General	Specify the hostname of the database server host or application server host.
Parallelism	Globals	Specify the maximum number of data streams that the database server or application server sends in parallel to the NetWorker server or NetWorker storage node during the backup.
Pool	Globals	Specify the pool to use for a backup. When selected, this attribute overrides the pool associated with the backup group.
Remote Access	Globals	Specify the username and fully qualified hostname of a remote system that can access backups of this host, in the following format: **user=***remote_username*,**host=***remote_hostname* Set this attribute in the following cases (otherwise, leave it blank): • Backup from a cluster • Backup in an Oracle Data Guard environment • Recovery to a host other than the one being backed up On a Solaris system with Solaris zones, the Remote Access attribute must contain the hostname of the zone in which NMDA operates.

Table 9 NetWorker Client resource attributes (page 3 of 5)

Attribute name	NMC tab	Attribute setting	
Retention Policy	General	Specify the minimum length of time that the NetWorker server maintains information about the backup data in the online media database. The value of the Retention Policy attribute must be greater than or equal to the value of the Browse Policy attribute. **Note:** If you set the NSR_SAVESET_RETENTION parameter as described in Appendix A, "NMDA Parameters and Configuration File," the parameter value overrides the Retention Policy attribute setting in the Client resource.	
Save Set	General	**Type of client**	**Save Set attribute setting**
		DB2	Specify the database and node to be backed up. Use the following format for the save set name: DB2:/*database_name*/*node_name* You can optionally specify multiple database instances on separate lines, for example: DB2:*DB1*/NODE0000 DB2:*DB2*/NODE0000 where *DB1* and *DB2* are databases of the same or different DB2 instances on the same client. You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in "Syntax rules for the NMDA configuration file" on page 355.
		Informix	Specify the Informix server instance and (optionally) dbspace to be backed up. Separate multiple dbspace names with a space: INFORMIX:/*instance*[/*dbspace1*[*dbspace2*]] To back up the entire instance: INFORMIX:/*instance* To back up selected dbspaces: INFORMIX:/*instance*/*dbspace1 dbspace2* You can optionally specify multiple instances and optionally dbspaces on separate lines, for example: INFORMIX:/INST1 INFORMIX:/INST2/dbspace2_1 dbspace2_2 where INST1 and INST2 are different Informix instances on the same client. You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in "Syntax rules for the NMDA configuration file" on page 355.

Table 9 NetWorker Client resource attributes (page 4 of 5)

Attribute name	NMC tab	Attribute setting	
Save Set	General	**Type of client**	**Save Set attribute setting**
		Lotus	Specify the name for the backup save set: NOTES:[*description*] For example, the following names are valid names: • NOTES: • NOTES: Monday Full Backup To specify the data to back up, set one of these parameters in the NMDA configuration file: • NSR_BACKUP_LOTUS_DIR—To back up the Domino or Notes data directory. • NSR_BACKUP_PATHS—To back up one or more specific directories, or files, or both. Appendix A, "NMDA Parameters and Configuration File," describes the parameters. You can optionally specify save sets for multiple Domino servers on separate lines, for example: NOTES:Domino_Server1 NOTES:Domino_Server2 where Domino_Server1 and Domino_Server2 are names of backups for different Domino servers on the same client. You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in "Syntax rules for the NMDA configuration file" on page 355.
		MySQL	Specify a name starting with MYSQL:/ for the backup: MYSQL:/*unique_backup_name* where *unique_backup_name* is a unique name that identifies the particular MySQL backup, for example, MYSQL:/Innodb_partial_backup. You must specify which MySQL instances and optionally databases and tables to include in the backup by setting the MYSQL_CFG_FILE, MYSQL_DATABASES, and MYSQL_INCLUDE parameters in the NMDA configuration file. Table 37 on page 380 describes the parameters. **Note:** You cannot use the same backup name in different NMDA backup configurations for MySQL. "MySQL backup information in NetWorker indexes" on page 162 and "MYSQL_BACKUP_NAME" on page 381 provide details on MySQL backup names and their importance for restores. You can optionally specify the backup names for multiple instances on separate lines, for example: MYSQL:*DBSID1* MYSQL:*DBSID2* where *DBSID1* and *DBSID2* are the unique backup names for different MySQL instances on the same client. You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in "Syntax rules for the NMDA configuration file" on page 355.

Table 9 NetWorker Client resource attributes (page 5 of 5)

Attribute name	NMC tab	Attribute setting	
Save Set	General	**Type of client**	**Save Set attribute setting**
		Oracle	Specify the complete pathname of each RMAN script to use for the scheduled backup: RMAN:*RMAN_script_pathname* Do not include spaces between the prefix RMAN: and the script name. On Windows, the pathname can include forward slashes, for example, RMAN:F:/scripts/incr_1_bkup. If you create two separate RMAN backup scripts in the files /disk/rman_scripts/archlogbkup and /disk/rman_scripts/fullbkup, specify each script pathname on a separate line: • RMAN:/disk/rman_scripts/archlogbkup • RMAN:/disk/rman_scripts/fullbkup You can optionally specify RMAN scripts for multiple databases on separate lines, for example: RMAN:/disk/rman_backup_PROD RMAN:/disk/rman_backup_SALES where /disk/rman_backup_PROD and /disk/rman_backup_SALES are the rman scripts for different databases, PROD and SALES respectively, on the same client. You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in "Syntax rules for the NMDA configuration file" on page 355.
		Sybase	Specify the Sybase server and optionally the Sybase databases for backup: SYBASE:/*ASE_server_name*[/*database_name*] To back up the entire server, specify the following setting: SYBASE:/*ASE_server_name* You can optionally specify multiple servers and optionally databases (of the same server or different servers) on separate lines, for example: SYBASE:/ASE_15_5_Server SYBASE:/ASE_15_7_Server/SALES This example backs up the entire server ASE_15_5_Server and backs up only the database SALES of ASE_15_7_Server. You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in "Syntax rules for the NMDA configuration file" on page 355
Schedule	General	For scheduled backups only. Specify the NetWorker backup schedule to use for the scheduled backup. "Configuring the Schedule resource" on page 87 describes backup schedules.	
Scheduled Backup	General	For scheduled backups only. Select the checkbox to enable the resource for scheduled backups.	
Storage Nodes	Globals	Specify the name of each storage node to which the database server can back up data. The database server backs up to the first active enabled storage node in the order listed in the attribute. The default storage node name, nsrserverhost, represents the NetWorker server.	

Configuring manual backups

You must complete the required steps to configure an NMDA manual backup.

1. Ensure that you have configured the required NetWorker resources as described in "Verifying the basic NetWorker resource configurations" on page 75.
2. Set the required NMDA parameters:
   - For a DB2, Lotus, MySQL, or Sybase manual backup, set the parameters in the NMDA configuration file.
   - For an Informix manual backup, set the parameters in the environment.
   - For an Oracle manual backup, create an RMAN script and set the parameters in the RMAN backup script by using the **send** command where possible. "Oracle backup considerations" on page 128 describes when you cannot use the **send** command.

   Set NSR_SERVER to the hostname of the NetWorker server for the backup if the server host is different from the NMDA client host.

   Set NSR_DATA_VOLUME_POOL (all NMDA applications) or NSR_LOG_VOLUME_POOL (DB2, Informix, MySQL, Sybase) to ensure that the backup is stored on the preferred volume. Otherwise, the NetWorker server assigns a volume according to the media pool selection criteria on the server.

   Ensure that you set the mandatory parameters for the manual backup of the database or application, as shown in the following table.

   Appendix A, "NMDA Parameters and Configuration File," describes the parameters.

Table 10 Mandatory parameters for manual NMDA backups

Database or application	Mandatory parameters for manual NMDA backups
Lotus	• Notes_ExecDirectory • NSR_RESOURCE_DIR (UNIX only) • PATH (UNIX only)
MySQL	• MYSQL_BACKUP_NAME • MYSQL_CFG_FILE • MYSQL_MEB_PATH (required only if MEB is not installed in the default location or if you have 32-bit and 64-bit MEB on this same system) • MYSQL_SBT_LIB_PATH (required only if you have 32-bit and 64-bit MEB on this same system)
Oracle	• NSR_DATA_VOLUME_POOL* (for backups that generate backup copies)
Sybase	• NSR_BACKUP_PATHS • SYBASE • SYBASE_USER • USER_PSWD (when Sybase server has a password) **Note:** To set the encrypted Sybase user password in the USER_PSWD parameter, you must use the **nsrdaadmin -P** command, as described in "USER_PSWD" on page 396.

3. Perform any additional configurations required for the specific database backup or application backup:

   - "DB2 backup considerations" on page 116
   - "Informix backup considerations" on page 117
   - "Lotus backup considerations" on page 118
   - "MySQL backup considerations" on page 125
   - "Oracle backup considerations" on page 128
   - "Sybase backup considerations" on page 137

Configuring deduplication backups with Avamar

To prepare for Avamar deduplication backups, review the requirements, best practices, and configuration procedures for scheduled deduplication backups and manual deduplication backups with Avamar.

Requirements for Avamar deduplication backups

Ensure that you meet the following requirements before you configure a deduplication backup with Avamar:

- The NetWorker client release and NetWorker server release support NMDA deduplication. The *EMC NetWorker Module for Databases and Applications Release Notes* provides details.
- You have installed and configured the Avamar server (DataGrid) as a NetWorker deduplication node. The *EMC NetWorker Avamar Integration Guide* provides details.
- For the NetWorker backup device that receives only backup metadata during the deduplication backup, you have configured the device as an advanced file type device (AFTD) as described in the *EMC NetWorker Avamar Integration Guide*.
- The user that will perform deduplication backups or restores has the Monitor NetWorker user group privilege. "NetWorker User Group resource" on page 75 describes user group privileges.
- You have set the NSR_NWPATH parameter for these Avamar deduplication backups:
  - Sybase deduplication backup on 64-bit Windows
  - Any deduplication backup with 32-bit NMDA installed on 64-bit Windows

Best practices for Avamar deduplication backups

Avamar deduplication is most beneficial in a data warehouse environment where the data does not change frequently. Deduplication can also be beneficial for databases where only a small percentage of data is updated repeatedly, or where new data is added to a database but the old data does not change much.

You can improve the performance of Avamar deduplication backups with the following best practices:

- Do not include a deduplication client in the same Group resource as nondeduplication clients.
- After you select a deduplication node (Avamar server) for an initial full backup of a client, continue to use that same node for all backups of the client. This practice takes advantage of the client deduplication information stored on the node.
- Schedule deduplication backups to avoid the Avamar server read-only periods. An Avamar server spends periods of time in maintenance mode, where the server might be unavailable for backup or have limited bandwidth.

  Note: A deduplication NMDA backup that runs during an Avamar maintenance period might become suspended until the Avamar server resources become available.

- Do not use database backup encryption because the encryption decreases the data deduplication ratio.
- Do not use database backup checksumming because the checksumming decreases the data deduplication ratio.
- Do not use database backup compression because the compression decreases the data deduplication ratio.
- NMDA applies AES encryption, checksumming, or compression, respectively, to only the metadata stored on the NetWorker backup media if you do both of the following actions:
  - You use NMDA backup encryption, checksumming, or compression.
  - You set the NSR_AES_ENCRYPTION, NSR_CHECKSUM, or NSR_COMPRESSION parameter.
- Do not use more than four to six concurrent backup sessions. Setting the backup parallelism to a high value can cause the deduplication backup to fail.

  For DB2 backups, do not use the DB2 PARALLELISM parameter. DB2 parallelism is different from the NMDA DB2 parallel (multisession) backups. The DB2 documentation describes DB2 parallelism.
- Do not use deduplication for incremental backups. An incremental backup backs up only the database blocks that have changed (DB2, Informix, or Oracle) or the transaction logs (Lotus or Sybase), which guarantees a low rate of duplication. With the overhead of deduplicating data added to the overhead of an incremental backup, the result is decreased performance and insignificant benefits for the data stored.

The application of browse and retention policies and the selection of media pools is the same for a deduplication backup as for a regular NMDA backup.

Avamar and NetWorker documentation describes the requirements for deduplication backups.

DB2 considerations for Avamar deduplication

You can improve the performance of DB2 deduplication backups with Avamar by using an appropriate setting:

- For scheduled backups, you can use the DB2BACKUP_DEDUP_DEVICE setting with the DB2_OPTIONS parameter as described in "DB2_OPTIONS" on page 368.
- For manual backups, you can use the **dedup_device** option of the DB2 **backup** command.

You can also improve the deduplication ratio by setting a larger DB2 backup buffer size with an appropriate setting:

- For scheduled backups, you can set the DB2_BUFFER_SIZE parameter as described in "DB2_BUFFER_SIZE" on page 367.
- For manual backups, you can use the **buffer** *buffer_size* option with the DB2 **backup** command.

The DB2 documentation provides more details about DB2 tuning parameters for deduplication devices.

Oracle considerations for Avamar deduplication

You can improve the performance of Oracle deduplication backups with the following additional best practices:

- For database full backups, you can improve the deduplication ratio by not using RMAN multiplexing and by ensuring that the Oracle software does not concatenate data from different files in the same backup piece. To disable multiplexing and prevent concatenation, ensure that **filesperset** is set to 1 in the RMAN configuration. The Oracle documentation describes RMAN multiplexing and the number of files in each backup set.
- If you perform database incremental backups or archived log backups to Avamar, the Oracle software sends only the changed blocks. As a result, the deduplication ratio is typically low. To increase the backup throughput for these backups, set **filesperset** to a larger value or do not set **filesperset**. If you do not set **filesperset**, then the Oracle software uses the default setting. The Oracle documentation provides more details about **filesperset**.
- For a manual backup or scheduled backup configured without the wizard (client-side configuration), set the NSR_DEDUP_CACHE_TAG parameter to a unique value for each allocated channel when the backup uses multiple channels. "NSR_DEDUP_CACHE_TAG" on page 361 provides details.
- When the database structure does not change often, associate the backup of specific tablespaces with a specific channel. This practice ensures that Oracle does not distribute the data to a different channel when the database structure or database size changes. Group tablespaces that contain similar (duplicated) data and associate them with the same channel. This practice requires familiarity with the database data. You must add a tablespace to the backup script when you create a tablespace.

For example, the following RMAN backup script shows how to set the cache tag per channel and to associate tablespaces with a channel:

```
run {
   allocate channel c1 type 'SBT_TAPE';
   send channel c1 'NSR_ENV=(NSR_DEDUP_CACHE_TAG=orcl102_c1)';
   allocate channel c2 type 'SBT_TAPE';
   send channel c2 'NSR_ENV=(NSR_DEDUP_CACHE_TAG=orcl102_c2)';

   send 'NSR_ENV=(NSR_DEDUP_BACKUP=TRUE,
   NSR_DEDUP_NODE=avamar.emc.com)';

   backup filesperset=1
   (tablespace tbs1, tbs5 channel c1)
   (tablespace tbs2, tbs3, tbs4 channel c2);

   release channel c1;
   release channel c2;
}
```

Configuring scheduled Avamar deduplication backups

To configure a scheduled deduplication backup, use either the configuration wizard (server-side configuration) or the NMC method without the wizard (client-side configuration).

Server-side configuration

To use the configuration wizard to configure a scheduled deduplication backup, follow the instructions in "Configuring scheduled backups with the wizard" on page 82.

On the screen to specify deduplication options:

- Select the attribute to enable an Avamar deduplication backup.
- Select the hostname of the deduplication node (Avamar server) that will store the deduplicated backup data.

Client-side configuration

To use the NMC method without the wizard to configure a scheduled deduplication backup, follow the instructions in "Configuring scheduled backups without the wizard" on page 85.

You must complete the required steps to apply additional settings for a deduplication backup.

1. In the NetWorker Client resource for the NMDA client:
   - Select the attribute to enable an Avamar deduplication backup.
   - Select the hostname of the deduplication node (Avamar server) that will store the deduplicated backup data.
2. If the **nsravtar** binary (part of the NetWorker client package) is in a nondefault location, set NSR_NWPATH according to the method in "NMDA configuration file" on page 354. "NSR_NWPATH" on page 364 describes the parameter.
3. For an Oracle deduplication backup only, set the NSR_DEDUP_CACHE_TAG parameter in the RMAN script according to "Oracle considerations for Avamar deduplication" on page 97.

4. For normal operations, use default settings for the cache usage and the chunk size. However, if you require nondefault settings, you can modify the following parameters with special care in the NMDA configuration file or in the Oracle RMAN script:

   - "NSR_DEDUP_CACHE_ENABLED" on page 360
   - "NSR_DEDUP_CHUNK_SIZE" on page 361

Configuring manual Avamar deduplication backups

You must complete the required steps to configure NMDA manual backups that use Avamar deduplication.

1. In at least one of the NetWorker Client resources for the NMDA client, use NMC to:

   - Select the attribute to enable an Avamar deduplication backup.
   - Select the hostname of the deduplication node (Avamar server) that will store the deduplicated backup data.

2. Set the following parameters in the NMDA configuration file for a DB2, Informix, Lotus, or Sybase backup, or in the RMAN script for an Oracle backup:

   - "NSR_DEDUP_BACKUP" on page 360
   - "NSR_DEDUP_NODE" on page 362

   The NSR_DEDUP_NODE parameter setting must have exactly the same value as the De-duplication Node attribute in step 1. For example, when you use a short name or fully-qualified name in the attribute setting, you must use the same short name or fully-qualified name in the parameter setting.

3. If the **nsravtar** binary (part of the NetWorker client package) is in a nondefault location, set NSR_NWPATH according to the method in "NMDA configuration file" on page 354. "NSR_NWPATH" on page 364 describes the parameter.

4. If you are configuring an Oracle manual deduplication backup, review the "Oracle considerations for Avamar deduplication" on page 97.

Configuring Data Domain backups

This topic provides specific information about using DD Boost with NMDA. The *EMC NetWorker and EMC Data Domain Deduplication Devices Integration Guide* provides additional information about the DD Boost configuration and initial setup.

Deduplication backups and restores with a Data Domain system are preferable to those with an Avamar server for large databases with high data change rates. A daily change rate of more than 5 percent is a high change rate.

"Deduplication backups and restores with EMC Data Domain" on page 30 describes the NMDA support of deduplication backups with a Data Domain system configured as a NetWorker AFTD, VTL device, or DD Boost device.

You can use the following information to configure a Data Domain backup to a supported device:

- NetWorker documentation describes the configuration of a Data Domain backup to an AFTD or VTL device.
- "Configuring backups to AFTD storage" on page 104 describes the configuration of a backup to an AFTD through either the Client Direct method or a storage node.
- "Configuring DD Boost backups to use Client Direct" on page 101 and "Configuring DD Boost backups to use a storage node" on page 103 describe the configuration of a Data Domain backup to a DD Boost device.

Best practices for Data Domain backups to any device

You can improve the performance of Data Domain backups to any type of device (AFTD, VTL device, or DD Boost device) with the following best practices:

- Do not use backup compression or NMDA backup compression because the compression decreases the data deduplication ratio.
- Do not use backup encryption or NMDA backup encryption because the encryption decreases the data deduplication ratio.
- Follow the best practices for DB2 backups to Data Domain in "DB2 considerations for Data Domain backups" on page 100.
- Follow the best practices for Oracle backups to Data Domain in "Oracle considerations for Data Domain backups" on page 101.

DB2 considerations for Data Domain backups

You can improve the performance of DB2 backups to Data Domain by using an appropriate setting:

- For scheduled backups, you can use the DB2BACKUP_DEDUP_DEVICE setting with the DB2_OPTIONS parameter as described in "DB2_OPTIONS" on page 368.
- For manual backups, you can use the **dedup_device** option of the DB2 **backup** command.

You can also improve the deduplication ratio by setting a larger DB2 backup buffer size with an appropriate setting:

- For scheduled backups, you can set the DB2_BUFFER_SIZE parameter as described in "DB2_BUFFER_SIZE" on page 367.
- For manual backups, you can use the **buffer** *buffer_size* option with the DB2 **backup** command.

The DB2 documentation provides more details about DB2 tuning parameters for deduplication devices.

Oracle considerations for Data Domain backups

You can improve the performance of Oracle backups to Data Domain as follows:

- Database full backups—You can improve the database full backup deduplication ratio by not using RMAN multiplexing and by ensuring that the Oracle software does not concatenate data from different files in the same backup piece. To disable multiplexing and prevent concatenation, ensure that **filesperset** is set to 1 in the RMAN configuration. The Oracle documentation describes RMAN multiplexing and the number of files in each backup set.
- Database incremental backups or archived log backups—The Oracle software sends only the changed blocks in these backups. As a result, the deduplication ratio is typically low. To increase the backup throughput for these backups, set **filesperset** to a larger value or do not set **filesperset**. If you do not set **filesperset**, then the Oracle software uses the default setting. The Oracle documentation provides more details about **filesperset**.
- Multisection backups—You can further improve the backup and restore throughput if you have large datafiles by using multisection backups and specifying **section size**. In a multisection backup, multiple channels can back up a single file. During the restore of a multisection backup, multiple channels can restore a single file. There is no impact to the deduplication ratio, assuming that **section size** is not set to a very small value (less than 1 GB). The Oracle documentation provides details about multisection backups and performance.

Best practices for DD Boost backups

Before you configure a Data Domain deduplication backup with a DD Boost device, ensure that the NetWorker client and server releases support the DD Boost device. The *EMC NetWorker Module for Databases and Applications Release Notes* provides details.

Use Client Direct access for DD Boost backups to eliminate data transfer between the NMDA client and NetWorker storage node. However, using this feature requires an additional 24 MB of memory per data stream on the NMDA client and also increases the NMDA CPU usage.

Data Domain and NetWorker documentation describes the requirements for deduplication backups.

Configuring DD Boost backups to use Client Direct

You can configure a scheduled backup or manual backup to a DD Boost device that uses the Client Direct feature to deduplicate data on the NMDA host and pass only unique data blocks directly to the device, bypassing the storage node.

1. Review "Best practices for DD Boost backups" on page 101.
2. Configure the DD Boost device and enable the Data Domain system for NetWorker operations. The *EMC NetWorker and EMC Data Domain Deduplication Devices Integration Guide* provides details.

3. For a Client Direct backup to a DD Boost device over Fibre Channel, ensure that the database-specific operating system user has the correct device permissions as described in the following documents:

   - Fibre Channel Devices with Products using DD Boost in Linux Environment (Document ID 95007)
   - Fibre Channel Devices with Products using DD Boost in Windows Environment (Document ID 95005)

   Note: This step is not required for backups through a storage node or for snapshot-based backups performed with NSM. Table 5 on page 77 provides a definition of the database-specific operating system user.

4. Associate the backup with a pool that contains the DD Boost device configured in step 2:

   - For a scheduled backup, associate the required media pool with the NMDA Group resource or Client resource by using the NMC program. The *EMC NetWorker Administration Guide* and NMC online help describe how to use the NMC program.
   - For a manual backup, set the NSR_DATA_VOLUME_POOL parameter to the required pool. "NSR_DATA_VOLUME_POOL" on page 359 describes the parameter.

5. Follow the configuration instructions in the following table.

6. For a Client Direct backup, use the default configuration settings. The Client Direct feature is enabled by default.

   With Client Direct enabled, NMDA tries to perform a Client Direct backup. If a Client Direct backup is not possible, NMDA reverts to a backup through the storage node.

   Note: If you upgrade from NetWorker server release 7.6.*x* to 8.0 or later, and you want to disable Client Direct for scheduled backups, you must clear the Client Direct checkbox in the wizard or disable the Client Direct attribute in the Client resource. With NetWorker server 8.0 or later, the Client Direct setting in the wizard field or Client resource takes precedence over the NSR_DIRECT_ACCESS parameter setting.

Table 11 Configuring a Data Domain backup (page 1 of 2)

Type of Data Domain deduplication backup	Configuration instructions	Additional instructions to use DD Boost over Fibre Channel or if the pool contains Data Domain and other types of devices
Scheduled backup configured with the wizard	"Configuring scheduled backups with the wizard" on page 82	On the Client Backup Options wizard page: • Select the Data Domain backup . • Select the proper Data Domain interface. Ensure that you select the Fibre Channel setting if you want to use an FC interface.

Table 11 Configuring a Data Domain backup (page 2 of 2)

Type of Data Domain deduplication backup	Configuration instructions	Additional instructions to use DD Boost over Fibre Channel or if the pool contains Data Domain and other types of devices
Scheduled backup configured without the wizard	"Configuring scheduled backups without the wizard" on page 85	In the Client resource of the NMDA host: • Select the Data Domain Backup attribute. • Select the proper setting of the Data Domain Interface attribute. Ensure that you select the Fibre Channel setting if you want to use an FC interface.
Manual backup	"Configuring manual backups" on page 94	In the NMDA configuration file: • Set NSR_DEVICE_INTERFACE=DATA_DOMAIN. "NSR_DEVICE_INTERFACE" on page 362 provides details. • Set NSR_DATA_DOMAIN_INTERFACE to the proper interface value. Ensure that you set the Fibre Channel value if you want to use an FC interface. "NSR_DATA_DOMAIN_ INTERFACE" on page 359 provides details.

Configuring DD Boost backups to use a storage node

You can configure a scheduled backup or manual backup to a DD Boost device, where the backup does not try to use the Client Direct method but instead deduplicates data on a storage node.

1. Follow step 1 to step 5 in "Configuring DD Boost backups to use Client Direct" on page 101.
2. Ensure that the Client Direct feature is disabled according to the appropriate procedure:
   - Set NSR_DIRECT_ACCESS=No in the appropriate location if required:
     - For a manual backup with any NetWorker server release, set the parameter in the NMDA configuration file (DB2, Lotus, MySQL, Sybase), Oracle RMAN script, or environment variable (Informix).
     - For a wizard configuration of a scheduled backup with a NetWorker server release earlier than 8.0, set the parameter in the advanced option table in the wizard.
     - For a client-side configuration of a scheduled backup with a NetWorker server release earlier than 8.0, set the parameter in the NMDA configuration file (DB2, Informix, MySQL, Lotus, Sybase) or Oracle RMAN script.

     "NSR_DIRECT_ACCESS" on page 362 describes the parameter.
   - For a scheduled backup with NetWorker server release 8.0 or later, clear the Client Direct checkbox in the wizard or disable the Client Direct attribute in the NetWorker Client resource.

Configuring transaction log backups to traditional NetWorker devices

You can selectively send data to both Data Domain devices and traditional NetWorker devices within the same backup. You can configure a single Informix, MySQL, or Sybase backup that backs up database data to a Data Domain device and backs up transaction logs to a traditional NetWorker device.

1. Follow step 1 to step 5 in "Configuring DD Boost backups to use Client Direct" on page 101 with the following condition:

   Ensure that the pool selected for a database backup contains only Data Domain devices. Do not follow the instructions in the third column of Table 11 on page 102.

2. Configure a pool with traditional devices for the log backups and set the NSR_LOG_VOLUME_POOL parameter to the pool name.

3. Do not set NSR_DIRECT_ACCESS=Yes or the backup will fail. Ensure the correct NSR_DIRECT_ACCESS setting:

   - To back up the database data directly from the NMDA host, set NSR_DIRECT_ACCESS to the default value.
   - To back up the database data through the storage node, set NSR_DIRECT_ACCESS=No.

Configuring backups to AFTD storage

You can configure an NMDA scheduled or manual backup to an AFTD as either a Client Direct backup or a backup that sends data through a storage node.

Configuring Client Direct backups to AFTD

You must complete the required steps to configure a scheduled or manual backup that uses Client Direct to an AFTD.

1. Configure an AFTD by using either the Device Wizard or the NMC device properties window. Ensure that you meet all the AFTD requirements for the storage node and backup client. The *EMC NetWorker Administration Guide* provides complete details about AFTD configurations, including considerations for Client Direct clients.

2. Ensure that the NMDA client has the required access to the AFTD through CIFS. The *EMC NetWorker Administration Guide* provides details.

3. Associate the backup with a pool that contains the AFTD configured in step 1.

4. Follow the appropriate configuration instructions:

   - "Configuring scheduled backups with the wizard" on page 82
   - "Configuring scheduled backups without the wizard" on page 85
   - "Configuring manual backups" on page 94

5. For a Client Direct backup, use the default configuration settings. The Client Direct feature is enabled by default.

   With Client Direct enabled, NMDA tries to perform a Client Direct backup. If a Client Direct backup is not possible, NMDA reverts to a traditional backup through the storage node.

Configuring storage node backups to AFTD

You must complete the required steps to configure a scheduled or manual backup to an AFTD where the backup does not try to use the Client Direct method but instead sends data through a storage node.

1. Follow step 1 to step 4 in "Configuring Client Direct backups to AFTD" on page 104.
2. Ensure that the Client Direct feature is disabled according to the appropriate procedure:
   - For a manual backup, set NSR_DIRECT_ACCESS=No in the NMDA configuration file (DB2, Lotus, MySQL, Sybase), Oracle RMAN script, or environment variable (Informix). "NSR_DIRECT_ACCESS" on page 362 provides details.
   - For a scheduled backup, clear the Client Direct checkbox in the wizard or disable the Client Direct attribute in the NetWorker Client resource.

Configuring probe-based backups

Note: NMDA supports probe-based backups of Oracle disk backups in addition to regular Oracle data.
NMDA does not support probe-based backups of MySQL data.

Before you configure a probe-based backup, ensure that you have installed the required NetWorker releases, as described in the *EMC NetWorker Module for Databases and Applications Release Notes*.

An NMDA probe-based backup starts when both of these conditions are true:

- (Condition 1) The current time is within a specified window of time (the backup window, as defined by the probe start time and probe end time in the probe-enabled backup group resource).
- The backup meets one of the following conditions:
  - (Condition 2) A specified amount of time has elapsed since the previous probe-based backup.
  - (Condition 3) One or all of the probes associated with the backup are successful, depending on the probe success criteria specified in the backup configuration.

    You can set the probe success criteria in the NetWorker Group resource to the value Any or All.

At specified probe intervals, the NetWorker server performs the following steps:

1. The server checks for condition 1, to determine if the current time is within the backup window.
2. If the backup meets condition 1, then the server checks for condition 2, to determine if a specified amount of time has elapsed since the last probe-based backup:
   - If the backup meets condition 2, then the server starts the probe-based backup.
   - If the backup does not meet condition 2, then the server checks for condition 3, to determine if one or all of the probes are successful:
     - If you set the probe success criteria to Any and any probe is successful, then the server starts the probe-based backup.
     - If you set the probe success criteria to All and all the probes are successful, then the server starts the probe-based backup.

NMDA probe

The NMDA **nsrdaprobe** program checks for one of the following conditions for the database or application:

- For DB2, Informix, or Oracle, the **nsrdaprobe** program checks for the number of logs generated since the last probe-based backup.
- For Lotus, the **nsrdaprobe** program checks for the size of Domino transaction logs generated since the last backup of the Domino database specified by LOTUS_NSF_FILE in the NetWorker Probe resource.
- For Sybase, the **nsrdaprobe** program checks for the number of transaction log pages generated since the last probe-based backup.

Note: After a point-in-time restore to an earlier time or any other database operation that causes a log reset, the probe triggers a new backup.

NOTICE

(Sybase only) Due to a Sybase limitation, NMDA does not support probe-based backups with **nsrdaprobe** of a Sybase database that has both data and log segments on the same device. NMDA cannot determine the used log pages for such a configuration.
For example, assume that database SYBDB1 on Sybase server SYBSERVER1 has data and log segments on the same device. If you configure a probe-based backup of SYBDB1 only (SYBASE:/SYBSERVER1/SYBDB1), **nsrdaprobe** always requests that the backup runs, regardless of the amount of log pages generated.
If you configure a probe-based backup of the whole server (SYBASE:/SYBSERVER1), **nsrdaprobe** does not use the log pages for SYBDB1 when calculating the amount of log pages generated for the whole SYBSERVER1.

User-defined probe

If you want to check for a user-defined condition (other than the number or size of generated logs) that triggers a probe-based backup, create a script or program that meets these requirements:

- The script or program name starts with **nsr** or **save**.
- The script or program is in the same directory as the NetWorker client binaries on the host where the condition is checked.
- The script or program file includes the "execute" permission.
- The script or program returns one of these code values when finished running the probe:
  - 0—Signifies that the backup condition has been met. The backup runs.
  - 1—Signifies that the backup condition has not been met. The backup does not run.
  - Other than 0 or 1—Signifies that an error occurred during the probe. The backup does not run.

Example 5 on page 111 shows an NMDA probe-based backup with a user-defined probe.

Configuring a probe-based backup

To configure a probe-based backup, you must configure the required NetWorker Probe, Group, and Client resources and complete the scheduled backup configuration.

Note: For MySQL probe-based backups, NMDA supports only user-defined probes, not the NMDA probe implemented through the **nsrdaprobe** program.

1. Create a separate NetWorker Probe resource for the **nsrdaprobe** program and any other script or program that checks for a user-defined condition.

   Set the Probe resource attributes as described in the following table.

   NOTICE

   Create a separate NetWorker Probe resource for each database that the **nsrdaprobe** program will probe.

Table 12 NetWorker Probe resource attributes (page 1 of 3)

Attribute	Description
Name	Specify a name to identify the Probe resource. Each Probe resource must have a unique name, which does not have to be the same as the probe script or program name.
Probe Command	Specify the name of the probe script or program that checks ("probes") for the condition that triggers a probe-based backup. Specify **nsrdaprobe** to use the NMDA probe, or specify the user-created script or program.

Table 12 NetWorker Probe resource attributes (page 2 of 3)

Attribute	Description		
Command Options	Required for the **nsrdaprobe** program only. Specify a comma-separated list of parameters with their settings. "Command Options settings for the nsrdaprobe program" on page 110 describes the possible parameter settings.		
	Parameter	**Type of client**	**Description**
	DBA_DISK_BACKUP	Oracle	**Optional. Set to TRUE if** both of the following conditions are true on the Oracle client: • A backup administrator has used the wizard to configure NMDA scheduled backups of Oracle disk backups of the database. • You want **nsrdaprobe** to **check if any new Oracle disk backups have been created for the database since the last successful probe-based backup.** With DBA_DISK_BACKUP=TRUE, if the probe finds any new Oracle disk backups and no Oracle disk backup is currently running for the database, the probe triggers an NMDA scheduled backup of the new disk backups. **Note:** Do not set DBA_DISK_BACKUP in the same Probe resource as LOG_THRESHOLD.
	LOG_THRESHOLD	DB2, Informix, Oracle	Mandatory unless DBA_DISK_BACKUP=TRUE is set for an Oracle client. Specify the change threshold, which is the minimum number of logs required to trigger a new probe-based backup. **Note:** Do not set LOG_THRESHOLD in the same Probe resource as DBA_DISK_BACKUP.
		Lotus	Mandatory. Specify the change threshold, which is the minimum size in KB of Lotus Domino transaction logs required to trigger a new probe-based backup. **Notice:** For Lotus circular and linear transaction logging, the change threshold value must be sufficiently smaller than the value specified by "Maximum log space" in the Domino server "Transactional Logging" configuration. This requirement ensures that a probe-based backup is triggered before the transaction log size exceeds the maximum log space value and causes transaction logs to be overwritten.
		Sybase	Mandatory. Specify the change threshold, which is the minimum number of Sybase transaction log pages required to trigger a new probe-based backup.
	LOTUS_NSF_FILE	Lotus	Recommended. Specify the full pathname of a Domino database file used for transaction log querying. The specified database file can be any file backed up by the associated scheduled backup. If you do not set this parameter, the default value used is admin4.nsf in the Lotus data directory (UNIX) or C:\Program Files\IBM\Lotus\Domino\data\admin4.nsf (Windows).
	NSR_DEBUG_LEVEL	All	Optional. Specify the level of debug information generated by the probe and written to the debug log. "NSR_DEBUG_LEVEL" on page 360 describes debug levels.
	NSR_DIAGNOSTIC_DEST	All	Optional. Specify the directory location of the NMDA debug logs, including the debug logs generated by the probe. "NSR_DIAGNOSTIC_DEST" on page 362 provides details.

Table 12 NetWorker Probe resource attributes (page 3 of 3)

Attribute	Description		
Command Options	**Parameter**	**Type of client**	**Description**
	NSR_ORACLE_ CONNECT_FILE	Oracle	Optional. Set if both of the following conditions are true: • You have not configured the Client resource with the wizard. • You have not set up the NWORA resource file with the Oracle home and database connection information. Specify the pathname of the RMAN connection file, which contains the connection strings required to connect to the Oracle database. "Command Options settings for the nsrdaprobe program" on page 110 provides a sample setting of this parameter.
	ORACLE_SERVICE	Oracle	Optional. Set if both of the following conditions are true: • You have not configured the Client resource with the wizard. • You have set up the NWORA resource file with the Oracle home and database connection information by using the command **nsroraadmin –r add sid=***Net_service_name* **home=***Oracle_home* **connect=***connect_filepath*. "Configuring the NWORA resource file with the nsroraadmin program" on page 329 describes the **nsroraadmin** command. Specify the Net service name for the Oracle database. The ORACLE_SERVICE setting must be the same as the NSR_ORACLE_SID setting in the NWORA resource file. "Command Options settings for the nsrdaprobe program" on page 110 provides a sample setting of this parameter.

At the end of a successful probe-based backup, the **nsrdaprobe** program stores information about the current transaction log and some other database information in the State attribute of the Probe resource. User-defined probes do not use the State attribute. The State attribute is visible only in diagnostic mode.

2. Configure the NetWorker Group resource for a probe-enabled backup group. Set the probe-specific attributes in the Group resource, as described in the *EMC NetWorker Administration Guide* (the information about creating and scheduling a probe group).

   When you enable probing through the Group resource attributes, each probe runs once when the probe interval is reached, which is the time window defined by the Probe Start Time and Probe End Time attributes. A probe-based backup ignores the group's Start Time attribute.

   Note: If you start a probe-enabled backup group manually, probing occurs immediately (only once, not repeatedly at intervals) and the backup starts only if the probe conditions are met.

3. Configure the scheduled backup with or without the wizard:

   - "Configuring scheduled backups with the wizard" on page 82
   - "Configuring scheduled backups without the wizard" on page 85

In addition, in the Client resource:

- For the Probe attribute, specify the name of the Probe resource from step 1. You can associate a Client resource with only one probe.

  Note: The configuration wizard does not display the Probe field. If you configure a Client resource with the wizard, you must then use NMC to edit the Client resource and set the Probe attribute.

- For the Group attribute, specify the probe-enabled group from step 2.

  Note: A probe-based backup group must include at least one probe-enabled client, which is a Client resource associated with a probe.

Command Options settings for the nsrdaprobe program

To use the NMDA **nsrdaprobe** program, you must include specific parameter settings in the Command Options attribute in the Probe resource. The parameter settings depend on the particular scenario:

- The LOG_THRESHOLD parameter is mandatory unless DBA_DISK_BACKUP is set for an Oracle client. Do not set both DBA_DISK_BACKUP and LOG_THRESHOLD in the same Probe resource.
- For debugging purposes only, set the NSR_DEBUG_LEVEL and NSR_DIAGNOSTIC_DEST parameters. Set NSR_DIAGNOSTIC_DEST only if you do not have enough space for the logs in the default location. For example:

  ```
  LOG_THRESHOLD=10, NSR_DEBUG_LEVEL=9, NSR_DIAGNOSTIC_DEST=/tmp
  ```

- For a Lotus backup only, set the LOTUS_NSF_FILE parameter.
- For an Oracle backup only, three possible scenarios dictate the required settings in the Command Options attribute:

  a. You have configured the Client resource with the wizard.

     In this case, the Command Options attribute must include the LOG_THRESHOLD parameter unless you set DBA_DISK_BACKUP=TRUE. For example:

     ```
     LOG_THRESHOLD=10
     ```

     or

     ```
     DBA_DISK_BACKUP=TRUE
     ```

  b. You have configured the Client resource without the wizard (client-side configuration) and you have not set up the NWORA resource file with the Oracle home information and database connection information.

     In this case, the Command Options attribute must include the parameters LOG_THRESHOLD and NSR_ORACLE_CONNECT_FILE. For example:

     ```
     LOG_THRESHOLD=10, NSR_ORACLE_CONNECT_FILE=/RMAN/rmanpw
     ```

c. You have configured the Client resource without the wizard and you have set up the NWORA resource file to retrieve Oracle home information and database connection information.

The Command Options attribute must include the parameters LOG_THRESHOLD and ORACLE_SERVICE, with ORACLE_SERVICE set to the same Net service name as NSR_ORACLE_SID in the NWORA file. For example:

```
LOG_THRESHOLD=10, ORACLE_SERVICE=proddb.world
```

You must set up the NWORA resource file by using the command **nsroraadmin -r add sid=***Net_service_name* **home=***Oracle_home* **connect=***connect_filepath*.

"Configuring the NWORA resource file with the nsroraadmin program" on page 329 describes the **nsroraadmin** command.

Example 5 Multiple probes for a probe-based backup

You can configure a probe-based backup with multiple probes. Depending on the Probe Success Criteria setting in the Group resource, a backup starts when the backup meets the conditions for any or all of the probes.

In the following Oracle example, the trigger for the probe-based backup is when both these conditions are true:

- At least 25 Oracle log files are generated on an NMDA client named dbhost.
- More than two tape drives are idle in a jukebox, attached to a NetWorker storage node named jukeboxhost. You use the Jukebox for the probe-based backup.

The probe runs every hour between 6 p.m. and 2 a.m. or until the preceding conditions are met, whichever occurs first.

Both the NMDA client and the storage node are Solaris machines.

The **nsrdaprobe** program is in the /usr/sbin directory on the NMDA client. The **nsrdaprobe** program checks for the number of Oracle log files generated on the NMDA client.

You created a script named **nsrjukeboxprobe** with "execute" permissions in the /usr/sbin directory on the storage node. The script checks for the number of idle tape drives in the jukebox, and returns either of two values:

- 0—Signifies that more than two tape drives are idle in the jukebox.
- 1—Signifies that two or fewer tape drives are idle in the jukebox.

You can configure this probe-based backup with multiple probes as follows:

1. Create a Probe resource for the **nsrdaprobe** program with the following attribute settings:
   - Name: NMDA Oracle probe
   - Probe Command: **nsrdaprobe**
   - Command Options: LOG_THRESHOLD=25
2. Create a Probe resource for the user-defined probe with the following attribute settings:
   - Name: Jukebox probe
   - Probe Command: **nsrjukeboxprobe**

3. Create a Group resource with the required attribute settings for the probe-enabled backup group, including the following settings:

   - Name: probe_group
   - Probe Interval: 60
   - Probe Start Time: 18:00
   - Probe End Time: 2:00
   - Probe Based Group: Enabled

     Note: The Probe Based Group attribute is a checkbox in NMC.

   - Probe Success Criteria: All

4. Configure an NMDA scheduled backup through the wizard. The Client resource includes the following attribute settings:

   - Name: dbhost
   - Backup Command: **nsrdasv -T oracle**
   - Group: probe_group
   - Probe: NMDA Oracle probe
   - Save Set: RMAN:/orcl102_FULL

5. Create a Schedule resource named SkipAll, and set the level to skip for each day in the schedule. This resource enables a jukebox probe to run on the storage node without backing up the storage node host:

   - Name: SkipAll
   - Period: Either Week or Month
   - Calendar: Skip

     Note: Select the Skip level for every day in the period.

6. Create a generic Client resource without the wizard for the storage node host, with the following attribute settings:

   - Name: jukeboxhost
   - Backup Command: (blank)
   - Group: probe_group
   - Probe: Jukebox probe
   - Save Set: SKIP

     Note: The Save Set attribute requires a keyword.

   - Schedule: SkipAll

Configuring parallel backups

NMDA supports backups performed with multiple parallel sessions (also known as multiple streams, multiple channels, or multinode). This technology extracts multiple streams of data in parallel from a database, and writes them in parallel to one or more storage devices. This method can enhance performance when you back up or restore a large amount of data.

A parallel backup may lead to data multiplexing. Multiplexing is the NetWorker ability to write multiple backup streams simultaneously to the same storage device. NetWorker multiplexes backup streams when all the following conditions are true:

- The NetWorker device type is not AFTD or Data Domain.
- The number of parallel backup sessions is greater than the number of devices available for backup.
- The target session value for the device is not 1.

The use of NetWorker multiplexing can improve backup performance, but it can adversely affect restore performance for Informix, Lotus, and Oracle.

NOTICE

You must disable NetWorker multiplexing for DB2 backups and Sybase backups. Otherwise, the restore will become suspended or will fail.

You can disable NetWorker multiplexing by setting the number of backup sessions equal to the number of available storage devices, and by setting the target session value to 1 for those devices.

Parallel backups to AFTD or Data Domain devices are not multiplexed. However, the use of multiple concurrent streams to the same device can impact the overall throughput and deduplication ratio. Minimize the number of sessions to any device that includes deduplication by setting the target session attribute value of the device to 1. The *EMC NetWorker and EMC Data Domain Deduplication Devices Integration Guide* provides details.

To configure an NMDA parallel backup as either a manual or scheduled backup:

1. Determine the number of available devices, the target sessions, and the maximum sessions per device for the backup.
2. On the NetWorker server:
   a. Set the server parallelism in the NetWorker Server resource and the client parallelism in the NetWorker Client resource to a value greater than the number of sessions used for NMDA backups. "NetWorker Server resource" on page 75 and "Configuring the Client resource" on page 89 describe the parallelism setting.
   b. If required, create a Device resource for each device according to "NetWorker Device resource" on page 78.
   c. For DB2 and Sybase backups, disable NetWorker multiplexing.

3. Use the instructions in the appropriate topic to specify the number of sessions (parallelism) for a backup:

   Note: The backup parallelism must not be greater than the total number of device sessions available for the backup through the media pool configuration, as determined in step 1. Otherwise, the backup will become suspended because it is waiting for an available device to be mounted for the proper pool, or the backup will fail if not enough devices are available.

   - "Configuring DB2 parallel backups" on page 114
   - "Configuring Informix parallel backups" on page 114
   - "Configuring Lotus parallel backups" on page 115
   - "Configuring Oracle parallel backups" on page 115
   - "Configuring Sybase parallel (multistripe) backups" on page 115

Configuring DB2 parallel backups

Specify that the backup will use the number of sessions equal to the number of available backup devices:

- For a scheduled backup configured with the configuration wizard, specify the number of DB2 backup sessions on the corresponding wizard page.
- For a scheduled backup that uses an NMDA configuration file, set the number of sessions in the DB2_SESSIONS parameter.
- For a manual backup, specify the number of sessions with the **open sessions** option of the **db2** command on the operating system command line, for example:

  ```
  db2 backup db sample load /usr/lib/libnsrdb2.xx open # sessions
  options @/pathname/nmda_db2.cfg
  ```

  where:

  - *sample* is the name of the DB2 database.
  - *xx* is the extension for the operating system.
  - *#* is the number of sessions.
  - *pathname*/nmda_db2.cfg is the complete pathname of the NMDA configuration file for the DB2 backup.

Configuring Informix parallel backups

For both manual backups and scheduled backups, set the Informix parameter BAR_MAX_BACKUP in the $ONCONFIG file.

The Informix IDS documentation on the IBM website describes the $ONCONFIG file and parameter settings.

Configuring Lotus parallel backups

Specify that the backup will use the number of sessions equal to the number of available backup devices:

- For a scheduled backup configured with the configuration wizard, specify the number of Lotus backup sessions on the corresponding wizard page.
- For a manual backup or scheduled backup that uses an NMDA configuration file, set the number of sessions in the NSR_PARALLELISM parameter.

Configuring Oracle parallel backups

For a scheduled backup configured with the configuration wizard, specify the number of Oracle backup sessions (channels) on the corresponding wizard page.

For a scheduled backup that uses an NMDA configuration file, and for a manual backup, allocate multiple RMAN channels. The number of channels determines the backup parallelism.

Do not use NetWorker multiplexing with RMAN multiplexing. "Configuring parallel backups" on page 112 describes how to disable NetWorker multiplexing.

To guarantee that a backup does not use NetWorker multiplexing no matter what the storage device settings are on the server, you can set the NSR_NO_MULTIPLEX parameter, as described in "NSR_NO_MULTIPLEX" on page 389.

If you want to use NetWorker multiplexing in addition to RMAN multiplexing, run the **set parallelmediarestore off** command during the Oracle restore, to avoid restore performance degradation. For example:

```
set parallelmediarestore off;
run {
   allocate channel c1 type 'SBT_TAPE';
   restore database;
   release channel c1;
}
```

Configuring Sybase parallel (multistripe) backups

NMDA does not support multistripe backups for the backup of transaction logs (incremental backups). If you use the multistripe backup option for incremental backups, NMDA automatically converts the backup to a nonmultistripe (single-stripe) incremental backup.

Specify the backup to use the number of multistripe sessions:

- For a scheduled backup configured with the configuration wizard, specify the number of sessions on the corresponding wizard page.
- For a manual backup or scheduled backup that uses an NMDA configuration file, set the number of sessions in the NSR_PARALLELISM parameter.

DB2 backup considerations

The following DB2 operations require specific considerations:

- Automatic deletion of DB2 recovery objects
- Automatic backups of DB2 transaction logs

Configuring automatic deletion of DB2 recovery objects

When you set the DB2 database configuration parameter AUTO_DEL_REC_OBJ=ON, the DB2 database manager automatically performs these operations:

- Prunes the database history.
- Deletes the corresponding backup images, load copy images, and log files.

If you set the parameter AUTO_DEL_REC_OBJ=ON, the DB2 system might perform the maintenance operations as part of the backup operation, for example, after a successful full backup.

For successful deletion of the requested objects from NetWorker, ensure that the DB2 user that runs the backup has the required privileges for backup deletion on the NetWorker server. Table 4 on page 76 describes the required privileges.

You can also prune the history file and delete the backups manually. "Deleting DB2 backups" on page 164 provides details.

Configuring automatic backups of DB2 transaction logs

When you have configured the automatic backup of DB2 transaction logs, NMDA performs the log backups based on DB2 database policy settings. NMDA has no control of when the logs are backed up, how often, and so on. For successful log backups, ensure that a device is always available for the backups.

Complete the required steps to configure the automatic backup of DB2 transaction logs when they become full.

1. Create an NMDA configuration file for backing up the transaction logs only. For example, you can name this configuration file nmda_db2_tlogs.cfg.

   Appendix A, "NMDA Parameters and Configuration File," describes the parameters set in the configuration file:

   - The NSR_SERVER parameter is mandatory. Set this parameter to the hostname of the NetWorker server that will back up the logs.
   - If the configuration file is for a cluster environment, set the NSR_CLIENT parameter to the virtual cluster hostname.

   The following sample shows the configuration file content for transaction log backups:

   ```
   NSR_SERVER=TURBO
   NSR_LOG_VOLUME_POOL=DB2INST1_Logs
   ```

2. Configure the database with the command and options appropriate for the client operating system:

   - On UNIX:

     ```
     $ db2 update db cfg for sample using logarchmeth1
     VENDOR:/usr/lib/libnsrdb2.so logarchopt1
     @pathname/nmda_db2_tlogs.cfg
     ```

   - On Windows:

     ```
     $ db2 update db cfg for sample using logarchmeth1
     VENDOR:NetWorker_install_dir\nsr\bin\libnsrdb2.dll logarchopt1
     @pathname\nmda_db2_tlogs.cfg
     ```

     where:

     – *sample* is the name of the database to be backed up.

     – *pathname*/nmda_db2_tlogs.cfg is the complete pathname of the configuration file. Do not specify a relative pathname.

     – *NetWorker_install_dir* is the path on Windows systems that contains the NetWorker software.

3. When you complete steps 1 and 2, manually perform a full backup of the database, as described in "Performing DB2 manual backups with the db2 backup command" on page 148 or "Performing DB2 manual backups with the DB2 GUI" on page 149.

NOTICE

You must perform an initial full backup of the database.

Informix backup considerations

Automatic (continuous) backups of Informix logical logs require specific considerations.

Configuring automatic (continuous) backups of Informix logical logs

To configure NMDA to automatically back up the IDS logical logs when they become full, modify the Informix automatic log backup script, log_full.sh (UNIX) or log_full.bat (Windows), on the IDS host. Modify the script to include the following lines:

- On UNIX:

  ```
  NSR_LOG_VOLUME_POOL=NetWorker_pool_name
  NSR_SERVER=NetWorker_server_name
  export NSR_LOG_VOLUME_POOL
  export NSR_SERVER
  ```

- On Windows:

  ```
  set NSR_LOG_VOLUME_POOL=NetWorker_pool_name
  set NSR_SERVER=NetWorker_server_name
  ```

You can use the Informix ALARMPROGRAM configuration option to start the backups on demand when the logical logs fill.

After you successfully back up a log file, ON-Bar closes the file, frees the space used by the file, and opens a new file for transaction logging.

NMDA always performs logical log backups as level full (ON-Bar level 0) backups.

NOTICE

Dedicate a backup device for continuous log backups. The dedicated device on the backup server must be always available to receive logical log data.

Lotus backup considerations

The following Lotus operations require specific considerations:

- Lotus manual backups on Solaris or Linux
- Lotus database or directory link backups
- Lotus transaction log backups
- Lotus incremental backups with the comfort span option
- Partitioned Domino server backups
- Lotus DAOS backups
- Lotus restartable scheduled backups

Setting the environment for Lotus manual backups on Solaris or Linux

For Lotus manual backups with a Domino server on Solaris or Linux only, set the environment variable LD_LIBRARY_PATH:

- Set LD_LIBRARY_PATH to the complete pathname of the Lotus directory that contains the library files libnotes.so and libndgts.so.
- Set LD_LIBRARY_PATH in the same shell in which you perform the operation.

Considerations for Lotus database or directory link backups

To prevent NMDA from following the directory links and database links during Lotus backups, set the following parameter in the wizard or in the NMDA configuration file:

```
NSR_FOLLOW_LINKS = FALSE
```

If a database link, *link*.nsf, or directory link, *link*.dir, has a bad reference, the destination database or destination directory either does not exist or cannot be opened. NMDA cannot determine whether the file is a link or a database or directory. In this case, NMDA does not back up the link because NMDA does not know whether to back it up as a database, directory, or regular operating system file. The backup fails unless you set the parameter NSR_SKIPDBERRORS to TRUE in the NMDA configuration file.

Considerations for Lotus transaction log backups

"Lotus full and incremental backups" on page 40 and "Lotus transaction log backups" on page 40 describe the conditions for the backup of Lotus transaction logs.

During a full backup, NMDA does not back up transaction logs by default. You can set the NSR_BACKUP_LOGS_MODE parameter or the corresponding wizard option to back up the transaction logs and mark them reusable.

During an incremental backup, NMDA backs up the transaction logs unless you set the NSR_INCR_BACKUP_LOGS_MODE parameter or the corresponding wizard option to specify otherwise. By default, NMDA marks each successfully backed-up log as reusable to enable the Domino server to free space in the log repository.

You can configure a manual or scheduled backup of transaction logs only:

- For a manual backup of transaction logs, set NSR_BACKUP_LEVEL to the value txnlog or 1 in the NMDA configuration file.
- For a scheduled backup configured with the wizard, select **Back up archived transaction logs only** on the **Select the Backup Objects** page of the wizard.
- For a scheduled backup configured without the wizard, set backup level 1 in the Schedule resource or Group resource to specify a transaction log backup. You can select the schedule in the Group resource that contains the Client resource configured for the NMDA Lotus transaction log backup.

  The *EMC NetWorker Administration Guide* describes how to configure Schedule, Group, and Client resources for different backup scenarios and the precedence rules.

Note: If NSR_BACKUP_LOTUS_DIR is FALSE or not set and NSR_BACKUP_PATHS is not set, NMDA backs up the transaction logs, no matter what the settings are for other parameters.

You cannot perform the transaction logs only backup with a Lotus DAOS backup. For example, if you have configured a transaction logs only backup, a DAOS backup does not run if the DAOS backup is also configured through settings in the LOTUS_DAOS {} section of the configuration file.

For any backup that includes transaction logs, you can set the NSR_RETAIN_NUM_LOGS parameter or the corresponding wizard option to back up a specified number of logs without marking them reusable. The recovery process can be faster when you retain logs on the Domino system because you do not need to restore the logs from NetWorker. Example 6 on page 119 shows how the NSR_RETAIN_NUM_LOGS setting can affect log backups.

You can set the NSR_MAX_TXN_LOGS parameter or corresponding wizard option to specify the number of logs to store per save set.

"NMDA Lotus parameters" on page 371 describes the parameters.

Example 6 Using NSR_RETAIN_NUM_LOGS to control reusable logs

If you set NSR_RETAIN_NUM_LOGS = 5 and there are 20 archived transaction logs to be backed up, NMDA performs the following actions:

1. Backs up the first 15 logs and marks them reusable.
2. Backs up the last 5 logs without marking them reusable.

During a subsequent backup, after the system generates an additional 11 logs, NMDA performs the following actions:

1. Recognizes that the first (oldest) 5 logs have already been backed up, and does not back them up again but marks them reusable.
2. Backs up 6 of the new logs and marks them reusable.
3. Backs up the last 5 logs without marking them reusable.

Lotus incremental backups with the comfort span option

NMDA supports the NSR_COMFORT_SPAN parameter for incremental backups only when you enable the Domino server for transaction logging in archive mode. The parameter specifies the acceptable amount of logs in kilobytes that can be applied to a database during the recovery. If that amount is exceeded at backup time, NMDA backs up the database file as a full backup and also backs up the logs. A full backup reduces future recovery time. Fewer transaction logs are required to recover the logged database.

Note: When you enable the Domino server for transaction logging in archive mode, the default NMDA behavior during an incremental backup is to back up the transaction logs only.

For example, the NMDA configuration contains the following parameter settings for a Lotus incremental database backup with the comfort span option:

```
NSR_BACKUP_LEVEL = incr
NSR_COMFORT_SPAN = 196608
```

If NMDA determines that more than 196608 KB of logs need to be applied to recover the specified database, NMDA backs up the database in addition to backing up the logs.

Configuring partitioned Domino server backups

For backups of partitioned Domino servers, configure the backup for each partition separately. The configuration is the same as for the backup of a regular Domino server except that the NSR_LOTUS_DATA_DIR parameter must specify which partitioned Domino server to back up. For manual backups and scheduled backups configured without the wizard, set NSR_LOTUS_DATA_DIR in the configuration file. The NSR_LOTUS_DATA_DIR setting is mandatory in the configuration wizard.

Note: The NSR_LOTUS_DATA_DIR parameter does not specify the data to be backed up. To specify the backup data, set NSR_BACKUP_LOTUS_DIR or NSR_BACKUP_PATHS in the configuration file. Alternatively, set the equivalents in the NetWorker User for Lotus GUI or in the configuration wizard (select the directories or files to back up).

For a scheduled backup configured without the wizard, in the Save Set attribute of the Client resource, specify a descriptive name for the backup save set stored on the media. The Save Set value must start with the NOTES: prefix. For example, you could specify the following value for the Save Set attribute:

- When you back up the entire data directory for the partition:

  ```
  NOTES:partition1_/disk2/notesdata1
  ```

- When you back up a database named db.nsf for the partition:

  ```
  NOTES:partition1_/disk2/notesdata1/db.nsf
  ```

Configuring Lotus DAOS backups

You can use the **nsrdaadmin -W** command to convert an NMDA client-side DAOS configuration to the NMDA server-side (wizard) DAOS configuration. "Converting client-side to server-side configurations" on page 79 provides details.

Note: After the configuration conversion, the DAOS backups will run successfully. However, when you use the wizard for the first time to modify the DAOS configuration, you must type the path for **DAOS base directory** on the wizard page **Specify the Lotus Domino/Notes Information**.

In an *integrated* DAOS backup, NMDA backs up the Domino data first and then backs up the DAOS repository or part of it during the same backup session. In a *stand-alone* DAOS backup, NMDA backs up only the DAOS files without the Domino database files. IBM recommends an integrated DAOS backup as preferable to a stand-alone backup.

To configure an integrated backup of Domino database files and DAOS (NLO files):

1. Review "Best practices for Lotus DAOS backups" on page 121.
2. Complete "Configuring integrated Lotus DAOS backups" on page 122.

Best practices for Lotus DAOS backups

Review the IBM documentation for details about the required DAOS configuration and backup practices.

The following list provides specific best practices for reference purposes:

- Set the DAOS deferred deletion interval to a period longer than the backup cycle, which is the period between full backups.

  Note: The Domino server deletes an attachment in a DAOS directory only when the last database that references it is deleted and after a user-defined delay time called the *deferred deletion interval*.

- Do not prune the NLO files from the DAOS repository before you have backed up the files. This practice ensures that the NLO files will be recoverable.
- Do not include the daos.cfg and daoscat.nsf files (stored in the Lotus data directory) in Lotus DAOS backups.

The IBM documentation describes the DAOS setup procedures.

Note: A DAOS backup also backs up transaction logs if the backup level is incremental or you set the NSR_BACKUP_LOGS_MODE parameter for a full backup.

Configuring stand-alone Lotus DAOS backups

To configure a stand-alone DAOS backup that backs up only DAOS files without backing up the Domino database files:

1. Follow the steps for configuring a regular Lotus backup.
2. Specify the following parameter settings in the LOTUS{} section of the configuration file:

```
NSR_BACKUP_PATHS=DAOS_directory
NSR_BACKUP_ALL_EXTENSIONS=TRUE
```

Note: The NMDA wizard does not support the configuration of a stand-alone DAOS backup.
If you want to use the NetWorker User for Lotus GUI for the backup, you must not set the NSR_BACKUP_PATHS parameter. Instead, select the DAOS directory through the GUI.

Configuring integrated Lotus DAOS backups

A manual or scheduled integrated backup backs up the Domino data first and then backs up the DAOS repository or part of it during the same backup session.

Note:
- The DAOS backup runs only if the Domino data backup succeeds.
- The NetWorker User for Lotus GUI does not support an integrated backup. If you want to use the GUI, you must run a stand-alone database backup first, followed by a stand-alone DAOS backup.

To configure an integrated backup, use one of these methods:

- "Wizard method (server-side configuration)" on page 122
- "Nonwizard method (client-side configuration)" on page 123

Wizard method (server-side configuration)

To configure a scheduled integrated backup with the configuration wizard:

1. On the wizard page **Specify the Lotus Domino/Notes Information**, select **Perform DAOS backup** and specify the root path of the DAOS directory in **DAOS base directory**.

   Note: Set the **DAOS base directory** field even if you will back up a subset of the DAOS repository.

2. On the page **Select the Backup Objects**, select at least one Lotus Domino data file to be backed up.

3. On each wizard page that appears, specify the options and values required to configure the Domino server backup first, followed by the DAOS backup configuration. Note the following recommendations:

   - On the page **Specify the Database and File Options**, if the DAOS directory is in the Domino data path to be backed up, add the DAOS directory to the exclude path list. This setting excludes the DAOS directory from the Domino data backup.
   - On the page **Select the Domino Attachment Object Service (DAOS) Additional Parameters**, you might want to set the NSR_SAVESET_BROWSE and NSR_SAVESET_RETENTION parameters. You can use these parameters to keep the DAOS backups for a longer time than the corresponding Domino data backup.

Nonwizard method (client-side configuration)

To configure a manual or scheduled integrated backup without the configuration wizard, set the required parameters in a single configuration file:

1. Set the parameters required for the Domino data backup in the LOTUS{} section of the configuration file. For example, the following parameter in the LOTUS{} section specifies the backup of the Lotus data directory:

   ```
   NSR_BACKUP_LOTUS_DIR=TRUE
   ```

   The LOTUS{} section must include either NSR_BACKUP_LOTUS_DIR or NSR_BACKUP_PATHS, but not both, to specify at least one Domino data file for the backup.

2. If the DAOS directory is in the Domino data path to be backed up, add the DAOS directory to the exclude list in the NSR_EXCLUDE_LIST parameter in the LOTUS{} section. This setting excludes the DAOS directory from the Domino data backup.

3. Set the following mandatory parameters in the LOTUS_DAOS{} section of the configuration file:

   ```
   NSR_BACKUP_PATHS=DAOS_base_dirpath_or_list_of_subdirectory_paths
   NSR_BACKUP_ALL_EXTENSIONS=TRUE
   ```

   The parameter settings in the LOTUS_DAOS{} section apply only to the DAOS backup, not to the Domino data backup.

   Note: To back up the daos.cfg file in the Lotus data directory and the parameter settings in the LOTUS{} section, do not include the backup of the file. You must add the file path to the NSR_BACKUP_PATHS setting in the LOTUS_DAOS{} section.

NOTICE

In the configuration file, the LOTUS_DAOS{} section must appear after the LOTUS{} section, to ensure that NMDA backs up the DAOS files after the Domino database files.

The LOTUS{} and LOTUS_DAOS{} sections can include different parameter settings that specify different backups levels and so on for the Domino data and DAOS backups.

The DAOS backup inherits most of the parameters from the LOTUS{} section. If NMDA requires the same parameter for both backups (for example, Notes_ExecDirectory), then you can set the parameter in the LOTUS{} section only. For the DAOS backup, a parameter setting in the LOTUS_DAOS{} section overrides the setting of the same parameter in the LOTUS{} section.

The following parameters are not inherited from the LOTUS{} section, and apply only to the section in which they are specified:

- NSR_BACKUP_PATHS
- NSR_EXCLUDE_FILE
- NSR_EXCLUDE_LIST

The following parameters do not apply to DAOS backups:

- NSR_BACKUP_LOTUS_DIR
- NSR_COMFORT_SPAN
- NSR_FOLLOW_LINKS
- NSR_SKIPDBERRORS
- PRECMD

After the DAOS backup ends, NMDA updates the catalog file and runs the post-command. As a result, if you set NSR_CATALOGFILE and POSTCMD to different values in the LOTUS_DAOS{} section and LOTUS{} sections, then NMDA uses the values from the LOTUS_DAOS{} section.

You can set the following parameter values in the LOTUS_DAOS{} section to different values than in the LOTUS{} section:

- NSR_BACKUP_LEVEL=incr, in case NMDA backs up the databases as full.
- NSR_SAVESET_BROWSE and NSR_SAVESET_RETENTION, to keep the DAOS backups for a longer time.
- NSR_SAVESET_NAME, to specify a different name for DAOS save sets that differentiates them from Domino data save sets. For example, set NSR_SAVESET_NAME=NOTES_DAOS.

  If you do not set NSR_SAVESET_NAME, the DAOS backup save set has the same name as the Domino data save set from the same integrated backup, such as NOTES_*number*. A save set with the same name appears twice in the scheduled backup details in the NMC interface.

Appendix A, “NMDA Parameters and Configuration File,” describes the NMDA configuration file and NMDA Lotus parameters.

Configuring Lotus restartable scheduled backups

You can configure a Lotus restartable scheduled backup either through the wizard or the nonwizard method.

For either method, use the NMC program to optionally set the Client Retries attribute in the Group resource to a value greater than zero. This setting enables the automatic restart of a

failed backup. Otherwise, you must restart the backup manually through the group restart option.

- To configure a restartable scheduled backup with the wizard, complete "Configuring scheduled backups with the wizard" on page 82 and select the Checkpoint Enabled checkbox on the proper wizard page.
- To configure a restartable scheduled backup without the wizard, complete "Configuring scheduled backups without the wizard" on page 85 and select the Checkpoint Enabled checkbox in the NMC Client resource.

Note: NMDA ignores the Checkpoint Granularity attribute during a checkpoint restartable backup. Changing the attribute setting has no impact on the NMDA Lotus backup.

"Lotus restartable backup information in NetWorker indexes" on page 161 describes the information about Lotus restartable backups that the NetWorker server maintains in the NetWorker indexes.

MySQL backup considerations

The following MySQL backup topics require specific considerations:

- Configuration parameter settings
- MySQL backup granularity
- MySQL incremental backups
- MySQL binary log backups
- MySQL backup privileges

Configuration parameter settings

You must set the required configuration parameters to enable an NMDA MySQL backup:

- You must have a valid MySQL configuration file, my.cnf, which is used to connect to the MySQL instance. The MySQL configuration file is different from the NMDA configuration file.
- For a manual backup, set backup parameters in the NMDA configuration file.
- For a scheduled backup, set backup parameters in the NMDA configuration file for a client-side configuration, or specify settings in the configuration wizard for a server-side configuration.

You can set MEB backup parameters in the MYSQL configuration file, such as the MySQL user name and password and the list of databases to be backed up. However, set backup-specific parameters in the NMDA configuration file or wizard instead, if possible. Although NMDA supports the MEB backup parameter settings in the MySQL configuration file, you can simplify the configuration by using the NMDA parameters.

If you set the same parameters in both the MySQL configuration file and the NMDA configuration file or wizard, the NMDA parameters take precedence. For example, the MYSQL_DATABASES setting in the NMDA configuration file takes precedence over the --databases setting in the MySQL configuration file.

Appendix A, "NMDA Parameters and Configuration File," describes all the supported configuration parameters. Online help in the wizard describes the wizard settings.

Considerations for MySQL whole and partial backups

By default, NMDA performs a whole instance backup of a MySQL server instance.

To configure a partial backup, you can set the additional required parameters in the NMDA MySQL configuration file for a manual backup or for a scheduled backup configured without the wizard:

- With a MyISAM storage engine, you can set MYSQL_DATABASES to back up any combination of specified databases and specified tables within different databases.
- With an InnoDB storage engine, you can set the configuration parameters for the particular type of partial backup as described in the following table.

  Note: A partial backup of specified tables with an InnoDB storage engine always includes the system tablespace and all the tables within it.

For a scheduled partial backup configured with the wizard, you must specify the corresponding configuration settings in the wizard.

Table 13 Configuration parameters for InnoDB partial backups

Types of InnoDB partial backups	Configuration parameters to enable the backups
Backup of specified databases	MYSQL_DATABASES = “db1[.tbl1] db2[.tbl2] db3[.tbl3] ...” where [ ] contains an optional value
Backup of specified tables when the file-per-table option is enabled	MYSQL_INCLUDE =” *‹regular expression matching the per-table data file names›*”
Backup of only the tables in a specified instance	MYSQL_ONLY_INNODB_OPTIONS = NO_FRM
With MEB 3.7 or later, backup of all the tables and their associated .frm files in a specified instance	MYSQL_ONLY_INNODB_OPTIONS = WITH_FRM_ALL
With MEB 3.7 or later, backup of specified tables and their associated .frm files when MYSQL_INCLUDE is set and the file-per-table option is enabled	MYSQL_ONLY_INNODB_OPTIONS = WITH_FRM_RELATED

For example, the NMDA MySQL configuration file for a partial backup contains the following parameter settings:

```
MYSQL_INCLUDE="innodb-sales*"
MYSQL_ONLY_INNODB_OPTIONS=NO_FRM
```

Due to an MEB limitation, if you set either of the following parameters for an InnoDB backup, you must run the InnoDB backup as an OS user with write permissions to the parent directory of the MySQL data directory:

- MYSQL_ONLY_INNODB_OPTIONS = WITH_FRM_ALL
- MYSQL_ONLY_INNODB_OPTIONS = WITH_FRM_RELATED

Considerations for MySQL incremental backups

With MEB 3.7 or later, you can set MYSQL_INCR_OPTIONS=REDO_LOG_ONLY to configure a redo log only incremental backup. A redo log only incremental backup is a differential incremental backup that backs up the redo log changes since the last full or incremental backup (differential or cumulative).

Considerations for MySQL binary log backups

To enable a MySQL binary log backup, ensure that binary logging is enabled and the MySQL instance is online.

To configure a binary log backup as a manual backup or a client-side scheduled backup (configured without the wizard), you can set the MYSQL_LOG_OPTIONS parameter:

- Set MYSQL_LOG_OPTIONS = LOGS_ONLY_BACKUP to back up only the binary logs for the instance.
- Set MYSQL_LOG_OPTIONS = INCLUDE_LOGS to back up the binary logs after a whole instance backup.

If you set the parameter to both values as follows, NMDA applies only the last value (in this case, INCLUDE_LOGS) to the parameter:

MYSQL_LOG_OPTIONS = LOGS_ONLY_BACKUP, INCLUDE_LOGS

You can optionally include the PURGE_LOGS value in the MYSQL_LOG_OPTIONS setting to delete the binary logs from the disk after the log backup. For example:

MYSQL_LOG_OPTIONS = INCLUDE_LOGS, PURGE_LOGS

For a scheduled log backup configured with the wizard, you must specify the corresponding configuration settings in the wizard.

Note: You cannot set MYSQL_LOG_OPTIONS for a partial instance backup, for example, when you set MYSQL_DATABASES or MYSQL_INCLUDE. You can set MYSQL_LOG_OPTIONS only for a stand-alone log backup or a whole instance backup with a log backup.

MySQL backup privileges

Before you run an NMDA MySQL backup, ensure that the backup user has the following MySQL privileges:

- RELOAD on all databases and tables.
- CREATE, INSERT, and DROP on the tables mysql.ibbackup_binlog_marker, mysql.backup_progress, and mysql.backup_history.
- SUPER, used to optimize locking and minimize disruption to database processing.
- CREATE TEMPORARY TABLES for the mysql database.
- REPLICATION CLIENT to retrieve the binlog position, which is stored with the backup.

The **mysqlbackup** information in the MySQL Enterprise Backup User's Guide describes the most up-to-date MySQL privileges required for backups.

Configuring MySQL 5.6 features

You must complete the required configurations if you use specific features of MySQL release 5.6 for NMDA operations:

- You can store InnoDB tables in a specified directory outside the MySQL data directory. When you restore an NMDA backup of InnoDB tables stored outside the MySQL data directory, if any .ibd file from the backup already exists on the target system, you must set the **--force** option in the MYSQL_MEB_OPTIONS parameter.
- You can store the InnoDB undo logs or rollback segments in one or more separate tablespaces outside of the system tablespace.

  For an NMDA backup of the InnoDB undo logs, you must set the following parameters in the MySQL configuration file: **innodb_undo_directory**, **innodb_undo_logs**, **innodb_undo_tablespaces**. In the NMDA configuration file, you must also set the MYSQL_CFG_FILE parameter to the pathname of the MySQL configuration file.
- For an NMDA restore, you can specify the checksum algorithm of InnoDB tablespaces in the target database by setting the **innodb_checksum_algorithm** value in the MySQL configuration file.
- For an NMDA restore with MySQL 5.6, MEB requires the **innodb_data_file_path** parameter to be set in the MySQL configuration to enable a copy back operation. Without the parameter setting, the operation fails.

"NMDA MySQL parameters" on page 380 provides details about setting the parameters for MySQL backup and restore operations.

Oracle backup considerations

For Oracle scheduled backups, you must set certain NMDA parameters in the configuration file. You must set all the other NMDA parameters for Oracle backups in the RMAN backup script as described in "NMDA Oracle parameters" on page 387.

You can store RMAN backup scripts as flat ASCII files. Alternatively, if you use a Recovery Catalog, you can store backup scripts in the Recovery Catalog database. The Oracle backup and recovery documentation describes how to store backup scripts in the Recovery Catalog database.

If you use automatic channel allocation and persistent settings, you can run the **backup** command as a stand-alone command. "Configuring automatic channel allocation" on page 132 provides details.

To perform a backup to NetWorker by using NMDA, set the **type** option in the **allocate channel** command to **SBT** or **SBT_TAPE**.

In the RMAN backup script in Example 7 on page 129, the **format** string **FULL_%d_%U** specifies the name of each backup piece. This name can be anything, provided that each backup piece has a unique name on the NetWorker server. You can use substitution variables, such as **%d** and **%U**, to guarantee unique names:

- **%d** specifies the name of the database.
- **%U** specifies a unique Oracle system-generated file name.

A **format** string such as **FULL** or **FULL_%d** will not generate unique names. Similarly, the **format** string **FULL_%U** will not generate unique names for two databases that you back up to the same NetWorker server.

NOTICE

If a backup piece name is not unique, the Oracle backup fails.

RMAN scripts for manual backups

For Oracle manual backups, you must set all the parameters in the RMAN backup script. Use the **send** command to set any parameter in the RMAN script if you do not use the backup copies feature. "The send command" on page 398 and "Backup copies" on page 49 provide details.

Example 7 on page 129 provides a sample RMAN script for a manual backup.

Example 7 RMAN script for a manual backup

The following RMAN script is for a manual backup of an entire Oracle database to the volume pool MondayFulls of the (remote) NetWorker server mars.emc.com:

```
run {
   allocate channel t1 type 'SBT_TAPE';
   allocate channel t2 type 'SBT_TAPE';

   send 'NSR_ENV=(NSR_SERVER=mars.emc.com,
   NSR_DATA_VOLUME_POOL=MondayFulls)';
   backup full filesperset 4 format 'FULL_%d_%U' (database);

   release channel t1;
   release channel t2;
}
```

During an Oracle manual backup, the prefix RMAN: automatically precedes the backup piece name in the NetWorker media database. For example, if the backup piece name specified in the RMAN script is accounts_data_file, the manual backup records the save set name as RMAN:accounts_data_file in the media database. The **mminfo** command output includes the save set name in this form.

The following sources provide more information:

- The Oracle backup and recovery documentation describes how to write RMAN scripts.
- The Oracle Enterprise Manager documentation describes how to use the Oracle Enterprise Manager Backup Wizard to generate RMAN scripts.
- Appendix B, "Oracle RMAN Commands," describes RMAN commands.
- "Verifying backup information in NetWorker indexes" on page 159 describes the backup information stored in the NetWorker indexes.

RMAN scripts for scheduled backups

When you configure a scheduled backup without the wizard (client-side configuration), you must set certain required parameters, such as ORACLE_HOME, in the configuration file as described in "NMDA Oracle parameters" on page 387. For all the other parameters that you must set in the RMAN backup script, use the **send** command if you are not using automatic channels. "The send command" on page 398 provides details.

"Common NMDA parameters" on page 358 and "NMDA Oracle parameters" on page 387 include lists of the common parameters and Oracle parameters, respectively.

NOTICE

For scheduled backups, do not include **send** as part of the **allocate channel** command. The **send** command must be separate.

For example, NMDA does not support the following command for a scheduled backup:

```
allocate channel t1 type 'SBT_TAPE' send
   'NSR_ENV=(NSR_SERVER=mars.emc.com)';
```

The following commands are correct for a scheduled backup:

```
allocate channel t1 type 'SBT_TAPE';
send channel t1 'NSR_ENV=(NSR_SERVER=mars.emc.com)';
```

"Configuring automatic channel allocation" on page 132 provides information about automatic channel allocation.

Example 8 on page 130 provides a sample RMAN script for a scheduled backup.

Example 8 RMAN script for a scheduled backup

The following RMAN script is for a scheduled backup of an entire Oracle database to the volume pool MondayFulls. The Recovery Catalog is used in this case:

```
connect target target_user/target_passwd@target_Netservicename;
connect rcvcat rcvcat_user/rcvcat_passwd@rcvcat_Netservicename;

run {
   set command id to 'xxx';

   allocate channel t1 type 'SBT_TAPE';
   allocate channel t2 type 'SBT_TAPE';

   send 'NSR_ENV=(NSR_DATA_VOLUME_POOL=MondayFulls)';
   backup full filesperset 4 format 'FULL_%d_%U' (database);

   release channel t1;
   release channel t2;
}
```

If you use automatic channel allocation and persistent settings, you must still create a scheduled RMAN backup script to contain the following commands:

- **connect target**
- **connect rcvcat** (if using a Recovery Catalog)
- **backup**

The command **connect target** *target_user*/*target_passwd*@*target_Netservicename* is mandatory in each RMAN script for a scheduled backup that does not use Oracle operating system authentication (ORACLE_USER). This command establishes the proper connection to the target database.

Specify the correct values in the **connect target** command:

- *target_user* is the user with SYSDBA privileges for the target database.
- *target_passwd* is the password of the *target_user* (for connecting as SYSDBA), specified in the target database's **orapwd** file.
- *target_Netservicename* is the Net service name of the target database. This name is mandatory in the **connect target** command.

You must use a password file for the target database. The Oracle documentation provides details.

Note: Because each scheduled backup RMAN script requires a **connect target** command, each Oracle instance requires a separate scheduled backup RMAN script, unless a backup uses Oracle operating system authentication (ORACLE_USER).

The command **connect rcvcat** *rcvcat_user*/*rcvcat_passwd*@*rcvcat_Netservicename* is mandatory if you use the Recovery Catalog for the scheduled Oracle backup. This command establishes the proper connection to the Recovery Catalog database.

Specify the correct values in the **connect rcvcat** command:

- *rcvcat_user* is the owner of the Recovery Catalog database.
- *rcvcat_passwd* is the password of the *rcvcat_user*.
- *rcvcat_Netservicename* is the Net service name of the Recovery Catalog database.

To enable cancellation of the scheduled backup, the scheduled Oracle backup script must include **set command id to '***xxx***'** (where *xxx* can be any string of characters enclosed in single quotes). "Canceling scheduled backups" on page 146 describes how to cancel a scheduled backup.

The remainder of the scheduled backup script in Example 8 on page 130, starting with the first **allocate channel** command, is similar to the manual backup script in Example 7 on page 129 except the scheduled backup script excludes the NSR_SERVER parameter.

NOTICE

Do not set the parameter NSR_SERVER or NSR_GROUP in a scheduled RMAN backup script. NMDA sets these two parameters to the values specified in the Client resource for the Oracle scheduled backup, and you cannot override these values. However, you can set NSR_SERVER and NSR_GROUP in the scheduled backup script when you manually test-run the script as described in "Testing RMAN scripts for scheduled backups" on page 132.

You must store each scheduled backup RMAN script as a text file. The database administrator must give minimal permissions to the scheduled backup script file. This way, unauthorized users cannot see the sensitive user IDs and passwords of the target and Recovery Catalog databases.

A single Oracle instance can use multiple RMAN scripts, for example, to perform tablespace-level, file-level, full, or incremental backups, and so on. The database administrator might place the two common **connect** commands in a single file and then run those two **connect** commands in all RMAN scripts by using the **@** command.

Testing RMAN scripts for scheduled backups

When you create an RMAN script, test the script before using it for scheduled backups.

To test the RMAN script, type the following command:

```
rman cmdfile 'script_name' send
   '"NSR_ENV=(NSR_SERVER=NetWorker_server_name,
   NSR_GROUP=group_name)"'
```

where:

- *script_name* is the RMAN script file pathname.
- *NetWorker_server_name* is the name of the server that starts the backup.
- *group_name* is the name of the scheduled backup group as specified in the Client resource.

Configuring automatic channel allocation

NOTICE

Manual and automatic channels are mutually exclusive. You cannot mix manual and automatic channels in an RMAN session.
The format of an automatic channel name of the device type for NMDA backups and restores is ORA_SBT_*n* or ORA_SBT_TAPE_*n*, where *n* is the channel number. Do not use this name format for manual channel allocation for NMDA. Otherwise, RMAN reports an error.
With automatic channel allocation, specification of the **send** command before the **backup** or **restore** command causes the following error:

```
RMAN-06422: no channels found for SEND command
```

You must use the **configure channel...parms...** command to set the NMDA parameters for automatic channels for an NMDA backup. Do not use the **send** command or option to set the NMDA parameters for automatic channels if you plan to use scheduled backups.

The following tables lists all the NMDA parameters and their requirements for Oracle operations:

- Common NMDA parameters: Table 33 on page 358
- NMDA Oracle parameters: Table 38 on page 388

Example 9 Using the configure channel command with parms option for automatic channels

You can set an NMDA parameter to the same value for all automatic channels by typing the following **configure channel** command:

- With Oracle version 11gR1 or earlier:

```
configure channel device type 'sbt_tape' parms
'ENV=(NSR_CLIENT=mars)'
```

- With Oracle version 11gR2 or later:

```
configure channel device type 'sbt_tape' parms
'SBT_PARMS=(NSR_CLIENT=mars)'
```

This command sets the NSR_CLIENT parameter to the value mars for all the automatic channels.

Example 10 Specifying parameter values per automatic channel

You can set specific NMDA parameter values for different channels by typing the **configure channel** *n* **device type...parms...** command, where *n* represents a channel number. For example, you can set the NSR_DATA_VOLUME_POOL parameter separately for each channel.

You specify a NetWorker data volume pool for the second automatic channel by typing the following **configure channel** command:

- With Oracle version 11gR1 or earlier:

```
configure channel 2 device type 'sbt_tape' parms
'ENV=(NSR_DATA_VOLUME_POOL=Oracle2)'
```

- With Oracle version 11gR2 or later:

```
configure channel 2 device type 'sbt_tape' parms
'SBT_PARMS=(NSR_DATA_VOLUME_POOL=Oracle2)'
```

Creating Oracle backup copies

Due to Oracle limitations, the backup copies feature is supported with manual backups only. For Oracle backup copies, set the NMDA parameters with the **parms** option, not with the **send** command or option.

To create multiple copies of manual backups (up to four copies):

1. Use one of the following RMAN commands:
   - The **configure...backup copies for device type sbt_tape to...** command specifies persistent settings for duplexing backups through NMDA.

     For example, specify persistent settings for duplex copies of datafiles and archived redo logs (respectively) in NMDA backups with the following types of **configure** commands:

     ```
     configure datafile backup copies for device type 'sbt_tape' to 2
     configure archivelog backup copies for device type 'sbt_tape' to 2
     ```

- The **backup** command with the **copies** option applies to objects within the **backup** command. The **backup...copies** setting takes precedence over the persistent settings in the **configure...backup copies** command.
- The **set backup copies** command applies to all backup objects in the same **run** job.

2. Define a separate NetWorker pool for each copy and set the following parameters with the **parms** option:
   - NSR_DATA_VOLUME_POOL (if **copies** is set to 1)
   - NSR_DATA_VOLUME_POOL1 (if **copies** is set to 2)
   - NSR_DATA_VOLUME_POOL2 (if **copies** is set to 3)
   - NSR_DATA_VOLUME_POOL3 (if **copies** is set to 4)

   Appendix A, "NMDA Parameters and Configuration File," describes these parameters.

Example 11 Using the set backup copies command in the RMAN script

The following RMAN script uses the **set backup copies** command to generate the backup copies. The script sets the parameters with the **parms** option, as required:

Note: With Oracle version 11gR2 or later, use **parms 'SBT_PARMS=(...)'** instead of **parms 'ENV=(...)'**. "Precedence rules" on page 402 provides details.

```
run {
   set backup copies 4;

   allocate channel ch1 parms 'ENV=(NSR_SERVER=server_name,
   NSR_DATA_VOLUME_POOL=nmda1, NSR_DATA_VOLUME_POOL1=nmda2,
   NSR_DATA_VOLUME_POOL2=nmda3, NSR_DATA_VOLUME_POOL3=nmda4)';

   backup format '%d_%U'
   tag tag_name
   (tablespace 'SYSTEM' );

   release channel ch1;
}
```

Example 12 Using automatic channels for backup copies

The following **configure** commands configure the RMAN automatic channels. You can include the **configure** commands in the RMAN script. The **configure...backup copies** command generates the backup copies. The command that sets the parameters uses the **parms** option, as required:

Note: With Oracle version 11gR2 or later, use **parms 'SBT_PARMS=(...)'** instead of **parms 'ENV=(...)'**. "Precedence rules" on page 402 provides details.

```
configure default device type to 'sbt_tape';
configure datafile backup copies for device type 'sbt_tape' to 4;
configure channel device type 'sbt_tape' parms
   'ENV=(NSR_SERVER=server_name, NSR_DATA_VOLUME_POOL=nmda1,
   NSR_DATA_VOLUME_POOL1=nmda2, NSR_DATA_VOLUME_POOL2=nmda3,
   NSR_DATA_VOLUME_POOL3=nmda4)';
```

The RMAN script invoked for the manual backup is as follows:

```
connect target sys/oracle@test;
run {
   backup format '%d_%U'
   tag tag_name
   (tablespace 'SYSTEM');
}
```

Configuring save set bundling for scheduled Oracle backups

You can use either the configuration wizard or the nonwizard method to configure save set bundling for scheduled Oracle backups:

- If you use the configuration wizard:
  - Set NSR_BUNDLING to TRUE in the Advanced Environment Options table on the corresponding wizard page.
  - Ensure that the user group privileges for the root user or administrative user on the NMDA client include the privileges required for save set bundling, as described in Table 4 on page 76. Ensure that you have configured the corresponding User Group resource on the NetWorker server, as described in "NetWorker User Group resource" on page 75.
- If you use the nonwizard configuration method:
  1. Set the NSR_BUNDLING parameter by typing the following command:

     nsroraadmin -r add NSR_BUNDLING enabled

     The default value of the NSR_BUNDLING parameter is "disabled".

     "Configuring the NWORA resource file with the nsroraadmin program" on page 329 describes the **nsroraadmin** command. The command sets the parameter value in the NWORA resource file, as described in "NWORA resource file" on page 324.
  2. Ensure that you have configured the NMDA scheduled backups according to the "Configuring NMDA backups" on page 70.
  3. Ensure that the user group privileges for the root user or administrative user on the NMDA client include the privileges required for save set bundling, as described in Table 4 on page 76. Ensure that you have configured the corresponding User Group resource on the NetWorker server, as described in "NetWorker User Group resource" on page 75.
  4. If you do not include the username and password in the RMAN script (for example, you include the connection strings as a command file in the RMAN script, such as @*connection_file*), ensure that you meet these requirements:
     - The ORACLE_SID parameter is set in the NMDA configuration file, as described in "NMDA Oracle parameters" on page 387.
     - You have created an NWORA SID resource with the NSR_ORACLE_CONNECT_FILE parameter setting in the NWORA resource file (nwora.res) for the ORACLE_SID, as described in "NWORA SID resources" on page 327.

NMDA cannot retrieve the connection strings from the RMAN script when you include the connection strings as a command file in the script. In this case, NMDA must retrieve the connection strings from the connection file specified by the parameter in the NWORA resource file.

5. In a RAC system, ensure that you allocate all the channels on the same NMDA client node where the backup is initiated. Save set bundling does not support load balancing across different RAC nodes.

To disable save set bundling, set the NSR_BUNDLING parameter value to disabled by typing the following command:

```
nsroraadmin -r update NSR_BUNDLING disabled
```

Configuring policy uniformity for scheduled Oracle backups

If you enable save set bundling as described in "Configuring save set bundling for scheduled Oracle backups" on page 135, NMDA automatically enables policy uniformity.

If you are not using save set bundling, you can use either the configuration wizard or the nonwizard method to enable policy uniformity for scheduled Oracle backups:

- If using the configuration wizard, set NSR_INCR_EXPIRATION to TRUE in the Advanced Environment Options table on the corresponding wizard page.
- If using the nonwizard configuration method:
  - Set the NSR_INCR_EXPIRATION parameter value to enabled by typing the following command:

    ```
    nsroraadmin -r add NSR_INCR_EXPIRATION enabled
    ```

    The default value of the NSR_INCR_EXPIRATION parameter is "disabled".

    "Configuring the NWORA resource file with the nsroraadmin program" on page 329 describes the **nsroraadmin** command. The command sets the parameter value in the NWORA resource file, as described in "NWORA resource file" on page 324.
  - Complete step 2 to step 4 under "Configuring save set bundling for scheduled Oracle backups" on page 135.
  - In a RAC system, ensure that you allocate all channels on the same NMDA client node where the backup is initiated. Policy uniformity does not support load balancing across different RAC nodes.

To disable policy uniformity, set the NSR_INCR_EXPIRATION parameter value to "disabled" by typing the following command:

```
nsroraadmin -r update NSR_INCR_EXPIRATION disabled
```

Configuring operations in an Oracle Data Guard environment

To configure NMDA backups and restores in an Oracle Data Guard environment:

1. Follow the instructions in Oracle documentation about how to set the required RMAN configurations, for example, to use a Recovery Catalog and the DB_UNIQUE_NAME parameter.
2. Install and configure the NMDA and NetWorker client software on the primary database host, and on each physical standby database host included in the backups and restores.
3. Configure a Client resource on the NetWorker server for the primary database host and each physical standby database host involved in the backups and restores. In the Client resource of the primary database host, add the hostname of the physical standby host in the Remote Access attribute if you set NSR_CLIENT to the primary database hostname in step 4.
4. Create an RMAN script for the primary database and the standby database. Set the same NSR_CLIENT parameter value in both. Ensure that the NSR_CLIENT value used for a backup is the same as the NSR_CLIENT value used for the restore of the backup. Setting NSR_CLIENT to the primary database hostname might be preferable.

Sybase backup considerations

The following Sybase operations require specific considerations:

- Setting up Sybase roles and permissions
- Verification of Sybase database backups
- Sybase incremental backups to specific volume pools
- Sybase transaction log backups
- Sybase threshold procedure
- Sybase backups on HP-UX
- Sybase 15.7 operations

Note: NMDA does not support probe-based backups with **nsrdaprobe** of a Sybase database that has the data and the log segments on the same device. "NMDA probe" on page 106 provides details.

Review "Sybase full and incremental backups" on page 59 for considerations for the NMDA Sybase full and incremental backups.

Setting up Sybase roles and permissions

The Sybase administrator is the person who is responsible for Sybase backup and recovery. NMDA relies on the Sybase administrator's ability to dump databases and load databases when performing Sybase backup operations and recovery operations. The administrator must have the appropriate Sybase roles and permissions.

The following table lists the Sybase roles and permissions that NMDA requires for performing the Sybase administrative actions.

Table 14 Sybase roles and permissions

Role or permission	Action	NMDA command
SA_role or create database privileges	Create a database	Not applicable
SA_role, DBO (database ownership), or OPER_role	Backup and restore databases	**nsrdasv** **nsrsybrc**
SA_role, DBO	Run a database consistency check	**nsrsybcc**

When you set the NMDA parameter USE_CONSISTENCY_CHECK to TRUE for a scheduled backup, the SA_role or DBO must run the backup to ensure that the **nsrsybcc** command succeeds. "USE_CONSISTENCY_CHECK" on page 396 describes the parameter.

NMDA requires that the Sybase user, for example, Sybase OPER_role, be a member of the database to check whether the database and the log are on separate segments. If the Sybase administrator is not a member of the database, then the backup fails. However, this limitation does not apply to recovering the Sybase database.

Specifying verification of Sybase database backups

NMDA supports verification of Sybase database backups at the header verification and full verification levels.

Set the NSR_ASE_VERIFY parameter either with the wizard or in the NMDA configuration file to specify the backup verification level. Set the parameter to one of the following values:

- header—Specifies to verify the page header information only.
- full—Specifies to verify both the header information and the rows structure (full verification of the backup).

For example, the following NSR_ASE_VERIFY setting specifies to perform a full verification of the backup:

```
NSR_ASE_VERIFY=full
```

If you do not specify a verification value, then NMDA does not perform any verification but adds a message to the log file.

Specifying volume pools for Sybase incremental backups

An NMDA Sybase full backup or incremental backup stores the backup data in one save set and creates a separate save set for the backup metadata. For example, an NMDA Sybase backup of the database SYBASE:/SERVER/sybdb produces two save sets:

- SYBASE:/SERVER/sybdb—Contains the metadata for the backup.
- SYBASE:/SERVER/sybdb.1—Contains the data for the backup.

During a Sybase incremental backup, NMDA backs up only the transaction log of each database. You can send the backup of logs to a special NetWorker log volume pool by setting the NSR_LOG_VOLUME_POOL parameter, as described in "NSR_LOG_VOLUME_POOL" on page 395.

However, NMDA backs up the metadata from a Sybase incremental backup to a regular (not log) volume pool, for example, as specified by NSR_DATA_VOLUME_POOL. In the preceding example, NMDA stores the incremental backup of sybdb as follows:

- SYBASE:/SERVER/sybdb—Stored in the NSR_DATA_VOLUME_POOL pool.
- SYBASE:/SERVER/sybdb.1—Stored in the NSR_LOG_VOLUME_POOL pool.

Sybase transaction log backups

For manual and scheduled Sybase backups, NMDA backs up the Sybase transaction log during an incremental backup.

NMDA supports the setting of the NSR_DUMP_LOG_OPT parameter to deal with the transaction log in an emergency situation only, such as a lack of free log space or a corrupted media database.

NOTICE

Use the NSR_DUMP_LOG_OPT parameter with care because it can prevent you from restoring the ASE server or specific databases.

Table 15 on page 139 describes the full backup behavior for the different NSR_DUMP_LOG_OPT settings when the database and the transaction log are on the same or separate devices.

Table 16 on page 140 describes the incremental backup behavior for different NSR_DUMP_LOG_OPT settings when the database and the transaction log are on separate devices.

Note: NMDA does not support the NSR_DUMP_LOG_OPT parameter for incremental backups when the database and the transaction log are on the same device.

The Sybase documentation describes when and how to use the "no_log", "no_truncate", and "truncate_only" options.

Table 15 Full backups with NSR_DUMP_LOG_OPT (page 1 of 2)

NSR_DUMP_LOG_OPT parameter setting	Database and log on separate devices	Database and log on the same device
For non-read-only databases:		
"no_log"	1. Truncates the transaction log *without* logging the transaction. 2. Backs up the database.	1. Displays an error message stating that the setting is invalid and the "no_log" operation is ignored. 2. Backs up the database. 3. Truncates the transaction log.

Table 15 Full backups with NSR_DUMP_LOG_OPT (page 2 of 2)

NSR_DUMP_LOG_OPT parameter setting	Database and log on separate devices	Database and log on the same device
“no_truncate”	1. Backs up the transaction log. 2. Backs up the database.	Backs up the database.
“truncate_only”	1. Truncates the transaction log. 2. Backs up the database.	1. Truncates the transaction log. 2. Backs up the database.
(No setting)	Backs up the database.	1. Backs up the database. 2. Truncates the transaction log.
For read-only databases:		
“no_log”	1. Truncates the transaction log *without* logging the transaction. 2. Backs up the database.	1. Displays an error message stating that the setting is invalid and the “no_log” operation is ignored. 2. Backs up the database.
“no_truncate”	1. Backs up the transaction log. 2. Backs up the database.	Backs up the database.
“truncate_only”	Displays an error message stating that the setting is invalid for read-only databases.	Displays an error message stating that the setting is invalid for read-only databases.
(No setting)	Backs up the database.	Backs up the database.

Table 16 Incremental backups with NSR_DUMP_LOG_OPT (page 1 of 2)

NSR_DUMP_LOG_OPT parameter setting	Database and log on separate devices
For non-read-only databases:	
“no_log”	1. Displays an error message stating that the setting is not supported with incremental backups, and that a full backup must be performed instead. 2. Backs up the transaction logs.
“no_truncate”	Backs up the transaction logs.
“truncate_only”	1. Displays an error message stating that the setting is not supported with incremental backups, and that a full backup must be performed instead. 2. Backs up the transaction logs.
(No setting)	Backs up the transaction logs.
For read-only databases:	
“no_log”	1. Displays an error message stating that the setting is not supported with incremental backups, and that a full backup must be performed instead. 2. Backs up the transaction logs.

Table 16 Incremental backups with NSR_DUMP_LOG_OPT (page 2 of 2)

NSR_DUMP_LOG_OPT parameter setting	Database and log on separate devices
"no_truncate"	Backs up the transaction logs.
"truncate_only"	Displays an error message stating that the setting is invalid for read only databases.
(No setting)	Backs up the transaction logs.

Configuring a Sybase threshold procedure

Sybase software enables you to register a threshold procedure to free up the log space for a database by performing the dump of the log when a threshold is reached. If there is not enough log space, the transactions will be either terminated or suspended.

NMDA supports the threshold procedure through the threshold.sql file, as described in "Sample threshold procedure" on page 142.

A threshold procedure provides the following features:

- When the system reaches the threshold of a Sybase database, NMDA backs up the transaction log with the **dump** command.
- If the system does not support a backup (dump) of the transaction log, then NMDA performs a full database backup and truncates the transaction log.

NOTICE

If you use the threshold procedure for transaction log backups, set the following as environment variables beforehand:
- NSR_CLIENT
- NSR_DATA_VOLUME_POOL
- NSR_LOG_VOLUME_POOL
- NSR_SERVER

For UNIX or Linux, as the Sybase user, set the environment variables in the shell that launches the Sybase Backup Server.
For Windows, as the Sybase user, set the environment variables as Windows system environment variables, then reboot the system and start the Sybase Backup Server.
Appendix A, "NMDA Parameters and Configuration File," describes these parameters.

If NMDA cannot perform a full database dump, then perform either of these actions:

- Add space to the transaction log.
- Terminate processes that were suspended when the threshold was crossed.

The Sybase documentation provides information about thresholds.

Threshold procedure versus probe-based backup

The NMDA threshold procedure frees the log space for the Sybase database.

The NMDA probe-based backup backs up a server (all databases) or an individual database, based on the number of transaction log pages generated. The main purpose of a probe-based backup is not to free the log space, but to determine if there has been enough database activity to start a backup.

"Configuring probe-based backups" on page 105 describes probe-based backups.

Sample threshold procedure

Use the sample threshold procedure described in this topic to implement transaction log backups to free the log space. Edit the sample threshold procedure to suit the environment.

The following table lists the default location for the sample threshold procedure.

Table 17 Threshold procedure location

Operating system	Location
AIX	/usr/bin/threshold.sql
HP-UX	/opt/networker/bin/threshold.sql
Solaris	/usr/sbin/threshold.sql
Linux	/usr/sbin/threshold.sql

Installing the sample threshold procedure in a database

To use the sample threshold procedure:

1. Run the **isql** command with threshold.sql as an input file:

   ```
   isql -Usa -P password -S Sybase_server -D database_name
   -i threshold.sql
   ```

   where:

   - *password* is the password of the Sybase SA account.
   - *Sybase_server* is the name of the Sybase server.
   - *database_name* is the name of the Sybase database.

2. Start an **isql** session and verify that the threshold procedure is in place:

   ```
   isql -Usa -P password -S Sybase_server -D database_name
   1> sp_help
   2> go
   ```

When you run the **sp_help** command, the procedure **sp_thresholdaction** must appear. If the procedure does not appear, then verify that the database used is the correct one.

The Sybase documentation describes how to use the **sp_addthreshold** command to perform these actions:

- Add the NMDA threshold procedure to the Sybase server.
- Manage free space with thresholds.

Setting the environment for Sybase backups on HP-UX

For Sybase backups on HP-UX Itanium, ensure the required setting of the LD_PRELOAD environment variable:

1. In the shell where the Sybase backup server runs, set the variable as the Sybase user:

   ```
   LD_PRELOAD=$SYBASE/$SYBASE_ASE/lib/libnsrsyb.so
   ```

   Note: For ASE 15.0.*x*, use $SYBASE/$SYBASE_ASE/lib/libnsrsyb32.so in the preceding command.

2. Start the Sybase backup server for the NMDA backup in the same shell where you set the LD_PRELOAD variable in step 1.

Configuring Sybase ASE 15.7 backup features

If you use Sybase database compression, then do not set the NMDA parameter NSR_COMPRESSION when you back up a Sybase compressed database.

You can have NMDA perform a shrink log operation during a Sybase 15.7 database backup by using one of the following methods:

- Specify the NSR_DUMP_DATA_OPT setting in the Advanced Options table in the configuration wizard.
- Set the NSR_DUMP_DATA_OPT parameter for a manual backup or a scheduled backup configured without the wizard.

"NSR_DUMP_DATA_OPT" on page 394 provides details.

CHAPTER 3
Backup Procedures

This chapter includes the following topics:

- Performing scheduled backups 146
- Performing manual backups 148
- Verifying backup information in NetWorker indexes 159
- Synchronizing backup catalogs and deleting backups 164

Performing scheduled backups

A scheduled backup is the type of backup that the NetWorker server initiates by running the **savegrp** program according to a configured backup schedule.

Complete these procedures to prepare for a scheduled backup:

1. Ensure that the required backup configurations are in place, as described in "Configuring NMDA backups" on page 70.
2. Run a test scheduled backup manually, as described in "Testing scheduled backups" on page 146.

At the end of a scheduled backup, the software automatically backs up additional files by default:

- Specific files for the application as described in Chapter 5, "Disaster Recovery"
- NetWorker bootstrap and client indexes

Chapter 5, "Disaster Recovery," provides details about any other files to back up in preparation for disaster recovery.

Testing scheduled backups

Use NMC to manually test the scheduled backup of a selected backup group. This test backup confirms that the scheduled backup runs as expected.

1. In the web browser, go to the URL http://*NMC_server_name*:9000, and log in to the NMC console server.

   The default NMC connection port is 9000. If you configured a different port, then use the different number.
2. In the **Enterprise** window, select the NetWorker server host and start the administrative GUI.
3. In the **Monitoring** window, select the group name for the backup.
4. Start the scheduled backup for the specified group. Right-click the group name and select **Start** from the drop-down menu.

The *EMC NetWorker Administration Guide* and NMC online help describe how to use the NMC interface.

Canceling scheduled backups

You can cancel a scheduled backup by using the following procedures.

NOTICE

If you cancel a backup, some of the backed-up data might not be recoverable.
To restart a cancelled scheduled backup, follow the instructions in "Restarting failed scheduled backups" on page 147.

1. Start the NMC program.
2. Select the backup group in the **Monitoring** window.

3. For an Oracle backup, ensure that the Oracle user has the required privilege for backup deletion, as described in Table 4 on page 76.

   When you cancel an Oracle backup, the Oracle software might ask NMDA to remove some of the completed backups in the same backup session if those backups are not recoverable.

4. Right-click the group and select **Stop**.

Restarting failed scheduled backups

You can enable the automatic restart of a failed scheduled backup. You can also manually restart a failed backup.

- To enable the automatic restart of a scheduled backup after a failure, set the **Client Retry** attribute in the Group resource in the NMC GUI:
  - For a Lotus checkpoint restart (CPR) backup that restarts within the group restart window, or an Oracle restartable backup, NMDA restarts the backup from the point-of-failure.

    Note: A Lotus CPR backup is also called a Lotus restartable scheduled backup.

  - For other backups or a Lotus CPR backup that restarts outside the restart window, NMDA restarts the backup from the beginning.
- To manually start a cancelled or failed backup from the beginning, right-click the group to start in the NMC GUI, and select **Start**.
- To manually restart a Lotus CPR backup from the point-of-failure, right-click the group to start in the NMC GUI, and select **Restart**.

Monitoring scheduled backups

You can monitor a scheduled backup by using the Monitoring window in NMC.

Note: The names and locations of tabs and windows might vary between different NMC releases.

The Group Details window in NMC displays messages and backup status information during and after a scheduled backup.

To view the Group Details window, right-click the group name and select Show Details in the NMC Monitoring window of the NetWorker Administration window.

The *EMC NetWorker Administration Guide* provides details about viewing scheduled backups in NMC.

In addition to monitoring a database backup in NMC, you can use the database activity logs, if supported, to monitor the backup results from the database server.

Performing manual backups

You can perform a manual backup after you have completed the backup configurations and determined the files to back up.

Unlike scheduled backups, a manual backup does not automatically back up these files:

- NetWorker server bootstrap and client indexes
- Additional files for the application that a scheduled backup automatically backs up, as described in Chapter 5, "Disaster Recovery"

NOTICE

After running a manual backup, back up the server bootstrap and client indexes, as described in the *EMC NetWorker Administration Guide*. Also, back up any additional files required to prepare for disaster recovery, as described in Chapter 5, "Disaster Recovery."

Perform a manual backup by using the appropriate procedure from the following table.

Table 18 Manual backup procedures

Database or application	Manual backup procedures
DB2	• "Performing DB2 manual backups with the db2 backup command" on page 148 • "Performing DB2 manual backups with the DB2 GUI" on page 149
Informix	• "Performing Informix manual backups with the onbar command" on page 150
Lotus	• "Performing Lotus manual backups with the nsrdasv command" on page 150 • "Performing Lotus manual backups with NetWorker User for Lotus" on page 151
MySQL	• "Performing MySQL manual backups with the nsrdasv command" on page 153
Oracle	• "Performing Oracle manual backups with the rman command" on page 153 • "Performing Oracle manual backups with Oracle Enterprise Manager" on page 154
Sybase	• "Performing Sybase manual backups with the nsrdasv command" on page 155 • "Performing Sybase manual backups with NetWorker User for Sybase" on page 156 **Note:** Prior to a Sybase manual backup, perform a database consistency check according to "Performing Sybase database consistency checks before backups" on page 154.

Performing DB2 manual backups with the db2 backup command

Before you perform the manual backup, ensure that the required backup configurations are completed from "Configuring NMDA backups" on page 70.

You can run the appropriate **db2 backup** command to perform a DB2 manual backup from the command line.

1. Log in to the DB2 client host as the DB2 operating system user.
2. Run the **db2 backup** command with the appropriate options as described in the DB2 documentation. For example:
   - On UNIX, type the command:

     ```
     db2 backup db sample load /usr/lib/libnsrdb2.so options
     @pathname/nmda_db2.cfg
     ```

- On Windows, type the command:

  db2 backup db *sample* **load**
  *NetWorker_install_dir***nsr\bin\libnsrdb2.dll options**
  @*pathname***nmda_db2.cfg**

  where:

  – *sample* is the name of the database to back up.

  – *pathname*/nmda_db2.cfg is the complete pathname of the NMDA configuration file that contains the parameter settings for the DB2 manual backup.

  – *NetWorker_install_dir* is the NetWorker software installation directory on Windows systems.

After completing the manual backup, back up the server bootstrap and client indexes to prepare for disaster recovery of the NetWorker server. Refer to Chapter 5, "Disaster Recovery," for details about any additional files to back up.

Performing DB2 manual backups with the DB2 GUI

Before you perform the manual backup, ensure that the required backup configurations are completed from "Configuring NMDA backups" on page 70.

You can run the appropriate DB2 GUI as the DB2 user with permissions to perform the DB2 manual backup:

- With DB2 version 10.1 or later, run the IBM Data Studio GUI.
- With previous versions of DB2, run the DB2 Control Center GUI.

Specify the NMDA configuration file in the GUI by setting VENDOROPT to the value @*configuration_file_pathname*. For example:

```
@d:\nmda_db2.cfg
```

Set the Vendor DLL to the NMDA DB2 library name in the GUI.

Note: After you select the DLL path with the DB2 Control Center for a Windows client, enclose the path with quotes or use a short file name (8.3 format). Otherwise, the backup returns an error similar to the following example:

```
SQL0104N  An unexpected token
"Files\Legato\nsr\bin\libnsrdb2.dll" was found following "<identifier>".
Expected tokens may include:  "INCLUDE".
```

After completing the manual backup, back up the server bootstrap and client indexes to prepare for disaster recovery of the NetWorker server. Refer to Chapter 5, "Disaster Recovery," for details about any additional files to back up.

Performing Informix manual backups with the onbar command

Before you perform the manual backup, ensure that the required backup configurations are completed from "Configuring NMDA backups" on page 70.

You can run the appropriate **onbar** command to perform an Informix manual backup from the command line.

1. Log in to the Informix client host as the Informix operating system user.
2. Ensure that the required NMDA parameters, such as NSR_SERVER, are set as environment variables. Appendix A, "NMDA Parameters and Configuration File," describes the NMDA parameters.
3. To manually back up dbspaces, blobspaces, or logical log files, run the appropriate **onbar** backup command at the command line. For example:

   ```
   onbar -b -L 0 dbspace01
   onbar -l -c
   ```

   These commands perform the following tasks:

   - Level 0 (NetWorker level full) backup of the dbspace named dbspace01
   - Closure of the active logical log
   - Backup of the logical logs, not including the newly activated log

After completing the manual backup, back up the server bootstrap and client indexes to prepare for disaster recovery of the NetWorker server. Refer to Chapter 5, "Disaster Recovery," for details about any additional files to back up.

Performing Lotus manual backups with the nsrdasv command

Before you perform the manual backup, ensure that the required backup configurations are completed from "Configuring NMDA backups" on page 70.

You can run the appropriate **nsrdasv** command to perform a Lotus manual backup from the command line.

1. Set either the NSR_BACKUP_PATHS or NSR_BACKUP_LOTUS_DIR parameter in the NMDA configuration file to specify the Lotus directories or the files to back up. Do not set both parameters simultaneously. These descriptions provide more details:

   - "NSR_BACKUP_LOTUS_DIR" on page 372
   - "NSR_BACKUP_PATHS" on page 373
2. Set the NSR_NOTES_INI_PATH parameter in the NMDA configuration file, as described in "NSR_NOTES_INI_PATH" on page 375.
3. Log in to the Domino or Notes host as the Lotus operating system user.
4. Run the appropriate **nsrdasv** backup command at the command line:

   nsrdasv(**.exe**) **-z** *configuration_file_path*

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the parameter settings for the Lotus manual backup.

After completing the manual backup, back up the server bootstrap and client indexes to prepare for disaster recovery of the NetWorker server. Refer to Chapter 5, "Disaster Recovery," for details about any additional files to back up.

Performing Lotus manual backups with NetWorker User for Lotus

Before you perform a manual backup, ensure that the required backup configurations are completed from "Configuring NMDA backups" on page 70.

On Windows systems only, you can run the NetWorker User for Lotus GUI (**nwbml.exe**) to perform manual backups of Lotus Domino/Notes databases located on the same host.

1. Start the GUI. For example, select **NetWorker User for Lotus** from **Start** › **Programs** › **EMC NetWorker**.
2. To connect to a different NetWorker server, complete these steps:
   a. Select **Select NetWorker Server** from the **Operation** menu.

      The **Change Server** dialog box appears.

   b. Click **Update List** to refresh the list of NetWorker servers.

   c. Select or type the name of the server.

   d. Select **Save as Default Server** to use the server as the default NetWorker server.

   e. Click **OK**.
3. In the NetWorker User for Lotus GUI, select **Backup** from the **Operation** menu.

   The **Backup** window appears as shown in the following figure. The online help describes the toolbar buttons.

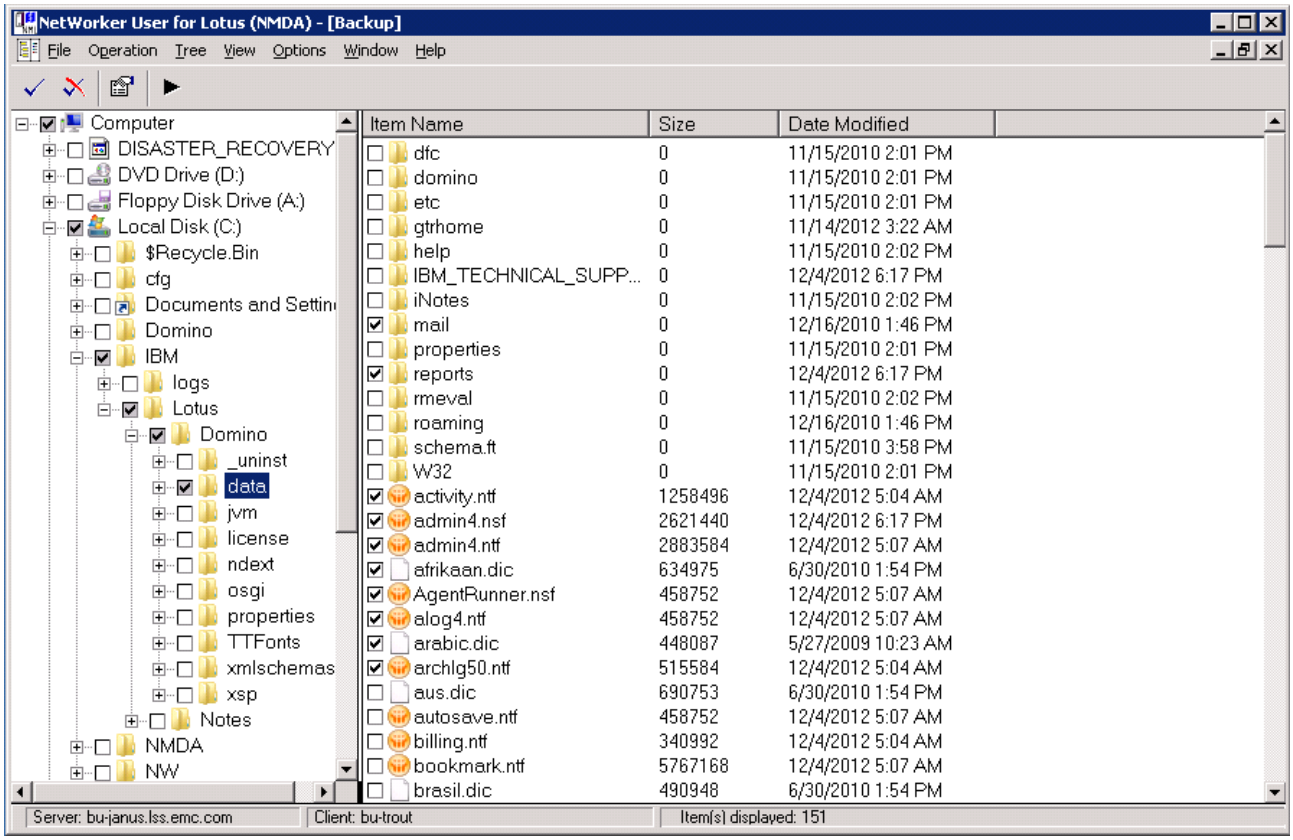

Figure 2 Backup window in NetWorker User for Lotus

4. To view a list of files or databases available for backup, select a Lotus directory in the left pane of the **Backup** window. The Lotus directory contents appear in the right pane.
5. Select the checkbox next to each file, database, or directory to be backed up. If you select a directory, all the files and subdirectories in that directory are backed up. You must select at least one item for backup.

   Note: Do not set NSR_BACKUP_PATHS or NSR_BACKUP_LOTUS_DIR in the NMDA configuration file. Instead, use the GUI to select the objects to back up.
   During a backup of transaction logs only, any selected databases are ignored and are not backed up. However, you must select at least one database to enable the log backup.

6. From the **Options** menu, select **Backup Options**.

   The **Backup Options** dialog box appears as shown in the following figure.

   Note: Parameters specified in the **Backup Options** dialog box take precedence over the corresponding parameters specified in the NMDA configuration file.
   To enable Lotus incremental backups, set NSR_BACKUP_LEVEL=incr in the configuration file. To enable backups of Lotus transaction logs only, set NSR_BACKUP_LEVEL=txnlog in the configuration file. "NSR_BACKUP_LEVEL" on page 372 provides details.

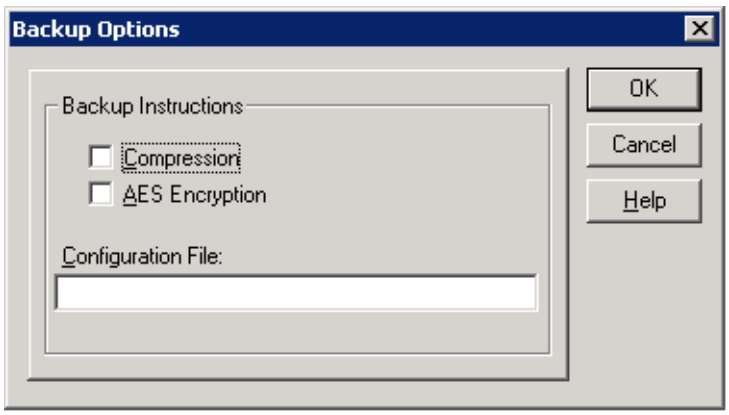

Figure 3 Backup Options dialog box in NetWorker User for Lotus

7. In the **Backup Options** dialog box, specify the required options:
   a. To specify compression or encryption:
      - For the current backup only, select **Compression**, **AES Encryption**, or both.
      - For all backups from the GUI, set NSR_COMPRESSION, NSR_AES_ENCRYPTION, or both to TRUE in the NMDA configuration file.
   b. To specify the location of the NMDA configuration file, type the complete pathname in the **Configuration File** box. "NMDA configuration file" on page 354 provides details about the NMDA configuration file.
   c. Click **OK**.

8. Click **Start** in the **Backup** window.

   The **Backup Status** window displays information about the backup being performed.

After completing the manual backup, back up the server bootstrap and client indexes to prepare for disaster recovery of the NetWorker server. Refer to Chapter 5, "Disaster Recovery," for details about any additional files to back up.

Performing MySQL manual backups with the nsrdasv command

Before you perform the manual backup, ensure that the required backup configurations are completed from "Configuring NMDA backups" on page 70.

You can run the appropriate **nsrdasv** command to perform a MySQL manual backup from the command line.

1. Log in to the MySQL server host as the MySQL operating system user.
2. Run the appropriate **nsrdasv** backup command at the command line:

   ```
   nsrdasv -z configuration_file_path
   ```

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the parameter settings for the MySQL manual backup.

After completing the manual backup, back up the server bootstrap and client indexes to prepare for disaster recovery of the NetWorker server. Refer to Chapter 5, "Disaster Recovery," for details about any additional files to back up.

Performing Oracle manual backups with the rman command

Before you perform the manual backup, ensure that the required backup configurations are completed from "Configuring NMDA backups" on page 70.

You can run the appropriate **rman** command to perform an Oracle manual backup from the command line.

1. Log in to the Oracle Server host as the Oracle operating system user.
2. Run the appropriate **rman** command at the command line to start the RMAN backup script. For example:

   a. Store the RMAN script for a manual backup from Example 7 on page 129 in the /disk1/scripts/full_backup.txt file on a UNIX system that runs the Oracle Server.

   b. Configure the Net service to connect to the payroll and rcvcatdb databases.

   c. Start the manual backup by running this command:

   ```
   rman target sys/oracle@payroll rcvcat rman/rman@rcvcatdb cmdfile
   \'/disk1/scripts/full_backup.txt\'
   ```

After completing the manual backup, back up the server bootstrap and client indexes to prepare for disaster recovery of the NetWorker server. Refer to Chapter 5, "Disaster Recovery," for details about any additional files to back up.

Performing Oracle manual backups with Oracle Enterprise Manager

Before you perform the manual backup, ensure that the required backup configurations are completed from "Configuring NMDA backups" on page 70.

You can run the Oracle Enterprise Manager to perform an Oracle manual backup from the GUI.

Note: If you schedule a backup by using Oracle Enterprise Manager, NMDA considers the backup to be a manual backup.

1. Log in to the Oracle Server host as the Oracle operating system user.
2. To manually back up Oracle data by using the GUI, run the Oracle Enterprise Manager Backup Management Tools that run the RMAN backup script.

   The Backup Management Tools include a graphical user interface to RMAN for generating the required RMAN commands and performing backup and restore operations.

   NOTICE

   When a backup or restore completes successfully, the status of the job appears as failed in the job queue history of the Oracle Enterprise Manager. This incorrect status is a known issue with Oracle Enterprise Manager. View the job output to confirm that the backup or restore completed successfully.

After completing the manual backup, back up the server bootstrap and client indexes to prepare for disaster recovery of the NetWorker server. Refer to Chapter 5, "Disaster Recovery," for details about any additional files to back up.

Performing Sybase database consistency checks before backups

Before you start a Sybase manual backup, you can perform a Sybase database consistency check with the appropriate **nsrsybcc** command at the command line.

1. Log in to the Sybase host as the Sybase operating system user.
2. Type the following command:

   ```
   nsrsybcc -U user_ID -P password [-o dbcc_option]
   SYBASE:/ASE_server_name[/database_name]
   ```

   where:

   - *user_ID* is the username of the Sybase database user account.
   - *password* is the password of the Sybase database user account.
   - *dbcc_option* is the option that specifies the type of database consistency check, as described in the following table.
   - *ASE_server_name* is the Sybase server name.
   - *database_name* is the name of the database on the Sybase server.

Note: If you specify the Sybase server name only, SYBASE:/*ASE_server_name*, with the **nsrsybcc** command (without specifying a database name), the command performs a database consistency check for every database on the Sybase server.

If you do not include the **-o** option to specify the type of database consistency check, then the **nsrsybcc** command performs the following checks:

- The **dbcc checkstorage** check when the **dbccdb** database is set up
- The **dbcc checkcatalog**, **dbcc checkalloc**, and **dbcc checkdb** checks when the **dbccdb** database is not set up

Table 19 Database consistency check options of the nsrsybcc command

The -o option	Type of database consistency check
-o ckdb	dbcc checkdb
-o ckal	dbcc checkalloc
-o ckcat	dbcc checkcatalog
-o ckdbnoidx	dbcc checkdb (skip_ncindex)
-o ckstor	dbcc checkstorage **Note:** Ensure that the **dbccdb** database is set up.

The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrsybcc** command.

Performing Sybase manual backups with the nsrdasv command

Before you perform the manual backup, ensure that the required backup configurations are completed from "Configuring NMDA backups" on page 70.

You can run the appropriate **nsrdasv** command to perform a Sybase manual backup from the command line.

1. Set the NSR_BACKUP_PATHS parameter in the NMDA configuration file to specify the Sybase server or the Sybase databases to back up. "NSR_BACKUP_PATHS" on page 394 provides details.
2. Log in to the Sybase host as the Sybase operating system user.
3. Run the appropriate **nsrdasv** backup command at the command line:

   ```
   nsrdasv(.exe) -z configuration_file_path
   ```

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the parameter settings for the Sybase manual backup.

After completing the manual backup, back up the server bootstrap and client indexes to prepare for disaster recovery of the NetWorker server. Refer to Chapter 5, "Disaster Recovery," for details about any additional files to back up.

Performing Sybase manual backups with NetWorker User for Sybase

Before you perform a manual backup, ensure that the required backup configurations are completed from "Configuring NMDA backups" on page 70.

On Windows systems only, you can run the NetWorker User for Sybase GUI (**nwbms.exe**) to perform manual backups of local Sybase databases.

1. Start the GUI. For example, select **NetWorker User for Sybase** from **Start** › **Programs** › **EMC NetWorker**.
2. To connect to a different NetWorker server, complete these steps:
   a. Select **Select NetWorker Server** from the **Operation** menu.

      The **Change Server** dialog box appears.

   b. Click **Update List** to refresh the list of NetWorker servers.

   c. Select or type the name of the server.

   d. Select **Save as Default Server** to use the server as the default NetWorker server.

   e. Click **OK**.
3. In the NetWorker User for Sybase GUI, select **Backup** from the **Operation** menu.

   The **Backup** window appears. The online help describes the toolbar buttons.
4. Type the required field values in the **Sybase Server Login** dialog box if it appears and click **OK**.

   The online help describes the fields in the dialog box.
5. To view a list of files or databases available for backup, select a Sybase server in the left pane of the **Backup** window. The Sybase server contents appear in the right pane.
6. Select the checkbox next to each file or database to be backed up.
7. From the **Options** menu, select **Backup Options**.

   The **Backup Options** dialog box appears, as shown in the following figure.

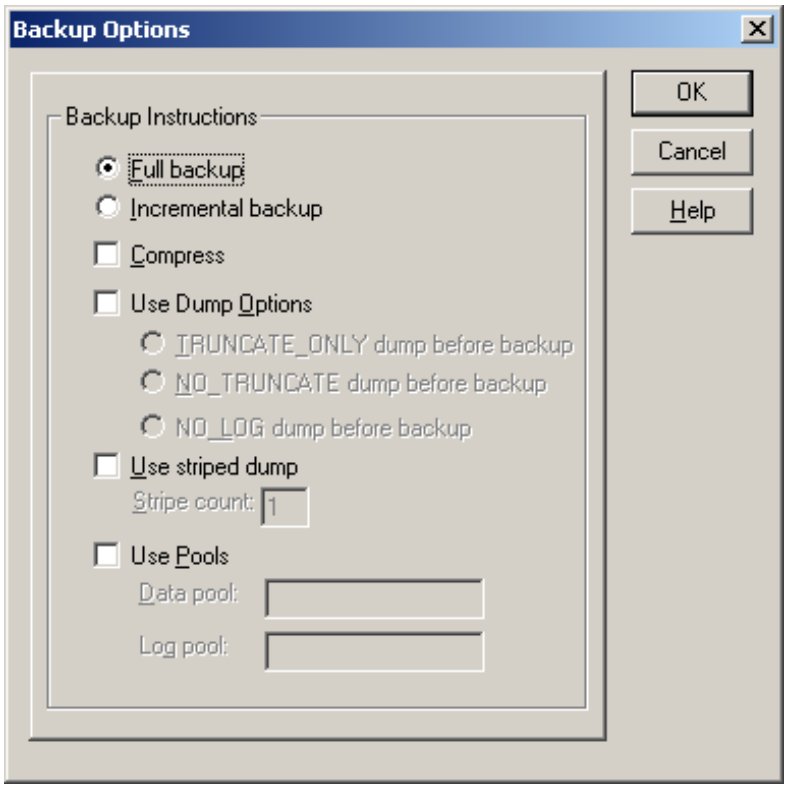

Figure 4 Backup Options dialog box in NetWorker User for Sybase

8. In the **Backup Options** dialog box, select the required options and click **OK**.

   The online help describes the options in the dialog box.

9. Click **Start** in the **Backup** window.

   The **Backup Status** window displays information about the backup being performed.

After completing the manual backup, back up the server bootstrap and client indexes to prepare for disaster recovery of the NetWorker server. Refer to Chapter 5, "Disaster Recovery," for details about any additional files to back up.

Canceling manual backups

You can cancel a manual backup by using the following procedures.

NOTICE

If you cancel the backup, some of the backed-up data might not be recoverable. Restart the cancelled backup process from the beginning, and ensure that the backup successfully completes without interruption.

NMDA does not support other procedures for canceling a manual backup. For example, the **kill -9** command or Windows equivalent might cause a Lotus Domino server to crash.

- To cancel a backup that you started from a command line, press **Ctrl-c** or the equivalent "attention" keyboard shortcut.
- To cancel a backup that you started from a database GUI such as Oracle Enterprise Manager, follow the instructions specific to the GUI.
- To cancel a Lotus or Sybase backup from the NetWorker User for Lotus or NetWorker User for Sybase GUI, select **File** › **End Backup**. This method applies to Windows only.

Canceling Oracle manual backups

Before you cancel an Oracle manual backup, ensure that the Oracle user has the required privilege for backup deletion, as described in Table 4 on page 76. When you cancel the backup, the Oracle software might ask NMDA to remove the NetWorker index entries for previous backups.

To cancel an unresponsive Oracle manual backup, follow the instructions in "Canceling unresponsive Oracle backups on UNIX" on page 427 or "Canceling unresponsive Oracle backups on Windows" on page 427.

Monitoring manual backups

You can monitor a manual backup from the command line or the backup GUI program. You can also monitor a manual backup in NMC, which provides a centralized view of all the NetWorker server backup and restore activity.

After you have started a manual backup from the command line or the backup GUI program, view the operational messages that appear about the status of the backup.

View the following information in the Monitoring window in NMC:

- Backup status information for running backups
- Alerts, log messages, and information about device operations

The *EMC NetWorker Administration Guide* describes how to use NMC.

Verifying backup information in NetWorker indexes

The NetWorker server maintains information about each backup in the online indexes.

- To query the client file index, use the **nsrinfo** command on the NMDA client host:

  ```
  nsrinfo -n NMDA_application -s NetWorker_server_hostname
  NMDA_client_hostname
  ```

  where *NMDA_application* is **db2**, **informix**, **mysql**, **notes**, **oracle**, or **sybase**.

- To query the media database, use the **mminfo** command on the NMDA client host:

  ```
  mminfo -v -s NetWorker_server_hostname -c NMDA_client_hostname
  ```

The *EMC NetWorker Command Reference Guide* and the UNIX man pages describe these NetWorker commands.

The following examples include the output of the **nsrinfo** and **mminfo** commands for an Oracle backup with the backup object (backup piece) named t1ld0g32_1_1:

```
# nsrinfo -n oracle -s nw-server ca-oracle

t1ld0g32_1_1, date=1273166244 Thu May 06 13:17:24 2010

# mminfo -v -s nw-server -c ca-oracle

volume          client        date          time          size
Oracle.001      ca-oracle     05/06/10      13:17:24      5633 KB

ssid            fl            level         name
4108515830      cb            full          RMAN:t1ld0g32_1_1
```

Cross-check the client file index and media database by using the backup save time through the **-t** option of the **nsrinfo** and **mminfo** commands:

```
mminfo -s nw-server -c ca-oracle -t 1273166244
nsrinfo -s nw-server -n oracle -t '05/06/10 13:17:24' ca-oracle
```

The NMC interface also provides queries and reports about the backup information.

Verifying Avamar backup information in NetWorker indexes

You can verify information about Avamar deduplication backups in the NetWorker online indexes by using the index query results from the **nsrinfo** and **mminfo** commands.

- To query the client file index, use the **nsrinfo** command, as described in "Verifying backup information in NetWorker indexes" on page 159.
- To query the media database, use the **mminfo** command with the following options:
  - **-q dedupe**—Displays only the save sets created through Avamar deduplication.
  - **-S**—Lists the extended options for each save set in the save set completion report.

The following **nsrinfo** query example includes the information from the client file index for a deduplication backup of Lotus data:

```
nsrinfo -n notes -vV win3e10-nml80

XBSA file 'NOTES:/C:/Program Files/IBM/Lotus/Domino/data/events4.nsf',
   size=404, off=404, app=notes(17), date=1236960727 3/13/2009
   11:12:07 AM, copyID = 1236960727.1236960729
```

In this case, the size of the backup file, 404, is the size of the metadata (hash information) from the NetWorker backup media for the deduplication backup.

The following **mminfo** query example includes the information from the media database for a single save set of Lotus data deduplicated to an Avamar server:

```
ssid=3585770967 savetime=3/13/2009 11:12:06 AM (1236960726)
   win3e10-nml80:NOTES:  level=full  sflags=vF  size=1608  files=4
   insert=3/13/2009  create=3/13/2009  complete=3/13/2009
   browse=4/13/2009 11:59:59 PM  retent=3/13/2010 11:59:59 PM
clientid=292ecd0b-00000004-49b194cd-49b194cb-00020c00-729b1c29
                *Client path: /NetWorker/win3e10-nw75/win3e10-nml80;
              *Data set size: 71041152;
          *De-Dup session id: 158;
         *De-Dup snapup time: 2009-03-13;
             *De-duplication: Yes;
        *De-duplication host: bu-doppelganger.lss.emc.com;
                    *Domain: /NetWorker/win3e10-nw75;
    *New data on De-Dup Node: 10483991.00;
                 *New files: 4;
        *Size on De-Dup Node: 71041152.00;
                       group: nml80;
Clone #1: cloneid=1236960727  time=3/13/2009 11:12:07 AM
   retent=3/13/2010  flags=F  frag@  0 volid=3602548168
   file/rec=3585770967/0  rn=0  last=3/13/2009
```

The save file size, 1608, is the size of the NetWorker save set on the backup media that contains the metadata only for the deduplication backup.

The extended attributes in the **mminfo** listing describe the following data:

- Total data protected by NetWorker and stored on the Avamar server.
- New data passed to the Avamar server during the deduplication backup.

Deletion of a deduplication backup produces the following results:

- Immediate deletion of the backup information from the NetWorker indexes.
- Queued request to delete the backup from the Avamar server.

Note: After a deduplication save set passes its retention time and its data chunks are deleted from the Avamar server, the save set might not be recoverable even with the **scanner** program.

Verifying Data Domain backup information in NetWorker indexes

You can verify information about Data Domain deduplication backups in the NetWorker online indexes by using the index query results from the **nsrinfo** and **mminfo** commands.

- To query the client file index, use the **nsrinfo** command, as described in "Verifying backup information in NetWorker indexes" on page 159.
- To query the media database, use the **mminfo** command with the following options:
  - **-q DataDomain**—Displays only the save sets created through deduplication to a Data Domain device.
  - **-S**—Lists the Data Domain information for each save set, such as the Data Domain clone ID and Data Domain statistics.

The NetWorker client index information for a Data Domain deduplication backup is the same as for a regular backup.

The following **mminfo** query lists only the backups to a Data Domain device:

```
mminfo -S -q DataDomain
```

The following **mminfo** query example includes the information from the media database for a single save set of a DB2 backup to a Data Domain device:

```
ssid=2253717478  savetime=08/24/2011 09:42:55 AM (1314193375)
   bu-pluto:DB2:/SMS1/NODE0000
level=full  sflags=vF  size=52099434476  files=1  insert=08/24/2011
create=08/24/2011  complete=08/24/2011  browse=09/24/2011 11:59:59 PM
   retent=08/24/2012 11:59:59 PM
clientid=7d32da8e-00000004-4e42944b-4e42944a-01565000-d2648f56
*ss data domain backup cloneid: 1314193376, 1314193377;
*ss data domain dedup statistics: \
"v1:1314193376:52271619528:52218920124:41558381793",
"v1:1314193377:52271619528:52218920124:41558381793";
                        group: db2_pluto_dfa;
Clone #1: cloneid=1314193376  time=08/24/2011 09:42:56 AM
   retent=08/24/2012  flags=F
    frag@  0  volid=2337602992  file/rec=0/0  rn=0  last=08/24/2011
Clone #2: cloneid=1314193377  time=08/24/2011 09:42:57 AM
   retent=08/24/2012  flags=F
    frag@  0  volid=2354380207  file/rec=0/0  rn=0  last=08/24/2011
```

Lotus restartable backup information in NetWorker indexes

You can use the NMC Console or the **mminfo** command to query information about Lotus restartable backups in the NetWorker indexes.

The following figure provides an example.

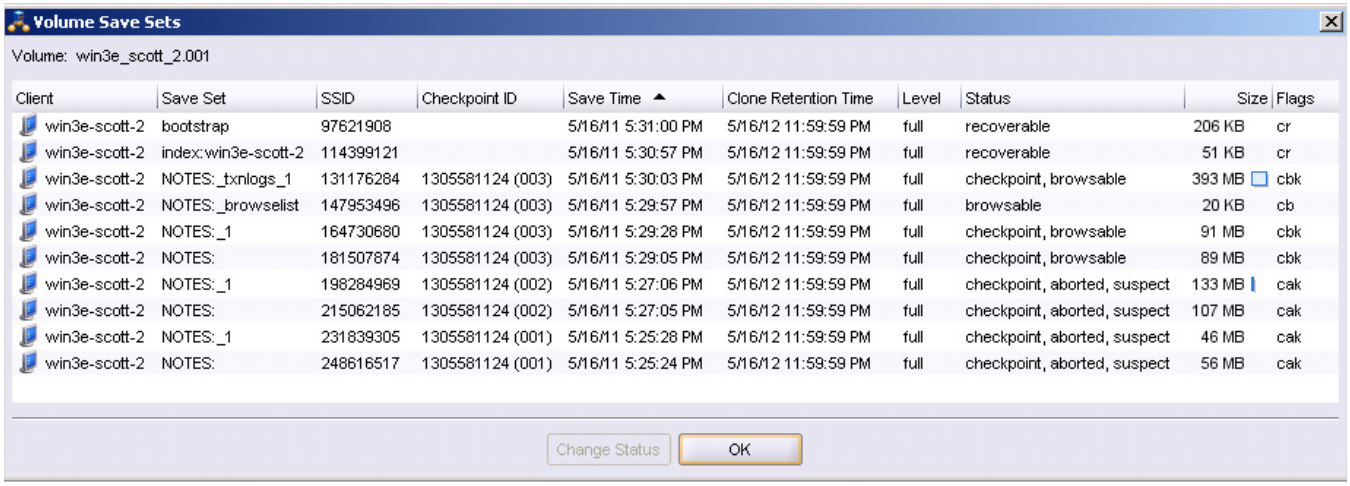

Figure 5 Lotus restartable backup information in NMC Console

The following **mminfo** query provides an example of Lotus restartable backup information in the media database. The **k** flag, checkpoint ID, and sequence number identify the checkpoint-enabled save sets. If the checkpoint restart backup fails, the partial save sets that have the **k** flag are not removed from the client file index or media database and can be used for restore:

```
mminfo -q "checkpoint_id=1305581124" -r "client, name, ssid(11), checkpoint_id, checkpoint_seq,
nsavetime, sumflags(3)"

client        name                      ssid       chkpt_id   chkpt_seq save time  fl
win3e-scott-2 NOTES:_txnlogs_1          131176284  1305581124 3         1305581403 cbk
win3e-scott-2 NOTES:_browselist         147953496  1305581124 3         1305581397 cb
win3e-scott-2 NOTES:_1                  164730680  1305581124 3         1305581368 cbk
win3e-scott-2 NOTES:                    181507874  1305581124 3         1305581345 cbk
win3e-scott-2 NOTES:_1                  198284969  1305581124 2         1305581226 cak
win3e-scott-2 NOTES:                    215062185  1305581124 2         1305581225 cak
win3e-scott-2 NOTES:_1                  231839305  1305581124 1         1305581128 cak
win3e-scott-2 NOTES:                    248616517  1305581124 1         1305581124 cak
```

The *EMC NetWorker Administration Guide* describes how to query information about partial save sets.

MySQL backup information in NetWorker indexes

When NMDA performs a MySQL backup, the software stores extra information in the NetWorker client file index about the MySQL log sequence number (LSN). NMDA uses the LSN information from previous backups to perform MySQL incremental backups.

NMDA generates one of the following names for the backup piece file that is stored in the NetWorker client file index for a MySQL backup:

MYSQL:/*unique_backup_name*_‹*backup_level*›_whole_‹*timestamp*›
MYSQL:/*unique_backup_name*_‹*backup_level*›_partial_‹*timestamp*›

where:

- MYSQL:/*unique_backup_name* is the unique name that you specify for the MySQL backup save set in one of these settings:
  - MYSQL_BACKUP_NAME parameter for a manual backup
  - Save Set attribute in the Client resource for a scheduled backup
- ‹*backup_level*› is the MySQL backup level:
  - full specifies a full backup.
  - differential specifies a differential incremental backup.
  - cumulative specifies a cumulative incremental backup.
- whole specifies a whole instance backup.
- partial specifies a partial backup.
- ‹*timestamp*› is the save time generated by the NetWorker server.

NMDA also generates an extra file named *backup_piece_file_name*_meta that contains meta information about the backup and is also stored in the client file index. NMDA uses this meta backup piece file during incremental backups and restores.

For example, the following **nsrinfo** command provides information from the client file index about the backup pieces for a manual full backup of an instance named myinstance. NSR_BACKUP_NAME was set to MYSQL:/myinstance for the manual backup:

```
# nsrinfo -n mysql -N MYSQL:/myinstance -s nw-server mysqlntx64

scanning client `mysqlntx64' for all savetimes from the oracle
   namespace on server nw-server
MYSQL:/myinstance_full_whole_1322094812, size=20971860, off=0,
   app=mysql(99), date=1322094812 11/17/2011 7:33:32 PM
MYSQL:/myinstance_full_whole_1322094812_meta, size=2048, off=0,
   app=mysql(99), date=1322094812 11/17/2011 7:33:34 PM
```

NMDA uses the save set name MYSQL:/*unique_backup_name*_logs for MySQL binary log backups.

For example, the following **mminfo** command provides save set information from the media database about the log backup included with the manual backup of the instance named myinstance:

```
# mminfo -s nw-server -c mysqlntx64

volume     client     date       size   level   name
mysql.001  mysqlntx64 11/17/2011 10 KB  full    MYSQL:/myinstance_logs
mysql.001  mysqlntx64 11/17/2011 19 MB  manual  MYSQL:/myinstance
```

Synchronizing backup catalogs and deleting backups

You can use the NetWorker policies to manage the lifecycle of an NMDA backup. In normal circumstances, these policies match any policies stored in the respective backup application catalogs. A DB2 database, Informix database, or Oracle database has its own backup catalog.

The NetWorker index entries can become out of sync with entries in an application catalog. It is a good practice to keep the NetWorker indexes synchronized with the backup application catalogs.

For example, a backup might expire in the NetWorker indexes or you might delete a backup. You might need to synchronize the NetWorker indexes and the DB2, Informix, or Oracle backup catalog. In addition, certain databases can have retention policies and can remove a backup based on the policies. You must synchronize the retention policies with NetWorker.

Note: Ensure that the database user that deletes the backup entries has the required privileges for backup deletion, as described in Table 4 on page 76.

To keep the backup catalogs synchronized, use the database-specific interface or instructions when deleting a DB2, Informix, or Oracle backup from both the database backup catalog and the NetWorker indexes.

Lotus and Sybase do not maintain their own backup catalogs. You can delete a Lotus backup or Sybase backup from the NetWorker indexes only by using the **nsrmm** command.

Deleting DB2 backups

The **db2 prune** command deletes backup entries for DB2 databases or tablespaces from both the DB2 server and NetWorker server. Deletion of backup entries might be necessary if the NetWorker index and DB2 recovery history files become too large and the retention period is long.

Note: You cannot use the **db2 prune** command to delete snapshot backups. "Deleting DB2 snapshots" on page 307 provides details.

To delete DB2 backup entries on both the DB2 server and NetWorker server:

1. Set the DB2 database configuration VENDOROPT parameter to the pathname of the NMDA configuration file (nmda_db2.cfg) for the DB2 database or tablespace whose backups are to be deleted. For example:

   ```
   db2 update db cfg for sample using vendoropt
   @/db/pathname/nmda_db2.cfg
   ```

   where:

   - *sample* is the name of the database or tablespace whose backups are to be deleted.
   - *pathname*/nmda_db2.cfg is the complete pathname of the NMDA configuration file that contains the parameter settings for the DB2 backup.

2. Enable the automatic deletion of physical backup images and log files by the **db2 prune** command:

   ```
   db2 update db cfg for sample using AUTO_DEL_REC_OBJ ON
   ```

   where *sample* is the name of the database whose backups are to be deleted.

   Note: Without this step, the **db2 prune** command removes entries only in the DB2 history file and does not remove the associated database backups and log files.

3. Remove unwanted backup entries with the **db2 prune** command. For example:

   ```
   $ db2 connect to sample
   $ db2 prune history timestamp and delete
   $ db2 terminate
   ```

   where:

   - *sample* is the name of the database whose backups are to be deleted.
   - *timestamp* (in format *yyyymmddhhmmss*, with minimum *yyyy*) specifies deletion of entries that are less than or equal to the timestamp value.

   Note: The **prune** command does not remove the most recent full backup entry regardless of the timestamp value, unless you include **with force option** after the timestamp.

4. Inspect the DB2 history file and the NetWorker index to verify that the backup objects are removed:

   Note: The NetWorker indexes might not update immediately.

   - On the DB2 server, use the following command:

     ```
     $ db2 list history backup all for sample
     ```

   - On the NetWorker server, use any one of the following commands:

     ```
     nsrinfo -v -s NetWorker_server -n db2 -X all DB2_server_hostname
     nsrinfo -n db2 DB2_server_hostname
     mminfo -c DB2_server_hostname
     ```

     where:

     – *NetWorker_server* is the hostname of the NetWorker server that contains the index with the backup entries.

     – *DB2_server_hostname* is the hostname of the DB2 server used to store the NMDA backups in the NetWorker indexes.

     The *EMC NetWorker Command Reference Guide* and the UNIX man pages describe the **nsrinfo** and **mminfo** commands.

The DB2 documentation describes the **db2 prune** command and the configuration parameters that you can set to maintain the backup history, particularly the REC_HIS_RETENTN and NUM_DB_BACKUPS parameters.

Deleting Informix backups

Deleting Informix backup entries might be necessary if the NetWorker index and Informix recovery history files become excessively large and the retention period is high. You can delete Informix backup entries for dbspaces from both the Informix server and NetWorker server with the **onsmsync** command.

Run the Informix **onsmsync** utility with the appropriate command options to remove the following items from the sysutils database and the emergency boot file:

- Backups that the NetWorker server has expired.
- Old backups based on the age of backup.
- Old backups based on the number of occurrences of the backups.

Note: To successfully delete the NetWorker index entries associated with the Informix backup entries, ensure that the user who runs the **onsmsync** utility has the required privileges for backup deletion, as described in Table 4 on page 76.

Cross-checking and deleting Oracle backups

Use the appropriate Oracle commands to keep the NetWorker indexes and RMAN catalog synchronized and to delete the entries for Oracle backups, as required.

- To keep the RMAN catalog and NetWorker indexes synchronized, run the **crosscheck** command regularly. For example, run the **crosscheck** command after you delete a backup from the NetWorker side by relabeling a tape or device.

  If you do not keep the RMAN catalog and NetWorker indexes synchronized, you might have issues at backup time or restore time. For example, Oracle might try to restore a backup that does not exist on NetWorker. "Backup and restore optimization" on page 49 provides details.

- To change the status of Oracle backup pieces to expired in the RMAN catalog when the corresponding NetWorker client file index entries are no longer browsable, run the **change...crosscheck** or **crosscheck** command.

  In the RMAN catalog, an expired status for an Oracle backup piece indicates that the backup piece does not exist in the NetWorker indexes.

- To delete expired backups or delete backups that were manually deleted on the NetWorker server, run the **crosscheck** and **delete expired backup** commands.

- To delete a backup manually if required, run the **delete** command with other options. To use the **delete** command, you must have the required NetWorker privileges as described in "NetWorker User Group resource" on page 75. The Oracle backup and recovery documentation provides details about the **delete** command.

Note: When deleting a backup, set the NSR_NWPATH parameter in the NWORA resource file or in the RMAN script under these conditions:
- NetWorker client binaries are located in a nondefault directory on the Oracle Server host.
- NWORA resource file was not created.
"NSR_NWPATH" on page 364 provides details.

CHAPTER 4
Data Restore and Recovery

This chapter includes the following topics:

- Data restore and recovery terminology 168
- Performing NMDA data restore and recovery 169
- Performing DB2 data restore and recovery 171
- Performing Informix data restore and recovery 178
- Performing Lotus data restore and recovery 181
- Performing MySQL data restore and recovery 196
- Performing Oracle data restore and recovery 207
- Performing Sybase data restore and recovery 216
- Canceling a data restore or recovery 223

Data restore and recovery terminology

Unlike NetWorker software that uses the term *recover* for all backup retrieval activities, NMDA distinguishes between the restore and recovery of a database:

- *Restore* means to retrieve datafiles from a backup and store the files on a disk.
- *Recover* means to apply the transaction logs, redo logs, logical logs, or binary logs to bring a database to a specific point-in-time and to make the database consistent. When the logs are not available, NMDA can restore data only to the time of the backup. Database vendors use specific terminology to identify the application of logs:
  - DB2—Rollforward
  - Informix—Logical restore
  - Lotus, MySQL, Oracle, Sybase—Recovery

NMDA can restore only the following data:

- Data that you backed up according to instructions in Chapter 3, "Backup Procedures"
- Data that you backed up by using a legacy NetWorker module that NMDA supports

You cannot use the NetWorker server interface or the **recover** program to restore data that you backed up with NMDA. You can restore an NMDA backup only by using specific procedures for the database or application.

NetWorker indexes and policies used for restores

During an NMDA backup, the NetWorker server records specific information in the online indexes as follows:

- The server records information about each backup object in the client file index.
- The server records information about each save set in the media database.

NMDA queries the NetWorker online indexes during a restore.

The NetWorker server maintains a client file index entry until the browse policy specified for the client save set expires. The server maintains a media database entry until the retention policy specified for the client save set expires.

You define the lifecycle by setting the browse policy and retention policy in these locations:

- NetWorker Client resource for scheduled backups.
- NSR_SAVESET_BROWSE and NSR_SAVESET_RETENTION parameters for manual backups. If you do not set these parameters, the NetWorker server assigns the browse policy and retention policy for a manual backup according to settings in the NetWorker Client resource.

When the browse policy expires, the NetWorker software immediately removes the backup entry from the client file index. When the retention policies for all the save sets on a backup volume expire, the volume becomes recyclable and eligible for automatic relabeling by the NetWorker server. However, the save set entries remain in the media database until the volume is relabelled. When the volume is relabelled, the data on the volume becomes inaccessible and you can no longer restore the data.

NMDA uses both the client file index entries and media database entries to restore backup data. If any entries are missing, the restore fails. To prevent a restore failure:

- Set the browse policy to a period long enough to retain the client index entries for restoring the data.
- Set the retention policy to a period equal to or longer than the browse policy period.

After a browse policy expires or the NetWorker indexes are lost, you can use the NetWorker **scanner** program to rebuild the indexes.

To ensure a successful restore and recovery, you must keep the NetWorker online indexes synchronized with the application backup catalogs, as described in "Synchronizing backup catalogs and deleting backups" on page 164.

Note: For Oracle backups, the index entries regenerated by using the **scanner** program might cause the NetWorker indexes to become unsynchronized with the Oracle RMAN catalog, which might cause problems. To prevent problems, ensure that the Oracle backup pieces have unique names, as described in "Oracle backup considerations" on page 128.

The *EMC NetWorker Administration Guide* describes how the NetWorker server uses browse policies and retention policies to manage backup data and track data on volumes.

Performing NMDA data restore and recovery

Before you perform a data restore and recovery, you must meet the requirements for the configurations and volumes.

- You have configured the database system or application system for recovery according to the appropriate vendor documentation.
- The NetWorker server services and client services are running on the required hosts.
- The NetWorker backup volume required for the restore is online and available.

  Note: For tape type devices, you have set the restore parallelism to the number of available devices.

- If you are restoring data to a different destination host than the host you backed up (source client):
  - You have installed and configured NMDA on the destination host.
  - On the NetWorker server that contains the backup to be restored:
    - A Client resource exists for the destination host. The Client resource can contain these attribute settings:
      - Save Set: All
      - Group: (blank)
    - The Remote Access attribute in the Client resource of the source client contains the following value:

      user=*db_or_app_user*,host=*destination_host*

- For a restore from a DD Boost device over Fibre Channel, ensure that the database-specific operating system user has the correct device permissions as described in the following documents:
  - Fibre Channel Devices with Products using DD Boost in Linux Environment (Document ID 95007)
  - Fibre Channel Devices with Products using DD Boost in Windows Environment (Document ID 95005)

  Table 5 on page 77 provides a definition of the database-specific operating system user.

Note: You must perform the restore and recovery on the destination host.

◆ You have set the NMDA parameters for the restore according to Appendix A, "NMDA Parameters and Configuration File":

- You have set the parameters for a DB2 restore in the NMDA configuration file.
- You have set the parameters for an Informix restore in the environment.
- You have set the parameters for a Lotus restore in the NMDA configuration file.
- You have set the parameters for a MySQL restore in the NMDA configuration file.
- You have set the parameters for an Oracle restore in the RMAN script.
- You have set the parameters for a Sybase restore as the **nsrsybrc** command options, if possible. Otherwise, you have set the parameters in the environment.

These parameters are mandatory for specific restores:

- NSR_CLIENT—For a redirected restore to a new destination host, the parameter is set to the hostname of the backed-up (source) client.
- NSR_ENCRYPTION_PHRASES—The parameter is set if both of these conditions are true:
  - NMDA backed up the data with AES encryption (NSR_AES_ENCRYPTION=TRUE).
  - The encryption phrase on the NetWorker server has changed since NMDA backed up the data.

  Note: By default, if you have not set NSR_ENCRYPTION_PHRASES, NMDA obtains the encryption phrase from the NetWorker server for decrypting an AES-encrypted backup during a restore.

- NSR_NWPATH—If you installed the NetWorker client in a nondefault location, you have set the parameter for an Avamar deduplication restore.
- NSR_SERVER—If the NetWorker server is not located on the NMDA host, you have set the parameter to the hostname of the server that contains the backup.

Perform an NMDA data restore by using the appropriate procedure from the following table.

Table 20 Data restore procedures

Database or application	Data restore procedures
DB2	• **db2 restore** command • IBM Data Studio GUI or DB2 Control Center GUI "Performing DB2 data restore and recovery" on page 171 provides details.
Informix	• **onbar** restore command "Performing Informix data restore and recovery" on page 178 provides details.
Lotus	Database restore and DAOS file-level restore: • **nsrnotesrc** command • NetWorker User for Lotus GUI (Windows only) Document-level recovery: • **nsrdocrc** command • Notes client GUI (Windows only) "Performing Lotus data restore and recovery" on page 181 provides details.
MySQL	• **nsrmysqlrc** command for restore, recovery, and other advanced operations: list image, extract, extract and prepare, copy back, validate "Performing MySQL data restore and recovery" on page 196 provides details.
Oracle	• **rman** command to run an RMAN restore script • Oracle Enterprise Manager "Performing Oracle data restore and recovery" on page 207 provides details.
Sybase	• **nsrsybrc** command • NetWorker User for Sybase GUI (Windows only) "Performing Sybase data restore and recovery" on page 216 provides details.

Performing DB2 data restore and recovery

After you determine the number of restore devices and sessions to use, you can run the **db2 restore** command or the DB2 GUI to perform a DB2 data restore. After the restore completes, you can apply the transaction logs to recover the DB2 database.

The DB2 documentation provides details about the different commands and options used for restore and recovery.

Perform the DB2 data restore and recovery by using the appropriate procedures:

- "Determining how many restore devices and sessions to use" on page 172
- "Performing DB2 data restores with the db2 restore command" on page 173
- "Performing DB2 data recovery" on page 176
- "Performing DB2 data restore and recovery with the DB2 GUI" on page 178

Determining how many restore devices and sessions to use

If NMDA used multiple tape devices and multiple sessions to perform a DB2 backup, use the same number of tape devices and sessions for the DB2 restore.

NOTICE

Restore with only one session per tape device because restoring with multiple sessions per tape device can impede performance.

To determine the number of sessions that NMDA used for the DB2 backup, use the **nsrinfo** command:

```
# nsrinfo -s NetWorker_server -n db2 -X all DB2_client | grep
   database_name
```

For example:

```
# nsrinfo -s bu-llet -n db2 -X All bu-gingersnap | grep SAMPLE

version=1,  DB2, objectname=/SAMPLE/NODE0000
   /DB_BACKUP.20100621171932.3, createtime=Mon Jun 21 17:19:34 2010,
   copytype=BSACopyType_BACKUP, copyId=1277155174.1277155175,
   restoreOrder=1277155174.1, objectsize=0.0, resourcetype=database,
   BSAObjectType_DATABASE,  BSAObjectStatus_ACTIVE,
   description=NMDA_v11:DB2_v970:FULL_BACKUP:SAMPLE:TNE,
   objectinfo=db2inst1:3

version=1,  DB2, objectname=/SAMPLE/NODE0000
   /DB_BACKUP.20100621171932.2, createtime=Mon Jun 21 17:19:33 2010,
   copytype=BSACopyType_BACKUP, copyId=1277155173.1277155174,
   restoreOrder=1277155173.1, objectsize=0.0, resourcetype=database,
   BSAObjectType_DATABASE,  BSAObjectStatus_ACTIVE,
   description=NMDA_v11:DB2_v970:FULL_BACKUP:SAMPLE:TNE,
   objectinfo=db2inst1:3

version=1,  DB2, objectname=/SAMPLE/NODE0000
   /DB_BACKUP.20100621171932.1, createtime=Mon Jun 21 17:19:32 2010,
   copytype=BSACopyType_BACKUP, copyId=1277155172.1277155173,
   restoreOrder=1277155172.1, objectsize=0.0, resourcetype=database,
   BSAObjectType_DATABASE,  BSAObjectStatus_ACTIVE,
   description=NMDA_v11:DB2_v970:FULL_BACKUP:SAMPLE:TEQ,
   objectinfo=db2inst1:3

3 objects found
```

The objectinfo value shows the number of sessions. For example:

```
objectinfo=db2inst1:3
```

where:

- db2inst1 is the name of the backed-up database instance.
- 3 is the number of sessions that NMDA used for the backup.

A restore operation for this example would use three sessions and three tape devices with one session per tape device.

Performing DB2 data restores with the db2 restore command

You can run the appropriate **db2 restore** command from the command line to perform a DB2 data restore to either the same DB2 server host or a different host.

A DB2 restore can restore the data to the original database or to a different database under the same or different DB2 instance.

Performing DB2 data restores to the same instance

1. Log in to the DB2 client host as the DB2 operating system user.
2. If you are restoring the most recent backup of a database or tablespace, skip this step. Otherwise, if you are recovering the data to a point-in-time, note the timestamp of the backup to restore.

   If the timestamp of the backup is unknown, find the timestamp by querying all the backups with the following command:

   ```
   $ db2 list history backup all for sample
   ```

   where *sample* is the name of the database to be restored.
3. If NMDA used multiple sessions for the backup, note the number of sessions used. You must specify the number as the **open sessions** value in the **db2 restore** command. "Determining how many restore devices and sessions to use" on page 172 provides details.
4. Run the **db2 restore** command with the appropriate options, as described in the DB2 documentation. For example:

   - On UNIX, type the command:

     ```
     $ db2 restore db sample load /usr/lib/libnsrdb2.so open n sessions
     options @pathname/nmda_db2.cfg taken at yyyymmddtttt into sample2
     ```

   - On Windows, type the command:

     ```
     $ db2 restore db sample load
     NetWorker_install_dir\nsr\bin\libnsrdb2.dll open n sessions
     options @pathname\nmda_db2.cfg taken at yyyymmddtttt into sample2
     ```

     where:

     – *sample* is the name of the database to be restored.

     – *n* is the number of restore sessions, if NMDA used multiple sessions for the backup. This number can range from 1 to the number from step 3.

       Note: Use only one restore device per restore session.

     – *pathname*\nmda_db2.cfg is the complete pathname of the NMDA configuration file.

     – *yyyymmddtttt* is the timestamp of the backup to restore, as noted in step 2.

       Skip the **taken at** parameter if you are restoring only the most recent backup of a database or tablespace.

- *sample2* is the new name of the database, if you are restoring to a different database name.

  Skip the **into** parameter if you are restoring the database to the original database name.

- *NetWorker_install_dir* is the Windows system path of the NetWorker location.

Performing DB2 data restores to a different instance

1. Log in to the DB2 client host as the DB2 operating system user.
2. If you are restoring the most recent backup of a database or tablespace, skip this step. Otherwise, if you are recovering the data to a point-in-time, note the timestamp of the backup to restore.

   If the timestamp of the backup is unknown, find the timestamp by querying all the backups with the following command:

   ```
   $ db2 list history backup all for sample
   ```

   where *sample* is the name of the database to be restored.
3. If NMDA used multiple sessions for the backup, note the number of sessions used. You must specify the number as the **open sessions** value in the **db2 restore** command. "Determining how many restore devices and sessions to use" on page 172 provides details.
4. Grant the new instance, for example, db2inst2, permission to restore the database:

   a. Start the NMC program, and open the NetWorker Client resource for the original backed-up host.

   b. In the Application Information attribute, set the following value:

      ```
      DB2_R=sample:db2inst1:db2inst2:
      ```

      where:

      - *sample* is the database name.
      - *db2inst1* and *db2inst2* are the names of the instances with permissions to restore the database. *db2inst1* is the old instance, and *db2inst2* is the new instance.

        Note: Separate each instance with a colon (:) and insert a colon after the last instance.
        You need the Configure NetWorker privilege to modify the Application Information attribute.

5. From the new instance, generate a redirection script by running the **db2 restore** command with the **redirect generate script** option. For example:

   - On UNIX, type the command:

     ```
     $ db2 restore db sample load /usr/lib/libnsrdb2.so options
     @pathname/nmda_db2.cfg taken at yyyymmddtttt redirect generate
     script pathname/my_redirect.ddl
     ```

   - On Windows, type the command:

     ```
     $ db2 restore db sample load
     NetWorker_install_dir\nsr\bin\libnsrdb2.dll options
     @pathname\nmda_db2.cfg taken at yyyymmddtttt redirect generate
     script pathname\my_redirect.ddl
     ```

     where:

     – *pathname**my_redirect.ddl* is the complete pathname of the generated redirection script.

     – The other command line options are the same as described in "Performing DB2 data restores to the same instance" on page 173.

   Note: Ensure that the new instance has read and write permission to the script.

6. Edit the generated script, and define the following parameters:

   - OPTIONS (mandatory)—Complete pathname of the NMDA configuration file.
   - ON (mandatory)—Complete pathname of the new database instance.
   - INTO—New database name, if you are redirecting the recovery to a new name.
   - TAKEN AT—Timestamp of the backup to recover, *yyyymmddtttt*, if you are restoring the data to a point-in-time, as noted in step 2.
   - OPEN SESSIONS—Number of restore sessions, if NMDA used multiple sessions for the backup, as determined in step 3. Use only one restore device per restore session.

   For example:

   ```
   OPTIONS '@/bigspace/home/db2inst2/nmda_db2.cfg'
   ON '/bigspace/home/db2inst2'
   INTO sample2
   ```

   Note: If you created DMS tablespaces with the backup, you might need to set the SET TABLESPACE CONTAINERS parameter to the appropriate value.

   The DB2 documentation provides details.

7. On the DB2 server host, type the following command to run the redirection script under the redirected different instance where the data is to be restored:

   ```
   $ db2 -tvf my_redirect.ddl
   ```

   where *my_redirect.ddl* is the name of the generated redirection script.

Performing DB2 data recovery

You can run the **db2 rollforward** command to apply the transaction logs that are stored on the backup media to recover a DB2 database to either the current time or a specific point-in-time.

You can optionally prefetch the transaction logs from backups to local storage first and then run the **db2 rollforward** command to roll forward the logs from the local copy.

If you want to restore and recover your DB2 database in a single operation, you can run the **db2 recover** command instead.

Note: To use rollforward recovery, NMDA must have backed up the transaction logs. "DB2 transaction log backups" on page 35 provides details.

Performing DB2 recovery with the db2 rollforward command

When you perform a rollforward recovery of a restored DB2 database without prefetching the transaction logs, DB2 restores the logs from the NMDA backups and applies the logs to recover the database to either the current time or a point-in-time.

To apply all the transactions to the end of the logs, run the following command:

```
$ db2 "rollforward db sample to end of logs and complete"
```

where *sample* is the database name.

To apply all the transactions to a specific point-in-time, run the following command:

```
$ db2 "rollforward db sample to yyyy-mm-dd.hh.mm.ss using local time"
```

Performing DB2 recovery with fetched logs and the db2 rollforward command

A fetched logs rollforward recovery of a restored DB2 database provides the benefits of convenience and speed. You can prefetch the transaction logs from the NMDA backups to local storage and then roll forward the logs from the local copy. This type of rollforward recovery facilitates the selection of specific logs, and can be faster than the rollforward of logs stored on the backup media, especially tape media.

1. Run the **nsrdb2rlog** command to retrieve or fetch a copy of the DB2 transaction logs from the NetWorker server to a local file system.

   Note: To list the logs on the NetWorker server, use the **nsrinfo** command.

   For example, to retrieve all the transaction logs for a database to the end of the logs, run the following command:

   ```
   $ nsrdb2rlog -s server -a sample -d destination_dir
   -z configuration_file
   ```

   where:

   - *server* is the name of the host on which the database resides.
   - *sample* is the name of the database that the logs belong to.
   - *destination_dir* is the directory where the log files will be recovered.
   - *configuration_file* is the complete pathname of the NMDA configuration file.

The **nsrdb2rlog** man page and the *EMC NetWorker Module for Databases and Applications Command Reference Guide* describe the command.

2. Complete the recovery by updating the database with the retrieved transaction logs. The following examples describe how to apply the transaction logs.

   To apply all transactions to the end of the logs, run the following command:

   ```
   $ db2 rollforward db sample to end of logs and complete overflow log
   path (c:\log_path)
   ```

   To apply all transactions to a specific point-in-time, run the following command:

   ```
   $ db2 rollforward db sample to yyyy-mm-dd.hh.mm.ss using local time
   overflow log path (c:\log_path)
   ```

   where in both of these examples:

   - **c:***log_path* is the complete pathname of the retrieved transaction log file stored locally on the DB2 host.
   - *yyyy-mm-dd.hh.mm.ss* is the date format and time format.

Performing DB2 restore and recovery with the db2 recover command

The **db2 recover** command combines the functions of the **db2 restore** command and **db2 rollforward** command. You can run the **db2 recover** command to restore a backed-up database or tablespace, with the transaction logs applied to a specific point-in-time.

1. Set the DB2 database configuration VENDOROPT parameter to the pathname of the NMDA configuration file (nmda_db2.cfg) for the database to be recovered. For example:

   ```
   db2 update db cfg for sample using vendoropt
   @/db/pathname/nmda_db2.cfg
   ```

   where:

   - *sample* is the name of the database or tablespace to be recovered.
   - *pathname*/nmda_db2.cfg is the complete pathname of the NMDA configuration file. Do not use a relative pathname.

2. Run the **db2 recover** command with appropriate options to recover the database or tablespace to the end of the logs or to a specific point-in-time. The following examples describe the commands to use in each case.

   To apply all transactions to the end of the logs, run the following command:

   ```
   $ db2 recover db sample to end of logs
   ```

   To apply transactions to a specific point-in-time, run the following command:

   ```
   $ db2 recover db sample to yyyy-mm-dd.hh.mm.ss using local time
   ```

   The command line options in these examples are the same as described in "Performing DB2 data restores to the same instance" on page 173.

Note: The **db2 recover** command does not support the **load** syntax or **options** syntax that is available with **db2 backup** and **db2 restore** commands. Instead, the **db2 recover** command uses information in the DB2 history file to determine what file to load during the recovery and uses the VENDOROPT variable to pass the options file. For a dropped database, use the **db2 restore** and **db2 rollforward** commands to perform disaster recovery. You cannot use the **db2 recover** command for dropped databases.

Performing DB2 data restore and recovery with the DB2 GUI

Depending on the DB2 version, you can run the appropriate DB2 GUI to perform a DB2 data restore or recovery:

- With DB2 version 10.1 or later, run the IBM Data Studio GUI.
- With previous versions of DB2, run the DB2 Control Center GUI.

Specify the NMDA configuration file in the GUI by setting VENDOROPT to the value *@configuration_file_pathname*. For example:

```
@d:\nmda_db2.cfg
```

Note: The NMDA configuration file must contain the parameter settings for the restore.

Set the Vendor DLL to the NMDA DB2 library name in the GUI.

Note: After you select the DLL path with the DB2 Control Center for a Windows client, enclose the path with quotes or use a short file name (8.3 format). Otherwise, the restore returns an error similar to the following example:

```
SQL0104N  An unexpected token
"Files\Legato\nsr\bin\libnsrdb2.dll" was found following "<identifier>".
Expected tokens may include:  "INCLUDE".
```

Performing Informix data restore and recovery

With the Informix database server in the appropriate mode, you can run the **onbar** command from the command line to perform an Informix data restore and recovery.

Perform the Informix data restore and recovery by following these steps:

1. “Determining the Informix restore mode” on page 179
2. “Performing Informix data restores with the onbar command” on page 180

“Performing an Informix imported restore” on page 229 describes how to restore data from one Informix server instance to the same instance on a different destination host.

Determining the Informix restore mode

You can perform an Informix restore with the database server in one of three modes: cold, warm, or mixed. Each of the restore types consists of specific physical restores and logical restores with the **onbar** command.

- **Cold restore mode**—Restores both critical and noncritical data when the database server is offline. A cold restore performs the following operations:

  1. A backup of the logs that have yet to be backed up.
  2. A physical restore and logical restore of the critical dbspaces.
  3. A physical restore and logical restore of the noncritical dbspaces.

  After a cold restore completes, the database server remains in quiescent mode.

  Note: A cold restore of selected dbspaces succeeds only if the restore command line includes the critical dbspaces. Critical dbspaces are defined as the root dbspace and any dbspace that contains either physical logs or logical logs.

- **Warm restore mode**—Restores noncritical data while the database server is online or quiescent. A warm restore performs the following operations:

  1. A backup of the logs that have yet to be backed up.
  2. One or more physical restores.
  3. A closing and backup of the current logical log.
  4. A logical log restore.

- **Mixed restore mode**—Enables the quick recovery of critical dbspaces, plus any data to which users require immediate access. A mixed restore performs the following operations:

  1. A cold restore of the critical dbspaces, with the database server in offline mode.
  2. A warm restore of noncritical dbspaces, with the database server in online or quiescent mode.

  After the database server returns to quiescent mode, you must perform a warm restore of the other dbobjects.

Note: The **onbar** cold, warm, mixed, and point-in-time restore modes require enabled versions of the NetWorker and NMDA software to enable the backup of outstanding logs. The **onbar** restore first performs a backup of outstanding logs before starting the restore.

The ON-Bar utility maintains a history of backup and restore operations in the sysutils database, and stores an extra copy of the backup history in the emergency boot file. ON-Bar uses the sysutils database in a warm restore when only a portion of the data is lost. ON-Bar uses the emergency boot file in a cold restore when the sysutils database is inaccessible.

You can use the Informix **onsmsync** utility to regenerate the emergency boot file and expire old backups.

Performing Informix data restores with the onbar command

You can run the **onbar** command from the command line as the Informix user, which runs the ON-Bar utility, to perform an Informix restore of database server instances or database objects.

Use the appropriate environment variables and the **onbar** command to perform the required type of restore.

- **Physical restore**—Replaces lost or corrupted Informix dbobjects from the NetWorker backup media. You can perform this type of restore as a whole-system restore or selected-dbspace restore.

  For example:

  1. Set the following environment variables in the shell or in the command window:

     ```
     NSR_SERVER=mars
     INFORMIXDIR=C:\Progra~1\IBM\Informix
     ONCONFIG=ONCONFIG.ifxm1
     ```

     Note: On UNIX or (only with Informix 12.10 or later) on Windows, you must also set the INFORMIXSQLHOSTS variable.

  2. Run the following **onbar** command:

     onbar -r -p [*dbspace_name*]

- **Logical Restore**—Recovers the server transactions completed since the last dbobject backup and then rolls forward the logical log files backed up for the dbobjects. If different backup sessions are involved, the log rolls forward transactions made since the backup time recorded for each dbobject restored.

  For example:

  1. Set the following environment variables in the shell or in the command window:

     ```
     NSR_SERVER=mars
     INFORMIXDIR=C:\Progra~1\IBM\Informix
     ONCONFIG=ONCONFIG.ifxm1
     ```

     Note: On UNIX or (only with Informix 12.10 or later) on Windows, you must also set the INFORMIXSQLHOSTS variable.

  2. Run the following **onbar** command:

     onbar -r -l

- **Combined Restore**—Enables you to issue a single command to perform a physical restore, immediately followed by a logical restore.

  For example:

  1. Set the following environment variables in the shell or in the command window:

     ```
     NSR_SERVER=mars
     INFORMIXDIR=C:\Progra~1\IBM\Informix
     ONCONFIG=ONCONFIG.ifxm1
     ```

Note: On UNIX or (only with Informix 12.10 or later) on Windows, you must also set the INFORMIXSQLHOSTS variable.

2. Run the following **onbar** command:

```
onbar -r [dbspace_name]
```

- **Point-in-time restore**—Performs a whole-system physical restore of the database server data from a whole-system backup to a specified time instead of the default time. The default time is the time of the last database server backup.

  For example:

  1. Set the following environment variables in the shell or in the command window:

```
NSR_SERVER=mars
INFORMIXDIR=C:\Progra~1\IBM\Informix
ONCONFIG=ONCONFIG.ifxm1
```

  Note: On UNIX or (only with Informix 12.10 or later) on Windows, you must also set the INFORMIXSQLHOSTS variable.

  2. Run the following **onbar** command:

```
onbar -r -t "2010-06-24 00:00" -w -p
```

Performing Lotus data restore and recovery

Before you perform a Lotus data restore and recovery, you must meet the following requirements:

- For restore and recovery operations with a Domino server on Linux or Solaris, set the environment variable LD_LIBRARY_PATH to the complete pathname of the Lotus directory that contains the library files libnotes.so, libndgts.so, and libxmlproc.so. Set the variable in the same shell in which you perform the operation. For example:

```
export LD_LIBRARY_PATH=/opt/lotus/notes/latest/sunspa
```

- To prepare for an in-place recovery of a logged database when the original database is still present, perform one of the following actions:
  - Delete the original database.
  - Direct the recovery to a new directory with the NSR_RELOCATION_DEST parameter, and change (zap) the database instance ID (DBIID) with the NSR_DBIID parameter in the NMDA configuration file.

  NOTICE

  After you change the DBIID of a recovered database, you cannot recover subsequent changes to the database until the next full backup is complete. Perform a full backup after changing the DBIID of a recovered database. If an incremental backup is the next backup of a database after a DBIID change, NMDA automatically performs a full backup of the database instead.

- Lotus transactional logging is set on the Domino server for the required type of recovery operation:
  - If you have enabled Lotus archived logging, you can recover a database or document to any point-in-time, provided that a specific backup is available and the transaction logs for that period are available.

    "Setting the NSR_RECOVER_TIME parameter" on page 379 describes how to determine the recovery point-in-time.
  - If you have disabled Lotus archived logging, you can restore a database or document to the time of a specific backup only.

  If transaction logs are missing on the system and need to be restored from the NetWorker system, you can set the TRANSLOG_RECOVER_PATH parameter. Set the parameter in the notes.ini file to the full pathname of the alternative log location. Setting this parameter improves the restore performance by preventing the logs that will be restored from interfering with the logs generated by Domino operations. The IBM documentation describes the parameter.

NOTICE

When you recover a logged database to a different Domino server than the backed-up server, you cannot apply the transaction logs because the logs are not available in the new destination.
If you recover a logged database to a new destination host, either perform disaster recovery first to restore the transaction logs to the new destination or perform one of the following actions:
- In the NetWorker User for Lotus GUI, select **Do not apply Transaction Logs** in the **Recover Options** dialog box.
- In the NMDA configuration file, set NSR_APPLY_LOGS=FALSE.

Perform the Lotus data restore and recovery by using the appropriate procedures:

- "Performing Lotus database recovery with the nsrnotesrc command" on page 182
- "Performing Lotus database recovery with NetWorker User for Lotus" on page 185
- "Performing recovery of partitioned Domino servers" on page 190
- "Performing Lotus DAOS data recovery" on page 191
- "Performing Lotus document-level recovery with nsrdocrc command" on page 192
- "Performing Lotus document-level recovery with the Notes client GUI" on page 194

Performing Lotus database recovery with the nsrnotesrc command

You can run the **nsrnotesrc** command from the command line to restore and recover Lotus databases and DAOS files.

On UNIX or Linux, you must run the **nsrnotesrc** command as the Lotus user that starts the Domino server. Do not run the command as the root user.

You must set NSR_BACKUP_PATHS or NSR_RECOV_LIST_FILE, not both parameters, in the NMDA configuration file to specify the Lotus database files or the directories for restore.

You must also set the mandatory parameters Notes_ExecDirectory, NSR_NOTES_INI_PATH, and PATH for the restore. You can set other optional parameters as appropriate.

Note: File paths are case-sensitive for all operating systems, including Windows. The file paths must match the case of the entries in the NetWorker client index. If you are uncertain about the case of file paths, use the **nsrinfo** command to verify the backup entries in the NetWorker client file index.

Ensure that you set all the parameters required for recovery in the configuration file as described in Appendix A, "NMDA Parameters and Configuration File."

To perform the Lotus database recovery, run the **nsrnotesrc** command from the command line:

```
nsrnotesrc(.exe) -z configuration_file_path
```

where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the recovery parameter settings.

If the **nsrnotesrc** command prompts with a message that an existing file has the same name as a file being recovered, reply with the value **n**, **N**, **y**, **Y**, **r**, or **R**, as described in "NSR_RECOV_INTERACT" on page 376.

The following examples describe the parameters set in the NMDA configuration file for different types of Lotus database recoveries. After you set the parameters, you can run the **nsrnotesrc -z** command to perform the recovery.

Example 13 **Recovery of a database to a specific point-in-time**

By default, the NMDA software restores the most recent Lotus backup available. NMDA requests that Domino apply the logs up to the current time if both of these conditions are true:

- You have enabled archived transaction logging.
- The required transaction logs are on the Domino system or are available from a backup.

The NSR_RECOVER_TIME parameter enables you to recover the database to a point-in-time before the current time. If transaction logs are available, NMDA recovers the database to the time set by the parameter. Otherwise, NMDA restores the database to the last backup before or equal to the time set by the parameter. The parameter value is the time in **nsr_getdate** format. For example:

```
NSR_RECOVER_TIME = "Wed June 23 2012 14:23"
```

"Setting the NSR_RECOVER_TIME parameter" on page 379 describes how to determine the recovery point-in-time.

Example 14 **Recovery of specific Lotus data files or directories**

Use the NSR_BACKUP_PATHS parameter to specify a list of files and directories to recover. When you specify a directory, NMDA recovers all the files in that directory and its subdirectories that you backed up. You can include both data files and directories in a single list. For example:

- On UNIX:

```
NSR_BACKUP_PATHS = /lotusdata/account.nsf, /lotusdata/mail/
```

- On Windows:

```
NSR_BACKUP_PATHS = C:\Lotus\Domino\data\account.nsf,
C:\Lotus\Domino\data\mail\
```

Example 15 **Recovery of all Lotus database files**

The following parameter setting in the NMDA configuration file specifies the recovery of all the Lotus database files that you backed up from the given NetWorker client:

```
NSR_BACKUP_PATHS = NOTES:
```

NOTICE

Use the NOTES: option with caution. NMDA attempts to restore the data for all partitions of a partitioned Domino server, or for all Domino installations when there are multiple Domino installations on the client.

Example 16 **Relocated data restore**

By default, NMDA restores data files to the same location from which you backed them up. Set the NSR_RELOCATION_DEST parameter to specify a different destination directory. For example:

```
NSR_RELOCATION_DEST = /newdata/
```

The relocation of multiple files with NSR_RELOCATION_DEST preserves the relative directory structure of the files. When you use NSR_RELOCATION_DEST to specify the new location /newdata, NMDA restores the files /olddata/mail1/file1.nsf and /olddata/mail2/file2.nsf to /newdata/olddata/mail1/file1.nsf and /newdata/olddata/mail2/file2.nsf.

If you also set NSR_RECOVER_OPTIONS=REMOVE_COMMON_PATH for the relocated restore, NMDA removes the common path from the original file pathnames during the file restore to the new directory. If you set the parameter in this example, NMDA removes the common directory path, olddata, and restores the files to /newdata/mail1/file1.nsf and /newdata/mail2/file2.nsf.

Note: If you recover logged databases, set the NSR_DBIID parameter to change the DBIID of the recovered files. Example 19 on page 185 describes NSR_DBIID.

NOTICE

During a restore to a different directory, the **nsrnotesrc** command does not prompt when an existing file has the same name as the recovered file. Instead, the restore overwrites the existing file.

Example 17 **Restore of a logged database without applying transaction logs**

By default, NMDA restores and recovers a logged Domino database to the current time or to a time set through the NSR_RECOVER_TIME parameter. Set the NSR_APPLY_LOGS parameter to FALSE if you want to restore a logged database only *without* applying the transaction logs:

```
NSR_APPLY_LOGS = FALSE
```

Example 18 **Relocation of a linked database during Lotus recovery**

By default, when the **nsrnotesrc** command recovers a Lotus link file, the command also recovers the database or directory that the link points to. The recovery occurs to the location from where the database or directory was backed up.

If the NSR_RELOCATION_DEST parameter is set in the NMDA configuration file, the database file that a link points to is recovered to the specified relocation directory.

For example, the link file /space1/notes/data/link.nsf points to the database file /space2/notes/data.nsf. A backup includes both the link file and the database. NSR_RELOCATION_DEST is set to /space3/new.

The recovery relocates the link file and database file to these locations:

- Recovered link file: /space3/new/space1/notes/data/link.nsf
- Recovered database file (pointed to by the link file): /space3/new/space2/notes/data.nsf

Example 19 **PIT recovery of a database with a change of DBIID to a new directory**

The following parameter settings in the NMDA configuration file specify the point-in-time (PIT) recovery of the logged database named budget2010.nsf to a new directory C:\Lotus\Domino\Data\tmpdir. The operation requires a change of the DBIID if the original database exists on the system:

```
LOTUS {
   Notes_ExecDirectory = C:\Lotus\Domino\Data
   NSR_BACKUP_PATHS = C:\Lotus\Domino\Data\budget2010.nsf
   NSR_DBIID = 1
   NSR_RELOCATION_DEST = C:\Lotus\Domino\Data\tmpdir
   NSR_RECOVER_TIME = "Wed June 23 2012 14:23"
}
```

Since the DBIID of the database changes after the recovery, you must back up the new database if you plan to use it.

"Setting the NSR_RECOVER_TIME parameter" on page 379 describes how to determine the recovery point-in-time.

Performing Lotus database recovery with NetWorker User for Lotus

On Windows systems only, you can run the NetWorker User for Lotus GUI (**nwbml.exe**) to recover Lotus Domino/Notes database files to either the local host or a remote host.

Perform the Lotus data recovery by using the appropriate GUI procedure:

- "Performing Lotus recovery of local data with the GUI" on page 186
- "Performing Lotus directed recovery with the GUI" on page 189

Performing Lotus recovery of local data with the GUI

You can use the NetWorker User for Lotus GUI to recover data on the local Windows host.

1. Start the GUI. For example, select **NetWorker User for Lotus** from **Start** › **Programs** › **EMC NetWorker**.
2. To connect to a different NetWorker server, complete these steps:
   a. Select **Select NetWorker Server** from the **Operation** menu.

      The **Change Server** dialog box appears.

   b. Click **Update List** to refresh the list of NetWorker servers.

   c. Select or type the name of the server.

   d. Select **Save as Default Server** to use the server as the default NetWorker server.

   e. Click **OK**.
3. In the NetWorker User for Lotus GUI, select **Recover** from the **Operation** menu.

   The **Recover** window appears as shown in the following figure. The online help describes the toolbar buttons.

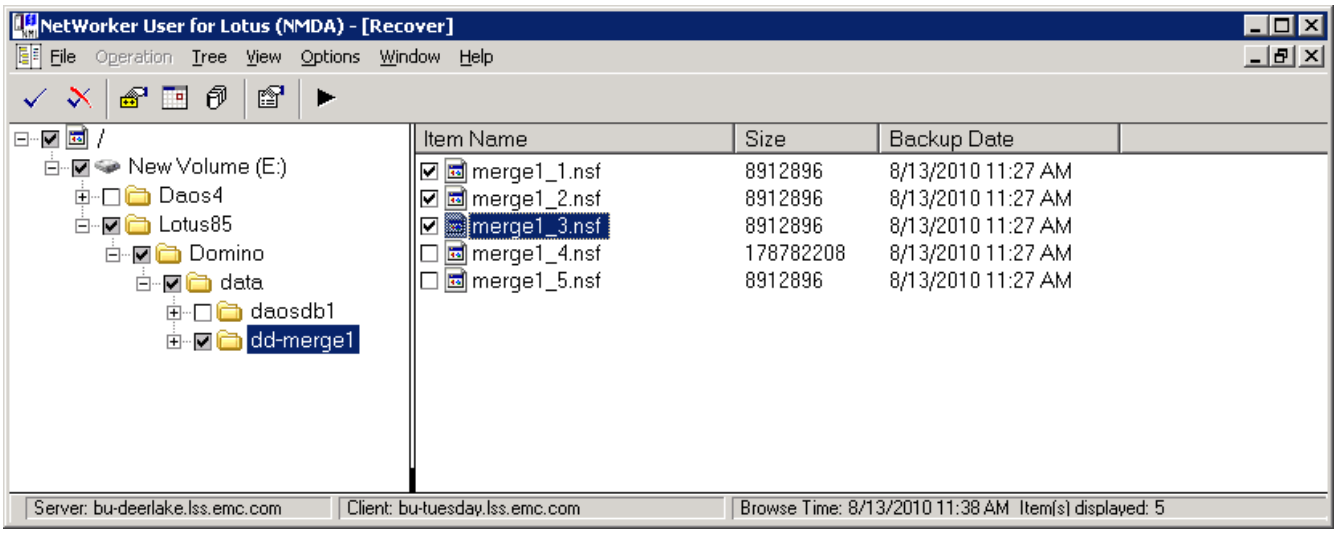

Figure 6 Recover window in NetWorker User for Lotus

4. To view a list of files or databases available for restore, select a Lotus directory in the left pane. The Lotus directory contents appear in the right pane. By default, the GUI shows the latest available backups.

   Note: When a Lotus file does not exist on the computer, the corresponding icon might appear incorrectly as a folder icon in the **Recover** window. The incorrect icon display does not affect the success of the recovery.

5. To view the available versions of backed-up data:
   a. In the **Recover** window, select a database file or a directory.

   b. Select **Versions** from the **View** menu.

The **Versions** window appears and includes the backup history of the selected object. The list sorts the versions according to backup time with the most recent backup at the top of the list.

6. To view previous backups for recovery to a previous point-in-time, change the browse time:
   a. In the **Recover** window, select **Change Browse Time** from the **View** menu.

      The **Change Browse Time** dialog box appears.

   b. Set a new date by selecting a day from the calendar.
   c. Click **Previous Month** or **Next Month** to change from the current month.
   d. In the **Time** text box, type a time to browse.

Note: The browse time cannot be earlier than the time of the first backup because the client file index does not have entries before that time. To verify the browse policy and retention policy, check the Client resources for the client by using the NetWorker administration program.

7. Select the checkbox next to each file or database to be recovered.

Note: Do not set NSR_BACKUP_PATHS in the NMDA configuration file. Instead, use the GUI to select the files and directories to be recovered.

8. From the **Options** menu, select **Recover Options**.

   The **Recover Options** dialog box appears as shown in the following figure.

Note: All the values set in the **Recover Options** dialog box take precedence over the corresponding parameters set in the NMDA configuration file.

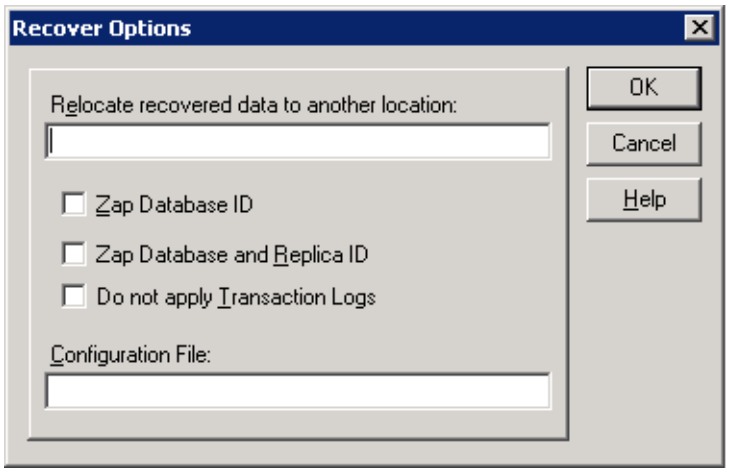

Figure 7 Recover Options dialog box in NetWorker User for Lotus

9. In the **Recover Options** dialog box, specify the required options:

   a. To specify the location of the NMDA configuration file, type the complete pathname in the **Configuration File** box. "NMDA configuration file" on page 354 provides details about the NMDA configuration file.

   NOTICE

   For Lotus data recovery, you must store the configuration file on the destination client where the database files are to be restored.

   The Notes_ExecDirectory parameter is the only Lotus restore parameter that is mandatory in the configuration file. You can set specific parameters in the configuration file so that you do not need to set values in the dialog box every time you run a recovery.

   b. Complete other fields in the **Recover Options** dialog box, if required. The fields correspond to parameters in the NMDA configuration file as described in the following table. "Performing Lotus database recovery with the nsrnotesrc command" on page 182 includes examples of the parameters.

Table 21 Configuration file parameters corresponding to Recover Options fields

Field in Recover Options dialog box	Corresponding parameter in configuration file
Relocate recovered data to another location	NSR_RELOCATION_DEST= *relocation_directory*
Zap Database ID	NSR_DBIID = 1
Zap Database and Replica ID	NSR_DBIID = 2
Do not apply Transaction Logs	NSR_APPLY_LOGS = FALSE

   c. Click **OK**.

10. Ensure that you have set the backup volumes for the recovery:

    a. In the **Recover** window, ensure that you have selected the required entries for recovery.

    b. From the **View** menu, select **Required Volumes**.

    The **Required Volumes** window appears with the backup volumes listed.

    c. Load and mount the required volumes, as appropriate.

11. Click **Start** in the **Recover** window.

    The **Recover Status** window appears with information about the recovery.

Performing Lotus directed recovery with the GUI

For a Lotus directed recovery, you can use the NetWorker User for Lotus GUI on one host to start the recovery of Lotus database files from a different host or to a different host.

The following terms refer to the client computers in the directed recovery:

- Source client—The NetWorker client whose data is to be restored. The source client name must match the name of the NetWorker client file index that contains the backup entries to be restored.
- Destination client—The host to which the Lotus data will be recovered.
- Performing (administrative) client—The Windows host where you run the NetWorker User for Lotus.

Complete these procedures to perform a Lotus directed recovery with the NetWorker User for Lotus GUI:

1. "Configuring Lotus directed recovery" on page 189
2. "Performing Lotus directed recovery" on page 190

Configuring Lotus directed recovery

- When the administrative client or destination client is different than the source client:
  1. On the NetWorker server that has the backup to restore, create a NetWorker Client resource for the administrative client or destination client, if one does not yet exist.
  2. In the Remote Access attribute of the Client resource for the source client, add one of the following lines for the valid username and hostname of the administrative client or destination client:

     ```
     user=username, host=administrative_client_hostname
     user=username, host=destination_client_hostname
     ```

- When the administrative client is different than the destination client:
  1. Edit the **nsrlotus_remrecov** script on the destination client. The comments in the script file provide details about editing the script.

     Note: The script file is in the NMDA installation directory. On Windows, the script file name is nsrlotus_remrecov.bat. On UNIX, the script file name is nsrlotus_remrecov. Do not move the script from the installation directory, but you can rename the script if you ensure that the name starts with nsr.

  2. On the administrative client, before starting the NetWorker User for Lotus, set the REMOTE_RECOVCMD environment variable to the file name of the **nsrlotus_remrecov** script that you modified on the destination client. If you have not set the variable, the NetWorker User for Lotus uses **nsrlotus_remrecov**, by default.

Note: Set REMOTE_RECOVCMD in the environment of the administrative client, not in the NMDA configuration file, and include any file name extension (for example, .bat), as required. For example:

```
REMOTE_RECOVCMD=nsrlotus_remrecov.bat
```

The script name must not be in quotes. On Windows, the script name must not contain any spaces.

3. In the /nsr/res/servers file on a destination client only, specify the hostname of the administrative client.

Performing Lotus directed recovery

1. On the administrative client, start the NetWorker User for Lotus program.
2. From the **Operation** menu, select **Directed Recover**.

   The **Source Client** window appears.
3. In the **Source Client** window, select the source client to recover from, and click **OK**.

   The **Destination Client** window appears.
4. In the **Destination Client** window, select the client to recover to.
5. In the **Recover** window, select the files to recover and specify the recovery settings.

   If required, set the recovery options by selecting **Recover Options** from the **Options** menu. In the **Recover Options** dialog box, if you specify the NMDA configuration file in the **Configuration File** box, type the pathname of the file on the destination client.

   Step 4 to step 10 in the section "Performing Lotus recovery of local data with the GUI" on page 186 provide details.

   NOTICE

   When you perform the directed recovery to a UNIX or Linux destination client, the total length of all pathnames marked for restore must not exceed the limit of 10 KB. Otherwise, the directed recovery fails.
   When the administrative client is different than the Windows destination client, if you select to restore any Lotus file that contains a space in its pathname, the directed recovery fails.
6. Click **Start** in the **Recover** window.

Performing recovery of partitioned Domino servers

To enable recovery of partitioned Domino servers, set the required parameters in the NMDA configuration file:

1. Set the PATH parameter to include the directory of the partitioned Domino server for recovery that contains the notes.ini file.
2. Ensure that the PATH parameter from step 1 does not include the data directories of other partitioned Domino servers.

3. When you recover all the Lotus data for a partitioned Domino server, set the NSR_BACKUP_PATHS parameter to specify each top-level directory that contains the data for the partition for recovery. For example:

   ```
   NSR_BACKUP_PATHS = NOTES:M:\Lotus\p1data
   ```

   Note: Do not specify only NOTES: for the NSR_BACKUP_PATHS parameter because NMDA attempts to recover all the Lotus data for all the Domino servers on the host.

Appendix A, "NMDA Parameters and Configuration File," describes the parameters in the NMDA configuration file.

Performing Lotus DAOS data recovery

To restore Domino databases and the corresponding missing NLO files, complete these steps:

1. Restore the Domino databases first by using a regular NMDA Lotus restore, either through the **nsrnotesrc –z** *pathname_to_nmda_lotus.cfg* command or through the NetWorker User for Lotus program.

   Note: NMDA ignores the parameter settings in the LOTUS_DAOS{} section of the NMDA configuration file during restore operations.

2. Run the Domino **tell** command to determine the missing NLO files for the restored databases. For example:

   ```
   tell daosmgr listnlo -o output_file MISSING database.nsf
   ```

3. Create a configuration file for the DAOS recovery or modify the existing nmda_lotus.cfg file used in step 1. When restoring through the **nsrnotesrc** command, specify either one of the following parameters in the LOTUS{} section of the NMDA configuration file:

   ```
   NSR_RECOV_LIST_FILE = full_path_to_output_file_from_tell_command
   ```

   or

   ```
   NSR_BACKUP_PATHS =
   list_of_NLO_files_from_tell_command_separated_by_commas
   ```

   "NSR_BACKUP_PATHS" on page 373 and "NSR_RECOV_LIST_FILE" on page 376 describe the parameters.

   If you use the NetWorker User for Lotus program to restore the missing DAOS files, do not set either of the parameters in the configuration file.

4. Perform the second restore to restore the missing NLO files, either through the **nsrnotesrc –z** *pathname_to_nmda_lotus.cfg* command or through the NetWorker User for Lotus program.

5. After restoring the DAOS directory or NLO files, resynchronize the DAOS repository, for example, by running the following command:

   ```
   tell daosmgr resync [force]
   ```

   This command also updates or re-creates the daos.cfg and daoscat.nsf files in the Lotus data directory.

The IBM documentation describes how to restore and resynchronize the DAOS repository.

If you want to restore the DAOS directory or specific NLO files only, perform only step 3.

If you know the missing DAOS files or directories beforehand, you can skip step 2 and restore the DAOS files or directories simultaneously with the database files by using one of these methods:

- Use the **nsrnotesrc** command with the NSR_BACKUP_PATHS parameter set to a list of the database and DAOS files or directories.
- Use the NetWorker User for Lotus program and select the required files or directories for restore in the program.

Performing Lotus document-level recovery with nsrdocrc command

You can run the **nsrdocrc** command from the command line to recover deleted Notes documents (not modified Notes documents or design documents) in a local database. This document-level recovery can recover the documents to any point-in-time if the following conditions are true:

- The database is in archived log mode.
- Backups of the database and the transaction logs are available.

Note: The **nsrdocrc** command operates only on a single database file. To restore documents from multiple databases, you must run the **nsrdocrc** command multiple times.

To recover deleted documents from a database file, complete these steps as the Lotus user:

1. Ensure that the NMDA configuration file contains the required parameter settings. For example:

   ```
   NSR_BACKUP_PATHS = database_to_be_restored
   NSR_NOTES_INI_PATH = notes.ini_file_location
   NSR_RECOVER_TIME = point_in_time_for_database_recovery
   NSR_RELOCATION_DEST = directory_path_for_database_restore
   NSR_SERVER = NetWorker_server_name
   ```

   Note: If you have not set NSR_RELOCATION_DEST, the database is restored to /nsr/apps/tmp (UNIX) or *NetWorker_install_path*\apps\tmp (Windows).

   Appendix A, "NMDA Parameters and Configuration File," describes parameters in the NMDA configuration file.

   "Setting the NSR_RECOVER_TIME parameter" on page 379 describes the NSR_RECOVER_TIME parameter.

2. To prevent the Domino server from deleting the directory that contains the temporary database file when it deletes the database, perform one of the following actions:

   - Set the following parameter in the Lotus Domino notes.ini file:

     ```
     DISABLE_DIR_DEL_IF_EMPTY=1
     ```

     The IBM documentation describes this parameter.

   - Create a dummy file in the NSR_RELOCATION_DEST directory.

3. On an AIX (64-bit), Linux, or Solaris client, set the required environment variable for the library locations:

   - On a 64-bit AIX client, set the environment variable LIBPATH to the complete pathname of the Lotus directory that contains the library files libxmlproc_r.a, libndgts_r.a, and libnotes_r.a.

     For example, set LIBPATH to the pathname of the Lotus directory:

     ```
     export LIBPATH=/opt/ibm/lotus/notes/latest/ibmpow
     ```

   - On a Linux or Solaris client, set the environment variable LD_LIBRARY_PATH to the complete pathname of the Lotus directory that contains the library files libxmlproc.so, libndgts.so, and libnotes.so.

     For example, set LD_LIBRARY_PATH to the pathname of the Lotus directory:

     ```
     export LD_LIBRARY_PATH=/opt/lotus/notes/latest/sunspa
     ```

4. Type the **nsrdocrc** command at the command line:

   nsrdocrc(**.exe**) **-z** *configuration_file_path*

   where *configuration_file_path* is the complete pathname of the NMDA configuration file.

   The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrdocrc** command.

Example 20 on page 193 provides an example of document-level recovery with the **nsrdocrc** command.

"Performing document-level recovery of database links" on page 196 describes how to recover documents from a linked database.

Example 20 **Recovering deleted documents for a logged database**

To recover deleted documents in the database F:\Lotus\Domino\data\account.nsf residing on client host saturn, as backed up on June 15 to the NetWorker server mars:

1. Enter the **nsrinfo** command to obtain the backup time of the directory that contains account.nsf:

   ```
   nsrinfo.exe -s mars -n notes saturn | grep "Jun 15"

   NOTES:/F:/Lotus/Domino/data/, date=984492446 Tue Jun 15 08:07:26
   2010 ...
   NOTES:/F:/Lotus/Domino/data/account.nsf, date=984492440 Tue Jun 15
   08:07:20 2010
   ```

2. Ensure that the NMDA configuration file contains the following parameter settings, including the backup time of the F:\Lotus\Domino\data directory:

   ```
   LOTUS {
   Notes_ExecDirectory = F:\Lotus\Domino
   NSR_BACKUP_PATHS = F:\Lotus\Domino\data\account.nsf
   NSR_NOTES_INI_PATH = F:\Lotus\Domino\notes.ini
   NSR_RECOVER_TIME = 984492446
   NSR_RELOCATION_DEST = D:\tempdir
   NSR_SERVER = mars
   }
   ```

3. To perform the document recovery, type the **nsrdocrc** command:

   ```
   nsrdocrc.exe -z configuration_file_path
   ```

Note: Recovering deleted documents from one large database to another large database can take considerable time.

Performing Lotus document-level recovery with the Notes client GUI

On Windows systems only, you can use the Lotus Notes client GUI to recover deleted or modified Notes documents (not design documents) in either of the following databases:

- A local Notes or Domino database
- A remote Domino database

To enable the use of the Lotus Notes client GUI, you must add the document-level recovery feature to the Lotus Notes client as described in the *EMC NetWorker Module for Databases and Applications Installation Guide*.

Document-level recovery through the Notes client GUI can recover the documents only to the time of the selected backup.

The database with the documents for recovery can be physically located on either the local computer or a remote Domino server, but the database must be accessible through the Notes client.

Performing document-level recovery of deleted or modified documents

NOTICE

Before you restore a modified document, back up the current database to save the latest changes in the document. Otherwise, all unsaved changes in the document will be lost during the restore.

To recover deleted or modified documents from a database file:

1. If recovering documents in a remote Domino database, ensure that you meet the following requirements:
   - The user who runs the Notes client program on the local host is:
     - Listed in the Remote Access attribute in the NetWorker Client resource of the remote Domino server.
     - Granted administrative privileges on the remote Domino server.
   - You have configured a NetWorker Client resource on the same NetWorker server for the host where Notes client program runs.

2. To prevent the Domino server from deleting the directory that contains the temporary database file when it deletes the database, perform one of the following actions:
   - Set the following parameter in the Lotus Domino notes.ini file:

     ```
     DISABLE_DIR_DEL_IF_EMPTY=1
     ```

     The IBM documentation describes this parameter.
   - Create a dummy file in the directory used for restore.
3. Start the Lotus Notes client GUI.
4. Open the database that contains the documents to recover.
5. Select the documents to recover.

   Note: Skip this step if you are recovering deleted documents.

6. From the **Actions** menu, select either **NMDA Lotus - Restore Selected Documents** or **NMDA Lotus - Restore Deleted Documents**.

   The NetWorker Module dialog box appears. For a successful recovery, each field in the dialog box except Encryption Phrase must contain a valid value, as required.
7. In the NetWorker Module dialog box, specify the required values:
   a. In the **Database** text box, type the complete pathname of the database that contains the documents to recover.

      "Performing document-level recovery of database links" on page 196 describes how to recover documents from a linked database.

   b. In the **Temporary Directory for Restore** text box, type a temporary directory on the local host, where the database backup is to be restored.

   c. In the **NetWorker Server** text box, type the hostname of NetWorker server that contains the backup to restore.

   d. In the **NetWorker Client** text box, type the name of the NetWorker client file index that contains information about the backup to be restored.

      Note: If the database is on a remote host, the default name listed in the **NetWorker Client** text box might not be correct. This can result in a restore failure.

   e. In the **Show Backups in Last** text box, type the number of past days (since the current date) for which to display backup information about the database. For example, the value 5 causes the display of backups from the past five days.

   f. Click **Refresh** to display a list of the database backups under **List of Backups**.

   g. In the **List of Backups** display, select a database backup to be restored.

   h. In the **Encryption Phrase** text box, type the encryption phrase used to back up the database if the datazone pass phrase on the NetWorker server changed since the backup.

Note: Encryption keys typed in the text box are cached in memory and continue to be used for recoveries until the Notes client program exits and restarts.

8. Click **Restore**.
9. When the recovery is complete, press **F9** to refresh the Lotus Notes screen and display the recovered files.

Performing document-level recovery of database links

To perform a document-level recovery of a linked database, specify the complete pathname of the link file that points to the database by one of these methods:

- Type the link file path in the **Database** field of the Lotus Notes client as described in "Performing document-level recovery of deleted or modified documents" on page 194.
- Specify the link file path with the NSR_BACKUP_PATHS parameter in the NMDA configuration file as described in "NSR_BACKUP_PATHS" on page 373 for a **nsrdocrc** restore.

Note: NMDA does not support document-level recovery of directory links.

If the directory link file C:\Domino\data\link.dir points to the directory D:\Lotus\dir that contains a database db22.nsf, perform a document-level recovery of the database. Specify the complete pathname of the database D:\Lotus\dir\db22.nsf, not the directory link.

Performing MySQL data restore and recovery

Before you perform a MySQL data restore or recovery, you must meet the following requirements:

- You have set the required parameters in the MySQL configuration file. The MEB and MySQL documentation provides details about the MySQL configuration file.
- You have set the required parameters in the NMDA configuration file. Use a different NMDA configuration file for a MySQL restore or recovery than the NMDA configuration file used for a MySQL backup. Appendix A, "NMDA Parameters and Configuration File," describes all the supported parameters.
- The local directory specified by MYSQL_BACKUP_DIR has enough space to contain the temporary backup files extracted from the backup image. You can use the NetWorker **nsrinfo** and **mminfo** commands and the output from a list image operation to determine the size of the backup files to be extracted.
- The directory specified by MYSQL_DATADIR has enough space to contain the restored backup data. You can use the **nsrinfo** and **mminfo** commands and the output from a list image operation to determine the space required for the data.
- You have determined the binary log range required for the recovery as described in "Determining the binary logs for MySQL recovery" on page 197. The binary logs required for the recovery of data must be on the system. If the logs are not on the

system, restore the logs from the binary log backups as described in "Performing MySQL restores of binary log backups" on page 201.

You can run the **nsrmysqlrc** command from the command line to perform a MySQL data restore or recovery. You do not need to specify a MySQL database username or password for the operation.

Note: The operating system user that runs the **nsrmysqlrc** command must be the owner of the directory specified by MYSQL_DATADIR.

The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrmysqlrc** command.

Perform the MySQL restore and recovery operations by using the appropriate procedures:

- "Determining the binary logs for MySQL recovery" on page 197
- "Performing MySQL recovery of whole instance backups" on page 198
- "Performing MySQL restores of partial backups" on page 199
- "Performing MySQL restores of binary log backups" on page 201
- "Performing MySQL redirected restores" on page 201
- "Performing MySQL list image operations" on page 202
- "Performing MySQL extract operations" on page 203
- "Performing MySQL extract and prepare operations" on page 204
- "Performing MySQL copy back operations" on page 205
- "Performing MySQL validate operations" on page 206

Determining the binary logs for MySQL recovery

If you do not specify binary logs by using MYSQL_BINLOG during the recovery, NMDA restores only the appropriate full backup and associated incremental backups to bring the instance to the time of the last incremental backups.

For a recovery to a point-in-time or the current time, you must specify the binary logs by using the MYSQL_BINLOG parameter setting in the NMDA configuration file. Use the following steps to determine the binary log information.

1. Run the **nsrinfo** command to identify the binary logs required for the last backup. Obtain the binary log start number as shown in bold text in the following example:

   nsrinfo -v -n mysql -s *server_name client_name*

   ```
   UNIX file `full_whole_1358798275', NSR size=180356628,
   date=1358798275 Mon 21 Jan 2013 02:57:55 PM EST, (unknown fid), file
   size=0
   MYSQL file `full_whole_1358798275_meta', NSR size=320,
   date=1358798275 Mon 21 Jan 2013 02:57:55 PM EST
   start LSN=150623744, end LSN=150626108, backup level=full, port
   number=3306, bin log=mysql-bin.000036, binlog position=2113, db
   size=0 bytes, whole/partial=whole, redo log only incr=no,
   compressed=no, compresse level=0, InnoDB only=no, backup
   dir=/nsr/apps/tmp/MYSQL_BACKUP_DIR_1358798274, data
   dir=/var/lib/mysql, InnoDB log file in group=2, InnoDB log file
   size=5M
   ```

2. Obtain the latest binary log number on the system as shown in bold text in the following example:

```
mysql@bu-selma:~/cfg> ls -l /var/lib/mysql

total 315408
-rw-r----- 1 mysql mysql 134217728 Feb 12 14:48 ib_logfile0
-rw-r----- 1 mysql mysql 134217728 Feb 12 11:29 ib_logfile1
-rw-r--r-- 1 mysql mysql      2397 Feb  8 16:53 my.cnf
drwx--x--x 2 mysql mysql      4096 Jul  4  2012 mysql
-rw-rw---- 1 mysql mysql       150 Feb  6 17:48 mysql-bin.000042
-rw-rw---- 1 mysql mysql       150 Feb 11 15:46 mysql-bin.000043
-rw-rw---- 1 mysql mysql       150 Feb 11 15:47 mysql-bin.000044
```

Record the log information from the preceding steps. You will need to set MYSQL_BINLOG accordingly during the recovery of MySQL databases.

Performing MySQL recovery of whole instance backups

By default, the **nsrmysqlrc** command restores and recovers a whole backup of a MySQL instance to the current (latest) time. You can optionally specify a point-in-time for the recovery.

For a current-time recovery of a whole instance backup, you can specify all the binary log backups to apply to bring the restored instance to the current time. The software restores the latest full backup of the instance, restores any associated incremental backups, and applies all the available binary log backups to recover the instance to the current time.

Note: For an InnoDB instance, the software uses an apply-log operation to prepare the data restored from the full backup before restoring the incremental backups.

For a point-time-time recovery of a whole instance backup, you can specify all the binary log backups to apply to bring the restored instance to the point-in-time. The software restores the full backup of the instance that is closest to (and no later than) the point-in-time, and restores any associated incremental backups performed before the point-in-time. The software then applies the appropriate binary log backups to recover the instance to the point-in-time.

1. Set the required parameters in the NMDA configuration file for the recovery. For example:

```
MYSQL_BACKUP_DIR = /nsr/apps/tmp/backup
MYSQL_BACKUP_NAME = MYSQL:/myinstance_whole
MYSQL_BINLOG = [/var/lib/mysql/mysql-bin.000036;
/var/lib/mysql/mysql-bin.000044]
MYSQL_CFG_FILE = /etc/my.cnf
MYSQL_DATADIR = /var/lib/mysql
MYSQL_MEB_OPTIONS = uncompress
NSR_RECOVER_TIME = "Wed Feb 6 2013 14:23"
NSR_SERVER = NetWorker_server_hostname
```

   In this example, the MYSQL_BACKUP_NAME and MYSQL_BACKUP_DIR parameters specify to extract the backup files from the backup image named MYSQL:/myinstance_whole into the /nsr/apps/tmp/backup directory, respectively. NMDA first extracts a full backup into a “full” subdirectory and then each incremental backup into a separate “incr#” subdirectory under /nsr/apps/tmp/backup/, where #

is the sequence number starting at 1. MYSQL_MEB_OPTIONS specifies to uncompress the compressed backup. A copy back operation restores the prepared backup to the /var/lib/mysql directory specified by MYSQL_DATADIR.

You can set MYSQL_BINLOG to specify binary logs to apply to the restored instance to bring the instance to the current time or a point-in-time. If you set MYSQL_BINLOG, the recovery creates a binlog_trx file, under a directory specified by MYSQL_BACKUP_DIR, which contains all the SQL transactions from the binary logs. These transactions occurred between the last available backup and the last transaction in the binary logs. In step 4, you will play back these transactions to bring the restored instance closer to the required time.

Set NSR_RECOVER_TIME for a point-in-time recovery only. Set NSR_SERVER to the hostname of the NetWorker server host.

Appendix A, "NMDA Parameters and Configuration File," provides details on each parameter in the NMDA configuration file.

2. Run the **nsrmysqlrc** command at the command line to restore the whole instance backup:

   ```
   nsrmysqlrc -z configuration_file_path
   ```

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the required parameter settings.

3. If MYSQL_DATADIR specifies the MySQL data directory, shut down the database server to enable the copy back operation to restore the prepared backup to the data directory. The restore prompts you to shut down the database server if you do not disable the prompt by setting NSR_RECOV_INTERACT.

4. If MYSQL_BINLOG specifies one or more binary logs, restart the database server as required and run the following command to play back the SQL transactions from the binary logs, as the **nsrmysqlrc** program instructs. This action brings the restored instance to the required time and completes the recovery:

   ```
   mysql < <MYSQL_BACKUP_DIR>/binlog_trx
   ```

Performing MySQL restores of partial backups

You can run the **nsrmysqlrc** command to restore a partial backup to the time of the last backup (full or incremental). You cannot use **nsrmysqlrc** to apply binary logs to bring the restored partial backup to a more recent time.

The **nsrmysqlrc** program restores only the databases and tables included in the backup image specified by MYSQL_BACKUP_NAME. Do not restore a partial backup to the original MySQL instance, to prevent the overwriting of other databases and tables that were not included in the partial backup. This is a MYSQL limitation. Instead, restore the partial backup to a new MySQL instance.

For a current-time restore of a partial backup, the software restores the database objects from the latest full backup and restores any associated incremental backups. However, the software does not apply any binary log backups to the restored partial backup.

Note: For InnoDB data, the software uses an apply-log operation to prepare the data restored from the full backup before restoring the incremental backups.

For a point-in-time restore of a partial backup, the software restores the database objects from the full backup that is closest to (and no later than) the point-in-time, and restores any associated incremental backups performed before the point-in-time. However, the software does not apply any binary log backups to the restored partial backup.

1. Set the required parameters in the NMDA configuration file for the restore. For example:

   ```
   MYSQL_BACKUP_DIR = /nsr/apps/tmp/backup
   MYSQL_BACKUP_NAME = MYSQL:/mydb1
   MYSQL_CFG_FILE = /etc/my.cnf
   MYSQL_DATADIR = /var/lib/mysql
   MYSQL_MEB_OPTIONS = uncompress
   NSR_RECOVER_TIME = "Wed Feb 6 2013 14:23"
   NSR_SERVER = NetWorker_server_hostname
   ```

   These parameters specify to extract the backup files from the backup image named MYSQL:/mydb1, specified by MYSQL_BACKUP_NAME, into the directory /nsr/apps/tmp/backup specified by MYSQL_BACKUP_DIR. NMDA first extracts a full backup into a "full" subdirectory and then each incremental backup into a separate "incr#" subdirectory under /nsr/apps/tmp/backup/, where #is the sequence number starting at 1. MYSQL_MEB_OPTIONS specifies to uncompress the compressed backup. A copy back operation restores the prepared backup to the /var/lib/mysql directory specified by MYSQL_DATADIR.

   Set NSR_RECOVER_TIME for a point-in-time recovery only. Set NSR_SERVER to the hostname of the NetWorker server host.

   Note: If you set MYSQL_BINLOG to specify one or more binary log backups, the restore operation ignores the setting. You cannot apply binary logs to a restored partial backup.

   Appendix A, "NMDA Parameters and Configuration File," describes parameters in the NMDA configuration file.

2. Run the **nsrmysqlrc** command at the command line to restore the partial backup:

   ```
   nsrmysqlrc -z configuration_file_path
   ```

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the required parameter settings.

3. If MYSQL_DATADIR specifies the MySQL data directory, shut down the database server to enable the copy back operation to restore the prepared backup to the data directory. The restore prompts you to shut down the database server if you do not set NSR_RECOV_INTERACT to disable the prompt.

Note: Due to an MEB limitation, a restore of a partial InnoDB backup (a backup performed with MYSQL_INCLUDE or the **include** parameter in the MySQL configuration file) does not perform the copy back operation to complete the restore of the prepared backup. The restore operation displays a message about how to complete the restore.

Performing MySQL restores of binary log backups

You can run the **nsrmysqlrc** command to restore one or more binary logs from binary log backups.

1. Set the required parameters in the NMDA configuration file for the binary log restore. For example:

   ```
   MYSQL_RESTORE_OPERATION = binlog_restore
   MYSQL_BACKUP_DIR = /nsr/apps/tmp/backup
   MYSQL_BINLOG = [/var/lib/mysql/bin.001; /var/lib/mysql/bin.005]
   MYSQL_CFG_FILE = /etc/my.cnf
   NSR_CLIENT = NetWorker_client_hostname
   NSR_SERVER = NetWorker_server_hostname
   ```

   These parameters specify to restore the latest backup of the range of binary logs from /var/lib/mysql/bin.001 to /var/lib/mysql/bin.005. The operation restores the binary logs into the /nsr/apps/tmp/backup directory.

   Set MYSQL_RESTORE_OPERATION=binlog_restore to specify the binary log restore. Set MYSQL_BINLOG to specify a single binary log or a range of binary logs. If you do not set MYSQL_BACKUP_DIR, the binary logs are restored to the original location where the logs were backed up.

   Set NSR_CLIENT if the backup client is different from the client that runs **nsrmysqlrc**. Set NSR_SERVER if the NetWorker server host is different from the client host.

   Appendix A, "NMDA Parameters and Configuration File," describes parameters in the NMDA configuration file.

2. Run the **nsrmysqlrc** command at the command line to restore the binary log backup:

   nsrmysqlrc -z *configuration_file_path*

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the required parameter settings.

Performing MySQL redirected restores

You can perform two types of redirected restores:

- You can restore MySQL data backed up from one data directory to a different data directory using the same or different instance on the same host.
- You can restore a MySQL backup created on one server to a different server on a different host, for example, in a disaster recovery scenario. "Preparing for MySQL disaster recovery" on page 235 provides details on disaster recovery.

Performing a MySQL redirected restore to a different data directory

You can use the **nsrmysqlrc** command to restore backup data to a different data directory on the same host.

1. Set MYSQL_DATADIR to the pathname of the new data directory.
2. In the MySQL configuration file, my.cnf, locate all the parameter settings that include the original data directory and change those settings to point to the new data directory.

3. Use one of the following procedures to complete the restore:
   - "Performing MySQL recovery of whole instance backups" on page 198
   - "Performing MySQL restores of partial backups" on page 199

Performing a MySQL redirected restore to a different host

You can use the **nsrmysqlrc** command to restore and recover a backup of MySQL server A to MySQL server B on a different host.

Note: You must run the **nsrmysqlrc** command on the server B host.

1. Ensure that you specify the following settings for the restore:
   - Specify the username and fully qualified hostname for the remote server B host in the Remote Access attribute in the Client resource of the server A host. Use the following format for the Remote Access attribute value:

     user=*remote_username*,**host=***remote_hostname*

   - Set MYSQL_DATADIR to the data directory pathname on the server B host.
   - Set NSR_CLIENT to the hostname of the server A host.
   - Set NSR_SERVER to the hostname of the NetWorker server used for the backup.
2. Use one of the following procedures to complete the restore:
   - "Performing MySQL recovery of whole instance backups" on page 198
   - "Performing MySQL restores of partial backups" on page 199

Performing MySQL list image operations

You can run the **nsrmysqlrc** command to perform a "list image" operation that lists all the backup files or a specific file or directory from a backup image created by an NMDA MySQL backup. The list image operation does not alter the backup image.

You do not need to extract files from the backup image prior to the list image operation. You must ensure that the device containing the backup image is mounted before you perform a list image operation.

1. Set the required parameters in the NMDA configuration file for the list image operation. For example:

   ```
   MYSQL_RESTORE_OPERATION = list_image
   MYSQL_BACKUP_NAME = MYSQL:/myinstance_whole
   MYSQL_SRC_ENTRY = meta/backup_var.txt
   NSR_SERVER = NetWorker_server_hostname
   ```

   These parameters specify to list the information about the single file meta/backup_var.txt from the backup image named MYSQL:/myinstance_whole.

   Set MYSQL_BACKUP_NAME to specify a backup image name. As an alternative, you can specify a backup piece name, obtained by querying the client file index with the **nsrinfo** command.

   Set MYSQL_SRC_ENTRY only to list the information about a specific file or directory. Set NSR_SERVER to the hostname of the NetWorker server host.

Appendix A, "NMDA Parameters and Configuration File," describes parameters in the NMDA configuration file.

2. Run the **nsrmysqlrc** command at the command line to complete the list image operation:

   nsrmysqlrc -z *configuration_file_path*

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the required parameter settings.

A list image operation that lists all the backup files from a backup image produces the following type of output, which shows the size of each block in bytes and the total backup size in bytes:

```
mysqlbackup: INFO: Backup Image MEB version string: 3.6.1 [2011/09/28]
[File]: [Size:         188]: backup-my.cnf
[File]: [Size:        5678]: meta/backup_create.xml
[File]: [Size:    16777216]: datadir/ibdata1
[File]: [Size:     2097152]: datadir/ibdata1.$_append_$.1
[Dir]: datadir/mysql
[File]: [Size:          35]: datadir/mysql/backup_history.CSM
[File]: [Size:        2358]: datadir/mysql/backup_history.CSV
[File]: [Size:       71260]: datadir/mysql/backup_history.frm
[File]: [Size:          35]: datadir/mysql/backup_progress.CSM
[File]: [Size:        5474]: datadir/mysql/backup_progress.CSV
[File]: [Size:       33370]: datadir/mysql/backup_progress.frm
[File]: [Size:           0]: datadir/mysql/columns_priv.MYD
[File]: [Size:        4096]: datadir/mysql/columns_priv.MYI
[File]: [Size:        8820]: datadir/mysql/columns_priv.frm
[File]: [Size:         880]: datadir/mysql/db.MYD
[File]: [Size:        5120]: datadir/mysql/db.MYI
[File]: [Size:        9582]: datadir/mysql/db.frm
...
...
[Dir]: datadir/test
[File]: [Size:        2560]: datadir/ibbackup_logfile
[File]: [Size:         182]: meta/backup_variables.txt
[File]: [Size:       31727]: meta/backup_content.xml
[File]: [Size:       12881]: meta/image_files.xml
 mysqlbackup: INFO:  Backup image contents listed successfully.
    Source Image Path= sbt:fiftytry
111116 11:36:01 mysqlbackup: INFO: meb_sbt_restore_close: blocks: 20
   size: 1048576  bytes: 20067081
mysqlbackup completed OK!
```

Performing MySQL extract operations

You can run the **nsrmysqlrc** command to perform an "extract" operation that extracts either all the backup files or only a single file or directory from a backup image. For example, you can extract just the meta backup piece file from a backup image to obtain specific information about the backup from the file.

Note: Before you perform an extract operation, you must ensure the local device contains enough space for the extracted files. To determine the size of the backup files to be extracted, you can use the NetWorker **nsrinfo** and **mminfo** commands and the output from a list image operation.

You can extract part of the data from a backup image, but you cannot then restore the data to a database.

1. Set the required parameters in the NMDA configuration file for the extract operation. For example:

   ```
   MYSQL_RESTORE_OPERATION = extract
   MYSQL_BACKUP_DIR = /nsr/apps/tmp/backup
   MYSQL_BACKUP_NAME = MYSQL:/myinstance_whole
   NSR_SERVER = NetWorker_server_hostname
   ```

   These parameters specify to extract all the backup files from the backup image named MYSQL:/myinstance_whole to the local directory /nsr/apps/tmp/backup.

   MYSQL_BACKUP_NAME specifies a backup image name. As an alternative, you can specify a backup piece name, obtained by querying the client file index with the **nsrinfo** command.

   To extract only a single file or directory from a backup image, set MYSQL_EXTRACT_PATHS and do not set MYSQL_BACKUP_DIR.

   Set NSR_SERVER if the NetWorker server host is different from the client host.

   Appendix A, "NMDA Parameters and Configuration File," describes parameters in the NMDA configuration file.

2. Run the **nsrmysqlrc** command at the command line to complete the extract operation:

   nsrmysqlrc -z *configuration_file_path*

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the required parameter settings.

Performing MySQL extract and prepare operations

An NMDA MySQL backup is a raw backup that must be processed to prepare it for restore. You can run the **nsrmysqlrc** command to perform an "extract and prepare" operation that extracts all the backup files from a backup image and processes the files to produce a prepared backup in the MYSQL_BACKUP_DIR/full directory.

The **nsrmysqlrc** program performs the following steps to complete an extract and prepare operation:

1. Locates the latest full backup and all the associated incremental backups (if any) for the backup image specified by the MYSQL_BACKUP_NAME parameter.
2. Extracts the full backup and incremental backups to subdirectories under the directory specified by the MYSQL_BACKUP_DIR parameter.
3. Uses an apply-log operation to prepare the full backup as required.
4. Applies the incremental backups to the prepared full backup, which resides in the MYSQL_BACKUP_DIR/full directory.

By default, the extract and prepare operation produces a prepared backup that corresponds to the current (latest) time.

Alternatively, you can use the NSR_RECOVER_TIME parameter to specify a point-in-time, so that the operation produces a prepared backup closest to (and no later than) the specified time.

Note: The extract and prepare operation does not apply any binary log backups in preparing the backup.

1. Set the required parameters in the NMDA configuration file for the extract and prepare operation. For example:

   ```
   MYSQL_RESTORE_OPERATION = extract_and_prepare
   MYSQL_BACKUP_DIR = /nsr/apps/tmp/backup
   MYSQL_BACKUP_NAME = MYSQL:/mydb1
   MYSQL_CFG_FILE = /etc/my.cnf
   MYSQL_MEB_OPTIONS = uncompress
   NSR_RECOVER_TIME = "Wed Feb 6 2013 14:23"
   NSR_SERVER = NetWorker_server_hostname
   ```

   MYSQL_RESTORE_OPERATION=extract_and_prepare and MYSQL_BACKUP_DIR specify to extract the backup files from the backup image named MYSQL:/mydb1 and produce a prepared backup in the /nsr/apps/tmp/backup/full directory. The setting MYSQL_MEB_OPTIONS=uncompress specifies to uncompress the compressed backup.

   Set NSR_RECOVER_TIME to produce a prepared backup closest to (and no later than) a point-in-time. Set NSR_SERVER if the NetWorker server host is different from the client host.

   Appendix A, "NMDA Parameters and Configuration File," describes parameters in the NMDA configuration file.

2. Run the **nsrmysqlrc** command at the command line to complete the extract and prepare operation:

   ```
   nsrmysqlrc -z configuration_file_path
   ```

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the required parameter settings.

You can perform a copy back operation on the prepared backup to complete the restore. For example, you can perform the extract and prepare operation on a different machine than the production server and then copy back the prepared backup to the production server to complete the restore.

Performing MySQL copy back operations

You can run the **nsrmysqlrc** command to perform a "copy back" operation that copies a prepared backup to a specified directory. The prepared backup can be produced either by the **nsrmysqlrc** program through an "extract and prepare" operation or by the MEB utilities.

For example, you can perform a copy back operation to the MySQL data directory to restore a prepared backup to the production database server. A copy back operation to the data directory requires a shutdown of the database server.

You can also use a copy back operation to copy a prepared backup to an alternative directory, for example, to use for testing, reporting, or deployment in a replication environment.

Ensure that the database server is shut down before you copy back a prepared backup to a MySQL data directory.

1. Set the required parameters in the NMDA configuration file for the copy back operation. For example:

   ```
   MYSQL_RESTORE_OPERATION = copy_back
   MYSQL_BACKUP_DIR = /nsr/apps/tmp/backup
   MYSQL_CFG_FILE = /etc/my.cnf
   MYSQL_DATADIR = /var/lib/mysql
   MYSQL_INNODB_LOG_FILE_SIZE = 108576
   MYSQL_INNODB_LOG_FILES_IN_GROUP = 2
   NSR_SERVER = NetWorker_server_hostname
   ```

   These parameters specify to perform a copy back operation that restores the prepared backup from the /nsr/apps/tmp/backup/full directory to the /var/lib/mysql directory.

   Set NSR_SERVER if the NetWorker server host is different from the client host.

   Appendix A, "NMDA Parameters and Configuration File," describes parameters in the NMDA configuration file.

2. Run the **nsrmysqlrc** command at the command line to complete the copy back operation:

   nsrmysqlrc -z *configuration_file_path*

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the required parameter settings.

Performing MySQL restores of InnoDB tables outside data directory

With MySQL 5.6, you can optionally store InnoDB tables in a specified directory outside the MySQL data directory.

You must complete the required steps to restore an NMDA backup of InnoDB tables stored outside the data directory if you want to restore the tables to a directory structure that is different than the directory structure at the backup time.

1. Extract and prepare the backup to a staging area. "Performing MySQL extract and prepare operations" on page 204 provides details.
2. Edit all the .bl files in the backup directory so that each file includes the correct pathnames for the new directory structure.
3. Copy back the data to the MySQL data directory. "Performing MySQL copy back operations" on page 205 provides details.

Performing MySQL validate operations

With MEB 3.7 or later, you can run the **nsrmysqlrc** command to perform a validate operation that validates the integrity of the backup image. The validate operation does not alter the backup image.

Ensure that the device containing the backup image is mounted before you perform a validate operation.

1. Set the required parameters in the NMDA configuration file for the validate operation. For example:

   ```
   MYSQL_RESTORE_OPERATION = validate
   MYSQL_BACKUP_NAME = myinstance_full_whole_1349901207
   NSR_SERVER = NetWorker_server_hostname
   ```

   These parameters specify to validate the integrity of the backup image with the backup piece name myinstance_full_whole_1349901207. You can obtain the backup piece name by querying the client file index with the **nsrinfo** command.

   Set NSR_SERVER to the hostname of the NetWorker server host.

   Appendix A, "NMDA Parameters and Configuration File," describes parameters in the NMDA configuration file.

2. Run the **nsrmysqlrc** command at the command line to complete the validate operation:

   nsrmysqlrc -z *configuration_file_path*

   where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the required parameter settings.

Performing Oracle data restore and recovery

Prepare for an Oracle data restore and recovery by using the following information:

- "Determining the volumes for Oracle restores" on page 207
- "Preventing restore performance degradation with Oracle 10.2 or later" on page 211
- "RMAN scripts for restore and recovery" on page 211
- "Using the Oracle recovery configuration wizard" on page 212

Perform the Oracle data restore by using the appropriate procedures:

- "Performing Oracle restores with the rman command" on page 215
- "Performing Oracle restores with the Oracle Enterprise Manager" on page 215

Complete the Oracle data recovery, if required, according to "Performing Oracle data recovery" on page 215.

Determining the volumes for Oracle restores

With Oracle 10gR2 and later, you can use the **restore...preview** command to identify backups that RMAN needs for a restore operation. The output of **restore...preview** also indicates which media (NetWorker volumes) you need for the restore and if any of the NetWorker volumes are remote. A volume is remote if the volume requires operator intervention to become accessible to NetWorker.

Note: Ensure that you have set NSR_VOLUMES_INFO=TRUE before you run the **restore...preview** command.

If a required backup is on a remote volume, perform one of the following operations:

- Run the **change...unavailable** command to prevent RMAN from selecting the remote backups. Then retry the **restore...preview** operation to see if RMAN selects another remote backup. When RMAN does not select any remote backups, you can perform the **restore** operation.
- Run the **restore...preview** command with the **recall** option. The **restore...preview recall** operation automatically requests the retrieval of the remote volumes by generating a media notification on the NetWorker server. If a backup piece required for the restore spans multiple volumes, the operation generates a media notification for each volume sequentially.

Note: Due to an Oracle limitation, the Oracle software does not display the volume information if a volume is remote on Windows.

The volumes listed in the output of **restore...preview** are the volumes used at the backup time. Due to an Oracle limitation, if the volume that NetWorker intended to use for the restore at the preview time is different from the original volume used at the backup time (for example, due to NetWorker staging), the **restore...preview** operation still lists the volume name used at the backup time.

Example 21 Using restore...preview to determine volumes for restore

The following RMAN script identifies the backups that RMAN needs for restoring datafiles 1 and 2. The NetWorker server name is server1:

```
run {
   allocate channel t1 type SBT send 'ENV=(NSR_SERVER=server1,
   NSR_VOLUMES_INFO=TRUE)';
   restore datafile 1, 2 preview;
}
```

The following output of the **restore...preview** command shows that the backup of datafile 1 is on the volume DBMIData.808 and the backup of datafile 2 is on the volume DBMIData.802. This output also shows that the volume DBMIData.808 is remote:

```
List of Backup Sets
===================
BS Key  Type LV Size       Device Type Elapsed Time Completion Time
------- ---- -- ---------- ----------- ------------ ---------------
96      Full    127.00M    SBT_TAPE    00:00:05     17-MAY-10
        BP Key: 99   Status: AVAILABLE  Compressed: NO  Tag:
   TAG20100517T144317
        Handle: 06ldtl86_1_1   Media: DBMIData.808
  List of Datafiles in backup set 96
  File LV Type Ckp SCN    Ckp Time  Name
  ---- -- ---- ---------- --------- ----
  1       Full 225701     17-MAY-10 /space/oradata/tartst/Sys1.ora

BS Key  Type LV Size       Device Type Elapsed Time Completion Time
------- ---- -- ---------- ----------- ------------ ---------------
109     Full    46.25M     SBT_TAPE    00:00:03     17-MAY-10
        BP Key: 111   Status: AVAILABLE  Compressed: NO  Tag:
   TAG20100517T144415
        Handle: 08ldtla0_1_1   Media: DBMIData.802
  List of Datafiles in backup set 109
```

```
  File LV Type Ckp SCN    Ckp Time  Name
  ---- -- ---- ---------- --------- ----
  2       Full 225811     17-MAY-10 /space/oradata/tartst/utbs1.ora

List of remote backup files
===========================
        Handle: 06ldtl86_1_1   Media: DBMIData.808
```

If you have a list of Oracle backup piece names for restore, you can use the **nsrorainfo** command instead to identify the NetWorker volumes required for restoring the backup pieces.

The **nsrorainfo** command provides the following volume information for the restore of each backup piece:

- The name and location of the volume.
- The save time of the backup piece.

The **nsrorainfo** command does not provide the status of a volume for the restore. The command cannot recall the remote volumes.

You can use the following **nsrorainfo** command syntax to identify the volumes for restore of specified backup pieces:

```
nsrorainfo[.exe] [-c NetWorker_client_name] [-s NetWorker_server_name]
   [-f file_name] [backup_piece_name1 [backup_piece_name2 ...]]
```

where:

- *NetWorker_client_name* is the hostname of the NetWorker client whose index contains information about the Oracle backup pieces. By default, the client is on the local host.
- *NetWorker_server_name* is the hostname of the NetWorker server to query for the volumes. By default, the server is on the local host.
- *file_name* is the name of a text file that contains a list of one or more backup piece names for restore:
  - The file must contain each backup piece name on a separate line.
  - The file cannot contain spaces or comments, for example, comment lines preceded with the # symbol.
- *backup_piece_name1* and *backup_piece_name2* are backup piece names for restore.

Command options in brackets ([ ]) are optional. Do not include the brackets when you type the command.

With the **nsrorainfo** command, specify backup piece names by either or both of the following methods:

- List the backup piece names as options of the command.
- List the backup piece names in a text file, and specify the name of the file with the **-f** option of the command.

The listed volumes are the most accessible volumes, which the NetWorker server intends to use for the restore at the time that you type the command. For example, the command lists the clones of volumes if the original volumes are not accessible.

If you remove any listed volumes from the NetWorker devices or you delete any volumes after you type the **nsrorainfo** command, the server can perform the restore by using different volumes that are accessible.

Example 22 on page 210 and Example 23 on page 210 describe how to use the **nsrorainfo** command and the command output.

The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrorainfo** command.

Example 22 Sample nsrorainfo commands for Oracle restores

Each of the following **nsrorainfo** commands provides a list of the volumes required to restore the specified backup pieces:

- The following command searches in the NetWorker index of the client mars on the server server1 for information about the volumes that contain the backup pieces backupc_1 and backupc_2:

  ```
  nsrorainfo -c mars -s server1 backupc_1 backupc_2
  ```

- The following command searches in the NetWorker index of the local host for information about the volumes that contain the backup pieces listed in the file backup2.txt. The command assumes that the NetWorker client and server are both on the local host:

  ```
  nsrorainfo -f backup2.txt
  ```

- The following command searches in the NetWorker index of the client mars for information about the volumes that contain both of these backup pieces:

  - The backup piece backupc_3.
  - The backup pieces listed in the file backup3.txt.

  The command assumes that the NetWorker server is on the local host:

  ```
  nsrorainfo -c mars backupc_3 -f backup3.txt
  ```

Example 23 Volume information that the nsrorainfo command provides

The following **nsrorainfo** command searches in the NetWorker index of the local host on the server mars for information about the volumes that contain the backup pieces backup1 and backup2:

```
nsrorainfo -s mars backup1 backup2
```

The **nsrorainfo** command provides the following type of information:

```
backup1:
   mars.003 at /space/nw_volume1 (save time 1098886937)
   mars.004 at /space/nw_volume2 (save time 1098883454)

backup2:
   mars.005 at /dev/rmt/0cbn (save time 1098883452)
```

According to this command display:

- You require volumes mars.003 and mars.004 to restore the backup piece backup1.
- You require volume mars.005 to restore the backup piece backup2.

Preventing restore performance degradation with Oracle 10.2 or later

Due to an Oracle limitation, degradation of NMDA restore performance might occur with Oracle 10.2 or later if you use NetWorker multiplexing with NMDA for Oracle backups.

If you have enabled NetWorker multiplexing for NMDA Oracle backups, you can prevent the restore performance degradation by including the **set parallelmediarestore off** command in the RMAN script used for the Oracle restore.

For example, the following RMAN restore script contains the required Oracle command to disable the multiplexing during the Oracle restore:

```
set parallelmediarestore off;
run {
   allocate channel c1 type 'SBT_TAPE';
   restore database;
   release channel c1;
}
```

RMAN scripts for restore and recovery

You need an appropriate RMAN script to perform the preferred type of Oracle restore and recovery operation on the Oracle Server host. You can create the RMAN script either manually or by using the recovery configuration wizard. "Using the Oracle recovery configuration wizard" on page 212 describes the wizard.

You can store the RMAN restore scripts as text files. Alternatively, if you use a Recovery Catalog, you can store the restore scripts in the Recovery Catalog database. The Oracle backup and recovery documentation provides details.

You must set all the parameters in the RMAN restore script, preferably with the **send** command. "The send command" on page 398 provides details.

"Common NMDA parameters" on page 358 and "NMDA Oracle parameters" on page 387 describe the common parameters and Oracle parameters, respectively.

Example 24 RMAN script for a tablespace restore

The following RMAN script performs a restore of an Oracle tablespace by using the NetWorker server mars.emc.com. The script restores the Oracle data to the NetWorker client server1.emc.com. The script also includes the recovery operation as described in "Performing Oracle data recovery" on page 215:

```
run {
   allocate channel t1 type 'SBT_TAPE';
   allocate channel t2 type 'SBT_TAPE';
   send 'NSR_ENV=(NSR_SERVER=mars.emc.com,
   NSR_CLIENT=server1.emc.com)';
   sql 'alter tablespace users offline immediate';
   restore tablespace users;
   recover tablespace users;
   sql 'alter tablespace users online';
   release channel t1;
   release channel t2;
}
```

"RMAN scripts for manual backups" on page 129 describes the setting of NMDA parameters in an RMAN script.

Example 25 RMAN script for a restore from a specified pool

By default, NMDA and NetWorker software use the configuration settings and the information in the media database to determine the backup volume to use for an NMDA restore.

Optionally, you can use the NSR_RECOVER_POOL parameter in the RMAN restore script to restore data from a specific volume pool if there are multiple copies (clones) of the backup on different volume pools. "NSR_RECOVER_POOL" on page 364 provides details.

The following RMAN script performs a restore of the database from the specified volume pool named OracleClonePool2, where the pool contains a clone of the original backup volume:

```
shutdown immediate;
startup mount;
run {
   allocate channel c1 type 'SBT_TAPE';
   send channel c1 'NSR_ENV=(NSR_SERVER=backup01,
   NSR_RECOVER_POOL=OracleClonePool2)';
   restore database;
   release channel c1;
}
```

Using the Oracle recovery configuration wizard

NMDA supports an Oracle recovery configuration wizard, integrated with NMC. You can use the wizard to create an RMAN script for restore and recovery of Oracle data backed up by NMDA.

You can run the recovery configuration wizard from the NetWorker Console Administration window, which you can start on any supported host by using a web browser session and by specifying the Console server URL.

To configure an Oracle restore and recovery with the wizard:

1. Review "Features of the Oracle recovery configuration wizard" on page 212.
2. Meet the "Requirements for using the Oracle recovery configuration wizard" on page 213.
3. Follow the steps in "Creating an RMAN restore script with the wizard" on page 214.

Features of the Oracle recovery configuration wizard

The Oracle recovery configuration wizard can create an RMAN script for the following types of restore and recovery:

- Current-time restore and recovery of a whole or partial Oracle database, where a partial database includes tablespaces or datafiles.

  The wizard can configure a tablespace restore if the control file contains information about the tablespace.

- Point-in-time restore and recovery of a whole Oracle database.
- Restore of individual archived redo logs.
- Restore and recovery of Oracle data to a different database through the creation of a duplicate database on either the local host or a remote host, by using backups of the original target database.

The database duplication script (created by the wizard) uses the RMAN **duplicate** command to create a duplicate database while the original database remains. The duplicate database can either be an identical copy of the original database or contain only a subset of the original tablespaces. For example, the duplicate database can run independently on a remote host for practicing restore and recovery operations while the production database remains in operation on the local host:

- If the duplicate database is to be created on the same host as the original database, the wizard generates the RMAN script on the local host. The wizard requests names for the duplicate database, datafiles, and redo logs that differ from the names for the original database.
- If the duplicate database is to be created on a remote host, the wizard generates the RMAN script on either the local host or a remote host, as specified in the wizard. The wizard requests a name for the duplicate database that differs from the name for the original database. The datafile and redo log names can be the same as for the original database.

The recovery configuration wizard can only create an RMAN script for restore and recovery. The wizard cannot modify an existing RMAN script. You must use a text editor to modify an RMAN script that the wizard created.

"RMAN scripts for restore and recovery" on page 211 describes RMAN scripts for restore and recovery.

The recovery configuration wizard does not support the following features:

- Cluster or Oracle RAC systems
- Proxy backups
- RMAN automatic channels

The following sources describe the configuration wizard:

- Descriptive inline text in the wizard
- Online help in the wizard
- *EMC NetWorker Module for Databases and Applications Release Notes*

Requirements for using the Oracle recovery configuration wizard

Before you use the Oracle recovery configuration wizard, you must meet the following requirements:

- The NMC user that starts the wizard (the wizard user) has the Remote Access NetWorker privileges on the NetWorker server that contains the NMDA client configuration.
- Communication between the NMC server, NetWorker server, and NMDA client uses NetWorker **nsrauth** authentication. The NetWorker documentation provides requirements for **nsrauth** authentication.

- You have created the NetWorker Client resource for the NMDA client by using one of the following methods:
  - Backup configuration wizard in NMDA
  - Conversion of a legacy NetWorker module configuration with the **nsrdaadmin** command
  - Client-side configuration method (without the wizard), where the value of the Save Set attribute of the Client resource has the RMAN: prefix
- Before the wizard creates a database duplication script, the AUXILIARY instance exists on the local host or a remote host, and is accessible through Oracle Net. The *Oracle Database Backup and Recovery Advanced User's Guide* describes how to create an AUXILIARY instance.

Creating an RMAN restore script with the wizard

1. Start the **NetWorker Management Console** software.
2. Open the **Administration** window:
   a. In the **Console** window, click **Enterprise**.
   b. In the left pane, select a NetWorker server in the **Enterprise** list.
   c. In the right pane, select the application.
   d. From the **Enterprise** menu, click **Launch Application**.

   The **Administration** window appears as a separate application.
3. In the **Administration** window, click **Configuration**.
4. In the **Configuration** window, click **Clients**.
5. To start the wizard, right-click the NMDA client in the right pane, and select **Recover**.
6. On each wizard screen that appears, specify the required values for the RMAN script configuration.

   Each wizard screen includes an online help button that you can click to access descriptions of all the fields and options on the screen:
   - On all screens but the last screen, click **Next** to continue.
   - On the last screen titled **Review and Accept the Script Creation**, click **Create** to create the RMAN restore script.

NOTICE

If you create an RMAN script with the wizard to perform a tablespace restore when the database is not open, do not select these options:
- Options to place the tablespace in offline mode before the restore
- Options to place the tablespace in online mode after the recovery.

Oracle requires the database to be open to alter the availability of a tablespace.

Performing Oracle restores with the rman command

You can run the **rman** command at the command line on the Oracle Server host, which runs the RMAN utility, to perform an Oracle data restore.

If the RMAN restore script on page 211 is in the file /disk1/scripts/restore.txt and you configured the Net service to connect to the databases payroll and rcvcatdb, type this command to perform the Oracle restore:

```
rman target internal/oracle@payroll rcvcat rman/rman@rcvcatdb cmdfile
   \'/disk1/scripts/restore.txt\'
```

On Windows systems, the command to run the RMAN script is **rman.exe**.

Performing Oracle restores with the Oracle Enterprise Manager

The Oracle Enterprise Manager Backup Management Tools provide a graphical user interface to RMAN.

Note: NMDA does not support the use of the Oracle recovery configuration wizard with the Oracle Enterprise Manager Backup Management Tools.

Use the Oracle Enterprise Manager to perform the following operations:

- Generate the required RMAN commands.
- Perform Oracle backup and restore operations.

NOTICE

After the completion of an NMDA backup or restore, the job queue history of the Oracle Enterprise Manager might display the status of the job as failed, even if the backup or restore completed successfully. This incorrect status is due to an existing problem with Oracle Enterprise Manager. View the job output to confirm that the backup or restore completed successfully.

Performing Oracle data recovery

After you use the RMAN utility to restore the NMDA backups of Oracle data, you can complete the Oracle data recovery, if required.

To recover the Oracle data, 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. A sample RMAN script appears on page 211. With this method, RMAN automatically restores archived redo logs that are required for the recovery but do not exist on the Oracle server host.
- After the RMAN restore script completes successfully, use the Oracle command line (for example, SQL* Plus) or graphical interfaces to run the recovery.

Performing Sybase data restore and recovery

Before you perform a Sybase data restore and recovery, you must meet the following requirements:

- The Sybase server and Sybase Backup Server are running.
- The target database exists, to which data will be restored. The database is at least as large as the size of the database backup.

  Note: To create a database for restore, use the **for load** option.

- The target database to which data will be restored is not in use. The database is taken offline during the restore.

Perform the Sybase data restore and recovery by using the appropriate procedures:

- "Specifying verification of Sybase database restores" on page 216
- "Performing Sybase data restores with the nsrsybrc command" on page 216
- "Performing Sybase data restores with NetWorker User for Sybase" on page 220

Note: If NMDA backed up the Sybase data as a multistripe backup, NMDA automatically enables a multistripe restore and uses the same session number as the multistripe backup.

Specifying verification of Sybase database restores

NMDA supports verification of Sybase database restores at the header verification level and the full verification level.

Run the **nsrsybrc** restore command with the **-V** option to specify the restore verification level:

```
nsrsybrc -V [header | full | verifyonly]
```

where:

- **header**—Specifies to verify the page header information only.
- **full**—Specifies to verify both the header information and the rows structure (full verification of the backup).
- **verifyonly**—Specifies to verify minimal header information without restoring the database.

If you do not specify a verification value, the restore performs no verification and adds a message to the log file.

Performing Sybase data restores with the nsrsybrc command

You can run the **nsrsybrc** command from the command line, as the operating system user that launched the Sybase Backup Server, to perform a Sybase restore of the whole Sybase server or of individual databases.

NOTICE

In the **nsrsybrc** command, the Sybase server name and database name are case-sensitive and the names must be in the same case as recorded in the corresponding backup entries in the NetWorker indexes.

By default, the **nsrsybrc** command restores the most recent database backup and the command rolls transactions forward by recovering the transaction logs. The operation brings the database back online at the end of the restore.

The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrsybrc** command.

After the restore operation completes, perform the following steps:

1. Check the database to ensure that the data has been restored.
2. Run a database consistency check as described in "Performing Sybase database consistency checks before backups" on page 154.
3. Run a full backup of the database.

Example 26 on page 217 describes a restore of the Sybase server and a single database or multiple databases.

Example 27 on page 218 describes a point-in-time recovery.

The following topics describe the other supported types of Sybase data restores:

- "Performing Sybase redirected restores" on page 218
- "Performing Sybase restores to a different host" on page 220

Example 26 Restores of a single Sybase database and multiple databases

- This command restores a single Sybase database:

```
nsrsybrc -U user_ID -P password -s NetWorker_server_name
SYBASE:/ASE_server_name/database_name
```

  where:

  - *user_ID* is the username of the Sybase user account.
  - *password* is the password of the Sybase user account.
  - *NetWorker_server_name* is the hostname of the NetWorker server.
  - *ASE_server_name* is the Sybase server name.
  - *database_name* is the name of the database on the Sybase server.

- This command restores multiple Sybase databases:

```
nsrsybrc -U user_ID -P password -s NetWorker_server_name
SYBASE:/ASE_server_name/database1_name
SYBASE:/ASE_server_name/database2_name
SYBASE:/ASE_server_name/database3_name
```

- This command restores all the databases on the Sybase server, except the master database:

```
nsrsybrc -U user_ID -P password -s NetWorker_server_name
SYBASE:/ASE_server_name
```

Example 27 Point-in-time recovery of Sybase data

A point-in-time recovery recovers the Sybase data to a specific time in the past, without restoring the entire transaction log. The **nsrsybrc** command uses the time specified with the **-t** option to restore data to a specific point-in-time. The operation loads the most recent full backup before the designated time and then applies transaction log backups up to the designated time.

The following command performs a point-in-time recovery:

```
nsrsybrc -U user_ID -P password -s NetWorker_server_name
   -t "MM/DD/YY HH:MM:SS" SYBASE:/ASE_server_name/database_name
```

where *MM/DD/YY HH:MM:SS* indicates the month, day, year, hour, minute, and seconds to recover data to. The *ASE_server_name* and *database_name* must be in the same case as recorded in the backup entries in the NetWorker indexes.

Note: Because the NetWorker server and client can be in different time zones, set the **-t** option to a value of the local time on the Sybase server. A successful point-in-time recovery brings the database online.

NOTICE

After you perform a point-in-time recovery, the Sybase server restarts the database log sequence. Performing an incremental backup before a full backup causes future restore and recovery operations to fail. Perform a full backup after a point-in-time recovery.

Performing Sybase redirected restores

A redirected restore loads the backup of an old database to a new database.

NMDA supports the following types of redirected restores of Sybase data:

- "Same Sybase server but a different database name" on page 218
- "Different Sybase server but the same database name" on page 219
- "Different Sybase server and a different database name" on page 219
- "Different database types" on page 219

Same Sybase server but a different database name

To restore data to the same Sybase server but to a different database name:

1. Ensure that you have created the new database with the proper device allocations.
2. Type the following command as the Sybase user:

   ```
   nsrsybrc -U user_ID -P password -s NetWorker_server_name
   -d SYBASE:/ASE_server_name/new_database_name
   SYBASE:/ASE_server_name/old_database_name
   ```

   where:

   - *user_ID* is the username of the Sybase user account.
   - *password* is the password of the Sybase user account.
   - *NetWorker_server_name* is the hostname of the NetWorker server.
   - *ASE_server_name* is the Sybase server name.

- *new_database_name* is the name of the database on the Sybase server to which the data is restored.
- *old_database_name* is the name of the database on the Sybase server that was backed up.

Different Sybase server but the same database name

To restore data to a different Sybase server but to the same database name:

1. Ensure that you have created the new database with the proper device allocations.
2. Type the following command as the Sybase user:

   ```
   nsrsybrc -U user_ID -P password -s NetWorker_server_name
   -d SYBASE:/new_Sybase_server_name/database_name
   SYBASE:/old_Sybase_server_name/database_name
   ```

   where:

   - *new_Sybase_server* is the name of the Sybase server to which the data is restored.
   - *old_Sybase_server* is the name of the Sybase server from which the data was backed up.

   Note: The username and password are for the new Sybase server.

Different Sybase server and a different database name

To restore data to a different Sybase server and to a different database name:

1. Ensure that you have created the new database with the same device allocations as the original database.
2. Type the following command as the Sybase user:

   ```
   nsrsybrc -U user_ID -P password -s NetWorker_server_name
   -d SYBASE:/new_Sybase_server_name/new_database_name
   SYBASE:/old_Sybase_server_name/old_database_name
   ```

   where:

   - *new_database_name* is the name of the new database on the new Sybase server.
   - *old_database_name* is the name of the old database on the old Sybase server.

   Note: The username and password are for the new Sybase server.

Different database types

With Sybase ASE 15.5, NMDA supports the restore of any supported type of database to any other supported type of database. For example, you can restore a regular database to an in-memory database, and you can restore a relaxed durability database to an in-memory database.

The restore instructions are the same as in the preceding topics for a redirected restore.

Performing Sybase restores to a different host

To restore Sybase data to a new destination host:

1. Follow the instructions in "Performing NMDA data restore and recovery" on page 169 to set up the NMDA software.
2. Run the **nsrsybrc** command with the **-c** option to specify the NetWorker client where the backup was performed. For example:

   ```
   nsrsybrc -U user_ID -P password -s NetWorker_server_name
   -c original_client_name SYBASE:/ASE_server_name/database_name
   ```

   Note: Although the **-d** option indicates the destination for restore, this command does not use the option because the destination server name and database name are the same as the original NetWorker client.

When you restore data to a new destination host, you can also relocate the data to a new Sybase server or database. For example, run the following command on the destination host to relocate the data to a new Sybase server:

```
nsrsybrc -U user_ID -P password -s NetWorker_server_name -c
   original_client_name -d
   SYBASE:/new_Sybase_server_name/database_name
   SYBASE:/old_Sybase_server_name/database_name
```

Performing Sybase data restores with NetWorker User for Sybase

On Windows systems only, you can run the NetWorker User for Sybase GUI (**nwbms.exe**) to restore Sybase data.

You must run the Sybase GUI on the destination host where the data will be restored.

1. Start the GUI. For example, select **NetWorker User for Sybase** from **Start** › **Programs** › **EMC NetWorker**.
2. To connect to a different NetWorker server, complete these steps:
   a. Select **Select NetWorker Server** from the **Operation** menu.

      The **Change Server** dialog box appears.

   b. Click **Update List** to refresh the list of NetWorker servers.

   c. Select or type the name of the server.

   d. Select **Save as Default Server** to use the server as the default NetWorker server.

   e. Click **OK**.
3. In the NetWorker User for Sybase GUI, select **Recover** from the **Operation** menu.

   The **Recover** window appears. The online help describes the toolbar buttons.

4. Complete the **Recover Sybase Server** dialog box if it appears:
   a. Type the required values in the **Recover Sybase Server** dialog box:
      – For **Server name**, type the name of the backed-up Sybase server.
      – For **Host name**, type the operating system hostname of the backed-up Sybase server.
   b. Click **OK**.

   Note: For restores, the Sybase server name and database name that you type in the GUI are case-sensitive. The names must be in the same case as recorded in the corresponding backup entries in the NetWorker indexes.

5. To view a list of files or databases available for restore, select the Sybase server in the left pane. The Sybase server contents appear in the right pane. By default, the GUI shows the latest available backups.
6. To view previous backups for recovery to a previous point-in-time, change the browse time:
   a. In the **Recover** window, select **Change Browse Time** from the **View** menu.

      The **Change Browse Time** dialog box appears.
   b. Set a new date by selecting a day from the calendar.
   c. Click **Previous Month** or **Next Month** to change from the current month.
   d. In the **Time** text box, type a time to browse.

   Note: The browse time cannot be earlier than the time of the first backup because the client file index does not have entries before that time. To verify the browse policy and retention policies, check the Client resources for the client by using the NetWorker administration program.

7. Select the checkbox next to each file or database to be restored.
8. If you want to relocate the restored data to a new Sybase server or a new database, ensure that you have created any new database with the proper device allocation.
9. From the **Options** menu, select **Recover Options**.

   The **Recover to Sybase Server** dialog box appears as shown in the following figure.

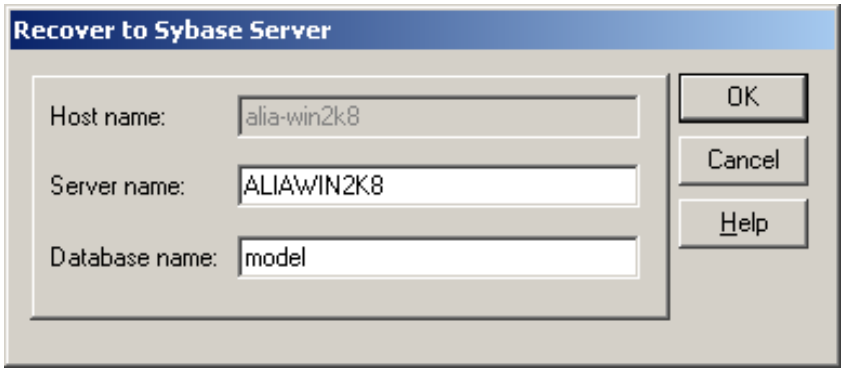

Figure 8 Recover to Sybase Server dialog box in NetWorker User for Sybase

10. In the **Recover to Sybase Server** dialog box, specify the required options:
    a. Type the required values for the Sybase server name and database name.

    Note: The Sybase server name and database name are case-sensitive. You must type the names in the same case as recorded in the backup entries in the NetWorker indexes.
    If you leave either text box blank, NMDA uses the same Sybase server name or database name as existed at the backup time.

    - To restore data to the same Sybase server but to a different database name:
      - In the **Server name** text box, optionally type the old Sybase server name.
      - In the **Database name** text box, type the new database name.
    - To restore data to a different Sybase server but to the same database name:
      - In the **Server name** text box, type the new Sybase server name.
      - In the **Database name** text box, optionally type the old database name.
    - To restore data to a different Sybase server and to a different database name:
      - In the **Server name** text box, type the new Sybase server name.
      - In the **Database name** text box, type the new database name.

    b. Click OK.
11. Ensure that you have set the backup volumes for the restore:
    a. In the **Recover** window, ensure that you have selected the required entries for recovery.
    b. From the **View** menu, select **Required Volumes**.

    The **Required Volumes** window appears with the backup volumes listed.
    c. Load and mount the required volumes, as appropriate.
12. Click **Start** in the **Recover** window.

    The **Recover Status** window appears with information about the recovery.

Canceling a data restore or recovery

You can cancel a data restore or recovery by using the following procedures.

NOTICE

Using any other procedures to cancel a data restore or recovery in progress can corrupt the database being recovered. If you cancel a data restore or recovery, rerun the operation to ensure that the database files are successfully restored.

NMDA does not support other procedures for canceling a data restore or recovery. For example, the **kill -9** command or Windows equivalent might cause a Lotus Domino server to crash.

- To cancel a restore or recovery that you started from the command line, press **Ctrl-c**.
- To cancel a restore and recovery from the NetWorker User for Lotus or the NetWorker User for Sybase program, select **End Recover** from the **File** menu.

  Note: This method does not cancel a directed recovery of a Lotus Windows client when the performing client is different than the destination client. As a workaround, terminate the **nsrlotus_remrecov** process by using the Task Manager on the remote Windows client.

CHAPTER 5
Disaster Recovery

This chapter includes the following topics:

- Preparing for disaster recovery 226
- Preparing for DB2 disaster recovery 227
- Preparing for Informix disaster recovery 228
- Preparing for Lotus disaster recovery 231
- Preparing for MySQL disaster recovery 235
- Preparing for Oracle disaster recovery 236
- Preparing for Sybase disaster recovery 240

Note: This chapter describes how to recover from a disaster. Use this chapter with Chapter 4, "Data Restore and Recovery" and the *EMC NetWorker Disaster Recovery Guide*.

Preparing for disaster recovery

A disaster is any loss of data where the computing environment is not available to restore the data. Ordinary data recovery procedures are not sufficient to recover the computing environment and the data to normal day-to-day operations.

The following examples are possible causes of a disaster:

- Computer viruses that corrupt the computing system
- Hardware failures and software failures
- Infrastructure interruptions, inconsistencies, or loss of services, such as communications or network services, which result in damage to the computing environment

To develop a plan for recovering from a disaster on the computer system, determine the required frequency of backups. Consider that the backup frequency is a trade-off between the time spent backing up data and the time spent later recovering the data after a disaster.

To prepare the environment of a database server or application server for disaster recovery, implement a backup strategy that includes regular backups of the following files:

- Database files
- Transaction logs
- Any other control files and configuration files that you need for the database server or application server operations

  The topics in this chapter describe the specific files to back up for the database or application.
- NetWorker server bootstrap and index records (media database, resource database, and server index that reside on the NetWorker server)

  NetWorker server disaster recovery requires the NetWorker bootstrap records. The *EMC NetWorker Disaster Recovery Guide* provides details.

  At the end of an NMDA scheduled backup, the software automatically backs up the NetWorker bootstrap and index records. If you run only manual backups, perform regular backups of the NetWorker server bootstrap and client file indexes according to instructions in the *EMC NetWorker Administration Guide*.

Depending on the type of disaster, you might need to perform any of the following operations:

Note: In cluster environments (both active-passive clusters and active-active clusters), you must perform these operations on all the required cluster hosts.

- Reinstall the operating system.
- Reinstall the database or application software.
- Configure the database server or instance.
- Reinstall the NetWorker client and NMDA software. Reinstall additional software if NMDA backups use the software, such as the Replication Manager software.

The operating system documentation and the database vendor documentation describe procedures for the operating system and database or application.

The *EMC NetWorker Installation Guide* describes how to reinstall the NetWorker client.

The *EMC NetWorker Module for Databases and Applications Installation Guide* describes how to reinstall NMDA.

When you perform a disaster recovery to a new destination host, perform all the operations described in the following topics on the new host. In addition, to enable the restore of backups done on the original host to the new host, set the Remote Access attribute. Set the attribute in the NetWorker Client resource of the original (source) host (in NMC or the NetWorker server) to the following value:

database_username@destination_hostname

When you perform a restore to a remote destination host, you might not want to transfer data over a slow network. You could use the NetWorker storage node feature to improve restore performance. After physically transferring backup media (for example, tapes) containing the backup to the remote location, install the storage node on the remote computer. Configure the NetWorker server to use the remote device of the storage node. The *EMC NetWorker Administration Guide* describes storage nodes.

NOTICE

After disaster recovery is complete, perform a full backup of the environment.

Preparing for DB2 disaster recovery

To prepare a DB2 server environment for disaster recovery, perform scheduled backups of the DB2 server instances and DB2 transaction logs with the NSR_DR_BACKUP_INFO parameter set to TRUE. When NSR_DR_BACKUP_INFO is set to TRUE, the NMDA **nsrdasv** command performs the following operations after the database backup.

1. Runs the following DB2 commands to generate the system config file:

   ```
   db2 get admin cfg
   db2 get dbm cfg
   db2 get db cfg for DATABASE_NAME
   db2 get instance
   db2 list db directory
   db2level
   db2set -all
   ```

2. Runs the following commands to generate the registry file:

   - On UNIX:

     ```
     db2greg -dump
     ```

   - On Windows:

     ```
     regedit /e db2_registry_config_timestamp.txt
     HKEY_LOCAL_MACHINE\Software\IBM\
     ```

3. Stores the command output from the preceding steps in temporary files such db2_system_config_1257779556.txt and db2_registry_config_1257779580.txt in the following directory:
   - On UNIX: /nsr/apps/tmp
   - On Windows: *NetWorker_install_path*\apps\tmp
4. Backs up the temporary file from step 3 by using the NetWorker **save** interface.

In addition, you can set the NSR_DR_FILE_LIST parameter for scheduled backups to specify a file that contains a list of extra files for the NetWorker **save** interface to back up. "NSR_DR_FILE_LIST" on page 369 provides details.

After a manual backup, you can manually back up the critical DB2 files by using the NetWorker file system backup interface, **save**. The *EMC NetWorker Administration Guide* provides details.

During a disaster recovery, you can use the NetWorker **recover** command to restore the backed-up files. For example:

```
recover -s NetWorker_server_name temporary_file
```

The DB2 backup and recovery documentation describes how to recover a DB2 database. "Performing DB2 data restore and recovery" on page 171 provides details.

Note: You must use the **db2 restore** command and **db2 rollforward** command to perform disaster recovery. You cannot use the **db2 recover** command for disaster recovery.

Disaster recovery of a DB2 pureScale system with DB2 release 10.1 requires that the pureScale system in the disaster recovery environment has the same topology as the production system. Addition or removal of any member nodes in the recovered pureScale system affects the use of the **rollforward** command for the recovery as described in the DB2 documentation. This limitation is removed in DB2 10.5 as described in "DB2 10.5 features" on page 36.

Preparing for Informix disaster recovery

To prepare an Informix server environment for disaster recovery, set the NSR_DR_BACKUP_INFO parameter to TRUE and perform regular scheduled backups of the Informix server instances and logical logs. The NSR_DR_BACKUP_INFO setting ensures that the scheduled backup backs up the following critical information:

- Informix ONCONFIG file
- ixbar file
- onconfig boot file
- sqlhosts file (UNIX only)
- sm_versions file
- Copy of Windows registry information stored under "HKEY_LOCAL_MACHINE\SOFTWARE\Informix" (Windows only)

"NSR_DR_BACKUP_INFO" on page 371 provides details.

In addition, you can set the NSR_DR_FILE_LIST parameter for scheduled backups to specify a file that contains a list of extra files to back up, instead of the preceding default list. You must also set NSR_DR_BACKUP_INFO to TRUE to enable the NSR_DR_FILE_LIST setting. "NSR_DR_FILE_LIST" on page 371 provides details.

The NetWorker **save** interface backs up the files created with the NSR_DR_BACKUP_INFO setting and NSR_DR_FILE_LIST setting. You must restore the files by using the NetWorker **recover** interface.

After a manual backup, you can manually back up the critical Informix files by using the NetWorker file system backup interface, **save**. The *EMC NetWorker Administration Guide* provides details.

Performing an Informix imported restore

You can use the Informix imported restore feature to transfer all the data from one Informix server instance to the same instance on a new destination host. You can perform imported restores by using either whole system (serial) backups or storage-space (parallel) backups. You also need to have a backup of the disaster recovery files described in "Preparing for Informix disaster recovery" on page 228.

You can use the following steps to perform the Informix imported restore.

1. Follow the instructions in "Preparing for disaster recovery" on page 226 to install the software, configure the database on the destination host, and set the Remote Access attribute in the Client resource on the NetWorker server, as required.

   Note: Create the data spaces in the same path location as on the source server.

2. On Windows systems, use **regedit** to copy the **sqlhosts** information from the source computer to the destination computer. Use the following registry entry:

   ```
   HKEY_LOCAL_MACHINE/SOFTWARE/Informix/SQLHOSTS/...
   ```

3. If the source INFORMIXDIR does not match the destination INFORMIXDIR, create a symbolic link to recover the bootstrap from the source computer. For example, if INFORMIXDIR on the source computer was /usr2/informix and INFORMIXDIR on the destination computer is /usr/local/informix, create the /usr2 directory on the destination computer and the symbolic link as follows:

   - On UNIX:

     ```
     mkdir /usr2
     ln -s /usr/local/informix /usr2/informix
     ```

   - On Windows:

     Create a shortcut from the source INFORMIXDIR to the destination INFORMIXDIR. The operating system documentation provides details on creating a shortcut.

NOTICE

To perform an imported restore, you must use the same database server number on the destination computer as used on the source computer. You can change the database server name in an imported restore.

4. Shut down the destination database server before doing a restore. For an IDS Dynamic Server, type the following command:

   onmode -ky

5. Set the required environment variables. For example:

   ```
   NSR_CLIENT = source_machine_hostname
   NSR_SERVER = backup_server
   ```

6. Restore the following Informix critical files:

   - For the Informix server:
     - $INFORMIXDIR/etc/oncfg_*original_dbserver_name.server_number*
     - ixbar.*server_number*
     - oncfg_*source_dbserver_name.server_number*
     - ixbar.*server_number*
     - $ONCONFIG
     - sqlhosts (UNIX only)

       On Windows, use **regedit** to copy the sqlhosts information from the source computer to the destination computer. Use the following registry entry:

       ```
       HKEY_LOCAL_MACHINE/SOFTWARE/Informix/SQLHOSTS/...
       ```

   - For all coserver numbers, copy the oncfg files:
     - oncfg_*source_dbserver_name.server_number.coserver_number*
     - sqlhosts
     - xcfg_*source_computer.server_number*

   For example, use the NetWorker software to recover the emergency boot file and configuration file for the Informix server instance, with $INFORMIXDIR set to the same value as at backup time. With the **recover** command, you must use the **-c** option to specify the original NetWorker client hostname:

   ```
   recover -a -s source_NetWorker_client_hostname\
   $INFORMIXDIR/etc/sqlhosts \
   $INFORMIXDIR/etc/onconfig.std \
   $INFORMIXDIR/etc/ixbar.server_number \
   $INFORMIXDIR/oncfg_server_name.server_number
   ```

7. Rename the $INFORMIXDIR/etc/oncfg_*original_dbserver_name.server_number* file and replace the source server name with the destination server name. For example:

   ```
   $INFORMIXDIR/etc/oncfg_ol_destination_dbserver_name.server_number
   ```

8. Update the sqlhosts file and include the proper shared memory and proper network settings for the destination Informix server.

9. Update the ONCONFIG file and replace the source server name with the destination server name. For example:

   ```
   DBSERVERNAME ol_destination_dbserver_name
   ```

10. Perform a full system restore with the following **onbar** command:

    ```
    onbar -r
    ```

11. Update the sqlhosts file and replace the source hostname with the destination hostname.

Recovering from an Informix server disk crash

You can perform an Informix cold restore to restore a damaged primary disk that contains critical Informix server dbobjects and NetWorker client binaries.

You can use the following steps to perform an Informix cold restore.

1. Follow the instructions in "Preparing for disaster recovery" on page 226 to install the software, configure the database on the destination host, and perform other configurations as required.
2. Use the NetWorker software to recover the emergency boot file and configuration file for the Informix server instance.

   Note: $INFORMIXDIR must have the same value as the parameter value at backup time.

3. If you must replace the physical media that contains the logical logs before you start the restore, manually salvage the current logical log file with the following **onbar** command:

   ```
   onbar -b -l -s
   ```

4. Restore data from the most recent backup with the following **onbar** command:

   ```
   onbar -r
   ```

   Once the restore completes, the Informix server remains in quiescent mode.

The Informix server documentation describes how to use the **onbar** command to restore data from the NetWorker backup media.

Preparing for Lotus disaster recovery

To prepare a Lotus server environment for disaster recovery, perform regular scheduled backups of the following items:

- Lotus data directory, unless you need to protect only certain subdirectories in the data directory
- notes.ini file
- Transaction logs

Note: NMDA backs up the transaction logs only if you have enabled Domino transaction logging and set the logging to archive mode.

By default, NMDA backs up only specific types of files, as described in "Files backed up during Lotus backups" on page 38. If required, enable the backup of all types of files by setting the parameter NSR_BACKUP_ALL_EXTENSIONS=TRUE. You can explicitly list files with nondefault extensions in the NSR_BACKUP_PATHS parameter when you have not set NSR_BACKUP_ALL_EXTENSIONS. "NSR_BACKUP_ALL_EXTENSIONS" on page 372 provides details.

Perform Lotus disaster recovery by using the appropriate procedures:

- "Recovering a nonlogged Lotus environment" on page 232
- "Recovering a logged Lotus environment" on page 233

NOTICE

Do not use the NetWorker User for Lotus GUI for disaster recovery.

Recovering a nonlogged Lotus environment

You can use the following steps to recover a Lotus Domino environment that is not in archived log mode.

Note: For a partitioned Domino server, repeat these steps for each partition.

1. Follow the instructions in "Preparing for disaster recovery" on page 226 to install the software and perform other configurations as required.

   Note: Reinstall the Lotus Notes client or Domino server software in the same location as before, but do not configure software.

2. When you recover a partitioned Domino server, ensure that the PATH parameter in the NMDA configuration file lists the data directory of the partition to be recovered before the data directory of any other partition. Appendix A, "NMDA Parameters and Configuration File," describes the NMDA configuration file.
3. Recover the notes.ini file by using the **nsrnotesrc** command with the parameter setting NSR_NO_NOTES_INIT = TRUE in the NMDA configuration file.

   For example, to recover the notes.ini file on a Windows system:

   a. Ensure that the NMDA configuration file contains the following parameter settings:

   ```
   NSR_SERVER = NetWorker_server_name
   NSR_BACKUP_PATHS = C:\Lotus\Domino\notes.ini
   NSR_NO_NOTES_INIT = TRUE
   ```

   NSR_BACKUP_PATHS specifies the case-sensitive path of the notes.ini file as recorded in the NetWorker indexes.

   b. Type the **nsrnotesrc** command to perform the recovery:

   ```
   nsrnotesrc -z configuration_file_path
   ```

c. When prompted whether to overwrite the current notes.ini file, type **y**.

Note: To prevent the prompting, you can set the parameter NSR_RECOV_INTERACT = Y in the configuration file.

4. Recover all the databases by using the proper NSR_BACKUP_PATHS setting:

   - For a nonpartitioned Domino server or Notes client:

     ```
     NSR_BACKUP_PATHS = Lotus_top_data_directory
     ```

   - For a partitioned Domino server:

     ```
     NSR_BACKUP_PATHS = Lotus_partition_top_data_directory
     ```

   You must recover the databases to a new location by using the **nsrnotesrc** command with the parameter setting NSR_RELOCATION_DEST = *destination_path* in the NMDA configuration file.

5. After the disaster recovery process is complete, perform a full backup of the Domino server to prevent any future loss of data.

Recovering a logged Lotus environment

To recover database backups to the last committed transaction in the archived transaction logs, you must meet the following requirements:

- The Domino server or the partition to be recovered had Lotus transactional logging enabled and set to Archive style.
- A backup of an up-to-date notes.ini file is available for the Domino server.
- A recoverable backup of the database files is available.
- You backed up the archived log extents (the transaction log files), which are available from the time of the last full backup.

You can use the following steps to recover a Domino environment that is in archived log mode:

Note: For a partitioned Domino server, repeat these steps for each partition.

1. Follow the instructions in "Preparing for disaster recovery" on page 226 to install the software and perform other configurations as required.

   Note: Reinstall the Domino server software in the same location as before, but do not configure it.

2. When you recover a partitioned Domino server, ensure that the PATH parameter is set in the NMDA configuration file. The parameter must list the Domino data directory of the partition to be recovered before the data directory of any other Domino server partition. Appendix A, "NMDA Parameters and Configuration File," describes the NMDA configuration file.

3. Recover the Domino notes.ini file by using the **nsrnotesrc** command with the parameter setting NSR_NO_NOTES_INIT = TRUE in the NMDA configuration file.

   For example, to recover the notes.ini file on a Windows system:

   a. Ensure that the NMDA configuration file contains the following parameter settings:

      ```
      NSR_SERVER = NetWorker_server_name
      NSR_BACKUP_PATHS = C:\Lotus\Domino\notes.ini
      NSR_NO_NOTES_INIT = TRUE
      ```

   b. Type the **nsrnotesrc** command to perform the recovery:

      ```
      nsrnotesrc -z configuration_file_path
      ```

   c. When prompted whether to overwrite the current notes.ini file, type **y**.

      Note: To prevent the prompting, you can set the parameter NSR_RECOV_INTERACT = Y in the configuration file.

4. Check the Domino notes.ini file to determine the original log directory for the server, as specified by the TRANSLOG_Path setting. Ensure that the directory exists and contains no old files.

5. Restore the last archived log extent backed up since the most recent full backup. Restore the log file to a temporary directory by using the **nsrnotesrc** command with the following restore parameters set in the NMDA configuration file:

   ```
   NSR_NO_NOTES_INIT = TRUE
   NSR_NUMBER_LOGS = 1
   NSR_RELOCATION_DEST = temporary_directory_path
   ```

   Note: For a partitioned Domino server, also set NSR_LOG_DIR for the restore of log files. Set the parameter to the original log directory pathname of the recovered partition.

   For example, to restore the archived log file to the temporary directory D:\temp\Lotus directory on Windows:

   a. Ensure that the NMDA configuration file contains these parameter settings:

      ```
      NSR_NO_NOTES_INIT = TRUE
      NSR_NUMBER_LOGS = 1
      NSR_RELOCATION_DEST = D:\temp\Lotus
      ```

   b. Type the **nsrnotesrc** command to perform the restore:

      ```
      nsrnotesrc -z configuration_file_path
      ```

6. Copy the restored log files from the temporary directory to the original log directory for the server as specified by the TRANSLOG_Path setting.

7. Enable the creation of the control file by setting the following parameter in the Domino notes.ini file:

   ```
   TRANSLOG_Recreate_Logctrl=1
   ```

8. Recover all the databases into a temporary directory by using the proper NSR_BACKUP_PATHS setting:

   - For a nonpartitioned Domino server or Notes client:

     ```
     NSR_BACKUP_PATHS = Lotus_top_data_directory
     ```

   - For a partitioned Domino server:

     ```
     NSR_BACKUP_PATHS = Lotus_partition_top_data_directory
     ```

   - For a Domino server with DAOS files:

     If the backup includes DAOS files and the DAOS directory is outside of the Lotus data directory, add the DAOS directory to NSR_BACKUP_PATHS (separated by a comma):

     – For a nonpartitioned Domino server:

       ```
       NSR_BACKUP_PATHS = Lotus_top_data_directory,
       Lotus_top_DAOS_directory
       ```

     – For a partitioned Domino server:

       ```
       NSR_BACKUP_PATHS = Lotus_partition_top_data_directory,
       Lotus_partition_top_DAOS_directory
       ```

   You must recover the databases to a new location by using the **nsrnotesrc** command with the parameter setting NSR_RELOCATION_DEST = *destination_path* in the NMDA configuration file.

9. After you recover the databases or DAOS files, copy the files to the Lotus data directory or DAOS directory of the Domino server or a specific partition.
10. For a nonpartitioned Domino server on a Windows system, copy the recovered notes.ini file to the original directory. By default, the notes.ini file is not in the data directory.
11. Start the Domino server.
12. After the disaster recovery process is complete, perform a full backup of the Domino server to prevent any future loss of data.

Preparing for MySQL disaster recovery

To prepare a MySQL server environment for disaster recovery, perform frequent backups of the critical data components:

- Perform regularly scheduled full backups and incremental backups of the whole MySQL instance.
- If binary logging is enabled, perform regularly scheduled backups of the MySQL binary logs. Set MYSQL_LOG_OPTIONS in the NMDA configuration file to specify the binary log backups.

- Perform manual or scheduled backups of any MySQL configuration files, also called option files, and other important files:
  - To back up the files manually, run the NetWorker **save** command.
  - To include the MySQL configuration files and other important files in scheduled backups, set NSR_DR_BACKUP_INFO and NSR_DR_FILE_LIST:
    - Set NSR_DR_BACKUP_INFO to include the MySQL configuration file specified by MYSQL_CFG_FILE in the scheduled backups.
    - Set NSR_DR_FILE_LIST to include a list of additional configuration files and other important files in the scheduled backups.

MySQL documentation describes the MySQL configuration files and other important files. The *EMC NetWorker Administration Guide* describes the **save** command.

Performing a MySQL disaster recovery

You can use the following steps to perform a MySQL disaster recovery.

1. Perform the required procedures to set up the system hardware, install the operating system, and install the MySQL server software. Ensure that the MySQL server version is the same as used to create the MySQL backups.
2. Install the NetWorker and NMDA software according to the appropriate installation guides.
3. Run the NetWorker **recover** command to restore the MySQL configuration files and other important files backed up with the **save** command or the NSR_DR_BACKUP_INFO and NSR_DR_FILE_LIST settings. Restore the files to the appropriate directories. For example:

   ```
   recover -s NetWorker_server_hostname -c original_client_hostname
   -d /etc -a /etc/my.cnf
   ```

   The *EMC NetWorker Administration Guide* describes the **recover** command.
4. Restore any binary log backups according to "Performing MySQL restores of binary log backups" on page 201.
5. Recover the whole MySQL instance to the current time according to "Performing MySQL recovery of whole instance backups" on page 198. To perform a redirected restore of the instance to a different server than the server from which the data was backed up, follow the instructions in "Performing MySQL redirected restores" on page 201.

Preparing for Oracle disaster recovery

To prepare an Oracle server for disaster recovery, back up this minimum list of files:

- Oracle database (all the datafiles)
- Archived redo logs
- Control file
- Initialization parameter files, including one or both of the following files:
  - PFILE (user-managed parameter file)
  - SPFILE (server-managed parameter file)

- Network files, including listener.ora, sqlnet.ora, tnsnames.ora
- Text file that contains the Oracle DBID
- Password file, in the following location by default:
  - On UNIX, $ORACLE_HOME/dbs/orapw$ORACLE_SID
  - On Windows, %ORACLE_HOME%\database\PWD%ORACLE_SID%.ora
- Registry files:
  - On UNIX, oratab is typically in /var/opt/oracle or /etc
  - On Windows, *My Compute*r\HKEY_LOCAL_MACHINE\SOFTWARE\Oracle
- Recovery Catalog, if applicable
- RMAN scripts, if applicable

Note: At the end of a successful Oracle scheduled backup, NMDA automatically backs up the NWORA resource file if the file exists.

The Oracle documentation provides an exhaustive list of the files (other than the Oracle database) that you must back up.

Follow these guidelines to prepare for disaster recovery:

- Institute mirrored control files.

  Refer to Oracle documentation for recommendations on whether to institute mirrored online redo logs.
- Back up the archived redo logs frequently between database backups.
- Back up the Recovery Catalog after every target database backup.

You can use the following steps to back up the required files to prepare for disaster recovery.

1. Record the Oracle DBID in a text file and back up the file. "Creating the DBID text file" on page 238 provides details.
2. Use a postcommand script to back up files that Oracle RMAN does not back up. "Setting up a postcommand script for backup of Oracle files" on page 238 provides details.
3. Use an RMAN backup with NMDA to back up the Oracle database and related files. "Setting up RMAN backups of the database and related files" on page 239 provides details.
4. Use an RMAN backup to back up the Recovery Catalog. "Setting up RMAN backups of Recovery Catalog" on page 240 provides details.

Perform an Oracle disaster recovery according to "Performing an Oracle disaster recovery" on page 240.

Creating the DBID text file

The Oracle DBID is an internal Oracle ID that helps Oracle find the autobackup of the SPFILE if the Recovery Catalog is not accessible.

Before you can back up the Oracle DBID, you must manually record the DBID in a text file. The simplest way to find the DBID of an Oracle database is to connect to the database through RMAN after the database is mounted.

After you have recorded the DBID in a text file, you can store the text file containing the DBID in any directory where you have the proper operating system permissions. You can use a postcommand script to back up the DBID text file, as described in "Setting up a postcommand script for backup of Oracle files" on page 238.

Setting up a postcommand script for backup of Oracle files

You can use a postcommand script to back up the files that Oracle RMAN does not back up, such as the following files:

- Initialization parameter file PFILE (user-managed parameter file)
- Network files, including listener.ora, sqlnet.ora, tnsnames.ora
- Text file that contains the Oracle DBID, as described in "Creating the DBID text file" on page 238
- Password file
- Registry
- RMAN scripts, if applicable

You can either create a postcommand script from scratch or modify the postcommand script included with NMDA.

On UNIX systems, you can use any name for the postcommand script. On Windows systems, the script name must end in .bat.

In a scheduled NMDA backup, include the postcommand script by using one of the following methods:

- If you configure the scheduled backup with the NMDA wizard, specify the postcommand script in the wizard.
- If you configure the scheduled backup without the NMDA wizard, set the POSTCMD parameter in the NMDA configuration file.

"Sample Oracle postcommand script" on page 238 describes the postcommand script included with NMDA.

After a manual backup, you can manually back up these files through the NetWorker file system backup program, **save**. The *EMC NetWorker Administration Guide* provides details.

Sample Oracle postcommand script

The NMDA installation provides a sample postcommand script that is specific to UNIX or Windows, depending on the operating system. The sample script is in the bin subdirectory under the NetWorker software directory, for example, under /usr/sbin.

You must customize the settings in the sample postcommand script for the specific environment. At a minimum, you must set the ORACLE_HOME and ORACLE_SID parameters in the script. If you do not set these two parameters, the postcommand script fails at runtime.

During a scheduled NMDA backup, the **nsrdasv** process passes the options **-s** *server_name* **-g** *group_name* to the postcommand script if the script name begins with nsrnmodr. This action ensures that the additional files are backed up to the same devices and have the same browse policy and retention policy as the Oracle datafiles backed up through RMAN.

If the script name does not begin with the nsrnmodr prefix, you must provide the options **-s** *server_name* **-g** *group_name* to the postcommand script.

Example 28 on page 239 shows the sample postcommand script installed with NMDA.

Example 28 Sample postcommand script on UNIX

```
##!/bin/ksh
#
ORACLE_HOME=
ORACLE_SID=
GRP=no
SRV=no
complete=0
shift
shift
while [ "$#" -gt "0" ]
do
   if [ "$1" = "-g" ]; then
      GRP=$2
      if [ "$SRV" != "no" ]; then
         complete=1
      fi
   elif [ "$1" = "-s" ]; then
      SRV=$2
      if [ "$GRP" != "no" ]; then
         complete=1
      fi
   fi
   shift
done

if [ $complete -eq 1 ]; then
   save -s $SRV -g $GRP $ORACLE_HOME/network/admin
   save -s $SRV -g $GRP $ORACLE_HOME/dbs/orapw$ORACLE_SID
   save -s $SRV -g $GRP /var/opt/oracle/oratab
   save -s $SRV -g $GRP $ORACLE_HOME/dbid.txt
fi
```

Setting up RMAN backups of the database and related files

Set up an RMAN backup with NMDA to back up the following files:

- Oracle database (all the datafiles)
- Archived redo logs
- Control file
- Initialization parameter file SPFILE (server-managed parameter file)

Follow the instructions in the preceding chapters of this guide to configure and run the RMAN backup with NMDA.

For example, to include the control file and SPFILE in the backup, you can add the following commands to the RMAN backup script:

- **backup current control file**
- **backup spfile**

The RMAN documentation describes RMAN commands and scripts.

If you want to back up PFILE (user-managed parameter file) or other files that Oracle RMAN does not back up, you can use a postcommand script. "Setting up a postcommand script for backup of Oracle files" on page 238 describes the postcommand script.

Setting up RMAN backups of Recovery Catalog

Set up an RMAN backup of the Recovery Catalog by using the same method as for the target database backup, as described in "Setting up RMAN backups of the database and related files" on page 239.

The Oracle documentation describes Recovery Catalog backups.

Performing an Oracle disaster recovery

You can use the following steps to perform an Oracle disaster recovery.

1. Follow the instructions in "Preparing for disaster recovery" on page 226 to install the software, configure the database on the destination host, and set the Remote Access attribute in the Client resource on the NetWorker server, as required.
2. To recover Oracle files backed up through a postcommand script, use either the NetWorker User GUI or the **recover** command.

   For example, a typical **recover** command is as follows:

   ```
   recover -s NetWorker_server -c client_name_of_original_host
   -d /var/opt/oracle -a /var/opt/oracle/oratab
   ```

   Note: On a Windows system, you might need to reinsert the oracle.reg file into the registry after recovering the file, for example, with the following command:

   ```
   regedit /S C:\temp\oracle.reg
   ```

   The Oracle documentation provides details.
3. To perform the rest of the disaster recovery, follow the instructions in the *Oracle Database Backup and Recovery User's Guide*. In the RMAN script, set the NSR_CLIENT parameter to the name of the original host.

Preparing for Sybase disaster recovery

To prepare a Sybase server environment for disaster recovery, perform regular scheduled backups of the Sybase server, databases, and transaction logs.

Perform the following additional tasks to prepare for disaster recovery:

- Keep up-to-date printouts of the Sybase system tables.
- Keep up-to-date printouts of the scripts for disk init and create databases.

- Do not store user databases or any databases other than master, tempdb, model, and sybsystemdb on the master device.
- Back up the master database after performing actions such as initializing database devices, creating or altering databases, or adding a new server login.

The Sybase documentation provides details.

You can use the following steps to recover the Sybase server after a disaster.

1. Use the printout of database device allocations to re-create the databases. The Sybase documentation describes the information to track for disaster recovery.
2. Recover the Sybase system databases and user databases. The Sybase documentation provides details.
3. Use the **nsrsybrc** command to recover the Sybase data. "Performing Sybase data restore and recovery" on page 216 provides details.

Recovering the master database

The master database might be lost or corrupted in a disaster. The master database controls the operation of the Sybase server, and stores information about all user databases and associated database devices. The Sybase documentation describes how to recover the master database in different scenarios.

You can use the following example procedure to recover the master database on Sybase ASE version 15.0 and later under these conditions:

- The master device is lost.
- A valid dump exists and has a default sort order.
- All other devices are undamaged and do not require inspection.

NOTICE

The recovery of the master database to a different Sybase server copies all the device allocations to the new Sybase server. Therefore, if you recover the master database to another Sybase server on the same computer as the original, both servers try to use the same database files. To prevent this issue, follow the Sybase documentation to recover the master database to a different Sybase server.

You can use the following steps to recover the master database.

1. Rebuild the lost master device by using the **dataserver** command.
2. Start the Sybase server in single-user mode, also called master-recover mode.
3. Ensure that the Sybase server has the correct name for the Sybase Backup Server in the **sysservers** table.
4. Use the following **nsrsybrc** command to recover the master database from the backup:

   ```
   nsrsybrc -U user_ID -P password -s NetWorker_server_name
   SYBASE:/ASE_server_name/master
   ```

   where:

   - *user_ID* is the username for the Sybase user account.
   - *password* is the password for the Sybase user account.

- *NetWorker_server_name* is the hostname of the NetWorker server.
- *ASE_server_name* is the name of the Sybase server.

Note: In the **nsrsybrc** command, the *ASE_server_name* is case-sensitive and must be in the same case as recorded in the backup entries in the NetWorker indexes.

The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrsybrc** command.

After the master database is loaded, the Sybase server performs postprocessing checks and validations and then shuts down.

5. Restart the Sybase server.
6. If required, recover the model database and other databases that were on the master device.
7. Log in as Systems Administrator and inspect the databases on the Sybase server to ensure that all the databases are present.

Recovering user databases after database device failure

You can use the following example procedure to recover Sybase user databases after the database device, not the log device, fails. The Sybase documentation describes how to recover user databases.

You can use the following steps to recover a database after the database device fails.

1. Perform an incremental backup of each database on the failed device to back up the transaction logs:

   nsrdasv -z *configuration_file_path*

   where *configuration_file_path* is the pathname of the NMDA configuration file. The configuration file must contain the following parameter settings:

   ```
   NSR_BACKUP_LEVEL=incr
   NSR_BACKUP_PATHS=SYBASE:/ASE_server_name/database_name
   NSR_DUMP_LOG_OPT="no_truncate"
   NSR_SERVER=NetWorker_server_name
   SYBASE_USER=Sybase_username
   USER_PSWD=encrypted_password
   ```

   Appendix A, "NMDA Parameters and Configuration File," describes the NMDA parameters.

2. Determine the space usage of each database on the failed device. For example, on the **isql** command line, type the following commands:

   select segmap, size from sysusages where dbid = db_id("*database_name***")**

   sp_helpdb *database_name*

   where *database_name* is the name of the database on the failed device.

3. After you have obtained the information for all databases on the failed device, drop each database by using the **drop database** command.

   If the system reports errors due to database damage, use the **dropdb** option with the **dbcc dbrepair** command. On the **isql** command line, type the following command:

   ```
   dbcc dbrepair (database_name, dropdb)
   ```

   where *database_name* is the name of a database on the failed device.

4. Drop the failed device by using the **sp_dropdevice** system procedure.
5. Initialize the new devices by using the **disk init** command.
6. Re-create each database, one at a time, by using the **create database** command.
7. Recover each damaged database from the most recent database backup. For example:

   ```
   nsrsybrc -U user_ID -P password -s NetWorker_server_name
   SYBASE:/ASE_server_name/database_name
   ```

   where:

   - *user_ID* is the username for the Sybase user account.
   - *password* is the password for the Sybase user account.
   - *NetWorker_server_name* is the hostname of the NetWorker server.
   - *ASE_server_name* is the name of the Sybase server.
   - *database_name* is the name of a database on the Sybase server.

   Note: In the **nsrsybrc** command, *ASE_server_name* and *database_name* are case-sensitive and must be in the same case as recorded in the backup entries in the NetWorker indexes.

   The **nsrsybrc** command recovers the last full database backup of the specific database and applies all the associated transaction log backups in the order of their creation. The recovery also brings the database online.

   The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrsybrc** command.

CHAPTER 6
Cluster Systems

This chapter includes the following topics:

- Active-passive clusters 246
- Active-active application clusters (overview) 254
- DB2 DPF systems 256
- DB2 pureScale systems 260
- Informix MACH systems 264
- Oracle RAC systems 267
- Sybase ASE Cluster Edition systems 272

Active-passive clusters

An active-passive cluster is a group of linked virtual hosts or physical hosts with shared storage, called cluster nodes, which work together and represent themselves as a single host called a virtual cluster host. In an active-passive cluster, some nodes are active and others are stand-by, waiting to take over the processing when an active node fails. You can connect to the cluster by using a virtual cluster name or IP address, regardless of which nodes are active. You usually implement active-passive clusters for high availability solutions or to improve computer performance.

The following figures display examples of supported active-passive cluster configurations:

Note: Each database node in the figures corresponds to one of the following items:
- DB2 node
- Informix server
- Lotus Domino server or partition
- Oracle instance
- Sybase server

- Figure 9 on page 246—Single database node with failover capability
- Figure 10 on page 247—Multiple database nodes with failover capability
- Figure 11 on page 247—Multiple database nodes with mutual failover capability

Note: Some cluster documents or applications documents might describe the two active-passive clusters configured as shown in Figure 11 on page 247 as active-active clusters, for example, with a Lotus Domino partition server in a cluster. However, NMDA documentation refers to these clusters as active-passive clusters.

The *EMC NetWorker Software Compatibility Guide* describes all the supported cluster environments.

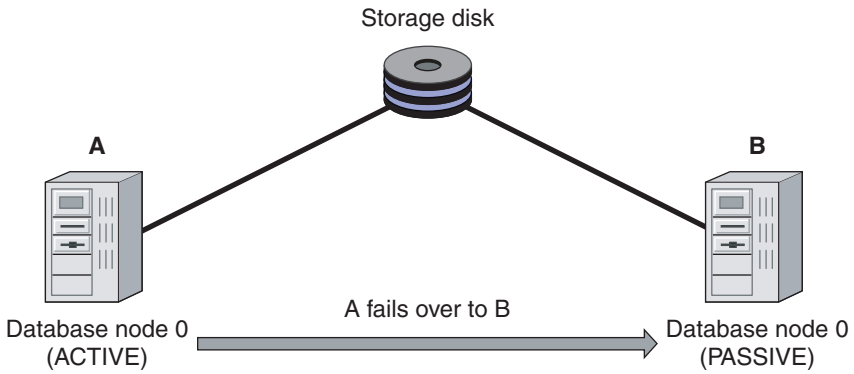

Figure 9 Single database node with failover capability

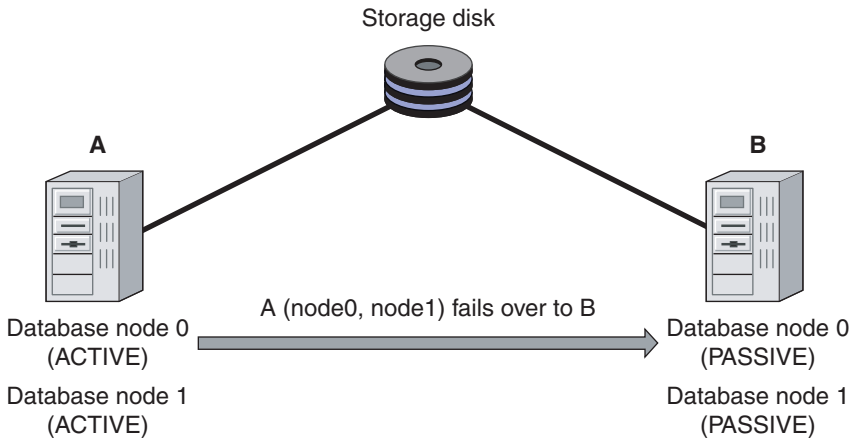

Figure 10 Multiple database nodes with failover capability

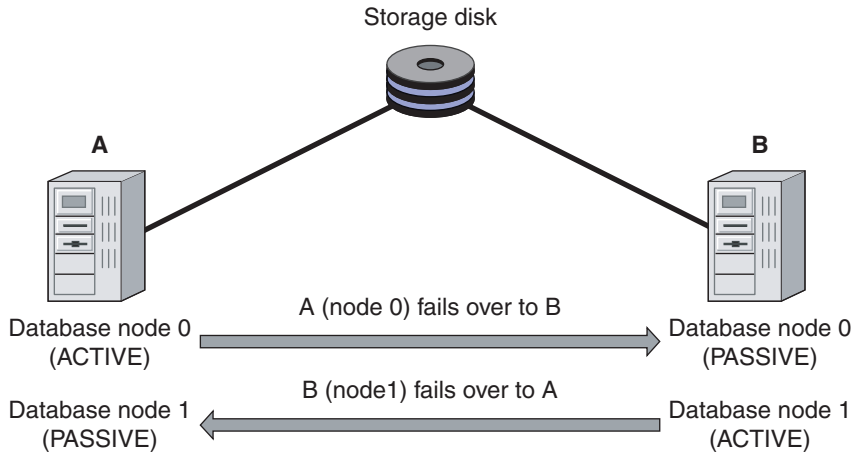

Figure 11 Multiple database nodes with mutual failover capability

Backup failover

When a node failure occurs during a manual backup, a DBA must restart the backup after cluster recovery.

When a node failure occurs during a scheduled backup, the NetWorker server restarts the backup if the Client Retries attribute in the NetWorker Group resource is set to a value greater than zero. If you configured a backup for a checkpoint restart, the NetWorker server restarts the backup from the beginning or from the point of failure. The restarted backup runs on the node that takes control of the cluster service.

Cluster configuration requirements

Before you configure backups in a cluster system, ensure that you meet the following requirements:

- You have set up the database in the cluster according to the appropriate database server documentation.
- You have installed NMDA and the NetWorker software on each node of the cluster that will participate in the backup or restore. The *EMC NetWorker Module for Databases and Applications Installation Guide* provides details.

Lotus cluster requirements

If you specify NSR_CLIENT in the configuration file when you use the NetWorker User for Lotus program to perform a manual backup on a cluster, the saved index uses the NSR_CLIENT name.

"Performing Lotus data restore and recovery" on page 181 provides details on how to configure and run a directed recovery.

Sybase cluster requirements

On UNIX or Linux systems, you must create the required symbolic link to the Sybase library:

- If the Sybase ASE binaries are on the shared disk, create the symbolic link to the Sybase library on the shared disk.
- If the Sybase ASE binaries are on each node in the cluster, create the symbolic link to the Sybase library on each node where you installed NMDA and you will run backups.

The *EMC NetWorker Module for Databases and Applications Installation Guide* describes how to create the symbolic link to the Sybase library.

NSR_CLIENT setting for cluster systems

During an NMDA backup, the NetWorker server creates entries about the backed-up data in the online client file index. During an NMDA restore, the software retrieves the data by first searching the client file index.

The NSR_CLIENT parameter provides the following information to the NetWorker server:

- During a backup, the name of the NetWorker client whose index file has the record of backup information.
- During a restore, the name of the NetWorker client whose index file is used to search for the data for restore.

If you do not set NSR_CLIENT, NMDA uses the name of the local host, not the virtual cluster host, for the operations during cluster backups and restores.

The NSR_CLIENT value (either the default value or an explicitly defined value) used for a backup must be the same as the NSR_CLIENT value used for the restore of the backup.

You must set the NSR_CLIENT parameter to the virtual cluster hostname in a cluster environment to enable the restore of backup data, no matter which cluster node was used during the backup.

Note: For parallel backups or restores, you must set NSR_CLIENT to the same value for all sessions.

NSR_PHYSICAL_HOST_LICENSE setting for cluster systems

By default, NMDA is licensed per virtual host and you require a license for each virtual cluster hostname for the backup. You cannot change the default behavior when you back up to NetWorker server 7.6 or earlier.

When you back up to NetWorker server 7.6 SP1, you must set the parameter NSR_PHYSICAL_HOST_LICENSE=TRUE to enforce that NMDA is licensed per cluster node, not per virtual cluster hostname.

"NSR_PHYSICAL_HOST_LICENSE" on page 364 describes the parameter.

Converting client-side to server-side configurations in a cluster

To convert a client-side backup configuration that uses an NMDA configuration file to the wizard (server-side) configuration in a cluster, follow the instructions for "Converting client-side to server-side configurations" on page 79. You can run the **nsrdaadmin** conversion utility from any of the cluster nodes configured for backup.

During the conversion, if **nsrdaadmin** prompts for the names of users to add to the Lockbox resource, specify the names of operating system users on each cluster node used for the backup.

Note: If you do not grant the Lockbox ACL rights during the conversion, you can add users later by using NMC to edit the Lockbox resource for the virtual cluster host.

The *EMC NetWorker Module for Databases and Applications Release Notes* describes the conversion limitations.

Configuring scheduled backups in a cluster with the wizard

You can configure a scheduled backup in a cluster either with or without the wizard. You can use the configuration wizard to create a server-side configuration.

Ensure that you meet the cluster configuration requirements as described in "Cluster configuration requirements" on page 248.

Follow the procedure for "Configuring scheduled backups with the wizard" on page 82 to create the configuration with the wizard. Keep in mind the following cluster requirements.

Note: To enable configuration with the wizard, configure the NetWorker client software to be highly available. The *EMC NetWorker Cluster Installation Guide* provides details.

1. On the **Specify the Client Name** page, type the name of the virtual host for the cluster.
2. On the page where you set the advanced configuration parameters, ensure the following settings:
   - The NSR_CLIENT parameter is set to the virtual cluster hostname. "NSR_CLIENT setting for cluster systems" on page 248 provides details.
   - The NSR_PHYSICAL_HOST_LICENSE parameter is set to specify the correct NMDA licensing, if required. "NSR_PHYSICAL_HOST_LICENSE setting for cluster systems" on page 249 provides details.
3. On the **Select the NetWorker Client Properties** page, add *database_username*@*cluster_node* to the **Remote Access** field for each cluster node that can perform backups or restores.

   Note: Do not use the following syntax in the **Remote Access** field: user=*database_username*,host=*cluster_node*

   On Windows clusters only, add **SYSTEM**@*cluster_node* to the **Remote Access** field for each cluster node that participates in backups and restores.

With these settings, the wizard performs the following actions:

- Creates a NetWorker Client resource for the virtual cluster host with the NMDA settings for scheduled backups and the following additional settings:
  - The Backup Command attribute contains **-c** *cluster_virtual_hostname*.
  - The Remote Access attribute contains the user entries as described in step 3.
- Creates a generic NetWorker Client resource for each cluster node, if it does not already exist.
- If the configuration contains sensitive information such as passwords, stores the information in the Lockbox resource. The wizard grants access (ACL rights) to the Lockbox resource to the users on the virtual cluster host and cluster node, described in step 3.

  Note: If additional users require ACL rights to the Lockbox, you can add them to the Lockbox resource directly in NMC.

Configuring scheduled backups in a cluster without the wizard

You can configure a scheduled backup in cluster either with or without the wizard. You can perform the configuration without the wizard by using settings in NMC to create a client-side configuration.

Ensure that you meet the cluster configuration requirements as described in "Cluster configuration requirements" on page 248.

Follow the procedure for "Configuring scheduled backups without the wizard" on page 85 to create the configuration for the cluster without the configuration wizard. Keep in mind the following cluster requirements:

1. Create a generic Client resource for each cluster node that will run backups and restores, if a Client resource does not yet exist for a given host.

   The following example includes the attribute settings in a generic Client resource:

   - Scheduled Backup: (clear)

     Note: Do not select the checkbox for Scheduled Backup, so the scheduled backup does not back up the generic client.

   - Backup Command: (blank)
   - Group: (blank)
   - Save Set: SKIP
   - Schedule: SkipAll

2. Create the NMDA configuration file in one of the following locations:

   - On the shared disk of the virtual host.
   - In the same location on the local disk of each physical node.

     A copy of the configuration file must exist in the same location on each physical node.

3. Create an NMDA-specific Client resource on the NetWorker server for each virtual cluster host that will run backups and restores, according to "Configuring the Client resource" on page 89:

   - The Save Set attribute contains the required value as described in Table 9 on page 89.
   - The Remote Access attribute contains the name of each cluster node that can store and retrieve the backups.

     The Remote Access attribute contains this value for each cluster node:

     user=*database_username*,**host=***cluster_node*

     On Windows clusters only:

     user=SYSTEM,host=*cluster_node*

     For Sybase scheduled backups on UNIX only:

     user=*root_username*,**host=***cluster_node*

   - The Backup Command attribute contains the following value:

     nsrdasv -z *configuration_file_path*

     where *configuration_file_path* is the pathname of the NMDA configuration file created in step 2.

For example, if physical hosts *clus_phys1* and *clus_phys2* form a Windows cluster that contains a DB2 database, the configuration of the Client resource for the virtual cluster host contains the following attribute settings:

Backup Command: **nsrdasv -z** *pathname*/**nmda_db2.cfg**

Group: db2group

Remote Access: user=*dbinst*,host=*physical_node_1*

user=*dbinst*,host=*physical_node_2*

user=SYSTEM,host=*clus_phys1*

user=SYSTEM,host=*clus_phys2*

Save Set: DB2:/SAMPLE/NODE0001

4. On the NMDA client, set the NSR_CLIENT parameter to the virtual cluster hostname in the NMDA configuration file (for other than Oracle backups) or in the RMAN script (for Oracle backups). You specify the configuration file or the RMAN script in the attributes of the virtual Client resource. "NSR_CLIENT setting for cluster systems" on page 248 provides details.
5. On the NMDA client, set the NSR_PHYSICAL_HOST_LICENSE parameter to specify the correct NMDA licensing, if required, in the NMDA configuration file (for other than Oracle backups) or in the RMAN script (for Oracle backups). You specify the configuration file or the RMAN script in the attributes of the virtual Client resource. "NSR_PHYSICAL_HOST_LICENSE setting for cluster systems" on page 249 provides details.

Configuring manual backups in a cluster

To configure a manual NMDA backup in a cluster, follow the same instructions as used to configure a scheduled backup without the wizard. "Configuring scheduled backups in a cluster without the wizard" on page 250 provides details.

The only difference is that you can configure a generic NetWorker Client resource for the virtual cluster host instead of the NMDA-specific Client resource.

In addition, ensure that you set the NSR_SERVER=*NetWorker_server_hostname* parameter when you back up to a remote NetWorker server.

Configuring Avamar deduplication backups in a cluster

Use the following procedures to configure an NMDA deduplication backup with Avamar in a cluster:

- "Configuring scheduled backups in a cluster with the wizard" on page 249 or "Configuring scheduled backups in a cluster without the wizard" on page 250 or "Configuring manual backups in a cluster" on page 252
- "Configuring scheduled Avamar deduplication backups" on page 98 or "Configuring manual Avamar deduplication backups" on page 99 for the virtual cluster host

Ensure that there is at least one NetWorker Client resource enabled for deduplication backups:

- For a scheduled deduplication backup, enable deduplication for the Client resource that is part of the NetWorker group for the backup.

  Note: The Client resource can be for a virtual cluster host (recommended) or physical node.

- For both scheduled and manual deduplication backups, enable deduplication for the Client resource of the host that is set in the NSR_CLIENT parameter in the NMDA wizard or in the configuration file.

Configuring Data Domain deduplication backups in a cluster

Use the following procedures to configure a Data Domain deduplication backup in a cluster:

- "Configuring scheduled backups in a cluster with the wizard" on page 249 or "Configuring scheduled backups in a cluster without the wizard" on page 250 or "Configuring manual backups in a cluster" on page 252
- "Configuring DD Boost backups to use Client Direct" on page 101 or "Configuring DD Boost backups to use a storage node" on page 103 for the virtual cluster host

Configuring probe-based backups in a cluster

Use the following procedures to configure a probe-based NMDA backup in a cluster:

- Configure a scheduled backup in the cluster according to "Configuring scheduled backups in a cluster with the wizard" on page 249 or "Configuring scheduled backups in a cluster without the wizard" on page 250.
- Configure the probe-based backup according to "Configuring probe-based backups" on page 105. Associate the Probe resource with the Client resource for the virtual cluster host.

Setting up nodes to back up to a local storage node

Typically, the NetWorker server backs up the data in a cluster to the first storage node listed in the Storage Nodes attribute of the Client resource. The host is set in the NSR_CLIENT parameter, usually the virtual cluster host. You can configure a virtual cluster client to direct the backups to the storage node located on the cluster node on which the backup runs (local storage node).

You can use the following steps to set up the cluster nodes to back up to a local storage node.

1. Install the NetWorker storage node on each node used for the NMDA backup.
2. Create a Storage Node resource on the NetWorker server for each storage node.
3. Create a NetWorker Device resource for the device on each node used for the backup. Label and mount a NetWorker volume for each device.

4. Ensure that the Groups and the selection criteria (such as Clients) of the media pool used for the devices match the settings in the NMDA backup configuration.
5. In the NetWorker Client resource (configured for the virtual cluster client according to "Configuring scheduled backups in a cluster with the wizard" on page 249 or "Configuring scheduled backups in a cluster without the wizard" on page 250), set the Storage Nodes attribute in NMC to the following values in this order:

   ```
   curphyhost
   nsrserverhost
   ```

Example 29 on page 270 describes how to set up three Oracle RAC nodes as storage nodes for NMDA backups.

Configuring recoveries in a cluster

For recovery in a cluster, ensure that the required configurations are in place:

- In the Client resource of the virtual cluster host, you have set the Remote Access attribute as described in "Performing NMDA data restore and recovery" on page 169.
- You have set the NSR_CLIENT parameter to the same hostname as set in NSR_CLIENT during the backup, usually the virtual cluster hostname. "NSR_CLIENT setting for cluster systems" on page 248 provides details.

To run a restore and recovery, use the procedures described in Chapter 4, "Data Restore and Recovery."

Active-active application clusters (overview)

An active-active application cluster is a group of linked virtual hosts or physical hosts with shared storage called cluster nodes, which can access the database data from multiple nodes concurrently.

Note: DB2 DPF systems do not adhere to the preceding definition of an active-active application cluster because each DB2 partition has its own set of computing resources, such as CPU and storage. However, for the purposes of this guide and NMDA licensing, DB2 DPF is an active-active application cluster.

NMDA supports the following types of active-active application cluster software:

- DB2 DPF
- DB2 pureScale
- Informix MACH
- Oracle RAC
- Sybase ASE Cluster Edition

Note: Lotus Domino documentation might use the term "active-active cluster" for a Lotus partitioned Domino server configured in a cluster, similar to the one shown in Figure 11 on page 247. However, the Lotus partitioned Domino server is not an active-active application cluster. The NMDA documentation refers to it as an active-passive cluster.

This topic provides information that is common to all the supported types of active-active application clusters.

The *EMC NetWorker Software Compatibility Guide* describes all the supported cluster environments.

The following topics describe specific types of active-active application clusters:

- "DB2 DPF systems" on page 256
- "DB2 pureScale systems" on page 260
- "Informix MACH systems" on page 264
- "Oracle RAC systems" on page 267
- "Sybase ASE Cluster Edition systems" on page 272

Active-active cluster configuration requirements

Before you configure a backup or recovery in an active-active application cluster system, ensure that you meet the following requirements on the nodes that will participate in the backup or recovery:

- You have set up the cluster according to the appropriate database server documentation.
- You have installed NMDA and the NetWorker software on each node of the cluster that will participate in a backup or restore. The *EMC NetWorker Module for Databases and Applications Installation Guide* provides details.
- To enable NMDA operations on UNIX systems, you have ensured that the /nsr/apps and /nsr/apps/tmp directories have the drwxrwxrwt access permissions.

Select one node from the active-active cluster as the node that will have the backup data stored under its client name. To ensure that each node in the backups stores the data under the same client name, set the NSR_CLIENT parameter to the same value for each node. This setting also facilitates recovery because you do not need to remember which node backed up what data.

Configure backups and restores in the active-active cluster according to the corresponding instructions for an active-passive cluster:

- "Configuring scheduled backups in a cluster without the wizard" on page 250

  Note: You cannot use the NMDA wizard to configure a scheduled backup in an active-active cluster.

- "Configuring manual backups in a cluster" on page 252
- "Configuring Avamar deduplication backups in a cluster" on page 252
- "Configuring Data Domain deduplication backups in a cluster" on page 253
- "Configuring probe-based backups in a cluster" on page 253
- "Configuring recoveries in a cluster" on page 254

Ensure that you configure the required nodes in the active-active application cluster:

- Configure the node that stores the indexes in the same way as the virtual cluster client in an active-passive cluster.
- Configure the other active-active nodes in the same way as the cluster nodes in an active-passive cluster.

Restarting backups after a system failure

If an active-active application cluster supports a virtual IP and you configured the cluster with a virtual IP, then you will have the same backup failover capabilities as an active-passive cluster, as described in "Backup failover" on page 247.

Otherwise, an active-active application cluster does not have the same backup failover capabilities as an active-passive cluster. If a system failure occurs on the node used to start an NMDA scheduled backup, the backup might fail. In this case, you must use manual intervention to configure and restart the backup on a different node that is available:

1. On the available node, ensure that you have installed the following software:
   - NetWorker client
   - NetWorker storage node (optional)
   - NMDA
2. Configure a new NMDA-specific Client resource for the available node.
3. Replace the original Client resource with the new Client resource from step 2 for the NMDA backup.

DB2 DPF systems

The DB2 Database Partitioning Feature (DPF) offers an environment where a single database divided by logical nodes can reside on multiple partitions, either on the same host or on multiple hosts. A partitioned database can manage high volumes of data and provides benefits such as increased performance and high availability.

The following figures display examples of the supported DPF configurations:

- Figure 12 on page 256—Database nodes reside in partitions on a single physical host
- Figure 13 on page 257—Multiple DPF nodes reside in partitions on separate hosts

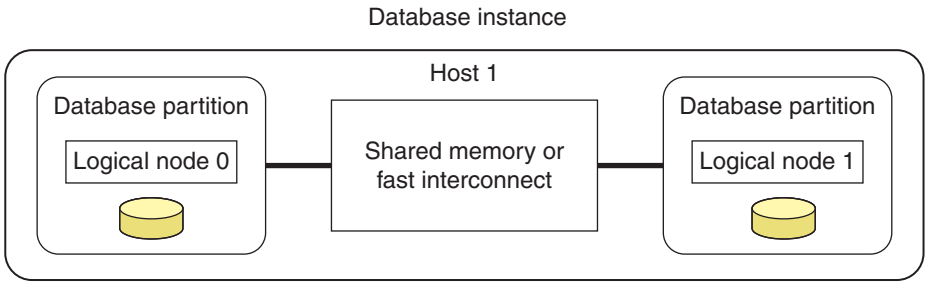

Figure 12 Single DPF host with shared memory

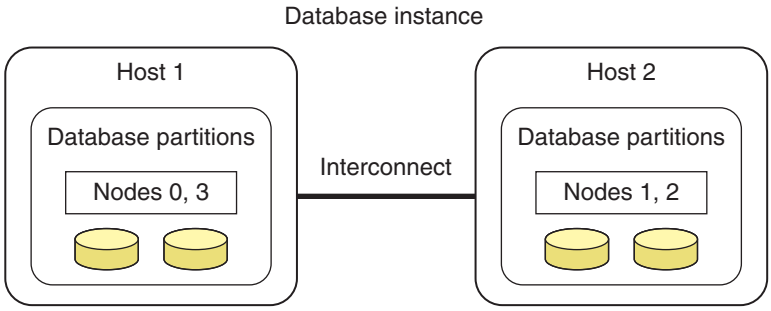

Figure 13 Two DPF hosts with multiple nodes

Configuring backups and restores in a DPF environment

To configure database backups with NMDA in a DB2 DPF system, follow the instructions in "Active-active cluster configuration requirements" on page 255.

The following additional requirements apply to DB2 DPF systems:

- For all DB2 DPF backups, the total number of target sessions for the backup pool must be equal to or greater than the number of database partitions.
- To configure a manual DPF backup for the supported DB2 versions, follow the instructions in "Configuring manual DPF backups" on page 258.
- To configure a scheduled DPF backup, follow the instructions in "Configuring scheduled DPF backups" on page 259.
- To restore and recover a DPF backup, follow the instructions in "Restoring and recovering DPF backups" on page 260.
- For a probe-based backup, configure a DPF backup as usual. Complete the probe-based backup configuration according to "Configuring probe-based backups" on page 105. Associate the Probe resource with the Client resource configured for the scheduled backup.

Note: Ensure that you meet the following special requirements in a DPF environment:
- The LOG_THRESHOLD parameter setting in the Probe resource must specify the total number of logs generated by all the DPF nodes since the last probe-based backup.
- The scheduled backup must run on the catalog partition.

Configuring manual DPF backups

DPF backups with NMDA must use an NMDA configuration file, which resides on the DB2 server host.

You can use the following steps to configure a manual backup of a DPF database that resides on two hosts, host1 and host2.

1. Ensure that you have configured the DPF database partitions according to the appropriate DB2 documentation.
2. Create an NMDA configuration file (for example, nmda_db2.cfg) in the database instance directory, so that it is accessible to both hosts. You must set the following parameters:
   - NSR_SERVER—Hostname of the NetWorker server.
   - NSR_CLIENT—Name of the host with the database catalog node partition (NODE0000).

   Note: Setting NSR_CLIENT enables NetWorker to store backups from all the nodes under the same index as described in "Active-active cluster configuration requirements" on page 255.
3. Use the NMC program (**Configuration** › **Clients**) to create a basic NetWorker Client resource for both host1 and host2.

   In each Client resource, set the Remote Access attribute to the following value:

   user=*database_username*,**host=***other_node_hostname*

Performing manual DPF backups

Supported DB2 versions enable a manual DPF backup of either all the partitions or specific partitions.

To perform a manual DPF backup of all the partitions, run the following command:

```
db2 backup db sample on all dbpartitionnums load /usr/lib/libnsrdb2.xx
   options @pathname/nmda_db2.cfg
```

where:

- *sample* is the name of the database to be backed up.
- *xx* is the extension for the operating system.
- *pathname*/nmda_db2.cfg is the full pathname of the NMDA configuration file.

To perform a manual DPF backup of specific partitions, run the following command:

```
db2 backup db sample on dbpartitionnums (1,3...) load
   /usr/lib/libnsrdb2.xx options @pathname/nmda_db2.cfg
```

where:

- 1, 3... are numeric values of each partition to back up, for example, 1 and 3.
- Other options are the same as in the preceding command.

Configuring scheduled DPF backups

The DB2_PARTITION_LIST parameter enables the simultaneous backup of multiple partitions.

A scheduled DPF backup must use an NMDA configuration file, for example, nmda_db2.cfg, which resides on the host with the DPF catalog node partition.

For all the supported DPF configurations shown in Figure 12 on page 256 and Figure 13 on page 257, configure only one NetWorker Client resource, one Group resource, and one NMDA configuration file.

You can use the following steps to configure a scheduled backup of a DB2 DPF database that resides on two host DB2 servers, for example, host1 and host2.

1. Ensure that you have configured the DPF database partitions according to the appropriate DB2 documentation.
2. Create an NMDA configuration file, for example, nmda_db2.cfg, in the database instance directory. The configuration file must contain the following settings:
   - NSR_SERVER—Hostname of the NetWorker server.
   - NSR_CLIENT—Name of the client with the database catalog node partition (NODE0000).
   - INSTHOME—Pathname of the DB2 instance home directory.
   - DB2_NODE_NAME—Name of the DB2 instance.
   - DB2_USER—Name of the DB2 user.
   - DB2_OPTIONS—Either DB2BACKUP_ONLINE or DB2BACKUP_OFFLINE.
   - DB2_PARTITION_LIST is set as follows:
     – To back up all the DPF partitions, specify the value all:

       DB2_PARTITION_LIST=all

     – To back up specific DPF partitions only, specify the partitions in a comma-separated list, for example:

       DB2_PARTITION_LIST=0,1,3
3. Encrypt a DB2 user password with the **nsrdaadmin -P** command. "USER_PSWD" on page 370 provides details.
4. Create only one Group resource. Use the NMC program to create a Group resource, for example, dpf_group1.
5. Use the NMC program to create a NetWorker Client resource for the database host (host1) that has the DPF catalog node partition (NODE0000). For example, the attribute settings are as follows:

   Name: host1
   Save Set: DB2:/*sample*/NODE0000
   Group: dpf_group1
   Backup Command: nsrdasv -z *pathname*/nmda_db2.cfg
   Remote Access: user=*database_username*, host=host2

6. Use the NMC program to create a basic NetWorker Client resource for the other database host (host2) if one does not already exist. For example, the attribute settings are as follows:

   Name: host2
   Save Set: DB2: SKIP
   Group: (blank)
   Backup Command: (blank)

The scheduled backup backs up the partitions specified in the DB2_PARTITION_LIST parameter.

Restoring and recovering DPF backups

The **db2 recover** command combines the functions of the **db2 restore** command and **db2 rollforward** command.

To restore and recover a database on all the database partitions, you can run the **db2 recover** command as shown in this example:

```
db2 recover db sample to end of logs on all DBPARTITIONNUMS
```

where *sample* is the name of the database for restore.

To restore and recover a database on specific partitions only, you can also run the **db2 recover** command as shown in this example:

```
db2 recover db sample to end of logs on DBPARTITIONNUMS (2, 4)
```

where:

- *sample* is the name of the database for restore.
- 2 and 4 are the database partition numbers.

You can run the **db2 restore** command and **db2 rollforward** command separately in a DPF environment, in the same way that you run the commands in a non-DPF environment.

The DB2 documentation provides details about the different commands and options used for restore and recovery.

DB2 pureScale systems

A DB2 pureScale system is an active-active application cluster with a shared-disk architecture for high availability, scalability, and application transparency. The cluster includes a single database partition that is shared by the entire group of all the cluster member nodes.

Note: The term node as used in other active-active application clusters is referred to as a member in a DB2 pureScale system.

NMDA does not support automatic failover in a DB2 pureScale system. After a node failure, NMDA does not automatically restart a backup from the point of failure on another node. You must manually restart the backup on a different available node for a manual backup, or change the NetWorker group to the group of an available node.

DB2 pureScale does not support delta or incremental backups. As a result, NMDA does not support the DB2BACKUP_DELTA or DB2BACKUP_INCREMENTAL setting with the DB2_OPTIONS parameter for a pureScale backup.

NOTICE

To prepare for disaster recovery, ensure that you back up all the files that are specific to the DB2 pureScale system.

Configuring backups and restores in a DB2 pureScale system

Before you configure a backup or recovery in a DB2 pureScale system, ensure that you meet the requirements in "Active-active cluster configuration requirements" on page 255.

You must select one DB2 pureScale member as the node that will have the backup data stored under its client name. To ensure that each member node performing a backup stores the data under the same client name, you must set the NSR_CLIENT parameter to the same value for each member node. This setting facilitates recovery because you do not need to remember which member node backed up what data.

The following additional requirements apply to DB2 pureScale systems:

- You have set up the DB2 pureScale cluster system according to the appropriate IBM documentation.
- You have installed NMDA on the required member hosts. The *EMC NetWorker Module for Databases and Applications Installation Guide* provides installation information.

  If you will configure automatic backups of the transaction logs to a NetWorker backup device, then NMDA must be installed on all the member hosts.

  If you will not configure the automatic log backups, then NMDA can be installed on a subset of the member hosts for backups and restores. Install the software on at least two member hosts so that, if one host fails during a backup, you can complete the backup on the other host.
- If you make topology changes in a pureScale system by adding or removing any members, then ensure that you perform any configuration steps on each member.
- To configure a manual backup, follow the instructions in "Configuring manual backups in a DB2 pureScale system" on page 262.
- To configure a scheduled backup, follow the instructions in "Configuring scheduled backups in a DB2 pureScale system" on page 262.
- To configure a transaction log backup, follow the instructions in "Configuring automatic backups of transaction logs in a DB2 pureScale system" on page 263.

- For a probe-based backup, configure a scheduled backup as usual and follow the additional instructions in "Configuring probe-based backups" on page 105. Associate the Probe resource with the Client resource configured for the scheduled backup.

  Note: Ensure that the LOG_THRESHOLD parameter setting in the Probe resource specifies the total number of transaction logs generated by all the active member nodes since the last probe-based backup.

- To restore and recover a pureScale backup, follow the instructions in "Restoring and recovering backups in a DB2 pureScale system" on page 263.

Configuring manual backups in a DB2 pureScale system

A manual backup in a DB2 pureScale system must use an NMDA configuration file. Create the configuration file with the appropriate parameter settings as described in "Configuring manual backups in a cluster" on page 252.

Ensure that all the members use the same NMDA configuration file by storing the configuration file in a shared folder, which is a file system folder that is accessible to all the member hosts.

Performing manual backups in a DB2 pureScale system

A user on any active member host in a DB2 pureScale system can run a single backup command to back up the database for the entire pureScale system.

To perform a manual backup, run the **db2 backup** command as described in "Performing DB2 manual backups with the db2 backup command" on page 148. For example:

1. Log in to an active DB2 pureScale member host as the DB2 operating system user.
2. Run the **db2 backup** command with the appropriate options. In the options, specify the NMDA configuration file that you created in the preceding topic. For example:

   - On UNIX:

     ```
     db2 backup db sample load /usr/lib/libnsrdb2.xx options
     @pathname/nmda_db2.cfg
     ```

   - On Windows:

     ```
     db2 backup db sample load
     NetWorker_install_dir\nsr\bin\libnsrdb2.dll options
     @pathname\nmda_db2.cfg
     ```

Configuring scheduled backups in a DB2 pureScale system

A scheduled backup in a DB2 pureScale system must use an NMDA configuration file. Follow the instructions in "Configuring scheduled backups in a cluster without the wizard" on page 250 to create the configuration file with the appropriate parameter settings and create the NMDA-specific Client resource for the scheduled backup.

You cannot use the NMDA wizard to configure a scheduled backup in a DB2 pureScale system.

The Backup Command attribute in the Client resource must include the configuration file name, for example:

```
nsrdasv -z /db2sd_data/nmda_cfg/nmda_db2s.cfg
```

In this example, /db2sd_data is a shared folder that is accessible to all the members hosts in the DB2 pureScale system.

The NetWorker server starts a scheduled backup automatically, based on the NetWorker backup schedule. You can also manually start a scheduled backup from the NMC console.

Configuring automatic backups of transaction logs in a DB2 pureScale system

You can configure the automatic backup of transaction logs to a specified shared directory or a NetWorker backup device or both. For successful log backups to a NetWorker device, ensure that a device is always available for the backups.

Configure the automatic backup of the transaction logs whenever required by DB2, for example, when the logs become full, by following the instructions in "Configuring automatic backups of DB2 transaction logs" on page 116.

Restoring and recovering backups in a DB2 pureScale system

In a DB2 pureScale system, you can log in to any active pureScale member and run the **db2 restore** and **db2 rollforward** commands to restore and roll forward a backed-up database as described in "Performing DB2 data restore and recovery" on page 171.

For example, you can run the following command to perform a point-in-time restore of a database named *sample* from a backup that used *n* number of sessions:

```
db2 restore db sample load /usr/lib/libnsrdb2.xx open n sessions
   options @pathname/nmda_db2.cfg taken at yyyymmddtttt
```

You also have an option to prefetch the transaction logs from NMDA backups to local storage and then roll forward the logs from the local copy. "Performing DB2 recovery with fetched logs and the db2 rollforward command" on page 176 describes how to run the **nsrdb2rlog** command to retrieve a copy of the DB2 transaction logs from the NetWorker server to a local file system.

In the DB2 pureScale system, each member has its own transaction logs. NMDA backs up each member's logs under the specific member node name, NODE*nnnn*. For example, the **nsrinfo** command displays the following archived log backups for member nodes 0 and 1:

```
/SAMPLE/NODE0000/DB2LOG/:/C0000020_S0000112.LOG, date=1368466701
   5/13/2013 1:38:21 PM
/SAMPLE/NODE0001/DB2LOG/:/C0000019_S0000054.LOG, date=1368466492
   5/13/2013 1:34:52 PM
```

You can use the **-N** option with the **nsrdb2rlog** command to specify the member node for which to retrieve the logs. To fetch the logs backed up for multiple members, you must run the **nsrdb2rlog** command separately for each member.

In the pureScale system, the same NSR_CLIENT setting is used for each node during backups so that all the logs are backed up under the same client name. To specify that client name during the log restore, you can specify either the **-c** *client_name* option or the **-z** *configuration_file* option with the **nsrdb2rlog** command. If you specify the **-z** option, the the restore uses the client name from NSR_CLIENT in the configuration file.

The following example shows the **nsrdb2rlog** commands to retrieve the logs of two different member nodes, node 0 and node 1, from the NMDA backups under the client name, bu-purescale-3:

```
nsrdb2rlog -s serverA -a sample -c bu-purescale-3 -N 0 -C 20 -S 101
   -E 112 -d /db2sd_20130612185334/db2inst1/sqllib_shared/overflow
   -I db2inst1
nsrdb2rlog -s serverA -a sample -c bu-purescale-3 -N 1 -C 19 -S 50
   -E 54 -d /db2sd_20130612185334/db2inst1/sqllib_shared/overflow
   -I db2inst1
```

The **db2 recover** command combines the functions of the **db2 restore** and **db2 rollforward** commands. "Performing DB2 restore and recovery with the db2 recover command" on page 177 describes how to set the DB2 database configuration VENDOROPT parameter and use the **db2 recover** command for recovery.

The DB2 documentation also provides details about the different commands and options used for restore and recovery.

Informix MACH systems

Multi-node Active Clusters for High Availability or High Availability Clusters (MACH) is the Informix Dynamic Server (IDS) high-availability system, which provides increased failover capabilities, flexibility, and scalability in an Informix environment.

A MACH cluster includes the following database servers:

- The primary server, which receives updates.
- One or more secondary servers, each of which is a mirror image of the primary server and is in perpetual recovery mode that applies logical-log records from the primary server.

You can configure one of the secondary servers with the Informix Connection Manager to take over the primary server role if the primary server becomes unavailable.

NOTICE

NMDA operations do not support failover from primary to secondary MACH servers. "Restarting backups after a system failure" on page 256 provides details.

MACH supports the following features:

- Continuous Log Restore (CLR)—CLR provides a method to create a hot backup of a primary Informix server. To accomplish the hot backup, this method places the secondary server in rollforward mode and constantly applies the logical logs from the primary server through special **ontape** or **onbar** recovery modes.

  Note: If the primary server fails, the secondary server could be out-of-date by at least one log.

- High Availability Data Replication (HDR)—HDR enables the replication of a complete database instance including logs from the primary to a hot standby secondary instance (the HDR instance).

When the primary server is up and running, the HDR secondary server runs in read-only mode and typically offloads reporting and complex SQL queries from the primary server. If a catastrophic failure occurs, the HDR secondary can be automatically promoted to primary server and can take over the work of the failed primary server.

- Remote Standalone Secondary Servers (RSS)—RSS servers are similar to HDR servers. With RSS, the primary server does not wait for responses from the secondary server before committing database transactions.

  In a common RSS scenario, remote backup servers run in locations geographically distant from the primary server because RSS is not sensitive to network latency. RSS instances run in read-only mode and offload the database reporting. If a primary instance fails, the standby instance cannot directly take over the work of the primary instance. The standby instance must be first promoted from an HDR secondary instance, then it can be promoted to a primary instance. The IBM Informix documentation provides details.

- Shared Disk Secondary Servers (SDS)—SDS servers do not maintain a copy of the primary server, but share disks with the primary server, typically through some form of disk mirroring configuration.

  Similar to the other types of secondary servers, an SDS server runs in a read-only mode and typically offloads the database processing. An SDS server can be directly promoted to a primary server, unlike an RSS server. Because an SDS server is so close to the primary server (shares the same disk), it is often the best type of server to initially fail over to if the primary server encounters a problem.

Using Continuous Log Restore for backup and restore

You can access CLR functionality with NMDA through manual **onbar** backups.

You must meet the following requirements before you can use CLR with NMDA:

- CLR requires manual intervention to perform log restores. You must run the **onbar –r –l –C** command manually on the secondary server to restore the logs.
- You do not need to restore logs individually to the secondary server. You can restore the logs in bulk by running the **onbar –r –l** command.
- CLR with NMDA does not support the probe-based backups or other types of scheduled backups through either client-side configurations or server-side configurations.

You can use the following steps to implement CLR with NMDA.

1. Configure IDS to work with NMDA according to the instructions in the *EMC NetWorker Module for Databases and Applications Installation Guide* for both the primary server and secondary server.
2. On the primary system, perform a level 0 backup with the **onbar -b -L 0** command. Before you start the backup, ensure that you have set the required NMDA Informix parameters in the environment according to "Configuring manual backups" on page 94.

3. On the secondary system, perform an imported restore, but only perform a physical restore with the **onbar -r -p** command. After the physical restore completes on the secondary system, the database server waits in fast recovery mode to restore the logical logs.
4. On the primary system, back up the logical logs with the **onbar -b -l** command.
5. On the secondary system, set NSR_CLIENT to the hostname of the primary server, and restore the logical logs with the **onbar -r -l -C** command.
6. Repeat step 4 and step 5 for all the logical logs available for backup and restore. You can use the Informix **alarmprogram** script, configured through the Informix ALARMPROGRAM parameter, to automate the log backups.
7. On the secondary system, run the following commands to finish restoring the logical logs and quiesce the server:
   - If logical logs are available to restore, run the **onbar -r -l** command.
   - After you have restored all the available logical logs, run the **onbar -r -l -X** command only in case of a failover.

Backups of primary and HDR or RSS or SDS servers

IDS does not support the backup of any secondary servers of a MACH cluster. On the primary server of a MACH cluster, NMDA supports the following features:

- Manual **onbar** backups through the **onbar –b** command
- Manual log backups through the **onbar –b –l** command
- Scheduled backups through a client-side configuration
- Backup configuration and backup changes through the wizard
- Configuration of probe-based backups

Note: Although you can use the Informix wizard to configure the backup of a secondary server, the resulting scheduled backup of the secondary server would fail. The Informix BAR_ACT log would include an error code of 151 and would contain the following error message:

```
DR: This command is not valid on a secondary server.
```

Restores of primary and HDR or RSS or SDS servers

You can use an **onbar** restore with MACH in the following two cases:

- New secondary server setup
- Restore of a cluster

Set up the restore of an HDR or RSS secondary server the same way as an Informix imported restore. "Performing an Informix imported restore" on page 229 and IDS documentation provide details.

Set up the restore of a primary server by using the same method as for a typical NMDA restore. Perform additional steps to update the secondary servers during the restore as described in the IDS documentation.

Oracle RAC systems

An Oracle Real Application Cluster (RAC) is an active-active application cluster environment for parallelism and high availability. A node in an Oracle RAC system is a physical or virtual host with a hostname such as host1.emc.com. An Oracle instance is a memory structure and a group of Oracle Server processes running on a node.

An Oracle database (for example, named databs1) comprises datafiles used by the Oracle instances and shared between the nodes. All instances share the same datafiles and control file. Each instance must have its own set of redo log files and its own archived redo logs.

A RAC system enables multiple Oracle instances across multiples nodes to access the same Oracle database simultaneously. Oracle RAC is a cluster software infrastructure that provides concurrent access to the same storage and the same set of datafiles from all nodes in the cluster. All the database files reside on shared disks.

Configuring backups and restores in a RAC environment

To configure database backups with NMDA in an Oracle RAC system, follow the instructions in "Active-active cluster configuration requirements" on page 255.

The following additional considerations apply to Oracle RAC systems:

- Configure the Oracle Net services to enable connect-time (SQL Net) failover. "Connect-time failover" on page 268 provides guidelines. The Oracle documentation provides more details.
- To set up a local storage node for each RAC node involved in a backup, follow the instructions in "Setting up nodes to back up to a local storage node" on page 253.
- Ensure that the configuration is set up for parallel backups and restores. "Parallel RAC backups and restores" on page 269 provides details.
- Create the appropriate RMAN scripts for the preferred types of Oracle backups and restores on the RAC system. "Creating RMAN backup and restore scripts" on page 269 provides details.
- Review the additional issues with Oracle recovery operations in "Archived redo logs" on page 271.
- For a probe-based backup, configure a RAC backup as usual. Complete the probe-based backup configuration according to "Configuring probe-based backups" on page 105. Associate the Probe resource with the Client resource configured for the scheduled backup.

Note: In a RAC environment:
- The LOG_THRESHOLD parameter setting in the Probe resource must specify the total number of logs generated by all the RAC nodes in a cluster since the last probe-based backup.
- Select one node as the backup node, and use that node to perform all the probe-based backups in the RAC. If this node is lost, set up a new backup node to use the same Probe resource as the old backup node used.

Connect-time failover

Neither Oracle RMAN nor NMDA supports Transparent Application Failover (TAF). As a result, if a failure occurs during an Oracle backup, NMDA does not automatically restart the backup from the point of failure on another node.

Note: NMDA supports only connect-time failover.

When multiple listeners support a single service, a connect-time failover reroutes the connection request to another listener if the first listener is either down or cannot connect. To enable the connect-time failover in RAC, there must be a listener on each node, and each instance must use the same Net service name. Ensure that GLOBAL_DBNAME does not appear in the SID_LIST_LISTENER parameter of the listener.ora file since the setting disables the failover.

When you use the local Net service naming method, the client's tnsnames.ora file must include the following parameters:

```
o11pA.emc.com =
   (DESCRIPTION =
      (ADDRESS_LIST =
         (FAILOVER = ON)
         (ADDRESS = (PROTOCOL = tcp) (HOST = nodeA)
         (PORT = 1521))
         (ADDRESS = (PROTOCOL = tcp) (HOST = nodeB)
         (PORT = 1521))
      )
      (CONNECT DATA =
         (SERVICE_NAME = proddb)
      )
   )
```

- Set the FAILOVER parameter to ON. The default value is ON for an ADDRESS_LIST and OFF when you do not specify ADDRESS_LIST.
- Include the ADDRESS_LIST parameter:
  - If you specify multiple addresses but omit the ADDRESS_LIST parameter, the Oracle Net service reads the addresses sequentially and attempts to connect to the last one only.
  - If you specify the ADDRESS_LIST parameter, the service tries the addresses in the order they appear in the list.
- In the CONNECT_DATA section, use SERVICE_NAME instead of the system identifier (SID). SERVICE_NAME must be different from SID.

When a node or listener to which a client tries to connect is not available, the client contacts the next listener on the list. When the instance is down but the listener is running, the failover occurs only if you have configured the instance to dynamically register with the listener.

Dynamic instance registration

During dynamic instance registration, the database registers itself with the Oracle listener on startup and unregisters itself on shutdown.

You must meet the following requirements for dynamic instance registration:

- You must set the INSTANCE_NAME and SERVICE_NAME parameters in the initialization file (init*oracle_sid*.ora). There can be several services for a single instance.
- If the listener does not listen on the default port (1521), you must set the LOCAL_LISTENER parameter in the initialization file.
- The SID_LIST_LISTENER parameter in listener.ora must not include SID_DESC for the RAC instances.
- It is not necessary to have the listener.ora file when the listener listens on the default port.

When the instance is down, the listener does not know how to connect to it. As a result, the listener tries the next connect option specified in the ADDRESS_LIST in the tnsnames.ora file.

Note: Some applications such as Oracle Enterprise Manager still require static database registration with a listener.

Static instance registration

With static registration, you manually configure the information about the instance in the listener.ora file through SID_DES in the SID_LIST_LISTENER parameter. The listener contains continuous information about the instance even if the instance is down. As a result, when the instance is down, the listener still tries to connect to it rather than try the next connect option specified in the tnsnames.ora file.

Parallel RAC backups and restores

NMDA enables Oracle backups on either a single node or several nodes of the Oracle RAC system.

A parallel Oracle backup uses Oracle instances that run in parallel on multiple RAC nodes. In the RMAN backup script created for running a parallel Oracle backup, allocate multiple channels for the backup and specify that each channel run on a specific node. You must set the NSR_CLIENT parameter to the same value for each channel. “NSR_CLIENT setting for cluster systems” on page 248 describes the parameter.

NMDA enables restores of the Oracle data to any physical node in the RAC, regardless of which RAC node originally performed the backup.

Creating RMAN backup and restore scripts

You can use a single RMAN backup script to run a parallel Oracle backup with NMDA on a RAC system. In the RMAN script, allocate multiple channels for the backup and specify that each channel run on a specific node.

Note: NMDA does not support multiple RMAN restores that run simultaneously.

To run an Oracle restore on a RAC system, none of the nodes can be open. You need to mount only the node that is running the RMAN restore script.

Example 29 on page 270 describes a sample Oracle RAC configuration, including the required RMAN backup and restore scripts.

Example 29 Setting up an Oracle RAC configuration

A RAC system contains three nodes (hosts) named nodeA, nodeB, and nodeC. Each node has a Linux operating system and an attached tape drive for use during NMDA backups. Each node has the NetWorker storage node software.

In the NMC interface, you create a Storage Node resource for each node by right-clicking **Storage Nodes** in the **Devices** pane and selecting **New**.

After you create the Storage Node resources, you create a Device resource for each tape drive. You have attached the tape devices to storage nodes, so the device names must have the format rd=*hostname*:*device_name*. For example:

- You have attached tape device /dev/rmt/tape0 to node A. In the Device resource, the device name is rd=nodeA:/dev/rmt/tape0.
- You have attached tape device /dev/rmt/tape3 to node B. In the Device resource, the device name is rd=nodeB:/dev/rmt/tape3.
- You have attached tape device /dev/rmt/tape1 to node C. In the Device resource, the device name is rd=nodeC:/dev/rmt/tape1.

In the tape device on each node, you label and mount a volume. You assign all the volumes to the Default pool in this example.

You select node A to store the index entries for the NMDA backups and start the backups. The choice of node A is arbitrary. You could select node B or node C instead. In all the RMAN backup scripts and restore scripts, you must set NSR_CLIENT to the node A hostname.

In the NetWorker Client resource for node A:

- The Remote Access attribute is set to the hostnames of node B and node C:

  ```
  user=database_username, host=nodeB
  user=database_username, host=nodeC
  ```

- The Storage Nodes attribute is set to:

  ```
  curphyhost
  nsrserverhost
  ```

- You have set the remaining attributes as required. For example:
  - The Backup Command attribute is set to the command used for the backup:

    nsrdasv(**.exe**) **-z** *configuration_file_path*

    where *configuration_file_path* is the complete pathname of the configuration file that contains the NMDA parameter settings for the backup.
  - The Group attribute is set to the backup group name.
  - The Save Set attribute is set to the RMAN script pathname.

  Table 9 on page 89 describes the Client resource attributes.

The following RMAN script uses all three nodes to perform the backup. Each node backs up data to its local tape drive. The NetWorker client file index of node A stores the backup information from all three nodes:

```
connect target sys/oracle@connect_identifier;
run {
   allocate channel t1 type 'SBT_TAPE'
   connect 'sys/oracle@Net_service_name_of_instance_A';
   allocate channel t2 type 'SBT_TAPE'
   connect 'sys/oracle@Net_service_name_of_instance_B';
   allocate channel t3 type 'SBT_TAPE'
   connect 'sys/oracle@Net_service_name_of_instance_C';

   send 'NSR_ENV=(NSR_CLIENT=nodeA)';
   backup database;
   release channel t1;
   release channel t2;
   release channel t3;
}
```

To enable restores, you must set NSR_CLIENT to node A hostname. The NetWorker server then obtains the backup information from the client file index of node A. To restore the backup data to a different node such as node B, add the node B hostname to the Remote Access attribute in the NetWorker Client resource of node A.

For example, the following RMAN script restores the database. You can run the script on any host:

```
connect target sys/oracle@connect_identifier;
run {
   allocate channel t1 type 'SBT_TAPE';
   allocate channel t2 type 'SBT_TAPE';

   send 'NSR_ENV=(NSR_SERVER=NetWorker_server, NSR_CLIENT=nodeA)';
   restore database;

   release channel t1;
   release channel t2;
}
```

Archived redo logs

Each node in a RAC system maintains a separate set of redo logs. Redo logs that become full are archived on the local node. As a result, the archived redo logs are divided among the nodes of the system.

To enable RMAN to back up and recover a RAC system, make all the archived redo log files accessible by all nodes in the backup or recovery. The Oracle RAC documentation describes how to share the archived redo logs.

The following topics provide sample scripts to back up and restore all the archived redo log files in a RAC system:

- "Backing up all archived logs from each node" on page 272
- "Restoring all archived logs from each node" on page 272

Note: The archived logs must use the same NSR_CLIENT setting used to back up the database.

Backing up all archived logs from each node

You can back up all the archived log files in a RAC system from a single node, such as a node named ops1.emc.com, by using the following type of RMAN script:

```
run {
   allocate channel t1 type 'SBT_TAPE'
   connect 'user_name/user_passwd@connect_string_of_ops1';
   allocate channel t2 type 'SBT_TAPE'
   connect 'user_name/user_passwd@connect_string_of_ops2';

   send 'NSR_ENV=(NSR_CLIENT=ops1.emc.com)';
   backup filesperset 10
   (archivelog all delete input format 'al_%s_%p');

   release channel t1;
   release channel t2;
}
```

Restoring all archived logs from each node

You can restore all the archived log files in a RAC system from a single node, such as a node named ops1.emc.com, by using the following type of RMAN script:

```
run {
   allocate channel t1 type 'SBT_TAPE'
   connect 'user_name/user_passwd@connect_string_of_ops1';
   allocate channel t2 type 'SBT_TAPE'
   connect 'user_name/user_passwd@connect_string_of_ops2';

   send 'NSR_ENV=(NSR_SERVER=mars.emc.com, NSR_CLIENT=ops1.emc.com)';
   restore (archive log all);

   release t1;
   release t2;
}
```

Sybase ASE Cluster Edition systems

An ASE Cluster Edition system is an active-active application cluster that consists of multiple ASE servers. The servers are located across multiple nodes that connect to a shared set of disks and run as a shared-disk cluster, including multiple physical hosts.

The following conditions exist in an ASE Cluster Edition system:

- Each machine is a node.
- Each ASE server is an instance.
- Each node contains one instance.
- All the nodes must be on a storage area network (SAN).

The connected instances that form the cluster together manage databases that reside on shared disks. All the databases are accessible from each instance. The multiple instances that make up the cluster appear to clients as a single system.

Clients connect logically to the shared-disk cluster while remaining connected physically to individual instances. If one instance fails in the cluster, clients connected to that instance can fail over to any of the remaining active instances.

The database devices in the cluster (except for private devices used by local user temporary databases) must be raw devices, not block devices. The quorum device that includes configuration information for the cluster and is shared by all the instances must be on a raw partition that is accessible to all the nodes of the cluster.

The ASE Cluster Edition documentation describes setting up the required devices and configuring the nodes and instances of an ASE Cluster Edition system.

NMDA supports two types of Sybase backup server configurations:

- Single backup server configuration
- Multiple backup server configuration

In a single backup server configuration, the ASE cluster contains only one backup server at any time. The cluster administration can move the backup server from one node to another node, but only one node at a time contains a running backup server.

In a multiple backup server configuration, the ASE cluster contains two or more backup servers, running simultaneously on separate nodes. Any of the backup servers can perform the backup or restore of the databases.

A Sybase ASE Cluster Edition system supports two types of multiple backup server configurations:

- Dedicated—The cluster assigns a specific backup server to each instance in the cluster.
- Round-robin—The cluster does assign a specific backup server to an instance. The cluster assigns the least busy backup server from a group according to availability.

NOTICE

NMDA supports only the dedicated type of multiple backup server configuration, not the round-robin type.

The ASE Cluster Edition documentation describes Sybase backup servers in an ASE cluster.

Configuring backups and restores for ASE Cluster Edition

You must run the **nsrdasv** and **nsrsybrc** commands on the same node as the Sybase backup server. You can run the **nsrsybcc** command on any node in the cluster.

Select one node in the cluster to use for all backups. If the backup node goes down, you can switch to a different node in the cluster for the backups. "Restarting backups after a system failure" on page 256 describes how to restart a backup on a different node after a failure. However, from that point on, use the new node as the backup node. Do not constantly switch the backup servers.

NOTICE

NMDA does not support failover on ASE cluster databases.

You must select one node to store all the backups for the whole ASE Cluster Edition. Specify the hostname of the node with the NSR_CLIENT parameter as described in "NSR_CLIENT setting for cluster systems" on page 248. The client file index of the node

records all the backups of the ASE cluster. Use the same NSR_CLIENT value even if the original node that performed the backup failed and you configured a new node to run the backups.

Ensure that you meet the configuration requirements in "Active-active cluster configuration requirements" on page 255.

For a probe-based backup, configure a scheduled backup for ASE Cluster Edition as usual. Complete the probe-based backup configuration according to "Configuring probe-based backups" on page 105. Associate the Probe resource with the Client resource configured for the scheduled backup.

Skipping the backup of temporary databases

An ASE Cluster Edition system supports the following types of temporary databases:

- Local system
- Local user
- Global user
- Global system

NMDA skips the backup of all these types of temporary databases in an ASE Cluster Edition system. NMDA also skips the backup of tempdb and user-created temp databases on unclustered Sybase installations.

Restoring an ASE Cluster Edition backup

To restore backups, use the **nsrsybrc -c** option to specify the NetWorker client hostname of the node used to store the NetWorker indexes for the backup. The value for the **-c** option must be same as the NSR_CLIENT parameter used during a backup.

Use the **nsrsybrc -d** option to specify a redirected restore. For example:

```
nsrsybrc -U user_ID -P password -s NetWorker_server_name
   -c index_node_name -d
   SYBASE:/backup_server_node_Sybase_instance_name/database_name
   SYBASE:/index_node_Sybase_instance_name/database_name
```

Note: In the **nsrsybrc** command, the instance names and database names are case-sensitive and must be in the same case as recorded in the backup entries in the NetWorker indexes.

Performing backup and recovery of the quorum device

NMDA does not provide backups of the quorum device.

You must back up the quorum device by using the Sybase **qrmutil** utility. The **qrmutil** utility creates a backup of the quorum device as a file on the file system. The Sybase ASE Cluster Edition documentation describes the utility.

You can then use the NetWorker **save** program to back up the quorum backup file.

Note: It is the responsibility of the backup administrator to create the quorum device backup file and to protect that backup file with the NetWorker **save** program.

To recover the quorum backup file, you must use the NetWorker **recover** program. The **recover** program restores the quorum backup file as a file on the file system. You must then load the quorum backup file into the database by using the appropriate Sybase ASE Cluster Edition utility.

Recovering the master database on clusters

The instructions for recovering the master database on Sybase clusters differ from the instructions for recovering the master database on Sybase single node installations. Contact Sybase Technical Support to obtain the instructions.

Disregard Sybase instructions that recommend the use of the **load** command for recovering database backups. You must run the **nsrsybrc** command to recover the Sybase master database, not the **load** command. Run the **nsrsybrc** command for the recovery after starting the instance in single-user mode, also known as master-recover mode.

CHAPTER 7
Multiple Installations on a Single Host

This chapter includes the following topics:

- Multiple database or application installations 278
- Multiple databases or applications with the same bitness 278
- Multiple databases or applications with 32-bit and 64-bit coexistence 278

Multiple database or application installations

NMDA supports multiple database installations or application installations on the NMDA host, including the same or different types of databases or applications with the same or different bitness. For example:

- 32-bit Lotus Domino and 32-bit Sybase ASE servers on 32-bit or 64-bit AIX system
- 32-bit Lotus Domino and 64-bit DB2 servers on a 64-bit Sun SPARC system
- 32-bit Informix IDS and 64-bit IDS servers on a 64-bit HP PA-RISC system
- 32-bit Informix and 64-bit Oracle servers on a 64-bit Windows system

The documentation for the database or application that you are using describes the versions of 32-bit and 64-bit database or application software that can coexist on the same system.

Multiple databases or applications with the same bitness

You can use the following steps to operate NMDA in an environment that contains multiple databases or multiple applications with the same bitness on a single NMDA host.

1. Install the NMDA software according to the instructions in the *EMC NetWorker Module for Databases and Applications Installation Guide*.
2. Configure and perform a backup or recovery according to the instructions in this guide for each database or application. Ensure that you meet the following requirements:
   - You set the required backup or restore parameters to the correct values for the specific database server or application server only.

     Appendix A, "NMDA Parameters and Configuration File," describes the NMDA parameters.
   - The user that runs the backup or recovery is the same user that operates the database server or application server that is backed up or recovered.

Multiple databases or applications with 32-bit and 64-bit coexistence

The *EMC NetWorker Module for Databases and Applications Installation Guide* describes the NMDA installation for 32-bit and 64-bit database or application coexistence on 64-bit systems.

Note: In an environment where 32-bit applications and 64-bit applications coexist, you cannot use the wizard to configure a 32-bit application. You can only use the wizard to configure a 64-bit application.

On a 64-bit system, you might install NMDA to enable the coexistence of 32-bit and 64-bit NMDA databases or applications. You must use the correct command or correct library name in specific commands and specific attributes as shown in the following table.

Table 22 Commands or library names for 32-bit and 64-bit coexistence (page 1 of 2)

Name to use for 64-bit database or 64-bit application	Name to use for 32-bit database or 32-bit application	Location where you use the command or library name	Topic with details about the command or library name
For all databases and applications:			
nsrdaprobe(.exe)	**nsrdaprobe32**(.exe)	In Probe Command attribute of the Probe resource, required for a probe-based backup	"Configuring probe-based backups" on page 105
nsrdasv(.exe)	**nsrdasv32**(.exe)	In Backup Command attribute of the Client resource, required for a scheduled backup	"Configuring the Client resource" on page 89
For DB2 on 64-bit Windows only:			
libnsrdb2.dll	**libnsrdb232.dll**	In **db2 backup** command	• "Performing DB2 manual backups with the db2 backup command" on page 148 • "Performing manual DPF backups" on page 258
		In **db2 restore** command	• "Performing DB2 data restores to the same instance" on page 173 • "Performing DB2 data restores to a different instance" on page 174
		In **db2 update db cfg** command	"Configuring automatic backups of DB2 transaction logs" on page 116
		In DB2_VENDOR_LIB_PATH parameter setting for scheduled backups	"DB2_VENDOR_LIB_PATH" on page 369
For Lotus only:			
nsrdasv(.exe)	**nsrdasv32**(.exe)	In Lotus manual backup command	"Performing Lotus manual backups with the nsrdasv command" on page 150
nsrdocrc(.exe)	**nsrdocrc32**(.exe)	In Lotus document-level recovery command	"Performing Lotus document-level recovery with nsrdocrc command" on page 192
nsrnotesrc(.exe)	**nsrnotesrc32**(.exe)	In Lotus database recovery command	"Performing Lotus database recovery with the nsrnotesrc command" on page 182
		In **nsrlotus_remrecov**(.bat) script for Lotus directed recovery	"Performing Lotus directed recovery with the GUI" on page 189
For MySQL only:			
nsrdasv	**nsrdasv32**	In MySQL manual backup command	"Performing MySQL manual backups with the nsrdasv command" on page 153
nsrmysqlrc	**nsrmysqlrc32**	In MySQL restore and recovery commands	"Performing MySQL data restore and recovery" on page 196

Table 22 Commands or library names for 32-bit and 64-bit coexistence (page 2 of 2)

Name to use for 64-bit database or 64-bit application	Name to use for 32-bit database or 32-bit application	Location where you use the command or library name	Topic with details about the command or library name
For Sybase only:			
nsrdasv(**.exe**)	**nsrdasv32**(**.exe**)	In Sybase manual backup command	"Performing Sybase manual backups with the nsrdasv command" on page 155
nsrsybcc(**.exe**)	**nsrsybcc32**(**.exe**)	In Sybase command for database consistency check	"Performing Sybase database consistency checks before backups" on page 154
nsrsybrc(**.exe**)	**nsrsybrc32**(**.exe**)	In Sybase restore command	"Performing Sybase data restores with the nsrsybrc command" on page 216

CHAPTER 8
Oracle DBA and NetWorker Backup Administrator Collaboration

This chapter includes the following topics:

- DBA disk backup issues in Oracle environments 282
- NMDA solution for DBA disk backup issues 282
- Configuring backups of DBA disk backups 284
- Reporting on backups of DBA disk backups 288
- Purging DBA disk backups 290
- Restore and recovery of DBA disk backups 291

DBA disk backup issues in Oracle environments

In many Oracle environments, the Oracle DBA performs RMAN native disk backups to a Fast Recovery Area (FRA) on a file system or ASM. The DBA cannot keep those disk backups for the long term due to limited space in the FRA. The DBA also cannot use the native disk backups for disaster recovery. As a result, backup administrators, such as NetWorker administrators, are frequently asked to perform the following tasks:

- Move the disk backup pieces to backup devices for long-term storage.
- Perform optional cloning to near-line or offline devices to prepare for disaster recovery.

This type of Oracle environment enables the separation of the DBA and backup administrator roles and empowers the DBA to protect their own data. However, this environment includes the following major issues:

- During an Oracle data recovery, the DBA performs the following tasks before starting the restore and recovery:

  1. Manually determines which backup piece files the Oracle software needs for the recovery of a corrupted or lost whole database or database objects, for example, a tablespace named HR.
  2. Asks the backup administrator to restore the corresponding backup piece files from the NetWorker server to the original file system directory.

  As a result, the Oracle data recovery is an error-prone two-step process.
- After the backup administrator backs up the disk backups to the NetWorker server, the DBA must manually purge the disk backups to reclaim space on the disk for the next disk backup, archived logs, and so on.
- The backup administrator does not know what is included in the Oracle backup piece files because the files are in a database proprietary format, and whether a specific database has been protected. As a result, the backup administrator cannot report on these backups and provide Service Level Agreements (SLAs).
- The backup administrator must determine the correct time to start the backup, based on when the DBA disk backup job completes every day. This task might become tedious for the backup administrator because it depends on any variance in the DBA disk backup completion time. For example, if the DBA backups to disk do not complete on time or start running at 5 a.m. instead of 2 a.m. every night, the disk backups are not copied to the backup device on time.

NMDA solution for DBA disk backup issues

NMDA solves all the issues in the DBA disk backup environment and empowers both the Oracle DBA and NetWorker backup administrator to perform their specific role tasks by enabling the following functionality:

- The DBA continues to perform RMAN disk backups to the FRA, without requiring any NMDA or NetWorker knowledge.
- The DBA performs a one-step recovery, without knowing whether the backups are on the FRA or on a NetWorker device.

- The backup administrator moves the disk backups to the NetWorker server and catalogs the backups, without requiring any Oracle knowledge.
- The backup administrator reports on what was backed up, without requiring any Oracle knowledge.
- The Oracle software automatically purges the DBA disk backups on the FRA when the disk backups are moved to a backup device. The DBA does not need to manually purge disk backups on the FRA to reclaim disk space.
- The backup schedules of the DBA and backup administrator are automatically synchronized. The synchronized schedules ensure that the NetWorker software starts the copy of Oracle disk backups to a backup device as soon as Oracle disk backups are completed successfully.
- All archived redo logs are automatically backed up, even if the logs are not in the FRA, to ensure the recovery of the data backed up by the DBA.

The NMDA solution environment includes the following workflows:

- The DBA continues to perform RMAN disk backups to a FRA on a file system or ASM, without requiring any knowledge of NMDA or NetWorker software. The RMAN disk backups do not require the use of any third-party backup software.
- The backup administrator uses the new simplified NMDA workflow to move DBA disk backups to a NetWorker backup device and catalog these backups in the NetWorker indexes as Oracle backups, without having any Oracle knowledge.
- The backup administrator consults the DBA to obtain specific Oracle information, such as the Oracle SID and OS username for NMDA to connect to the database. The backup administrator enters the information in the backup configuration wizard to configure the NMDA backups of the DBA disk backups to NetWorker backup devices. "Configuring backups of DBA disk backups" on page 284 provides details.
- NMDA automatically discovers the DBA disk backup information by querying the RMAN catalog. NMDA internally uses RMAN SBT to copy the disk backup images to the NetWorker server, including full or incremental backup sets, image copies, control file auto backups, SPFILE backups, and archived redo logs.
- If the archived transaction log is unavailable or corrupted, NMDA looks outside the FRA directory to find a good copy to back up.
- NMDA ensures that the backup copies are registered as "Oracle" backup copies to both the NetWorker online indexes and Oracle backup catalog (RMAN catalog).
- NMDA saves metadata information at the end of each backup, which includes details about the database files and archived redo log files included in the NMDA backup.

If the DBA wants to use NMDA directly to send Oracle backups on the FRA to NetWorker backup devices, then the DBA can use the regular NMDA workflow and include the **backup recovery area** command in the RMAN script as described in "Other Oracle features" on page 56.

Note: NMDA scheduled backups of Oracle disk backups do not support the following regular NMDA Oracle features:
- Backup copies
- Oracle RAC
- Oracle Data Guard
- Oracle Exadata
- Restartable backups
- Snapshot backups and restores

Configuring backups of DBA disk backups

Note: This topic and the remaining topics in this guide apply only to NetWorker administrators, not to Oracle DBAs.

You must complete the required settings in the NMC backup configuration wizard for NMDA to configure the backup of DBA disk backups. The wizard field help describes each field on the wizard pages. You cannot use any other method to configure this type of backup.

"Configuring scheduled backups with the wizard" on page 82 provides basic tips on using the configuration wizard.

Before you run the configuration wizard, you must ensure that the required NMDA and NetWorker software is installed on the client host according to the instructions in the *EMC NetWorker Module for Databases and Applications Release 1.6 Installation Guide*.

1. Start the backup configuration wizard:
   a. In the NMC **Enterprise** window, select the NetWorker server name and double-click the **NetWorker** application to launch the application.
   b. In the **Configuration** window under the NetWorker server name, click **Clients**.
   c. Start the wizard by using the appropriate method:
      – To create a backup configuration, right-click **Clients** and select **Client Backup Configuration** › **New**.
      – To modify a backup configuration created with the wizard, right-click the NMDA client in the right pane and select **Client Backup Configuration** › **Modify**.
2. Select Oracle from the list of available applications.
3. On the **Select the Configuration Type** page, select **Scheduled backup of Oracle disk backups**.
4. On the **Specify the Oracle Information** page, complete the field settings. Consult the Oracle DBA to obtain the required information for the fields:
   - **Oracle installation directory (ORACLE_HOME)**
   - **Oracle tnsnames.ora directory (TNS_ADMIN)**
   - **Oracle locale (NLS_LANG)**

- If the DBA uses OS authentication on UNIX or Linux to connect to the Oracle database, select **Use operating system authentication** and complete the following fields:

  – **OS username**
  – **Oracle instance SID**

  Complete the RMAN catalog fields only if a Recovery Catalog is used:

  – **Use RMAN catalog** (select the checkbox)
  – **Username**
  – **Password**
  – **Net Service name**

- If the DBA uses database authentication to connect to the Oracle database, select **Use RMAN connection file** and complete the **Connection file** field.

  An RMAN connection file contains "connect target" and "connect catalog" strings used to connect to the production database and Recovery Catalog, for example:

```
connect target sys/oracle@proddb;
connect catalog rman/rman@oracat;
```

  Note: Only the Oracle DBA should have read, modify, and execute permissions to the RMAN connection file, described in "The connection file" on page 328.

5. On the **Specify the Performance and Additional Backup Options** page, complete the field settings:

   - Specify a parallelism value for **Number of backup sessions (parallelism)**.

     The parallelism value determines the number of streams that NMDA will use in parallel to perform the backup of the data. NMDA writes the streams of backup data in parallel to one or more storage devices.

     Note: The backup parallelism must not be greater than the total number of device sessions available for the backup through the media pool configuration. Otherwise, the backup will become suspended because it is waiting for an available device to be mounted for the proper pool, or the backup will fail if not enough devices are available.

   - Select any required options under **NetWorker Data Options**.
   - Specify the values of any required NMDA parameters under **Advanced Backup Options** as described in "Advanced NMDA parameters" on page 286.
   - Select **Delete expired backups from RMAN catalog** to have NMDA automatically delete expired backups from the RMAN catalog.

     The Delete expired backups option enables you to delete any expired backup in the NetWorker index from the Oracle backup catalog, called the RMAN catalog. This option ensures that the Oracle catalog is up-to-date. Consult the DBA about whether to select this option.

6. On the **Specify the Preprocessing, Postprocessing, and Advanced Environment Options (Optional)** page, complete the field settings according to the field help:
   - Specify a preprocessing script pathname for **Preprocessing script** if you want a preprocessing script to run before the scheduled backup.
   - Specify a postprocessing script pathname for **Postprocessing script** if you want a postprocessing script to run after the scheduled backup.

     To prepare the Oracle server for disaster recovery, you should set up a postprocessing script to back up files that the Oracle DBA cannot back up with RMAN, as described in "Setting up a postcommand script for backup of Oracle files" on page 238. The Oracle DBA must provide values used in the postprocessing script, such as the ORACLE_HOME and ORACLE_SID settings.
   - Specify the values of any required NMDA parameters under **Advanced Environment Options** as described in "Advanced NMDA parameters" on page 286.

The NMDA backup configuration wizard validates the settings and ensures that the following conditions are true to enable the scheduled backup of Oracle disk backups:

- The Oracle database is mounted.
- The Oracle disk backups are located on an enabled FRA.

NOTICE

Ensure that the Oracle DBA does not have the Operate NetWorker privilege (with NetWorker 7.*x*) or the Change Application Settings privilege (with NetWorker 8.0 or later). With these privileges, the DBA could accidentally delete some NMDA backup entries from the NetWorker index by running Oracle commands that delete obsolete Oracle disk backups.

Note: When the DBA starts protecting a new database by using Oracle disk backups, the DBA must notify the backup administrator about the new database to include in the NMDA backups.

Advanced NMDA parameters

The following table describes the NMDA parameters that you can set in the configuration wizard during the backup configuration procedure in "Configuring backups of DBA disk backups" on page 284.

Table 23 Advanced NMDA parameters for backups of DBA disk backups (page 1 of 2)

Parameter	Description	Default and valid values
NSR_CLIENT	Specifies the NetWorker client name to use for a backup. Recommended for a backup in a cluster. The wizard automatically sets this parameter to the virtual hostname for the backup configuration of a cluster virtual host.	• Hostname of the physical host on which the backup session runs (default). • Valid NetWorker client hostname.
NSR_DATA_VOLUME_POOL	Specifies the name of the NetWorker volume pool to use for a backup. Optional for a backup.	• Most appropriate pool as selected by the NetWorker server (default). • Valid name of a NetWorker volume pool.

Table 23 Advanced NMDA parameters for backups of DBA disk backups (page 2 of 2)

Parameter	Description	Default and valid values
NSR_DATA_VOLUME_POOL‹*n*»	Specifies an additional NetWorker pool if you want to send different parallel streams to different volumes. Optional for a backup.	• Undefined (default). • Valid NetWorker pool name that is different from the name used by the parameter NSR_DATA_VOLUME_POOL or a different NSR_DATA_VOLUME_POOL‹*n*» parameter.
NSR_MAX_START_RETRIES	Specifies how many times NMDA tries to connect to the NetWorker server before the operation fails. NMDA waits for 30 seconds between each try to connect. For example, the connection to the NetWorker server might fail for one of these reasons: • The NetWorker server is not ready because the devices are not mounted. • The **nsrindexd** service of the NetWorker server is busy due to other client sessions. Optional for a backup.	• 4 (default). • Integer number of tries to connect to the NetWorker server.
NSR_NO_MULTIPLEX	Specifies whether to disable multiplexing on the NetWorker device. A TRUE setting can improve the Oracle restore and recovery performance for tape devices. Optional for a backup.	• FALSE (default) = Enable multiplexing on the device. • TRUE = Disable multiplexing on the device. **Note:** Do not set this parameter to TRUE for nontape devices, such as AFTD or DD Boost devices.
NSR_SERVER_NIC	Specifies the name of a network interface card (NIC) on a NetWorker server. Optional for a backup.	• Undefined (default). • Valid name of a NetWorker server NIC.

Configuring optional probe-based backups of DBA disk backups

You can optionally configure probe-based backups that detect any new Oracle disk backups and trigger the corresponding NetWorker backups by using NMDA.

The probe-based backups enable the automatic synchronization of the Oracle disk backup schedule to the FRA and the NetWorker backup schedule.

To configure a probe-based backup, you must complete the required NetWorker Probe, Group, and Client resource configurations.

1. Create a separate NetWorker Probe resource and set the Probe resource attributes as described in the following table.

   NOTICE

   Create a separate NetWorker Probe resource for each database in the backup.

Table 24 NetWorker Probe resource attributes (page 1 of 2)

Attribute	Description
Name	Specify a name to identify the Probe resource. Each Probe resource must have a unique name.
Probe Command	Specify the name of the NMDA probe program, **nsrdaprobe**, which checks (“probes”) for the condition that triggers a probe-based backup.

Table 24 NetWorker Probe resource attributes (page 2 of 2)

Attribute	Description	
Command Options	Specify a comma-separated list of parameters with their settings.	
	Parameter	**Description**
	DBA_DISK_BACKUP	**Set to TRUE.** With DBA_DISK_BACKUP=TRUE, if the probe finds any new Oracle disk backups and no Oracle disk backup is currently running for the database, the probe triggers an NMDA scheduled backup of the new disk backups.
	NSR_DEBUG_LEVEL	Optional. Specify the level of debug information generated by the probe and written to the debug log. "NSR_DEBUG_LEVEL" on page 360 describes debug levels.
	NSR_DIAGNOSTIC_DEST	Optional. Specify the directory location of the NMDA debug logs, including the debug logs generated by the probe. "NSR_DIAGNOSTIC_DEST" on page 362 provides details.

2. Use NMC to configure the NetWorker Group resource for a probe-enabled backup group. Set the probe-specific attributes in the Group resource as described in the *EMC NetWorker Administration Guide* (the information about creating and scheduling a probe group).

   When you enable probing through the Group resource attributes, each probe runs once when the probe interval is reached, which is the time window defined by the Probe Start Time and Probe End Time attributes. A probe-based backup ignores the group's Start Time attribute.

   Note: If you start a probe-enabled backup group manually, probing occurs immediately (only once, not repeatedly at intervals) and the backup starts only if the probe conditions are met.

3. Use NMC to edit the NetWorker Client resource:
   - Set the Probe attribute to the name of the Probe resource from step 1. You can associate a Client resource with only one probe.
   - Set the Group attribute to the probe-enabled group from step 2.

     Note: A probe-based backup group must include at least one probe-enabled client, which is a Client resource associated with a probe.

Reporting on backups of DBA disk backups

Note: This section and the remaining sections in this guide apply only to NetWorker administrators, not to Oracle DBAs.

For reporting purposes, you can use the NMDA utility **nsrorainfo** to query what is included in the backup of DBA disk backups, such as the database name and the database file names.

At the end of a scheduled backup of DBA disk backups, NMDA creates a metadata save set that contains the names and attributes of the database files and archived redo logs included in the backup. The **nsorainfo** command uses this metadata, which facilitates the creation of reports about the NMDA scheduled backups of Oracle disk backups.

For example, the backup administrator that configures NMDA scheduled backups of DBA disk backups can run the **nsrorainfo** command to generate a daily or weekly report about the backups. The backup administrator can provide the report to the Oracle DBA, with details about the database files and archived redo logs included in the backups.

The backup administrator can run the **nsrorainfo** command directly on the backup client host. Alternatively, the backup administrator can use tools to run the command remotely from a different host.

You can use the following **nsrorainfo** command syntax to generate a report about all the database files and archived redo logs in an NMDA backup of Oracle disk backups:

```
nsrorainfo[.exe] [-c NetWorker_client_name] [-s NetWorker_server_name]
   -L [-D database_name] [-t time]
```

where:

- *NetWorker_client_name* is the hostname of the NetWorker client whose index contains backup information. By default, the client is the local host.
- *NetWorker_server_name* is the hostname of the NetWorker server to query for the backup information. By default, the server is the local host.
- *database_name* is the name of the Oracle database. If specified, the command reports on the NMDA backup of Oracle disk backups of the database. If not specified, the command reports on the latest backup of any database.
- *time* is a time in **nsr_getdate(3)** format. If specified, the command reports on the latest backup before or at the specified time. If not specified, the command reports on the most recent backup.

Command options in brackets ([ ]) are optional. Do not include the brackets when you type the command.

For example, the following **nsrorainfo** command generates a report about the latest NMDA backup of Oracle disk backups of the database TARTST:

```
nsrorainfo -c bu-galaxy -s bu-galaxy -L -D TARTST

Database name: TARTST
Backup time: Mon Apr  1 14:28:21 2013
List of objects included in the same backup session:
List of database files
   Level:FULL,     SCN:701262,
   Name:/bigspace/oradata/SAP/sapdata1/btabd_1/btabd.data1,
   Tablespace:PSAPBTABD
   Level:FULL,     SCN:701262,
   Name:/bigspace/oradata/SAP/sapdata1/temp_1/temp.data1,
   Tablespace:PSAPTEMP
   Level:FULL,     SCN:701262,
   Name:/bigspace/oradata/SAP/sapdata4/user1d_1/user1d.data1,
   Tablespace:PSAPUSER1D
   Level:FULL,     SCN:701262,
   Name:/bigspace/oradata/SAP/sapdata2/user1i_1/user1i.data1,
   Tablespace:PSAPUSER1I
```

```
   Level:FULL,    SCN:701262,    Name:/bigspace/oradata/SAP/rbs1.ora,
   Tablespace:RBS
   Level:FULL,     SCN:701262,
   Name:/bigspace/oradata/SAP/sysaux1.ora, Tablespace:SYSAUX
   Level:FULL,     SCN:701262,
   Name:/bigspace/oradata/SAP/sapdata1/system_1/system.data2,
   Tablespace:SYSTEM
   Level:FULL,    SCN:701262,    Name:/bigspace/oradata/SAP/Sys1.ora,
   Tablespace:SYSTEM
   Level:FULL,     SCN:701262,
   Name:/bigspace/oradata/SAP/tools1.ora,  Tablespace:TOOLS
   Level:FULL,     SCN:701262,
   Name:/bigspace/oradata/SAP/undotbs1.ora,        Tablespace:UNDOTBS1
   Level:FULL,    SCN:701262,    Name:/bigspace/oradata/SAP/usr1.ora,
   Tablespace:USR
   Control file is included.
List of archived redo logs
   Thread:1,       Low SCN:516533, Next SCN:527810
   Thread:1,       Low SCN:527810, Next SCN:544344
   Thread:1,       Low SCN:544344, Next SCN:552877
   Thread:1,       Low SCN:552877, Next SCN:566651
   Thread:1,       Low SCN:566651, Next SCN:575414
   Thread:1,       Low SCN:575414, Next SCN:615642
   Thread:1,       Low SCN:615642, Next SCN:635734
   Thread:1,       Low SCN:635734, Next SCN:656480
   Thread:1,       Low SCN:656480, Next SCN:676574
   Thread:1,       Low SCN:676574, Next SCN:676966
   Thread:1,       Low SCN:676966, Next SCN:676969
   Thread:1,       Low SCN:676969, Next SCN:676971
   Thread:1,       Low SCN:676971, Next SCN:677006
   Thread:1,       Low SCN:677006, Next SCN:701144
   Thread:1,       Low SCN:701144, Next SCN:701450

List of savesets generated by the same backup session:
   TARTST_03o60949_1_1
   TARTST_01o608p1_1_2
   TARTST_02o608pk_1_2
   TARTST_04o6096h_1_1
(SCN is Oracle System Change Number)
```

The *EMC NetWorker Module for Databases and Applications Command Reference Guide* provides details about the **nsrorainfo** command.

Purging DBA disk backups

Note: This section and the remaining sections in this guide apply only to NetWorker administrators, not to Oracle DBAs.

When you have configured NMDA scheduled backups of Oracle disk backups, the disk backups and archived redo logs on the FRA must be regularly purged after the scheduled backups. Regular purging ensures that the FRA contains enough space for new disk backups and archived redo logs.

With this NMDA feature, after NMDA backs up an Oracle disk backup, the disk backup becomes eligible for deletion. The Oracle software automatically deletes the disk backup from the FRA when space is required.

If the Oracle DBA uses an Oracle retention policy to manage the Oracle disk backups, the DBA runs the **report obsolete** and **delete obsolete** commands to report and delete obsolete disk backups. If the DBA uses automatic channels for backups, the DBA must run the following commands to ensure that the commands affect only the disk backups and not the corresponding NMDA backups:

```
report obsolete device type disk
delete obsolete device type disk
```

If the DBA does not include **device type disk** in the commands, the commands use both the RMAN catalog and the NetWorker index for the **report** and **delete** operations. As a result, if the DBA has the Operate NetWorker privilege or Change Application Settings privilege, the **delete** command might delete the corresponding NMDA backups from the NetWorker index.

NOTICE

Ensure that the Oracle DBA does not have the Operate NetWorker privilege (with NetWorker 7.*x*) or the Change Application Settings privilege (with NetWorker 8.0 or later).

Restore and recovery of DBA disk backups

The DBA typically performs the restore and recovery of Oracle data. When you use NMDA scheduled backups of Oracle DBA disk backups, the DBA can restore and recover the data with a single-step operation, no matter where the backup is located (FRA or NetWorker device). The DBA writes the RMAN recovery script as if the Oracle disk backup is still in the FRA.

If the DBA uses manual channel allocation, the DBA must allocate both the DISK and SBT types of channels as instructed in the RMAN user guide because some of the backups might already be moved to a NetWorker device. In all other cases, the Oracle software automatically allocates both types of channels.

The Oracle software automatically retrieves the backup from the copy on the NetWorker device if the backup is not available in the FRA.

Note: The Oracle RMAN software favors disk backups during a restore. RMAN retrieves a backup automatically from the FRA even if another copy of the backup exists in an NMDA backup.

As a backup administrator, you are typically not involved in the recovery of an Oracle database. NMDA automatically determines the parameter settings used during the backup (NetWorker server name, client name, and so on) and uses the correct recovery settings accordingly. The DBA does not need to set any NMDA or NetWorker specific parameters in the recovery script.

The DBA can perform the restore and recovery of the Oracle disk backups to either the original data host or to a remote host, as long as the NMDA and NetWorker software is installed on the destination host.

Complete the following steps to enable a recovery to a remote host:

1. Ensure that NMDA is installed and configured on the remote host.

2. Ensure that a NetWorker Client resource exists on the NetWorker server for the remote host.
3. Ensure that the Remote Access attribute in the Client resource of the original data host contains the following value:

   ```
   user=db_or_app_user,host=destination_host
   ```

If the recovery to a remote host uses the recovery catalog or the original control file, the recovery retrieves the NSR_CLIENT and NSR_SERVER settings from the automatic channel persistent setting. If the recovery cannot retrieve the settings or if the recovery script uses manual channels, the DBA receives the following type of error message:

- If the recovery locates the NetWorker server:

  ```
  ORA-19511: Error received from media manager layer, error text:
  Could not locate the backup piece 'backup_piece_name' on the
  NetWorker server 'server_name'. Notify the backup administrator.
  ```

- If the recovery does not locate the NetWorker server:

  ```
  ORA-19511: Error received from media manager layer, error text:
  Could not locate the NetWorker server. The server name must be
  provided. Notify the backup administrator.
  ```

If this happens, create a text file named nmda_oracle_defaults in the directory /nsr/apps/res (UNIX or Linux) or *NetWorker_install_dir*\apps\res (Windows) and include the following parameter settings in the text file:

```
NSR_CLIENT=NetWorker_client_hostname
NSR_SERVER=NetWorker_server_hostname
```

CHAPTER 9
Snapshot Backups and Restores

This chapter includes the following topics:

- Snapshot operations with NetWorker Snapshot Management 294
- Configuring NSM snapshot backups 299
- Configuring NSM snapshot restore and recovery 303
- Setting the NSM parameters 304
- DB2 considerations for NSM snapshot operations 305
- Oracle considerations for NSM snapshot operations 308
- Replication Manager snapshot operations with Oracle ASM 338

Snapshot operations with NetWorker Snapshot Management

NetWorker Snapshot Management (NSM) backups and restores provide continuous snapshot-based protection and continuous availability of data on supported types of primary storage.

Note: With NetWorker release 8.1 or later, the NSM feature integrates and replaces the PowerSnap Modules. NSM is available as part of the NetWorker client software.

NMDA supports NSM snapshot backups and restores of DB2 and Oracle data on the primary storage, for example, EMC VNX Block (CLARiiON) or EMC VMAX (Symmetrix). The *EMC NetWorker Software Compatibility Guide* describes the supported software and the primary storage platforms.

NSM supports scheduled backups only, not manual (client-initiated) backups. The NSM documentation provides details.

The following topics describe the software requirements and the supported types of NSM snapshot backups and restores:

- "Software requirements for NSM snapshot operations" on page 294
- "Types of NSM backups" on page 295
- "Types of NSM restores" on page 296
- "NSM backup processes" on page 297

The following topics describe how to configure and perform NSM snapshot operations:

- "Configuring NSM snapshot backups" on page 299
- "Configuring NSM snapshot restore and recovery" on page 303
- "DB2 considerations for NSM snapshot operations" on page 305
- "Oracle considerations for NSM snapshot operations" on page 308

Software requirements for NSM snapshot operations

The following table lists the software requirements for a typical network environment that uses NMDA for NSM snapshot operations.

Table 25 Typical configuration for NSM snapshot operations with NMDA

Computer or device	Required software or configuration
Database server host	DB2 or Oracle server, NetWorker client, NMDA
Proxy client host (mount host)	NetWorker client, NetWorker storage node
Storage array	Storage array that NMDA supports with NSM, for example, VMAX or VNX Block
NetWorker server host	NetWorker server, NetWorker storage node, NetWorker client

The following figure shows the software components used in the NMDA environment for NSM snapshot backups and restores.

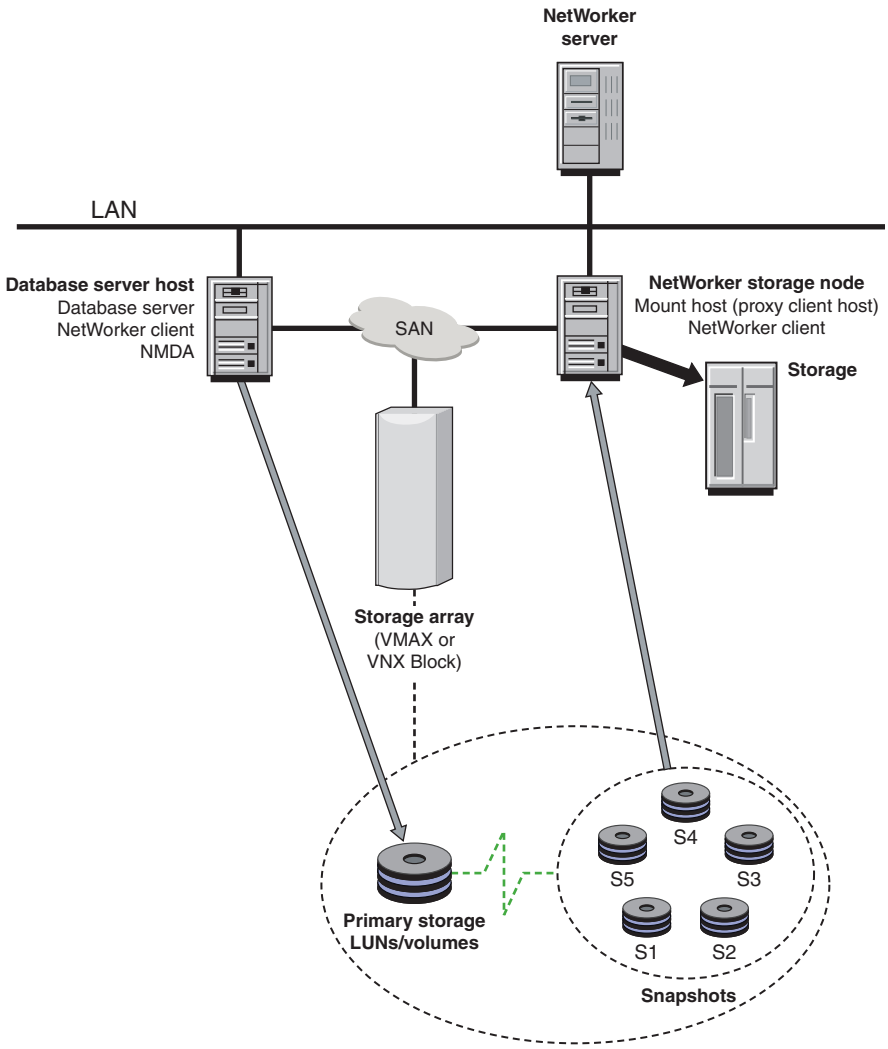

Figure 14 Software used in the NMDA environment for NSM snapshot operations

Types of NSM backups

NMDA supports the types of NSM backups described in the following table.

Table 26 Supported types of NSM backups

Backup type	Snapshot retention	Location of stored data
"Snapshot backup (instant backup)" on page 296	Permanent	Snapshot on primary storage
"Rollover-only backup (immediate live backup)" on page 296	Temporary	Secondary backup media only
"Snapshot and rollover backup (deferred live backup)" on page 296	Permanent	Both primary and secondary storage

The Snapshot Policy attribute in the NetWorker Group resource specifies the NSM backup type. "Configuring Group resources for NSM backups" on page 301 provides details.

Snapshot backup (instant backup)

A snapshot backup, formerly known as an instant backup, creates a permanent point-in-time copy or snapshot of the data on the primary storage system. The snapshot is available to NMDA for performing snapshot restores or rollbacks. You can schedule a snapshot backup to occur many times in a single day with minimal impact to the database server or network.

You must configure a snapshot policy to control the lifecycle of the snapshot. The policy specifies the frequency of snapshot backups and how long to retain snapshots before recycling the snapshots.

Rollover-only backup (immediate live backup)

A rollover-only backup, formerly known as an immediate live backup, is also called a serverless snapshot backup. For this type of backup, NSM creates a temporary snapshot and immediately backs up the snapshot to secondary storage. The software then automatically deletes the snapshot from the primary storage so that NMDA cannot use the snapshot for performing snapshot restores or rollbacks.

NSM can use a proxy client host, also called the mount host, that is separate from the database server host to move the snapshot to the secondary storage. The use of a proxy client reduces the impact on the database server. The proxy client or mount host is typically a NetWorker storage node.

Snapshot and rollover backup (deferred live backup)

A snapshot and rollover backup, formerly known as a deferred live backup, creates a permanent snapshot and uses NSM to back up the snapshot to secondary storage. The permanent snapshot is retained on the primary storage and is available to NMDA for performing snapshot restores or rollbacks for the period specified by either of the following policies:

- Snapshot expiration policy
- Retain Snapshots attribute of the snapshot policy

Similar to a rollover-only backup, NSM can use a proxy client host, also called the mount host, to roll over the backup to secondary storage.

Types of NSM restores

NMDA supports the types of NSM restores described in the following table.

Table 27 Supported types of NSM restores

Restore type	Data retrieved from
"Snapshot restore (instant restore)" on page 297	Mounted snapshot
"Restore from rollover (restore from secondary storage)" on page 297	Secondary storage
"Rollback" on page 297	Unmounted snapshot, entirely restored to source location

The RESTORE_TYPE_ORDER parameter setting specifies the types of NSM snapshot restore. "Configuring NSM snapshot restore and recovery" on page 303 provides details.

Snapshot restore (instant restore)

A snapshot restore, formerly known as an instant restore, retrieves the saved data from a mounted snapshot created through a snapshot backup. This type of restore requires a minimal amount of time.

Restore from rollover (restore from secondary storage)

A restore from rollover, formerly known as a restore from secondary storage, restores a snapshot (saved to the secondary storage system) from the storage system to the database server. You can redirect this type of restore to a different host.

NSM supports two types of restore from secondary storage: conventional restore and file-logical image restore (FLIR). The NSM documentation provides details.

Note: The restore of snapshot data from the secondary storage requires the use of NSM.

Rollback

A rollback restores an entire snapshot back to the source location on the database server by using the hardware capabilities. A rollback does not support relocation of the database to a different host because the relocation requires the reverse synchronization of the data between the snapshot and its original source.

Use the NSM documentation to determine if a rollback is supported on a specific type of hardware.

NOTICE

A rollback is a destructive restore because the rollback overwrites the entire contents of a snapshot unit, such as a volume or disk.

NSM backup processes

Note: You can start an NSM backup only by automatic or manual invocation of the NetWorker scheduled backup group. For example, you cannot schedule an NSM backup of Oracle data through Oracle Enterprise Manager or start the backup from Oracle RMAN. The NSM documentation describes how to manually start a scheduled backup.

In general terms, an NSM snapshot backup includes the following processes.

1. At the time of a scheduled snapshot backup, the NetWorker server starts the NMDA **nsrdasv** process, which invokes the DB2 or Oracle RMAN backup.
2. The DB2 or Oracle backup process loads the NMDA libnsrdb2.*xx* or libnsrora.*xx* library, respectively. Each library communicates with NSM.
3. On the database server host, NSM takes a point-in-time (PIT) snapshot of the database data on the primary storage. NSM uses an application programming interface (API) specific to the storage system to take the snapshot. The snapshot then becomes available to the proxy client (mount host).

4. If the backup type has a rollover, NSM on the proxy client moves the snapshot data on the primary storage to the NetWorker server or the storage node. The server or storage node stores the data on secondary storage, such as tape or disk. If required, NSM deletes the snapshot from the primary storage.
5. At the end of the snapshot backup, the NetWorker server updates the online client index and the media database with information about the backup.

The following figure illustrates the data flow of NSM snapshot backups and restores in an NMDA environment with NetWorker client release 8.1 or later. The NSM feature provides the snapshot backup and restore functionality as part of the NetWorker client software.

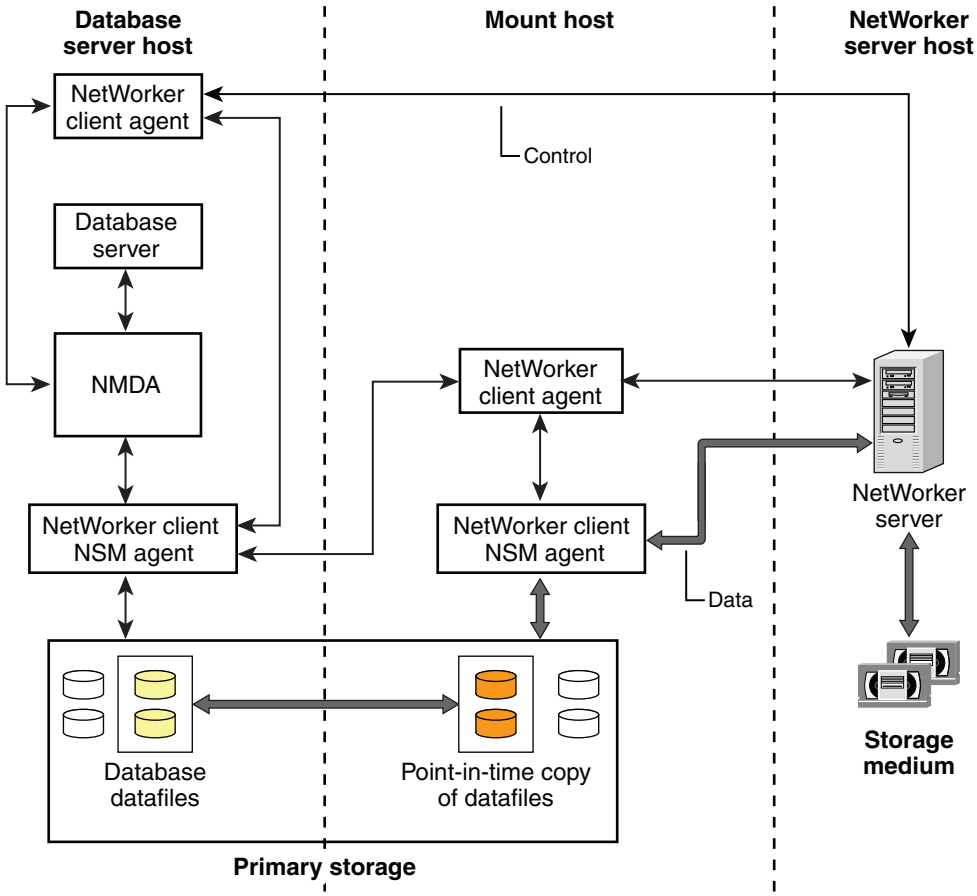

Figure 15 NSM snapshot backup and restore data flow

Configuring NSM snapshot backups

You must configure the required NetWorker resources and parameters for NMDA snapshot backups with NSM by using either the configuration wizard or the nonwizard method.

1. Ensure that you have installed both NMDA and the required NetWorker client (including the NSM feature) on the database host according to the instructions in the following documents:

   - *EMC NetWorker Module for Databases and Applications Installation Guide*
   - NetWorker documentation

   NOTICE

   To enable NMDA operations on UNIX systems, ensure that the /nsr/apps and /nsr/apps/tmp directories have the drwxrwxrwt access permissions.

2. Ensure that you have installed the required NetWorker client on the mount host.
3. Review "Types of NSM backups" on page 295 to determine which type of snapshot backup to perform.
4. Ensure that you have completed the basic database server configuration and NetWorker configuration according to "Configuring NMDA backups" on page 70.

   You must configure the required NetWorker Server, Client, Device, and other resources. For backups that include a rollover, you must configure a Device resource for each secondary storage device, such as a tape drive, used during the backups. You must mount the devices before the backups.

   Ensure that you have configured any required Snapshot Pool resources according to "Configuring Snapshot Policy resources for NSM backups" on page 301.
5. If required, configure internationalization (I18N) support according to "Configuring internationalization (I18N) support" on page 300.
6. For Oracle snapshot backups, the wizard supports the snapshot backup configuration. This is the recommended method of configuring snapshot backups of Oracle data. Follow the information in "Configuring a scheduled backup with the wizard" on page 83 except skip any sentence that includes a cross-reference.

   On the **Select the Backup Application Type** page in the wizard, select the checkbox that enables snapshot management for the selected Oracle application.
7. Review the considerations and perform any required procedures for DB2 or Oracle snapshot backups:

   - "DB2 considerations for NSM snapshot operations" on page 305
   - "Oracle considerations for NSM snapshot operations" on page 308
8. If you want to configure DB2 snapshot backups or if you cannot use the wizard method to configure Oracle snapshot backups for some reason, complete the following steps:

   a. Configure the required NetWorker resources:

      – "Configuring Snapshot Policy resources for NSM backups" on page 301
      – "Configuring Group resources for NSM backups" on page 301
      – "Configuring Client resources for NSM backups" on page 301

b. Set the required NSM parameters according to "Setting the NSM parameters" on page 304.

If the **nsrsnapck** binary is not in the default installation location, set the NSR_NWPATH parameter. On Linux, the default installation location is /usr/sbin. "NSR_NWPATH" on page 364 provides details.

9. If you are not using the configuration wizard, complete any additional configuration procedures:

- For an Avamar deduplication backup with NSM, follow the information in "Configuring deduplication backups with NSM" on page 302.
- For a Client Direct snapshot backup to an AFTD, follow the information in "Configuring Client Direct backups to AFTD" on page 104 except skip the step about configuring a scheduled backup.
- For a Client Direct snapshot backup to a DD Boost device, follow the information in "Configuring DD Boost backups to use Client Direct" on page 101 except skip the configuration instructions in the second column of Table 11 on page 102.

Test the NSM snapshot backup configuration according to "Testing an NSM snapshot backup" on page 302.

Configuring Snapshot Pool resources for NSM backups

You must configure a separate pool to support NSM snapshot backups. NSM stores the metadata from the point-in-time copy (snapshot) in this pool. You configure the pool by using the same method as for a regular NMDA backup. However, the specified backup device must be a file or advanced file type.

Note: Specify the pool name in the Snapshot Pool attribute of the NetWorker Group resource, as described in "Configuring Group resources for NSM backups" on page 301.

The NSM documentation describes the configuration of this extra pool.

Configuring internationalization (I18N) support

In a non-English environment, NMDA supports internationalization (I18N) of snapshot backups and restores with NSM. To configure I18N support for NSM snapshot backups, follow the instructions in "Configuring internationalization (I18N) support" on page 71.

As an additional Oracle requirement, enable catalog synchronization for NSM backups by using the **nsroraadmin** command to set the NSR_ORACLE_NLS_LANG parameter to the same value as the environment variable NLS_LANG. "Configuring I18N support for Oracle operations" on page 73 describes the NLS_LANG variable.

For example, in a Japanese locale, set the parameter as follows:

```
nsroraadmin -r add NSR_ORACLE_NLS_LANG JAPANESE_JAPAN.JA16EUC
```

"Configuring the NWORA resource file with the nsroraadmin program" on page 329 describes the **nsroraadmin** command. The command sets the parameter value in the NWORA resource file.

Configuring Snapshot Policy resources for NSM backups

You need a special NetWorker snapshot policy to perform NSM snapshot backups. You can either specify a preconfigured policy or create a snapshot policy.

Configure a NetWorker Snapshot Policy resource by using the instructions in the NSM documentation.

Configuring Group resources for NSM backups

For NSM snapshot backups, configure a NetWorker Group resource by using the instructions in the NSM documentation.

Ensure that you configure the NetWorker Group resource to which the database server host belongs with the snapshot attribute settings. The following table provides a sample.

Table 28 Sample NetWorker Group resource attributes for a snapshot backup

Attribute	Setting
Snapshot	True.
Snapshot Policy	Serverless backup or rollover-only backup. You can set daily or other customized policies later.
Snapshot Pool	A Pool resource dedicated to the storage of snapshot metadata. Use file-type volume devices instead of tape.
Start Time	Set in relation to the Number of Snapshots attribute for the snapshot policy: (Interval *x* Number of Snapshots) must be less than or equal to (24:00 h - Start Time).
Interval	Set in relation to the Number of Snapshots attribute for the snapshot policy.

Configuring Client resources for NSM backups

Configure the NetWorker Client resource for the database server host according to "Configuring the Client resource" on page 89.

The following Client resource requirements for the database server host apply to NSM snapshot backups:

- The Browse Policy attribute in the Client resource applies only to the NetWorker client file index entries for backups on secondary storage.
- The Retention Policy attribute in the Client resource applies only to the NetWorker media database entries for backups on secondary storage.
- The snapshot policy determines the lifecycle of a point-in-time copy (snapshot). You specify the snapshot policy in the Group resource to which the given client belongs.
- For the Group attribute, specify the name of the NetWorker Group resource created for the NSM backups. "Configuring Group resources for NSM backups" on page 301 provides details.

- For the Backup Command attribute, type the following value:

  nsrdasv(**.exe**) **-z** *configuration_file_path*

  where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the NMDA parameter settings for the NSM backup.

- For NSM backups that use a proxy client host, the Remote Access attribute must include the proxy client hostname.
- For DB2 backups, use the Application Information attribute in the Client resource to set the NSM parameters for the appropriate primary storage system as described in "Setting the NSM parameters" on page 304.

Create a NetWorker Client resource for the mount host according to the NSM documentation.

Configuring deduplication backups with NSM

NMDA supports deduplication NSM backups of only the data that is rolled over to secondary storage, not the data in snapshot (PIT) backups.

The following restrictions apply to deduplication backups with NSM:

- NMDA does not support deduplication NSM backups of raw devices or volumes.
- (Avamar deduplication backups only) If NSR_PS_SAVE_PARALLELISM is set to a value greater than 4, then NSM automatically reduces the parameter setting to 4 during the deduplication NSM backup.

You can use the following steps to configure a deduplication backup with NSM.

1. Perform the required NSM backup configurations according to "Configuring NSM snapshot backups" on page 299.
2. Perform the required deduplication backup configurations according to "Configuring scheduled Avamar deduplication backups" on page 98 by using the NMC method for a client-side configuration (without the wizard).

The following sources describe additional requirements and limitations of deduplication backups and restores with NSM:

- *EMC NetWorker Module for Databases and Applications Release Notes*
- NSM documentation

Testing an NSM snapshot backup

NOTICE

You can start an NMDA backup with NSM only by automatic or manual invocation of the scheduled NetWorker backup group.

To verify the scheduled backup setup, follow the instructions for regular backups in "Testing scheduled backups" on page 146.

To determine if the Oracle RMAN script for an Oracle NSM backup contains any errors, log the RMAN output into a file. Specify this logging by setting the parameter NSR_RMAN_ARGUMENTS in the NMDA configuration file used for the backup.

Canceling an NSM snapshot backup

Cancel an NSM snapshot backup by using the same methods as used for nonsnapshot scheduled backups. "Canceling scheduled backups" on page 146 provides details.

Configuring NSM snapshot restore and recovery

To configure an NSM snapshot restore, set the required parameters for the restore by using the proper method:

- For a DB2 NSM restore, set the parameters in the NMDA configuration file as described in "NMDA configuration file" on page 354.
- For an Oracle NSM restore, set the parameters as described in "NSM parameter settings" on page 312.

Set the following parameters for an NSM snapshot restore:

- NSR_DATA_MOVER—Hostname of the proxy host used for the backup.
- NSR_SERVER—Hostname of the NetWorker server.
- RESTORE_TYPE_ORDER—One or more of the following values, with each value separated from the others by a colon (:):
  - pit—Specifies a snapshot (PIT) estore.
  - conventional—Specifies a snapshot restore from secondary storage media.
  - flir—Specifies a file-logical image restore from secondary storage media.
  - rollback—Specifies a rollback restore from a point-in-time snapshot copy.

  The default value of RESTORE_TYPE_ORDER is pit:conventional.

  If you specify multiple values for RESTORE_TYPE_ORDER, the software attempts each type of restore in the order specified until a restore operation succeeds.

  NOTICE

  For the RESTORE_TYPE_ORDER parameter, NMDA does not support the force_rollback option, which is supported by NSM. If you specify the option, the restore fails, even if you also specify other valid restore options.

  "Types of NSM restores" on page 296 describes the supported restore types.

The NSM documentation describes the parameters used to restore snapshots.

Setting the NSM parameters

The following table provides a basic list of supported NSM parameters. The list is not exhaustive. The NSM documentation provides a complete list of NSM parameters.

Set the required NSM parameters by using the proper method:

- For DB2 NSM operations, set the parameters in the Application Information attribute in the NetWorker Client resource for a backup and in the NMDA configuration file for a restore.
- For Oracle NSM operations, set the parameters as described in "NSM parameter settings" on page 312.

Table 29 NSM parameters

Parameter	Description	Default and valid values
NSR_DATA_MOVER	Mandatory for an NSM backup or restore that uses a proxy client host. Specifies the hostname of the proxy client host.	• None (default). • The valid hostname of the proxy client host.
NSR_MAX_STREAMS	Optional. Specifies the maximum number of restore streams.	• 16 (default). • An integer value.
NSR_PS_SAVE_PARALLELISM	Optional. Specifies the number of concurrent backup streams on the proxy client host.	• 16 (default). • An integer value less than or equal to the Parallelism attribute value in the NetWorker Client resource.
NSR_SNAPSHOT_POSTCMD	Optional for an Oracle NSM backup only. Specifies a script to run immediately after the snapshot is complete but before the backup goes to NetWorker devices, if applicable. The script output goes to the file $NSR_DIAGNOSTIC_DEST/nmda_oracle_proxy.messages.	• None (default). • Full pathname of the script to run immedately after the snapshot is complete but before the backup goes to NetWorker devices. If the script pathname is incorrect or the user does not have execute permissions, the backup fails before starting and generates an error. If the script pathname is correct but the script execution fails, the backup does not fail.
RESTORE_TYPE_ORDER	Optional. Specifies the type of NSM restore to perform. **Note:** If you specify multiple values, the software attempts each type of restore in the order specified until a restore operation succeeds.	• pit:conventional (default). • One or more of the following values, each value separated from others by a colon (:): - pit—Specifies a snapshot (PIT) restore. - conventional—Specifies an NSM restore from secondary storage media. - flir—Specifies a file-logical image restore from secondary storage media. - rollback—Specifies a rollback restore from a point-in-time copy. "Configuring NSM snapshot restore and recovery" on page 303 provides details.

Example 30 NSM parameter settings for different types of primary storage

The NSM documentation for the given type of primary storage describes the NSM parameters in the following examples:

- VNX Block:

```
NSR_DATA_MOVER=datamover.emc.com
NSR_SNAP_TYPE=emcclar
EMCCLAR_SNAP_SUBTYPE=Clone
FRAME_IP=10.5.167.17:10.5.167.18
```

- VMAX:

```
NSR_DATA_MOVER=datamover.emc.com
NSR_SNAP_TYPE=symm-dmx
```

DB2 considerations for NSM snapshot operations

For NSM snapshot operations with DB2 data, the IBM DB2 software provides a feature called Advanced Copy Services (ACS) that enables snapshots of DB2 data. The NMDA software works with ACS, NSM, and NetWorker software to back up snapshots of DB2 data.

NOTICE

By default, DB2 ACS includes log files in the snapshot backups, in addition to data files. As a result, ensure that both data and log files are located on a snapshotable device. Otherwise, the snapshot backup fails.

Due to a DB2 ACS limitation, NMDA with NSM only supports the snapshot backup and restore of the whole DB2 database. NMDA does not support the snapshot backup and restore of selected DB2 tablespaces, logs, or other files.

It is recommended that you also set up the automatic backup of transactions logs to protect the logs. The logs are required to run the **db2 rollforward** command to apply the transaction logs and recover a DB2 database to either the current time or a specific point-in-time. The automatic backup of transaction logs does not use snapshots. The logs are backed up through regular backups. "Configuring automatic backups of DB2 transaction logs" on page 116 describes how to configure the DB2 automatic backup of transaction logs to NetWorker devices.

Restoring a DB2 NSM backup

You can use the following steps and referenced documentation to restore an NSM backup of a DB2 database.

1. Configure the DB2 NSM restore according to "Configuring NSM snapshot restore and recovery" on page 303.

   The *DB2 Data Recovery and High Availability Guide* provides details.

2. To restore a database, run the **db2 restore** command. For example, on Linux:

   ```
   db2 restore db SAMPLE use snapshot library /usr/lib/libnsrdb2.so
   options @pathname/nmda_db2.cfg logtarget include force
   ```

   where *pathname*/nmda_db2.cfg is the full pathname of the NMDA configuration file.

   The DB2 documentation describes the **db2 restore** command.

3. To apply the transaction logs and recover a DB2 database to either the current time or a specific point-in-time, run the **db2 rollforward** command.

Managing and deleting DB2 NSM backups

DB2 includes a binary named **db2acsutil** used to perform the following tasks:

- List the valid DB2 snapshot backups on the primary storage
- Delete DB2 snapshot backups and release the associated resources

The IBM DB2 documentation describes the **db2acsutil** utility.

NMDA also supports the synchronous removal ("pruning") of snapshot entries from the DB2 history file through the **nsrdb2cat** binary.

Querying DB2 snapshots

Query of DB2 snapshots with the **db2acsutil** command produces a list of valid snapshots retained in the repository.

Note: You cannot monitor the status of snapshots created with NMDA.

The following examples are snapshot queries on AIX:

```
db2acsutil load /usr/lib/libnsrdb2.so options @pathname/nmda_db2.cfg
   query older than 10 days snapshot db SAMPLE

db2acsutil load /usr/lib/libnsrdb2.so options @pathname/nmda_db2.cfg
   query snapshot instance db2inst1

db2acsutil load /usr/lib/libnsrdb2.so options @pathname/nmda_db2.cfg
   query snapshot taken at 20100612121212

db2acsutil load /usr/lib/libnsrdb2.so options @pathname/nmda_db2.cfg
   query snapshot dbpartitionnum 0

db2acsutil load /usr/lib/libnsrdb2.so options @pathname/nmda_db2.cfg
   query older than 5 days ago instance db2inst1
```

where *pathname*/nmda_db2.cfg is the full pathname of the NMDA configuration file.

Deleting DB2 snapshots

Deletion of DB2 snapshots created with NMDA supports only the **taken at** *yyyymmddhhmmss* option of the **db2ascutil** command. The software deletes snapshot entries from both the DB2 backup history and the NetWorker server indexes.

Note: If the **nsrsnapck** binary, required for deletion operations, is not in the default installation location, set the NSR_NWPATH parameter in the NMDA configuration file. On Linux, the default installation location is /usr/sbin. Appendix A, "NMDA Parameters and Configuration File," describes the NMDA configuration file.

The following example is a snapshot deletion on AIX:

```
db2acsutil LOAD /usr/lib/libnsrdb2.so options @pathname/nmda_db2.cfg
   delete snapshot older than 10 days db SAMPLE taken at 20100612121212
```

where *pathname*/nmda_db2.cfg is the full pathname of the NMDA configuration file.

The DB2 documentation provides details.

Pruning DB2 snapshots

Through the **nsrdb2cat** binary, NMDA supports the synchronous removal ("pruning") of snapshot entries from the DB2 history file when the entries expire and are removed from the NetWorker indexes.

To enable the synchronous pruning, set the following parameters in the NMDA DB2 resource file (/nsr/apps/res/nmdb2.res):

- DB2PATH
- NSR_DB2CAT_MODE

You use the NMDA DB2 resource file to enable the pruning function of DB2 snapshot backups. The NMDA installation includes a template file, /nsr/apps/res/nmdb2.res, that you must modify to enable the pruning.

The following table describes parameters that you can set in the DB2 resource file.

The DB2 resource file uses the same syntax rules as the NMDA configuration file. "Syntax rules for the NMDA configuration file" on page 355 provides details.

Table 30 DB2 resource file parameters (page 1 of 2)

DB2 resource file parameter	Definition	Default and valid values
DB2PATH	Mandatory only if NSR_DB2CAT_MODE is enabled. Specifies the location of the DB2 binary directory.	• Undefined (default) • Pathname of the DB2 binary directory.
NSR_DB2CAT_MODE	Mandatory. Specifies whether automatic catalog synchronization is enabled or disabled for snapshot backups. When enabled, NMDA removes the corresponding entries in the DB2 history catalog file whenever snapshot backups are expired and removed from the NetWorker indexes.	• Undetermined (default) • Enabled = Snapshot backup with catalog synchronization (DB2 history pruning). • Disabled = Snapshot backup without catalog synchronization **Note:** Snapshot backups fail if this parameter value is not enabled or disabled.

Table 30 DB2 resource file parameters (page 2 of 2)

DB2 resource file parameter	Definition	Default and valid values
NSR_DEBUG_LEVEL	Optional. Appendix A, "NMDA Parameters and Configuration File," provides details.	• 0 (default) = Do not generate debug messages. • 1 to 9 = Write debug messages to the debug log file (name has .log extension). The level of detail in the debug messages increases with the debug level.
NSR_REMOVE_ON_FAILURE	Optional. Specifies whether expired NetWorker index entries are removed if corresponding backup entries are not successfully removed from the DB2 history.	• FALSE (default) = Remove expired NetWorker index entries only if they are successfully removed from the DB2 history. • TRUE = Remove all expired NetWorker index entries, even if they are not successfully removed from the DB2 history.

Oracle considerations for NSM snapshot operations

The following topics describe how to configure and perform Oracle operations with NSM.

Note: Certain Oracle RMAN features, such as checking for corrupt blocks, are not applicable to NSM snapshot operations.

NMDA supports proxy backups and restores of archived redo logs.

Oracle does not support proxy backups of datafiles or archived redo logs that reside on Oracle Automated Storage. Oracle Automated Storage is also known as Oracle Automated Storage Management (ASM). "Replication Manager snapshot operations with Oracle ASM" on page 338 describes how to use EMC Replication Manager with NMDA to perform snapshot backups and restores of Oracle ASM data.

Perform the Oracle configuration procedures according to the following topics:

- "Configuring the required Oracle settings" on page 308
- "Configuring the NWORA resource file" on page 310
- "Creating RMAN scripts for NSM snapshot backups" on page 310

Review the following information about performing Oracle backups and restores with NSM:

- "Checking configuration consistency" on page 313
- "Performing Oracle NSM backups" on page 315
- "Verifying Oracle NSM backup information in NetWorker indexes" on page 317
- "Restoring an Oracle NSM backup" on page 319

"NSM backups and restores on cluster systems" on page 335 describes Oracle operations with NSM in a cluster environment.

Configuring the required Oracle settings

For Oracle backups with NSM, do not locate the database control files and online redo log files on the same volume (snapshot unit) as the datafiles that will be backed up through NSM backups.

If the Oracle database will probably have a lot of read or write activity, or an error such as "skgfdisp: async read/write failed" appears, specify the following values in the Registry and Initialization Parameter file:

- In the Registry, specify the following parameters under HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE:
  - ORA_*oracle_sid*_WORKINGSETMAX
  - ORA_*oracle_sid*_WORKINGSETMIN

  Possible values to set for the parameters are as follows:
  - ORA_*oracle_sid*_WORKINGSETMAX = 1600
  - ORA_*oracle_sid*_WORKINGSETMIN = 1200

  More information on these parameters and Oracle memory management on Windows is available in the Oracle document number 46001.1, "Oracle Database and the Windows NT Memory Architecture, Technical Bulletin."
- In the Initialization Parameter file (such as init*oracle_sid*.ora), increase the value of LARGE_POOL_SIZE to a large value that is appropriate for the particular system.

File types not supported for Oracle NSM backups

Oracle software supports snapshot backups of only Oracle database files and archived redo logs. RMAN uses regular (nonsnapshot) backups to back up other files, such as the SPFILE and control file.

Note: Due to Oracle limitations, you cannot perform NSM snapshot backups of Oracle datafiles or archived redo logs on ASM. "Replication Manager snapshot operations with Oracle ASM" on page 338 describes how to back up this environment by using snapshots with Replication Manager.

Control file versus recovery catalog

Note: For NSM snapshot backups, use an RMAN recovery catalog instead of a control file.

The control file of an Oracle database can store only a limited number of backup entries. When the number of entries exceeds the limit, new entries in the control file overwrite old ones. You can determine the number of entries in a control file from the appropriate Oracle dynamic view. The Oracle documentation provides details.

Snapshot backups use control file entries of type PROXY COPY. For snapshot backups, you can use an RMAN recovery catalog instead of a control file because there is no limit to the number of entries a recovery catalog can contain.

NOTICE

If you use a control file as the RMAN catalog during a snapshot backup, ensure that the control file contains enough free entries for the backup. RMAN creates a new entry in the control file for each file backed up in a snapshot backup. The backup of a large database with many files can quickly use all the free entries in the control file and can start overwriting old entries. With overwritten entries, you cannot restore the corresponding backups.

Configuring the NWORA resource file

To enable NSM snapshot backups, set the NSR_ORACLECAT_MODE parameter resource to either enabled or disabled in the NWORA resource file, as described in "NWORA resource file" on page 324. If you do not set the resource value, NSM snapshot backups fail.

To enable catalog synchronization, perform the configuration procedures in "Catalog synchronization for Oracle NSM backups" on page 322. You must configure catalog synchronization before you perform any NSM snapshot backups of a database.

If you enable catalog synchronization, the NWORA resource file must contain a SID resource for each Oracle database to be backed up. Use the NSR_ORACLE_SID parameter in the NWORA resource file to specify the Oracle SID value. Set the NSR_ORACLE_SID parameter to the same value as the ORACLE_SID parameter in the NMDA configuration file.

Creating RMAN scripts for NSM snapshot backups

The information on RMAN backup scripts in "RMAN scripts for manual backups" on page 129 also applies to RMAN scripts for Oracle backups with NSM.

The following added requirements apply to RMAN scripts for Oracle backups with NSM:

- You must set the appropriate parameters as described in "Setting the parameters for Oracle NSM operations" on page 312.
- You must specify the **proxy** or **proxy only** option with each RMAN **backup** command.

  Note: You cannot use certain options of the RMAN **backup** command, such as **maxsetsize**, **filesperset**, and **diskratio**, with the **proxy** option. Contact Oracle Corporation for more information about these RMAN options.

- For Oracle NSM backups, you must include the **%p** variable in the **format** string, either explicitly or implicitly within **%U**. The Oracle backup and recovery documentation provides details.
- Allocate only one channel in the RMAN script. Do not allocate more than one channel in the RMAN script, to try to distribute the NSM snapshot backup over more than one channel.

  Note: The NSM parameter NSR_PS_SAVE_PARALLELISM defines the NSM backup parallelism. Table 29 on page 304 provides details.

The following sample RMAN script performs an NSM backup of an entire Oracle database that resides on one or more primary storage devices:

```
run {
   allocate channel t1 type 'SBT_TAPE';
   send 'NSR_ENV=(
   NSR_PROXY_PFILE=/oracle/rman/proxy.cfg)';
   backup full proxy only
   format 'FULL_%d_%U'
   (database);
   release channel t1;
}
```

NSR_PROXY_PFILE is an optional NMDA parameter for NSM backups. "Setting the parameters for Oracle NSM operations" on page 312 provides details.

Multiple channels in RMAN scripts

The allocation of multiple channels in an RMAN script does not control the degree of backup or restore parallelism. Oracle uses only one of the allocated channels for the backup or restore, unless you use specific backup options.

Example 31 RMAN scripts with multiple channels

The proxy backup performed with the following RMAN script backs up data to either the OracleVolume1 or OracleVolume2 volume pool, not to both volume pools. The Oracle software uses only one of the allocated channels for the proxy backup:

```
run {
   allocate channel c1 type 'SBT_TAPE';
   allocate channel c2 type 'SBT_TAPE';
   send channel c1 'NSR_ENV=(NSR_DATA_VOLUME_POOL=OracleVolume1)';
   send channel c2 'NSR_ENV=(NSR_DATA_VOLUME_POOL=OracleVolume2)';
   backup proxy only tablespace tbs1, tbs2, tbs3, tbs4;
   release channel c1;
   release channel c2;
}
```

The following RMAN script shows a configuration that NMDA does not support. The script distributes proxy backups over two channels, but NMDA does not support this configuration as Oracle uses only one of the channels for the proxy backup:

```
run {
   allocate channel c1 type 'SBT_TAPE';
   allocate channel c2 type 'SBT_TAPE';
   send channel c1 'NSR_ENV=(NSR_DATA_VOLUME_POOL=OracleVolume1)';
   send channel c2 'NSR_ENV=(NSR_DATA_VOLUME_POOL=OracleVolume2)';
   backup proxy
   (tablespace tbs1, tbs2 channel c1)
   (tablespace tbs3, tbs4 channel c2);
   release channel c1;
   release channel c2;
}
```

To ensure that the proxy backup succeeds, use the following RMAN script to replace both of the preceding two backup scripts:

```
run {
   allocate channel c1 type 'SBT_TAPE';
   send channel c1 'NSR_ENV=(NSR_DATA_VOLUME_POOL=OracleVolume1)';
   backup proxy tablespace tbs1, tbs2, tbs3, tbs4;
   release channel c1;
}
```

You might want to allocate more than one channel if you know that some of the data does not reside on supported primary storage devices. In this case, one channel is for NSM snapshot backups and all the others are for nonsnapshot backups.

Setting the parameters for Oracle NSM operations

You can set two types of parameters for the Oracle backups and restores with NSM:

- The NMDA parameters described in Appendix A, "NMDA Parameters and Configuration File."

  You must set the parameters by using one of the methods in "Setting the NMDA parameters" on page 354.

- The NSM parameters described in "NSM parameter settings" on page 312.

NSM parameter settings

You must set the NSM parameters by using one of the following methods:

- By setting the parameters in the **send** command in one of the following ways:
  - With the **rman** command on the operating system command line.
  - In the RMAN backup or restore script.

  "The send command" on page 398 describes the **send** command.

- By setting the parameters in a user-defined configuration file. You must specify the complete pathname of the file in the parameter NSR_PROXY_PFILE, as described in "NSR_PROXY_PFILE" on page 390.

  The configuration file consists of a separate line for each parameter setting:

  ```
  parameter_name=parameter_value
  ```

  where:

  - *parameter_name* is the parameter name, such as RESTORE_TYPE_ORDER.
  - *parameter_value* is the parameter value, such as pit.

Use the following guidelines to set NSM parameters:

- A parameter setting in the configuration file takes precedence over a parameter setting in the **send** command.

  If the same NSM parameter is set to different values in the configuration file and **send** command, the value in the configuration file is the one used for the NSM operation.

- In the configuration file, the first valid occurrence of an NSM parameter takes precedence over any other occurrences of the same parameter in the same file.

- NMDA does not support the following methods:
  - Use of the **parms** option in the **configure channel** command to set NSM parameters.
  - Use of the **setenv** command on the operating system command line to set NSM parameters.

The following examples include NSM parameter settings: Example 30 on page 305, Example 32 on page 313, Example 33 on page 313.

Table 29 on page 304 provides a list of supported NSM parameters.

Example 32 Setting the NSM parameter NSR_SNAPSHOT_POSTCMD

The NMDA Oracle parameter NSR_SNAPSHOT_POSTCMD can optionally specify a user-provided script that the software runs immediately after completing the snapshot but before the Oracle NSM backup moves to NetWorker devices, if applicable.

"NSR_SNAPSHOT_POSTCMD" on page 304 describes the parameter, which you can set with the **send** command, the **parms** options, or in a user-defined configuration file through the NSR_PROXY_PFILE parameter.

For example, if you perform an Oracle NSM backup in mount mode, you can use the NSR_SNAPSHOT_POSTCMD script to open the database. As a result, the database is offline for a shorter period during the backup. This RMAN backup script includes the parameter setting with the **send** command:

```
run {
   allocate channel t1 device type sbt;

   send channel t1
   'NSR_ENV=(NSR_PROXY_PFILE=/space/myorcl/rman/pfileclar.txt,
   NSR_DATA_VOLUME_POOL=Oracle,
   NSR_SNAPSHOT_POSTCMD=/space/myorcl/rman/snap.ksh)';

   backup proxy database;
   release channel t1;
}
```

Example 33 Setting the NSM parameter RESTORE_TYPE_ORDER

To set the NSM parameter RESTORE_TYPE_ORDER for a snapshot restore, you can create a configuration file named /oracle/rman/proxy.cfg that consists of this line:

```
RESTORE_TYPE_ORDER=rollback:pit:conventional
```

In this case, you must set the NMDA parameter NSR_PROXY_PFILE to /oracle/rman/proxy.cfg by using the **send** command. For example, the following command sets the parameter correctly:

```
allocate channel t1 device type 'SBT_TAPE';
send 'NSR_ENV=(NSR_PROXY_PFILE=/oracle/rman/proxy.cfg)';
```

"Configuring NSM snapshot restore and recovery" on page 303 provides details.

Checking configuration consistency

During a scheduled backup, NMDA checks for consistency between the NetWorker Group resource configuration and the RMAN backup session.

If NMDA finds a discrepancy between the Group resource configuration and the RMAN session, warning messages appear or the backup fails as described in the following topics:

- "With a group configured for NSM snapshot backups" on page 314
- "With a group configured for nonsnapshot backups" on page 314

With a group configured for NSM snapshot backups

Setting the Snapshot attribute in the NetWorker Group resource to True configures the resource for NSM snapshot backups. However, this configuration does not guarantee the execution of an NSM snapshot backup. RMAN might still perform only nonsnapshot Oracle backups if either of the following conditions exist:

- None of the **backup** commands in the RMAN script include the **proxy** or **proxy only** option.
- The **backup** commands in the RMAN script include the **proxy** or **proxy only** option. However, none of the Oracle database objects (tablespaces or datafiles) specified in the **backup** commands reside on a primary storage device that NSM supports.

If RMAN performs only nonsnapshot Oracle backups due to one of these conditions, NMDA generates the following warnings in the savegroup completion report:

```
WARNING: Snapshot savegrp is completed but no Oracle proxy backup is
   detected.
WARNING: Either fix your RMAN script or reconfigure the group resource
   without snapshot flag.
```

While the resulting backups are valid nonsnapshot backups (not NSM snapshot backups), correct the RMAN script or relocate the Oracle datafiles to a supported primary storage device, as required to enable NSM snapshot backups.

The current EMC compatibility guides describes the primary storage devices supported for NSM snapshot backups.

If a **backup** command in the RMAN script includes the **proxy only** option and the Oracle data objects reside on volumes that do not support snapshots, the scheduled backup fails. RMAN cannot perform a regular backup of the objects. The Oracle documentation describes the difference between the **proxy** and **proxy only** options.

Note: During a backup, if NSM cannot determine whether a file is snapshotable, the NSM backup fails.

With a group configured for nonsnapshot backups

Setting the Snapshot attribute in the NetWorker Group resource to False configures the resource for nonsnapshot backups. In this case, you cannot use the **proxy** option or **proxy only** option with a **backup** command in the RMAN script. Any NSM backup specified in the RMAN script will fail. If there are nonsnapshot backups and NSM snapshot backups in the same RMAN script, RMAN might complete one or more nonsnapshot backups before an NSM snapshot backup fails.

Notes:

- If RMAN terminates any of the NSM backups in an RMAN script, the savegroup completion report lists failure of the scheduled backup.
- If any NSM backups in an RMAN script fail, RMAN still performs a nonsnapshot backup of the corresponding archived redo logs.

Example 34 Oracle NSM backup failure

A scheduled backup includes the following RMAN script, with the database files residing on volumes that support snapshots while archived logs reside on volumes that do not support snapshots. However, the Snapshot attribute in the Group resource is set to False. As a result, the NSM backup of the database fails:

```
run {
   allocate channel ch1 type 'SBT_TAPE';
   backup proxy database plus archivelog;
}
```

Despite the NSM backup failure, RMAN performs a nonsnapshot backup of the archived redo logs. The savegroup completion report lists failure of the scheduled backup.

Performing Oracle NSM backups

You can start an Oracle NSM backup only through automatic or manual invocation of the scheduled NetWorker backup group.

NMDA creates temporary files for processing purposes in the following directory:

- On UNIX systems, the directory is /nsr/apps/tmp.
- On Windows systems, the directory is *NetWorker_install_path*\apps\tmp, where *NetWorker_install_path* is the root directory of the NetWorker installation path.

Note: During RMAN operations, do not modify any files in this directory.

NSM backup summary line in savegroup report

The savegroup completion report for an NSM backup contains a summary line that includes the backup size and the number of files. The summary line refers to backup data written to secondary storage only.

The summary line for a snapshot backup includes the size of only the metadata stored for the backup, not the size of the files stored on the primary storage as a point-in-time copy. The number of files includes the number of entries generated for the metadata plus the number of entries generated for the backup pieces.

Savegroup completion status

When a snapshot and rollover backup runs as part of a scheduled group, the backup process includes two steps:

1. A snapshot backup is performed. At the end of the snapshot backup, the backup entries for the point-in-time copy are recorded in the NetWorker indexes and RMAN catalog.
2. The backup to secondary storage is performed. At the end of this backup, the backup entries for data stored on the secondary storage are recorded in the NetWorker indexes.

"NSM backup processes" on page 297 provides details.

If the snapshot backup succeeds but the backup to secondary storage fails, the entire scheduled backup fails. However, the point-in-time copy created during the snapshot backup is a valid backup, which you can use for an snapshot restore or a rollback restore.

Note: If RMAN performs only a regular Oracle backup during the snapshot backup step ("Checking configuration consistency" on page 313 describes when this can happen), the backup to secondary storage fails. There is no point-in-time copy to move to secondary storage. Although the entire backup is reported as failed, the backup stores the data on tape and you can use the data for restore.

NWORA resource file backup

If an NSM backup completes successfully, NMDA automatically backs up the NWORA resource file, described in "NWORA resource file" on page 324.

The NWORA resource file backup occurs at the backup level specified in the Schedule resource, for example, incremental. Oracle backups always occur at the full level. The NetWorker server selects the pool for the NWORA resource file backup based on existing resource configurations. The NSR_DATA_VOLUME_POOL parameter setting does not affect the pool selection.

The savegroup completion report contains a summary line for the backup that includes the phrase "NWORA Resource Backup." The information also appears in the backup debug log file in either of the following locations:

- Directory in the NSR_DIAGNOSTIC_DEST setting
- Default directory, /nsr/apps/logs (UNIX) or *NetWorker_install_path*\apps\logs (Windows)

In the NetWorker indexes, the save set name for the NWORA resource file backup is the same as the file pathname. You can use the NetWorker **mminfo** command to display the save set name. "NWORA resource file backup in the NetWorker indexes" on page 318 describes the backup in the NetWorker indexes.

You can restore the NWORA resource file backup by using the NetWorker **recover** command or the **nwrecover** GUI program. The *EMC NetWorker Administration Guide* provides details.

Note: The backup stores the file under the "backup" namespace, not the "oracle" namespace.

The browse policy and retention policy applied to the NWORA resource file backup are the most conservative policies of the given NetWorker client, not the policies applied to the Oracle backups. As a result, you may see a difference between the policies assigned to the NWORA resource file backup and the Oracle backups.

Enabling point-in-time recovery without a Recovery Catalog

Note: If you perform a point-in-time recovery with an RMAN Recovery Catalog, the information in this topic does not apply.

During an Oracle backup with NSM, the Oracle software backs up the control file after the NSM backup of the datafiles is complete. In a large database production environment, there might be a delay between the end time of the datafile backup and the start time of the control file backup.

If the database structure changes during the time delay (for example, you add a new datafile), you must back up the control file in a separate RMAN session before the change occurs. You must perform the separate backup because the control file backup from the NSM database backup session will include information about the new database structure.

Verifying Oracle NSM backup information in NetWorker indexes

The NetWorker server maintains information about each backup in its online indexes. "Terminology used in this guide" on page 25 provides details.

The index entry for an NSM backup is stored in the NetWorker client file index of the database server host. For example, the entry is stored under the "db2" or "oracle" namespace, similar to a regular backup.

The NetWorker client file index and the media database each contain a different value for the save set name of an Oracle NSM backup, similar to a nonsnapshot scheduled backup.

Query the online NetWorker indexes by using the NetWorker commands **nsrinfo** and **mminfo**:

- Type the **nsrinfo** command to query the NetWorker client file index. For example:

  ```
  nsrinfo -n oracle -s NetWorker_server Oracle_Server_hostname
  ```

- Type the **mminfo** command to query the NetWorker media database. For example:

  ```
  mminfo -v -s NetWorker_server -c Oracle_Server_hostname
  ```

The *EMC NetWorker Command Reference Guide* and the UNIX man pages describe these NetWorker commands.

Entries in the client file index

For a backup piece created through an Oracle NSM backup, the client file index contains three types of backup entries under the "oracle" namespace:

- One entry is for the backup piece name that RMAN assigns, such as /PROXY_O901JB_811_1/ in Example 35 on page 317.
- The second entry is for the point-in-time metadata, such as /brcmeta.1/ in Example 35 on page 317. Only a snapshot backup creates this entry.
- The third entry is for the Oracle datafile backed up to secondary storage, for example, /JBOD13_NMDA11_MVOL3/tbspc4_data1.dbf in Example 35 on page 317. Only a backup to secondary storage creates this entry.

Example 35 Oracle NSM backup entries in the client file index

The **nsrinfo** command provides information about the NSM backup entries in the NetWorker client file index:

```
nsrinfo -n oracle marmaris

scanning client 'marmaris' for all savetimes from the oracle namespace
/PROXY_O901JB_811_1/, date=1279735274 Wed Jul 21 14:01:14 EDT 2010
/brcmeta.1/, date=1279735271 Wed Jul 21 14:01:11 EDT 2010
Physical files to rollover:
```

```
/JBOD13_NMDA11_MVOL3/tbspc4_data1.dbf
/JBOD13_NMDA11_MVOL3/tbspc4_data1.dbf, date=1279735277 Wed Jul 21
   14:01:17 EDT 2010
```

Entries in the media database

For a backup piece created through an NSM backup, the media database contains two types of entries:

- One entry is for the point-in-time metadata. Only a snapshot backup creates this entry.

  In the **mminfo** command output for this entry:

  - The Size field contains the size of the metadata stored on the NetWorker device.
  - The Flag field (fl) includes the letter P, representing the point-in-time copy.

  To list the entries for a snapshot backup only, type the following **mminfo** command:

  mminfo -v -c *Oracle_Server_hostname* **-q 'snap'**

  The NSM documentation provides details.

- The other entry is for the Oracle datafile backed up to secondary storage. Only a backup to secondary storage creates this entry.

Both entries in the media database include the name of the RMAN backup script used for the NSM backup, such as /space1/home/oracle/bp1 in the following example.

Example 36 Oracle NSM backup entries in the media database

The **mminfo** command provides information about the NSM backup entries in the NetWorker media database:

```
mminfo -v -c marmaris

volume          client         date          time          size
nmda.002        marmaris       07/21/10      14:01:11      102 MB
snap.001        marmaris       07/21/10      14:01:13        2 KB

ssid            fl             lvl           name
4064690015      cb             full          /space1/home/oracle/bp1
4098244417      cbP            full          /space1/home/oracle/bp1
```

NWORA resource file backup in the NetWorker indexes

In the NetWorker indexes, the NWORA resource file backup is stored under the "backup" namespace. As a result, you can use the NetWorker **recover** program or **nwrecover** program to restore the backup. The save set name for the backup is the same as the file pathname.

Query the NetWorker indexes for information about the NWORA resource file backup by using the NetWorker commands **nsrinfo** and **mminfo**.

Example 37 Resource file backup entry in the client file index

The **nsrinfo** *Oracle_Server_hostname* command provides information about the NWORA resource file backup entry in the NetWorker client file index:

```
nsrinfo marmaris

scanning client 'marmaris' for all savetimes from the backup namespace

/nsr/res/nwora.res, date=1279735271 Wed Jul 21 14:01:11 EDT 2010
/nsr/res/, date=1279735271 Wed Jul 21 14:01:11 EDT 2010
/nsr/, date=1279735271 Wed Jul 21 14:01:11 EDT 2010
/, date=1279735271 Wed Jul 21 14:01:11 EDT 2010
```

Note: The **nsrinfo -n oracle** command does not display the entry because the entry is in the "backup" namespace, not in the "oracle" namespace. The "backup" namespace is the default namespace for the **nsrinfo** command.

Example 38 Resource file backup entry in the media database

The **mminfo -v -c** *Oracle_Server_hostname* command provides information about the NWORA resource file backup entry in the NetWorker media database:

```
mminfo -v -c marmaris

volume          client          date            time            size
nmda.002        marmaris        07/21/10        14:01:11        4 KB

ssid            fl              lvl             name
3863367791      cb              full            /nsr/res/nwora.res
```

The *EMC NetWorker Command Reference Guide* and the UNIX man pages describe these NetWorker commands.

Restoring an Oracle NSM backup

Review the following information about restores of an Oracle NSM backup:

- "Creating RMAN scripts for Oracle NSM restores" on page 319
- "Performing Oracle NSM restores" on page 320
- "Relocating files during an Oracle NSM restore" on page 321
- "Catalog synchronization for Oracle NSM backups" on page 322

"NSM backups and restores on cluster systems" on page 335 describes the restores of NSM backups in a cluster environment.

Creating RMAN scripts for Oracle NSM restores

For an Oracle NSM restore, you can use the same RMAN script that you use for a nonsnapshot Oracle restore.

Note: The RMAN **restore** command does not include a **proxy** option.

To create an RMAN script for an NSM snapshot restore, follow the instructions in Chapter 4, "Data Restore and Recovery."

To perform an NSM restore, you must set the appropriate parameters as described in "Setting the parameters for Oracle NSM operations" on page 312.

Performing Oracle NSM restores

The following requirements apply to Oracle NSM restores:

- You must install the NetWorker client (including the NSM feature) according to the instructions in the NetWorker documentation. Refer to the NetWorker client version for the primary storage system.
- Each element of the restore path must exist. Otherwise, the restore fails. For example, to restore a file backup to /space1/oradata/file.dbf, the path /space1/oradata must exist.
- An Oracle NSM restore of a symbolic link restores the Oracle file to the location pointed to by the symbolic link. Both the symbolic link and the restore path must exist. Otherwise, the restore fails.
- For a rollback restore, you must set the psrollback.res file as described in "Rollback restore" on page 320.
- For a user-specified relocation of files during an NSM snapshot restore, you must specify the relocation path as described in "Relocating files during an Oracle NSM restore" on page 321.
- After an Oracle restore is complete, a database administrator must recover the database by using the standard Oracle **recover** command.

Concurrent restore streams

During an NSM snapshot restore, NSM creates concurrent restore streams to optimize the restore.

The NSR_MAX_STREAMS parameter defines the maximum number of concurrent restore streams. Table 29 on page 304 provides details.

Directory created for Oracle data restore

An NSM restore of Oracle data creates a .nworapc subdirectory with 0700 permissions under the restore directory for the temporary relocation of the restored files. This relocation is independent of a user-specified relocation. The empty .nworapc subdirectory persists after the restore. You can delete the subdirectory manually, if required.

If an NSM restore of Oracle data fails, the nonempty .nworapc subdirectory persists after the restore. You can delete the subdirectory manually, if required. Do not use any datafiles from this subdirectory for Oracle recovery, or database corruption might occur. If you restart the failed restore, NMDA automatically cleans this subdirectory.

Rollback restore

The psrollback.res file lists all the files, directories, partitions, and volumes to exclude from the rollback safety check. A rollback operation overwrites the items excluded from the safety check.

For a rollback restore, the psrollback.res file must contain the directory name .nworapc. The file location is as follows:

- On UNIX systems: /nsr/res/psrollback.res
- On Windows systems: *NetWorker_install_path*\res\psrollback.res, where *NetWorker_install_path* is the root directory of the NetWorker installation path

Add the directory name to the file by using a text editor as either the root user on UNIX or a member of the Microsoft Windows Administrators group.

The following sources describe the psrollback.res file:

- NSM documentation for the primary storage system
- Comments within the psrollback.res file itself

Note: The NSM documentation describes whether rollback is supported on a particular storage system.

Relocating files during an Oracle NSM restore

This topic describes the user-specified relocation of an Oracle NSM restore with NMDA.

NOTICE

A rollback restore does not support relocation. If the RESTORE_TYPE_ORDER parameter includes the rollback value and the RMAN restore script specifies relocation, the restore fails, even if the parameter includes other values.

During an Oracle NSM restore, NMDA supports relocation, which is the restore of datafiles (regular files or raw volumes) to a new location. You can specify the new location by using the RMAN **set newname** command.

Note: A regular Oracle restore supports relocation, but the Oracle Server controls the relocation.

To relocate a regular file or raw volume during an Oracle NSM restore, the **set newname** command must specify the name of the relocated file as one of the following pathnames:

- The complete pathname of the relocated file.
- The complete pathname of a symbolic link that points to the location where the file will be restored.

Example 39 Symbolic link specified in the set newname command

If the symbolic link /tmp/file1 points to /dbapps/proddb/file2 and the **set newname** command specifies the symbolic link /tmp/file1, the restore operation restores the backed-up file to /dbapps/proddb/file2.

NOTICE

The procedure to relocate a raw volume includes a restriction that does not apply when relocating a regular file.

To relocate a raw volume, the base file name (the file name without the directory path) of the original backed-up raw volume must be one of the following file names:

- The base file name of the relocation path specified in the **set newname** command.
- If the **set newname** command specifies a symbolic link, the base file name in the symbolic link.

Example 40 Relocation of a raw volume

A backed-up raw volume has the name /dev/volume_one/rvol1. You can specify the /dev/volume_two/rvol1 relocation path in the **set newname** command because the original path and the relocation path have the same base file name, rvol1. However, specifying /dev/volume_one/rvol2 in the **set newname** command would cause the NSM restore to fail because the original path and the relocation path have different base file names.

The following procedure is one way to relocate /dev/volume_one/rvol1 to /dev/volume_one/rvol2:

1. Create a symbolic link named /tmp/rvol1 that points to /dev/volume_one/rvol2.
2. Specify /tmp/rvol1 in the **set newname** command in the RMAN restore script.

In this case, the relocation succeeds because both the original path and symbolic link name have the same base file name, rvol1.

Catalog synchronization for Oracle NSM backups

During Oracle backups, RMAN stores information about each backup piece in the RMAN repository, also known as the "RMAN catalog". Similarly, NMDA stores information about each backup piece in the NetWorker indexes, or what Oracle documentation refers to as the "MML catalog."

During Oracle restores, the following actions occur:

- The RMAN catalog determines the data to be restored.
- The NetWorker indexes provide information that NMDA requires to perform the restore.

It is important to keep the RMAN catalog and NetWorker indexes synchronized, especially when performing snapshot backups.

The catalogs are unsynchronized when one of the following conditions is true:

- The RMAN catalog contains backup piece entries that do not have corresponding NetWorker index entries.
- The NetWorker indexes contain backup piece entries that do not have corresponding RMAN catalog entries.

Extra entries in the catalogs

Extra entries in the NetWorker indexes do not cause problems if the extra entries contain unique backup piece names that RMAN does not try to reuse for backups.

However, extra entries in the RMAN catalog can cause problems. When the RMAN catalog contains extra entries without corresponding entries in the NetWorker indexes, the following types of problems can occur:

- When you enable RMAN backup optimization, RMAN might skip backing up certain files.
- The RMAN catalog might expire backups required for restores.
- RMAN restores might fail when RMAN tries to restore backup pieces with no corresponding NetWorker index entries.

Extra entries can occur in the RMAN catalog when either expiration or NetWorker commands such as **nsrmm** remove the corresponding NetWorker index entries. This occurrence does not have much impact on nonproxy backups, but becomes a more serious problem for proxy backups. The severity of the problem increases with the frequency of snapshot backups. You might configure snapshot backups to expire quickly, within hours. When a snapshot backup expires, the NetWorker index entries are removed. Catalog synchronization addresses the problem.

Removing snapshot backup entries from the NetWorker indexes

Snapshot backup entries in the NetWorker indexes are removed in one of the following ways:

- At the start of a snapshot backup, if the number of existing snapshot backups equals the the Retain Snapshots attribute value in the Snapshot Policy resource, the following actions occur:
  - The oldest snapshot backup automatically expires.
  - The NetWorker index entries of the oldest snapshot backup are removed.

  Note: The automatic expiration and the index entry removal do not apply to snapshot backups performed through **nsrdasv -c** *different_client_name*, where *different_client_name* is a different name than used in the Client resource for the backup.
- When a snapshot backup expires due to reaching the expiration time defined in the NetWorker Snapshot Policy resource, the NetWorker process **nsrim** prunes the backup entries from the NetWorker indexes.
- The DBA uses a NetWorker command, such as **nsrmm**, to remove a snap set.

Automatic catalog synchronization for NSM snapshot backups

NMDA provides automatic catalog synchronization that resolves the issues described in "Extra entries in the catalogs" on page 323. When you enable catalog synchronization in NMDA, the proxy backup entries in the RMAN catalog and NetWorker indexes are synchronized automatically.

NOTICE

To enable automatic catalog synchronization for an NSM snapshot backup:
- Set the ORACLE_SID parameter in the NMDA configuration file for the NSM backup. "NMDA Oracle parameters" on page 387 provides details.
- Ensure that an NWORA resource file includes the required resources as described in "NWORA resource file" on page 324.

The NMDA program **nsroraclecat** uses the NWORA resources in the file to perform automatic synchronization of the RMAN catalog and NetWorker indexes.

Note: DBAs can also synchronize the catalogs manually by using RMAN commands.

The following topics describe how to configure and perform the catalog synchronization:

- "NWORA resource file" on page 324
- "Automatic catalog synchronization with the nsroraclecat program" on page 332

NWORA resource file

NSM backups require the NWORA resource file to exist in the following location:

- On UNIX systems: /nsr/apps/res/nwora.res
- On Windows systems: *NetWorker_install_path*\apps\res\nwora.res, where *NetWorker_install_path* is the root directory of the NetWorker installation path

The **nsroraadmin** program creates and maintains the NWORA resource file.

To enable automatic catalog synchronization for proxy backups with NSM, the NWORA resource file must include specific NWORA resources.

Note: You must not edit the NWORA resource file manually. You must add, modify, or delete all the resources in the file by using the **nsroraadmin** program only. Run the **nsroraadmin** program as either the root user on UNIX or a member of the Microsoft Windows Administrators group.

"Configuring the NWORA resource file with the nsroraadmin program" on page 329 describes the **nsroraadmin** program.

The NWORA resource file must contain two types of resources: NWORA parameter resources and NWORA SID resources. The following topics provide details:

- "NWORA parameter resources" on page 324
- "NWORA SID resources" on page 327

NWORA parameter resources

To enable NMDA proxy backups with NSM, the NWORA resource file must exist. You must set the NSR_ORACLECAT_MODE parameter to either enabled or disabled. Proxy backups require the following parameters:

- "NSR_NWPATH" on page 325
- "NSR_ORACLECAT_DEBUG_FILE" on page 325
- "NSR_ORACLECAT_LOG_FILE" on page 325

- "NSR_ORACLECAT_MODE" on page 325
- "NSR_REMOVE_ON_FAILURE" on page 325

Note: The parameter resources listed in the following table are the only ones supported. Do not try to add other parameter resources to the NWORA resource file.

Table 31 NWORA parameter resources

Parameter resource	Description	Default and valid values
NSR_NWPATH	Specifies the directory location of the NetWorker binary **nsrsnapck**. **Note:** If you use NMDA with Sun-branded NetWorker, you must set NSR_NWPATH by using the following **nsroraadmin** command: **nsroraadmin -r update NSR_NWPATH=/usr/sbin/nsr**	• Directory pathname for the location of **nsrsnapck** (default). • Valid directory pathname for the location of the NetWorker binary **nsrsnapck**.
NSR_ORACLECAT_DEBUG_FILE	Specifies the debug file used by the **nsroraclecat** program. Set this parameter only for debugging the **nsroraclecat** program. **Note:** The **nsroraclecat** debug file must be in a secure location because the file includes a copy of the strings from the RMAN connection file.	• Undefined (default). • Valid pathname of the **nsroraclecat** debug file. **Note:** If undefined, the **nsroraclecat** program does not generate debug information.
NSR_ORACLECAT_LOG_FILE	Specifies the operations log file used by the **nsroraclecat** program. The logged information includes the backup pieces successfully removed from the RMAN catalog, and those that failed to be removed during automatic catalog synchronization.	• Undefined (default). • Valid pathname of the **nsroraclecat** log file. **Note:** If undefined, the **nsroraclecat** program writes the logging information to the /nsr/applogs/nsroraclecat.log file by default.
NSR_ORACLECAT_MODE	Specifies whether automatic catalog synchronization is enabled or disabled during NSM snapshot backups.	• Undetermined (default). • Enabled. • Disabled. **Note:** Snapshot backups fail if the resource value is not enabled or disabled.
NSR_ORACLE_NLS_LANG	Required to enable catalog synchronization in a non-English environment only. Specifies the non-English locale value as set in the NLS_LANG environment variable. "Configuring internationalization (I18N) support" on page 71 provides details.	• Undefined (default). • Valid locale value, same as set in the NLS_LANG environment variable. **Note:** ,Catalog synchronization fails if the value is not the same as the NLS_LANG variable value in a non-English environment.
NSR_REMOVE_ON_FAILURE	Specifies whether the corresponding NetWorker index entries are removed when the **nsroraclecat** program fails to remove one or more RMAN catalog entries during automatic catalog synchronization. "Automatic catalog synchronization with the nsroraclecat program" on page 332 provides details.	• FALSE (default). • TRUE.

Using the nsroraadmin command to set parameter resources

When you use the **nsroraadmin** command (with any options) for the first time after the NMDA installation, the command automatically creates the NWORA resource file. The new NWORA.res file includes the parameter resources from Table 31 on page 325.

Note: The NWORA.res file also configures save set bundling and policy uniformity through the NSR_BUNDLING and NSR_INCR_EXPIRATION parameter settings, respectively. "Configuring save set bundling for scheduled Oracle backups" on page 135 and "Configuring policy uniformity for scheduled Oracle backups" on page 136 provide details.

Depending on the command options used, the **nsroraadmin** command sets the parameter resources to either default values or customized values.

Note: You cannot delete the NWORA parameter resources. However, you can modify the parameter values by using the **nsroraadmin** command.

To view the NWORA parameter resources in the resource file, use the **nsroraadmin -r list** command.

To modify NWORA parameter resource settings, use the **nsroraadmin -r update** command.

"Configuring the NWORA resource file with the nsroraadmin program" on page 329 describes how to use the **nsroraadmin** command.

Example 41 Default NWORA parameter resources

After the NMDA installation, if the first **nsroraadmin** command used is **nsroraadmin -r list** (to list the NWORA resource file contents), the command adds the following NWORA parameter resources to the resource file:

```
NSR_NWPATH=NetWorker_binary_path
NSR_ORACLECAT_MODE=enabled
NSR_REMOVE_ON_FAILURE=undetermined
NSR_ORACLE_NLS_LANG=
NSR_ORACLECAT_LOG_FILE=
NSR_ORACLECAT_DEBUG_FILE=
NSR_TMPDIR=
NSR_BUNDLING=disabled
NSR_INCR_EXPIRATION=disabled
```

NetWorker_binary_path is the pathname of the directory that contains the NetWorker binary **nsrsnapck**.

Note: You cannot use the NSR_TMPDIR, NSR_BUNDLING, and NSR_INCR_EXPIRATION parameters for snapshot backups with NSM.

To enable proxy backups with NSM, you must use the **nsroraadmin -r update** command to set NSR_ORACLECAT_MODE to either enabled or disabled.

The default NWORA resource file does not yet contain any NWORA SID resources as described in "NWORA SID resources" on page 327.

NWORA SID resources

An NWORA SID resource comprises a specific group of parameters for a single Oracle database. If you enable automatic catalog synchronization by setting NSR_ORACLECAT_MODE to enabled, the NWORA resource file must contain an NWORA SID resource for each Oracle database (ORACLE_SID). The NWORA SID resource can include only the parameters described in the following table.

However, you can add an unlimited number of NWORA SID resources to the resource file.

NOTICE

If you have enabled automatic catalog synchronization but you do not create an NWORA SID resource for an Oracle database, catalog synchronization of that database might fail. As a result, the catalogs can become unsynchronized unless you synchronize the catalogs manually by using RMAN commands. "Automatic catalog synchronization with the nsroraclecat program" on page 332 provides details.

Note: Each NWORA SID resource must have a unique NSR_ORACLE_SID value.

Table 32 NWORA SID resource components (page 1 of 2)

Parameter	Description	Default and valid values
NSR_ORACLE_CONNECT_FILE	Mandatory. Specifies the location of the file containing the connection strings required to create an RMAN session. "The connection file" on page 328 provides details.	• Undefined (default). • Valid pathname of the RMAN connection file.
NSR_ORACLE_HOME	Mandatory. Specifies the home directory of the Oracle installation. The RMAN executable must be in the subdirectory bin in this directory.	• Undefined (default). • Valid pathname of the Oracle home directory. **Note:** The value must equal the values of the Oracle parameter $ORACLE_HOME.
NSR_ORACLE_LIB_PATH	Optional. Specifies the pathname of the directory containing the Oracle shared libraries on UNIX, typically $ORACLE_HOME/lib.	• Undefined (default). • Valid pathname of the Oracle shared library directory on UNIX. **Note:** Do not set the parameter on Windows systems.
NSR_ORACLE_SID	Mandatory. Specifies the SID value of the Oracle database whose RMAN catalog is to be synchronized.	• Undefined (default). • Valid SID value of the Oracle database. **Note:** The value must be equal to the ORACLE_SID value in the NMDA configuration file used for the database backup. "NMDA Oracle parameters" on page 387 provides details.

Table 32 NWORA SID resource components (page 2 of 2)

Parameter	Description	Default and valid values
NSR_ORACLE_TNS_ADMIN	Optional. Specifies the pathname of the directory containing the Oracle Net configuration files.	• Undefined (default). • Valid pathname of Oracle network configuration directory. **Note:** The value must be equal to the Oracle parameter $TNS_ADMIN value.

Using the nsroraadmin command to set SID resources

To add an NWORA SID resource to the resource file, use the **nsroraadmin -r add** command. Example 43 on page 331 shows an example of how to use this command.

To modify NWORA SID resource settings, use the **nsroraadmin -r update** command.

"Configuring the NWORA resource file with the nsroraadmin program" on page 329 describes how to use the **nsroraadmin** command.

The connection file

Catalog synchronization requires the connection file for an Oracle database. The **nsroraclecat** program uses the information in the connection file to delete RMAN catalog entries.

In the NWORA SID resource for the target database, you must set the parameter NSR_ORACLE_CONNECT_FILE to the pathname of the connection file. "NWORA SID resources" on page 327 provides details.

NOTICE

A DBA must create the connection file in a secure location.

The connection file must include the following strings:

- The connection string required to connect to the target database.
- If you use an RMAN recovery catalog, the connection string required to connect to the RMAN recovery catalog.

Note: The connection file must not contain any lines that include the # symbol.

If the connection file does not contain a connection string for an RMAN recovery catalog, the **nsroraclecat** program assumes that a control file is used as the RMAN repository.

Example 42 Connection file contents

If the following lines exist in the connection file, an RMAN recovery catalog is used as the RMAN repository:

```
connect target sys/oracle@proddb;
connect rcvcat rman/rman@oracat;
```

Note: RMAN catalog deletions fail if the connection file does not exist or does not contain the valid connection strings.

Configuring the NWORA resource file with the nsroraadmin program

You must add, modify, or delete all resources in the NWORA resource file by using the **nsroraadmin** program only.

To run the program, type the **nsroraadmin** command at the operating system command line as the root user on UNIX or as a member of the Microsoft Windows Administrators group.

"The nsroraadmin command syntax for NSM snapshot backups" on page 329 describes the command syntax.

Windows Server 2008 and Windows Vista requirements for the nsroraadmin command

On Windows Server 2008 and Windows Vista, you must run the **nsroraadmin** command in the **Command Prompt** window as an administrator:

1. Click **Start**.
2. Right-click **Command Prompt**.
3. Select **Run as administrator**.
4. Run the **nsroraadmin** command in the open **Command Prompt** window.

The nsroraadmin command syntax for NSM snapshot backups

The **nsroraadmin** command syntax for configuring NSM backup settings is as follows:

```
nsroraadmin [-D debug_level] -r list [ResourceName | SidName]

nsroraadmin [-D debug_level] -r add ResourceName ResourceValue

nsroraadmin [-D debug_level] -r add sid=SidName home=OracleHome
   connect=ConnectFilePath [lib=LibraryPath] [tns=TNSPath]

nsroraadmin [-D debug_level] -r update ResourceName ResourceValue

nsroraadmin [-D debug_level] -r update sid=SidName [home=OracleHome]
   [connect=ConnectFilePath] [lib=LibraryPath] [tns=TNSPath]

nsroraadmin [-D debug_level] -r delete SidName
```

where:

- *debug_level* is the level of debug information generated.
- *ResourceName* is the name of an NWORA parameter resource.
- *SidName* is the value of the NSR_ORACLE_SID parameter of an NWORA SID resource.
- *ResourceValue* is the value of the NWORA parameter resource.
- *OracleHome* is the value of the NSR_ORACLE_HOME parameter of the NWORA SID resource.
- *ConnectFilePath* is the value of the NSR_ORACLE_CONNECT_FILE parameter of the NWORA SID resource.
- *LibraryPath* is the value of the NSR_ORACLE_LIB_PATH parameter of the NWORA SID resource.

- *TNSPath* is the value of the NSR_ORACLE_TNS_ADMIN parameter of the NWORA SID resource.

The **-D** and **-r** options are the only supported options:

- The **-D** option causes the **nsroraadmin** command to print debug information.
- The **-r** option must include the appropriate keywords for the NWORA resource operation.

Command options and settings in brackets ([ ]) are optional. Do not include the brackets when you type the command.

The following topics describe how to use the **nsroraadmin** command to list, add, update, and delete NWORA resources:

- "Listing the NWORA resources" on page 330
- "Adding the NWORA resources" on page 330
- "Updating the NWORA resources" on page 331
- "Deleting the NWORA SID resources" on page 331

The following sources describe the **nsroraadmin** command:

- **nsroraadmin** man page on a UNIX Oracle Server that contains the NMDA software
- **nsroraadmin** entry in the *EMC NetWorker Module for Databases and Applications Command Reference Guide*

Listing the NWORA resources

- To display the entire NWORA resource file contents, type this command:

  ```
  nsroraadmin -r list
  ```

- To display the NSR_ORACLECAT_MODE parameter resource only, type this command:

  ```
  nsroraadmin -r list NSR_ORACLECAT_MODE
  ```

- To display an NWORA SID resource with the NSR_ORACLE_SID value of proddb, type this command:

  ```
  nsroraadmin -r list proddb
  ```

Adding the NWORA resources

- To add the NSR_ORACLECAT_MODE parameter resource with the value of enabled, type one of the following commands:

  ```
  nsroraadmin -r add NSR_ORACLECAT_MODE enabled

  nsroraadmin -r add NSR_ORACLECAT_MODE=enabled
  ```

  When you add the NSR_ORACLECAT_MODE parameter resource with the value of enabled, you enable automatic catalog synchronization for proxy backups with NSM.

  Note: If the NWORA parameter resource already exists in the resource file, using the **add** keyword updates the resource value.

- To add a new NWORA SID resource with the NSR_ORACLE_SID value of proddb and other specific values, use the information in Example 43 on page 331.

Example 43 Adding an NWORA SID resource

Before you perform proxy backups of an Oracle database with an ORACLE_SID value of proddb, you use the following command to add an NWORA SID resource to the resource file:

```
nsroraadmin -r add sid=proddb
   home=/dbapps/proddb/app/oracle/product/10.2.0/Db_1
   connect=/dbapps/proddb/connect.file lib=/usr/lib
   tns=/dbapps/proddb/tns
```

Note:
- When adding an NWORA SID resource, the keywords **sid**, **home**, and **connect** are mandatory; the keywords **lib** and **tns** are optional.
- If an NWORA SID resource with the same NSR_ORACLE_SID value already exists, the command updates the values of the existing resource.

After you run the **nsroraadmin** command, the NWORA SID resource includes the following settings:

```
NSR_ORACLE_CONNECT_FILE=/dbapps/proddb/connect.file
NSR_ORACLE_HOME=/dbapps/proddb/app/oracle/product/10.2.0/Db_1
NSR_ORACLE_LIB_PATH=/usr/lib
NSR_ORACLE_SID=proddb
NSR_ORACLE_TNS_ADMIN=/dbapps/proddb/tns
```

In this sample, the RMAN connection file is /dbapps/proddb/connect.file and the Oracle home directory is /dbapps/proddb/app/oracle/product/10.2.0/Db_1.

Updating the NWORA resources

- To update the value of the NSR_ORACLECAT_MODE parameter resource to enabled, type one of the following commands:

  ```
  nsroraadmin -r update NSR_ORACLECAT_MODE enabled

  nsroraadmin -r update NSR_ORACLECAT_MODE=enabled
  ```

- To update the values of the parameters NSR_ORACLE_HOME and NSR_ORACLE_CONNECT_FILE in an NWORA SID resource with the NSR_ORACLE_SID value of proddb, type this command:

  ```
  nsroraadmin -r update sid=proddb home=/dbapps/proddb/10.2.0/Db_1
  connect=/dbapps/oracle/connect/proddb.connect
  ```

  Note: When updating an NWORA SID resource, the keyword **sid** is mandatory. The keywords **home**, **connect**, **lib**, and **tns** are optional.

Deleting the NWORA SID resources

Type this command to delete an NWORA SID resource with the NSR_ORACLE_SID value of proddb:

```
nsroraadmin -r delete proddb
```

Note: You can delete only the NWORA SID resources from the resource file. You cannot delete the NWORA parameter resources.

Automatic catalog synchronization with the nsroraclecat program

NMDA and the NetWorker server jointly manage the automatic catalog synchronization. To remove the NMDA Oracle proxy backup entries from the NetWorker indexes, the NetWorker server runs the **nsrsnapck** program. Before removing the index entries, **nsrsnapck** runs the **nsroraclecat** program to remove the corresponding RMAN catalog entries.

Note: To perform manual catalog synchronization, you can use specific RMAN commands, as described in "Cross-checking and deleting Oracle backups" on page 166. The Oracle documentation describes the RMAN commands.

Review the following information about automatic catalog synchronization:

- "RMAN catalog entry removals with nsroraclecat" on page 332
- "Failure of the nsroraclecat program" on page 333
- "NetWorker index entry removals with nsrsnapck" on page 334
- "Catalog synchronization after NSM Oracle backup volume is relabelled manually" on page 334

RMAN catalog entry removals with nsroraclecat

The **nsroraclecat** program runs on the Oracle Server host that NMDA and NSM backs up:

- Do not try to run the **nsroraclecat** program manually.
- The **nsrsnapck** program automatically runs the **nsroraclecat** program.
- Only one **nsroraclecat** program can run at a time. If two **nsroraclecat** programs are started, the one started first completes its operation before the second one runs.

To remove the RMAN catalog entries, **nsroraclecat** obtains information from the NWORA resource file and generates temporary RMAN scripts that include an RMAN **change...delete** command for each backup piece to be removed.

One script is created for all the backup pieces from the same database (or ORACLE_SID). A separate script is created for each database.

The **nsroraclecat** program names each RMAN script as follows:

- On UNIX systems:

  `/nsr/apps/tmp/.nworapc/nsroracat_`*`date_pid`*

- On Windows systems:

  `NetWorker_install_path``\apps\tmp\.nworapc\nsroracat_`*`date_pid`*

  where:

  - *NetWorker_install_path* is the root directory of the NetWorker installation path.
  - *date* is the current date.
  - *pid* is the **nsroraclecat** process ID.

The **nsroraclecat** program runs each script in an RMAN session. After the scripts have finished running, the program removes the scripts.

Note: The **nsroraclecat** program generates information about the backup piece entries removed from the RMAN catalog. The program writes the information to the **nsroraclecat** log and debug files. "NSR_ORACLECAT_LOG_FILE" and "NSR_ORACLECAT_DEBUG_FILE" on page 325 describe these files.

The following sources describe the **nsroraclecat** program:

- The **nsroraclecat** man page on a UNIX Oracle Server that contains the NMDA software.
- The **nsroraclecat** entry in the *EMC NetWorker Module for Databases and Applications Command Reference Guide* at EMC Online Support.

Failure of the nsroraclecat program

The **nsroraclecat** program fails to remove expired backup pieces from the RMAN catalog in the following cases:

- The **nsrsnapck** program passes invalid information to **nsroraclecat**, for example, an invalid NetWorker client name or an invalid save time of a backup piece.
- The **nsroraclecat** program cannot connect to the NetWorker server to query the NetWorker indexes.
- The **nsroraclecat** program cannot locate the required backup pieces in the NetWorker indexes.

To diagnose the cause of a **nsroraclecat** program failure, review the **nsroraclecat** log files specified by NSR_ORACLECAT_DEBUG_FILE and NSR_ORACLECAT_LOG_FILE. The operations log file is /nsr/applogs/nsroraclecat.log by default.

If the **nsroraclecat** program fails to remove the expired backup pieces, then the **nsrsnapck** program removes the corresponding NetWorker index entries by using the procedures described in "NetWorker index entry removals with nsrsnapck" on page 334.

If the following files exist, you must remove the files:

- Files in one of the following directories:
  - On UNIX systems: /nsr/apps/tmp/.nworapc
  - On Windows systems: *NetWorker_install_path*\apps\tmp\.nworapc, where *NetWorker_install_path* is the root directory of the NetWorker installation path
- Files in either the temporary directory /tmp on UNIX systems or the temporary directory specified by the TEMP system variable on Windows systems, where the files have the name nwora_bp_*sid_pid*:
  - *sid* is an ORACLE_SID value.
  - *pid* is a **nsroraclecat** process ID.

Note: If **nsroraclecat** fails continuously, disable catalog synchronization by setting NSR_ORACLECAT_MODE to disabled until you determine the cause of the problem.

NOTICE

After a **nsroraclecat** program failure occurs or while catalog synchronization is disabled, the DBA must synchronize the catalogs manually by using specific RMAN commands. The Oracle documentation provides details.

NetWorker index entry removals with nsrsnapck

Once the **nsroraclecat** program has successfully removed the expired backup pieces from the RMAN catalog, the **nsrsnapck** program removes the corresponding NetWorker index entries.

If **nsroraclecat** fails to remove some of the backup entries from the RMAN catalog, the **nsrsnapck** program performs the appropriate action:

- Removes the corresponding NetWorker index entries when NSR_REMOVE_ON_FAILURE is set to TRUE.
- Does not remove the corresponding NetWorker index entries when NSR_REMOVE_ON_FAILURE is set to FALSE.

Note: When NSR_REMOVE_ON_FAILURE is set to FALSE, **nsrsnapck** removes only those NetWorker index entries that correspond to removed RMAN catalog entries.

NOTICE

The NSR_REMOVE_ON_FAILURE setting controls whether a corresponding NetWorker index entry is removed when the **nsroraclecat** program fails to remove an expired RMAN backup piece.
- In general, you must set NSR_REMOVE_ON_FAILURE to TRUE to enable NetWorker index entries to be removed, even if the RMAN catalog entries are not removed. Otherwise, if entries are not removed from the NetWorker index, the DBA must synchronize the catalogs manually with the RMAN commands. The Oracle documentation provides details.
- If you enabled RMAN backup optimization, you must set NSR_REMOVE_ON_FAILURE to FALSE to prevent the removal of NetWorker index entries. Otherwise, RMAN might skip backing up certain files.

Catalog synchronization after NSM Oracle backup volume is relabelled manually

If you relabel a NetWorker volume containing NSM Oracle backups, the NMDA program **nsroraclecat** does not remove the corresponding entries from the RMAN catalog during the NetWorker label operation.

In this case, you must perform the following procedures to synchronize the RMAN and NetWorker catalogs:

- Ensure that NSR_REMOVE_ON_FAILURE is set to TRUE in the NWORA resource file.
- Synchronize the RMAN catalog entries manually by using the RMAN **crosscheck** command.

Example 44 on page 335 provides a manual synchronization example.

The Oracle documentation describes the RMAN **crosscheck** command.

Example 44 Using RMAN commands to synchronize the RMAN catalog entries

The following example shows the RMAN commands used to synchronize the RMAN catalog entries after you manually relabelled an NSM Oracle backup volume:

```
connect target;
allocate channel for maintenance type 'SBT_TAPE' parms
   'ENV=(NSR_SERVER=NetWorker_server)';
crosscheck backup;
```

NSM backups and restores on cluster systems

NMDA can perform NSM backups and restores of a database configured on an active-passive cluster system, or on active-active cluster systems only in the combinations listed in the *EMC NetWorker Software Compatibility Guide*. This topic describes the support of cluster and failover for the NSM operations.

NOTICE

You cannot use the parameter NSR_CLIENT for NSM backups in a cluster system. You can use the parameter for restores and nonsnapshot backups in a cluster system, as described in Chapter 6, "Cluster Systems."

Review the following information about NSM operations on a cluster system:

- "NSM backup failover" on page 335
- "Configuring NSM backups from a virtual cluster client" on page 336
- "Configuring NSM backups from a physical cluster client" on page 337
- "Configuring restores from NSM backups on a cluster system" on page 338

NSM backup failover

During an NSM scheduled backup, the NetWorker server retries a failed backup on the failover node if you have met the following requirements:

- You have configured the database server to fail over, for example, by using Oracle Fail Safe with MSCS on a Windows system.
- You have set the Client Retries attribute to a nonzero value in the Group resource.

The retry of a failed backup occurs at the RMAN script level, whereby the RMAN script restarts from the beginning.

Note: To avoid restarting the backups of all objects in the RMAN script during the NetWorker retry, you can use the Oracle restartable backups feature. This feature enables you to back up only the files that have not been backed up since a specified time, for example, by using the **'sysdate -1'** option. "Restartable backups" on page 51 provides details.

Configuring NSM backups from a virtual cluster client

An NSM backup from a virtual cluster client (virtual host) protects the data on shared cluster disks.

You must complete the required steps to configure an NSM backup from a virtual cluster client.

1. Install NMDA and the required NetWorker client on each physical node of the cluster.
2. Create a NetWorker Client resource for the virtual host and each physical host, as described in "Configuring Client resources for NSM backups" on page 301:
   - In the Backup Command attribute, always specify the **-c** *client_name* option in **nsrdasv -z** *configuration_file_path* **-c** *client_name* if the client is a virtual host.
   - In the Remote Access attribute in the Client resource for a virtual cluster client, specify the Oracle user from each physical client that can store and retrieve backups.
   - In the Save Set attribute, specify the complete pathname of the RMAN script to back up the Oracle data on the shared disk.
3. Configure the other NetWorker resources required for NSM backups as described in "Oracle considerations for NSM snapshot operations" on page 308:
   - To enable backup failover, specify a nonzero value in the Client Retries attribute in the NetWorker Group resource for the scheduled backup. This value causes the NetWorker server to restart the failed Oracle backup on the failover node.
   - Specify other recommended attribute settings in the Group resource, as described in the cluster support information of the *EMC NetWorker Administration Guide*.
4. Configure the NWORA resource file on each node of the cluster, as described in "Configuring the NWORA resource file" on page 310.
5. If the NSM backup entries are to be stored in a NetWorker client file index other than the virtual client index, ensure the correct setting of the Remote Access attribute in the Client resource. For example, the backup entries are to be stored in a physical client index. Specify the Oracle user from the virtual host in the Remote Access attribute in the Client resource for *client_name*.

   If the NSM backup entries are stored in a NetWorker index other than the virtual client index, the expiration of snapshot backups created with the **-c** *client_name* option setting is different than without the setting.

   At the start of a snapshot backup, if the number of snapshot backups created by a client equals the Retain Snapshots attribute value in the Snapshot Policy resource, **nsrsnapck** automatically expires the oldest snapshot backup. The **nsrsnapck** program also removes the NetWorker index entries of the backup based on the NetWorker snapshot policy. The **nsrsnapck** program performs the automatic expiration based on the hostname of the NetWorker Client resource. The program is not aware of the **-c** *client_name* option. As a result, the program does not automatically remove the snapshot backup of **-c** *client_name*.

Notes:

- The host specified by *client_name* must have access to snapshot backups.
- NMDA and the NetWorker client (including the NSM feature) must be installed and configured on the host specified by *client_name*.
- When the backup starts from the virtual cluster client, the backup entries are stored in the NetWorker client file index of the virtual client by default.

Example 45 NSM backup entries in the index of a physical cluster client

You want the backup entries to be stored in the index of the physical cluster client mars.emc.com. In this case, you specify **nsrdasv -z** *configuration_file_path* **-c mars.emc.com** in the Backup Command attribute of the NetWorker Client resource.

Configuring NSM backups from a physical cluster client

An NSM backup from a physical cluster client protects Oracle data on private disks. This type of backup is similar to a nonsnapshot scheduled Oracle backup on an unclustered system.

The following sources describe how to set up an NSM backup from a physical cluster client:

- "Configuring the required Oracle settings" on page 308
- "Checking configuration consistency" on page 313
- "Performing Oracle NSM backups" on page 315
- *EMC NetWorker Administration Guide* (chapter on cluster support)

When the backup starts from the physical client, the backup entries are stored in the NetWorker index of the physical client by default.

Note: The entries for the NWORA resource file backup are stored in the NetWorker index of the physical client by default.

To specify that the NSM backup entries be stored in a NetWorker client file index other than the physical client index, for example, in a virtual client index:

- Specify **nsrdasv -z** *configuration_file_path* **-c** *client_name* in the Backup Command attribute in the Client resource for *client_name*.

  Note: When the client is a physical host, the backup is indexed under the hostname of the physical host by default. You need to specify the **-c** *client_name* option only when you want the index to be stored under a different hostname.

- Specify the Oracle user from the physical host in the Remote Access attribute in the Client resource for *client_name*.

If the NSM backup entries are stored in a NetWorker index other than the virtual client index, the expiration of snapshot backups created with the **-c** *client_name* option setting is different than without the setting.

At the start of a snapshot backup, if the number of snapshot backups created by a client equals the Retain Snapshots attribute value in the Snapshot Policy resource, **nsrsnapck** automatically expires the oldest snapshot backup. The **nsrsnapck** program also removes

the NetWorker index entries of the backup based on the NetWorker snapshot policy. The **nsrsnapck** program performs the automatic expiration based on the hostname of the NetWorker Client resource. The program is not aware of the **-c** *client_name* option. As a result, the program does not automatically remove the snapshot backup of **-c** *client_name*.

Notes:

- The host specified by *client_name* must have access to snapshot backups.
- NMDA and the NetWorker client (including the NSM feature) must be installed and configured on the host specified by *client_name*.

Example 46 NSM backup entries in the index of a virtual cluster client

You want to specify that the backup entries be stored in the index of the virtual client monalisa.emc.com. In this case, you specify **nsrdasv -z** *configuration_file_path* **-c monalisa.emc.com** in the Backup Command attribute of the NetWorker Client resource.

Configuring restores from NSM backups on a cluster system

You must complete the required steps to configure a restore from an NSM backup on a cluster system.

1. Set the parameter NSR_CLIENT to the correct value by using one of the methods in "Setting the NMDA parameters" on page 354:
   - To restore a backup from a virtual cluster client, set NSR_CLIENT to the name of the virtual cluster client.
   - To restore a backup from a physical cluster client, set NSR_CLIENT to the name of the physical cluster client.
2. In the Remote Access attribute of the Client resource, specify the hostname of the client on which you will start the restore.

Note: When a failover occurs during a restore, you must restart the restore manually on the failover node.

Replication Manager snapshot operations with Oracle ASM

Oracle software does not support NSM snapshot backups of Oracle data that resides on the Oracle Automatic Storage Management (ASM). You cannot use NMDA and NSM to perform a snapshot backup of Oracle ASM data.

When the Oracle data storage is on the ASM, if you perform an NSM backup of the Oracle ASM data by using the methods in "Snapshot operations with NetWorker Snapshot Management" on page 294, one of the following results occurs:

- If you did not specify **proxy only** in the RMAN **backup** command, then the NSM snapshot backup fails over to a nonsnapshot RMAN backup.
- If you specified **proxy only**, then the NSM snapshot backup fails.

By using EMC Replication Manager, NMDA can perform snapshot backups or split-mirror-based backups of the Oracle ASM data through disk replication technology. This topic describes how to configure the required integration of NMDA with Replication Manager, and then use the integrated solution to perform snapshot backups and restores of Oracle ASM data.

Note: Replication Manager cluster support for Oracle is primarily for Oracle RAC. Replication Manager supports RAC with or without ASM, but converts the RAC database into a single instance database for the replica. While the solution described in this topic also applies to Oracle RAC, you must adjust the scripts as required. The Replication Manager documentation describes additional settings required for Oracle RAC.

NMDA integration with Replication Manager

Replication Manager is an EMC software product that manages the creation and expiration of point-in-time replicas (clones or snapshots) of databases and file systems that reside on supported storage arrays.

When you integrate NMDA with Replication Manager, NMDA can back up and restore the replicas created by Replication Manager. The data in the replicas can also be used for reporting, analysis, and other repurposing.

During the snapshot backups of Oracle ASM data, the software creates replicas of Oracle ASM volumes on the production (source) host, mounted to a separate mount host, and backed up from the mount host. This process reduces the backup overhead on the production host.

A snapshot backup of an Oracle database on ASM storage with Replication Manager includes the following processes:

1. Replication Manager creates and mounts the replica database from the Oracle production database host to the mount host.
2. Oracle RMAN connects to the “replica” database on the mount host, and backs up the data to NetWorker by using NMDA.

Configuration of the snapshot backups of Oracle ASM data includes separate configuration procedures on the four main nodes involved in the backups:

- Replication Manager server
- Oracle production database host (contains Replication Manager Agent software)
- Oracle mount host (contains Replication Manager Agent software)
- NetWorker server

Requirements for backups of Oracle database components

Ensure that you meet the requirements for backups of individual database components:

- Datafiles—You must back up the datafiles in the replica database on the mount host through RMAN scripts. Although the datafiles are from the replica database, you can use the backups to restore and recover the original database.
- Control file—When the replica database is mounted, the control file is a current control file but is a different version from the production current control file. You must convert the current control file to the "backup" type (backup control file) to obtain a new timestamp for the backup control file. The "RMAN and Split Mirror Disk Backups" Metalink note (302615.1) from Oracle provides details.

  RMAN does not back up the backup control file on the mount host. You must back up the control file through the regular NMDA Oracle backup procedures on the production database host by using one of the following options:

  - Set up autobackup of the control file to the NetWorker backup devices, as described in "Control file and server parameter file autobackups" on page 48.
  - Run the **backup current control file** command.
- Parameter files—You must back up the SPFILE through the regular NMDA Oracle backup procedures on the production database host.

  Note: If you do not use an SPFILE, RMAN cannot back up the parameter file or PFILE. You must back up a PFILE on the original database host through a file system backup.
- Archived redo logs—Replication Manager can replicate the archived redo log directory. You can then back up the archived redo logs on the mount host through RMAN scripts.

  To back up additional archived redo logs generated after the replication, you can use the regular NMDA Oracle backups on the original database host.

  If the archived redo logs do not reside on a supported storage array, then you must use the regular NMDA Oracle backups on the production database host to back up the archived redo logs.
- Fast Recovery Area (FRA)—All the procedures used for the archived redo logs backups also apply to the FRA backups.

Requirements for snapshot backups of Oracle ASM data

Ensure that you meet the requirements for snapshot backups of the Oracle ASM data:

- An Oracle RMAN recovery catalog is used for the Oracle database. Snapshot backups do not support the use of the Oracle control file as an RMAN catalog.
- All the Oracle and ASM configuration requirements and setup requirements are met as described in the *EMC Replication Manager Product Guide* and *EMC Replication Manager Release Notes*.

  For example, Replication Manager supports only the external redundancy level for Oracle ASM.

  Confirm that the particular system is supported with Oracle ASM.

If you use Oracle RAC, refer to the Replication Manager documentation for additional Oracle configurations for RAC in ASM environments.

Configuring snapshot backups of Oracle ASM data

To configure the snapshot backups of Oracle ASM data, perform the required configuration steps on the four different nodes involved in the snapshot backups:

- "Configuring snapshot backups on the Replication Manager server" on page 341
- "Configuring snapshot backups on the production database host (Replication Manager Agent)" on page 345
- "Configuring snapshot backups on the mount host (Replication Manager Agent)" on page 345
- "Configuring snapshot backups on the NetWorker server host" on page 348

Configuring snapshot backups on the Replication Manager server

The *EMC Replication Manager Product Guide* describes the configuration steps on the Replication Manager server.

You can perform the following configuration steps by using the Replication Manager console on the Replication Manager server.

1. Create an application set for the Oracle production database. Ensure that you set the application credentials for the application set. The following figure shows a sample of the application credentials.

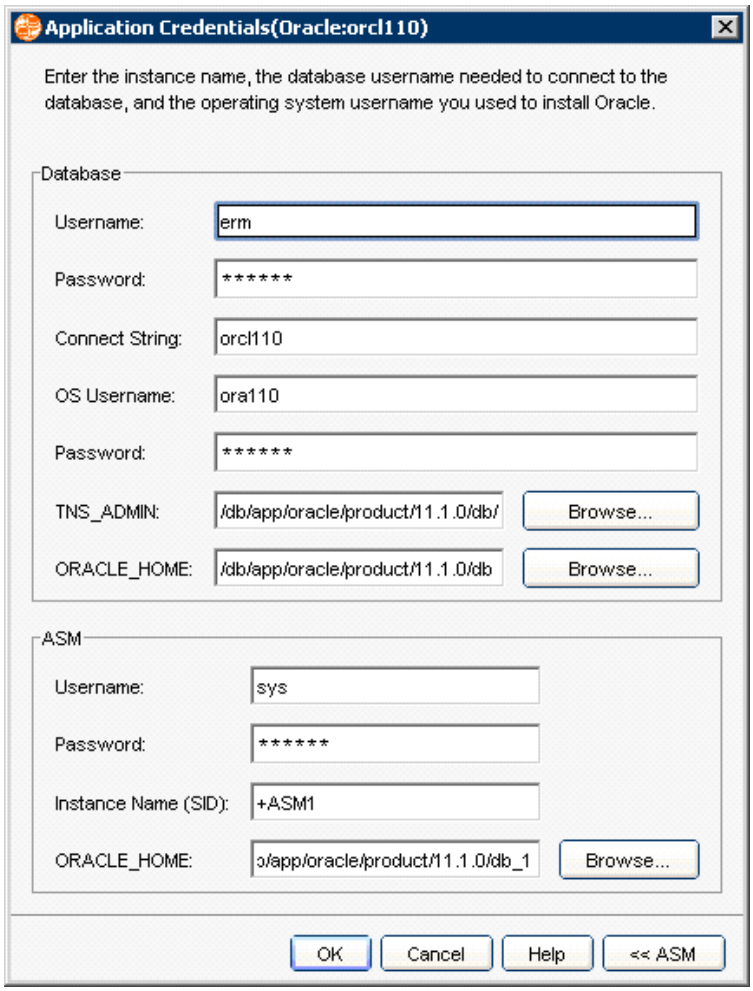

Figure 16 Application Credentials window in the Replication Manager console

2. If required, add storage and create a storage pool.
3. Create a job for the application set created in step 1:
   a. In the **Advanced Replication Settings** window, specify the proper settings for backup and recovery:
      – To perform an online backup, select **Online using hot backup mode**.

        Note: Do not select **Online without hot backup mode** for backup and recovery. The online hot backup mode does not require the **Use consistent split** option, but selecting the option has no impact.

      – To perform an offline backup, select **Offline by shutting down the database**.

        Replication Manager shuts down the database before creating a replica, and restarts the database after creating the replica. Ensure that the snapshot backup schedule corresponds to a time when the database can be placed offline.

      – To back up archived redo logs from the replica through RMAN scripts, select **Replicate archive log directory**.

- To back up the FRA from the replica through RMAN scripts, select **Replicate Flash Recovery Area**.
- If required, select the appropriate setting for **Copy parameter file to RM Server**.

The following figure shows a sample of the advanced replication settings.

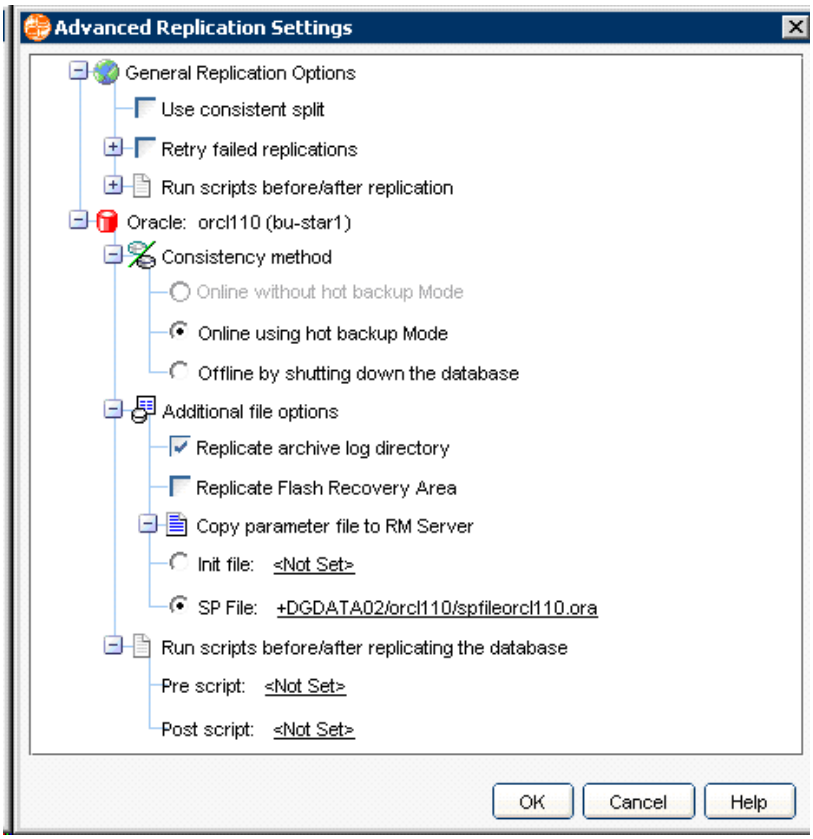

Figure 17 Advanced Replication Settings in the Replication Manager console

b. Specify the following mount options:

- Select a physical host that is different from the mount host.
- Under **Path options for data residing on file systems**, select **Original path**.
- Select **Recover the database**, and specify the following values under that option:
  - Select **Prepare only** for **Recovery type**.
  - Specify the **ORACLE_HOME** value and the user credentials.
- Under **ASM Options**, specify the value for **ORACLE_HOME**.

  You can specify a different ORACLE_HOME for the ASM instance from the one used by the database instance. Replication Manager requires the same operating system user to access both the ASM instance and the database instance. Before you run the job, ensure that you complete the instructions for "Configuring snapshot backups on the mount host (Replication Manager Agent)" on page 345.

– Under **Other Mount Options**, select **Fail the replica if the mount fails**.

Note: Do not select **Unmount the replica on job completion**. The replica must remain mounted during the RMAN backup. After the RMAN backup, the replica is unmounted gracefully through the POSTCMD script described in "Configuring snapshot backups on the mount host (Replication Manager Agent)" on page 345.

The following figure shows a sample of the required mount options.

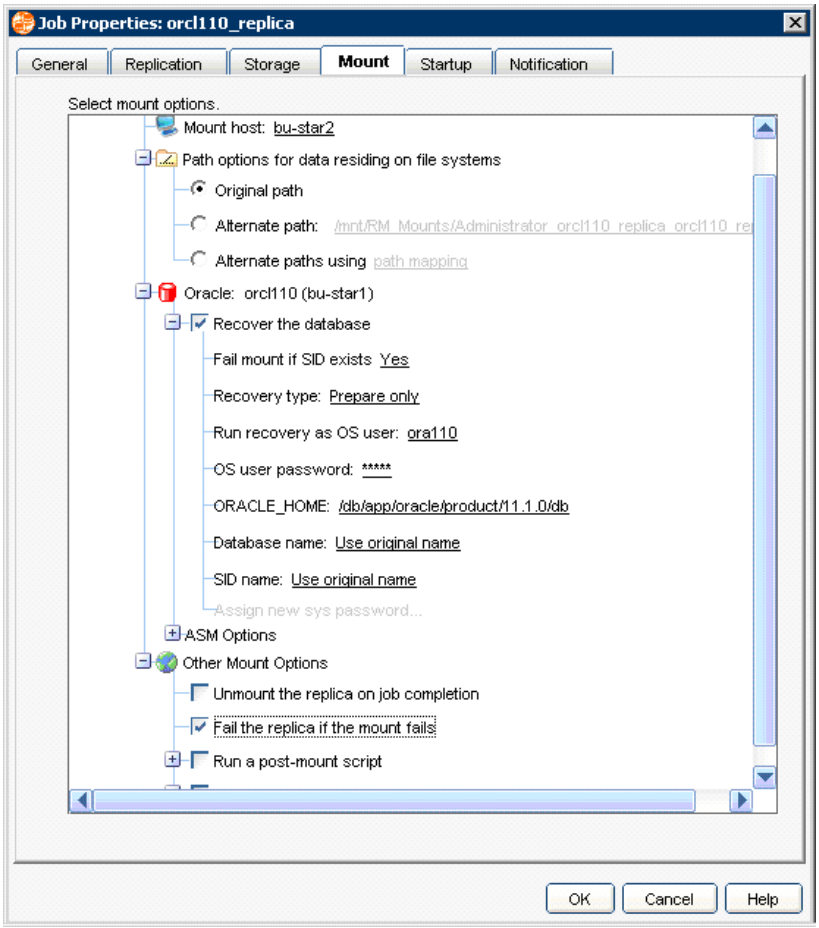

Figure 18 Mount options in the Replication Manager console

4. In the **Starting the Job** window of the job wizard, select **Manually, or using a third party scheduler**. This setting is a requirement for performing a manual or scheduled NetWorker backup.

   The following figure shows the required setting in the job wizard.

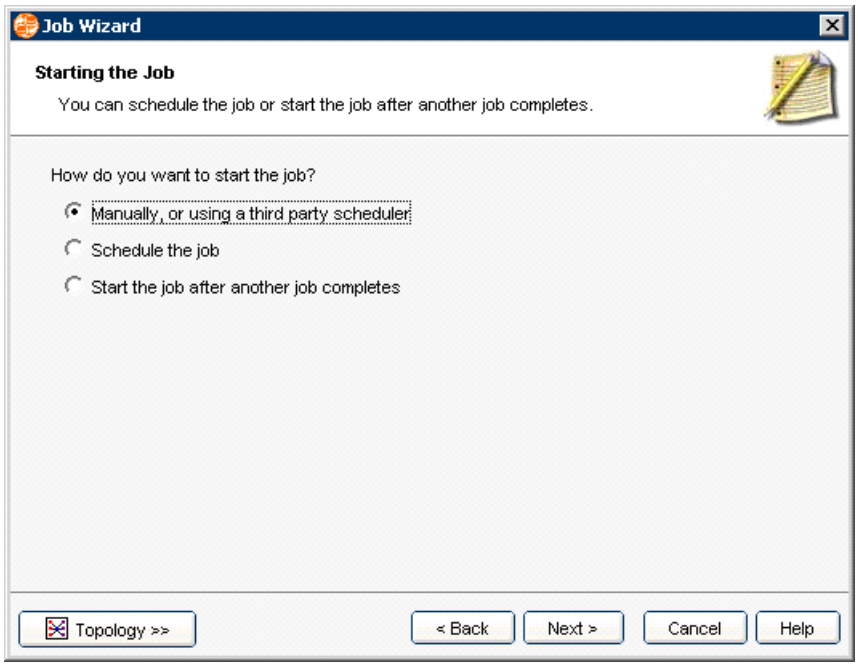

Figure 19 Starting the Job window in the job wizard

Configuring snapshot backups on the production database host (Replication Manager Agent)

You must perform the following configuration steps on the Oracle production database host.

1. Ensure that you have installed and configured the Replication Manager Agent software, and completed the Oracle configuration steps according to the Oracle procedures in the *EMC Replication Manager Product Guide*.

   For example, verify that the tnsnames.ora file contains at least one entry with a dedicated connection that the Replication Manager will use. You must use this dedicated connection when you define an Oracle application set.

2. Install the NetWorker client and NMDA according to the instructions in the *EMC NetWorker Module for Databases and Applications Installation Guide*.

Configuring snapshot backups on the mount host (Replication Manager Agent)

You must perform the following configuration steps on the Oracle mount host.

1. Ensure that you have installed the Oracle software:

   - You must install the same Oracle software version on both the production database host and mount host.
   - You can install and use a different ORACLE_HOME for ASM from the one for the database instance host, if you have used this setup on the original database host.

2. If you use ASM, ensure that the Cluster Synch Services (CSS) daemon is running.
3. Ensure that you have installed and configured the Replication Manager Agent software and all the other required software (for example, Solutions Enabler) according to the Replication Manager documentation.

4. (Optional) Set up the RMCLI initialization file to identify the following values:

   - Replication Manager server host and port
   - Username and password for the connection

   If you have not set up the initialization file, the PRECMD and POSTCMD scripts used by NMDA must include the Replication Manager server host, port, and credentials. Refer to information about the scripts in step 8 and step 9.

5. Install the NetWorker client and NMDA software according to the instructions in the *EMC NetWorker Module for Databases and Applications Installation Guide*.

   If you set up the mount host as a NetWorker storage node to avoid network data transfer, you do not need to install the NetWorker client separately.

6. Configure the local Oracle Net listener and Net service names to enable you to connect to the Recovery Catalog database and the replica database after the replica database is mounted.

   The following samples include the network files used to configure the Net listener and service names. In this case, bu-star1 is the host where the RMAN catalog database is running, and bu-star2.lss.emc.com is the mount host:

   - $ORACLE_HOME/network/admin/listener.ora:

     ```
     LISTENER =
       (DESCRIPTION_LIST =
         (DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST =
          bu-star2.lss.emc.com)(PORT = 1521)))
     )
     ```

   - $ORACLE_HOME/network/admin/tnsnames.ora:

     ```
     ORCL110REP =
       (DESCRIPTION =
         (ADDRESS_LIST = (ADDRESS = (PROTOCOL = TCP)(HOST =
          bu-star2)(PORT = 1521)))
         (CONNECT_DATA = (SERVICE_NAME = orcl110))
     )
     CATALOG =
       (DESCRIPTION =
         (ADDRESS_LIST = (ADDRESS = (PROTOCOL = TCP)(HOST =
          bu-star1)(PORT = 1521)))
         (CONNECT_DATA = (SERVICE_NAME = catalog))
     )
     ```

7. Copy the production database password file to the mount host under $ORACLE_HOME/dbs. For example:

   ```
   scp bu-star1:/db/app/oracle/product/11.1.0/db/dbs/orapworcl110
   bu-star2:/db/app/oracle/product/11.1.0/db/dbs/
   ```

8. Create a PRECMD script, which is a preprocessing script with its full pathname specified by the PRECMD parameter in the NMDA configuration file. Appendix A, "NMDA Parameters and Configuration File," describes the NMDA parameters and the NMDA configuration file.

The PRECMD script must perform the following operations:

- Create and mount the database replica by starting the Replication Manager job described in "Configuring snapshot backups on the Replication Manager server" on page 341.
- Mount the ASM instance.
- Convert the control file from current mode to the backup mode, and mount the database according to the instructions in the "RMAN and Split Mirror Disk Backups" Metalink note (302615.1).
- Register the replica instance explicitly to the listener by using **alter system register**. Catalog the archived redo logs generated during hot backup mode if you use online backups.

"Preprocessing script for Oracle ASM backups" on page 350 provides a sample PRECMD script.

9. Create a POSTCMD script, which is a postprocessing script with its full pathname specified by the POSTCMD parameter in the NMDA configuration file.

   The POSTCMD script must unmount the database replica gracefully by using RMCLI.

   "Postprocessing script for Oracle ASM backups" on page 351 provides a sample POSTCMD script.

10. Create an RMAN script to perform the required snapshot backup. In the script, set the NSR_CLIENT parameter to the hostname of the production database host. This NSR_CLIENT setting enables the production database host to access the backup entries generated by the backup on the mount host, for catalog maintenance operations and restore.

    The following sample RMAN script runs an NMDA Oracle manual backup after the "replica" database is mounted. This script backs up the full database to the NetWorker server, bu-rocky, including all the archived redo logs not backed for a day. The production database host is bu-star1. The script connects to an RMAN Recovery Catalog database named catalog:

```
connect target sys/oracle@orcl110rep;
connect catalog rman/rman@catalog;
run {
   allocate channel ch1 type 'SBT_TAPE';
   allocate channel ch2 type 'SBT_TAPE';
   send device type 'SBT_TAPE'
   'NSR_ENV=(NSR_SERVER=bu-rocky, NSR_CLIENT=bu-star1)';
   backup full
   format 'RM_NMDA_DB_%d_%U'
   database
   plus archivelog not backed up since time 'sysdate - 1';
   release channel ch1;
   release channel ch2;
}
```

11. If you will perform scheduled backups, create the NMDA configuration file with required parameter settings, as described in Appendix A, "NMDA Parameters and Configuration File." Ensure that you set the following parameters in the file:

    ```
    PRECMD=full_pathname_of_PRECMD_script
    POSTCMD=full_pathname_of_POSTCMD_script
    ```

    where:

    - *full_pathname_of_PRECMD_script* is the complete pathname of the preprocessing script created in step 8.
    - *full_pathname_of_POSTCMD_script* is the complete pathname of the postprocessing script created in step 9.

Configuring snapshot backups on the NetWorker server host

The *EMC NetWorker Administration Guide* and NMC online help describe how to configure NetWorker resources.

You must perform the required configuration steps on the NetWorker server host:

- For manual snapshot backups, use NMC to create generic Client resources for the production database host and mount host. Ensure that the Remote Access attribute in the Client resource of the production database host contains the following value:

  ```
  Oracle_user@mount_host
  ```

- For scheduled backups, complete the following steps:

  1. Configure the required NetWorker Group and Schedule resources.
  2. Create a NetWorker Client resource of the mount host that contains the following attribute settings:
     - Backup Command—**nsrdasv -z** *NMDA_configuration_file_path*, where the NMDA configuration file includes the parameters set in step 11 of "Configuring snapshot backups on the mount host (Replication Manager Agent)" on page 345.
     - Group—Name of the NetWorker backup group configured in step 1.
     - Save Set—RMAN:*full_path_of_RMAN_script*, where the RMAN script is the script created in step 10 of "Configuring snapshot backups on the mount host (Replication Manager Agent)" on page 345.
  3. Create another NetWorker Client resource for the production database host to back up the parameter file and control file. Ensure that the Remote Access attribute contains the following value:

     ```
     Oracle_user@mount_host
     ```

Performing backups and restores of Oracle ASM data

To perform the snapshot backups and restores of Oracle ASM data, follow the instructions in the appropriate topic:

- "Performing manual snapshot backups of Oracle ASM data" on page 349
- "Performing scheduled snapshot backups of Oracle ASM data" on page 349

- "Performing incremental backups on the mount host" on page 349
- "Performing restore and recovery of Oracle ASM backups" on page 349

Performing manual snapshot backups of Oracle ASM data

1. Log in to the mount host as an Oracle user.
2. Run the PRECMD script, described in step 8 of "Configuring snapshot backups on the mount host (Replication Manager Agent)" on page 345.
3. Run the following command to perform the backup:

   ```
   rman @RMAN_script_name
   ```

4. Run the POSTCMD script, described in step 9 of "Configuring snapshot backups on the mount host (Replication Manager Agent)" on page 345.

Performing scheduled snapshot backups of Oracle ASM data

The NetWorker server automatically starts scheduled backups, based on the specified backup schedule. "Performing scheduled backups" on page 146 describes scheduled backups.

To manually start and test a scheduled backup group, you can use the NMC NetWorker Administration GUI, or the **savegrp** command with the appropriate options.

The *EMC NetWorker Administration Guide* and NMC online help describe the NMC GUI program. The *EMC NetWorker Command Reference Guide* describes the **savegrp** command.

Performing incremental backups on the mount host

You can perform regular RMAN incremental backups on the mount host.

You can also use additional scripting to perform the new incremental backup with change tracking, available in Oracle 10g and known as Block Change Tracking (BCT). The EMC white paper titled "Reducing Backup Window and Recovery Time with Oracle Database 11g RMAN and EMC TimeFinder/Clone" provides details.

Performing restore and recovery of Oracle ASM backups

To restore and recover a snapshot backup of Oracle ASM data, you must perform an RMAN restore with NMDA as described in Chapter 4, "Data Restore and Recovery."

The restore and recovery of the Oracle ASM backups does not use the Replication Manager.

Sample preprocessing and postprocessing scripts

The following topics provide sample preprocessing and postprocessing scripts for the snapshot backups of Oracle ASM data on UNIX or Linux.

Preprocessing script for Oracle ASM backups

The following sample script performs preprocessing tasks on UNIX or Linux before a snapshot backup of Oracle ASM data starts. The main preprocessing script, **precmd**, calls each of the following scripts in turn:

- **create_replica.sh**
- **mountasm.sql**
- **convertcf.sql**
- **catalog.sql**

You must specify the complete pathname of the main preprocessing script with the PRECMD parameter in the NMDA configuration file:

```
$ more precmd

#!/bin/ksh
echo on
# This script creates and mounts the replica database
# by starting the job configured on the RM server.
# If the replica cannot be mounted, the script fails.
DATE=$(date +%Y%m%d)
LOGFILE=/db/home/ora110/rman/rm_rman_backup_$DATE.log
RM_HOME=/opt/emc/rm/gui
ORACLE_SID=orcl110
echo "*** Starting the Oracle database replica for $ORACLE_SID ****" >>
   $LOGFILE
RM_COMMAND_LINE="$RM_HOME/rmcli host=bu-caspian.lss.emc.com port=65432
   login user=Administrator password=password
   file=/db/home/ora110/rman/create_replica.sh >> $LOGFILE "
eval ${RM_COMMAND_LINE} &
Pid=$!
wait $Pid
rm_status=$?
if [ $rm_status != 0 ] ; then
   echo "The RM job returned status of "$rm_status
   echo  $0 "exiting."
   exit 1
fi
# Switch to the Oracle user
su - ora110
# Mount the ASM instance
export ORACLE_SID=+RM0001
eval "sqlplus  /nolog <  /db/home/ora110/rman/mountasm.sql >>$LOGFILE"
   &
# Change the CONTROL FILE type to backup type and mount the "replica"
   database
export ORACLE_SID=orcl110
eval "sqlplus /nolog < /db/home/ora110/rman/convertcf.sql >>$LOGFILE"
   &
# Catalog the archived redo logs generated during hot backup mode
eval "rman @/db/home/ora110/rman/catalog.sql >>$LOGFILE" &
exit 0


$ more /db/home/ora110/rman/create_replica.sh

if  run-job name=orcl110_replica appset=orcl110_appset then exit 0
   else exit 1
```

```
$ more /db/home/ora110/rman/mountasm.sql

connect sys/oracle as sysdba;
startup mount;
exit


$ more /db/home/ora110/rman/convertcf.sql

connect erm/oracle as sysdba;
startup mount;
recover database using backup controlfile until cancel;
CANCEL
shutdown immediate;
startup mount;
alter system register;


$ more /db/home/ora110/rman/catalog.sql

connect target  erm/oracle@orcl110rep;
# catalog the archived redo log generated during hot backup mode
catalog start with '/tmp/orcl110/arch/';
```

Postprocessing script for Oracle ASM backups

The following sample script performs postprocessing tasks on UNIX or Linux after a snapshot backup of Oracle ASM data completes. The main postprocessing script, **postcmd**, calls the **unmount_replica.sh** script.

You must specify the complete pathname of the main postprocessing script with the POSTCMD parameter in the NMDA configuration file:

```
$ more postcmd

#!/bin/ksh
echo on
# This script unmounts the replica database
# If the replica cannot be unmounted, the script fails.
DATE=$(date +%Y%m%d)
LOGFILE=/db/home/ora110/rman/rm_rman_backup_$DATE.log
RM_HOME=/opt/emc/rm/gui
ORACLE_SID=orcl110
echo "*** Unmounting the Oracle database replica for $ORACLE_SID ****"
   >> $LOGFILE
RM_COMMAND_LINE="$RM_HOME/rmcli host=bu-caspian.lss.emc.com port=65432
   login user=Administrator password=password
   file=/db/home/ora110/rman/unmount_replica.sh >> $LOGFILE "
eval ${RM_COMMAND_LINE} &
Pid=$!
wait $Pid
rm_status=$?
if [ $rm_status != 0 ] ; then
       echo "The RM job returned status of "$rm_status
       echo  $0 "exiting."
       exit 1
fi
exit 0

$ more /db/home/ora110/rman/unmount_replica.sh

if unmount-replica position=last appset=orcl110_appset then exit 0
   else exit 1
```

APPENDIX A
NMDA Parameters and Configuration File

This appendix includes the following topics:

- Setting the NMDA parameters 354
- NMDA configuration file 354
- Common NMDA parameters 358
- NMDA DB2 parameters..... 366
- NMDA Informix parameters 370
- NMDA Lotus parameters..... 371
- NMDA MySQL parameters 380
- NMDA Oracle parameters 387
- NMDA Sybase parameters..... 392

Setting the NMDA parameters

This appendix describes the parameters that you can set for the NMDA backups and restores.

Note: Unless noted otherwise, NMDA supports the parameters for both nonsnapshot backups and restores and snapshot backups and restores. NMDA supports snapshot operations with NSM for DB2 and Oracle only.

If you perform client-side configuration (without the configuration wizard), then you must set specific parameters for an NMDA scheduled or manual backup or an NMDA restore, typically in the NMDA configuration file.

"NMDA configuration file" on page 354 describes the NMDA configuration file and the exceptions to setting the parameters in the file.

"Common NMDA parameters" on page 358 describes the parameters that NMDA uses for backups and restores of all the supported databases and applications.

The following topics describe the parameters that NMDA uses for backups and restores of specific databases and specific applications only:

- "NMDA DB2 parameters" on page 366
- "NMDA Informix parameters" on page 370
- "NMDA Lotus parameters" on page 371
- "NMDA MySQL parameters" on page 380
- "NMDA Oracle parameters" on page 387
- "NMDA Sybase parameters" on page 392

NMDA configuration file

As part of a client-side configuration for backup or restore operations, you must typically set the required NMDA parameters in the NMDA configuration file.

The only exceptions to setting the NMDA parameters in the configuration file are as follows:

- For Informix manual backups and restores, set parameters in the environment.
- For Oracle backups and restores, set certain parameters in the RMAN script as described in "NMDA Oracle parameters" on page 387.
- For Sybase restores, set parameters as **nsrsybrc** command options, if possible. If you cannot set a parameter as a **nsrsybrc** command option, then set the parameter in the environment.

NMDA includes the following templates for the NMDA configuration file:

- nmda_db2.cfg—Template for DB2 parameters
- nmda_informix.cfg—Template for Informix parameters
- nmda_lotus.cfg—Template for Lotus parameters
- nmda_mysql_backup.cfg—Template for MySQL backup parameters
- nmda_mysql_restore.cfg—Template for MySQL restore parameters
- nmda_oracle.cfg—Template for Oracle parameters
- nmda_sybase.cfg—Template for Sybase parameters

The configuration file templates are in the following directory:

- On UNIX: /nsr/apps/config
- On Windows: *NetWorker_install_path*\apps\config, where *NetWorker_install_path* is the root directory of the NetWorker installation path, for example, C:\Program Files\Legato\nsr or C:\Program Files\EMC NetWorker\nsr

Make a copy of the required configuration file templates, for example, in the original directory or an alternate location. The uninstall of NMDA software removes the original templates.

To create the NMDA configuration file for backups and restores of a database or application:

1. Copy the appropriate template file to any location on the client host.
2. Customize the parameter settings in the file.

You can name the configuration file with any preferred name. You must specify the configuration file pathname with the **-z** option in the appropriate backup or restore command.

NOTICE

Ensure that the NMDA configuration file is in a secure location and is accessible by privileged users only. If possible, make this file readable and writeable by the administrative or application user that performs the operation.

Syntax rules for the NMDA configuration file

The NMDA configuration file must conform to the following syntax rules:

- Each parameter setting must be in one of the following formats:

  ```
  NAME = value
  NAME = value1, value2, value3
  ```

  where:

  - *NAME* is the parameter name.
  - *value*, *value1*, *value2*, *value3* are the assigned parameter values.

- Parameter names and values are case-sensitive, unless specified otherwise in this appendix.
- You can optionally group parameter settings within braces as follows:

  ```
  keyword {
  ...parameter_settings...
  }
  ```

  where *keyword* is one of the following case-insensitive keywords, to signify that the parameter settings apply to a particular type of database or application:

  - DB2
  - INFORMIX
  - LOTUS
  - LOTUS_DAOS
  - MYSQL

- ORACLE
- SYBASE

Note: You must use the LOTUS_DAOS keyword only to set parameters for a Lotus DAOS backup. The LOTUS_DAOS{} section must appear after the LOTUS{} section in the same configuration file. "Configuring integrated Lotus DAOS backups" on page 122 describes parameters in the LOTUS_DAOS{} section.

- If you back up multiple database or server instances in a single backup configuration, specify the parameter settings in the configuration file as follows:
  - For parameter settings common to all the instances, specify the parameter settings outside of the braces.
  - For parameter settings unique to specific instances on the NMDA client, group the parameter settings within braces.

  The following example shows the correct positions of parameter settings in the configuration file:

```
# Global parameters common to all the instances
parameter1 = value
parameter2 = value
parameter3 = value

DBSID1 {
# Parameter settings for DBSID1
parameter4 = value1
parameter5 = value
}

DBSID2 {
# Parameter settings for DBSID2
parameter4 = value2
}
```

  where *DBSID1* and *DBSID2* are the instance-specific values that appear on separate lines in the Save Set attribute of the Client resource for the scheduled backup. These values do not include the prefix DB2:, INFORMIX:, NOTES:, MYSQL:, RMAN:, or SYBASE:.

  A global parameter appears outside of all the braces. An instance-specific parameter appears within the braces after a particular instance value, such as *DBSID1*.

  Note: If you group parameters in braces for different instances, do not use the DB2, INFORMIX, LOTUS, MYSQL, ORACLE, or SYBASE keyword in the same configuration file. However, you can use the LOTUS_DAOS keyword for any Lotus DAOS settings.

- The following precedence rules apply to the parameter settings:
  - Global parameter settings apply to all the database or server instances included in the backup.
  - If a global parameter is set more than once, the last setting takes precedence over all the previous settings of the global parameter.

- Instance-specific parameter settings within braces apply only to the particular database or server instance and override the corresponding global parameter settings for the backup of that specific instance.
- If a parameter appears more than once within the same braces, the last setting takes precedence over all the previous settings of the parameter within the braces.

◆ The only supported type of parameter nesting is the nesting of a LOTUS_DAOS{} parameter group inside a group of Lotus server-specific parameters. For example, in a Lotus backup configuration, the Save Set attribute of the Client resource contains two server-specific save sets:

```
NOTES:server1
NOTES:server2
```

The corresponding NMDA configuration file contains the following nested groups of parameters:

```
server1 {
   NSR_BACKUP_PATHS = /lotus/path
   LOTUS_DAOS {
      NSR_BACKUP_PATHS = /lotus/daos/path
   }
}
server2 {
   NSR_BACKUP_PATHS = /another/lotus/path
   LOTUS_DAOS {
      NSR_BACKUP_PATHS=/another/lotus/daos/path
   }
}
```

The database backup of save set NOTES:server1 uses the first NSR_BACKUP_PATHS setting, /lotus/path. The Lotus DAOS phase of that backup uses the NSR_BACKUP_PATHS setting in the first LOTUS_DAOS group, /lotus/daos/path.

◆ Separate multiple values for a parameter with commas.

◆ You can specify the values of a parameter over multiple lines if each line ends in a comma. For example:

```
NAME = value1,
       value2,
       value3
```

◆ If the line specifying a parameter does not end in a comma, the next line must contain a new parameter setting.

◆ Use white space as preferred. NMDA ignores all the white space.

◆ Text on a line after the # symbol (where # is not enclosed in quotes) is a comment. NMDA ignores all the comments. You can place a comment after a parameter setting on the same line.

Common NMDA parameters

The following table describes the common NMDA parameters, which NMDA uses for backups and restores of all the supported databases and applications.

Table 33 Common NMDA parameters (page 1 of 9)

Parameter	Description	Default and valid values
NSR_AES_ENCRYPTION	Specifies whether the NetWorker server encrypts the backup data through 256-bit AES encryption, which uses the key or pass phrase set in the Datazone pass phrase attribute of the NetWorker Server resource. Optional for a backup. Be careful when you change the pass phrase on the NetWorker server. If the pass phrase on the server changes and you cannot remember the pass phrase originally used for an NMDA backup, you cannot recover the encrypted data. The *EMC NetWorker Administration Guide* provides details on pass phrases. **Note:** Record each key (pass phrase) used for 256-bit AES encryption. You need the key to restore the backup later.	• FALSE (default) = Do not encrypt data with 256-bit AES encryption during the backup. • TRUE = Encrypt data with 256-bit AES encryption during the backup.
NSR_CHECKSUM	Specifies whether the NetWorker software performs checksumming on the backup data. Optional for a backup. **Note:** When you restore an NMDA Sybase backup created with NSR_CHECKSUM set to TRUE, a "CRC mismatch" error during the restore indicates that the restored data does not match the backed-up data due to backup corruption. The error message does not appear on the screen as expected. The error message appears only in the nmda_sybase.messages.raw file and the Sybase Backup server error log.	• FALSE (default) = NetWorker software does not perform checksumming. • TRUE = NetWorker software performs checksumming.
NSR_CLIENT	Specifies the NetWorker Client resource to use for a backup or restore. Recommended for a backup or restore in a cluster, DB2 DPF, Informix MACH, Oracle RAC, or Sybase ASE Cluster Edition system, and for a redirected restore to a different host. Chapter 6, "Cluster Systems" provides details. For an Oracle client-side scheduled backup, set this parameter in the NMDA configuration file.	• Hostname of the physical host on which the session runs (default). • Valid NetWorker client hostname.

Table 33 Common NMDA parameters (page 2 of 9)

Parameter	Description	Default and valid values
NSR_COMPRESSION	Specifies whether the NetWorker software performs compression on the backup data. NMDA supports only the default NetWorker encryption algorithm. NMDA does not support backup compression with GZIP or BZIP2. Optional for a backup. **Note:** If you use database or application backup compression, do not set this NMDA parameter. There is no benefit to running NMDA compression on already compressed data.	• FALSE (default) = NetWorker software does not perform compression. • TRUE = NetWorker software performs compression.
NSR_DATA_DOMAIN_ INTERFACE	Specifies the network interface to use to send backup data to the DD Boost device. Optional for a manual backup only. Set this parameter if you have a Fibre Channel (FC) connection to the DD Boost device. You must set this parameter together with NSR_DEVICE_INTERFACE=DATA_DOMAIN. If you do not set NSR_DEVICE_INTERFACE to DATA_DOMAIN, then this parameter is ignored.	• IP (default) = Backup data is sent over an IP network to the DD Boost device. • Any = Backup data is sent over either an IP or FC network to the DD Boost device, depending on the available device. • Fibre Channel = Backup data is sent over an FC network to the DD Boost device.
NSR_DATA_VOLUME_POOL	Specifies the name of the NetWorker volume pool to use for the backup. For snapshot backups (DB2 and Oracle only), specifies the volume pool for backups to secondary storage only. Optional for a DB2, Informix, Lotus, MySQL, or Sybase manual backup. Mandatory for an Oracle manual backup that uses the **set duplex** command (with **duplex** set to 1, 2, 3, or 4) or other RMAN commands to generate backup copies. For an Oracle manual backup that generates backup copies, set this parameter with the **parms** option in the RMAN script, not with the **send** command or **send** option. **Note:** For a scheduled backup, this parameter overrides any pool associated with the scheduled backup group if you do not set the Pool attribute of the Client resource. The backup ignores the parameter if you set the Pool attribute.	• Most appropriate pool as selected by the NetWorker server (default). • Valid name of a NetWorker volume pool. • For a manual Oracle backup, the name must be different from the name used by the parameter NSR_DATA_VOLUME_POOL1, NSR_DATA_VOLUME_POOL2, or NSR_DATA_VOLUME_POOL3.

Table 33 Common NMDA parameters (page 3 of 9)

Parameter	Description	Default and valid values
NSR_DEBUG_LEVEL	Specifies the level of debug messages that NMDA writes to the debug log file, located in the directory specified by NSR_DIAGNOSTIC_DEST or in the default directory, /nsr/apps/logs (UNIX) or *NetWorker_install_path*\apps\logs (Windows). **Note:** Use this parameter for debugging purposes with assistance from EMC Customer Support only. Optional for a backup or restore. For an Oracle operation, set this parameter either in the configuration file or with the **parms** option in the RMAN script, not with the **send** command or **send** option.	• 0 (default) = NMDA does not generate debug messages. • 1 to 9 = NMDA writes debug messages to the debug log file with a .log filename extension. The level of detail in the generated debug messages increases with the debug level.
NSR_DEDUP_BACKUP	Specifies whether a manual NMDA backup performs Avamar deduplication. Mandatory for a manual Avamar deduplication backup. Do not set this parameter for a scheduled Avamar deduplication backup. **Note:** For a scheduled Avamar deduplication backup, set the De-duplication Backup attribute in the NetWorker Client resource, instead of setting this parameter.	• FALSE (default) = Manual NMDA backup does not perform Avamar deduplication. The backup ignores other NSR_DEDUP* parameters. • TRUE = Manual NMDA backup performs Avamar deduplication. If you do not set NSR_DEDUP_NODE, the backup fails.
NSR_DEDUP_CACHE_ENABLED	Specifies whether an Avamar deduplication backup uses a hash cache. Optional for an Avamar deduplication backup. The **nsravtar** process creates the cache in the /nsr/dedup/cache directory or the equivalent directory on Windows. Use of the cache increases both the Avamar deduplication backup performance and the disk usage in the cache directory. **Note:** Setting this parameter requires knowledge of the potential effects on Avamar server operations.	• TRUE (default) = An Avamar deduplication backup uses a hash cache to increase performance. Use this value, if possible. • FALSE = An Avamar deduplication backup does not use a hash cache. The backup ignores the NSR_DEDUP_CACHE_TAG parameter.

Table 33 Common NMDA parameters (page 4 of 9)

Parameter	Description	Default and valid values
NSR_DEDUP_CACHE_TAG	Specifies the tag used to generate the hash cache name for an Avamar deduplication backup. Mandatory for Informix parallel deduplication backups when you enable caching. Mandatory for an Oracle deduplication backup configured without the wizard. Optional for a DB2, Lotus, MySQL, or Sybase deduplication backup. For an Oracle deduplication backup: • If you use multiple channels on Windows, set this parameter with the **send** command, not the **parms** option, in the RMAN script. • If you use automatic channel allocation (and you do not use multiple channels on Windows), set this parameter with the **parms** option, not the **send** command, in the RMAN script. • Set this parameter to a different value for each channel; if you use the same tag value for more than one channel, the deduplication backup fails.	• Application-specific default value: - For DB2 (nonsnapshot backup): ‹*DB2_node_name*›_‹*backup_session_number*›, where the DB2 server supplies the backup session number. - For DB2 (snapshot backup): / (forward slash). - For Informix: ‹*XBSA objectName.pathName*›, supplied by the IDS server. - For Lotus: Same value as the save set name. - For MySQL: / (forward slash) - For Oracle: / (forward slash); all RMAN channels (concurrent backup processes will have the same cache tag, which will cause the backup to fail). - For Sybase: ‹*database*›_‹*stripe_number*› if a database is provided for backup; otherwise, ‹*server*›_‹*stripe_number*›. Default value is the recommended value for all application backups except Informix parallel backups and Oracle backups. • String value of the tag that will generate a deduplication cache name. Do not include the client name in the value. Set the parameter to a different value for each concurrent NMDA process. • Recommended value for an Oracle backup: ‹*ORACLE_SID or Net_service_name*›_‹*channel_ID*› For example: ORCL102_t1 **Note:** The software uses the tag value to generate the cache name through hashing. The cache name does not contain this parameter value.
NSR_DEDUP_CHUNK_SIZE	Specifies the size in bytes that the Avamar server uses for data chunks in a deduplication backup. If you specify a nonzero value, the Avamar server uses the fixed size for all the data chunks saved in the deduplication backup. Optional for an Avamar deduplication backup. **Note:** Setting this parameter requires knowledge of the potential effects on Avamar server operations.	• 0 (default) = Avamar server uses variable sizes for the data chunks. Specify this value if possible. • Size (greater than zero) in bytes = Avamar server uses the size for all the data chunks in a deduplication backup, for example, 1024, 2048, 5096, 8194, or a value recommended in the Avamar documentation.

Table 33 Common NMDA parameters (page 5 of 9)

Parameter	Description	Default and valid values
NSR_DEDUP_NODE	Specifies the hostname of the Avamar server that will store the deduplicated client data. The hostname must be the same as the Avamar server hostname set in the NetWorker De-duplication Node resource. Mandatory for a manual Avamar deduplication backup. Do not set this parameter for a scheduled Avamar deduplication backup. **Note:** For a scheduled Avamar deduplication backup, set the De-duplication Node attribute in the NetWorker Client resource instead of setting this parameter.	• Undefined (default). • Avamar server hostname set in the NetWorker De-duplication Node resource.
NSR_DEVICE_INTERFACE **Note:** This parameter is deprecated. It is supported in the current NMDA release, but will be unsupported in a future release.	Specifies whether to always store the backup on the Data Domain device from the specified pool. Optional for manual deduplication backups with the Data Domain media type only. Set this parameter when the backup pool contains a mixture of Data Domain devices and other types of devices, for example, AFTD, tape, and so on. Do not set for scheduled deduplication backups. For a scheduled deduplication backup, set the Data Domain Backup attribute in the NMDA Client resource when the pool has mixed devices.	• Undefined (default) = Backup data is stored on the most appropriate device from the pool selected by NetWorker. • DATA_DOMAIN = Backup data is always stored on the Data Domain device.
NSR_DIAGNOSTIC_DEST	Specifies the directory location of the NMDA debug logs except the configuration wizard debug logs. Optional for a backup or restore. **Note:** You cannot relocate the operational and error logs (nmda_*.messages.raw) and the wizard debug logs. For an Oracle operation, set this parameter either in the configuration file or with the **parms** option in the RMAN script (not with the **send** command or **send** option).	• By default, NMDA generates the debug log files in the directory /nsr/apps/logs or *NetWorker_install_path*\apps\logs. • Valid pathname of the directory where NMDA generates the debug logs except the wizard debug logs.
NSR_DIRECT_ACCESS	Specifies the method used to perform a deduplication backup to a Data Domain device. Restores ignore this parameter. "Deduplication backups and restores with EMC Data Domain" on page 30 describes the restore behavior.	• Default (default) = Data Domain backup tries to use the Client Direct method that deduplicates data on the NMDA host and sends unique blocks directly to the Data Domain device. If the software cannot use the Client Direct method, the software deduplicates the backup on the storage node. • No = Data Domain backup does not try to use the Client Direct method. The software deduplicates the Data Domain backup on the storage node or the Data Domain device. • Yes = Data Domain backup uses only the Client Direct method. If the software cannot use this method, the backup fails.

Table 33 Common NMDA parameters (page 6 of 9)

Parameter	Description	Default and valid values
NSR_DPRINTF	Specifies whether NetWorker core debug messages are written to the NMDA debug log files, as described for "NSR_DEBUG_LEVEL" on page 360. Optional for a backup or restore. Used for debugging purposes with assistance from EMC Customer Support only. For an Oracle operation, set this parameter either in the configuration file or with the **parms** option in the RMAN script, not with the **send** command or **send** option.	• FALSE (default) = NetWorker core debug messages are not written to the NMDA debug log files. • TRUE = NetWorker core debug messages are written to the NMDA debug log files.
NSR_ENCRYPTION_PHRASES	Specifies one or more encryption phrases for decrypting data during an NMDA restore. If you do not set this parameter, the NMDA restore obtains the encryption phrase from the NetWorker server. Optional for a restore. If both of the following conditions are true, set this parameter to the phrase used to originally back up the data: • NMDA backed up the data with 256-bit AES encryption. • The encryption phrase on the NetWorker server has changed since the backup. For an Oracle restore, set this parameter with the **send** command in the RMAN script.	• By default, an NMDA restore obtains the encryption phrase from the Datazone pass phrase attribute of the NetWorker Server resource, as described in "Datazone pass phrase" on page 75. • One or more encryption phrases to use during an NMDA restore. Each phrase must be a string enclosed in quotes. Separate multiple phrases with commas. For Oracle restores only, surround the entire group of phrases with outer quotes that are different from the inner quotes. For example, this parameter is for an Oracle restore: NSR_ENCRYPTION_PHRASES="'*key1*','*key2*'" - NMDA itself supports double ("), single ('), and backward (`) quotes. - Certain shells, databases, or applications might not support certain types of quotes.
NSR_GROUP	Specifies the name of a NetWorker backup group to use for a manual backup. Optional for a manual backup. Do not set this parameter for a scheduled backup. This parameter can be used to enable the following operations: • Test a backup manually on the client side before starting scheduled backups. • Direct a manual backup to a NetWorker device pool, based on the backup group criteria.	• Default (default). • Valid name of a NetWorker backup group.
NSR_MAX_START_RETRIES	Specifies how many times NMDA tries to connect to the NetWorker server before the operation fails. NMDA waits for 30 seconds between each try to connect. For example, the connection to the NetWorker server might fail for one of the following reasons: • The NetWorker server is not ready because the devices are not mounted. • The **nsrindexd** service of the NetWorker server is busy due to other client sessions. Optional for a backup or restore.	• 4 (default). • Integer number of tries to connect to the NetWorker server.

Table 33 Common NMDA parameters (page 7 of 9)

Parameter	Description	Default and valid values
NSR_NO_BUSY_ERRORS	Specifies whether a backup or restore fails immediately when the NetWorker server is busy or waits for the NetWorker server to accept the connection. Optional for a backup or restore. **Note:** This parameter has no effect for snapshot backups or restores. NSR_NO_BUSY_ERRORS is a deprecated parameter that will be unsupported in a future NMDA release. Use NSR_MAX_START_RETRIES instead.	• FALSE (default) = The backup or restore waits for the NetWorker server to accept the connection. • TRUE = The backup or restore fails immediately when the NetWorker server is busy.
NSR_NWPATH	Specifies the pathname of the directory that contains the NetWorker binaries. Mandatory for the following Avamar deduplication backups: • Sybase deduplication backup on 64-bit Windows • Any deduplication backup with 32-bit NMDA installed on 64-bit Windows Mandatory for MySQL or Oracle backup deletions if the NetWorker client binaries are in a nondefault directory on the server host. Recommended for Oracle operations on Windows if multiple database instances exist on the same host. Recommended for an NSM snapshot or Avamar deduplication backup or restore if either of the following conditions are true: • NSM (**nsrsnapck** binary) is in a nondefault location. • NetWorker client software (**nsravtar** binary) is in a nondefault location. **Note:** You cannot use NSR_NWPATH for deduplication backups or restores with NSM.	• System-specific default location of the NetWorker client binaries (default). • For an Avamar deduplication backup performed on 64-bit Windows, valid directory pathname of the NetWorker client binaries. • If NSM (**nsrsnapck** binary) is in a nondefault location, valid directory pathname of the NSM binary. • If NetWorker client software (**nsravtar** binary) is in a nondefault location, valid directory pathname of the **nsravtar** binary.
NSR_PHYSICAL_HOST_LICENSE	Specifies whether NMDA is licensed per cluster node or per virtual cluster name in an active-passive cluster environment with NetWorker server 7.6 SP1 or later. Optional in an active-passive cluster.	• FALSE (default) = NMDA is licensed per virtual cluster name in the active-passive cluster. • TRUE = NMDA is licensed per cluster node in the active-passive cluster.
NSR_RECOVER_POOL	Specifies the name of the NetWorker volume pool to use for a restore. You can use this option to restore data from a specified volume pool if there are multiple copies (clones) of the backup on different volume pools. Optional for a restore. Supported for a nonsnapshot restore only, not a snapshot restore.	• Pool determined by the NetWorker server (default). • Valid name of a NetWorker volume pool containing a cloned backup to use for a restore.

Table 33 Common NMDA parameters (page 8 of 9)

Parameter	Description	Default and valid values
NSR_SAVESET_BROWSE	Specifies the browse policy of a backup, as the date when the entry for the backup is removed from the NetWorker client index. Optional for a manual backup. **Note:** This parameter overrides the browse policy specified in the NetWorker Client resource.	• Browse policy specified in the NetWorker Client resource for the client (default). • Valid date in **nsr_getdate(3)** format.
NSR_SAVESET_RETENTION	Specifies the retention policy of a backup, as the date when the save set becomes recyclable. Optional for a manual backup. **Note:** This parameter overrides the retention policy specified in the NetWorker Client resource.	• Retention policy specified in the NetWorker Client resource for the client (default). • Valid date in **nsr_getdate(3)** format.
NSR_SERVER	Specifies the hostname of the NetWorker server to perform the backup or restore. Mandatory for a manual backup or restore if the NetWorker server host is different from the client host. For an Oracle operation: • If the operation generates backup copies, set this parameter with the **parms** option in the RMAN script, not with the **send** command or **send** option. • If the operation does not generate backup copies, set this parameter with the **send** command or **send** option in the RMAN script.	• Hostname of the NetWorker server detected on the client host (default). • Valid hostname of a NetWorker server.
POSTCMD	Specifies a postprocessing script to run after a scheduled backup: • The postprocessing script file might need permissions that enable execution by the root user. • The script must return a zero value when it succeeds, and a nonzero value when it fails. • On UNIX, the first line of the script must contain the following interpreter directive: #!/bin/sh Optional for a scheduled backup. Do not set this parameter for a manual backup. **Note:** If the scheduled backup fails, the postprocessing script still runs. If the postprocessing script fails, an error message appears but the scheduled backup succeeds.	• Undefined (default). • Valid pathname of a postprocessing script file. The pathname must not contain any spaces. For example, instead of setting POSTCMD to C:\Program Files\Legato\nsr\postcmd.bat, set the parameter to C:\Progra~1\Legato\nsr\postcmd.bat. Also, instead of setting POSTCMD to C:\Program Files\EMC NetWorker\nsr\postcmd.bat, set the parameter to C:\Progra~1\EMC NetWorker\nsr\postcmd.bat. • If the value is undefined or invalid, a postprocessing script does not run after the scheduled backup.

Table 33 Common NMDA parameters (page 9 of 9)

Parameter	Description	Default and valid values
PRECMD	Specifies a preprocessing script to run before a scheduled backup: • The preprocessing script file might need permissions that enable execution by the root user. • The script must return a zero value when it succeeds, and a nonzero value when it fails. The return of a nonzero value causes the scheduled backup to fail. • On UNIX, the first line of the script must contain the following interpreter directive: #!/bin/sh Optional for a scheduled backup. Do not set this parameter for a manual backup. **Note:** If the preprocessing script fails, NMDA does not perform the scheduled backup, an error message appears, and any postprocessing script does not run.	• Undefined (default). • Valid pathname of a preprocessing script file. The pathname must not contain any spaces. For example, instead of setting PRECMD to C:\Program Files\Legato\nsr\precmd.bat, set the parameter to C:\Progra~1\Legato\nsr\precmd.bat. Also, instead of setting PRECMD to C:\Program Files\EMC NetWorker\nsr\precmd.bat, set the parameter to C:\Progra~1\EMC NetWorker\nsr\precmd.bat. • If the value is undefined or invalid, a preprocessing script does not run before the scheduled backup.

NMDA DB2 parameters

For the following DB2 operations, you must set both the common parameters and the NMDA DB2 parameters in the NMDA configuration file:

- DB2 scheduled backups configured without the wizard (client-side configuration)
- DB2 manual backups
- DB2 restores

"Common NMDA parameters" on page 358 describes the common parameters.

The following table describes the NMDA DB2 parameters.

Set the parameters for DB2 transaction log backups or DB2 transaction log restores in a separate configuration file, for example, nmda_db2_tlogs.cfg, as specified in the LOGARCHOPT1 setting. "Configuring automatic backups of DB2 transaction logs" on page 116 provides details.

Table 34 NMDA DB2 parameters (page 1 of 5)

Parameter	Description	Default and valid values
DB2_ALIAS	Specifies the database alias of the DB2 database to back up. This alias is typically the same as the database name. If not set, NMDA derives this parameter and the partition number from the DB2 save set name. The DB2 save set retains its current form of DB2:/DB_NAME/NODE*XXXX*, where *XXXX* is the partition number. Optional for a DB2 scheduled backup only.	• Undefined (default). • Valid DB2 Alias Name.

Table 34 NMDA DB2 parameters (page 2 of 5)

Parameter	Description	Default and valid values
DB2_APPLY_NW_LEVELS	Specifies whether NMDA uses the backup levels specified in either the NMDA configuration file (for example, nmda_db2.cfg) or in the NetWorker Schedule resource. Optional for a DB2 scheduled backup only.	• FALSE (default) = NMDA uses the backup level set in the configuration file and maps the backup level to the NetWorker schedule as follows: - DB2BACKUP_FULL maps to the full level in the schedule. - DB2BACKUP_INCREMENTAL maps to the incr level in the schedule. - DB2BACKUP_DELTA maps to the level set in the NetWorker schedule. If the schedule does not specify a level from 1 to 9, then NMDA applies level 9. • TRUE = NMDA uses the backup level defined in the NetWorker schedule. NMDA ignores the backup level in the configuration file. NMDA maps the levels in the NetWorker schedule to the DB2 levels as follows: - The full level maps to DB2BACKUP_FULL. - The incr level maps to DB2BACKUP_INCREMENTAL. - The levels from 1 to 9 map to DB2BACKUP_DELTA.
DB2_BUFFER_SIZE	Specifies the DB2 backup buffer size in page units. This parameter corresponds to the **buffer** *buffer_size* option of the DB2 **backup** command. Optional for a DB2 scheduled backup only. DB2 automatically uses an optimal value for this parameter. However, when you perform a deduplication backup, you might improve the deduplication ratio by explicitly setting a larger buffer size. The DB2 documentation provides details about DB2 tuning parameters for deduplication devices.	• Optimal value set by DB2 (default). • Number of page units for the DB2 backup buffer size. The IBM documentation provides recommended settings.
DB2INSTANCE	Specifies the name (not the alias) of the DB2 instance that contains the database to be backed up, or the DB2 instance that contains the logs to be retrieved during a rollforward operation after a redirected recovery. For a DB2 scheduled backup, mandatory on UNIX only. For a DB2 rollforward operation after a redirected recovery, mandatory on all platforms. Set this parameter in the log configuration file specified by the DB2 database configuration parameter LOGARCHOPT1. For example, the following command specifies the file: **db2 update db cfg for** *database* **using LOGARCHMETH1 vendor:c:\NetWorker\nsr\bin\libnsrdb2.dll LOGARCHOPT1 @***log_config_file* **Note:** Ensure that this parameter is set correctly. The appropriate IBM DB2 documentation provides details.	• Undefined (default). • Valid name of the DB2 instance that contains the database, or the DB2 instance that contains the logs.

Table 34 NMDA DB2 parameters (page 3 of 5)

Parameter	Description	Default and valid values
DB2_NODE_NAME	Specifies the alias of the DB2 instance to which the user must connect for the backup. Mandatory for a DB2 scheduled backup only.	• Undefined (default). • Valid alias of the DB2 instance. If the node you are usiing is through a local connection, specify the instance name.
DB2_NUM_BUFFERS	Specifies the number of backup buffers used by DB2. This parameter corresponds to the **with** *num_buffers* **buffer** option of the DB2 **backup** command. Optional for a DB2 scheduled backup only. DB2 automatically uses an optimal value for this parameter.	• Optimal value set by DB2 (default). • Number of DB2 backup buffers. The IBM documentation provides more information.
DB2_OPTIONS	Specifies the DB2 backup options. Optional for a DB2 scheduled backup only. **Note:** At a minimum, specify either DB2BACKUP_DB or DB2BACKUP_TABLESPACE, not both. You cannot specify DB2BACKUP_DELTA or DB2BACKUP_INCREMENTAL in a DB2 pureScale system.	• Undefined (default). • One or more of the following values (case-sensitive), with multiple values separated by commas: - DB2BACKUP_COMPRESS - DB2BACKUP_DB - DB2BACKUP_DEDUP_DEVICE - DB2BACKUP_DELTA - DB2BACKUP_EXCLUDE_LOGS - DB2BACKUP_FULL - DB2BACKUP_INCLUDE_LOGS - DB2BACKUP_INCREMENTAL - DB2BACKUP_OFFLINE - DB2BACKUP_ONLINE - DB2BACKUP_TABLESPACE
DB2_PARALLELISM	Specifies the number of tablespaces that can be read in parallel during a DB2 backup. This parameter corresponds to the **parallelism** option of the DB2 **backup** command. Optional for a DB2 scheduled backup only. DB2 automatically uses an optimal value for this parameter.	• Optimal value set by DB2 (default). • Number of the maximum concurrent tablespaces read in parallel during a backup. The IBM documentation provides more information.
DB2_PARTITION_LIST	Specifies which nodes to back up for a DPF backup. Optional for a DB2 DPF backup only.	• Undefined (default). If not specified, the backup backs up a single node only. • Value “all” or any integer that specifies an individual node to back up. Use commas to separate multiple integers.
DB2PATH (Windows)	Specifies the path of the DB2 binaries location. Mandatory for a DB2 scheduled backup on Windows systems only.	• Undefined (default). • Valid path of the location of the DB2 binaries used for the backup.
DB2_QUIESCE	Specifies whether to quiesce the DB2 database during a backup. Optional for a DB2 scheduled backup only.	• FALSE (default) = Do not quiesce the DB2 database during a backup. • TRUE = Quiesce the DB2 database during a backup. **Note:** You must also set the DB2_ALIAS parameter.

Table 34 NMDA DB2 parameters (page 4 of 5)

Parameter	Description	Default and valid values
DB2_SESSIONS	Specifies the number of parallel NMDA sessions to run with the NetWorker server for the DB2 backup. Optional for a DB2 scheduled backup only.	• 1 (default). • Integer number of parallel NMDA sessions.
DB2_TBS_LIST	Specifies a list of tablespaces to back up. Mandatory for a DB2 scheduled backup only. Set this parameter for a tablespace backup only, not for a database backup.	• Undefined (default). • List of tablespace names, with multiple names separated by commas. For example: DB2_TBS_LIST = SYSCATSPACE, USERSPACE1
DB2_USER	Specifies the name of the DB2 user that connects to the DB2 instance for the backup. You must specify the user password with the USER_PSWD parameter. Mandatory for a DB2 scheduled backup only.	• Undefined (default). • Valid DB2 username.
DB2_VENDOR_LIB_PATH	Specifies the complete pathname of the NMDA shared library on the DB2 host. The path can point to various library versions to test and evaluate hotfixes. Optional for a DB2 scheduled backup only.	• On UNIX systems, if you do not specify a path, NMDA assumes the default location. For example, on Solaris systems: /usr/lib/libnsrdb2.so • On Windows systems, NMDA obtains the default path automatically from the registry.
INSTHOME (UNIX)	Specifies the path of the DB2 binaries location. Mandatory for a DB2 scheduled backup on UNIX only.	• Undefined (default). • Valid path of the location of the DB2 binaries used for the backup
NSR_DR_BACKUP_INFO	Specifies whether to back up additional disaster recovery information and support information along with a scheduled backup. If the backup of additional information fails, an error message appears but the backup report is successful. Optional for a DB2 scheduled backup only. **Note:** If you enable Avamar deduplication, the software does not deduplicate these additional files.	• TRUE (default) = Back up additional disaster recovery information and support information as detailed in "Preparing for DB2 disaster recovery" on page 227. • FALSE = Do not back up additional disaster recovery information and support information. **Note:** Set the following additional parameters: - DB2_ALIAS - DB2PATH (Windows) - INSTHOME (UNIX)
NSR_DR_FILE_LIST	Specifies a file that contains a list of files to back up in addition to the database backup. NMDA backs up the files as part of the scheduled backup, before any postprocessing script defined by POSTCMD runs. If the backup of extra files fails, an error message appears. Optional for a DB2 scheduled backup only.	• Undefined (default). • Valid complete pathname of the file that contains the list of extra files to back up. For example, the NSR_DR_FILE_LIST value is nmda_savelist.txt, which is a file that contains this list: /space12/vendor.cfg /space12/db2inst1/sqllib/db2nodes.cfg
NSR_LOG_VOLUME_POOL	Specifies the volume pool to use for a backup of the transaction logs. Optional for a DB2 backup.	• Most appropriate pool as selected by the NetWorker server (default). • Valid name of a NetWorker volume pool for the transaction logs.

Table 34 NMDA DB2 parameters (page 5 of 5)

Parameter	Description	Default and valid values
USER_PSWD	Specifies the encrypted password for the DB2 user that connects to the DB2 instance, as specified by the DB2_USER parameter. Mandatory for a DB2 scheduled backup only.	• Undefined (default). • Encrypted DB2 user password that you must set with the **nsrdaadmin -P** command, for example: **nsrdaadmin -P -z** *configuration_file_path* The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrdaadmin** command.

NMDA Informix parameters

For Informix scheduled backups configured without the wizard (client-side configuration), you must set both the common parameters and the NMDA Informix parameters in the NMDA configuration file.

For Informix manual backups and restores, you must set the parameters in the environment.

"Common NMDA parameters" on page 358 describes the common parameters.

The following table describes the NMDA Informix parameters.

Table 35 NMDA Informix parameters (page 1 of 2)

Parameter	Description	Default and valid values
DO_LOGFILE_BACKUPS	Specifies whether to perform the logical log file backup after the dbspace backup. Optional for an Informix scheduled backup only. **Note:** If you set DO_WHOLE_SYSTEM_BACKUP to TRUE, NMDA ignores DO_LOGFILE_BACKUPS.	• TRUE (default) = Perform the logical log file backup after the dbspace backup. • FALSE = Do not perform the logical log file backup after the dbspace backup. **Note:** If you set DO_LOGFILE_BACKUPS to TRUE, the **onbar -b -l -c** command runs after the backups of any Informix dbspaces.
DO_WHOLE_SYSTEM_BACKUP	Specifies whether to back up the logical log files during a scheduled backup. Optional for an Informix scheduled backup only. **Note:** If you set DO_WHOLE_SYSTEM_BACKUP to TRUE, NMDA ignores DO_LOGFILE_BACKUPS.	• TRUE (default) = Back up the logical log files during a scheduled backup. This action is equivalent to running **onbar** with the **-w** option. • FALSE = Do not back up the logical log files during a scheduled backup.
INFORMIXDIR	Specifies the directory pathname of the Informix RDBMS installation. Mandatory for an Informix scheduled backup only.	• Undefined (default). • Valid directory pathname of the Informix RDBMS installation.
INFORMIXSQLHOSTS	Specifies the name of the Informix SQL hosts file. Mandatory for an Informix scheduled backup on UNIX or (only with Informix 12.10 or later) on Windows. Optional for an Informix scheduled backup on Windows with Informix earlier than 12.10.	• Undefined (default). • Valid name of the Informix SQL hosts file.

Table 35 NMDA Informix parameters (page 2 of 2)

Parameter	Description	Default and valid values
NSR_DR_BACKUP_INFO	Specifies whether a scheduled backup backs up additional files as additional disaster recovery information and support information. If the backup of additional files fails, an error message appears but the backup report is successful. Optional for an Informix scheduled backup only. **Note:** If you enable Avamar deduplication, the software does not deduplicate these additional files.	• TRUE (default) = A scheduled backup backs up the following additional information: - Informix ONCONFIG file - ixbar file - onconfig boot file - sqlhosts file (UNIX only) - sm_versions file - Copy of Windows registry information under "HKEY_LOCAL_MACHINE\SOFTWARE\Informix" (Windows only) • FALSE = A scheduled backup does not back up the additional information.
NSR_DR_FILE_LIST	Specifies a file that contains a list of files to back up in addition to the database backup. NMDA backs up the files as part of the scheduled backup, before any postprocessing script defined by POSTCMD runs. If the backup of extra files fails, an error message appears but the backup report is successful. Optional for an Informix scheduled backup only.	• Undefined (default). • Valid complete pathname of the file that contains the list of extra files to back up. For example, the NSR_DR_FILE_LIST value is nmda_savelist.txt, which is a file that contains a list of the extra file pathnames. **Note:** Set NSR_DR_BACKUP_INFO to TRUE to enable this parameter setting.
NSR_LOG_VOLUME_POOL	Specifies the volume pool to use for a backup of the logical logs. Optional for an Informix scheduled backup only.	• Most appropriate pool as selected by the NetWorker server (default). • Valid name of a NetWorker volume pool for the logical logs.
ONCONFIG	Specifies the name of the Informix RDBMS configuration file. Mandatory for an Informix scheduled backup only.	• Undefined (default). • Valid name of the Informix RDBMS configuration file.

NMDA Lotus parameters

For the following Lotus operations, you must set both the common parameters and the NMDA Lotus parameters in the NMDA configuration file:

- Lotus scheduled backups configured without the wizard (client-side configuration)
- Lotus manual backups
- Lotus restores

"Common NMDA parameters" on page 358 describes the common parameters.

The following table describes the NMDA Lotus parameters.

Table 36 NMDA Lotus parameters (page 1 of 8)

Parameter	Description	Default and valid values
LOTUS_USER	Specifies the name of the Lotus Domino or Notes user. Mandatory for a Lotus scheduled backup on UNIX and Linux.	• Undefined (default). • Valid name of the Lotus Domino or Notes user.

Table 36 NMDA Lotus parameters (page 2 of 8)

Parameter	Description	Default and valid values
Notes_ExecDirectory	Specifies the complete pathname of the Lotus Domino or Notes directory that contains the application library. Mandatory for a Lotus backup or restore.	• Undefined (default). • Valid pathname of the Lotus Domino or Notes directory that contains the libnotes.*xx* or nnotes.dll library file.
NSR_APPLY_LOGS	Specifies whether to apply the transaction logs after a Lotus backup is restored. Optional for a Lotus restore.	• TRUE (default) = Apply the transaction logs after the backup is restored. • FALSE = Do not apply the transaction logs after the backup is restored.
NSR_AUTO_RESTORE	Specifies whether the Lotus database or file restore occurs automatically without user interaction. Optional for a Lotus restore.	• FALSE (default) = Lotus database or file restore occurs with user interaction. • TRUE = Lotus database or file restore occurs automatically without user interaction.
NSR_BACKUP_ALL_EXTENSIONS	Specifies whether to back up all the Lotus files or only the default set of Lotus files with specific file name extensions. Optional for a Lotus backup.	• FALSE (default) = Back up only Lotus files with names ending in .box, .dic, .dsk, .id, .ncf, .njf, .nrf, .nsf, and .ntf and the notes.ini file in the specific directory. • TRUE = Back up all the Lotus files with names ending in all extensions.
NSR_BACKUP_LEVEL	Specifies the level of Lotus backup to perform. Optional for a Lotus manual backup. Do not set this parameter for a scheduled backup.	• full (default) = Perform a full backup. • incr = Perform an incremental backup. • txnlog or 1 = Perform a backup of transaction logs only.
NSR_BACKUP_LOGS_MODE	Specifies the level of transaction log backup to perform during a full backup. Optional for a Lotus full backup only. Ignored for an incremental backup or transaction log backup. **Note:** NMDA always backs up the transaction logs during a Lotus incremental backup unless you set NSR_INCR_BACKUP_LOGS_MODE=1.	• 0 (default) = Do not process the transaction logs. • 1 = Back up the transaction logs and mark the logs as reusable. • 2 = Mark the transaction logs as reusable, but do not back up the logs. **Notice:** Use the NSR_BACKUP_LOGS_MODE=2 setting with extreme caution. With this setting, NMDA does not back up the transaction logs and the Domino server recycles the logs. When a log backup is missing, you might not be able to recover a database to any point-in-time; only the restore to the time of a given backup is guaranteed.
NSR_BACKUP_LOTUS_DIR	Specifies whether to back up files in the Lotus Domino or Notes data directory. • On UNIX systems, the data directory is the first Lotus data directory that NMDA finds in the parameter PATH. • On Windows systems, the data directory is the first Lotus data directory that NMDA finds in the Windows registry. Optional for a Lotus backup. **Note:** You cannot use this parameter with the NSR_BACKUP_PATHS parameter.	• FALSE (default) = Back up the Lotus directories and files specified with NSR_BACKUP_PATHS. • TRUE = Back up the Lotus data directory. On Windows, the backup includes the notes.ini file, whether or not the file resides in the default data directory. **Note:** The NSR_BACKUP_ALL_EXTENSIONS setting determines whether the backup includes all the Lotus files or only the default Lotus files with specific file name extensions in the Lotus data directory.

Table 36 NMDA Lotus parameters (page 3 of 8)

Parameter	Description	Default and valid values
NSR_BACKUP_PATHS	For a backup or restore, specifies the complete pathnames of one or more directories or files or both. "Wildcard support for Lotus operations" on page 378 describes the use of wildcards in the pathnames. Optional for a Lotus backup or restore. **Note:** For a backup, you cannot use this parameter with the NSR_BACKUP_LOTUS_DIR parameter. For a restore, you cannot use this parameter with the NSR_RECOV_LIST_FILE parameter.	• Undefined (default). **Note:** If NSR_BACKUP_LOTUS_DIR is FALSE or not set and NSR_BACKUP_PATHS is not set, NMDA backs up the transaction logs, no matter what the settings are for the backup level, NSR_BACKUP_LEVEL, NSR_BACKUP_LOGS_MODE, and NSR_INCR_BACKUP_LOGS_MODE. • For a backup or restore, valid pathnames of one or more directories or files or both, with multiple names separated by commas. The pathnames must not include the NOTES: prefix. • For a restore only, NOTES: keyword by itself, specifying to restore all the Lotus data backed up from the given client. Do not use the keyword for a partitioned Domino server or multiple Domino installations on a single UNIX host.
NSR_BROWSELIST_CACHE_DEST	Specifies the complete pathname of a directory to contain the browselist file, a temporary file that stores browselist data during a backup. The first 10 MB of generated browselist data is stored in a memory cache. Additional browselist data is stored in the browselist file on disk, named BrowselistCache_*PID*[_*PID*], where *PID* is the process ID. The browselist file is deleted when the backup ends or is terminated in a controlled manner. Optional for a Lotus backup.	• By default, the browselist file is in this directory: - /nsr/apps/tmp (UNIX) - *NetWorker_install_path*\apps\tmp (Windows) • Valid complete pathname of a directory on a writeable device to contain the browselist file.
NSR_CATALOGFILE	Specifies the complete pathname of the backup catalog file, which contains detailed information about each backed-up file. NMDA appends the information to the file after each backup. **Note:** The catalog file does not list the files from an integrated DAOS backup, but does list the files from a stand-alone DAOS backup. The catalog file also includes all the other details for a DAOS backup, same as for other Lotus backups. If NMDA cannot access the specified catalog file, the backup still runs as usual. At the end of the backup, an error message appears for the failed catalog file operation. Optional for a Lotus backup.	• If not specified, NMDA does not generate the backup catalog file (default). • Valid complete pathname of a backup catalog file. The directory path to the file must exist. NMDA creates the file during the backup if the file does not exist.
NSR_COMFORT_SPAN	Specifies the comfort span value to use for an incremental backup. "Lotus incremental backups with the comfort span option" on page 120 describes the comfort span. Optional for a Lotus incremental backup.	• Undefined (default). • Integer value between 65536 and 65536000, inclusive.

Table 36 NMDA Lotus parameters (page 4 of 8)

Parameter	Description	Default and valid values
NSR_CROSS_MOUNT_POINTS	Specifies whether NMDA crosses mount points during a Lotus backup. Optional for a Lotus backup.	• FALSE (default) = Lotus backup does not cross mount points. • TRUE = Lotus backup crosses mount points.
NSR_DBIID	Specifies that NMDA assigns either a new DBIID, or both a new DBIID and new replica ID, to a restored database. Optional for a Lotus restore.	• Undefined (default). • 1 = NMDA assigns a new DBIID to the restored database. • 2 = NMDA assigns a new DBIID and a new replica ID to the restored database.
NSR_EXCLUDE_FILE	Specifies the complete pathname of a file that lists file pathnames to exclude from the Lotus backup. Optional for a Lotus backup. "Wildcard support for Lotus operations" on page 378 describes using wildcards in the pathnames. **Note:** If you set both NSR_EXCLUDE_FILE and NSR_EXCLUDE_LIST, NMDA excludes all the files and objects specified through both parameters from the Lotus backup. NSR_EXCLUDE_FILE is a deprecated parameter that will be unsupported in a future NMDA release. Use the NSR_EXCLUDE_LIST parameter instead.	• If not specified, NMDA does not exclude pathnames from the backup paths specified by the user. • Valid complete pathname of a file that lists file paths to exclude from the Lotus backup.
NSR_EXCLUDE_LIST	Specifies the pathnames of database files or directories to exclude from the Lotus backup. If you specify a directory, NMDA excludes all its data including all subdirectories from the backup. Optional for a Lotus backup. "Wildcard support for Lotus operations" on page 378 describes the use of wildcards in the pathnames. **Note:** If you set both NSR_EXCLUDE_LIST and NSR_EXCLUDE_FILE, NMDA excludes all the files and objects specified through both parameters from the Lotus backup.	• If not specified, NMDA does not exclude pathnames from the backup paths specified by the user. • Valid pathnames of one or more database objects to exclude from the Lotus backup, with multiple names separated by commas.
NSR_FOLLOW_LINKS	Specifies which of the following actions occur when Lotus link files are to be backed up or restored: • NMDA backs up or restores both the Lotus link files and the data files or directories that the link files point to. • NMDA backs up or restores only the Lotus link files. Optional for a Lotus backup or restore.	• TRUE (default) = NMDA backs up or restores both the Lotus link files and the data files or directories that the links point to. • FALSE = NMDA backs up or restores only the Lotus link files.
NSR_INCR_BACKUP_LOGS_MODE	Specifies the level of transaction log backup to perform during an incremental backup. Optional for a Lotus incremental backup only. Ignored for a full backup or transaction log only backup.	• 0 = Do not process the transaction logs. • 1 (default) = Back up the transaction logs and mark the logs as reusable.

Table 36 NMDA Lotus parameters (page 5 of 8)

Parameter	Description	Default and valid values
NSR_LOG_DIR	Specifies the complete pathname of the log directory of a partitioned Domino server for disaster recovery only. Optional for a Lotus restore.	• Undefined (default). • Valid complete pathname of the log directory.
NSR_LOTUS_DATA_DIR	Specifies the complete pathname of the directory that contains the Lotus Notes data. Optional for a Lotus backup. Required for a partitioned Domino server or multiple Domino installations. **Note:** This parameter does not specify that the data will be backed up.	• Undefined (default). • Valid complete pathname of the directory that contains the Lotus Notes data.
NSR_MAX_TXN_LOGS	Specifies the number of transaction logs stored in a single save set during a Lotus backup. NMDA marks the logs reusable after the successful backup of all the logs in the save set. If a backup fails, NMDA marks none of the logs reusable from the incomplete save set. Optional for a Lotus backup of transaction logs.	• 10 logs per save set (default). • Integer number of logs per save set.
NSR_NO_NOTES_INIT	Specifies whether to initialize the Notes API during a disaster recovery. Mandatory for disaster recovery. Optional for a Lotus restore.	• FALSE (default) = Initialize the Notes API during the disaster recovery. • TRUE = Do not initialize the Notes API during the disaster recovery.
NSR_NOTES_CONNECT_TIMEOUT	Specifies a timeout value in seconds during which NMDA retries a Lotus backup in either of the following cases: • A Lotus database is offline while the Domino server runs a fixup command against the database. • In-place compaction of a Lotus database is in progress on the Domino server. NMDA cannot back up a database in either of these cases. NMDA retries the database backup after every five seconds, until either the database becomes accessible or the timeout is reached. If the timeout is reached first: • If NSR_SKIPDBERRORS = TRUE, NMDA skips the database and backs up the next database. • If NSR_SKIPDBERRORS = FALSE, NMDA fails the backup with an error. Optional for a Lotus backup.	• 30 (default); signifies a timeout of 30 seconds. • Integer value of timeout in seconds.
NSR_NOTES_INI_PATH	Specifies the complete pathname of the notes.ini file, including the file name. Recommended for a Lotus backup or restore.	• Undefined (default). • Valid complete pathname of the notes.ini file.

Table 36 NMDA Lotus parameters (page 6 of 8)

Parameter	Description	Default and valid values
NSR_NUMBER_LOGS	Specifies whether to restore the transaction logs during a disaster recovery only. This parameter is ignored during a regular Lotus recovery.	• 0 (default) = Do not restore the transaction logs. This is recommended for a disaster recovery of a nonlogged Domino environment. • 1 = Restore the transaction logs. This is recommended for a disaster recovery of a logged Domino environment.
NSR_PARALLELISM	Specifies the maximum number of concurrent backup or restore streams to send to or from the NetWorker server during a backup or restore. Optional for a Lotus backup or restore.	• Value determined by the NetWorker server, based on the NetWorker client and server parallelisms (default). • Integer number of the maximum concurrent backup or restore streams.
NSR_PREFETCH_LOGS	Specifies the number of transaction log files that the NMDA software retrieves in advance when NMDA applies logs to a restored Lotus database. Optional for a Lotus restore.	• 0 (default) = NMDA does not prefetch extra logs. NMDA restores only a log requested by Domino. • Integer number of transaction logs retrieved in advance, typically the number of logs backed up in a single backup.
NSR_RECOV_INTERACT	Specifies the default overwrite response when the name of a file being restored conflicts with an existing file name. The value of the parameter must be a single letter: • If the letter is lowercase, the value applies to the current file only and the overwrite prompt continues to appear for subsequent files. • If the letter is uppercase, the value applies to all the files being restored and no prompts appear unless as specified for the **R** value. Optional for a Lotus restore.	• Undefined (default). • **n** = Do not restore the current file. • **N** = Do not restore any files with conflicting names. No prompts appear. • **y** = Overwrite the existing file with the restored file. • **Y** = Overwrite all existing files with conflicting names. No prompts appear. • **r** = If restoring a logged database, do not rename the existing file, and restore the backed-up file with a name that begins with a tilde (~). If restoring a database that is not logged, rename the existing file by adding a tilde to the start of the filen ame, and restore the backed-up file with its original name. • **R** = Apply the actions of the **r** option to all existing files with conflicting names. No prompts appear.
NSR_RECOV_LIST_FILE	Specifies the complete pathname of a file that lists the files or directories to restore. Optional for a Lotus restore. **Note:** You cannot use this parameter with the NSR_BACKUP_PATHS parameter.	• Undefined (default). • Valid complete pathname of a file that lists Lotus Notes files or directories to restore. The file must contain one pathname per line, without any commas or other punctuation.
NSR_RECOVER_OPTIONS	Specifies additional recovery options. Optional for a Lotus recovery.	• Undefined (default). • REMOVE_COMMON_PATH = For a restore to a new destination directory, removes the common path of the files and directories specified with NSR_BACKUP_PATHS or specified in a file set with NSR_RECOV_LIST_FILE.

Table 36 NMDA Lotus parameters (page 7 of 8)

Parameter	Description	Default and valid values
NSR_RECOVER_TIME	Specifies the point-in-time to which NMDA recovers a database. Optional for a Lotus restore or recovery. "Setting the NSR_RECOVER_TIME parameter" on page 379 describes the setting of this parameter.	• By default, NMDA performs the following recovery: - Recovers a database in archived log mode to the current time. - Restores a database not in archived log mode to the most recent available backup. • Valid time in **nsr_getdate(3)** format.
NSR_RELOCATION_DEST	Specifies the pathname of a directory to which NMDA restores the Lotus database files. Optional for a Lotus restore. **Note:** You must set NSR_RELOCATION_DEST if the NSR_NUMBER_LOGS value is not zero.	• Undefined (default). • Valid pathname of a directory to which NMDA restores the database files. NMDA restores each database file to this location unless NSR_RECOVER_OPTIONS is set: *NSR_RELOCATION_DEST_path|original_file_path* "NSR_RECOVER_OPTIONS" on page 376 provides more information.
NSR_RESOURCE_DIR	Specifies the location of the directory that contains the Lotus resource files. Mandatory for a Lotus backup on UNIX only.	• Undefined (default). • Valid complete pathname of the Lotus directory that contains the resource files.
NSR_RETAIN_NUM_LOGS	Specifies the number of archived transaction logs to retain after a Lotus backup without marking them reusable. Optional for a Lotus backup.	• 0 = Mark all transaction logs reusable after a backup (default). • 1 or greater = Number of transaction logs to retain without marking them reusable. If the number is greater than the number of logs involved in the backup, retain all the logs without marking them reusable.
NSR_SAVESET_NAME	Specifies the base name for the save sets of a Lotus DAOS backup. If more than one save set is created, NMDA appends a numeric extension to the name to create the additional save set names. Optional for a Lotus DAOS backup only.	• Undefined (default). • Base name to use for DAOS backup save sets, for example, notes_DAOS.
NSR_SKIPDBERRORS	Specifies whether NMDA continues a Notes database backup if a noncritical error occurs while backing up Notes database files, not flat files. A noncritical error is one that allows NMDA to recover from the problem and to continue operations without compromising the backup data integrity. A noncritical error occurs when NMDA cannot access a Notes database while generating a list of files to back up. In this case, NMDA skips the file and the backup continues. A critical error occurs later if NMDA cannot access the file while trying to read the data to be saved. This error is critical because NMDA cannot safely skip the file and continue. During the backup of multiple databases, NSR_SKIPDBERRORS enables NMDA to skip the backup of problematic databases while continuing to back up good databases. Optional for a Lotus backup.	• FALSE (default) = NMDA does not continue a Notes database backup if a noncritical error occurs. • TRUE = NMDA continues a Notes database backup if a noncritical error occurs. **Note:** If you set NSR_SKIPDBERRORS to TRUE, check the output log after a backup to see if NMDA skipped any databases due to errors.

Table 36 NMDA Lotus parameters (page 8 of 8)

Parameter	Description	Default and valid values
NSR_VERBOSITY	Specifies whether NMDA writes a list of backed-up or restored files to standard output or to the NetWorker User for Lotus GUI. Optional for a manual backup or restore.	• TRUE (default) = NMDA writes successfully backed-up or recovered pathnames to standard output or to the NetWorker User for Lotus GUI. • FALSE = NMDA does not write successfully backed-up or recovered pathnames to standard output or to the NetWorker User for Lotus GUI.
NSR_XBSA_DEBUG	Specifies whether debug messages from the NetWorker XBSA library are written to the NMDA Lotus log at the level set by NSR_DEBUG_LEVEL, described in "NSR_DEBUG_LEVEL" on page 360. Optional for a backup or restore.	• FALSE (default) = XBSA library debug messages are not written to the NMDA Lotus log. • TRUE = XBSA library debug messages are written to the NMDA Lotus log.
PATH	Specifies the pathnames of the Domino data directory and Lotus software directory. Mandatory for a Lotus backup on UNIX only. Recommended for a Lotus backup on Windows and a Lotus restore.	• Undefined (default). • Valid directory pathnames of the Domino data directory and the installation directory of the Lotus binaries.

Wildcard support for Lotus operations

For a client-side configuration of a Lotus scheduled backup or manual backup, you can use wildcards to specify the pathnames of Lotus directories or files (or both) for backup. You specify the wildcards and pathnames with the NSR_BACKUP_PATHS parameter in the NMDA configuration file:

You can also use wildcards to specify the pathnames of Lotus directories or files (or both) to exclude from a backup. You specify the wildcards and pathnames in one of the following settings:

- NSR_EXCLUDE_LIST or NSR_EXCLUDE_FILE parameter in the NMDA configuration file
- Exclude Path field in the wizard

The following restrictions apply to using wildcards that NMDA expands:

- Wildcard expansion supports only asterisks (*) and question marks (?):
  - An asterisk (*) stands for any number of characters, including zero.
  - A question mark (?) stands for only one character.

  For example:
  - *.nsf matches abc.nsf and a.nsf, but not data.ntf.
  - *.n?f matches abc.nsf, a.nsf, and data.ntf, but not test.nf.
  - * matches all file names.
  - ?.nsf matches a.nsf, b.nsf, and c.nsf, but not ab.nsf.
- In NSR_BACKUP_PATHS, you can use a wildcard in the last component only of a pathname. The following example includes an invalid pathname that the NMDA software cannot expand:

```
NSR_BACKUP_PATHS = /local/*/*.nsf
```

This restriction does not apply to the exclusion of files from backup through the NSR_EXCLUDE_LIST or NSR_EXCLUDE_FILE parameter or the wizard. NMDA supports the following valid pathname:

```
NSR_EXCLUDE_LIST = /local/*/*.nsf
```

This setting excludes all the files under /local and its subfolders with file names such as address.nsf and nmda_d5.nsf.

Setting the NSR_RECOVER_TIME parameter

The information in this topic applies to both the **nsrnotesrc** and **nsrdocrc** commands.

Set the NSR_RECOVER_TIME parameter to restore a database that is not in archived log mode or to restore a nondatabase (flat) file from a backup earlier than the last one. Set the NSR_RECOVER_TIME parameter to the time of the backup or to a time after the backup but before the next backup. Use the **nsrinfo** command to determine the time. For example:

```
> nsrinfo -n notes -s bu-terminator mail1 | grep test_file

NOTES:/C:/IBM/Lotus/Domino/data/test_file.nsf, date=1304540291
   2011/05/04 16:18:11
NOTES:/C:/IBM/Lotus/Domino/data/test_file.nsf, date=1304453871
   2011/05/03 16:17:51
NOTES:/C:/IBM/Lotus/Domino/data/test_file.nsf, date=1304367455
   2011/05/02 16:17:35
```

The test_file.nsf database is in nonarchived log mode. To restore the database to its state at the time of the May 3 backup, set NSR_RECOVER_TIME to any time between 1304453871 (May 3 backup) and 1304540290 (1 second before the May 4 backup).

If the test_file.nsf database is in archived log mode, setting NSR_RECOVER_TIME to any time between 1304458034 and 1304540290 restores the database from the May 3 backup and applies transaction logs to the specified point-in-time. The 1304458034 time is the save time of the directory that contains test_file.nsf file and not the save time of the file itself. For example:

```
> nsrinfo -n notes -s bu-terminator –V mail1 | grep 05/03

NOTES:/C:/IBM/Lotus/Domino/data/, date=1304458034 2011/05/03 17:27:14
NOTES:/C:/IBM/Lotus/Domino/data/test_file.nsf, date=1304453871
   2011/05/03 16:17:51
```

If you set NSR_RECOVER_TIME to any time between the backup start time and end time (between 1304453871 and 1304458034 for the May 3 backup), this error might appear when you recover a logged database:

```
Backup was later than recovery point in time
```

To prevent this error, ensure that NSR_RECOVER_TIME is set to a time outside of the backup window.

Note: You cannot set the parameter to a value earlier than the time of the first backup because the client file index does not have any entries before that time.

NMDA MySQL parameters

For the following MySQL backup operations, set both the common parameters and the NMDA MySQL parameters in the NMDA configuration file dedicated to MySQL backup parameters:

- MySQL scheduled backups configured without the wizard (client-side configuration)
- MySQL manual backups

For MySQL restore operations, set both the common parameters and the NMDA MySQL parameters in the NMDA configuration file dedicated to MySQL restore parameters. Use a separate NMDA configuration file for the MySQL restore parameters.

"Common NMDA parameters" on page 358 describes the common parameters.

The following table describes the NMDA MySQL parameters.

NMDA MySQL backups and restores also require parameter settings in the MySQL configuration file, which is separate from the NMDA configuration file:

- Set the MYSQL_CFG_FILE parameter in the NMDA configuration file to the name of MySQL configuration file.
- The [mysqlbackup] or [client] section and the [mysqld] section of the MySQL configuration file must include the following parameters: **innodb_data_file_path**, **innodb_data_home_dir**, **innodb_log_group_home_dir**.
- The MySQL configuration file must not include the following parameters: **defaults-file**, **exec_when_locked**, **incremental**, **incremental_backup_dir**, **no_connection**, **no_history_logging**, **no_locking**, **start_lsn**, **suspend_at_end**.

The MySQL documentation describes the MySQL configuration file.

Parameter settings in the NMDA configuration file, described in the following table, take precedence over corresponding settings in the MySQL configuration file.

Table 37 NMDA MySQL parameters (page 1 of 8)

Parameter	Description	Default and valid values
MYSQL_BACKUP_DIR	Specifies the pathname of a directory to contain either the backup files extracted from a backup image or the binary logs restored from the NetWorker server. Mandatory for a MySQL restore operation except for a binary log restore or a list image operation. Optional for a binary log restore .	• /nsr/apps/tmp/BACKUP (default). • Valid pathname of a directory to which the **nsrmysqlrc** user has write permissions. **Note:** The parameter setting overrides the **backup_dir** setting in the MySQL configuration file.

Table 37 NMDA MySQL parameters (page 2 of 8)

Parameter	Description	Default and valid values
MYSQL_BACKUP_NAME	Specifies a name for a MySQL backup image or a backup name to use for a restore. For backups, this is a logical name of the backup. You cannot use the same backup image name for different MySQL backup configurations that back up different components of the same instance or back up different instances on the same machine. Mandatory for a MySQL manual backup or a restore operation. Do not set this parameter for a scheduled backup.	• Undefined (default). • Valid name for a MySQL backup: - For a MySQL backup, a backup image name that starts with MYSQL:/. For example: MYSQL_BACKUP_NAME = MYSQL:/myinstance_whole This example name is for the backup of a whole instance called myinstance. - For a MySQL restore, a backup image name or backup piece name, which is the unique save file name generated by NMDA for each backup.For example: MYSQL_BACKUP_NAME = myinstance_whole (backup image name) MYSQL_BACKUP_NAME = myinstance_full_whole_1322094812 (backup piece name as obtained with the **nsrinfo** command) **Note:** For an extract and prepare operation, you must specify a backup image name, not a backup piece name. The parameter setting overrides the **backup_image** setting in the MySQL configuration file.
MYSQL_BINLOG	Specifies one of the following values: • The pathname of a MySQL binary log to restore. • The pathnames of the first and last MySQL binary logs in a range of logs to restore. Optional for a MySQL instance recovery to the current time or a point-in-time.	• Undefined (default). • One or two valid pathnames of binary logs: - Pathname of a single binary log to restore. For example: MYSQL_BINLOG = /var/lib/mysql/bin.001 - Pathnames of the first and last logs in a log range, enclosed in square brackets and separated by a semi-colon. For example: MYSQL_BINLOG = [/var/lib/mysql/bin.001; /var/lib/mysql/bin.005] **Note:** In a log range, both log pathnames must be identical except for the numeric extension in the base file name.
MYSQL_CFG_FILE	Specifies the pathname of the MySQL configuration file, my.cnf. This file is different from the NMDA MySQL configuration file that contains the MYSQL_CFG_FILE parameter. The MySQL configuration file provides the parameter settings for a specific server instance. Mandatory for a MySQL backup. Mandatory for a restore with MySQL 5.6 or later. Optional for a restore with MySQL 5.5.	• Undefined (default). • Valid pathname of the MySQL configuration file for a server instance. For example: MYSQL_CFG_FILE = /etc/my.cnf
MYSQL_COMPRESS	Specifies whether the **mysqlbackup** process performs compression on the backup data. Optional for a MySQL backup.	• FALSE (default) = The **mysqlbackup** process does not perform compression. • TRUE = The **mysqlbackup** process performs compression.

Table 37 NMDA MySQL parameters (page 3 of 8)

Parameter	Description	Default and valid values
MYSQL_COMPRESS_LEVEL	Specifies the level of compression that the **mysqlbackup** process performs on the backup data. Optional for a MySQL backup.	• 0 (default) = The **mysqlbackup** process does not perform compression. • 1 to 9 = The **mysqlbackup** process performs the specified level of compression. The compression level increases with the numeric value.
MYSQL_DATABASES	Specifies a list of MySQL database names or table names for backup in the form "db1[.tbl1] db2[.tbl2] db3[.tbl3] ...". Optional for a MySQL backup.	• Undefined (default). • List of database names or table names or both, enclosed in double quotation marks with multiple names separated by spaces. For example: MYSQL_DATABASES = "db1 db2.tbl2 db3" **Note:** The parameter setting overrides both the **databases** setting in the MySQL configuration file and the list of names in the file specified by the **databases_list_file** setting.
MYSQL_DATADIR	Specifies the pathname of a MySQL data directory for backups. Specifies the pathname of the original MySQL data directory or a different MYSQL data directory for a recovery or copy back operation. If you do not specify this parameter for a recovery, the backup is restored to the original directory from which the data was backed up. Mandatory for a MySQL scheduled backup or a copy back operation. Optional for a MySQL recovery to the current time or a point-in-time.	• Data directory pathname obained from the mysqld section of the MySQL configuration file (default). • Valid pathname of the MySQL data directory. For a recovery or a copy back operation, the **nsrmysqlrc** user must have write permissions to the directory. For example: MYSQL_DATADIR = /var/lib/mysql **Note:** The parameter setting overrides the **datadir** setting in the MySQL configuration file. If this parameter is not set for a MySQL recovery to the current time or a point-in-time, the backup is restored to the original directory from which the data was backed up, and any **datadir** setting in the MySQL configuration file is ignored.
MYSQL_EXTRACT_PATHS	Specifies the pathname of a file or directory to be extracted from a backup image. Optionally also specifies the destination pathname to which the file or directory is extracted. If the destination pathname is not specified, the file or directory is extracted to the current working directory. Optional for an extract operation.	• Undefined (default). • *source_pathname* or *source_pathname* › *destination_pathname* where: - *source_pathname* is the pathname of a file or directory in the backup image, relative to the root of the image. - *destination_pathname* is the optional pathname on the local storage device to which the file or directory is extracted. For example: MYSQL_EXTRACT_PATHS = meta/backup_variables.txt › /tmp/backup_variables.txt

Table 37 NMDA MySQL parameters (page 4 of 8)

Parameter	Description	Default and valid values
MYSQL_INCLUDE	Specifies the per-table InnoDB datafiles to back up. Optional for a MySQL backup of InnoDB tables.	• Undefined (default). • Regular expression enclosed in double quotation marks for the per-table InnoDB datafiles to back up. For example: MYSQL_INCLUDE = “db2.tbl2” **Note:** The parameter setting overrides the **include** setting in the MySQL configuration file.
MYSQL_INCR_DIR	Specifies the pathname of the MEB incremental backup directory. Optional for a MySQL backup.	• /nsr/apps/tmp/BACKUP (default). • Valid pathname of the MEB incremental backup directory.
MYSQL_INCR_OPTIONS	Specifies the type of incremental backup to perform for an InnoDB database with MEB version 3.7 or later. An InnoDB incremental backup of only the redo logs is a differential incremental backup, which backs up the redo log changes since the last full or incremental backup. Optional for a MySQL incremental backup of an InnoDB database.	• Undefined (default). • REDO_LOG_ONLY = Perform an incremental backup of only the redo log of an InnoDB database. **Note:** The parameter setting corresponds to the **incremental-with-redo-log-only** setting in the MySQL configuration file. A setting of 0, 1, or FULL with NSR_BACKUP_LEVEL overrides this parameter.
MYSQL_INNODB_LOG_FILE_SIZE	Specifies the size of the InnoDB log file for a copy back operation. Mandatory for a copy back operation if the MySQL configuration file does not contain the corresponding setting.	• Undefined (default). • Integer value in the range of 108576 to 4294967295. **Note:** The parameter setting corresponds to the **innodb_log_file_size** setting in the MySQL configuration file.
MYSQL_INNODB_LOG_FILES_ IN_GROUP	Specifies the number of InnoDB log files in the log file group for a copy back operation. Mandatory for a copy back operation if the MySQL configuration file does not contain the corresponding setting.	• Undefined (default). • Integer value in the range of 2 to 100. **Note:** The parameter setting corresponds to the **innodb_log_files_in_group** setting in the MySQL configuration file.

Table 37 NMDA MySQL parameters (page 5 of 8)

Parameter	Description	Default and valid values
MYSQL_LOG_OPTIONS	Specifies options for a MySQL backup of binary logs. Based on the specfied options, the backup performs one of the following operations: • Backs up a whole instance and its binary logs. • Backs up only the binary logs for an instance. **Note:** Binary logging must be enabled for the instance, and the instance must be online. Otherwise, the binary log backup fails. Mandatory for a MySQL backup of binary logs.	• Undefined (default). • One or more of the following options, with multiple options separated by commas: - INCLUDE_LOGS specifies to back up the binary logs after the backup of a whole MySQL instance. If you specify both INCLUDE_LOGS and LOGS_ONLY_BACKUP, only the last value applies. - LOGS_ONLY_BACKUP specifies to back up only the binary logs for the instance. - PURGE_LOGS specifies to delete the binary logs after the backup. Optionally, specify this option with either INCLUDE_LOGS or LOGS_ONLY_BACKUP. For example: MYSQL_LOG_OPTIONS=INCLUDE_LOGS, PURGE_LOGS
MYSQL_MEB_OPTIONS	Specifies additional MySQL backup or restore options. For example, you can specify the uncompress option so that a restore uncompresses a compressed backup. **Note:** You can obtain information from the backup save set metadata about whether a backup is compressed. With this parameter, you can specify any of the options from the mysqlbackup section of the MySQL configuration file, my.cnf. Optional for a MySQL backup or restore.	• Undefined (default). • One or more of the option settings that you can specify in the mysqlbackup section of the MySQL configuration file, with multiple settings separated by commas. For example: MYSQL_MEB_OPTIONS=limit_memory=20, port=3306 **Note:** The parameter setting overrides the corresponding settings in the MySQL configuration file.
MYSQL_MEB_PATH	Specifies the pathname of the installation directory of the MEB binary **mysqlbackup**. Mandatory for the following operations: • MySQL backup or restore if the MEB binary programs are in a nondefault location • MYSQL backup or restore when you have 32-bit and 64-bit MEB on the same system.	• /opt/mysql/meb/bin (default). • Valid pathname of the installation directory of the MEB binary. For example: MYSQL_MEB_PATH=/opt/mysql/meb-3.7.1/bin

Table 37 NMDA MySQL parameters (page 6 of 8)

Parameter	Description	Default and valid values
MYSQL_ONLY_INNODB_OPTIONS	Specifies the type of MySQL backup to perform when the backup contains only InnoDB databases or tables. Due to a limitation with MEB 3.7 or later, when you specify the WITH_FRM_ALL or WITH_FRM_RELATED option, you must run the backup as an OS user with write permissions to the parent directory of the MySQL data directory. Optional for a MySQL backup that contains only InnoDB databases or tables.	• Undefined (default). • One of the following options for a backup of InnoDB databases or tables: - NO_FRM specifies to exclude .frm files from the backup. - WITH_FRM_ALL (MEB 3.7 or later only) specifies to include .frm files for all the InnoDB tables in the instance. - WITH_FRM_RELATED (MEB 3.7 or later only) specifies to include .frm files for only the tables that are included in the partial backup specified with MYSQL_INCLUDE. For example: MYSQL_ONLY_INNODB_OPTIONS=WITH_FRM_ALL **Note:** The parameter setting overrides these settings in the MySQL configuration file: **only-innodb**, **only-innodb-with-frm=all**, **only-innodb-with-frm=related**.
MYSQL_RESTORE_OPERATION	Specifies the restore operation to perform. Mandatory for a binary log restore, copy back, extract, extract and prepare, list image, or validate operation.	• Undefined (default). • One of the following values to specify a restore operation: - binlog_restore specifies to perform a binary log restore. - copy_back specifies to perform a copy back of prepared backup data to a specified directory. - extract specifies to extract data from a backup image. - extract_and_prepare specifies to extract data from a backup image and prepare the data as required (preparation is only required for InnoDB data). - list_image specifies to list files from a backup image. - validate specifies to validate the integrity of a backup image, only with MEB 3.7 or later. For example: MYSQL_RESTORE_OPERATION=extract_and_prepare
MYSQL_SBT_LIB_PATH	Specifies the pathname of the NMDA SBT library used for MySQL. Mandatory when you have 32-bit and 64-bit MEB on the same system.	• /usr/lib/libnsrmysql.so (default). • Valid pathname of the NMDA SBT library for MySQL, libnsrmysql.so. Set this parameter to the following value to back up 32-bit MySQL (uses 32-bit MEB) if you have 32-bit and 64-bit MEB on the same host: MYSQL_SBT_LIB_PATH=/usr/lib/libnsrmysql32.so **Note:** The parameter setting overrides the **sbt_lib_path** setting in the MySQL configuration file.

Table 37 NMDA MySQL parameters (page 7 of 8)

Parameter	Description	Default and valid values
MYSQL_SRC_ENTRY	Specifies the pathname of a single file or directory to be listed from a backup image. Optional for a list image operation.	• Undefined (default). • Valid pathname of a file or directory in the backup image. The pathname is relative to the root of the image. For example: MYSQL_SRC_ENTRY = meta/backup_var.txt
MYSQL_USER	Specifies the name of the MySQL backup user that connects to the MySQL instance. **Note:** You can specify the user password with the USER_PSWD parameter. Mandatory for a MySQL backup if the MySQL configuration file does not contain the username.	• Undefined (default). • Valid username of the MySQL backup user. **Note:** The parameter setting overrides the **user** setting in the MySQL configuration file.
NSR_BACKUP_LEVEL	Specifies the NetWorker backup level to use for a MySQL manual backup. Optional for a MySQL manual backup. Do not set this parameter for a scheduled backup.	• 0 or FULL (default) = Perform a full backup. • 1 = Perform a cumulative incremental backup, which backs up all data changed since the last full backup. • INCR = Perform a differential incremental backup, which backs up the data changed since the last full or incremental backup. **Note:** A setting of 0, 1, or FULL overrides the MYSQL_INCR_OPTIONS parameter.
NSR_DR_BACKUP_INFO	Specifies whether to back up additional information with a scheduled backup as disaster recovery and support information. The additional information includes the MySQL configuration file specified by MYSQL_CFG_FILE. If the backup of additional information fails, an error message appears but the backup report is successful. Optional for a MySQL scheduled backup only. **Note:** If you enable Avamar deduplication, the software does not deduplicate these additional files.	• TRUE (default) = Back up additional information with the scheduled backup. • FALSE = Do not back up additional information with the scheduled backup.
NSR_DR_FILE_LIST	Specifies a file that contains a list of MySQL files to back up in addition to the scheduled backup. NMDA backs up the files before any postprocessing script defined by POSTCMD runs. If the backup of extra files fails, an error message appears but the backup report is successful. Optional for a MySQL scheduled backup only.	• Undefined (default). • Valid pathname of a file that contains a list of additional files to back up during a scheduled backup. For example, NSR_DR_FILE_LIST is set to the pathname of a nmda_savelist.txt file that contains a list of file pathnames for backup. **Note:** Set NSR_DR_BACKUP_INFO=TRUE to enable this parameter setting.
NSR_LOG_VOLUME_POOL	Specifies the NetWorker volume pool to use for a backup of the MySQL binary logs. Optional for a MySQL binary log backup.	• Most appropriate pool as selected by the NetWorker server (default). • Valid name of a NetWorker volume pool for the MySQL binary logs.

Table 37 NMDA MySQL parameters (page 8 of 8)

Parameter	Description	Default and valid values
NSR_RECOV_INTERACT	Specifies whether to disable the prompt that the **nsrmysqlrc** program issues for shutting down the MySQL database server prior to a copy back of backup data to the data directory. Optional for a copy back operation to the data directory, which can be part of a MySQL instance recovery to the current time or a point-in-time.	• Undefined (default). • n = The **nsrmysqlrc** program does not issue the prompt to shut down the database server prior to a copy back operation to the data directory. For example: NSR_RECOV_INTERACT = n
NSR_RECOVER_TIME	Specifies the point-in-time to which the **nsrmysqlrc** program recovers the MySQL backup data. Mandatory for a MySQL point-in-time recovery.	• By default, the **nsrmysqlrc** program recovers the backup data to the most recent available backup. • Valid date in **nsr_getdate(3)** format.
USER_PSWD	Specifies the encrypted password for the MySQL backup user that connects to the MySQL instance. **Note:** You can specify the name of the MySQL backup user with the MYSQL_USER parameter. Mandatory for a MySQL backup if the MySQL configuration file does not contain the password.	• Undefined (default). • Encrypted MySQL backup user password that you must set with the **nsrdaadmin -P** command, for example: **nsrdaadmin -P -z** *configuration_file_path* The command adds the USER_PSWD setting to the NMDA MySQL configuration file. The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrdaadmin** command. **Note:** The parameter setting overrides the **password** setting in the MySQL configuration file.

NMDA Oracle parameters

For Oracle scheduled backups, you must set the following parameters in the NMDA configuration file as described in "NMDA configuration file" on page 354:

- NSR_DEBUG_LEVEL (enables debug messages for scheduled backups)
- NSR_DIAGNOSTIC_DEST (changes the debug log file locations for scheduled backups)
- NSR_DPRINTF (enables DPRINTF debug messages for scheduled backups)
- NSR_RMAN_ARGUMENTS
- ORACLE_HOME
- ORACLE_SID
- ORACLE_USER
- PRECMD
- POSTCMD
- TNS_ADMIN

You must set all the other NMDA parameters for Oracle operations in the RMAN script only.

When you set an NMDA parameter in the RMAN script, use one of the following methods, unless specified otherwise in Table 33 on page 358 (common parameters) or the following table (NMDA Oracle parameters):

- If you do not use the automatic channels feature or the backup copies feature, set the parameter with the RMAN **send** command in one of the following forms:
  - The **rman send** command on the operating system command line for manual backups.
  - The **send** command in the RMAN session or script.
- If you use the automatic channels feature or the backup copies feature, set the parameter with the **parms** option in the **configure channel** command.

  Note: With Oracle version 11gR2 or later, use **parms 'SBT_PARMS=(...)'** instead of **parms 'ENV=(...)'**.

  "Automatic channel allocation" on page 48 describes automatic channels. "Backup copies" on page 49 describes backup copies.

Do not set a parameter with the **parms 'ENV=(...)'**option. On Windows, if you set a parameter with **parms 'ENV=(...)'**, the parameter value remains in effect for all subsequent allocated channels and for all RMAN sessions until one of the following events occurs:

- You shut down the Oracle database.
- You set the parameter to a different value by using the **parms 'ENV=(...)'** option for subsequent allocated channels.
- You unset the parameter for the channel by using the **parms 'ENV=(...)'** option, as in the following example:

```
run {
   allocate channel t1 type 'SBT_TAPE'
   parms 'ENV=(NSR_SERVER=,NSR_DATA_VOLUME_POOL=)';
      :
      :
   release channel t1;
}
```

Note: On Windows systems, this situation does not occur if you set parameters with the **send** command in all RMAN sessions.

Table 38 NMDA Oracle parameters (page 1 of 5)

Parameter	Description	Default and valid values
NSR_DATA_VOLUME_POOL1	Specifies the name of the volume pool to use for a duplexed Oracle backup. Mandatory for an Oracle manual backup that uses the **set duplex** command (with **duplex** set to 2, 3, or 4) or other RMAN commands to generate two or more backup copies. Set this parameter with the **parms** option in the RMAN script (not with the **send** command or **send** option).	• Undefined (default). • Valid NetWorker pool name that is different from the name used by the parameter NSR_DATA_VOLUME_POOL, NSR_DATA_VOLUME_POOL2, or NSR_DATA_VOLUME_POOL3.

Table 38 NMDA Oracle parameters (page 2 of 5)

Parameter	Description	Default and valid values
NSR_DATA_VOLUME_POOL2	Specifies the name of the volume pool to use for a duplexed Oracle backup. Mandatory for an Oracle manual backup that uses the **set duplex** command (with **duplex** set to 3 or 4) or other RMAN commands to generate three or more backup copies. Set this parameter with the **parms** option in the RMAN script (not with the **send** command or **send** option).	• Undefined (default). • Valid NetWorker pool name that is different from the name used by the parameter NSR_DATA_VOLUME_POOL, NSR_DATA_VOLUME_POOL1, or NSR_DATA_VOLUME_POOL3.
NSR_DATA_VOLUME_POOL3	Specifies the name of the volume pool to use for a duplexed Oracle backup. Mandatory for an Oracle manual backup that uses the **set duplex** command (with **duplex** set to 4) or other RMAN commands to generate four backup copies. Set this parameter with the **parms** option in the RMAN script (not with the **send** command or **send** option).	• Undefined (default). • Valid NetWorker pool name that is different from the name used by the parameter NSR_DATA_VOLUME_POOL, NSR_DATA_VOLUME_POOL1, or NSR_DATA_VOLUME_POOL2.
NSR_ENV_LIST	Specifies a list of additional environment variables to set before starting RMAN during an Oracle scheduled backup. The Oracle Recovery Manager documentation describes the environment variables supported by RMAN. Optional for an Oracle scheduled backup. Set this parameter in the configuration file only.	• Undefined (default). • A list of one or more environment variable settings supported by RMAN. Each setting includes the variable name, an equal sign, and the variable value. Multiple settings are separate by commas, and the entire list is enclosed in double quotes, such as “name1=value1, name2=value2”. For example, the parameter specifies the NLS_DATE_FORMAT environment variable setting: NSR_ENV_LIST=”NLS_DATE_FORMAT='*dd-mmm-yyyy hh:mm:ss*'”
NSR_MMDB_RETRY_TIME	Specifies the number of minutes that NMDA tries to connect to the NetWorker media database before terminating the operation (backup, restore, or RMAN maintenance commands). When the media database is busy, NMDA tries to reconnect after sleeping for five seconds between attempts. Optional for an Oracle backup or retore.	• 0 (default) = NMDA does not try to reconnect to the media database if the first try fails. • Valid integer = Number of minutes that NMDA tries to connect to the media database.
NSR_NO_MULTIPLEX	When set for a specific RMAN channel, specifies whether to disable multiplexing during a backup on the NetWorker device that the RMAN channel uses. With multiplexing disabled, the backup cannot write other save sets to the device. Optional for an Oracle backup. To optimize restore operations, RMAN requires Oracle backups not to be multiplexed. Setting the parameter to TRUE can affect the backup performance. For example, the device can sit idle during part of the backup. If the setting adversely affects the performance, reset the parameter to FALSE.	• FALSE (default) = Enable multiplexing on the device that the RMAN channel uses. • TRUE = Disable multiplexing on the device that the RMAN channel uses. **Note:** Do not set this parameter to TRUE if you use a random access NetWorker device, such as an advanced file device.

Table 38 NMDA Oracle parameters (page 3 of 5)

Parameter	Description	Default and valid values
NSR_ORACLE_RETENTION	Specifies whether to use the Oracle RMAN retention policy to manage the backup lifecycle. Optional for an Oracle backup. Set this parameter to TRUE to use Oracle policies to manage the backup data lifecycle. NMDA supports this parameter only for a recovery window-based Oracle retention policy, not for a redundancy-based Oracle policy. **Note:** NMDA does not support this parameter for a scheduled backup of Oracle disk backups.	• FALSE (default) = Enable the NetWorker browse policy and retention policy . Use the policies to manage the lifecycle of the NMDA backup data. • TRUE = Use Oracle policies to manage the lifecycle of the NMDA backup data. NMDA internally sets the NetWorker browse and retention policy based on the Oracle retention policy and enables data uniformity to ensure that the backups are not expired if dependent backups are not expired.
NSR_PROXY_PFILE	Specifies the complete pathname of a configuration file that contains NSM parameter settings for a snapshot backup or restore. Mandatory for an Oracle backup or restore if a configuration file contains the NSM parameters settings for the operation. Supported for an NSM snapshot backup or restore only.	• Undefined (default). • Valid pathname of the configuration file. **Note:** If undefined or an invalid pathname, the operation ignores the parameter settings in the preferred configuration file.
NSR_RMAN_ARGUMENTS	Specifies any valid combination of options for the RMAN executable, **rman**(**.exe**). The Oracle Recovery Manager documentation describes the valid options. Optional for an Oracle scheduled backup. Set this parameter in the configuration file only. **Note:** NMDA does not support this parameter for a scheduled backup of Oracle disk backups.	• Undefined (default). • String that contains any valid combination of options for the RMAN executable, **rman**(**.exe**). For example, set the parameter to append RMAN output to the message log file /nsr/apps/logs/msglog.log if you do not use a Recovery Catalog: NSR_RMAN_ARGUMENTS=nocatalog msglog '/nsr/apps/logs/msglog.log' append
NSR_RMAN_OUTPUT	Specifies options to control how the RMAN output is redirected. Optional for an Oracle scheduled backup. Set this parameter in the configuration file only.	• Undefined (default). • SHELL = Start RMAN by using the operating system shell. Set this value if you use a shell-specific format for the RMAN output log file name. For example, set this parameter when NSR_RMAN_ARGUMENTS specifies to append RMAN output to a message log file with a date stamp in the log file name: NSR_RMAN_ARGUMENTS="nocatalog msglog '/nsr/apps/logs/dbid_arch_msglog_`date +%y%m%d_%H`.log' append'" NSR_RMAN_OUTPUT=SHELL
NSR_SERVER_NIC	Specifies the name of a network interface card (NIC) on a NetWorker server. Optional for an Oracle backup or restore. When you set this parameter with the RMAN **send** command for an allocated channel, the parameter value overrides the NSR_SERVER setting for that channel only. **Note:** You must explicitly set this parameter for each channel to which it applies. Setting this parameter is the only supported way to override the NSR_SERVER value for a scheduled backup.	• Undefined (default). • Valid name of a NetWorker server NIC.

Table 38 NMDA Oracle parameters (page 4 of 5)

Parameter	Description	Default and valid values
NSR_VOLUMES_INFO	Specifies whether NMDA obtains the latest NetWorker volume information (for example, remote or offline) for Oracle backups. Use this parameter only when you run the **restore...preview** command or **restore...preview recall** command Recommended for NMDA Oracle restore preview functionality.	• FALSE (default) = NMDA does not obtain the latest NetWorker volume information. • TRUE = NMDA obtains the latest NetWorker volume information.
ORACLE_HOME	Specifies the home directory pathname of the Oracle Server installation. Mandatory for an Oracle scheduled backup. Set this parameter in the configuration file only.	• Undefined (default). • Valid pathname of the home directory of the Oracle Server installation.
ORACLE_SID	Specifies the system identifier (SID) value of the Oracle database to be backed up. Mandatory for an Oracle scheduled backup in the following cases: • You store the **connect target** and **connect rcvcat** commands for the scheduled backup in a separate file, and you run the **connect** commands in the RMAN script by using the **@** command. • You enable save set bundling for the scheduled backup. • You perform an NSM snapshot backup with catalog synchronization enabled. Chapter 9, "Snapshot Backups and Restores," describes NSM snapshot backups. • You use Oracle operating system authentication on UNIX or Linux. You must also set ORACLE_USER as described in "ORACLE_USER (UNIX)" on page 391. Set this parameter in the configuration file only.	• Undefined (default). • Valid SID value of the Oracle database to be backed up. For example, if you enable catalog synchronization for NSM snapshot backups and orcl10 is the SID of the Oracle database to be backed up: ORACLE_SID=orcl10
ORACLE_USER (UNIX)	To enable a scheduled backup for operating system authentication, specifies the username of the Oracle operating system user that is set up to connect to the Oracle database through operating system authentication. You must also set ORACLE_SID as described in "ORACLE_SID" on page 391. Optional for an Oracle scheduled backup in a client-side configuration (configured with the NMC method, not the wizard) on UNIX systems only. Set this parameter in the configuration file only. **Note:** The following types of NMDA backups do not support the use of ORACLE_USER for performing an Oracle backup through operating system authentication: - Scheduled backup configured with the configuration wizard - Scheduled backup on Windows - Probe-based backup - NSM snapshot backup	• Undefined (default). • Valid username of the Oracle operating system user that you set up to connect to the Oracle database through operating system authentication.

Table 38 NMDA Oracle parameters (page 5 of 5)

Parameter	Description	Default and valid values
TNS_ADMIN	Specifies the directory pathname of the Oracle Net configuration files. Mandatory for an Oracle scheduled backup if the Oracle Net configuration files reside in a directory other than the default $ORACLE_HOME/network/admin directory. Set this parameter in the configuration file only.	• Undefined (default). • Valid pathname of the directory that contains the Oracle Net configuration files.

NMDA Sybase parameters

For the following Sybase operations, set both the common parameters and the NMDA Sybase parameters in the NMDA configuration file:

- Sybase scheduled backups configured without the wizard (client-side configuration)
- Sybase manual backups

For Sybase restores:

- Set a parameter as a **nsrsybrc** command line option, if possible. The *EMC NetWorker Command Reference Guide* and the UNIX man pages describe the **nsrsybrc** command.
- If you cannot set a parameter as a **nsrsybrc** command line option, then set the parameter in the environment.

"Common NMDA parameters" on page 358 describes the common parameters.

The following table describes the NMDA Sybase parameters.

Table 39 NMDA Sybase parameters (page 1 of 5)

Parameter	Description	Default and valid values
DBCCOPT	Specifies one or more **-o** options to pass to the **nsrsybcc** command, which performs a database consistency check for the Sybase backup. Optional for a Sybase backup. The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrsybcc** command. **Note:** If you set USE_CONSISTENCY_CHECK to TRUE and you do not set DBCCOPT, **nsrsybcc** performs all the possible checks before a backup.	• Undefined (default). • One or more of the following options of the **nsrsybcc** command, with a single option or multiple space-separated options enclosed in double quotes: -o ckdb -o ckal -o ckcat -o ckdbnoidx -o ckstor For example: DBCCOPT="-o ckcat -o ckal -o ckdb"

Table 39 NMDA Sybase parameters (page 2 of 5)

Parameter	Description	Default and valid values
LD_LIBRARY_PATH	Specifies the directory pathname of the Open Client Server (OCS) library. Mandatory for a Sybase scheduled backup on the following platforms: • HP-UX Itanium • Linux AMD64/EM64T • Solaris SPARC Optional for a Sybase manual backup on these platforms. For a Sybase manual backup, you can alternately set this parameter as an environment variable.	• Undefined (default). • Directory pathname of the OCS library. This value must be the same as set in the SYBASE.sh or SYBASE.csh script.
LD_LIBRARY_PATH_64	Specifies the directory pathname of the Open Client Server (OCS) library. Mandatory for a Sybase scheduled backup on Solaris AMD64/EM64T. Optional for a Sybase manual backup on Solaris AMD64/EM64T. For a Sybase manual backup, you can alternately set this parameter as an environment variable.	• Undefined (default). • Directory pathname of the OCS library. This value must be the same as set in the SYBASE.sh or SYBASE.csh script.
LIBPATH	Specifies the directory pathname of the Open Client Server (OCS) library. Mandatory for a Sybase scheduled backup on AIX. Optional for a Sybase manual backup on AIX. For a Sybase manual backup, you can alternately set this parameter as an environment variable.	• Undefined (default). • Directory pathname of the OCS library. This value must be the same as set in the SYBASE.sh or SYBASE.csh script.
NSR_ASE_PASSWORD	Specifies an unencrypted password to add to the Sybase **dump** command for password-protecting the Sybase backup data. Optional for a Sybase backup only. **Note:** The NMDA configuration file contains the unencrypted password value.	• Undefined (default). • Unencrypted password, from 6 to 30 characters in length, to add with the **passwd=** clause to the Sybase **dump** command.
NSR_ASE_VERIFY	Specifies one of the following options for backup verification: • full • header Optional for a Sybase backup only.	• Undefined (default). • One of the following values: - full = Verify both the header information and rows structure (full verification of the backup). - header = Verify the page header information only. For example, the following setting specifies a full verification of the backup: NSR_ASE_VERIFY=full

Table 39 NMDA Sybase parameters (page 3 of 5)

Parameter	Description	Default and valid values
NSR_BACKUP_LEVEL	Specifies the level of Sybase manual backup to perform. Optional for a Sybase manual backup. Do not set this parameter for a scheduled backup.	• full (default) = Perform a full backup, which backs up the database. • incr = Perform an incremental level backup, which backs up the transaction logs. **Notice:** A whole instance incremental backup skips the backup of any database that does not support incremental backups (for example, when the database data and transaction logs are on the same device).
NSR_BACKUP_PATHS	Specifies the backup of either the entire Sybase server or one or more Sybase databases. Mandatory for a Sybase manual backup only. **Note:** Do not specify both a server name and a list of databases.	• Undefined (default). • Valid pathnames in either of the following forms, with multiple database names separated by a comma: SYBASE:/*ASE_server_name* (backs up the entire server) SYBASE:/*ASE_server_name*/*database_name* [,SYBASE:/*ASE_server_name*/*database_name*...]
NSR_DUMP_DATA_OPT	Specifies options for the backup of a Sybase database.	• Undefined (default). • "shrink_log" = Perform a shrink log operation during the backup of a Sybase 15.7 or later database. For example, the following setting specifies to perform the shrink log operation during a Sybase database backup: NSR_DUMP_DATA_OPT="shrink_log"
NSR_DUMP_LOG_OPT	Specifies options for the transaction log backup in case of an emergency, such as a lack of free log space or a failed media database. Optional for a Sybase backup only. "Sybase transaction log backups" on page 139 provides details on using the options for transaction log backups. **Note:** Set this parameter only when there is an emergency during the transaction log backup.	• Undefined (default). • One of the following values: - "no_log" = Truncate the transaction log without recording the operation. - "no_truncate" = Back up the transaction log without truncating the log. - "truncate_only" =Truncate the transaction log without backing it up. For example, the following setting specifies to back up the transaction log without truncation: NSR_DUMP_LOG_OPT="no_truncate"
NSR_EXCLUDE_FILE	Specifies the complete pathname of a file that lists databases to exclude from a backup of a Sybase server. Optional for a Sybase backup of a server. **Note:** Do not specify this parameter for a Sybase backup of one or more databases.	• Undefined (default). • Valid complete pathname of an ASCII file that lists the databases to exclude from the backup of the server. The file lists each database on a separate line in the following format: SYBASE:/*ASE_server_name*/*database_name*

Table 39 NMDA Sybase parameters (page 4 of 5)

Parameter	Description	Default and valid values
NSR_LOCALE	Specifies the locale to use to connect to the Sybase server for a Sybase client-side scheduled backup or a manual backup configuration. Optional for a Sybase client-side scheduled backup or manual backup. **Note:** This parameter is not supported for a server-side scheduled backup or probe-based backup.	• Undefined (default). • Valid locale identifier.
NSR_LOG_VOLUME_POOL	Specifies the volume pool to use for a backup of the transaction logs. Optional for a Sybase backup. **Note:** NMDA backs up the metadata from a transaction log backup to a regular (nonlog) volume pool. "Specifying volume pools for Sybase incremental backups" on page 138 provides details.	• Most appropriate pool as selected by the NetWorker server (default). • Valid name of a NetWorker volume pool for the transaction logs.
NSR_PARALLELISM	Specifies the number of stripes to use to back up each database during a backup. A multistripe backup extracts multiple data streams in parallel from a database and writes the data streams in parallel to one or more media devices. Optional for a Sybase backup only. **Note:** NMDA does not support multistripe backups for the backup of transaction logs.	• 1 (default). • Integer value of 1 or greater for the number of multistripe sessions.
NSR_PROMOTE_FULL	Specifies whether to promote an incremental backup to a full backup when an incremental backup cannot be performed. **Notice:** This parameter only applies when the incremental backup is not an whole instance incremental backup. A whole instance icnremental backup ignores this parameter and skips the backup of any database that does not support incremental backups. Optional for a Sybase backup only.	• TRUE (default) = Promote an incremental backup to a full backup if an incremental backup cannot be performed. • FALSE = Do not promote an incremental backup to a full backup if an incremental backup cannot be performed.

Table 39 NMDA Sybase parameters (page 5 of 5)

Parameter	Description	Default and valid values
PATH	Specifies the following values: • On UNIX systems, the pathname of the directory that contains the NetWorker binaries. • On Windows systems, the pathnames of the directories that contain the Open Client Server (OCS) library and the NetWorker binaries. Mandatory for a Sybase scheduled backup on Windows only if the paths of the NetWorker client binaries and OCS library are not in the system path. Set on UNIX only if you relocated the NetWorker client binaries.	• Undefined (default). • Valid directory pathnames as follows: - On UNIX, the directory pathname of the NetWorker binaries. - On Windows, the directory pathnames of the OCS library and NetWorker binaries, with a semicolon separating the pathnames. For example, set the parameter as follows to add the OCS library pathnames on Windows: PATH=%PATH%;%SYBASE%\%SYBASE_OCS%\bin;%SYBASE%\%SYBASE_OCS%\dll
SYBASE	Specifies the pathname of the directory where you installed the Sybase ASE software. Mandatory for a Sybase backup or restore. Set the parameter with the required method: • For a scheduled backup, set the parameter in the configuration file. • For a manual backup, set the parameter in the configuration file or the environment. • For a restore, set the parameter in the environment.	• Undefined (default). • Valid directory pathname for the Sybase ASE software installation.
SYBASE_USER	Specifies the name of the Sybase user that connects to the Sybase server for the backup. You must specify the password of the user by setting the USER_PSWD parameter. *Mandatory* for a Sybase backup.	• Undefined (default). • Valid Sybase username.
USE_CONSISTENCY_CHECK	Specifies whether the **nsrsybcc** command performs a database consistency check before a backup occurs. Optional for a Sybase scheduled backup only.	• FALSE (default) = Do not run the database consistency check command, **nsrsybcc**, before a backup. • TRUE = Run the database consistency check command, **nsrsybcc**, before a backup. The specific checks performed depend on the DBCCOPT setting.
USER_PSWD	Specifies the encrypted password for the Sybase user that connects to the Sybase server, specified with the SYBASE_USER parameter. Mandatory for a Sybase backup only if the Sybase server has a password.	• Undefined (default). • Encrypted Sybase user password, which you must set by using the **nsrdaadmin -P** command, for example: **nsrdaadmin -P -z** *configuration_file_path* The *EMC NetWorker Module for Databases and Applications Command Reference Guide* describes the **nsrdaadmin** command.

APPENDIX B
Oracle RMAN Commands

This appendix includes the following topics, which describe Oracle RMAN commands that you can use in the RMAN scripts for Oracle backups and restores:

- The pool option of the backup command 398
- The send command 398

The pool option of the backup command

NOTICE

NMDA does not support the pool option of the RMAN backup command, with the exception of pool=0.

If you specify any nonzero value for the **pool** option of the RMAN **backup** command, the RMAN session terminates and NMDA returns the following error message:

```
sbtbackup: Oracle pools are not supported
```

"NMDA Oracle error messages" on page 428 describes this error message.

To specify the NetWorker volume pool that NMDA will use, set the parameter NSR_DATA_VOLUME_POOL in the RMAN script. Appendix A, "NMDA Parameters and Configuration File," provides details.

The send command

The NMDA implementation of the **send** command enables you to set both the common NMDA parameters and the NMDA Oracle parameters.

Set the parameter values by using the methods described in "NMDA Oracle parameters" on page 387. Use the **send** command where possible.

The following topics describe the **send** command syntax, the precedence rules, and how to use the **send** command to set the parameters:

- "Syntax rules" on page 398
- "Two ways to run the send command" on page 400
- "Precedence rules" on page 402

Note: In the following topics, brackets ([]) denote the optional portions of a command, for example, the command options and the corresponding settings. When you type the command, do not include the brackets.

Syntax rules

The **send** command must have the following format:

```
send [ device_type 'device_specifier' | channel channel_id ]
   'NSR_ENV=(name1=value1 [, name2=value2, ...])'
```

The following topics describe syntax rules for the two main parts of the **send** command:

- "The send command string" on page 399
- "The send command options" on page 399

The send command string

The command string in the **send** command is the string inside the quotes, 'NSR_ENV=(*name1*=*value1*...)'.

Follow these syntax rules for the **send** command string:

- Oracle software restricts the maximum length of the command string to 512 bytes, including the terminating NULL.
- The NSR_ENV keyword and the parameter names must be all uppercase.
- Between the NSR_ENV keyword and the opening parenthesis, you can optionally include an equal sign and one or more spaces. For example, the following commands are all correct:

```
send 'NSR_ENV = (NSR_SERVER=server1)'
send 'NSR_ENV=(NSR_SERVER=server1)'
send 'NSR_ENV (NSR_SERVER=server1)'
send 'NSR_ENV(NSR_SERVER=server1)'
```

- The parentheses in the command string are mandatory.
- Inside the parentheses, you must include one or more NMDA parameter names and the corresponding parameter values.
- Inside the parentheses, you must not include spaces around the equal signs. A space before an equal sign becomes part of the parameter name. A space after an equal sign becomes part of the parameter's value.
- Commas separating the *name*=*value* entries are mandatory.
- You must not include comments inside the quotes. In the following example, "# NSR_SERVER" is the first parameter name:

```
run {
   allocate channel t1 type 'SBT_TAPE';
   send 'NSR_ENV=(
#  NSR_SERVER=server1,
   NSR_CLIENT=oracle)';
      :
```

- A **send** command in an RMAN script can span multiple lines. For example:

```
send 'NSR_ENV=(
NSR_SERVER=server1,
NSR_CLIENT=oracle)';
```

The send command options

Run the **send** command with only one of the following options:

- **send** with no option (only the quoted command string) sets the parameters for all allocated channels.
- **send device_type SBT** sets the parameters for all channels allocated for NMDA.
- **send channel** sets the parameters for the specified channels only.

NOTICE

You can use the device_type or the channel option in the **send** command in an RMAN script only. You cannot use either option in the **send** command on the operating system command line. "The send command on the operating system command line" on page 401 provides details.

Example 47 A **send** command sets the parameters for a specified channel

In the following sample script, the **send** command sets the parameters for channel t1, not for channel t2:

```
run {
   allocate channel t1 type 'SBT_TAPE';
   allocate channel t2 type 'SBT_TAPE';
   send channel t1 'NSR_ENV=(NSR_SERVER=server1,
   NSR_DATA_VOLUME_POOL=MondayFulls)';
      :
}
```

The following table refers to the sample RMAN script in Example 47 on page 400, listing the values for options that the **send** command uses.

Table 40 Option values in the **send** command

Option value	Describes
device_specifier	The device type as specified in an **allocate channel** command in the RMAN script. For a backup tape device, use SBT or SBT_TAPE.
channel_id	The channel identifier as specified in an **allocate channel** command in the RMAN script. In the example, the identifier is t1.
name1	The first NMDA parameter name. In the example, the first parameter name is NSR_SERVER.
value1	The value assigned to the first parameter. In the example, the first value is server1.
name2	The second NMDA parameter name. In the example, the second parameter name is NSR_DATA_VOLUME_POOL.
value2	The value assigned to the second parameter. In the example, the second value is MondayFulls.

Two ways to run the send command

There are two different ways to run the **send** command:

- As an option of the **rman** command on the operating system command line, as described in "The send command on the operating system command line" on page 401.
- In the **run** job of the RMAN script, as described in "The send command in the RMAN script" on page 402.

The send command on the operating system command line

To run the **send** command as an option of the **rman** invocation on the operating system command line, type the command in the following format:

```
rman send "'NSR_ENV=(name1=value1[, name2=value2, ...])'"
```

- If more than one **send** option appears in the **rman** command, only the last **send** command runs.
- Follow all the **send** command syntax rules listed in "The send command string" on page 399, except for the last rule, which applies only to a **send** command in an RMAN script.
- Do not use either the **device_type** or **channel** option. "The send command options" on page 399 provides details.
- Use two sets of quotes around the command string, each set consisting of a single quote and a double quote. The single quote can be either before or after the double quote, but the second set of quotes must be opposite to the first set. For example, this command is also correct:

  ```
  rman send '"NSR_ENV=(name1=value1[, name2=value2, ...])"'
  ```

  Two sets of quotes are required to prevent some operating system shells (for example, **ksh**) from treating spaces inside the quotes as meta (special) characters and attempting to tokenize the string.
- The parameter values in the quoted string apply to all channels allocated during the RMAN session. These values are applied before any parameter values specified in **send** commands within the RMAN script. "Precedence rules" on page 402 provides details.

Example 48 An rman send command sets a parameter for all channels

In the following example, the NSR_SERVER parameter value (mars.emc.com) is applied to all three channels (t1, t2, t3) allocated in the RMAN script:

```
rman send "'NSR_ENV=(NSR_SERVER=mars.emc.com)'"
```

(RMAN script:)

```
run {
   allocate channel t1 type 'SBT_TAPE';
   allocate channel t2 type 'SBT_TAPE';
   allocate channel t3 type 'SBT_TAPE';
}
```

The send command in the RMAN script

To run the **send** command in the **run** job of the RMAN script, type the command in the following format, at the required point within the **run** command brackets:

```
send [ device_type 'device_specifier' | channel channel_id ]
   'NSR_ENV=(name1=value1 [, name2=value2, ...])'
```

- Follow all the **send** command syntax rules listed in "The send command string" on page 399.
- Use either the **device_type** or **channel** option (if required) with the **send** command in an RMAN script, as described in "The send command options" on page 399.
- Specify the correct option values in the **send** command, as described in "The send command options" on page 399.
- RMAN commands run in the order that they appear in the backup or restore script. For a parameter value to be in effect during a backup or restore, place the **send** command (setting the value) as follows in the script:
  - Before the **backup** or **restore** command.
  - After the **allocate channel** commands for those channels to which the parameter value applies.
- If no channel is allocated when the **send** command runs, an RMAN error appears.

The following sample RMAN script performs an Oracle backup of the entire database to the volume pool MondayFulls of the (remote) NetWorker server mars.emc.com:

```
run {
   allocate channel t1 type 'SBT_TAPE';
   allocate channel t2 type 'SBT_TAPE';
   send 'NSR_ENV=(NSR_SERVER=mars.emc.com,
   NSR_DATA_VOLUME_POOL=MondayFulls)';
   backup full filesperset 4
   format 'FULL_%d_%U'
   (database);
   release channel t1;
   release channel t2;
}
```

This script is the same as the sample script on page 129. The single **send** command sets the parameters for both channels.

Precedence rules

Parameters are set for channels allocated during the RMAN session in the following order:

1. In the **parms** option in the **allocate channel** or **configure channel** command (use **configure channel** only for automatic channel allocation).

   Note: With Oracle version 11gR2 or later, if you use both **'SBT_PARMS=(...)'** and **'ENV=(...)'** with the **parms** option, parameters set with **parms 'ENV=(...)'** are ignored.

2. In the **rman send** command on the operating system command line.
3. In the **send** command in the **run** job of the RMAN script.

NOTICE

If you simultaneously use the **send** command on the operating system command line and the **send** option in the **configure channel** command, Oracle software runs only the **send** option in the **configure channel** command.

To prevent confusion and simplify the task of setting parameters in a specific order, do not mix these different ways of setting parameters in the same RMAN session.

Example 49 Order of parameters set according to the precedence rules

The following example sets the parameters NSR_SERVER and NSR_CLIENT in this order:

- Sets NSR_SERVER to server1 (by **rman send**), changes NSR_SERVER to server2 (by the first **send** command), and finally changes NSR_SERVER to server3 (by **send channel**).
- Sets NSR_CLIENT to client1 (by **rman send**), changes NSR_CLIENT to client2 (by the first **send** command), and finally changes NSR_CLIENT to client3 (by **send channel**):

  ```
  rman send "'NSR_ENV=(NSR_SERVER=server1, NSR_CLIENT=client1)'"
  ```

  (RMAN script:)

  ```
  run {
     allocate channel t1 type 'SBT_TAPE';
     send 'NSR_ENV=(NSR_SERVER=server2, NSR_CLIENT=client2)';
     send channel t1 'NSR_ENV=(NSR_SERVER=server3,
     NSR_CLIENT=client3)';
  }
  ```

APPENDIX C
Troubleshooting and Error Messages

This appendix includes the following topics:

- General troubleshooting tips 406
- NMDA error messages 408
- DB2 troubleshooting tips and error messages 412
- Informix troubleshooting tips and error messages 415
- Lotus troubleshooting tips and error messages 417
- MySQL troubleshooting tips and error messages 421
- Oracle troubleshooting tips and error messages 425
- Sybase troubleshooting tips and error messages 440

General troubleshooting tips

Use the following list to troubleshoot basic problems in running backup and restore operations with NMDA.

To set up an NMDA backup and restore system, perform the following steps:

1. Ensure that you have configured the database server or application server system and any network services (if used) according to the instructions in the appropriate documentation.
2. Verify that the operating system, database server or application server, NetWorker server, NetWorker client, and NMDA software is supported and correctly installed and configured.

   The following sources describe the installation and configuration requirements:

   - *EMC NetWorker Module for Databases and Applications Installation Guide*
   - *EMC NetWorker Software Compatibility Guide*
   - *EMC NetWorker Installation Guide*
   - *EMC NetWorker Administration Guide*
   - Chapter 2, “Backup Configuration”

   Verify the version of NMDA installed by using one of the following commands to check the version of the **nsrdasv** program file, where *file_name* is the complete pathname of the **nsrdasv** binary:

   - On UNIX systems:

     ```
     what file_name
     ```

   - On Linux systems:

     ```
     strings file_name | grep "@(#)"
     ```

   - On Windows systems:

     1. Open Windows Explorer, and locate the **nsrdasv.exe** file in the *NetWorker_install_path*\bin directory.
     2. Right-click the file icon, and select **Properties**.
     3. In the **Properties** dialog box, select the **Version** tab to display the version information.

   For example, the following command displays version information for the **nsrdasv** binary on Solaris:

   ```
   what /usr/sbin/nsrdasv
   ```

   To display the NMDA version number, enter the appropriate command for the operating system:

   - On AIX:

     ```
     lslpp -L all | grep -i lgtonmda
     ```

   - On HP-UX:

     ```
     swlist -l product NMDA
     ```

- On Linux:

  ```
  rpm -qa | grep -i lgtonmda
  ```

- On Solaris:

  ```
  pkginfo -l LGTOnmda
  ```

3. Ensure that you can perform a manual backup with NMDA according to the instructions in "Performing manual backups" on page 148.
4. Ensure that you can perform a scheduled backup with NMDA according to the instructions in "Performing scheduled backups" on page 146.

Debug log files

You can set the following parameters in the configuration file (if not using the wizard) or in the wizard to specify settings for the NMDA debug logs:

- "NSR_DEBUG_LEVEL" on page 360
- "NSR_DIAGNOSTIC_DEST" on page 362

The NSR_DEBUG_LEVEL parameter specifies the level of debug messages that the NMDA software generates. The valid parameter values are 0 to 9:

- If you have not set the parameter or you have set the parameter to 0, NMDA does not generate debug messages.
- If you have set the parameter to a value between 1 and 9, NMDA generates debug messages in the debug log file on the NMDA host. The level of detail in the debug messages increases with the debug level.

The NSR_DIAGNOSTIC_DEST parameter specifies the directory location of the NMDA debug logs except the configuration wizard debug logs. If you have not set the parameter, the debug logs are in the following default directory:

- On UNIX: /nsr/apps/logs
- On Windows: *NetWorker_install_path*\apps\logs

Note: When you have installed 32-bit NMDA on 64-bit Windows, the 32-bit NMDA software generates logs in the *NetWorker_client_install_dir*\nsr\apps\logs directory instead of the directory where NMDA resides, which is C:\Program Files (x86)\Legato\nsr\apps\logs or C:\Program Files (x86)\EMC NetWorker\nsr\apps\logs.
As a workaround, set the NSR_DIAGNOSTIC_DEST parameter to the pathname of the preferred logs directory, for example:

```
NSR_DIAGNOSTIC_DEST=C:\Program Files (x86)\Legato\nsr\apps\logs
```

or

```
NSR_DIAGNOSTIC_DEST=C:\Program Files (x86)\EMC NetWorker\nsr\apps\logs
```

The NMDA log files have this file name format:

binaryname[*_app*].*date*.*time*.*pid*[*_threadid*].log

where *app* indicates the type of application:

- db2
- informix

- lotus
- mysql
- oracle
- sybase

The NetWorker server also writes diagnostic information from a manual or scheduled backup to specific log files on the NetWorker server. The *EMC NetWorker Administration Guide* describes these log files.

Wizard backup configuration fails, authentication denied

If the configuration wizard fails to create a scheduled backup, an error message might indicate that authentication is denied, or denied for *username*.

The lockbox with the database connection credentials was not accessible by the superuser on the client host where the backup failed and the message appeared.

As a solution, use the NMC program to ensure that the Lockbox resource is created for the given client and the Users attribute contains the superuser of the client.

The *EMC NetWorker Administration Guide* describes lockbox password management.

Backup becomes suspended

If the backup becomes suspended, the NetWorker server might be temporarily unavailable at the start of the backup. The backup waits until the NetWorker server becomes available.

As a solution, edit the NMDA configuration file or (for Oracle backups only) the RMAN backup scripts and set this parameter to TRUE:

```
NSR_NO_BUSY_ERRORS=TRUE
```

"NSR_NO_BUSY_ERRORS" on page 364 describes this parameter.

If you set the NSR_NO_BUSY_ERRORS parameter to TRUE and the backup still becomes suspended, use the following command to determine if the **nsrexecd** program is running:

```
# ps -ef | grep nsrexecd
```

If **nsrexecd** is not running, start the program with the following command:

```
# nsrexecd
```

NMDA error messages

During a backup or restore, the NMDA software records the NMDA error messages in an error log file on the NMDA host.

You can set the NSR_DIAGNOSTIC_DEST parameter to specify the directory location of the NMDA debug logs except the configuration wizard debug logs. "NSR_DIAGNOSTIC_DEST" on page 362 provides details.

If you have not set the NSR_DIAGNOSTIC_DEST parameter, the debug log file is in the following default directory:

- On UNIX: /nsr/apps/logs
- On Windows: *NetWorker_install_path*\apps\logs

NMDA generates a separate error log for each different application involved in the backups and restores. The error log file has the following file name:

`nmda_app.messages.raw`

where *app* indicates the type of application:

- db2
- informix
- lotus
- mysql
- oracle
- sybase

NMDA generates error messages in the nmda_*app*.messages.raw file in a language-independent form, readable by the **nsr_render_log** program only.

The *EMC NetWorker Administration Guide* describes how to use the **nsr_render_log** program to read any language-independent binary file, such as nmda_*app*.messages.raw.

NetWorker XBSA error messages

During a DB2, Informix, Lotus, or Sybase backup or restore, NMDA records error messages generated by the NetWorker X/Open Backup Services Application (XBSA) API in the NMDA error log or xbsa.messages file. The xbsa.messages file is in the same directory as the NMDA error logs.

Oracle operations do not generate NetWorker XBSA error messages.

NetWorker XBSA error messages appear in the following format:

```
XBSA-1.0 NMDA_release_branch. process_id day month date hh:mm:ss year
   function_name: BSA_RC_message_code: message
```

The following table lists the relevant NetWorker XBSA error messages.

Table 41 NetWorker XBSA error messages (page 1 of 3)

Error message	Description
BSA_RC_ABORT_SYSTEM_ERROR System detected error due to *explanation*. Operation aborted	A general system error occurred within a NetWorker XBSA function call. This error appears for all NetWorker errors that do not map cleanly to XBSA errors.
BSA_RC_BAD_CALL_SEQUENCE The sequence of API calls is incorrect. Must call item1 before item2	An API call sequence occurred that does not conform to the XBSA Data Movement API State Diagram document.
BSA_RC_BAD_HANDLE The handle used to associate this call with a previous BSAInit() call is invalid because explanation	The value passed into the function for bsaHandle contained a NULL pointer.
BSA_RC_BAD_PARAMETER received parameter parm with value value, which is invalid	The software received an invalid parameter.
BSA_RC_BUFFER_TOO_SMALL Buffer is too small to hold the object entry to be returned. n bytes required for the object entry	The buffer is too small to hold the object entry.

Table 41 NetWorker XBSA error messages (page 2 of 3)

Error message	Description
BSA_RC_DESCRIPTION_TOO_LONG The description field contained too many characters (n >= n)	The Description field in one of the supplied structures contained more than BSA_MAX_DESC characters, and the structure was not usable for the requested operation.
BSA_RC_INVALID_COPYTYPE the copyType field contained an unrecognized value of n	The copyType field in one of the supplied structures has a value that is not in the NetWorker XBSA libraries implementation of this enumerated type.
BSA_RC_INVALID_DATABLOCK the dataBlock parameter contained inconsistent values: bufferLength: n, bufferPtr: n, numBytes: n	The fields of a supplied DataBlock parameter are not internally consistent. This can occur under one of the following conditions: • The bufferLen field is less than the numBytes field while data is being sent. • The bufferLen field is nonzero and the bufferPtr field is NULL.
BSA_RC_INVALID_KEYWORD an entry in the environment structure is invalid (variable=value)	One of the environment strings passed into the function did not have a valid structure. The value structure of an environment keyword is KEYWORD = VALUE, where KEYWORD is a white space delimited string and VALUE is a white space delimited string followed by a null terminator. This can indicate several possible errors: • The KEYWORD was not in the reserved word list. The NetWorker XBSA libraries do not return this error because other environment variables might be passed into the library along with valid keywords. • The KEYWORD and VALUE strings were not separated by a '=' character. This type of error also detects environment vectors that are not terminated with a (char *)NULL entry, and invalid KEYWORD VALUE pair formats. • The VALUE string was invalid. • The VALUE string could not be validated, as in the case of a hostname string that the **gethostbyname**() function could not find.
BSA_RC_INVALID_OBJECTSTATUS the objectStatus field contained an unrecognized value of n	The objectStatus field in one of the supplied structures has a value that is not in the NetWorker XBSA libraries' implementation of this enumerated type.
BSA_RC_INVALID_OBJECTTYPE the objectType is invalid (n)	One of the object type parameters was either passed in directly or contained in one of the following structures: ObjectDescriptor. QueryDescriptor was not in the range of BSAObjectType_ANY to BSAObjectType_DIRECTORY.
BSA_RC_INVALID_TIME a time field contained an unrecognized value of n	The software received an invalid time value.
BSA_RC_MATCH_EXISTS object matching the specified predicate already exists	The object already exists in the NetWorker server that the NetWorker XBSA session uses, and the requested operation cannot be completed.
BSA_RC_MORE_DATA more data is available. Data can be obtained through BSAGetData() or BSAGetNextQueryObject()	This errorhas two meanings in the XBSA Data Movement API: • Object Data Retrieval – There is more data available for an object being read from the NetWorker server than is being used by the NetWorker XBSA session. Use BSAGetData to retrieve the next DataBlock from the NetWorker server. Refer to "BSA_RC_BUFFER_TOO_SMALL" on page 409 and "BSA_RC_NO_MORE_DATA" on page 411. The BSAGetObjectF function does not return this message because this function writes all the data for an object to a file descriptor. • Query Result Retrieval – There are more objects matching the requested query descriptor from the NetWorker server than is being used by the NetWorker XBSA session. Use BSAGetNextQueryObject to retrieve the next object descriptor from Backup Services. Refer to "BSA_RC_NO_MORE_DATA" on page 411.

Table 41 NetWorker XBSA error messages (page 3 of 3)

Error message	Description
BSA_RC_NO_MATCH The ResourceType predicate value of D does not match the reference value of L	The client index and media database are out of synch. To resynchronize the client index and media database, run the **nsrck -X** command. Alternatively, wait for NetWorker to run **nsrck** automatically.
BSA_RC_NO_MATCH The variable predicate value of value does not match the reference value of variable	No objects matching the specified QueryDescriptor were found in the NetWorker server that the NetWorker XBSA session uses.
BSA_RC_NO_MORE_DATA there is no more data for the current object	This error has two meanings in the XBSA Data Movement API: • Object Data Retrieval – This error is used when all the data for an object being retrieved from a NetWorker server was placed into the given DataBlock parameter for a function call. Refer to "BSA_RC_NO_MORE_DATA" on page 411. • Query Result Retrieval – This error is used when the last (or only) object matching a query is returned to the caller . Refer to "BSA_RC_NO_MORE_DATA" on page 411.
BSA_RC_NULL_APIVERSION an ApiVersion pointer is required	A pointer to an ApiVersion structure, passed into the function, was NULL and is required as input.
BSA_RC_NULL_DATABLOCK a data block pointer is required	The DataBlock pointer parameter for the called function was NULL. The caller is responsible for allocating and passing in the DataBlock structure to the NetWorker XBSA library. Refer to "BSA_RC_INVALID_DATABLOCK" on page 410.
BSA_RC_NULL_OBJECTNAME an object name is required	The ObjectName parameter passed into the called function was NULL.
BSA_RC_NULL_OBJECTOWNER an ObjectOwner pointer is required	A pointer to an object-owner structure was NULL and is required as input.
BSA_RC_OBJECTINFO_TOO_LONG The objectInfo field contained too many characters	The ObjectInfo parameter passed into the function, either directly or in one of the following data structures, was found to have more than BSA_MAX_OBJINFO characters: ObjectDescriptor
BSA_RC_OBJECTSPACENAME_TOO_LONG The objectSpaceName field contained too many characters	The string objectSpaceName contains more than BSA_MAX_OBJECTSPACENAME characters in an ObjectName structure.
BSA_RC_PATHNAME_TOO_LONG The pathName field contained too many characters	The string pathname contains more than BSA_MAX_PATHNAME characters in an ObjectName structure.
BSA_RC_RESOURCETYPE_TOO_LONG The resourceType field contained too many characters (n ›= n)	The string resourceType contains more than BSA_MAX_RESOURCETYPE characters and might be corrupt.
BSA_RC_SUCCESS the function was successful	The called function did not fail and is returned by all NetWorker XBSA function calls.
BSA_RC_TRANSACTION_ABORTED the transaction was aborted	The BSAEndTxn function call terminated the current transaction. A transaction can be terminated by an internal error or by a user request through the Vote parameter to this function.

DB2 troubleshooting tips and error messages

The following topics provide NMDA DB2 troubleshooting tips and error messages.

Delta or incremental backup fails in a DB2 pureScale system

If you attempt to use the DB2BACKUP_DELTA or DB2BACKUP_INCREMENTAL setting with the DB2_OPTIONS parameter to perform a delta or incremental backup in a DB2 pureScale system, the backup fails with the following error message:

```
NMDA backup failed.
64673:nsrdasv: Unable to backup SAMPLE database due to backup request
   failure, SQLCODE : -1419, SQL1419N  The statement, clause, command,
   API, or function is not supported in a DB2 pureScale environment.
   Reason code = "21".  SQLSTATE=56038
 SQLSTATE 56038: The requested feature is not supported in this
   environment.
```

A DB2 pureScale system does not support delta and incremental backup operations. You must perform a DB2 full backup only in a pureScale system.

Removal of failed backup

When a backup fails and does not roll back, a faulty record might be written to the NetWorker indexes and corrupted data might be saved. To prevent issues during a subsequent restore, remove the faulty record manually from the NetWorker indexes so NMDA cannot restore the failed backup.

Note: The NetWorker man pages, the *EMC NetWorker Command Reference Guide*, and the *EMC NetWorker Administration Guide* describe the NetWorker commands to remove a failed backup.

To remove a failed backup:

1. On the NetWorker server, type the following command to view backup records for the DB2 server in the media database:

   ```
   $ mminfo -v -c client_name.mydomain.com
   ```

   where *client_name.mydomain*.com is the hostname of the DB2 server where the database resides.

   Note: In a cluster environment, use the virtual hostname.

2. Inspect the **mminfo** command output to determine if a save set was created for a failed backup and was not automatically removed by the server. Note the save set ID (ssid).
3. Use the following command to remove the faulty save set from the media database:

   ```
   $ nsrmm -S ssid -d
   ```

   Ensure that you have the NetWorker user privileges required to perform the operation, as described in Table 4 on page 76.

The load libnsrdb2 command

The following table provides the path and the suffix information for the **load libnsrdb2** command.

Use the path and the suffix information to determine the correct shared library for the operating system.

Table 42 Path and suffix for the load libnsrdb2 command

Operating system	Path with suffix
AIX, Linux, Solaris, HP-UX Itanium	/usr/lib/libnsrdb2.so
Microsoft Windows	*NetWorker_install_path*\bin\libnsrdb2.dll

DB2 error messages

The following table describes DB2 SQL error messages that are specific to DB2 operations with NMDA.

The IBM DB2 reference documentation describes the SQL messages.

NOTICE

The table lists identical messages for different causes. View each cause for these multiple listings.

"NetWorker XBSA error messages" on page 409 describes the error messages generated by the NetWorker XBSA interface.

Table 43 DB2 SQL error messages (page 1 of 3)

Error message	Description
SQL1268N A rollforward recovery stopped due to error "SQL1042" while retrieving log file ‹logfile› for database ‹db› on node "0".	The NSR_ENCRYPTION_PHRASES parameter does not contain the datazone pass phrase used to back up the transaction logs. Set the NSR_ENCRYPTION_PHRASES parameter to the proper phrase. "NSR_ENCRYPTION_PHRASES" on page 363 provides details.
SQL2025N An I/O error "3" occurred on media "VENDOR".	You have not registered the client on the NetWorker server to which the NMDA software is backing up. Create a valid client on the NetWorker server. Test to ensure that the connection between the client and server is valid: **save -s** *servername/testfile*
SQL2025N An I/O error "3" occurred on media "VENDOR".	The NSR_CLIENT parameter is set to an invalid client name while running a backup. Set the NSR_CLIENT parameter to the name of the client from which the backup is running. "NSR_CLIENT" on page 358 provides details.
SQL2025N An I/O error "3" occurred on media "VENDOR".	In the vendor configuration file, the NSR_DATA_VOLUME_POOL parameter is set to a pool name that does not exist during a backup. If possible, remove the DB2 vendor.cfg file. "NMDA DB2 parameters" on page 366 provides details. Otherwise, do as follows: • Create a pool to the same name that the NSR_DATA_VOLUME_POOL is set. • Change the value of NSR_DATA_VOLUME_POOL to a valid pool.

Table 43 DB2 SQL error messages (page 2 of 3)

Error message	Description
SQL2025N An I/O error "3" occurred on media "VENDOR".	There is no NMDA license on the server. Each client requires a separate client license. Obtain a valid NMDA license.
SQL2025N An I/O error "3" occurred on media "VENDOR".	For a manual deduplication backup, you have not enabled the deduplication backup attribute in the corresponding NetWorker Client resource. Ensure that you have defined the corresponding NetWorker Client resource on the NetWorker server and you have enabled the deduplication backup attribute.
SQL2025N An I/O error "3" occurred on media "VENDOR".	For a manual or scheduled deduplication backup, the Avamar server is already in maintenance mode before the backup starts. Ensure that the Avamar server is available for the backup operation.
SQL2025N An I/O error "25" occurred on media "VENDOR".	The NSR_ENCRYPTION_PHRASES parameter does not contain the datazone pass phrase used to back up the database. Set the NSR_ENCRYPTION_PHRASES parameter to the proper phrase. "NSR_ENCRYPTION_PHRASES" on page 363 provides details.
SQL2025N An I/O error "25" occurred on media "VENDOR".	The user does not have the restore privilege on the NetWorker server. Add the "recover local data" privilege for the user.
SQL2025N An I/O error "25" occurred on media "VENDOR".	For the restore of a deduplication backup, the Avamar server is already in maintenance mode before the restore starts. Ensure that the Avamar server is available for the restore operation.
SQL2062N An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "0".	Permissions or ownership of debug files are incorrect for the database instance. Ensure that each database instance has a unique debug file name.
SQL2062N An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "4". Table 42 on page 413 provides the correct libnsrdb2 path and suffix information.	The NSR_CLIENT parameter value is an incorrect client name for the restore. Set the NSR_CLIENT parameter to the name of the client from which the restore is running. "NSR_CLIENT" on page 358 provides details.
SQL2062N An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "4". Table 42 on page 413 provides the correct libnsrdb2 path and suffix information.	You have not registered the client on the NetWorker server to which the module is restoring: 1. Create a valid client on the NetWorker server. 2. Test to enure that the connection between the client and server is valid: **save -s** *servername/testfile*
SQL2062N An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "11"	The software did not find matching backups. Restore of the deduplicated save set failed because the save set on the Avamar server was deleted. Restore the save set to the Avamar server before you restore the deduplicated save set.
SQL2062N An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "11".	You did not specify a valid timestamp for the object being restored. Specify a valid timestamp. Or: A database restore failed from one instance to another. The Applications Information attribute in the Client resource is missing instance information. Specify the Applications Information: DB2_R=database_name: db2inst1:db2inst2: "Performing DB2 data restores with the db2 restore command" on page 173 provides details.

Table 43 DB2 SQL error messages (page 3 of 3)

Error message	Description
SQL2062N An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "11". Table 42 on page 413 provides the correct libnsrdb2 path and suffix information.	The NSR_SERVER parameter value is an invalid server name. For example, the server might not exist and you cannot ping it. Set the NSR_SERVER parameter to a valid NetWorker server that has the DB2 server defined as a client. "NSR_SERVER" on page 365 provides details.
SQL2062N An error occurred while accessing media "libnsrdb2.so". Reason code: "25"	The backup of the *database_name* database failed due to a backup request failure. The BRC API call pb_open failed. Only serverless backups support a deduplication backup. Perform a serverless type of backup for the NSM deduplication backup of the database.
SQL2071N An error occurred while accessing the shared library. "c:\progra~1\legato\nsr\bin\libnsrdb2.dll". Reason code: "1".	The missing NMDA DB2 library, libnsrdb2.dll, is not in the correct place as indicated. Reinstall NMDA or place the libnsrdb2.dll library into the correct location as indicated.
SQL2071N An error occurred while accessing the shared library "/usr/lib/libnsrdb2.so". Reason code: "2". Table 42 on page 413 provides the correct libnsrdb2 path and suffix information.	An error message occurs when you use the following software: • 32-bit NMDA to back up a 64-bit database • 64-bit NMDA to back up a 32-bit database Use the correct version of NMDA for the database. For example, you must use a 64-bit version of NMDA to back up a 64-bit database.
SQL2079N On a Windows system, the backup failed for DB2 version 9.*x*. An error was reported by the shared library: "c:\progra~1\legato\nsr\bin\libnsrdb2.dll". Return code: "30".	The stack size is insufficient for the **db2syscs.exe** file. Increase the stack size as follows: 1. Stop the database engine with the **db2stop** command. 2. Use the **db2hdr.exe** utility to increase the stack size to a minimum of 1024. For example: `C:\Program Files\IBM\SQLLIB\BIN> ..\misc\db2hdr db2syscs.exe /s 1024,32` 3. Start the database engine with the **db2start** command.
SQL2079N An error was reported by the shared library during an attempted recovery: "/usr/lib/libnsrdb2.so". Return code "30".	The VENDOROPT parameter is null. Assign the VENDOROPT parameter a value that points to the nmda_db2.cfg file, for example: `db2 update db cfg using vendoropt @/db/nmda_db2.cfg`
SQL2079N An error was reported by the shared library "/usr/lib/libnsrdb2.so". Return code: "30".	The software did not find the configuration file for either backup or recovery. Specify the correct pathname for the NMDA DB2 configuration file (nmda_db2.cfg).

Informix troubleshooting tips and error messages

The following topics provide NMDA Informix troubleshooting tips and error messages.

No dbspaces/blobspaces found to back up or restore

If you try to back up a dbspace or blobspace that does not exist, the savegroup completion message indicates an ON-Bar error:

```
* mars:INFORMIX:/venus/bogus_space onbar returned status of 147
```

The ON-Bar BAR_ACT_LOG file displays a related list of messages:

```
2010-06-25 12:56:24 15612 15606 WARNING: DB/BLOBspace bogus_space does
   not exist.
2010-06-25 12:56:24 15612 15606 ERROR: There are no DB/BLOBspaces to
   backup/restore
```

You might also see these error messages if you try a point-in-time restore to a time before the first dbspace backup for the instance occurred.

To resolve the problem, ensure that you have the correct spelling, pathname, or point-in-time, then retry the backup or restore operation.

Unable to open connection to server

If you try to back up an Informix Dynamic Server instance that does not exist or is in offline mode during the backup, the savegroup completion message indicates an ON-Bar error:

```
* mars:INFORMIX:/venus onbar returned status of 151
```

The ON-Bar BAR_ACT_LOG file displays a related list of messages:

```
2010-06-25 13:07:29 15671 15665 onbar -b -L 0
2010-06-25 13:07:29 15671 15665
ERROR: Unable to open connection to server.
```

To resolve the problem, ensure that you have the correct spelling and pathname for the instance, check that the instance is in online mode, and then retry the backup.

Default value assigned to LTAPEDEV causes failure

Setting the LTAPEDEV configuration parameter in the ONCONFIG file to /dev/null causes logical logs to be erroneously marked as backed up (U-B----). This error occurs when the Dynamic Server switches to the next log before ON-Bar has a chance to send the logical log data to the NetWorker server. With the LTAPEDEV parameter assigned the value /dev/null, you can perform only whole system restores.

If LTAPEDEV is undefined or set to /dev/null in the ONCONFIG file, an ON-Bar logical log backup returns the error code 131 and a message is sent to BAR_ACT_LOG:

```
2010-06-25 10:50:00 12441  12404
ERROR: Unable to start the logical log backup: Log backup to device
   /dev/null not allowed
```

The NetWorker savegroup completion message also returns an error message:

```
--- Unsuccessful Save Sets ---
* mars:INFORMIX:/venus/rootdbs onbar returned status of 131
* mars:INFORMIX:/venus/rootdbs /usr/sbin/nsrdasv exiting.
```

To ensure the successful backup of the logical logs, set the LTAPEDEV parameter in the ONCONFIG file to anything other than /dev/null.

ON-Bar returned codes 2 through 34

The NMDA library libnsrifmx.*xx* or libxbsa.dll returns the error codes 2 through 34.

"NetWorker XBSA error messages" on page 409 describes the error messages generated by the NetWorker XBSA interface.

Lotus troubleshooting tips and error messages

The following topics provide NMDA Lotus troubleshooting tips and error messages.

"Invalid Time Specified" error

If you try to restore files by using the European date format (*dd*/*mm*/*yy*), an "Invalid Time Specified" error appears. The **nsrnotesrc** command is unable to interpret European date formats. When restoring files with the **nsrnotesrc** command, the date specified for the NSR_RECOVER_TIME parameter must be in the American format (*mm*/*dd*/*yy*).

For example, to restore files from a save set that is timestamped 06/26/10 15:08:34, ensure that the NMDA configuration file contains the following parameter settings:

```
NSR_BACKUP_PATHS = /notes/names.nsf
NSR_RECOVER_TIME = "06/26/10 15:08:34"
NSR_SERVER = spain
```

Backup failure due to time conversion

An NMDA Lotus backup fails when there is a problem with the time conversion or when opening a file to obtain the modification times:

- On Windows, the date format and time format (including DateOrder and DateSeparator) are read from the country settings specified on the operating system level.
- On UNIX and Linux, the Domino server ignores these settings in the locale that would affect the format used to display a date or time or both. Instead, the Domino server uses defaults that are coded in the appropriate Lotus Domino .res files.

You can use the following notes.ini settings on the Domino server to overwrite the defaults:

- DateOrder—Can be set to DMY, YMD, MDY
- ClockType—Can be set to 24_HOUR
- DateSeparator—Can be set to an arbitrary string and can be longer than one character, if required
- TimeSeparator—Can be set to an arbitrary string and can be longer than one character, if required

The Domino Administrator help provides more information.

Notes API 525 error

The call to the Notes API function NSFDBOpen() might return a 525 error, for example:

```
NSFDBOpen() failed for D:\Domino\Data\mail\tsk.nsf with error = 525
   (This database is currently in use by another person or process, and
   cannot be accessed at this time. In order to share a Notes database,
   it must be accessed via a Domino Server by all users of the
   database.)
```

Perform the following actions to prevent this error:

- On Windows, do not run the backup through Remote Desktop (Terminal) Services.
- On UNIX or Linux, ensure that the Domino installation directories and data directories do not contain operating system symbolic links.

  To store data files outside of the Domino data directory, use only the Lotus database links or directory links.
- Check with IBM Technical Support about setting NSF_DbCache_Disable=1 in the notes.ini file or issuing the appropriate command to disable caching from the Domino console.

NMC Group Details window displays “command not found”

The NMC Group Details window might display the following messages for the Lotus scheduled backup of a UNIX client:

```
sh: ps: command not found
sh: grep: command not found
```

The messages do not affect the backup of the Domino data. The Domino server generates the messages and passes them to the NetWorker server during the backup.

To prevent the messages, specify the paths of the **grep** command (/usr/bin) and **ps** command (/bin) in the PATH parameter setting in the NMDA configuration file on the UNIX Domino client.

For example, on a SuSE Linux Enterprise Server 10 client, the following PATH setting in the NMDA configuration file includes the **grep** and **ps** command paths:

```
PATH = /opt/ibm/lotus/notes/latest/linux:/usr/bin:/bin
```

Restore problem due to NSR_BACKUP_PATHS=NOTES:

If NSR_BACKUP_PATHS is set to the value NOTES: for a partitioned Domino server or multiple Domino installations on a single UNIX host, NMDA tries to restore the data for all the partitions or Domino installations. This operation can cause data corruption.

Do not set the restore parameter NSR_BACKUP_PATHS to the keyword NOTES:. Instead, set the parameter to a specific pathname.

“NSR_BACKUP_PATHS” on page 373 describes the restore parameter.

Failure of a document-level recovery on a remote Domino server

The NMDA software does not support document-level recovery on a remote Domino server through the **nsrdocrc** command.

On Windows only, NMDA supports document-level recovery of selected (modified) and deleted Notes documents on a remote Domino server through the Notes client program.

If a remote document-level recovery through the Notes client program fails with an "Authentication failure" error, ensure that you meet the following requirements:

- The user that runs the Notes client on the local host is:
  - Listed in the Remote Access attribute in the NetWorker Client resource of the remote Domino server.
  - Granted administrative privileges on the remote Domino server.
- You have configured a NetWorker Client resource on the same NetWorker server for the host where the Notes client program runs.

NMDA Lotus error messages

This topic describes error messages that might appear when you use NMDA for Lotus operations and the possible resolutions for the errors.

The NMDA programs for Lotus backups and restores generate error messages in the following format:

```
error_message_text (function_name | error_type | error_code |
   error_number)
```

where:

- *error_message_text* is the text of the NMDA error message.
- *function_name* is the name of the NMDA function that generated the error message.
- *error_type*, *error_code*, *error_number* are internal numbers representing an error type or error code to report to technical support, reached through EMC Customer Support at EMC Online Support.

Certain fields might be null or 0 if the information was not available at the time when the error message was generated.

"NetWorker XBSA error messages" on page 409 describes the error messages generated by the NetWorker XBSA interface. To ensure that XBSA error and debug messages are included in the NMDA Lotus debug log, set the parameter NSR_XBSA_DEBUG=TRUE as described in "NSR_XBSA_DEBUG" on page 378.

The following table lists error messages generated during Lotus backup and restore operations, in alphabetical order. The table first lists the messages that are common to both backups and restores, then lists the backup messages and restore messages.

Table 44 Error messages from Lotus backups and restores (page 1 of 3)

Error message	Description
Error messages common to both Lotus backups and restores:	
A NW server with a NSR client resource for *hostname* could not be located on the network.	NMDA could not access the NetWorker server for the backup. To resolve the error, ensure that you have spelled the NetWorker server name correctly and the server contains a Client resource for the NMDA client.
System error: error message is not available.	A problem occurred in the XBSA code. To obtain more debugging information about the problem, set NSR_XBSA_DEBUG=TRUE and NSR_DEBUG_LEVEL to a value greater than 3.

Table 44 Error messages from Lotus backups and restores (page 2 of 3)

Error message	Description
The call to NotesInitExtended failed: *Notes_API_error_number*.	NMDA could not initialize the Notes session. To resolve the error, locate the error message for the given *Notes_API_error_number* and perform the appropriate corrective action. If the error number is 421, set the NSR_NOTES_INI_PATH parameter.
The LoadLibrary() call failed.	NMDA could not find or load the libnotes.xx or nnotes.dll library file. To resolve the error, ensure that the Notes_ExecDirectory parameter is set to the directory path containing the Lotus library.
Error messages from Lotus backups:	
Could not open file: *file_name*.	NMDA could not open the file for the backup. To resolve the error, ensure that the file has the correct permissions.
Directory link points to its own directory.	NMDA detected a circular link. To resolve the error, delete the directory link file or correct the path set in the file.
Duplicate entry: *file_pathname*. Ignoring.	NMDA did not add another entry for the file to the backup list because the file was already on the list. This is an informational message only, and does not require a resolution.
Either nothing was specified to backup or system is down.	NMDA could not derive a list of files to back up. To resolve the error, ensure that you configure and perform the manual backup or scheduled backup according to the instructions in this administration guide.
Error finding last backup instance for *file_name*; error = *error_description*.	NMDA could not find an existing backup for the file, and performed a full backup of the file. This error typically occurs during an incremental backup. To resolve the error, ensure that the NetWorker server is up and accessible to the Lotus user on the client computer.
Failed to build exclude list, ignoring exclude list entries.	NMDA could not correctly process the values specified with the NSR_EXCLUDE_FILE or NSR_EXCLUDE_LIST parameter, so nothing was excluded from the backup. To resolve the error, ensure that the parameters specify the correct values.
Failed to open directory link file: *file_name*, errno = *error_number*.	NMDA could not open the file. To resolve the error, ensure that the file has the correct pathname or permissions or both.
Failed to open exclude list file, errno = *error_number*.	The exclude list file did not exist or did not have the correct access permissions. To resolve the error, ensure that you meet the following requirements: • The NSR_EXCLUDE file parameter specifies the correct pathname for the exclude list file. • The exclude list file has the correct permissions.
pathname is not a directory.	The *pathname* in the Domino directory link (.dir) file was not a directory. To correct the error, delete the .dir file from the directory if it is not used.
The password file lookup for the user *Notes_user* failed. The user name may be invalid.	The LOTUS_USER parameter specified an invalid Notes user. To resolve the error, ensure that the LOTUS_USER parameter specifies the name of a Lotus Notes user that is running the server to be backed up.
The ReadFile() call failed.	NMDA could not read the contents of a file to be backed up. To resolve the error, ensure that the file is a readable file.

Table 44 Error messages from Lotus backups and restores (page 3 of 3)

Error message	Description
Error messages from Lotus restores:	
Full backup not found, cannot recover without full backup *file_name*.	NMDA could not find a full backup for the file in the NetWorker index. To resolve the error, ensure that the path for the *file_name* matches the case-sensitive path in the NetWorker index.
No objects found for recover.	Either nothing was specified for recovery or the specified items were removed from the recovery list. To resolve the error, perform one of the following actions: • Specify the required items for recovery. • Set the parameter NSR_RELOCATION_DEST to enable relocated recovery, or set NSR_RECOV_INTERACT to **y** or **r** to enable renaming of the file.
The list of items to recover includes Notes databases. NSR_NO_NOTES_INIT parameter cannot be used when recovering Notes databases.	You specified the NSR_NO_NOTES_INIT parameter for a database recovery. To resolve the error, specify either the NSR_NO_NOTES_INIT parameter or a list of databases (not both) for the recovery. Use the NSR_NO_NOTES_INIT parameter for a disaster recovery only.
The log directory may only be *number* characters in length.	The NSR_LOG_DIR parameter specified a log directory pathname that is too long. To resolve the error, specify a different (shorter) pathname with the NSR_LOG_DIR parameter.
The recover time '-t' was either not specified or contained an invalid value.	During a document-level recovery, you did not specify the mandatory **-t** option correctly with the **nsrdocrc** command. To resolve the error, specify the **-t** *time* option correctly by using the **nsr_getdate()** format for the time. To recover the backup to the current time, use the **-t now** option.

MySQL troubleshooting tips and error messages

The following topics provide NMDA MySQL troubleshooting tips and error messages.

NMDA MySQL backup issues

An NMDA MySQL backup can encounter problems due to one of the following backup issues:

- A MySQL backup fails if you configure the backup as a cluster backup, a snapshot backup, or a probe-based backup that uses the **nsrdaprobe** program.
- A MySQL backup fails if you use the **mysqlbackup** command on the command line to perform the backup. You must use the **nsrdasv** command to perform a manual or scheduled backup:
  - For a manual backup, run the **nsrdasv** command on the command line.
  - For a scheduled backup configured without the wizard, specify the **nsrdasv** command in the Backup Command field in the Client resource.
- A MySQL backup fails if you do not set the required parameters in the NMDA configuration file and the MySQL configuration file, my.cnf. "NMDA MySQL parameters" on page 380 provides details.

For example, you must specify the username and password of the MySQL backup user by setting the required parameters in either the NMDA configuration file or the MySQL configuration file. If you do not specify the backup user credentials in one of the files, the backup fails.

- A MySQL backup fails if you specify the same backup image name as used for a previous backup in either of the following settings and the backup objects (database or tables) differ between the two backups:

  - BACKUP_NAME parameter for a manual backup
  - Save Set attribute in the Client resource for a scheduled backup

  Specify a unique backup name in these settings to ensure that each MySQL backup has a unique name in the NetWorker client file index. The backup also fails if you omit MYSQL:/ from the backup name in these settings.

- If you use NMDA with MEB 3.6.*x* to perform an InnoDB backup, the backup fails if you set any of the following NMDA parameters:

  - MYSQL_INCR_OPTIONS=REDO_LOG_ONLY
  - MYSQL_ONLY_INNODB_OPTIONS=WITH_FRM_ALL
  - MYSQL_ONLY_INNODB_OPTIONS=WITH_FRM_RELATED

  NMDA supports these parameters only with MEB 3.7 or later.

- If you set either of the following parameters for an InnoDB backup with MEB 3.7 or later, the backup fails if you do not run the backup with the required credentials:

  - MYSQL_ONLY_INNODB_OPTIONS=WITH_FRM_ALL
  - MYSQL_ONLY_INNODB_OPTIONS=WITH_FRM_RELATED

  You must run the InnoDB backup as an OS user with write permissions to the parent directory of the MySQL data directory.

- A partial instance backup fails if you set either of the following parameters:

  - MYSQL_LOG_OPTIONS=LOGS_ONLY_BACKUP
  - MYSQL_LOG_OPTIONS=INCLUDE_LOGS

  For example, you can run a partial instance backup by setting MYSQL_DATABASES or MYSQL_INCLUDE. You can only set MYSQL_LOG_OPTIONS for either a stand-alone log backup or a whole instance backup with a log backup.

- A MySQL binary log backup uses only the second value (in this case, INCLUDE_LOGS) for the MYSQL_LOG_OPTIONS parameter if you set the parameter to both of the following values:

  MYSQL_LOG_OPTIONS=LOGS_ONLY_BACKUP, INCLUDE_LOGS

- A MySQL binary log backup fails if binary logging is disabled or the MySQL instance is offline.

NMDA MySQL restore issues

An NMDA MySQL restore can encounter problems due to one of the following restore issues:

- A MySQL restore fails if you use the **mysqlbackup** command on the command line to perform the restore. You must use the **nsrmysqlrc** command to perform a MySQL restore.
- A MySQL restore or recovery might corrupt the InnoDB log files if you do not shut down the database server before a copy back operation that copies data to the MySQL data directory.
- A MySQL recovery fails if you try to use NMDA with MEB 3.6.*x* to recover an InnoDB backup performed with MEB 3.7 or later and any of the following parameter settings:
  - MYSQL_INCR_OPTIONS=REDO_LOG_ONLY
  - MYSQL_ONLY_INNODB_OPTIONS=WITH_FRM_ALL
  - MYSQL_ONLY_INNODB_OPTIONS=WITH_FRM_RELATED
- A MySQL restore might fail due to an error in a copy back operation. In this case, you do not need to restart the restore from the beginning. The prepared backup remains in the NSR_BACKUP_DIR/full directory, so you can just perform a copy back operation to complete the restore of the prepared backup to the data directory.
- Due to an MEB limitation, a restore of a partial InnoDB backup (a backup performed with MYSQL_INCLUDE or the **include** parameter in the MySQL configuration file) does not perform the copy back operation to complete the restore of the prepared backup. The restore operation displays a message about how to complete the restore.

NMDA MySQL error messages

This topic describes error messages that might appear when you use NMDA for MySQL operations and the possible resolutions for the errors.

The NMDA programs for MySQL backups and restores generate error messages in the following format:

```
error_message_text (function_name | error_type | error_code |
   error_number)
```

where:

- *error_message_text* is the text of the NMDA error message.
- *function_name* is the name of the NMDA function that generated the error message.
- *error_type*, *error_code*, *error_number* are internal numbers representing an error type or error code to report to technical support, reached through EMC Customer Support at EMC Online Support.

Certain fields might be null or 0 if the information was not available at the time when the error message was generated.

The following table lists error messages generated during MySQL backup and restore operations, in alphabetical order.

Table 45 Error messages from MySQL backups and restores

Error message	Description
mysqlbackup: ERROR: Log scan was only able to reach 4845560832, but a checkpoint was at 4845561254.	A MySQL incremental backup failed because the database server overwrote part of the circular InnoDB log file before the **ibbackup** process was able to read the log file. To resolve the error, perform one of the following actions: • Ensure that the database server has less load and then rerun the backup. • Reconfigure the database with larger InnoDB log files and then rerun the backup.
Error: operation failed. global_error_count: 2 backup_config.error_count: 0 92824:nsrdasv: The program 'mysqlbackup' terminated with the error code 11.	A MySQL restore of a database failed due to one of the following reasons: • The MySQL database server was not shut down prior to the restore. • The default InnoDB log size was too small. In either case, to resolve the error, complete the following steps: 1. At the mysql command prompt, set **innodb_fast_shutdown=0**. This setting ensures that all the data is flushed and committed to the database and not the logs. For example: mysql› **SET GLOBAL innodb_fast_shutdown=0** 2. Move or delete the ib_logfile file from the mysql directory. For example: **/var/lib/mysql/mv ib_logfile ./tmp/** 3. Only if the InnoDB log size was too small, update the **innodb_log_file_size** value in the MySQL configuration file. 4. Restart the database. The InnoDB engine creates a new set of logs. 5. Perform a MySQL full backup with NMDA.
Since the MYSQL_BACKUP_NAME parameter setting, '*backup_name*', is not an image name, the MYSQL_RESTORE_OPERATION parameter must be set to 'extract'.	A MySQL extract and prepare operation failed because the parameter MYSQL_BACKUP_NAME was set to a backup piece name instread of a image name. With MYSQL_BACKUP_NAME set to the backup piece name, you can perform an extract operation by setting MYSQL_RESTORE_OPERATION to the value “extract”. To perform an extract and prepare operation, change the MYSQL_BACKUP_NAME setting to a valid image name.
The first specified binary log does not match the binary log, '*binary_log_name*', for the last restored backup.	A MySQL current time restore or PIT restore failed because the first binary log specified by MYSQL_BINLOG does not correspond to the binary log in use when the last backup piece was created. Use the **nsrinfo** command with the **–v** option to find out which binary log was in use for the backup.
The nsrmysqlrc program could not find a save set for recovery.	A MySQL current time restore or PIT restore failed due to one of the following reasons: • MYSQL_BACKUP_NAME, NSR_SERVER, or NSR_RECOVER_TIME is set to an incorrect value. (NSR_RECOVER_TIME is used only for a PIT restore.) • The backup specified for the restore does not exist.
Warning! Cannot establish connection to mysql server: Can't connect to local MySQL server through socket '/var/lib/mysql/mysql.sock' (2)	NMDA could not access the MySQL server during an offline MySQL backup. To resolve the error, complete the following steps: 1. Add **no_connection** in the [mysqlbackup] section of the MySQL configuration file, my.cnf. 2. Specify **innodb_log_files_in_group** with the correct number of logs (2 is the default number) in the [mysqld] section of the MySQL configuration file.

Oracle troubleshooting tips and error messages

The following topics provide NMDA Oracle troubleshooting tips and error messages.

Diagnosing NMDA Oracle problems

In addition to "General troubleshooting tips" on page 406, use the following list to troubleshoot problems in running NMDA Oracle backup and restore operations:

- Without NMDA installed on the Oracle Server host, ensure that you can perform a backup and restore by using the **allocate channel t1 type disk** command.
- On UNIX or Linux, ensure that you have linked the correct libnsrora.* library file according to the linking commands in the *EMC NetWorker Module for Databases and Applications Installation Guide*.

  Compare the library file with the libnsrora.* file in the NMDA software package. The two files must be identical. Ensure that you have not linked Oracle to the wrong libnsrora.* or other library file.
- Perform a manual Oracle backup by using NMDA and the proper RMAN script.

  Set the required NMDA parameters in either the RMAN backup script or the **rman send** command on the operating system command line. Appendix A, "NMDA Parameters and Configuration File," describes how to set the parameters. "RMAN scripts for manual backups" on page 129 provides a simple startup RMAN script.

  If the manual backup fails, check the debug files for NMDA and the NetWorker server. "Debug log files" on page 407 provides details.

  If RMAN fails with the following error, ensure that NMDA and Oracle have the same bitness. Refer to the RMAN user guide for details on how to test that you have integrated the media management library correctly:

  ```
  ORA-19554: error allocating device, device type: SBT_TAPE, device
  name:
  ORA-27211: Failed to load Media Management Library
  Additional information: 25
  ```

  If RMAN fails with another error, check the RMAN output or the trace file as described in "The trace option of the backup command" on page 426.
- Perform a scheduled Oracle backup by using NMDA and the proper RMAN scripts.

  In the working RMAN manual backup script, add the **connect target** and **connect rcvcat** commands, as described in "RMAN scripts for scheduled backups" on page 130.
- If an Oracle RMAN session fails on Windows, check for any **nsrsbtcn.exe** processes that are still running and terminate the processes manually in the Task Manager.

RMAN error messages

Note: If the scheduling portion of an Oracle scheduled backup succeeds but the backup fails, error messages and debug information might be generated in the locations described in this topic.

RMAN stores information and RMAN error messages in the log file specified with the **msglog** option. Review the RMAN information in this log file after each backup.

To specify the name of the RMAN log file:

- For a manual Oracle backup, specify the **msglog** option in the **rman** command on the command line:

  ```
  rman target ... rcvcat ... msglog file_name
  ```

- For a scheduled Oracle backup, specify the **msglog** option in the parameter NSR_RMAN_ARGUMENTS in the NMDA configuration file. "NMDA Oracle parameters" on page 387 provides details.

The Oracle error messages guide describes specific RMAN error messages and recommended courses of action.

Note: During a backup on AIX or Windows, if an NMDA parameter is set to an invalid value, the resulting error message might be truncated in the RMAN output due to an Oracle RMAN limitation.

The trace option of the backup command

Set the **trace** option of the RMAN **backup** command to the value 0, 1, or 2. The default value of **trace** is 0.

The output of **trace** appears in the Oracle sbtio.log file.

The output of **trace** also appears in the log file nmda_oracle.messages.raw, located in the directory specified by the NSR_DIAGNOSTIC_DEST parameter (if set) or in the default directory:

- On UNIX systems: /nsr/apps/logs
- On Windows systems: *NetWorker_install_path*\apps\logs, where *NetWorker_install_path* is the root directory of the NetWorker installation path

These log files contain NMDA error messages when a failure occurs in the libnsrora.*x* library that is loaded by Oracle (known as the Media Management Library in Oracle documentation).

NMDA writes error messages in the nmda_oracle.messages.raw file in a language-independent binary form, readable by the **nsr_render_log** program only.

The *EMC NetWorker Administration Guide* describes how to use the **nsr_render_log** program to read a language-independent binary file, such as nmda_oracle.messages.raw.

The following table outlines the conditions that are traced when the **trace** option is set to each of the three valid values.

Table 46 Trace option values and conditions traced

Trace value	Conditions traced
0 (default)	All error conditions.
1	• All error conditions. • Entry and exit for each System Backup to Tape (SBT) function (the NMDA implementation of the Oracle SBT interface).
2	• All error conditions. • Entry and exit for each SBT function (the NMDA implementation of the Oracle SBT interface). • Values of all function parameters. • First 32 bytes of each read/write buffer.

Canceling unresponsive Oracle backups on UNIX

An Oracle backup on UNIX might not respond to the usual canceling steps, for example, if the backup is waiting for a device to be mounted. In this case, use the following steps.

Note: When you complete these steps, NMDA does not remove the backup entries from the NetWorker index. The NetWorker index and RMAN catalog might become unsynchronized, but this occurrence does not cause issues for subsequent RMAN operations.

1. Include the **set command id to '***xxx***'** command in the RMAN backup script used for the Oracle backup. Otherwise, the query in the next step will fail. Example 8 on page 130 provides a sample script with the command.
2. Run the following query in the Oracle **sqlplus** program to determine the Oracle process ID that corresponds to each RMAN channel:

   ```
   select spid, client_info from v$process p, v$session s where
   p.addr=s.paddr and client_info like '%id=%';
   ```

3. Type the **kill** command to cancel the backup process:

   ```
   kill -9 pid
   ```

   where *pid* is the appropriate backup process ID.

Canceling unresponsive Oracle backups on Windows

An Oracle backup on Windows might not respond to the usual canceling steps, for example, if the backup is waiting for a device to be mounted. In this case, use the following procedure.

To cancel an unresponsive Oracle backup on Windows, stop the **nsrsbtcn.exe** process in Task Manager.

Note: When you complete this procedure, NMDA does not remove the backup entries from the NetWorker index. The NetWorker index and RMAN catalog might become unsynchronized, but this occurrence does not cause issues for subsequent RMAN operations.

NMDA Oracle error messages

This topic describes error messages that might appear when you use NMDA for Oracle operations, and the possible resolutions for the errors:

- "Error messages from Oracle backups and restores" on page 428
- "Error messages from the nsroraadmin program" on page 437
- "Error messages from the nsrorainfo program" on page 439

Error messages from Oracle backups and restores

The following table lists error messages generated during Oracle backups and restores with NMDA, in alphabetical order. The table first lists the libnsrora.*xx* messages from nonproxy operations, then lists the libnsrora.*xx* messages from proxy operations and then the messages from scheduled backups (**nsrdasv** process).

The error messages appear in the following format:

```
function_name: error_message (error_type:error_code:error_number)
```

where:

- *function_name* is the name of the NMDA function that produced the error.
- *error_message* is the text of the error message, as shown in the table.
- *error_type*, *error_code*, *error_number* are internal numbers that represent an error type or error code. The significance for the user is as follows:
  - If *error_code* is 1, the system is out of memory.
  - If *error_code* is 3, 13, or 17, a code-level error occurred. Report the error message to Technical Support.

Table 47 Error messages from Oracle backups and restores (page 1 of 9)

Error message	Description
Error messages from nonproxy backups and restores:	
A connection to NW server '*server*' could not be established because '*reason*'.	NMDA could not connect to the NetWorker client file index due to the given reason. The client might not be configured as a client on the server. Take the corrective action suggested by the error message.
Could not create the LNM index lock file '*file_name*' (*errno*)	NMDA failed to create the lock file required for an index deletion operation. Report the error number (*errno*) to Technical Support.
Could not decode the 'sf_check' value: xdrs = 0x*value*	This error is an internal XDR error due to a network read or write operation. Report the error to Technical Support.
Could not decode the 'sf_magic' value: xdrs = 0x%*value*	This error is an internal XDR error due to a network read or write operation. Report the error to Technical Support.
Could not decode the 'sf_more' flag: xdrs = 0x*value*	This error is an internal XDR error due to a network read or write operation. Report the error to Technical Support.

Table 47 Error messages from Oracle backups and restores (page 2 of 9)

Error message	Description
Could not locate the LNM save file '*backup_piece_name*' on server '*server*'.	NMDA could not locate an index record for the backup piece. The index record is probably missing. Use the **mminfo** and **nsrinfo** commands to verify the status of the index record.
Could not locate the LNM save time '*save_time*' on server '*server*'.	NMDA could not locate an index record for the save time in the client file index. The index record is probably missing. Use the **mminfo** and **nsrinfo** commands to verify the status of the index record.
Could not lock '*file_name*' for index deletion. There were *number* attempts. (*errno*)	NMDA was able to create the lock file required for an index deletion operation, but could not lock the file after the given number of attempts. Report the error number (*errno*) to Technical Support.
Error in mmdb lookup by time: *reason*	A lookup in the media database failed for the given reason. Use the **mminfo** command to verify the status of the media database record. Take the corrective action suggested by the error message.
Exceeded the number of retries. The NetWorker server may be down or unreachable.	NMDA could not contact the NetWorker index service **nsrindexd**. This issue was probably due to the NetWorker services being shutdown. Restart the NetWorker services on the server, as required.
Exceeded the number of retries for nsr_init(). The NetWorker server may be down or unreachable.	After a maximum of five attempts, NMDA failed to call the NetWorker core function, nsr_init(). This issue was probably due to the NetWorker services being shutdown. Restart the NetWorker services on the server, as required.
Exceeded the number of retries for nsr_start(). The NetWorker server may be down or unreachable.	After a maximum of five attempts, NMDA failed to call the NetWorker core function, nsr_start(). This issue was probably due to the NetWorker services being shutdown. Restart the NetWorker services on the server, as required.
Invalid browse and retention policies. Values Ignored.	The NSR_SAVESET_BROWSE and NSR_SAVESET_RETENTION parameters both had invalid time values. Ensure that the parameters NSR_SAVESET_BROWSE and NSR_SAVESET_RETENTION in the RMAN script both have valid values in the NetWorker date format.
Invalid browse policy *browse_time*. Value Ignored.	The NSR_SAVESET_BROWSE parameter had an invalid time value, *browse_time*. Ensure that the parameter NSR_SAVESET_BROWSE in the RMAN script has a valid value in the NetWorker date format.
Invalid KEY word	The syntax of the string in the RMAN **send** command was incorrect. "The send command" on page 398 provides the correct **send** command syntax.
Invalid retention policy: *retention_time*. Value Ignored.	The NSR_SAVESET_RETENTION parameter had an invalid time value, *retention_time*. Ensure that the parameter NSR_SAVESET_RETENTION in the RMAN script has a valid value in the NetWorker date format.
NSR_DATA_VOLUME_POOL*n* is not set.	Multiple copies of the backup data were requested, but the required NSR_DATA_VOLUME_POOL parameters were not set. In the message, *n* was replaced by a number corresponding to the missing pool parameter. When multiple copies of backup data are requested, set the required NSR_DATA_VOLUME_POOL parameters. "NSR_DATA_VOLUME_POOL" on page 359 provides details.
ORA-19511: Error received from media manager layer, error text: Could not create the NWORA resource lock file (13) (103:105:13)	An NMDA backup failed because a valid NWORA resource file does not exist or is not available. If you do not use the wizard to configure a scheduled backup with save set bundling, use the **nsroraadmin** command to create a valid NWORA resource file, according to instructions in Chapter 2, "Backup Configuration" or Chapter 9, "Snapshot Backups and Restores."

Table 47 Error messages from Oracle backups and restores (page 3 of 9)

Error message	Description
Oracle pools are not supported	NMDA does not support Oracle pools. NMDA supports NetWorker pools only. Remove the **pool** option of the **backup** command in the RMAN script or set the **pool** option to zero. "The pool option of the backup command" on page 398 provides details.
'*string*' should be in format: KEY=(*xxxxx*)	The syntax of the string in the RMAN **send** command was incorrect. "The send command" on page 398 provides the correct **send** command syntax.
The ASDF body could not be unwrapped.	The incoming recover stream of data could not be decoded due to a possible network error or data corruption. Report the error to Technical Support.
The backup file already exists: *backup_piece_name*	NMDA could not complete the backup because the backup piece name already existed in the NetWorker client file index. Change the **format** option string of the RMAN command to produce a unique backup piece name, or remove obsolete backup pieces. Then restart the backup operation.
The call to nsr_init() failed with the message: *reason*	A call of the NetWorker core function, nsr_init(), failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The call to nsr_start() failed with the message: *reason*	A call of the NetWorker core function, nsr_rtart(), failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The current time could not be obtained (*errno*).	NMDA could not obtain the current time due to an operating system error. Report the operating system error (*errno*) to the appropriate vendor.
The data could not be XDR'd from the stream.	The incoming recover stream of data could not be decoded due to a possible network error or data corruption. Report the error to Technical Support.
The function mm_retrieve() failed with the error: *reason*	During a restore, a call of the NetWorker core function, mm_retrieve(), failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The function nsr_bind_recov_mm() failed with the error: *reason*	During a restore, a call of the NetWorker core function, nsr_bind_recov_mm(), failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The function nsr_end() failed with the error message: *reason*	A call of the NetWorker core function, nsr_end(), failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The function nsr_rstart() failed with the error: *reason*	During a restore, a call of the NetWorker core function, nsr_rstart(), failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The function sbtinit2() has already been called.	This error is an internal error due to Oracle calling the function sbinit2() twice. Report the error to Technical Support.
The functions sbtinit() or sbtinit2() have not been called.	This error is an internal error due to Oracle not calling the two SBT initialization routines. Report the error to Technical Support.
The index entry failed the cross check: cfx_name(*backup_piece_name*) save_time(*save_time*)	An index lookup located the entry in the client file index but not in the media database. Restart the NetWorker services, and use the **mminfo** and **nsrinfo** commands to verify the backup information in the indexes. Run the **nsrck** program to resolve any corruption of the indexes.

Table 47 Error messages from Oracle backups and restores (page 4 of 9)

Error message	Description
The lookup of '*backup_piece_name*' on server '*server*' failed - '*reason*'	NMDA could not locate *backup_piece_name* in the indexes due to the *reason*. The indexes might be corrupted. Run the **nsrck** program to resolve any corruption of the indexes.
The name of the NSR client could not be determined.	NMDA could not determine the name of the NetWorker client. Set the parameter NSR_CLIENT to the NetWorker client name by using the **send** command.
The name of the NSR server could not be determined.	NMDA could not determine the name of the NetWorker server. Set the parameter NSR_SERVER to the NetWorker server name by using the **send** command.
The NSR client name could not be determined.	NMDA could not determine the name of the NetWorker client. Set the parameter NSR_CLIENT to the NetWorker client name by using the **send** command.
The NSR server name could not be determined.	NMDA could not determine the name of the NetWorker server. Set the parameter NSR_SERVER to the NetWorker server name by using the **send** command.
The NSR_CLIENT parameter was not set.	NMDA could not determine the name of the NetWorker client. Set the parameter NSR_CLIENT to the NetWorker client name by using the **send** command.
The NSR_SERVER parameter was not set.	NMDA could not determine the name of the NetWorker server. Set the parameter NSR_SERVER to the NetWorker server name by using the **send** command.
The NW authentication for client '*client*' was refused by server '*server*' because '*reason*'.	NMDA could not obtain the required authentication to connect to the NetWorker client file index due to the given reason. The client might not be configured as a client on the server. Take the corrective action suggested by the error message.
The NW client has not been set.	NMDA could not determine the name of the NetWorker client. Set the parameter NSR_CLIENT to the NetWorker client name by using the **send** command.
The NW server does not have a valid NMDA proxy copy license.	The NetWorker server tried a proxy operation without the required license. Ensure that the NetWorker server has the required license for the proxy operation.
The NW server has not been set.	NMDA could not determine the name of the NetWorker server. Set the parameter NSR_SERVER to the NetWorker server name by using the **send** command.
The NWORA file ID could not be XDR'd. xdrm: 0x*value* NWORA fid: 0x*value* ssid: 0x*value* ssoff: 0x*value*	This error is an internal XDR error due to a network read or write operation. Report the error to Technical Support.
The record obtained has the wrong save time '*save_time1*'. The save time queried was '*save_time2*'.	NMDA located an index record in the client file index, but the index record had an unexpected save time. The indexes might be corrupted. Restart the NetWorker services, and run the **nsrck** program to resolve any corruption of the indexes.
The removal of SSID '*save_set_id*' failed with error: *reason*	An index deletion operation failed for the given reason. Use the **mminfo** and **nsrinfo** commands to verify the status of the index record. If required, report the error to Technical Support.

Table 47 Error messages from Oracle backups and restores (page 5 of 9)

Error message	Description
The savefile_fini() call failed. *reason*	During a restore, a call of the NetWorker core function, savefile_fini(), failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The sfhead could not be XDR'd.	This error is an internal XDR error due to a network write operation. Report the error to Technical Support.
The SS browse time is not in the future: current time: *current_time* browse: *browse_time*	The specified browse policy time was in the past. This issue might be due to a problem with the operating system time setting. Ensure that the browse policy time is set correctly. If required, ensure that the operating system time is set correctly.
The SS retention time is not in the future: current time: *current_time* retention: *rentention_time*	The specified retention policy time was in the past. This issue might be due to a problem with the operating system time setting. Ensure that the retention policy time is set correctly. If required, ensure that the operating system time is set correctly.
The UNIX attributes could not be XDR'd. xdrm: 0x*value* ua: 0x*value*	This error is an internal XDR error due to a network read or write operation. Report the error to Technical Support.
Error messages from proxy backups and restores:	
Attempted to restore file '*file_name*' to raw device '*device_name*'.	The software tried to perform a proxy restore of a regular file to a raw device. NMDA does not support this type of restore. Do not try to restore a regular file to a raw device.
Attempted to restore raw device '*device_name*' to file '*file_name*'.	The software tried to perform a proxy restore of a raw device file to a regular file. NMDA does not support this type of restore. Do not try to restore a raw device file to a regular file.
Cannot back up object *object_name* with proxy copy.	The RMAN **backup** command included the **proxy only** option, but the object *object_name* did not reside on a primary storage device that NSM supports. When the **backup** command includes the **proxy only** option, ensure that the object *object_name* resides on a primary storage device that NSM supports.
Could not find the nsrsnapck binary.	During an index removal for a proxy backup, NMDA could not locate the **nsrsnapck** binary, which is probably in a nondefault location. Ensure that the parameter NSR_NWPATH is set correctly.
Could not lstat - *file_name*	The lstat() system call failed. The file *file_name* either did not exist or had invalid permissions. Ensure that the file is an existing file with valid permissions.
Could not lstat secondary link - *file_name*	The lstat() system call failed. The file *file_name* was a symbolic link that pointed to a file that either did not exist or had invalid permissions. Ensure that the symbolic link points to an existing file with valid permissions.
Could not obtain NSR_ORACLECAT_MODE from NWORA resource file.	The error was due to one of the following conditions: • The NWORA resource file does not exist. • The NWORA resource file has incorrect permissions. • The NWORA resource file is corrupted. Based on the condition, perform one of the following actions: • If the NWORA resource file does not exist, create the file. • Ensure that the NWORA resource file has correct permissions. • If the NWORA resource file is corrupted, re-create the file. "NWORA resource file" on page 324 provides details.

Table 47 Error messages from Oracle backups and restores (page 6 of 9)

Error message	Description
Could not read link - *pathname*	A proxy backup failed due to the *pathname* that was an invalid symbolic link. Before a proxy backup, ensure that any symbolic link is a valid link.
Error creating staging directory '*directory*'.	During a proxy restore of a regular file, the permissions of the destination directory were possibly invalid. NMDA was unable to create the required staging subdirectory, .nworapc. Ensure that the destination directory has valid permissions for a proxy restore.
Invalid source path argument	A proxy backup failed due to an invalid source pathname. Perform a proxy backup with a valid source pathname only.
nsrsnapck_binary_name process failed with error - *reason*	During an index removal for a proxy backup, the **nsrsnapck** binary failed. The binary name is **nsrsnapck** on UNIX systems and **nsrsnapck.exe** on Windows systems. Report the error to Technical Support.
Path *pathname* is too long.	A proxy backup failed because the given pathname exceeded the limit of 1,024 bytes. Ensure that any pathname involved in a proxy backup does not exceed 1,024 bytes.
pb_init() failed with (*reason*): invalid BRCAPI version	The version number of the BRC API that was reported by NSM was corrupted. Report the error to Technical Support.
Proxy copy is not supported.	You tried to perform a proxy operation on a system that NMDA does not support for proxy operations. Do not try a proxy operation on an unsupported system. The *EMC NetWorker Software Compatibility Guide* at EMC Online Support describes the supported systems.
The BRC API did not return an error string for the SBTPC object: *object_name*	An unknown error occurred during a BRC API function call by NSM. Report the error to Technical Support.
The BRC status of logical object '*file_name*' was failure: *file_status*	NSM reported a failure during a proxy backup of the file *file_name*. Report the error to Technical Support.
The call to pb_environment() failed with error: *reason*	During a proxy operation, a pb_environment() function call failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The call to pb_open() failed with error: *reason*	During a proxy operation, a pb_open() function call failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The call to pb_prepare() failed with error: *reason*	During a proxy operation, a pb_prepare() function call failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The call to pb_status() failed for object '*object_name*' with the error: *reason*	During a proxy operation, a pb_status() function call failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The call to pb_status() for object '*object_name*' failed with error: *reason*	During a proxy operation, a pb_status() function call failed due to the given reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The canonical OS file name path is invalid: *file_name*	The operating system *file_name* specified for a proxy operation was not a valid pathname. Ensure that the file pathname specified for a proxy operation is a valid full pathname that is not a directory.
The data source is neither a file or a RAW volume - *file_name*	The file *file_name* involved in a proxy backup was not recognized as a regular file or raw volume. For proxy backups, NMDA supports only regular files and raw volumes. Ensure that *file_name* is either a regular file or raw volume, as required for proxy backups.

Table 47 Error messages from Oracle backups and restores (page 7 of 9)

Error message	Description
The destination does not have the same terminating name as the source '*device_name*'.	You tried to perform a proxy restore of a raw device to a location with a different basename from the backed-up source. For example, c1t2d0s2 is the basename (or terminating name) of /dev/rdsk/c1t2d0s2. Perform a proxy restore of the raw device to a location with the same basename as the backed-up source.
The file being recovered could not be found in its staging location: *file_name*	During a proxy restore, an error occurred at the point where the file *file_name* was to be moved from the staging directory .nworapc to the destination directory. Ensure that there are no permission problems or other problems with the destination directory and the staging directory .nworapc, and then restart the proxy restore. If the error recurs, report the error to Technical Support.
The file '*file_name*' cannot be removed from the staging directory (*errno*).	During a proxy restore of the file *file_name*, a file with the same name was found in the .nworapc subdirectory, probably left there by a previous failed restore. The *errno* is the error number from the failure to remove the existing file. Remove the file *file_name* from the .nworapc subdirectory, and restart the proxy restore.
The NMDA BRCAPI version *version* is outside the range supported by the BRC service: *earliest_version* - *latest_version*	NMDA does not support the NSM release that was used for a proxy operation. Ensure that a supported NetWorker client release with the NSM feature is installed. The *EMC NetWorker Software Compatibility Guide* at EMC Online Support describes the supported releases.
The NWORA resource file does not exist. Please create it with nsroraadmin.	A proxy backup failed because the NWORA resource file did not exist. Create the NWORA resource file by using the **nsroraadmin** command, and restart the proxy backup. "NWORA resource file" on page 324 provides details.
The NWORA resource lock file does not exist. Please create it by running 'nsroraadmin -r list'	A proxy backup failed because the NWORA resource lock file did not exist. Create the NWORA resource lock file by using the **nsroraadmin -r list** command, and restart the proxy backup. "NWORA resource file" on page 324 provides details.
The NWORA resource NSR_ORACLECAT_MODE is in the 'undetermined' state.	In the NWORA resource file, NSR_ORACLECAT_MODE was set to the default value of undetermined. Set the value of NSR_ORACLECAT_MODE to either enabled or disabled (as required) by using the **nsroraadmin** command.
The object '*file_name*' is not a file.	A proxy backup failed because the file *file_name* is not a datafile. The file is neither a raw file nor a regular file. Perform a proxy backup of a supported type of datafile only.
The ORACLE_SID must be set when performing proxy copy backups.	During a scheduled proxy backup, the parameter ORACLE_SID was not set in the configuration file. In the configuration file, set the parameter ORACLE_SID to the SID value of the Oracle database.
The OS file name has been specified multiple times by Oracle: *file_name*	This error is an internal Oracle error due to Oracle specifying the same file name twice during a proxy operation. Report the error to Technical Support.
The parameter file cannot be open: *file_name*	The configuration file specified by the parameter NSR_PROXY_PFILE could not be opened. The file must contain NSM parameter settings for a proxy backup or restore. Ensure that the value specified by the parameter NSR_PROXY_PFILE is a valid pathname of the configuration file.
The pb_cancel() call for object '*object_name*' returned the error message: *error*	The pb_cancel() function call failed during a proxy operation. Report the error to Technical Support.
The pb_inquiry() call failed for object '*object_name*': *error*	The pb_inquiry() function call failed during a proxy operation. Report the error to Technical Support.

Table 47 Error messages from Oracle backups and restores (page 8 of 9)

Error message	Description
The pb_inquiry() for object '*object_name*' failed because: *error*	The pb_inquiry() function call failed during a proxy operation. Report the error to Technical Support.
The pb_inquiry() of object '*object_name*' returned error: *error*	The pb_inquiry() function call failed during a proxy operation. Report the error to Technical Support.
The pb_restore() for object '*object_name*' failed with error: *error*	The pb_restore() function call failed during a proxy operation. Report the error to Technical Support.
The pb_save() of object '*object_name*' returned error: *error*	The pb_save() function call failed during a proxy operation. Report the error to Technical Support.
The pb_snapshot() call for object '*object_name*' failed with error: *error*	The pb_snapshot() function call failed during a proxy operation. Report the error to Technical Support.
The restore destination path is not valid: *file_name*	During a proxy restore operation, NMDA found the specified restore destination, *file_name*, to be invalid. Ensure that the specified restore destination is a valid pathname.
The restore operation for the file failed for an unknown reason: *file_name*	During a proxy restore, an error occurred at the point where the file *file_name* was to be moved from the staging directory .nworapc to the destination directory. Ensure that there are no permission problems or other problems with the destination directory and the staging directory .nworapc, and then retry the proxy restore. If the error occurs again, report the error to Technical Support.
The SBTPC object could not determine the destination of the restore.	During a proxy restore operation, NMDA could not determine where to restore the file. Report the error to Technical Support.
The SBTPC object is not in the PB_TYPE_PREPARE state: *object_name*	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object *object_name*. Report the error to Technical Support.
The SBTPC object is not in the SBTPCSTATUS_NOTREADY state: *object_name*	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object *object_name*. Report the error to Technical Support.
The SBTPC object '*object_name*' failed with the error message: *reason*	The proxy backup or restore of a file failed during an NSM operation, for the given *reason*. Report the error to Technical Support.
The SBTPC object '*object_name*' is entering the SBTPCSTART backup state but its BRC type is: *type*	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object *object_name*. Report the error to Technical Support.
The SBTPC object '*object_name*' is entering the SBTPCSTART restore state but its BRC type is: *type*	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object *object_name*. Report the error to Technical Support.
The SBTPC object '*object_name*' is entering the SBTPCSTART state but its status is: *status*	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object *object_name*. Report the error to Technical Support.
The SBTPC object '*object_name*' is leaving the BRC prepare state but its status is: *status*	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object *object_name*. Report the error to Technical Support.

Table 47 Error messages from Oracle backups and restores (page 9 of 9)

Error message	Description
The SBTPC object '*object_name*' is leaving the BRC save state but its status is: *status*	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object *object_name*. Report the error to Technical Support.
The SBTPC object '*object_name*' is leaving the BRC snapshot state but its status is: *status*	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object *object_name*. Report the error to Technical Support.
The SBTPC object '*object_name*' was aborted by the BRC service. Please check the PowerSnap logs for an explanation.	NSM terminated the proxy operation. Check the snapshot logs for a possible reason for the termination.
The sbtpccommit() function was called during restore.	This error is an internal Oracle error that occurred during a proxy restore. Report the error to Technical Support.
The staging directory '*directory*' has invalid permissions (*errno*).	During a proxy restore, NMDA was unable to write to the staging directory, *directory*. The *errno* is the error number from the function call that failed. Ensure that the staging directory has valid permissions for a proxy restore.
There are no SBTPC objects that have not returned their status.	This error is an internal error during a proxy operation due to Oracle expecting more files to be processed whereas NMDA has completed the file processing. Report the error to Technical Support.
This backup piece name is already used in the SBTPC session: *backup_piece_name*	This error is an Oracle error due to Oracle specifying the same backup piece name twice during a proxy operation. Report the error to Technical Support.
Error messages from scheduled backups (nsrdasv):	
client: WARNING! The NWORA resource file 'save' process output error messages. *client*: Please check the save log file for more information: *log_file*	The NWORA resource file could not be backed up after a successful RMAN backup. Analyze the *log_file* and if the log file includes an error message, take the corrective action suggested by the error message.
ORACLE_HOME is not defined. Cannot start RMAN.	You have not set ORACLE_HOME in the configuration file. Set ORACLE_HOME in the configuration file.
The backup config did not contain a *string*.	The **nsrdasv** program was run with the **-C** option, but the Backup Config attribute was not set in the Client resource. Remove this Client resource, and re-create the Client resource by using the backup configuration wizard.
The NSR client resource for *client_name* does not contain any backup configuration.	The **nsrdasv** program was run with the **-C** option, but the Backup Config attribute was not set in the Client resource. Remove this Client resource, and re-create the Client resource by using the backup configuration wizard.
The temporary file '*rman_script_path*' could not be created (*errno*).	The scheduled backup binary, **nsrdasv**, could not create the file *rman_script_path* to write the RMAN script generated by the backup configuration wizard. Ensure that the root user on UNIX or the Windows Administrator has "write" permissions on the directory path of the *rman_script_path* file.

Error messages from the nsroraadmin program

The following table lists error messages generated by the **nsroraadmin** program in alphabetical order.

The error messages appear in the following format:

```
nsroraadmin: error_message
```

where *error_message* is the text of the error message as shown in the table.

Table 48 Error messages from the nsroraadmin program (page 1 of 3)

Error message	Description
Command line arguments are not understood.	The **nsroraadmin** command included one or more invalid options. Use the **nsroraadmin** command with the correct options. "Configuring the NWORA resource file with the nsroraadmin program" on page 329 provides details.
Could not create the NWORA resource file (*errno*)	The **nsroraadmin** command could not create the NWORA resource file, possibly due invalid permissions. Ensure that valid permissions exist to enable the **nsroraadmin** command to create the NWORA resource file. "NWORA resource file" on page 324 provides details.
Could not create the NWORA resource lock file (*errno*)	The **nsroraadmin** command could not obtain the required lock file in the /nsr/tmp or *NetWorker_install_path*\tmp directory. The command needs the lock file tp access the NWORA resource file. Report the error to Technical Support.
Could not open resource file '*file_name*' (*errno*).	The **nsroraadmin** command could not open the NWORA resource file, possibly due invalid permissions. Verify that the NWORA resource file exists and has valid permissions. If required, create or repair the file by using the **nsroraadmin** command, or modify the file permissions.
No command line parameters are set.	The **nsroraadmin** command did not include the correct options. Use the **nsroraadmin** command with the correct options.
NSR_ORACLECAT_MODE can only be set to 'enabled', 'disabled' or 'undetermined'.	The **nsroraadmin** command included a NSR_ORACLE_CAT_MODE parameter resource setting other than enabled, disabled, or undetermined. In the **nsroraadmin** command, set the NSR_ORACLE_CAT_MODE parameter resource to enabled or disabled for snapshot backups. "Configuring the NWORA resource file with the nsroraadmin program" on page 329 provides details.
NSR_REMOVE_ON_FAILURE can only be set to 'TRUE' or 'FALSE'.	The **nsroraadmin** command included a NSR_REMOVE_ON_FAILURE parameter resource setting other than TRUE or FALSE. In the **nsroraadmin** command, set the NSR_REMOVE_ON_FAILURE parameter resource to either TRUE or FALSE only.
NWORA parameter resources must be specified in the 'ResourceName ResourceValue' format.	The **nsroraadmin** command did not include the name and value of the NWORA parameter resource in the correct format. In the **nsroraadmin** command, specify the name and value of the NWORA parameter resource in the correct format.
NWORA SID resource must be specified when doing deletion.	The **nsroraadmin** command with the **-r delete** option did not include the SID value of an Oracle database. In the **nsroraadmin** command with the **-r delete** option, specify the correct SID value.
The '-r' flag cannot be set multiple times.	The **nsroraadmin** command contained more than one **-r** option. Use the **nsroraadmin** command with only one **-r** option.
The '-r' option requires an NWORA resource specification.	The **nsroraadmin** command with the **-r** option did not include the required resource specification. In the **nsroraadmin** command with the **-r** option, specify the required resource name and resource value.

Table 48 Error messages from the nsroraadmin program (page 2 of 3)

Error message	Description
The '-r' option requires either an 'add', 'update', 'list' or 'delete' option.	In the **nsroraadmin** command, the **-r** option did not include one of the required keywords: **add**, **update**, **list**, or **delete**. In the **nsroraadmin** command, include one of the required keywords with the **-r** option. "Configuring the NWORA resource file with the nsroraadmin program" on page 329 provides details.
The first NWORA resource is not a header (*errno*).	The NWORA resource file is probably corrupted. Verify the contents of the NWORA resource file. If required, repair the resource file by using the **nsroraadmin** command.
The NWORA resource file does not contain the NSR_NWPATH resource.	The NWORA resource file does not contain the mandatory NSR_NWPATH parameter resource. The file might be corrupted. Verify the contents of the NWORA resource file. If required, repair the resource file by using the **nsroraadmin** command.
The NWORA resource file does not exist.	The NWORA resource file does not exist. Create the NWORA resource file by using the **nsroraadmin** command.
The NWORA resource named '*resource_name*' is not found.	The **nsroraadmin** command specified the name of a resource that does not exist in the NWORA resource file. In the **nsroraadmin** command, specify a valid resource name from the NWORA resource file.
The NWORA resource parameter list can only contain one entry.	The NWORA resource file includes multiple values for a resource, which NMDA does not support. The file might be corrupted. NMDA does not support manual editing of the file. Repair the NWORA resource file by using the **nsroraadmin** command.
The NWORA resource parameter list contains the invalid element '*resource_name*'.	The NWORA resource file contains an invalid resource name. The file might be corrupted. NMDA does not support manual editing of the file. Repair the NWORA resource file by using the **nsroraadmin** command. "Configuring the NWORA resource file with the nsroraadmin program" on page 329 provides details.
The NWORA resource parameter list for a SID requires the *item1*, *item2* and *item3* information.	The **nsroraadmin** command for creating or updating an NWORA SID resource did not include the required items. In the **nsroraadmin** command for creating or updating an NWORA SID resource, include the required items.
The NWORA resource '*resource_name*' is not a SID resource.	The **nsroraadmin** command with the **-r delete** option did not include a valid name of an NWORA SID resource. In the **nsroraadmin** command with the **-r delete** option, specify a valid name of an NWORA SID resource.
The NWORA resource specified is not supported: *resource_name* = *resource_value*	The **nsroraadmin** command included an invalid name or invalid value for an NWORA parameter resource. In the **nsroraadmin** command, specify a valid name and valid value for an NWORA parameter resource. "NWORA parameter resources" on page 324 provides details.
The NWORA SID resource for '*sid_value*' already exists.	The **nsroraadmin** command tried to add an NWORA SID resource that already existed. In the **nsroraadmin** command, specify the values for a new NWORA SID resource.
The SID token 'connect' is an empty string.	The **nsroraadmin** command did not include the required pathname of the RMAN connection file with the **connect** keyword. In the **nsroraadmin** command, specify a valid pathname of the RMAN connection file with the **connect** keyword. "Configuring the NWORA resource file with the nsroraadmin program" on page 329 provides details.
The SID token 'home' is an empty string.	The **nsroraadmin** command did not include the required pathname of the Oracle home directory with the **home** keyword. In the **nsroraadmin** command, specify a valid pathname of the Oracle home directory with the **home** keyword.

Table 48 Error messages from the nsroraadmin program (page 3 of 3)

Error message	Description
The SID token '*ORACLE_SID*' is invalid.	The **nsroraadmin** command with the **sid** keyword included an invalid SID value for the Oracle database. In the **nsroraadmin** command, specify a valid SID value with the **sid** keyword.
The SID token 'sid' is an empty string.	The **nsroraadmin** command did not include the required SID value of the Oracle database with the **sid** keyword. In the **nsroraadmin** command, specify a valid SID value with the **sid** keyword.
The tokens 'sid', 'home' and 'connect' must be set when adding a SID.	The **nsroraadmin** command to add an NWORA SID resource did not include the settings of the mandatory **sid**, **home**, and **connect** keywords. In the **nsroraadmin** command to add an NWORA SID resource, include the settings of the **sid**, **home**, and **connect** keywords.
The value of the NWORA resource is missing.	The **nsroraadmin** command with the **-r update** option did not include the NWORA resource value with the resource name. In the **nsroraadmin** command with the **-r update** option, specify the NWORA resource value with the resource name.
Unrecognized argument '*option*'.	The **nsroraadmin** command included the unrecognized option *option*. Use the **nsroraadmin** command with the correct options. "Configuring the NWORA resource file with the nsroraadmin program" on page 329 provides details.
You must be the super-user to update the NWORA resource file.	The wrong user typed the **nsroraadmin** command. Type the **nsroraadmin** command as the root user on UNIX or as a member of the Microsoft Windows Administrators group.

Error messages from the nsrorainfo program

The following table lists error messages generated by the **nsrorainfo** program, in alphabetical order.

The error messages appear in the following format:

```
The NW volume information lookup failed:
error_message
```

where *error_message* is the text of the error message as shown in the table.

Table 49 Error messages from the nsrorainfo program (page 1 of 2)

Error message	Description
A connection to NW server '*server*' could not be established because '*reason*'.	NMDA could not connect to the NetWorker client file index due to the given reason. The client might not be configured as a client on the server. Take the corrective action suggested by the error message.
Could not locate the LNM save file '*backup_piece_name*' on server '*server*'.	NMDA could not locate an index record for the backup piece. The index record is probably missing. Use the **mminfo** and **nsrinfo** commands to verify the status of the index record.
Could not locate the LNM save time '*save_time*' on server '*server*'.	NMDA could not locate an index record for the save time in the client file index. The index record is probably missing. Use the **mminfo** and **nsrinfo** commands to verify the status of the index record.
Error in mmdb lookup by time: *reason*	A lookup in the media database failed for the given reason. Use the **mminfo** command to verify the status of the media database record. Take the corrective action suggested by the error message.
Exceeded the number of retries. The NetWorker server may be down or unreachable.	NMDA could not contact the NetWorker index service **nsrindexd**. This issue was probably due to the NetWorker services being shutdown. Restart the NetWorker services on the server, as required.

Table 49 Error messages from the nsrorainfo program (page 2 of 2)

Error message	Description
The file '*file_name*' could not be opened.	The file specified with the **-f** option of the **nsrorainfo** command could not be accessed. Ensure that the specified file exists, and then type the **nsrorainfo** command again with the **-f** option.
The file name provided is NULL.	In the **nsrorainfo** command, the **-f** option did not include the required file name. In the **nsrorainfo** command, include the required file name with the **-f** option.
The index entry failed the cross check: cfx_name(*backup_piece_name*) save_time(*save_time*)	During an index lookup, the entry was located in the client file index but not in the media database. Restart the NetWorker services, and use the **mminfo** and **nsrinfo** commands to verify the backup information in the indexes. Run the **nsrck** program to resolve any corruption of the indexes.
The lookup of '*backup_piece_name*' on server '*server*' failed - '*reason*'	NMDA could not locate *backup_piece_name* in the indexes due to the *reason*. The indexes might be corrupted. Run the **nsrck** program to resolve any corruption of the indexes.
The NW authentication for client '*client*' was refused by server '*server*' because '*reason*'.	NMDA could not obtain the required authentication to connect to the NetWorker client file index due to the given reason. The client might not configured as a client on the server. Take the corrective action suggested by the error message.
The record obtained has the wrong save time '*save_time1*'. The save time queried was '*save_time2*'.	NMDA located an index record in the client file index, but the record had an unexpected save time. The indexes might be corrupted. Restart the NetWorker services, and run the **nsrck** program to resolve any corruption of the indexes.

Sybase troubleshooting tips and error messages

The following topics provide NMDA Sybase troubleshooting tips and error messages.

NMDA Sybase error messages

This topic describes error messages that might appear when you use NMDA for Sybase operations, and the possible resolutions for the errors.

"NetWorker XBSA error messages" on page 409 describes the error messages generated by the NetWorker XBSA interface.

The following topics describe the NMDA Sybase error messages:

- "Error messages from Sybase consistency checks, backups, and restores" on page 440
- "Sybase backup server and libnsrsyb error messages" on page 444

Error messages from Sybase consistency checks, backups, and restores

The following table lists error messages generated during Sybase operations with the NMDA software in alphabetical order. The table first lists the messages that are common to all the Sybase operations, and then lists the messages from consistency checks, backups, and restores.

Table 50 Error messages from Sybase consistency checks, backups, and restores (page 1 of 4)

Error message	Description
Error messages common to all the Sybase operations:	
CS-LIBRARY or CT-LIBRARY error: error_message. Operating system error number(*n*): *error_message*.	An error occurred in the Sybase Open Client library layer. The operating system part of the error message appears only if an operating system error occurred. These error messages normally appear during the master database recovery because this operation shuts down the Sybase server. The messages are not normal during other operations. The error message text describes the specific problem.
error from server *Sybase_server*: Msg number, Level number, State number	The Sybase server returned an error. Check the error message that follows this message to determine the reason for the error.
No database names were specified.	The **nsrdasv**, **nsrsybrc**, and **nsrsybcc** commands each operate on a database (or for **nsrsybrc** and **nsrsybcc**, a list of databases). You did not specify any database names at the command line.
no NetWorker server was specified	The NetWorker server was not specified or could not be found. You can use the **-s** *NetWorker_server* option to specify the NetWorker server to receive the command.
non fatal internal error from server *server_name*: Msg number, Level number, State number	The Sybase server returned a nonfatal error. This error does not stop the operation. Check the message to ensure that the error does not lead to future problems.
path needs to begin with SYBASE:. The command line has the form SYBASE:/instance_name[/database_name]	The database name option for the **nsrsybcc** program did not begin with the characters SYBASE:. All Sybase server save sets must begin with this name.
the command line may specify the entire instance or a list of individual databases, but not both	You can specify either the entire server (SYBASE:/*ASE_server_name*) or a list of databases (SYBASE:/*ASE_server_name*/*database_name1* SYBASE:/*ASE_server_name*/*database_name2*) at the command line. You cannot specify a server name and a list of databases simultaneously.
The command line specifies more than one Sybase instance. Only a single instance may be supplied with each command line.	Each invocation of the **nsrdasv**, **nsrsybcc**, or **nsrsybrc** program can operate on a single Sybase server because the user ID and password supplied are unlikely to be the same over multiple servers. Retry the command and run the command once for each Sybase server.
the database name database_name has a length greater than the maximum of 30	The database name supplied at the command line was longer than 30 characters. The maximum database name length is 30 characters.
The instance name was not provided in the command line command_line_value. The command line has the form SYBASE:/instance_name[/database_name].	You specifed the database as SYBASE: but you did not specify the server name.
unable to write environment variables to the temporary file	The system could not write to the temporary file used to pass environment variables between **nsrdasv**, **nsrsybrc**, and **libnsrsyb**. Check for file access or disk problems.
user name is required and was not supplied	You must supply a username for Sybase login. NMDA can obtain this username from the Client resource in the NetWorker server, from the command line, or from the environment variable $USER.
Error message from Sybase consistency checks:	
invalid check option -o value was supplied	The database consistency check option was not valid. The *EMC NetWorker Module for Databases and Applications Command Reference Guide* or the **nsrsybcc** man page provides a list of supported options.

Table 50 Error messages from Sybase consistency checks, backups, and restores (page 2 of 4)

Error message	Description
Error messages from Sybase backups:	
a full database backup is required and will be done before the transaction log backup	The incremental backup failed because a full backup must first be performed. Perform a full backup and then retry the transaction log backup.
An invalid backup level was supplied. Valid backup levels are full, incremental, and skip	NMDA does not support the backup level supplied to the **nsrdasv** command.
cannot find database *database_name* in instance *server_name*	The database to be backed up does not exist in the Sybase server.
PRECMD or POSTCMD did not return a result. It needs to return zero on success and nonzero on failure.	The PRECMD or POSTCMD did not return a status value.
process process_number running command PRECMD or POSTCMD completed with a result of n	The PRECMD or POSTCMD exited with a nonzero result code. Check the PRECMD or POSTCMD exit code for details. Also verify that the settings of PRECMD or POSTCMD are valid. "PRECMD" on page 366 and "POSTCMD" on page 365 provide details.
the exit status of process *process_number* could not be determined	The PRECMD or POSTCMD did not exit, but the process no longer exists.
The LNM level parameter value must be between 'FULL' and '9'	The environment variable NSR_BACKUP_LEVEL specified a level other than full, incremental, or skip.
the NSR_BACKUP_PATHS parameter was not set in the configuration file	The configuration files does not include the required setting of the NSR_BACKUP_PATHS parameter.
The Sybase user name was not set. Please specify the Sybase user in the configuration by setting the SYBASE_USER parameter	Supply a username for Sybase login through the SYBASE_USER parameter setting in the configuration file.
unable to create directory entries	NMDA could not create the directory entries. Check the xbsa.messages file for the specific reason for the entry creation failure.
unable to determine whether database and log are on separate segments	The database for backup is not in a state in which NMDA could query the database to determine whether incremental backups are possible. The error message from the Sybase server that appeared before this message indicates the reason that NMDA could not query the database.
unable to dump database *database_name* in instance *server_name*	The **dump database** command failed. The error message from the Sybase server that appeared before this message indicated the reason that NMDA could not dump the database.
unable to dump the transaction log for database *database_name* in instance *server_name*	The command to dump the transaction log failed. The error message from the Sybase server that appeared before this message indicated the reason that NMDA could not dump the transaction log.
unable to dump the transaction log without truncating it for database *database_name*	The command to dump the transaction log with the no_truncate option failed. The error message from the Sybase server that appeared before this message indicated the reason that NMDA could not truncate the transaction log.
unable to execute the command PRECMD or POSTCMD contents	NMDA could not find the PRECMD or POSTCMD command. Ensure that the command exists in one of the directories specified in $PATH.
unable to print savegrp completion message	After the backup occurred, the NetWorker software could not find the save sets in the media database.
Unable to print summary. One or more parameters are not set	The parameters that NetWorker software expected to find for the function that prints the **savegrp** summary were not supplied.

Table 50 Error messages from Sybase consistency checks, backups, and restores (page 3 of 4)

Error message	Description
unable to spawn process to issue the PRECMD or POSTCMD command	NMDA could not run the PRECMD or POSTCMD command because a process needed to run the command was not available.
unable to truncate the transaction log for database *database_name*	The command to truncate the transaction log failed. The error message from the Sybase server that appeared before this message indicated the reason that NMDA could not truncate the transaction log.
unable to truncate the transaction log for database *database_name* with the no_log option	The command to truncate the transaction log failed. The error message from the Sybase server that appeared before this message indicated the reason that NMDA could not truncate the transaction log.
Error messages from Sybase restores:	
cannot restore database "*database_name*" because it does not exist in "*instance_name*"	The **nsrsybrc** command could not find a backup of the database for recovery. Run the **nsrinfo** command to see if a backup exists, and ensure that the user ID used for the **nsrsybrc** command matches the object owner that is displayed. To prevent this problem, run the Sybase Backup server and the **nsrsybrc** command and **nsrdasv** command from the same user ID.
cannot restore to the destination database *database_name* because it does not exist in the instance *server_name*	The database to which the **nsrsybrc** command is recovering data does not exist. Create the database and retry the **nsrsybrc** command.
If master is being restored, no others can be restored in the same session. The database must be in master recover mode to recover master, and this precludes restoring any other database.	A list of databases to recover was specified, including the master database. Recovery of the master database shuts down the Sybase server, which prevents recovery of the other databases.
if the destination is an instance, the source must be an instance, too	You used the **-d destination** option to specify a server, but the item to be recovered is a single database. Retry the command and specify the destination database. For example: **nsrsybrc -U sa -P** *xxx* **-d** **SYBASE:/***destination_server***/***destination_database* **SYBASE:/***source_server***/***source_database* **Note:** The Sybase server name and database name are case-sensitive and must be in the same case as recorded in the NetWorker backup indexes.
if the source is an instance, the destination must be an instance, too	The object to be recovered is an entire Sybase server, but the destination specified for the server recovery is a database name. Retry the command and specify the destination as a server. For example: **nsrsybrc -U sa -P** *xxx* **-d SYBASE:/***destination_server* **SYBASE:/***source_server* **Note:** The Sybase server names are case-sensitive and must be in the same case as recorded in the NetWorker backup indexes.
internal error. Full backup expected but not found.	A full backup was found, but was then no longer available before the **nsrsybrc** command recovered the database. For example, this error occurs when you manually relabel the volume containing the full backup at the same time that the incremental backup depending on that full backup is recovered.
invalid time specification: time value	The **-t time** option supplied with the **nsrsybrc** command was not valid. This option must be in the **nsr_getdate** form. The **nsr_getdate** man page provides details.
no NetWorker server was specified	The NetWorker server was not specified or could not be found. Use the **-s** *server_name* option to specify the NetWorker server to receive the command.

Table 50 Error messages from Sybase consistency checks, backups, and restores (page 4 of 4)

Error message	Description
Recover option validation error.	You must specify either the entire server (SYBASE:/*ASE_server_name*) or a list of databases (SYBASE:/*ASE_server_name*/*database_name1* SYBASE:/*ASE_server_name*/*database_name2*) at the command line. You cannot specify both a server name and a list of databases simultaneously.
the command line did not specify a database or an instance to restore	You must specify the name of the database or Sybase server to be recovered with the **nsrsybrc** command.
there are no databases to restore in instance *server_name*	NMDA could not find any databases in the directory entry for the Sybase server database.
there is no backup of the instance for the time supplied	NMDA could not find a backup for the Sybase server name. Ensure that you run the **nsrsybrc** command with the same user ID as used to run the **nsrdasv** command. Otherwise, ensure that the time used is correct. If you do not enter a time, NMDA uses the current time.
there is no full backup of database *database_name* in instance *server_name* for the time supplied	Backups of this database exist, but there was not a full backup available for the time requested. Try an earlier time or run the **nsrinfo** command to determine when the last full backup occurred. For example, if the full backup has passed the browse policy, the full backup might be listed in the media database but not in the client index. In this case, re-create the entry in the client index with the **scanner -i** command and then recover the database with the **nsrsybrc** command.
unable to query backup	There was an error querying the backup from the server. Check the xbsa.messages file for the specific error text.

Sybase backup server and libnsrsyb error messages

When the Sybase Backup server encounters an error or condition requiring a warning, the server writes a message to the Sybase Backup server error log.

The default error log location is $SYBASE/$SYBASE_ASE/install.

The following table lists **libnsrsyb** error messages logged in the Sybase Backup server error log. The Sybase documentation describes other Sybase Backup server errors.

Table 51 Sybase backup server and libnsrsyb error messages (page 1 of 2)

Error message	Description
libnsrsyb opened with an unknown mode: internal error	The **libnsrsyb** shared library was opened with a mode other than read or write.
there is insufficient memory to continue	There is not enough memory to complete the operation.
The time stamp ddddddddd has non digits in it. Timestamps are composed of digits in the form YYYYMMDDhhmmsslll.	The timestamp supplied for the **load** command from the **isql** command line has a timestamp with an incorrect format. The timestamp must have the format *YYYYMMDDhhmmsslll*, where: • *YYYY* indicates the year. • *MM* indicates the month. • *DD* indicates the day. • *hh* indicates the hour. • *mm* indicates the minutes. • *ss* indicates the seconds. • *lll* indicates the milliseconds. The millisecond position is optional. Alternatively, you can enter **000** for the milliseconds.

Table 51 Sybase backup server and libnsrsyb error messages (page 2 of 2)

Error message	Description
time stamps are not valid for dump command	The **isql** command line specified a timestamp for a **dump** command. Timestamps are not valid with the **dump** command.
unable to close and create save set	The BSA call to create and close the save set for a database or transaction dump failed. Check the xbsa.messages file for specific details.
unable to close save set	The call to close the save set failed during a load of a database or a transaction log. Check the xbsa.messages file for specific details.
unable to create environment variables	The resources required to create the internal environment variable array were not available. This might be due to access problems in the /nsr/tmp directory.
Unable to create save set. There is likely a configuration or enabler problem. Set the debug level to at least 2, retry the operation, and check the /nsr/applogs/xbsa.messages file for the underlying reason.	The save set creation on the NetWorker server failed. If the debug level is at least 2 (the default), check the xbsa.messages file for the error text. If the debug level is not 2, change the setting to 2 and retry the operation. Check the xbsa.messages file for specific details.
unable to create the save set on the server	The call to create the save set on the NetWorker server failed. Check the xbsa.messages file for specific details.
unable to end the current read session	During a load database or load transaction log operation, the read session of the data from the NetWorker software could not be closed. Check the xbsa.messages file for specific details.
Unable to find backup of the (database or transaction log) SYBASE:/server_name/database_name. Check the command line for errors in the instance or database name or use nsrinfo to see which save sets are available.	The item to be loaded could not be found. Use the **nsrinfo** command to check that the object-owner for the backup is the same as the process that launched the Sybase Backup server and that backups exist for this database.
Unable to find full backup of the database database_name for the time supplied. Unable to find incremental backup of the database database_name for the time supplied. Unable to find backup of the database database_name for the time supplied.	No backup could be found in the NetWorker server. If no time was supplied, the time used is the current time, which means that no backup exists. Use the **nsrinfo** command to check which backups are available and ensure that the object owner shown is the same as the user ID that launched the Sybase Backup server.
unable to parse stripe specifier	The **isql** command line had a poorly formatted stripe specifier.
unable to read the requested number of bytes from the save set	A load database or load transaction log operation could not read the save set. Check the xbsa.messages file for specific details.
unable to send data to save set	A database dump or transaction log dump could not write the data to the save set. Check the xbsa.messages file for specific details.
unknown backup type supplied	The backup type supplied from the NetWorker server was not a database or a transaction log.

GLOSSARY

This glossary contains the definitions of terms found in this manual. Most of the terms are specific to NMDA. For terms specific to the NetWorker software, refer to the latest *EMC NetWorker Administration Guide*.

A

active-active application cluster Type of cluster configuration where a group of linked virtual or physical hosts with shared storage, called cluster nodes, can access the database data from multiple nodes concurrently.

active-passive cluster Type of cluster configuration where the data server runs on the active physical node, and other nodes are passive nodes that maintain data updates and wait to take over if the active node fails.

administrator Person who normally installs, configures, and maintains software on the network computers, and who adds users and defines user privileges.

attribute Feature of a NetWorker resource. A setting or information that the resource provides.

auto media management Feature that enables the storage device controlled by the NetWorker server to automatically label, mount, and overwrite an unlabelled volume. Volumes that are eligible for reuse are also automatically recycled.

autochanger See “library.”

B

backup
1. Duplicate of database data or application data or an entire computer system stored separately from the original, which you can use to recover the original if it is destroyed or damaged.
2. Operation that saves data to a volume for use during a recovery.

backup cycle Full or level 0 backup and all the subsequent incremental backups that are dependent on that backup.

backup level See “level.”

backup volume See “volume.”

blobspace (Binary large object space) Informix logical unit of storage comprised of one or more data chunks, used to store large objects, such as multimedia images.

bootstrap Save set that is essential for disaster recovery procedures with the NetWorker server. The bootstrap consists of three components that reside on the NetWorker server: the media database, the resource database, and a server index.

browse policy NetWorker policy that specifies the time during which backup entries are retained in the client file index. The index makes the associated backed-up data readily accessible for recovery. See also "retention policy."

C

catalog synchronization Process that removes an NSM snapshot backup entry from the database server catalog when the corresponding backup entry is removed from the NetWorker indexes. See also "NWORA resource file."

client Database or application server whose data can be backed up and restored with the NMDA software.

Client Direct Feature that enables clients to deduplicate backup data and send the deduplicated data directly to AFTD or DD Boost storage devices, bypassing the NetWorker storage node. The storage node manages the backup devices but does not handle the backup data.

client file index Database maintained by the NetWorker server that tracks every database object, file, or file system backed up. The NetWorker server maintains a single index file for each client computer. The tracking information for each backup is deleted from the index when the browse policy of the backup expires.

client-initiated backup See "manual backup."

client-side configuration NMDAbackup configuration, created without the configuration wizard, that is performed by using the NetWorker Management Console and the configuration files or parameters stored on the client host, as compared to server-side configuration. See also "server-side configuration."

clone Duplicate copy of backed-up data, which is indexed and tracked by the NetWorker server. You can clone single save sets or entire volumes.

clone volume Exact duplicate of a backup or an archive volume. NetWorker software can track four types of volumes (backup, archive, backup clone, and archive clone). You cannot intermix the save sets of these different types on one volume. You can use clone volumes in the same way as the original backup volume or an archive volume.

cluster nodes A group of linked virtual or physical hosts with shared storage in a cluster, which work together and represent themselves as a single host called a virtual cluster host.

cold backup See "offline backup."

configuration wizard NMDA wizard integrated with the NetWorker Management Console GUI, which you can use to configure the following operations:

- NMDA scheduled backup for all applications as either a nonsnapshot backup or a snapshot backup
- NMDA recovery for Oracle only

Console server See "NetWorker Management Console (NMC)."

D

daemon Process on UNIX systems that runs in the background and performs a specified operation at predefined times or in response to certain events.

Data Domain device Logical storage device created on a Data Domain system, used to store deduplicated NetWorker backups. Each device appears as a folder on the Data Domain system and appears with a storage volume name in NMC. Also known as a DD Boost device.

dbobject Informix database object, a term that can refer to a blobspace, dbspace, or logical log file.

dbspace Informix logical unit of storage that consists of one or more chunks. An IDS instance might consist of one or more dbspaces.

DD Boost An optimized library and communication framework with a special Data Domain API that allows the NetWorker software to define and interact with storage devices on the Data Domain system.

deduplication backup Type of backup in which redundant data blocks are identified and only unique blocks of data are stored. When the deduplicated data is restored, the restore returns the data to its original native format.

deprecated feature Feature that is supported in the current release of the product but will be unsupported and removed in a future release.

destination client Computer to which a directed recovery restores the database files.

device
1. Storage unit that can contain a backup volume. A storage unit can be a tape device, optical drive, autochanger, or disk connected to the server or to a storage node.
2. General term that refers to storage hardware.
3. Access path to the physical drive, when dynamic drive sharing (DDS) is enabled.

direct file access (DFA) See “Client Direct.”

directed recovery Method that recovers data that originated on one client host and re-creates it on a different client host, known as the destination client.

disaster recovery Restore and recovery of business operations and data if a hardware failure or software corruption occurs.

distributed segment processing (DSP) Part of the DD Boost interface, which enables data deduplication on a host before the data is sent to the Data Domain system for storage.

E

emergency boot file Informix ON-Bar ASCII file that contains all the information in the ON-Bar catalog tables that pertain to critical dbspaces.

event-based backup See “probe-based backup.”

expiration date Date when a volume changes from read/write to read-only.

expired save set Save set that has exceeded its browse time and has been removed from the NetWorker client file index. You cannot browse expired save sets.

F

firewall A system designed to prevent unauthorized access to or from a private network.

full backup See “level.”

G

group Client computer or group of clients configured to back up files during a NetWorker scheduled backup, according to a single designated schedule or set of conditions.

H

high-availability system System of multiple computers configured as cluster nodes on a network that ensures the application services continue despite a hardware failure or a software failure.

host Computer on a network.

hot backup See “online backup.”

I

incremental backup See “level.”

internationalization (I18N) Process of adapting software to accept input and output of data in various languages and locales.

J

jukebox See “library.”

L

label Electronic header on a volume used for identification by NetWorker or other data mover application.

legacy NetWorker module One of the NetWorker modules that NMDA replaces: NMDB2, NMI, NML, NMO, or NMS.

level Backup configuration option that specifies how much data is saved during a scheduled or manual backup:

- A full backup backs up all data objects, regardless of when they last changed.
- An incremental backup backs up only data objects that have changed since the previous backup.

Table 8 on page 87 describes the backup levels supported for each type of database or application.

library Hardware device that contains one or more removable media drives, slots for pieces of media, media access ports, and a robotic mechanism for moving pieces of media between these components. Libraries automate media loading and mounting functions during backup and recovery. The term library is synonymous with autochanger, autoloader, carousel, datawheel, jukebox, and near-line storage.

localization (L10N) Process of translating software for a user language or locale.

M

manual backup Backup that a user performs from the client, also known as an unscheduled, on-demand, or ad hoc backup.

mean time to recover (MTTR) Time specified to perform a recovery. For example, you might set 10 minutes as the goal for a recovery from a disk failure.

media Physical storage, such as magnetic tape, optical disk, or file system, to which a backup writes data. See also “volume.”

media index Database that contains indexed entries of storage volume location and the lifecycle status of all data and volumes managed by the NetWorker server. An alternate term is media database.

media pool See “pool.”

multiple session backup and restore Method of backing up or restoring multiple parallel streams of data simultaneously between a database and multiple media devices. An alternate term is multistripe.

multiplex To simultaneously write data from more than one save set to the same storage device.

N

NetWorker Management Console (NMC) Software program used to manage NetWorker servers and clients. The NMC server also provides reporting and monitoring capabilities for all NetWorker processes.

NetWorker Module for Databases and Applications (NMDA) Add-on module for NetWorker software that provides backup, restore, and storage management solutions for supported database software or application software. NMDA is an aggregate of individual NetWorker modules for DB2 (NMDB2), Informix (NMI), Lotus Domino/Notes (NML), Oracle (NMO), and Sybase (NMS) data.

NetWorker server Computer on a network that runs the NetWorker server software, contains the online indexes, and provides backup and restore services to the clients and the storage nodes on the same network.

NetWorker Snapshot Management (NSM) EMC NetWorker software integrated into the NetWorker client with NetWorker release 8.1, which provides point-in-time snapshots of data to be backed up. Applications that are running on the host system continue to write data during the snapshot operation. The snapshots include data from open files.

notification Message sent to the NetWorker administrator about important NetWorker events.

NWORA resource file NMDA Oracle resource file in which you must define resources to enable Oracle proxy backups and (optionally) catalog synchronization.

O

offline backup Backup of database objects performed while the corresponding database or instance is shutdown and unavailable to users.

online backup Backup of database objects performed while the corresponding database or instance is running and available to users.

online indexes Databases located on the NetWorker server that contain all the information pertaining to the client backups (client file index) and backup volumes (media index).

Oracle Recovery Catalog Collection of Oracle database tables maintained by RMAN, which includes information about the following items:

- Oracle backup sets and backup pieces
- Image copies and proxy copies
- Archived redo logs
- Stored scripts
- Target database schema

Oracle Recovery Manager (RMAN) Oracle utility that acts as an intelligent interface to Oracle databases for the backup and restore of Oracle database objects.

P

parallelism Method that simultaneously backs up or recovers data for multiple clients or multiple save sets for one client.

pathname Set of instructions to the operating system for accessing a file:

- An absolute pathname indicates how to find a file starting from the root directory and working down the directory tree.
- A relative pathname indicates how to find a file starting from the current location.

performing client Host where a directed recovery is initiated by using the NetWorker User for Lotus program.

physical cluster client NetWorker client bound to a physical host in the cluster. The client can have its own resources (private or local).

point-in-time copy (PIT copy) Fully usable copy of a defined collection of data, such as a consistent file system, database, or volume, which contains an image of the data as it appeared at a single point-in-time. An alternate term for a PIT copy is a shadow copy or a snapshot.

policy Set of defined rules for client backups that you can name and apply to multiple groups. Groups have dataset policies, schedule policies, browse policies, and retention policies.

policy uniformity Consistency of the browse policies and retention policies in a group of codependent Oracle save sets from the same scheduled backup cycle or save set bundle. NMDA enforces policy uniformity to ensure that incremental Oracle backups do not persist after other Oracle backups that they depend on have expired.

pool

1. NetWorker sorting feature that assigns specific backup data for storage on selected media volumes.
2. Collection of NetWorker backup volumes to which specific data was backed up.

primary storage Server storage subsystem that contains application data and any persistent snapshot backups of data. See also "secondary storage."

probe-based backup Type of scheduled backup where the NetWorker server initiates the backup only when certain conditions are met, as determined by one or more probes. An alternate term for probe-based backup is event-based backup.

proxy backup Backup of Oracle data that creates a point-in-time (snapshot) copy on primary storage through the NSM feature. The backup optionally backs up the snapshot to secondary storage. See also "snapshot backup," "snapshot and rollover backup," and "rollover-only backup."

proxy client host Host used in NSM snapshot backups that is separate from the database server host, with access to the primary storage unit. During an NSM backup to secondary storage, either the database server host or proxy client host backs up a point-in-time copy (snapshot) from the primary storage to secondary storage.

proxy restore Restore of Oracle data from a proxy backup through NSM. See also "snapshot restore," "rollback restore," and "restore from secondary storage."

Q

quiesce Database state or application state that enables the performance of a consistent backup.

R

recover To restore data files from backup media to a client disk and apply transaction logs, redo logs, or logical logs to make the data consistent with a given point-in-time.

recyclable save set Save set whose browse policy and retention policy have expired in the NetWorker indexes. Recyclable save sets are removed from the media database.

recyclable volume NetWorker volume whose data has exceeded both the browse policy and retention policy and is now available for relabeling and reuse.

regular backup or restore NMDA backup or restore that does not use snapshot technologies through NSM.

remote device Storage device attached to a storage node that is separate from the NetWorker server.

resource NetWorker software component that describes details of the NetWorker server or its clients. Clients, devices, schedules, groups, and policies are all NetWorker resources. Each resource has configurable attributes that define its properties.

resource database NetWorker database of information about each configured resource.

restore To retrieve individual data files from backup media and then copy the files to disk, without applying transaction logs. See also "recover."

restore from secondary storage Type of DB2 or Oracle restore that restores a proxy backup from a secondary storage medium.

retention policy NetWorker setting that determines the minimum time to retain backup data on a volume available for recovery. After this time, the data becomes eligible for overwriting. See also "browse policy."

RMAN See "Oracle Recovery Manager (RMAN)."

roll forward To apply transaction logs to a recovered database to restore the database to a state that is consistent with a given point-in-time.

rollforward recovery Type of DB2 database recovery that applies transaction logs to restore the database to a given point-in-time.

rollback restore NSM process by which a specific point-in-time copy (snapshot) of data is restored to the source location by using the hardware's particular capabilities. A rollback restore is a destructive save set restore.

rollover-only backup Type of snapshot backup in which an existing point-in-time copy (snapshot) of data, which a snapshot backup created on the primary storage, is immediately backed up to secondary storage. The snapshot is automatically deleted from the primary storage.

S

save NetWorker command that backs up client files to backup media volumes and makes data entries in the online index.

save set Collection of data from a single client computer backed up on storage media.

save set bundle Group of codependent Oracle save sets from the same scheduled backup cycle, assembled into a bundle according to configuration settings.

save set consolidation Process that performs a level 1 backup and merges the backup with the last full backup of a save set to create a full backup.

save set ID (ssid) Internal identification number assigned to a save set.

save stream Data and save set information written to a storage volume during a backup. A save stream originates from a single save set.

scanner NetWorker command used to read a backup volume when the online indexes are not available.

scheduled backup Type of backup configured to start automatically at a specified time for a group of one or more NetWorker clients. A scheduled backup generates a bootstrap save set.

secondary storage Storage library attached to the NetWorker server or to a storage node, used to store regular backups or snapshot backups. See also "primary storage."

server index See "client file index."

server-side configuration NMDA backup configuration created by using the configuration wizard, with settings saved on the NetWorker server, as compared to a client-side configuration. See also "client-side configuration."

shared disk Storage disk connected to multiple nodes in the cluster.

snap set Group of files, volumes, or file systems from a single client that describes the collection of data for which a point-in-time copy is created on an external disk subsystem, such as a storage array.

snapshot Point-in-time, read-only copy of data created on a primary storage system during a snapshot backup.

snapshot and rollover backup Type of snapshot backup in which an existing point-in-time copy (snapshot) of data, which a snapshot backup created on the primary storage unit, is backed up to secondary storage. The snapshot is retained on the primary storage.

snapshot backup Type of backup that creates a point-in-time copy (snapshot) of data on the primary storage unit, which can be immediately recovered as a backup copy.

snapshot policy Set of rules that control the lifecycle of a snap set. The snapshot policy specifies the frequency of snapshots and how long to retain snapshots before recycling.

snapshot restore Type of restore that copies data created during a snapshot backup to its original location or to an alternate location.

ssid See “save set ID (ssid).”

stage To move data from one storage medium to a less costly medium, and later removing the data from its original location.

storage device See “device.”

storage node Storage device physically attached to a computer other than the NetWorker server. The controlling NetWorker server administers the backup operations of the storage node.

T

tablespace Oracle database structure that consists of one or more data files.

target database Database that the NetWorker server backs up as a safeguard against data loss.

transaction log Record of named database transactions or a list of changed files in a database, stored in a log file to enable quick restore and rollback transactions.

U

unscheduled backup See “manual backup.”

V

virtual cluster client NetWorker client not permanently bound to one physical host but managed by a cluster manager. An alternate term for a virtual cluster client is a logical cluster client or a virtual client.

volume

1. Unit of physical storage medium, such as a magnetic tape, optical disk, or file system used to store data. You must store backup data on a backup volume, not on an archive volume or a clone volume.
2. Identifiable unit of data storage that may reside on one or more computer disks.

volume ID (volid) Internal identification that NetWorker software assigns to a backup volume.

volume name Name that you assign to a backup volume when you label the volume. See also “label.”

INDEX

Symbols

(Avamar) Deduplication Backup attribute
    in Client resource 90
(Avamar) Deduplication Node attribute
    in Client resource 90

Numerics

32-bit and 64-bit coexistence 278

A

active-active application cluster 254
active-passive cluster 246
ADDRESS_LIST parameter 268
Administrator, NetWorker
    Client resource 301
    Device resource 78
    Group resource 87, 301
    Pool resource 78, 300
    Schedule resource 87
    Server resource 75
    Snapshot Policy resource 296, 301
    User Group resource 75
Advanced Copy Services (ACS) 305
AES encryption 358
Aliases attribute
    in Client resource 89
allocate channel command
    parms option 402
archive logging 36
archived redo log
    backup 237, 272
    restore 272
    sharing across RAC nodes 271
attributes for resources
    Client resource 78
    Server resource 75
automatic catalog synchronization for proxy backups 332
automatic channel allocation 48, 128, 130
Avamar deduplication
    backup 29, 95
    backup in a cluster 252
    restore 29

B

backup
    Client resource 301
    command 310, 402
    command (pool option) 398
    command (trace option) 426
    copies 133
    delete 164
    error messages 408, 426
    from NetWorker User for Lotus 151
    from NetWorker User for Sybase 156
    Group resource 87, 301
    levels supported 40
    multiple session 112
    partitioned Domino server 120
    required Sybase roles 138
    Schedule resource 87
    transaction logs 116, 263
Backup Command attribute
    in Client resource 89, 302
backup configuration road map 70
backup copies during manual backups 133
backup current control file command 240
backup failover 247
backup hangs, troubleshooting 408
backup spfile command 240
backup types
    archived redo log 237, 272
    Avamar deduplication 29, 95
    control file 236, 239
    Data Domain deduplication 30
    DD Boost 101
    deferred live 315
    instant 315, 323
    manual 28, 148
    NetWorker bootstrap 66
    NWORA resource file 316, 318
    password file 237
    probe-based 28
    proxy 294
    registry files 237
    scheduled 28, 146
    snapshot 296
bootstrap
    backup 28
    NetWorker 66
    records 226
Browse Policy attribute
    in Client resource 89, 301
browse policy uniformity 55
browse time 187, 221
browselist save set 39
bundling, save set 52

C

canceling
    proxy backup 303
    scheduled backup 131, 147
catalog synchronization for proxy backups 332
change...crosscheck command 166
change...delete command 332
channel option, send command 399
Checkpoint enabled attribute
    in Client resource 89

Checkpoint granularity attribute
in Client resource 89
circular logging 35
Client Direct attribute
in Client resource 90
Client Direct method 101, 362
client file index 28, 317, 319
Client resource 301
Client resource attributes
(Avamar) Deduplication Backup 90
(Avamar) Deduplication Node 90
Aliases 89
Backup Command 89
Browse Policy 89
Checkpoint enabled 89
Checkpoint granularity 89
Client Direct 90
Data Domain Backup 90
Data Domain Interface 90
Group 90
Name 90
Parallelism 90
Remote Access 90
Retention Policy 91
Save Set 91, 92, 93
Schedule 93
Scheduled Backup 93
Storage Nodes 93
Client Retries attribute
in Group resource 336
cluster
active-active application 254
active-passive 246
Avamar deduplication backup 252
Data Domain deduplication backup 253
manual backups 252
probe-based backup 253
restoring backups 254
systems 33, 335
coexistence, 32-bit and 64-bit 278
collaboration
Oracle DBA and NetWorker Backup Administrator 48
command
allocate channel, parms option 402
backup 310, 402
backup (pool option) 398
backup (trace option) 426
backup current control file 240
backup spfile 240
change...crosscheck 166
change...delete 332
configure channel, parms option 312, 388
crosscheck 166
delete expired backup 166
nsrorainfo 209
nsrsybcc 154
nsrsybrc 138
restore 402
rman 153, 215
rman send 388, 401, 402
rman.exe 153, 215
savefs 65
savegrp 65
send 312, 388, 398, 399
send (channel option) 399
send (device_type option) 399
send (NSR_ENV keyword) 399
send (precedence rules) 402
Command attribute
in Probe resource 107, 287
Command Options attribute
in Probe resource 108, 109, 288
configuration
Client resource 301
Device resource 78
Group resource 87, 301
manual backup 148
Pool resource 78, 300
Schedule resource 87
scheduled backup 146
Server resource 75
Snapshot Policy resource 296, 301
User Group resource 75
configuration file
transaction logs 116
configuration template
nmda_db2.cfg 354
nmda_informix.cfg 354
nmda_lotus.cfg 354
nmda_mysql_backup.cfg 354
nmda_mysql_restore.cfg 354
nmda_oracle.cfg 354
nmda_sybase.cfg 354
configure channel command
parms option 312, 388
configuring NMDA backups 70
connection file, for catalog synchronization 325, 327, 328
control file
backup 236, 239
control files, mirrored 237
convert backup settings from earlier releases 249
crosscheck command 166

D

Data Domain Backup attribute
in Client resource 90
Data Domain deduplication
backup 30
backup in a cluster 253
restore 30
Data Domain Interface attribute
in Client resource 90
database
manual backup 148
scheduled backup 146
Database Partitioning Feature (DPF) 256, 260
Datazone pass phrase attribute
in Server resource 75
DB2
DPF systems 256

history file 36, 165
history pruning 307
prune command 164
resource file 307
restore and recovery 171
restore to same or different instance 173
DB2 transaction logs 35
DB2_ALIAS parameter 366
DB2_APPLY_NW_LEVELS parameter 367
DB2_BUFFER_SIZE parameter 367
DB2_NODE_NAME 259
DB2_NODE_NAME parameter 368
DB2_NUM_BUFFERS parameter 368
DB2_OPTIONS 259
DB2_OPTIONS parameter 368
DB2_PARALLELISM parameter 368
DB2_PARTITION_LIST 259
DB2_PARTITION_LIST parameter 368
DB2_QUIESCE parameter 368
DB2_SESSIONS parameter 369
DB2_TBS_LIST parameter 369
DB2_USER 259
DB2_USER parameter 369
DB2_VENDOR_LIB_PATH parameter 369
DB2INSTANCE parameter 367
DB2PATH parameter 307, 368
DBA_DISK_BACKUP parameter 108, 288
DBCCOPT parameter 392
DBID 237, 238
DD Boost backup 101
deferred live backup 296, 315
delete
backup 164
snapshots 307
delete expired backup command 166
destructive restores 297
Device resource 78
device_type option, send command 399
direct file access (DFA) 29
disaster recovery 226
preparation 236
distributed segment processing (DSP) 30
DO_LOGFILE_BACKUPS parameter 370
DO_WHOLE_SYSTEM_BACKUP parameter 370
DPF partitions 256, 260

E

EMC Online Support 17
enabling
policy uniformity 136
save set bundling 135
encryption
AES encryption 358
environment variable
NLS_LANG 74
error messages 408, 413
RMAN 426

F

failover
backup 247
connect-time 268
proxy backup 335
FAILOVER parameter 268
fetched logs rollforward recovery 176

G

generic Client resource 250
Group attribute
in Client resource 90, 301
Group resource 87, 301

H

high availability 246
history pruning 36, 307

I

I18N (internationalization) 71
immediate live backup 296
incremental backups
threshold procedure 142
Informix
restore and recovery 178
Informix Multi-node Active Clusters for High Availability (MACH) 264
Informix restore through onbar
combined restore 180
logical restore 180
physical restore 180
point-in-time restore 181
INFORMIXDIR parameter 370
INFORMIXSQLHOSTS parameter 370
initialization parameter file
initoracle_sid.ora 269
PFILE 236, 238, 240
SPFILE 236, 239
INSTANCE_NAME parameter 269
instant backup 315, 323
instant restore 297
INSTHOME 259
INSTHOME parameter 369
internationalization (I18N) 71
isql commands
threshold procedure 142

L

LD_LIBRARY_PATH parameter 393
LD_LIBRARY_PATH_64 parameter 393
LIBPATH parameter 393
live backup 296
deferred 315
LOCAL_LISTENER parameter 269
Lockbox resource, in cluster 250
LOG_THRESHOLD parameter 108
logs, transaction
backup levels 40

Lotus
restore types 42
Lotus browselist backups 39
Lotus restartable backups 41, 124, 161
LOTUS_NSF_FILE parameter 108
LOTUS_USER parameter 371

M

MACH systems 264
manual backup 148
cluster 252
configuration 94
from NetWorker User for Lotus 151
from NetWorker User for Sybase 156
levels supported 40
monitoring 158
road map 148
manual catalog synchronization for proxy backups 332
master database, recovering 241
media database 28, 318, 319
media management
storage devices 78
volume pools 78
mirrored
control files 237
online redo logs 237
MML catalog 322
monitoring
manual backup 158
scheduled backup 147
Multi-node Active Clusters for High Availability 264
multiple
applications on same host 34
devices and sessions 172
session backups 112, 172
MySQL
restore and recovery 196
MYSQL_BACKUP_DIR parameter 380
MYSQL_BACKUP_NAME parameter 381
MYSQL_BINLOG parameter 381
MYSQL_CFG_FILE parameter 381
MYSQL_COMPRESS parameter 381
MYSQL_COMPRESS_LEVEL parameter 382
MYSQL_DATABASES parameter 382
MYSQL_DATADIR parameter 382
MYSQL_EXTRACT_PATHS parameter 382
MYSQL_INCLUDE parameter 383
MYSQL_INCR_DIR parameter 383
MYSQL_INCR_OPTIONS parameter 383
MYSQL_INNODB_LOG_FILE_SIZE parameter 383
MYSQL_INNODB_LOG_FILES_IN_GROUP parameter 383
MYSQL_LOG_OPTIONS parameter 384
MYSQL_MEB_PATH parameter 384
MYSQL_ONLY_INNOB_OPTIONS parameter 385
MYSQL_RESTORE_OPERATION parameter 385
MYSQL_SBT_LIB_PATH parameter 385
MYSQL_SRC_ENTRY parameter 386
MYSQL_USER parameter 386

N

Name attribute
in Client resource 90
in Probe resource 107, 287
in Server resource 75
network files
listener.ora 237, 238, 268
sqlnet.ora 237, 238
tnsnames.ora 237, 238, 268
NetWorker
bootstrap backup 66
client file index 28, 317, 319
media database 28, 318, 319
staging 52
NetWorker configuration
Client resource 78, 301
Device resource 78
Group resource 87, 301
Pool resource 78, 300
Schedule resource 87
Server resource 75
Snapshot Policy resource 296, 301
NetWorker PowerSnap Modules 31
NetWorker Snapshot Management 294
NetWorker User for Lotus
backup 151
restore 185
NetWorker User for Sybase
backup 156
restore 220
NLS_LANG environment variable 74
nmda_app.messages.raw file 409
nmda_db2_tlogs.cfg configuration file 116
nmda_db2.cfg template file 354
nmda_informix.cfg template file 354
nmda_lotus.cfg template file 354
nmda_mysql_backup.cfg template file 354
nmda_mysql_restore.cfg template file 354
nmda_oracle.cfg template file 354
nmda_oracle.messages.raw file 426
nmda_sybase.cfg template file 354
Notes_ExecDirectory parameter 372
NSM
backup configuration 299
NSM feature 294
NSM parameters 304, 310, 312
NSR_AES_ENCRYPTION parameter 358
NSR_APPLY_LOGS parameter 372
NSR_ASE_PASSWORD parameter 393
NSR_ASE_VERIFY parameter 393
NSR_AUTO_RESTORE parameter 372
NSR_BACKUP_ALL_EXTENSIONS parameter 372
NSR_BACKUP_LEVEL parameter 372, 386, 394
NSR_BACKUP_LOGS_MODE parameter 372
NSR_BACKUP_LOTUS_DIR parameter 372
NSR_BACKUP_PATHS parameter 373, 394
NSR_BROWSELIST_CACHE_DEST parameter 373
NSR_BUNDLING parameter 135
NSR_CATALOGFILE parameter 373
NSR_CHECKSUM parameter 358

NSR_CLIENT 259
NSR_CLIENT parameter 286, 358
NSR_COMFORT_SPAN parameter 373
NSR_COMPRESSION parameter 359
NSR_CROSS_MOUNT_POINTS parameter 374
NSR_DATA_DOMAIN_INTERFACE parameter 359
NSR_DATA_MOVER parameter 304
NSR_DATA_VOLUME_POOL parameter 286
NSR_DATA_VOLUME_POOL parameters 287
NSR_DATA_VOLUME_POOL* parameters 79, 359, 388, 389
NSR_DB2CAT_MODE parameter 307
NSR_DBIID parameter 374
NSR_DEBUG_LEVEL parameter 108, 288, 360
NSR_DEDUP_BACKUP parameter 360
NSR_DEDUP_CACHE_ENABLED parameter 360
NSR_DEDUP_CACHE_TAG parameter 361
NSR_DEDUP_CHUNK_SIZE parameter 361
NSR_DEDUP_NODE parameter 362
NSR_DEVICE_INTERFACE parameter 362
NSR_DIAGNOSTIC_DEST parameter 108, 288, 362
NSR_DIRECT_ACCESS parameter 362
NSR_DPRINTF parameter 363
NSR_DR_BACKUP_INFO parameter 369, 371, 386
NSR_DR_FILE_LIST parameter 369, 371, 386
NSR_DUMP_DATA_OPT parameter 394
NSR_DUMP_LOG_OPT parameter 394
NSR_ENCRYPTION_PHRASES parameter 363
NSR_ENV keyword in send command 399
NSR_ENV_LIST parameter 389
NSR_EXCLUDE_FILE parameter 374, 394
NSR_EXCLUDE_LIST parameter 374
NSR_FOLLOW_LINKS parameter 374
NSR_GROUP parameter 131, 363
NSR_INCR_BACKUP_LOGS_MODE parameter 374
NSR_INCR_EXPIRATION parameter 136
NSR_LOCALE parameter 395
NSR_LOG_DIR parameter 375
NSR_LOG_VOLUME_POOL parameter 369, 371, 386, 395
NSR_LOTUS_DATA_DIR parameter 375
NSR_MAX_START_RETRIES parameter 287, 363
NSR_MAX_STREAMS parameter 304
NSR_MAX_TXN_LOGS parameter 375
NSR_MMDB_RETRY_TIME parameter 389
NSR_NO_BUSY_ERRORS parameter 364
NSR_NO_MULTIPLEX parameter 287, 389
NSR_NO_NOTES_INIT parameter 375
NSR_NOTES_CONNECT_TIMEOUT parameter 375
NSR_NOTES_INI_PATH parameter 375
NSR_NUMBER_LOGS parameter 376
NSR_NWPATH parameter 300, 325, 364
NSR_ORACLE_CONNECT_FILE parameter 109, 327
NSR_ORACLE_HOME parameter 327
NSR_ORACLE_LIB_PATH parameter 327
NSR_ORACLE_NLS_LANG parameter 300
NSR_ORACLE_NLS_LANG parameter resource 325
NSR_ORACLE_RETENTION parameter 390
NSR_ORACLE_SID parameter 327
NSR_ORACLE_TNS_ADMIN parameter 328
NSR_ORACLECAT_DEBUG_FILE parameter resource 325
NSR_ORACLECAT_LOG_FILE parameter resource 325
NSR_ORACLECAT_MODE parameter resource 325, 333
NSR_PARALLELISM parameter 376, 395
NSR_PHYSICAL_HOST_LICENSE parameter 364
NSR_PREFETCH_LOGS parameter 376
NSR_PROMOTE_FULL parameter 395
NSR_PROXY_PFILE parameter 390
NSR_PS_SAVE_PARALLELISM parameter 304
NSR_RECOV_INTERACT parameter 376, 387
NSR_RECOV_LIST_FILE parameter 376
NSR_RECOVER_OPTIONS parameter 376
NSR_RECOVER_POOL parameter 364
NSR_RECOVER_TIME parameter 377, 387
NSR_RELOCATION_DEST parameter 377
NSR_REMOVE_ON_FAILURE parameter resource 325
NSR_RESOURCE_DIR parameter 377
NSR_RETAIN_NUM_LOGS parameter 377
NSR_RMAN_ARGUMENTS parameter 390
NSR_RMAN_OUTPUT parameter 390
NSR_SAVESET_BROWSE parameter 89, 365
NSR_SAVESET_NAME parameter 377
NSR_SAVESET_RETENTION parameter 91, 365
NSR_SERVER 259
NSR_SERVER parameter 131, 365
NSR_SERVER_NIC parameter 287, 390
NSR_SKIPDBERRORS parameter 377
NSR_SNAPSHOT_POSTCMD parameter 304
NSR_VERBOSITY parameter 378
NSR_VOLUMES_INFO parameter 391
NSR_XBSA_DEBUG parameter 378
nsrd service 65
nsrdaprobe program 28, 110
nsrdasv program 65
nsrexecd service 65
nsrindexd service 67
nsrnotesrc options
    skip Notes API 232, 234
nsrnotesv options
    number of transaction logs 234
nsroraadmin program 326, 328, 329
nsroraclecat program 324, 332
nsrorainfo command 209
nsrsnapck program 332, 333
nsrsybcc
    command syntax 154
    Sybase roles command nsrsybcc 138
nsrsybrc
    Sybase roles 138
nsrsybsv
    Sybase roles 138
NWORA parameter resources
    NSR_ORACLE_NLS_LANG 325
    NSR_ORACLECAT_DEBUG_FILE 325
    NSR_ORACLECAT_LOG_FILE 325
    NSR_ORACLECAT_MODE 325, 333
    NSR_REMOVE_ON_FAILURE 325
NWORA resource file 324, 326, 327, 328
    backup 316, 318
NWORA SID resources 327, 328
nworapc directory 320

O

ONCONFIG parameter 371
online indexes 159
online NetWorker indexes 28
online redo logs, mirrored 237
Oracle
  DBID 237, 238
  manual backup script 129, 310, 402
  mirrored
    control files 237
    online redo logs 237
  password file 131
  recover 26, 168, 215
  Recovery Catalog database
    connection to 131
  restore 26, 168, 211, 272
  restore and recovery 207
  sbtio.log file 426
  scheduled backup script 130
Oracle Automated Storage Management (ASM) 308
Oracle DBA and NetWorker Backup Administrator 48
Oracle disk backups 48
Oracle Enterprise Manager Backup Management Tools 215
Oracle Net 267, 268
Oracle Real Application Cluster (RAC) 267
ORACLE_HOME parameter 391
ORACLE_SERVICE parameter 109
ORACLE_SID parameter 391
ORACLE_USER parameter 391
ORACLER_USER parameter 391

P

Parallelism attribute
  in Client resource 90
  in Server resource 75
parameter
  ADDRESS_LIST 268
  FAILOVER 268
  INSTANCE_NAME 269
  LOCAL_LISTENER 269
  SERVICE_NAME 268, 269
  SID_LIST_LISTENER 268, 269
parameter resources
  NSR_ORACLE_NLS_LANG 325
  NSR_ORACLECAT_DEBUG_FILE 325
  NSR_ORACLECAT_LOG_FILE 325
  NSR_ORACLECAT_MODE 325, 333
  NSR_REMOVE_ON_FAILURE 325
parameters
  DB2_ALIAS 366
  DB2_APPLY_NW_LEVELS 367
  DB2_BUFFER_SIZE 367
  DB2_NODE_NAME 368
  DB2_NUM_BUFFERS 368
  DB2_OPTIONS 368
  DB2_PARALLELISM 368
  DB2_PARTITION_LIST 368
  DB2_QUIESCE 368
  DB2_SESSIONS 369
  DB2_TBS_LIST 369
  DB2_USER 369
  DB2_VENDOR_LIB_PATH 369
  DB2INSTANCE 367
  DB2PATH 307, 368
  DBA_DISK_BACKUP 108, 288
  DBCCOPT 392
  DO_LOGFILE_BACKUPS 370
  DO_WHOLE_SYSTEM_BACKUP 370
  INFORMIXDIR 370
  INFORMIXSQLHOSTS 370
  INSTHOME 369
  LD_LIBRARY_PATH 393
  LD_LIBRARY_PATH_64 393
  LIBPATH 393
  LOG_THRESHOLD 108
  LOTUS_NSF_FILE 108
  LOTUS_USER 371
  MYSQL_BACKUP_DIR 380
  MYSQL_BACKUP_NAME 381
  MYSQL_BINLOG 381
  MYSQL_CFG_FILE 381
  MYSQL_COMPRESS 381
  MYSQL_COMPRESS_LEVEL 382
  MYSQL_DATABASES 382
  MYSQL_DATADIR 382
  MYSQL_EXTRACT_PATHS 382
  MYSQL_INCLUDE 383
  MYSQL_INCR_DIR 383
  MYSQL_INCR_OPTIONS 383
  MYSQL_INNODB_LOG_FILE_SIZE 383
  MYSQL_INNODB_LOG_FILES_IN_GROUP 383
  MYSQL_LOG_OPTIONS 384
  MYSQL_MEB_PATH 384
  MYSQL_ONLY_INNODB_OPTIONS 385
  MYSQL_RESTORE_OPERATION 385
  MYSQL_SBT_LIB_PATH 385
  MYSQL_SRC_ENTRY 386
  MYSQL_USER 386
  Notes_ExecDirectory 372
  NSM 304, 310, 312
  NSR_AES_ENCRYPTION 358
  NSR_APPLY_LOGS 372
  NSR_ASE_PASSWORD 393
  NSR_ASE_VERIFY 393
  NSR_AUTO_RESTORE 372
  NSR_BACKUP_ALL_EXTENSIONS 372
  NSR_BACKUP_LEVEL 372, 386, 394
  NSR_BACKUP_LOGS_MODE 372
  NSR_BACKUP_LOTUS_DIR 372
  NSR_BACKUP_PATHS 373, 394
  NSR_BROWSELIST_CACHE_DEST 373
  NSR_BUNDLING 135
  NSR_CATALOGFILE 373
  NSR_CHECKSUM 358
  NSR_CLIENT 286, 358
  NSR_COMFORT_SPAN 373
  NSR_COMPRESSION 359
  NSR_CROSS_MOUNT_POINTS 374
  NSR_DATA_DOMAIN_INTERFACE 359

NSR_DATA_MOVER 304
NSR_DATA_VOLUME_POOL 286, 287
NSR_DATA_VOLUME_POOL* 79, 359, 388, 389
NSR_DB2CAT_MODE 307
NSR_DBIID 374
NSR_DEBUG_LEVEL 108, 288, 360
NSR_DEDUP_BACKUP 360
NSR_DEDUP_CACHE_ENABLED 360
NSR_DEDUP_CACHE_TAG 361
NSR_DEDUP_CHUNK_SIZE 361
NSR_DEDUP_NODE 362
NSR_DEVICE_INTERFACE 362
NSR_DIAGNOSTIC_DEST 108, 288, 362
NSR_DIRECT_ACCESS 362
NSR_DPRINTF 363
NSR_DR_BACKUP_INFO 369, 371, 386
NSR_DR_FILE_LIST 369, 371, 386
NSR_DUMP_DATA_OPT 394
NSR_DUMP_LOG_OPT 394
NSR_ENCRYPTION_PHRASES 363
NSR_ENV_LIST 389
NSR_EXCLUDE_FILE 374, 394
NSR_EXCLUDE_LIST 374
NSR_FOLLOW_LINKS 374
NSR_GROUP 131, 363
NSR_INCR_BACKUP_LOGS_MODE 374
NSR_INCR_EXPIRATION 136
NSR_LOCALE 395
NSR_LOG_DIR 375
NSR_LOG_VOLUME_POOL 369, 371, 386, 395
NSR_LOTUS_DATA_DIR 375
NSR_MAX_START_RETRIES 287, 363
NSR_MAX_STREAMS 304
NSR_MAX_TXN_LOGS 375
NSR_MMDB_RETRY_TIME 389
NSR_NO_BUSY_ERRORS 364
NSR_NO_MULTIPLEX 287, 389
NSR_NO_NOTES_INIT 375
NSR_NOTES_CONNECT_TIMEOUT 375
NSR_NOTES_INI_PATH 375
NSR_NUMBER_LOGS 376
NSR_NWPATH 300, 325, 364
NSR_ORACLE_CONNECT_FILE 109, 327
NSR_ORACLE_HOME 327
NSR_ORACLE_LIB_PATH 327
NSR_ORACLE_NLS_LANG 300
NSR_ORACLE_RETENTION 390
NSR_ORACLE_SID 327
NSR_ORACLE_TNS_ADMIN 328
NSR_PARALLELISM 376, 395
NSR_PHYSICAL_HOST_LICENSE 364
NSR_PREFETCH_LOGS 376
NSR_PROMOTE_FULL 395
NSR_PROXY_PFILE 390
NSR_PS_SAVE_PARALLELISM 304
NSR_RECOV_INTERACT 376, 387
NSR_RECOV_LIST_FILE 376
NSR_RECOVER_OPTIONS 376
NSR_RECOVER_POOL 364
NSR_RECOVER_TIME 377, 387
NSR_RELOCATION_DEST 377
NSR_RESOURCE_DIR 377
NSR_RETAIN_NUM_LOGS 377
NSR_RMAN_ARGUMENTS 390
NSR_RMAN_OUTPUT 390
NSR_SAVESET_BROWSE 89, 365
NSR_SAVESET_NAME 377
NSR_SAVESET_RETENTION 91, 365
NSR_SERVER 131, 365
NSR_SERVER_NIC 287, 390
NSR_SKIPDBERRORS 377
NSR_SNAPSHOT_POSTCMD 304
NSR_VERBOSITY 378
NSR_VOLUMES_INFO 391
NSR_XBSA_DEBUG 378
ONCONFIG 371
ORACLE_HOME 391
ORACLE_SERVICE 109
ORACLE_SID 391
ORACLE_USER 391
PATH 378, 396
RESTORE_TYPE_ORDER 304, 321
SYBASE 396
SYBASE_USER 396
TNS_ADMIN 392
USE_CONSISTENCY_CHECK 396
USER_PSWD 370, 387, 396
parms option
allocate channel command 402
configure channel command 312, 388
partitioned Domino server
backup 120
recovery 190
password file 131
backup 237
PATH parameter 378, 396
permissions, Sybase 138
persistent settings 48, 130
PFILE 236, 238, 240
physical cluster client, proxy backups from 337
point-in-time (PIT) copies 294
point-in-time copy 31, 296, 300, 301, 318
policy
browse 89
retention 91
uniformity 55
pool option, backup command 398
Pool resource 78, 300
pool, volume
defined 78
label template for 79
postcommand script 238
PowerSnap Modules 31
precedence rules for send command 402
primary storage 31, 294, 314
Probe resource attributes
Command 107, 287
Command Options 108, 109, 288
Name 107, 287
probe-based backup 28

cluster 253
program
nsrd 65
nsrdaprobe 28, 110
nsrdasv 65
nsrexecd 65
nsrindexd 67
nsroraadmin 326, 328, 329
nsroraclecat 324, 332
nsrsnapck 332, 333
savefs 65
savegrp 65
scanner 169
proxy
backup 294
restore 296, 319
prune unwanted backups 165
pruning of DB2 snapshots 307
pruning of snapshots 36
psrollback.res file 321

R

RAC nodes as storage nodes 254
RAC systems 267
recover 26, 168, 215
rollforward 176
recoveries, types of
redirected restore 218
recovering
databases not on the master device 242
master database 241
recovery
partitioned Domino server 190
Recovery Catalog database
connection to 131
Recovery Catalog, backup 237, 240
Recovery Manager (RMAN)
commands 397
error messages 426
proxy backup scripts 310, 311
proxy restore scripts 319
restore scripts 211
scheduled backup scripts 130, 131
recovery options
number of transaction logs 234
skip Notes API 232, 234
redo logs, mirrored 237
registry files
backup 237
relocating files during proxy restores 321, 322
Remote Access attribute
in Client resource 90, 336
resource
types of
Client 301
Device 78
Group 87, 301
Pool 78, 300
Schedule 87
Server 75
Snapshot Policy 296, 301
User Group 75
restartable backups 41, 124, 161
restore
archived redo log 272
Avamar deduplication 29
cluster data 254
command 402
Data Domain deduplication 30
DB2 data 171
determining required volumes 209
from NetWorker User for Lotus 185
from NetWorker User for Sybase 220
Informix data 178
MySQL data 196
Oracle data 26, 168, 207, 211
physical restore 180
proxy 296, 319
redirected restore 218
rollback 321
Sybase data 216
restore from rollover 297
restore from secondary storage 297
RESTORE_TYPE_ORDER 303
RESTORE_TYPE_ORDER parameter 304, 321
Retain Snapshots attribute
in Snapshot Policy resource 323
Retention Policy attribute
in Client resource 91, 301
retention policy uniformity 55
RMAN catalog 322
rman command 153, 215
RMAN repository 322
rman send command 388, 401, 402
rman.exe command 153, 215
road map
manual backup 148
scheduled backup 146
rollback restore 297, 321
rollforward recovery 176
with DB2 transaction logs 35
rollover-only backup 296

S

Save Set attribute
in Client resource 91, 92, 93, 336
save set bundling 52
savefs command 65
savegroup completion report for proxy backup 315
savegrp command 65
sbtio.log file 426
scanner program 169
Schedule attribute
in Client resource 93
Schedule resource 87
scheduled backup 146
canceling 131, 147
cluster environment 250, 251
configuring Group resource 87
configuring Schedule resource 87

monitoring 147
road map 146
Scheduled Backup attribute
in Client resource 93
send command 312, 388, 398, 399
channel option 399
device_type option 399
NSR_ENV keyword 399
precedence rules 402
Server resource
attributes
Datazone pass phrase 75
Name 75
Parallelism 75
service
nsrd 65
nsrexecd 65
nsrindexd 67
SERVICE_NAME parameter 268, 269
SID_LIST_LISTENER parameter 268, 269
snapshot 31, 296, 300, 301, 318
backup and restore operations 294
backup processes 297
deletion 307
restore 306
snapshot and rollover backup 296
Snapshot attribute in Group resource 314
snapshot backup 296
Snapshot Policy resource 296, 301
Snapshot Pool attribute in Group resource 300
snapshot restore 297
SPFILE 236, 239
SQL error messages 413
staging 52
stopping
scheduled backup 131
storage devices
Device resource 78
volume pool 78, 79
Storage Nodes attribute
in Client resource 93
Sybase
restore and recovery 216
Sybase ASE Cluster Edition systems 272
SYBASE parameter 396
SYBASE_USER parameter 396
synchronization
automatic catalog 332
manual catalog 332

T

tablespace
manual backup 148
scheduled backup 146
target database
connection to 131
template file
nmda_db2.cfg 354
nmda_informix.cfg 354
nmda_lotus.cfg 354
nmda_mysql_backup.cfg 354
nmda_mysql_restore.cfg 354
nmda_oracle.cfg 354
nmda_sybase.cfg 354
threshold procedure, sample 142
TNS_ADMIN parameter 392
trace option, backup command 426
transaction logs
backup 116, 263
backup levels 40
configuration file 116
recovery with 176
rollforward recovery 116
troubleshooting 406
manual backup 407
scheduled backup 132, 407
types of backup
archived redo log 237, 272
manual 148
NetWorker bootstrap 66
scheduled 146

U

uniformity, policy 55
USE_CONSISTENCY_CHECK parameter 396
User Group resource 75
USER_PSWD parameter 370, 387, 396

V

VENDOROPT 164, 177
virtual cluster client, proxy backups from 336
virtual host 250
volume pool
Default pool 79
defined 78
resource 78, 300
specifying 79
volumes, determining for restore 209

W

wizard fails to create backup, troubleshooting 408