

Dell EMC NetWorker Module for Databases and Applications

Version 19.1

Administration Guide

302-005-541

REV 02

August 2019

Copyright © 2009-2019 Dell Inc. or its subsidiaries. All rights reserved.

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

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

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

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

CONTENTS

Tables		13
Figures		15
Preface		17
Chapter 1	Overview of NMDA Features	21
	Road map for NMDA operations.....	22
	Terminology that is used in this guide.....	23
	Importance of backups.....	24
	NMDA features for all supported applications.....	24
	Basic backup and recovery features.....	24
	Scheduled backups versus manual backups.....	25
	Client Direct feature.....	26
	Deduplication backups and restores with Data Domain.....	26
	Backups and restores with CloudBoost.....	27
	Snapshot backups and restores.....	28
	Virtualization support.....	28
	Cluster and high-availability system operations.....	29
	Configuration wizards.....	30
	Internationalization (I18N).....	30
	IPv6 support for backups and restores.....	31
	Multiple database or application installations on the same host.....	31
	Simplified scheduled backup configuration of multiple databases or applications.....	31
	NMDA features specific to DB2.....	31
	DB2 full, incremental, and delta backups.....	31
	DB2 transaction log backups.....	32
	DB2 history pruning.....	32
	DB2 10.5 features.....	32
	NMDA DB2 recovery wizard.....	33
	Restores of Fujitsu NetWorker Module for DB2 backups.....	34
	DB2 data recovery with the HPU utility.....	34
	DB2 load command with the copy yes option.....	34
	NMDA features specific to Informix.....	35
	Informix terminology.....	35
	Informix full and incremental backups.....	35
	Logical log backups.....	35
	Restores of Fujitsu NetWorker Module for Informix backups.....	36
	Informix 12.10 features.....	36
	NMDA features specific to Lotus.....	36
	Files that are backed up during Lotus backups.....	36
	Lotus full and incremental backups.....	37
	Lotus transaction log backups.....	37
	Lotus DAOS backups and restores.....	38
	Restartable scheduled backups.....	38

Types of Lotus restores.....	39
NetWorker User for Lotus program.....	39
NMDA Lotus recovery wizard.....	39
Partitioned Domino servers.....	40
NMDA features specific to MySQL.....	40
MySQL online and offline backups.....	40
MySQL full and incremental backups.....	40
MySQL binary log backups.....	41
MySQL backup granularity.....	41
MySQL restore and recovery operations.....	42
MySQL validate operation.....	44
MySQL 5.6 features.....	44
MySQL replicated slave server operations.....	44
NMDA features specific to Oracle.....	45
Oracle full and incremental backups.....	45
Backups of archived redo logs.....	45
Oracle DBA and NetWorker backup administrator collaboration.....	46
Control file and server parameter file autobackups.....	46
Automatic channel allocation.....	46
Backup and restore optimization.....	47
Backup copies.....	47
Backup and deletion of backup sets.....	47
Oracle Data Guard support.....	47
Oracle Exadata support.....	48
Restartable backups.....	48
Retention policies.....	49
NMDA Oracle recovery wizard.....	50
Restores of Fujitsu NetWorker Module for Oracle backups.....	51
Save set bundling.....	51
Policy uniformity.....	53
Other Oracle features.....	54
NMDA features specific to SAP IQ.....	57
Files that are backed up during SAP IQ backups.....	57
SAP IQ full, incremental, incremental since full, and transaction log backups.....	58
SAP IQ selective and all-inclusive backups.....	58
SAP IQ transaction log backups.....	58
SAP IQ restores.....	59
SAP IQ point-in-time recovery using log backups.....	60
NMDA features specific to Sybase.....	61
Sybase full, cumulative, and incremental backups.....	61
Sybase restore parameters.....	62
Restore of Sybase transaction logs only.....	62
Sybase up-to-the-minute recovery.....	62
Password-protected database backups and restores.....	63
Database consistency check.....	63
Database backup and restore verification.....	63
Exclusion of multiple user-defined temporary databases from backup....	63
Sybase in-memory and relaxed durability databases.....	64
Sybase ASE instance backups and restores with a large number of databases.....	64
Sybase ASE 15.7 features.....	64
NetWorker User for Sybase.....	65
NMDA Sybase recovery wizard.....	65
NMDA features specific to Orchestrated Application Protection.....	66

	NMDA features specific to MongoDB.....	67
	NMDA features specific to MySQL.....	67
	NMDA features specific to PostgreSQL.....	67
	NMDA components.....	69
	NMDA backup and restore processes.....	72
	Regular scheduled backup processes.....	72
	Regular restore processes.....	73
Chapter 2	Backup Configuration	75
	Configuring NMDA backups.....	76
	Configuring internationalization (I18N) support.....	77
	I18N support requirements.....	79
	Configuring I18N support for Informix operations.....	79
	Configuring I18N support for Oracle operations.....	79
	Configuring I18N support for SAP IQ operations.....	80
	Configuring I18N support for Sybase operations.....	81
	Verifying the database or application server configuration.....	81
	Verifying the basic NetWorker resource configurations.....	81
	Verifying the NetWorker Server resource.....	82
	Verifying the NetWorker User Group resource.....	82
	Verifying the NetWorker Client resource.....	85
	Verifying the NetWorker Device resource.....	85
	Verifying the NetWorker Pool and Label Template resources.....	85
	Configuring the firewall support.....	85
	Configuring scheduled backups with the wizard.....	86
	Requirements for using the backup configuration wizard.....	87
	Configuring a scheduled backup with the wizard.....	87
	Configuring scheduled backups without the wizard.....	89
	Configuring the NMDA parameters without the wizard.....	89
	Configuring the data protection policy with NMC.....	91
	Configuring the Client resource.....	92
	Configuring the backup group.....	98
	Configuring the backup policy.....	99
	Configuring the backup workflow.....	99
	Configuring the backup action.....	100
	Configuring manual backups.....	103
	Configuring Data Domain backups.....	105
	Best practices for Data Domain backups to any device.....	105
	Configuring DD Boost backups to use Client Direct.....	106
	Configuring DD Boost backups to use a storage node.....	108
	Configuring transaction log backups to traditional NetWorker devices....	108
	Configuring CloudBoost backups.....	109
	Best practices for CloudBoost backups.....	109
	Configuring CloudBoost backups to use Client Direct.....	109
	Configuring CloudBoost backups to use a storage node.....	110
	Configuring backups to AFTD storage.....	110
	Configuring Client Direct backups to AFTD.....	110
	Configuring storage node backups to AFTD.....	111
	Configuring probe-based backups.....	111
	NMDA probe.....	112
	User-defined probe.....	112
	Configuring a probe-based backup.....	113
	Configuring parallel backups.....	118
	Configuring DB2 parallel backups.....	120

Configuring Informix parallel backups.....	120
Configuring Lotus parallel backups.....	120
Configuring Oracle parallel backups.....	120
Configuring Sybase parallel (multistripe) backups.....	121
DB2 backup considerations.....	121
Best practices for DB2 backups.....	121
Configuring automatic deletion of DB2 recovery objects.....	122
Configuring automatic backups of DB2 transaction logs.....	122
Informix backup considerations.....	123
Configuring automatic (continuous) backups of Informix logical logs.....	123
Lotus backup considerations.....	124
Setting the environment for Lotus manual backups on Solaris or Linux....	124
Considerations for Lotus database or directory link backups.....	124
Considerations for Lotus transaction log backups.....	124
Lotus incremental backups with the comfort span option.....	126
Configuring partitioned Domino server backups.....	126
Configuring Lotus DAOS backups.....	127
Configuring Lotus restartable scheduled backups.....	130
MySQL backup considerations.....	130
Requirements of MySQL Enterprise Backup.....	130
Configuration parameter settings.....	130
Considerations for MySQL whole and partial backups.....	131
Considerations for MySQL incremental backups.....	132
Considerations for MySQL binary log backups.....	132
MySQL backup privileges.....	133
Configuring MySQL 5.6 features.....	133
Considerations for backups of MySQL replicated slave servers.....	134
Oracle backup considerations.....	134
RMAN scripts for manual backups.....	135
RMAN scripts for scheduled backups.....	135
Configuring automatic channel allocation.....	138
Creating Oracle backup copies.....	139
Configuring save set bundling for scheduled Oracle backups.....	141
Configuring policy uniformity for scheduled Oracle backups.....	142
Configuring operations in an Oracle Data Guard environment.....	143
Preventing possible performance issues with Oracle Exadata backups....	143
SAP IQ backup considerations.....	144
Setting up the SAP IQ interfaces file.....	144
Configuring the SAP IQ selective backups.....	145
Configuring the SAP IQ transaction log backups.....	145
Sybase backup considerations.....	146
Best practices for Sybase backups.....	146
Setting up Sybase roles and permissions.....	147
Specifying verification of Sybase database backups.....	148
Specifying volume pools for Sybase incremental backups.....	148
Sybase transaction log backups in an emergency.....	149
Configuring a Sybase threshold procedure.....	150
Setting the environment for Sybase backups on HP-UX.....	152
Configuring Sybase ASE 15.7 backup features.....	152
Orchestrated Application Protection backup considerations.....	153
MongoDB backup considerations.....	156
MySQL backup considerations.....	158
PostgreSQL backup considerations.....	161

Chapter 3	Backup Procedures	165
	Performing scheduled backups.....	166
	Testing scheduled backups.....	166
	Canceling scheduled backups.....	166
	Restarting failed scheduled backups.....	167
	Monitoring scheduled backups.....	167
	Performing manual backups.....	168
	Performing DB2 manual backups with the db2 backup command.....	168
	Performing DB2 manual backups with the DB2 GUI.....	169
	Performing Informix manual backups with the onbar command.....	170
	Performing Lotus manual backups with the nsrdasv command.....	170
	Performing Lotus manual backups with NetWorker User for Lotus.....	171
	Performing MySQL manual backups with the nsrdasv command.....	173
	Performing Oracle manual backups with the rman command.....	174
	Performing Oracle manual backups with Oracle Enterprise Manager..	174
	Performing SAP IQ manual backups with the nsrdasv command.....	175
	Performing Sybase database consistency checks before backups.....	175
	Performing Sybase manual backups with the nsrdasv command.....	176
	Performing Sybase manual backups with NetWorker User for Sybase....	176
	Performing Orchestrated Application Protection manual backups with	
	the nsroappbackup command.....	178
	Canceling manual backups.....	179
	Monitoring manual backups.....	180
	Verifying backup information in NetWorker indexes.....	180
	Verifying Data Domain backup information in NetWorker indexes.....	181
	Lotus restartable backup information in NetWorker indexes.....	181
	MySQL backup information in NetWorker indexes for backups without	
	Orchestrated Application Protection.....	182
	MongoDB backup information in NetWorker media database.....	183
	MySQL backup information in NetWorker media database for backups	
	with Orchestrated Application Protection.....	184
	PostgreSQL backup information in NetWorker media database.....	184
	Synchronizing backup catalogs and deleting backups.....	185
	Deleting DB2 backups.....	185
	Deleting Informix backups.....	187
	Cross-checking and deleting Oracle backups.....	187
Chapter 4	Data Restore and Recovery	189
	Data restore and recovery terminology.....	190
	NetWorker indexes and policies that are used for restores.....	190
	Performing NMDA data restore and recovery.....	191
	Performing DB2 data restore and recovery.....	193
	Determining how many restore devices and sessions to use.....	194
	Performing DB2 data restores with the db2 restore command.....	194
	Performing DB2 data recovery.....	198
	Performing DB2 data restore and recovery with the NMDA DB2	
	recovery wizard.....	201
	Performing DB2 data restore and recovery with the DB2 GUI.....	202
	Performing DB2 data recovery with the DB2 HPU utility.....	202
	Performing DB2 rollforward recovery after a load with copy yes option..	
	203	
	Performing Informix data restore and recovery.....	204
	Determining the Informix restore mode.....	204
	Performing Informix data restores with the onbar command.....	205

Performing Lotus data restore and recovery.....	207
Performing Lotus database recovery with the nsrnotesrc command..	208
Performing Lotus database restore and recovery with the NMDA Lotus recovery wizard.....	211
Performing Lotus database recovery with NetWorker User for Lotus.	212
Performing recovery of partitioned Domino servers.....	217
Performing Lotus DAOS data recovery.....	217
Performing Lotus document-level recovery with the nsrdocrc command.....	218
Performing Lotus document-level recovery with the Notes client GUI....	220
Performing MySQL data restore and recovery.....	222
Determining the binary logs for MySQL recovery.....	223
Performing MySQL restore and recovery with the NMDA MySQL recovery wizard.....	224
Performing MySQL recovery of whole instance backups.....	225
Performing MySQL restores of partial backups.....	227
Performing MySQL restores of binary log backups.....	228
Performing MySQL redirected restores.....	229
Performing MySQL list image operations.....	230
Performing MySQL extract operations.....	231
Performing MySQL extract and prepare operations.....	232
Performing MySQL copy back operations.....	233
Performing MySQL restores of InnoDB tables outside data directory	234
Performing MySQL validate operations.....	234
Performing MySQL recovery of a replicated slave server.....	235
Performing Oracle data restore and recovery.....	235
Determining the volumes for Oracle restores.....	235
Preventing restore performance degradation with Oracle 10.2 or later....	239
RMAN scripts for restore and recovery.....	239
Performing Oracle restore and recovery with the NMDA Oracle recovery wizard.....	240
Performing Oracle restores with the rman command.....	243
Performing Oracle restores with the Oracle Enterprise Manager.....	243
Performing Oracle data recovery.....	243
Performing SAP IQ data restore.....	244
Prerequisites for SAP IQ data restore.....	244
Performing the SAP IQ data restore with the nsriqrc command.....	245
Performing Sybase data restore and recovery.....	251
Specifying verification of Sybase database restores.....	251
Obtaining backup information with the listonly option.....	252
Performing restores of Sybase ASE backups with a large number of databases.....	254
Performing Sybase data restores with the nsrsybrc command.....	254
Performing Sybase data restore and recovery with the NMDA Sybase recovery wizard.....	260
Performing Sybase data restores with NetWorker User for Sybase...	262
Performing Orchestrated Application Protection data restore.....	264
Performing the restore of MongoDB data.....	264
Performing the restore of MySQL data.....	266
Performing the restore of PostgreSQL data.....	268
Canceling a data restore or recovery.....	272

- Preparing for disaster recovery..... 274
- Preparing for DB2 disaster recovery..... 275
 - Performing a DB2 disaster recovery..... 276
- Preparing for Informix disaster recovery..... 276
 - Performing an Informix imported restore..... 277
 - Recovering from an Informix server disk crash..... 279
- Preparing for Lotus disaster recovery..... 280
 - Recovering a nonlogged Lotus environment..... 280
 - Recovering a logged Lotus environment..... 281
- Preparing for MySQL disaster recovery..... 284
 - Performing a MySQL disaster recovery..... 284
- Preparing for Oracle disaster recovery..... 285
 - Creating the DBID text file..... 286
 - Setting up a postcommand script for backup of Oracle files..... 286
 - Setting up RMAN backups of the database and related files..... 287
 - Setting up RMAN backups of Recovery Catalog..... 287
 - Performing an Oracle disaster recovery..... 287
- Preparing for Orchestrated Application Protection disaster recovery..... 292
 - Performing an Orchestrated Application Protection disaster recovery.... 292
- Preparing for SAP IQ disaster recovery..... 293
- Preparing for Sybase disaster recovery..... 293
 - Recovering the master database..... 294
 - Recovering user databases after database device failure..... 295

- Chapter 6 Cluster and High-Availability (HA) Systems 297**
- Active-passive clusters..... 298
 - Backup failover..... 299
 - Cluster configuration requirements..... 299
 - Configuring scheduled backups in a cluster with the wizard..... 300
 - Configuring backups in a cluster without the wizard..... 301
 - Configuring probe-based backups in a cluster..... 303
 - Setting up nodes to back up to a local storage node..... 303
 - Enabling recoveries in a cluster..... 304
- Active-active application clusters and HA systems..... 304
 - Configuring an active-active application cluster..... 305
 - Restarting backups after a system failure..... 306
- DB2 DPF systems..... 306
 - Configuring backups and restores in a DPF environment..... 307
 - Configuring manual DPF backups..... 308
 - Performing manual DPF backups..... 308
 - Configuring scheduled DPF backups..... 309
 - Restoring and recovering DPF backups..... 310
- DB2 pureScale systems..... 310
 - Configuring backups and restores in a DB2 pureScale system..... 311
 - Configuring manual backups in a DB2 pureScale system..... 312
 - Performing manual backups in a DB2 pureScale system..... 312
 - Configuring scheduled backups in a DB2 pureScale system..... 312
 - Configuring automatic backups of transaction logs in a DB2 pureScale system..... 313
 - Restoring and recovering backups in a DB2 pureScale system..... 313
 - Restoring between a DB2 pureScale instance and Enterprise Server Edition..... 314
- DB2 HADR systems..... 315
 - Configuring backups and restores in a DB2 HADR system..... 316

	Configuring backups and restores with DB2 HADR in a DB2 pureScale system.....	317
	Performing DB2 HADR data recovery.....	317
	Informix MACH systems.....	318
	Using Continuous Log Restore for backup and restore.....	319
	Backups of primary and HDR or RSS or SDS servers.....	320
	Restores of primary and HDR or RSS or SDS servers.....	320
	Oracle RAC systems.....	320
	Configuring backups and restores in a RAC environment.....	320
	Creating RMAN backup and restore scripts.....	321
	Archived redo logs.....	323
	Oracle RAC One Node systems.....	324
	Sybase ASE Cluster Edition systems.....	325
	Configuring backups and restores for ASE Cluster Edition.....	325
	Skipping the backup of temporary databases.....	326
	Restoring an ASE Cluster Edition backup.....	326
	Performing backup and recovery of the quorum device.....	327
	Recovering the master database on clusters.....	327
Chapter 7	Multiple Installations on a Single Host	329
	Multiple database or application installations.....	330
	Multiple databases or applications with the same bitness.....	330
	Multiple databases or applications with 32-bit and 64-bit coexistence.....	330
Chapter 8	Oracle DBA and NetWorker Backup Administrator Collaboration	333
	DBA disk backup issues in Oracle environments.....	334
	NMDA solution for DBA disk backup issues.....	334
	Configuring backups of DBA disk backups.....	336
	Advanced NMDA parameters.....	338
	Configuring optional probe-based backups of DBA disk backups.....	339
	Reporting on backups of DBA disk backups.....	340
	Purging DBA disk backups.....	342
	Restore and recovery of DBA disk backups.....	342
Chapter 9	Snapshot Backups and Restores	345
	Snapshot operations with NetWorker Snapshot Management.....	346
	Software requirements for NSM snapshot operations.....	346
	Types of NSM backups.....	347
	Types of NSM restores.....	348
	NSM backup processes.....	349
	Configuring NSM snapshot backups.....	351
	Configuring Pool resources for NSM backups.....	353
	Configuring internationalization (I18N) support.....	353
	Configuring Client resources manually for NSM backups.....	354
	Testing an NSM scheduled backup.....	354
	Canceling an NSM snapshot backup.....	354
	Configuring NSM snapshot restore and recovery.....	355
	Setting the NSM parameters.....	358
	DB2 considerations for NSM operations.....	362
	Performing manual ProtectPoint NSM backups.....	362
	Restoring a DB2 NSM backup.....	363
	Managing and deleting DB2 NSM backups.....	364
	NSM backups and restores on cluster systems.....	367
	Oracle considerations for NSM operations.....	368

	Configuring the required Oracle settings.....	369
	Configuring the NWORA resource file.....	370
	Creating RMAN scripts for NSM snapshot backups.....	370
	Checking configuration consistency.....	374
	Performing Oracle NSM backups.....	375
	Verifying Oracle NSM backup information in NetWorker indexes.....	376
	Restoring an Oracle NSM backup.....	379
	Catalog synchronization for Oracle NSM backups.....	382
	NSM backups and restores on cluster systems.....	394
	Snapshot operations with Oracle ASM.....	397
Appendix A	NMDA Parameters and Configuration File	399
	Setting the NMDA parameters.....	400
	NMDA configuration file.....	400
	NMDA configuration file syntax.....	402
	Common NMDA parameters.....	406
	NMDA DB2 parameters.....	414
	NMDA Informix parameters.....	421
	NMDA Lotus parameters.....	424
	Wildcard support for Lotus operations.....	433
	Setting the NSR_RECOVER_TIME parameter.....	434
	NMDA MySQL parameters.....	435
	NMDA Oracle parameters.....	444
	NMDA Orchestrated Application Protection parameters.....	451
	NMDA parameters for Orchestrated Application Protection backups.....	451
	NMDA parameters for Orchestrated Application Protection restores.....	456
	NMDA SAP IQ parameters.....	457
	NMDA Sybase parameters.....	464
Appendix B	Oracle RMAN Commands	475
	The pool option of the backup command.....	476
	The send command.....	476
	Syntax rules.....	476
	Two ways to run the send command.....	478
	Precedence rules.....	480
Appendix C	Troubleshooting and Error Messages	483
	General troubleshooting tips.....	484
	Debug log files.....	485
	Wizard backup configuration fails, authentication denied.....	486
	Authentication error, unable to obtain the user credentials from the lockbox.....	487
	Backup becomes suspended.....	487
	NMDA error messages.....	487
	NetWorker XBSA error messages.....	488
	DB2 troubleshooting tips and error messages.....	492
	DB2 rollforward might fail on Windows with DB2 9.7 due to the logarchopt2 value.....	492
	Debug log issue with load command and copy yes option.....	492
	Deletion of DB2 snapshot save sets might fail with an error.....	492
	Delta or incremental backup failure in a DB2 pureScale system.....	492
	Delta or incremental restore failure.....	493
	Log retrieval issue for re-created database.....	493
	Pruning of DB2 log history after a rollforward might produce errors..	494

Pruning of log backups fails due to the logarchoptn value.....	495
Removal of failed backup.....	495
SOURCE_CLIENT prevents issue in rollforward with archived log backup.....	496
Wizard configuration or scheduled backup on AIX might fail due to DB2 10.5 and 10.1 issues.....	496
The load libnsrdb2 command.....	498
DB2 SQL error messages.....	498
Informix troubleshooting tips and error messages.....	502
No dbspaces/blobspaces found to back up or restore.....	502
Unable to open connection to server.....	503
Default value that is assigned to LTAPEDEV causes failure.....	503
ON-Bar returned codes 2 through 34.....	503
Timeout might occur for an NMDA Informix manual onbar backup after upgrade to NetWorker 9.x.....	504
Lotus troubleshooting tips and error messages.....	504
"Invalid Time Specified" error.....	504
Backup and recovery of Domino server fails with secured console.....	504
Backup failure due to time conversion.....	504
Child process termination for canceled Lotus backups.....	505
Notes API 525 error.....	505
NMC Group Details window displays "command not found".....	505
Prevent memory allocation failure with 32-bit Lotus Domino restores on AIX.....	506
Recovery report from NetWorker User for Lotus contains unexpected characters.....	506
Restore problem due to NSR_BACKUP_PATHS=NOTES:.....	506
Failure of a document-level recovery on a remote Domino server.....	506
NMDA Lotus error messages.....	507
MySQL troubleshooting tips and error messages.....	510
NMDA MySQL backup issues.....	510
NMDA MySQL restore issues.....	511
NMDA MySQL error messages.....	511
Oracle troubleshooting tips and error messages.....	513
Diagnosing NMDA Oracle problems.....	513
RMAN error messages.....	514
The trace option of the backup command.....	514
Canceling unresponsive Oracle backups on UNIX.....	515
Canceling unresponsive Oracle backups on Windows.....	516
Oracle rollback restore to a new database might fail when OMF is enabled.....	516
Rollback restore fails for Oracle NSM backup in FRA.....	516
NMDA Oracle error messages.....	517
Orchestrated Application Protection troubleshooting tips.....	531
Orchestrated Application Protection backup might fail on Linux with a fusermount error.....	532
Orchestrated Application Protection backup might become suspended on Linux with database password authentication.....	532
SAP IQ troubleshooting tips and error messages.....	532
Sybase troubleshooting tips and error messages.....	533
Sybase backup failure due to .lock files.....	533
NMDA Sybase error messages.....	533

TABLES

1	Revision history.....	17
2	Style conventions.....	18
3	NMDA components.....	69
4	NMDA parameters that support non-ASCII values.....	78
5	NetWorker Server resource attributes.....	82
6	User group privileges for NMDA operations.....	83
7	Database-specific user for each database or application.....	83
8	Mandatory parameters for scheduled NMDA backups.....	89
9	NetWorker Client resource attributes.....	93
10	Backup levels that are specified in the backup action.....	100
11	Mandatory parameters for manual NMDA backups.....	103
12	Configuring a Data Domain backup.....	107
13	NetWorker Probe resource attributes.....	113
14	Configuration parameters for InnoDB partial backups.....	131
15	Sybase roles and permissions.....	147
16	Backups with NSR_DUMP_LOG_OPT.....	149
17	Threshold procedure location.....	151
18	Command line options for Orchestrated Application Protection backups.....	154
19	NMDA MySQL operations through MEB and Orchestrated Application Protection.....	158
20	Mapping of NetWorker level backups to PostgreSQL backups.....	161
21	Manual backup procedures.....	168
22	Database consistency check options of the nsrsybcc command.....	176
23	Data restore procedures.....	192
24	Configuration file parameters corresponding to Recover Options fields.....	215
25	Command line options for PostgreSQL restores.....	271
26	Commands or library names for 32-bit and 64-bit coexistence.....	331
27	Advanced NMDA parameters for backups of DBA disk backups.....	338
28	NetWorker Probe resource attributes.....	339
29	Typical configuration for NSM operations with NMDA.....	346
30	Supported types of NSM backups.....	347
31	Supported types of NSM restores.....	348
32	Common NSM parameters.....	358
33	NSM parameters for ProtectPoint for VMAX operations only.....	359
34	NSM parameters for ProtectPoint with RecoverPoint operations only.....	360
35	DB2 resource file parameters.....	366
36	NWORA parameter resources.....	384
37	NWORA SID resource components.....	386
38	Common NMDA parameters.....	406
39	NMDA DB2 parameters.....	415
40	NMDA Informix parameters.....	421
41	NMDA Lotus parameters.....	424
42	NMDA MySQL parameters.....	436
43	NMDA Oracle parameters.....	445
44	NMDA parameters for Orchestrated Application Protection backups.....	451
45	NMDA parameters for Orchestrated Application Protection restores.....	456
46	NMDA SAP IQ parameters.....	458
47	NMDA Sybase parameters.....	464
48	Option values in the send command.....	478
49	NetWorker XBSA error messages.....	488
50	Path and suffix for the load libnsrdb2 command.....	498
51	DB2 SQL error messages	498
52	Error messages from Lotus backups and restores.....	507
53	Error messages from MySQL backups and restores.....	512

54	Trace option values and conditions traced.....	515
55	Error messages from Oracle backups and restores.....	517
56	Error messages from the nsroraadmin program.....	527
57	Error messages from the nsrorainfo program.....	530
58	SAP IQ issues and resolutions.....	533
59	Error messages from Sybase consistency checks, backups, and restores.....	534
60	Sybase backup server and libnsrsyb error messages.....	538


FIGURES

1	Regular scheduled NMDA backup.....	73
2	Backup window in NetWorker User for Lotus.....	172
3	Backup Options dialog box in NetWorker User for Lotus.....	173
4	Backup Options dialog box in NetWorker User for Sybase.....	177
5	Lotus restartable backup information in NMC Console.....	182
6	Recover window in NetWorker User for Lotus.....	213
7	Recover Options dialog box in NetWorker User for Lotus.....	214
8	Recover to Sybase Server dialog box in NetWorker User for Sybase.....	263
9	Single database node with failover capability.....	298
10	Multiple database nodes with failover capability.....	299
11	Multiple database nodes with mutual failover capability.....	299
12	Single DPF host with shared memory.....	307
13	Two DPF hosts with multiple nodes.....	307
14	NSM snapshot and clone to secondary storage data flow with NMDA.....	350
15	NSM snapshot restore data flow with NMDA.....	351

Preface

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

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

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

Purpose

This document describes how to configure and use the NetWorker Module for Databases and Applications (NMDA) version 19.1.

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 the following topics:

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

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
02	August 30, 2019	Updated the following topics: <ul style="list-style-type: none">• PostgreSQL full and transaction log backups (Chapter 1)—Added information about transaction log backups in the first Note. Added the second Note about transactions and WAL segment files.• PostgreSQL restores (Chapter 1)—Added the Note about transactions and WAL segment files. Updated the paragraph about registering the <code>nsroapprecover</code> program through the <code>restore_command</code> setting in the <code>recovery.conf</code> file.• NMDA components (Chapter 1)—In Table 3, updated the description of the NMDA Orchestrated Application Protection component <code>nmda_oapp_postgresql.example</code>.• PostgreSQL transaction log backup (Chapter 2)—Updated the third paragraph about command settings in the

Table 1 Revision history (continued)

Revision	Date	Description
		<code>postgresql.conf</code> file. Updated the Note with information about transactions and WAL segment files.
01	May 20, 2019	Initial release of this document for NMDA version 19.1.

Related documentation

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


The *NetWorker E-LAB Navigator* at <https://elabnavigator.emc.com/eln/elhome> provides details about supported environments and platforms.


The following additional documentation might be useful:

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

Special notice conventions that are used in this document

The following conventions are used for special notices:

 **NOTICE** Identifies content that warns of potential business or data loss.

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

Typographical conventions

The following type style conventions are used in this document:

Table 2 Style conventions

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

Table 2 Style conventions (continued)

... Ellipses indicate non-essential information that is omitted from the example.

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

Where to find product documentation

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

Where to get support

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

To access a product-specific page:

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

Knowledgebase

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

To search the Knowledgebase:

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

Live chat

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

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

Service requests

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

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

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

To review an open service request:

1. Go to <https://www.dell.com/support>.

2. On the **Support** tab, click **Service Requests**.
3. On the **Service Requests** page, under **Manage Your Service Requests**, click **View All Dell Service Requests**.

Online communities

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

How to provide feedback

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

CHAPTER 1

Overview of NMDA Features

This chapter includes the following topics:

• Road map for NMDA operations	22
• Terminology that is used in this guide	23
• Importance of backups	24
• NMDA features for all supported applications	24
• NMDA features specific to DB2	31
• NMDA features specific to Informix	35
• NMDA features specific to Lotus	36
• NMDA features specific to MySQL	40
• NMDA features specific to Oracle	45
• NMDA features specific to SAP IQ	57
• NMDA features specific to Sybase	61
• NMDA features specific to Orchestrated Application Protection	66
• NMDA components	69
• NMDA backup and restore processes	72

Road map for NMDA operations

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

About this task

Before you perform NMDA operations, review the *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 *NetWorker Module for Databases and Applications Installation Guide*.

i **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.

Procedure

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

- [Terminology that is used in this guide](#) on page 23
- [Importance of backups](#) on page 24
- [NMDA features for all supported applications](#) on page 24
- Information about the NMDA features for a specific database or application:
 - [NMDA features specific to DB2](#) on page 31
 - [NMDA features specific to Informix](#) on page 35
 - [NMDA features specific to Lotus](#) on page 36
 - [NMDA features specific to MySQL](#) on page 40
 - [NMDA features specific to Oracle](#) on page 45
 - [NMDA features specific to SAP IQ](#) on page 57
 - [NMDA features specific to Sybase](#) on page 61
 - [NMDA features specific to Orchestrated Application Protection](#) on page 66

If you have multiple database installations on the NMDA host, review [Multiple Installations on a Single Host](#) on page 329.

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

- [NMDA components](#) on page 69
- [NMDA backup and restore processes](#) on page 72

i **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 [Oracle DBA and NetWorker Backup Administrator Collaboration](#) on page 333.

2. To configure the NMDA backups, use the following information:

- In a stand-alone environment, configure the regular (not snapshot) backups according to [Configuring NMDA backups](#) on page 76.

- In a cluster environment, configure the regular backups according to the information in [Cluster and High-Availability \(HA\) Systems](#) on page 297.
- Configure the DB2 or Oracle snapshot backups according to the information in [Snapshot Backups and Restores](#) on page 345.

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

3. To perform the NMDA backups, use the following information:
 - In a stand-alone environment, perform the regular backups according to [Performing scheduled backups](#) on page 166 or [Performing manual backups](#) on page 168.
 - In a cluster environment, perform the regular backups according to the information in [Cluster and High-Availability \(HA\) Systems](#) on page 297.
 - Perform the DB2 or Oracle snapshot backups according to the information in [Snapshot Backups and Restores](#) on page 345.

You can verify information about NMDA backups according to [Verifying backup information in NetWorker indexes](#) on page 180.

You can synchronize information in the application backup catalog and NetWorker indexes according to [Synchronizing backup catalogs and deleting backups](#) on page 185.

4. To perform the restore and recovery of NMDA backups, use the following information:
 - In a stand-alone environment, restore the regular backups according to [Performing NMDA data restore and recovery](#) on page 191.
 - In a cluster environment, restore the regular backups according to the information in [Cluster and High-Availability \(HA\) Systems](#) on page 297.
 - Restore the DB2 or Oracle snapshot backups according to the information in [Snapshot Backups and Restores](#) on page 345.
5. Use the information in [Disaster Recovery](#) on page 273 to perform disaster recovery procedures.
6. Use the following information to troubleshoot any issues with NMDA operations:
 - [Troubleshooting and Error Messages](#) on page 483
 - *NetWorker Module for Databases and Applications Release Notes*

Terminology that is used in this guide

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

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.

The glossary provides details about the terms that are used in this guide.

Importance of backups

You must perform regular backups of specific data to prepare for the recovery of the database or application system.

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, SAP IQ, and Sybase ASE data. NMDA also provides data protection for MongoDB, MySQL, and PostgreSQL data through the Orchestrated Application Protection feature. NMDA operates with the NetWorker software and the supported database software or application software.

 **Note:** The product name Sybase ASE is synonymous with SAP ASE.

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, CloudBoost devices, and tape devices.

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

Basic backup and recovery features

NMDA supports basic backup and recovery features in addition to specific backup and recovery granularities for the different databases and applications.

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 that are 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:** Also, Oracle software can perform block-level recovery.
 - SAP IQ database, dbspace, and dbfile backup and 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 118.

Scheduled backups versus manual backups

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

The two types of backups are initiated through different methods:

- The NetWorker server initiates a scheduled backup:
 - A regular (time-based) scheduled backup starts at a time that is specified in the backup configuration settings.
 - 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 number of logs that were 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 111.

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

A backup with NetWorker Snapshot Management (NSM) is supported through a scheduled backup. A ProtectPoint backup with NSM is the only type of NSM backup that is also supported through a manual backup.

- A user initiates a manual backup on the application host through a database or application-specific command or procedure that runs the appropriate backup utility.

[Backup Configuration](#) on page 75 describes the configuration of both scheduled and manual backups.

[Backup Procedures](#) on page 165 describes the procedures for both scheduled 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 that is required for all backed-up data.

NOTICE

The NetWorker server backup by default backs up the bootstrap and client index information. The NetWorker disaster recovery documentation describes how to prepare for NetWorker server disaster recovery.

A NetWorker data protection policy applies only to a scheduled backup, and defines the workflow, action, backup group, and other required settings for a scheduled backup. NetWorker features that are related to a data protection policy, such as the clone action and the protection period, apply only to scheduled backups and do not apply to manual backups. [Configuring the data protection policy with NMC](#) on page 91 provides more details about data protection policies.

Client Direct feature

NMDA supports the Client Direct feature, which enables client backups to bypass the NetWorker storage node and send the backup data directly to the supported devices.

You can perform any of the following operations with the Client Direct feature:

- Send deduplicated backup data directly to Data Domain Boost (DD Boost) devices.
- Send backup data directly to advanced file type devices (AFTDs).
- Send backup data directly to CloudBoost devices.

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.

 **Note:** NMDA does not support Client Direct Lotus backups that are checkpoint restart enabled.

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 Data Domain

NMDA supports deduplication backups and restores with a 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

NMDA can use the Client Direct feature or the NetWorker storage node to perform 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:

- **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.
- **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.
- **Deduplication on a local storage node (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 that is used, the Data Domain system converts the stored data to its original nonduplicated 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 105.

Backups and restores with CloudBoost

NMDA supports backups and restores with a CloudBoost appliance. You can configure the CloudBoost system as a NetWorker CloudBoost device.

A CloudBoost backup to a CloudBoost device can be a manual backup or a scheduled backup, including a probe-based backup.

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

- **CloudBoost backups using Client Direct**—With the Client Direct workflow, NMDA uses the CloudBoost components to send the data 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 data directly from the NMDA client to the CloudBoost appliance.

- CloudBoost backups using a remote storage node—With the NetWorker storage node workflow, NMDA sends all the data to a remote storage node, which sends the data to the target device. NMDA uses this workflow if the NMDA client does not have direct access to the CloudBoost appliance or if the Client Direct workflow is not possible for other reasons.

By default, NMDA tries to perform a Client Direct backup. If the NMDA client does not have direct access to the CloudBoost appliance 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 CloudBoost appliance, 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.

You must configure a CloudBoost backup with a supported CloudBoost device according to [Configuring CloudBoost backups](#) on page 109.

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. 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 VMAX (Symmetrix), VNX Block (CLARiiON), and XtremIO. You can optionally perform a clone operation after a snapshot backup, to create a cloned copy of the snapshot data on secondary storage, such as a Data Domain device or tape.
- 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 replication management software 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 Client Direct backup to an AFTD or DD Boost device. The Client Direct feature applies only to a snapshot clone operation and sends the backup of the snapshot from the mount host to the device directly, bypassing the storage node.

[Snapshot Backups and Restores](#) on page 345 provides more details about snapshot backups and restores.

Virtualization support

NMDA supports several types of virtualization software, such as VMware, Solaris zones, and Microsoft Hyper-V. The *NetWorker E-LAB Navigator* at <https://elabnavigator.emc.com/eln/elhome> provides details about supported environments and platforms.

 **Note:** Ensure that you 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 Retries setting in the backup action resource has a nonzero value.

Solaris zones support

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

Cluster and high-availability system operations

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

The two types of supported clusters include the following features:

- 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. NMDA supports the following active-active application clusters:
 - DB2 Database Partitioning Feature (DPF)
 - DB2 pureScale
 - Informix Multi-node Active Clusters for High Availability or High Availability Clusters
 - ⓘ **Note:** NMDA documentation uses the MACH acronym for both Informix technologies: Informix Multi-node Active Clusters for High Availability and High Availability Clusters.
 - Oracle Real Application Clusters (RAC)
 - Sybase ASE Cluster Edition

NMDA also supports backups and restores in a DB2 High Availability Disaster Recovery (HADR) system. The DB2 servers in this system use the HADR feature, which automatically replicates the data changes from a primary node database to an identical database on one or more standby nodes.

[Cluster and High-Availability \(HA\) Systems](#) on page 297 describes the supported cluster and high-availability systems and their configurations for NMDA backups and restores.

Configuration wizards

NMDA provides 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 provides 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 86 describes how to use the backup configuration wizard.
- Recovery wizard that configures and runs the restore and recovery of DB2, Lotus Domino/Notes, MySQL, Oracle, and Sybase data from NMDA backups. Application-specific sections later in this chapter and in [Data Restore and Recovery](#) on page 189 provide more information about the recovery wizard.

The scheduled backup configuration wizards do not support the following features:

- Configuration of probe-based (event-based) backups.
 - Note:** When you create a NetWorker Client resource with the wizard, you can associate the resource with a probe by selecting the Probe attribute in the Client resource in NMC, outside of the wizard.
- Configuration of backups in active-active application clusters, such as DB2 Database Partitioning Feature (DPF), DB2 pureScale, Informix Multi-node Active Clusters for High Availability or High Availability Clusters, Oracle Real Application Clusters (RAC), and Sybase ASE Cluster Edition.
 - Note:** NMDA documentation uses the MACH acronym for both of the following Informix technologies: Informix Multi-node Active Clusters for High Availability and High Availability Clusters.
- Configuration of backups in a DB2 High Availability Disaster Recovery (HADR) cluster, which is an active-passive cluster.
- Configuration of split-mirror-based backups of Oracle ASM data with replication management software.
- Configuration of two different database backups in the same Client resource.
- Modification of a backup configuration that was created with the wizard from a legacy NetWorker module.
- Modification of a client-side NMDA configuration that was created with the NMC method without a 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.

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 77, 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.

NMDA supports the Data Domain IPv6, CloudBoost IPv6, and mixed IPv4/IPv6 networks.

The *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.

[Multiple Installations on a Single Host](#) on page 329 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 [Table 9](#) on page 93. 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 following settings determine whether NMDA backs up the multiple databases simultaneously or sequentially:

- Parallelism attribute setting in the NetWorker Client resource.
- NSR_PARALLELISM parameter setting.
- Setting of the number of backup sessions/channels in the backup configuration wizard.

During NetWorker scheduled backups, if the number of current running backup sessions is greater than the Parallelism attribute setting in the NetWorker Client resource, a newly started scheduled backup runs sequentially. The *NetWorker Administration Guide* provides details about the NetWorker client 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 full, incremental, and delta backups of DB2 database data and logs.

The supported types of DB2 backups include the following features:

- 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 122.

Note: For a re-created database, to prevent issues during a DB2 log retrieval as in a rollforward operation, remove the old log backups before you perform an archived log backup. Using a different log volume pool for the backup might not prevent the issues because the backup index search is not currently based on the pool setting.

Indexing of the archived log backups is based on the instance, the database name, and the log chain and sequence numbers. Therefore, an incorrect log might be retrieved when a database has multiple sets of archived log backups and the same index names appear in different sets.

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 365 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 185 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.

The following new features of DB2 10.5 are supported by the latest NMDA release:

- Increased availability support for topology changes in a DB2 pureScale environment. For example, after you add a pureScale node, DB2 software no longer requires an offline immediate database backup. You can restore backup images that were 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 311.

- Delta and incremental backups in a DB2 pureScale environment starting with DB2 10.5 fp4.
- 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.

NMDA DB2 recovery wizard

NMDA supports the DB2 recovery wizard, which is a Java application that you run from NMC to configure and run a restore and recovery of DB2 data that is backed up by NMDA. If the wizard supports the environment, use the wizard whenever possible to perform the DB2 restore and recovery operations.

You can run the recovery 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.

The NMDA DB2 recovery wizard can configure and run the following types of restore and recovery:

- Complete database recovery to the time before a failure.
- Point-in-time recovery of a whole database or selected tablespaces.
- Redirected restore of a database, including relocation of data if required.
- Recovery of DB2 data to a new database on either the same host or a different host.
- Restore of the history file only.
- Restore of the log files only.

NMDA 9.1 introduced the following options in the NMDA DB2 recovery wizard:

- Restore of the log files only.
- Query of the available backups for restore from either the DB2 history file or the backup storage.
- `NEWLOGPATH` option whereby you specify a directory to use for the log files after the database restore.
- `OVERFLOW LOG PATH` option whereby you specify an overflow log path to be searched for the logs that will be used for a rollforward operation.

The recovery wizard enables you to start a restore or recovery immediately or schedule the operation to start at a future time.

The recovery wizard does not support the following features:

- DB2 DPF systems.
- DB2 pureScale systems.

The following additional sources describe the wizard:

- Descriptive inline text in the wizard
- Online help in the wizard

- *NetWorker Module for Databases and Applications Release Notes*

Restores of Fujitsu NetWorker Module for DB2 backups

NMDA can restore backups that are performed with Fujitsu NetWorker Module for DB2 version 2.1 or 4.0 on an operating system that this NMDA release supports. The *NetWorker E-LAB Navigator* at <https://elabnavigator.emc.com/eln/elhome> 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.

DB2 data recovery with the HPU utility

NMDA supports the IBM Optim High Performance Unload (HPU) utility, which is a high-speed stand-alone utility that unloads, extracts, and repartitions backup data into DB2 databases.

The HPU utility can unload data directly from an NMDA DB2 full, incremental, or delta backup or from a tablespace container. By reading the data directly from the backup image instead of through the DB2 software layer or database engine, the utility provides quick and efficient recovery of discrete volumes of data.

For example, if an application has a corrupted customer address, HPU can unload the customer ID and address from a table in a backup image and then load the extracted data into the database to correct the customer address.

To minimize the impact on a production system, you can use HPU to unload data for dropped or corrupted tables into a non-production system, where you can cleanse and prepare the data for load into the production system.

You can use the HPU utility for a data unload by creating a control file and running the `db2hpu` command with the control file. The utility unloads data from the backup image to staging files and then writes the data to output files for use with the DB2 `load` utility.

The IBM DB2 documentation provides details about HPU and how to create the control file for HPU operations.

DB2 load command with the copy yes option

NMDA supports the DB2 `load` command with the `copy yes` option.

The `copy yes` option specifies that a copy of the changes will be saved to a specified location and can be used in the database recovery. The rollforward operation loads the copy of the saved changes directly into the database.

The load-copy image can be stored in the same location as log files (recommended) or in any other location.

The NetWorker index indicates a load-copy image with the `LOAD_COPY` type in the description field. You can use the `nsrinfo` command to obtain information about the load copy data that is

stored in the NetWorker indexes. In the following example, the `nsrinfo` command returns the line that includes `description...LOAD_COPY`:

```
nsrinfo -v -s NW_Server -n db2 -X all NW_Client
```

```
scanning client 'NW_Client' for all savetimes from the db2 namespace on server
NW_Server
version=1, DB2, objectname=/SAMPLE/NODE0000 /DB_BACKUP.20141114144126.1,
createtime=Fri Nov 14 14:41:26 2014, copytype=3 BSACopyType_BACKUP,
copyId=1415994086.1415994087, restoreOrder=1415994086.1, objectsize=0.0,
resourcetype=database, objecttype=4 BSAObjectType_DATABASE, objectstatus=2
BSAObjectStatus_ACTIVE, description=NMDA_v82:DB2_v1051:LOAD_COPY:SAMPLE:TEQ,
objectinfo=db2inst1:1, NSR size=278900
```

A load operation with the `copy yes` option enables the database to be restored with a rollforward operation without any special handling. The IBM documentation provides complete details about the `load` command and its features and options.

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 release earlier than 11.70, 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 123 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 that are performed with Fujitsu NetWorker Module for Informix version 2.0 on an operating system that this NMDA release supports. The *NetWorker E-LAB Navigator* at <https://elabnavigator.emc.com/eln/elhome> 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 that are 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)
 **Note:** The `daos.cfg` and `daoscat.nsf` files are not backed up, which follows the IBM recommendation.

[Configuring Lotus DAOS backups](#) on page 127 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 37 describes which files are backed up and at which backup level.

By default during Lotus backups, NMDA follows directory links (with a `.dir` file name 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 124 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 208 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](#) 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 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 can run a transaction logs only backup with multiple concurrent streams.

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 127 describes how to configure backups of Domino with DAOS enabled.

[Performing Lotus data restore and recovery](#) on page 207 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 known good point that is 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 *NetWorker Administration Guide* describes restarted backups for a backup group.

Note: NMDA does not support Client Direct Lotus backups that are checkpoint restart enabled due to the dependence of the checkpoint restart feature on the storage node `nsrmmmd` process to ensure that data is written to a device.

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. The checkpoint option is disabled for browselist save sets.

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 130.

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 the nsrdocrc command](#) on page 218 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 220 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 171
- [Performing Lotus database recovery with NetWorker User for Lotus](#) on page 212


NMDA Lotus recovery wizard

NMDA supports the Lotus recovery wizard, which is a Java application that you run from NMC to configure and run a restore and recovery of Lotus Domino/Notes data that is backed up by NMDA. If the wizard supports the environment, use the wizard whenever possible to perform the Lotus restore and recovery operations.

You can run the recovery 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.

The NMDA Lotus recovery wizard can configure and run the following types of restore and recovery:


- Complete database recovery to the current time or a selected backup time.
- Point-in-time recovery of a whole database.
- Restore of selected files (not DAOS files), including a relocation of data if required.
- Restore of selected DAOS files or all missing DAOS files.
- Restore and recovery of Lotus Domino/Notes data (not DAOS files) to a different database on either the same host or a different host, by using backups of the original database.

 **Note:** Due to a Domino server limitation, DAOS data can be restored to its original path only, not to a different host or a different partitioned Domino server.

The recovery wizard enables you to start a restore or recovery immediately or schedule the operation to start at a future time.

The recovery wizard does not support the following features:

- Disaster recovery.
- Lotus document-level recovery.
- Recovery of a backup that used NSR_SAVESET_NAME to specify a base name other than NOTES: for the save set. The recovery wizard can recover only backups with a save set name that starts with NOTES:.

 **Note:** You must manually update the link file of a restored relocation linked database.

The following additional sources describe the wizard:

- Descriptive inline text in the wizard.
- Online help in the wizard.
- *NetWorker Module for Databases and Applications Release Notes*.

Partitioned Domino servers

NMDA supports the backup and recovery of partitioned Domino servers.

[Configuring partitioned Domino server backups](#) on page 126 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.
 - ① **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.

- ① **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 that were created since the last successful log backup.

- ① **Note:** NMDA always performs the MySQL binary log backups at the level full, even during an incremental backup or a cumulative backup. NMDA does not back up any binary logs during 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.
 - ① **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 that are performed with NMDA. You can use either the NMDA MySQL recovery wizard or the NMDA program `nsrmysqlrc` to perform the restore and recovery.

The following topics provide details about the two different methods of MySQL restore and recovery.

MySQL restore and recovery with the NMDA MySQL recovery wizard

NMDA supports the MySQL recovery wizard, which is a Java application that you run from NMC to configure and run a restore and recovery of MySQL data that is backed up by NMDA. If the wizard supports the environment, use the wizard whenever possible to perform the MySQL restore and recovery operations.

You can run the recovery 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.

The NMDA MySQL recovery wizard can configure and run the following operations:

- Recovery of a whole instance backup or partial backup, with or without the binary logs, to the current time or a point-in-time.
 - ⓘ **Note:** You can select to have the binary logs replayed to the data after the restore of data and binary logs is complete.
- Restore of one or more binary log backups.
- Redirected restore and recovery.
- For the restore of data, one of the following combinations of operations:
 - Extract—Extracts backup files from the backup image.
 - Extract and prepare—Extracts backup files from the backup image and processes the files to produce a prepared backup that is ready for restore.
 - Extract and prepare and copy back—Extracts backup files from the backup image, processes the files to produce a prepared backup, and copies the prepared backup to the MySQL data directory, for example, to complete a restore to the database server.

For the restore of data, you must select one of the three supported combinations of extract, prepare, and copy backup operations, as listed in the wizard. If you select to include the copy back operation, NMDA shuts down the MySQL instance before the recovery. You can optionally select to restart the instance after the restore.

NMDA restores all the data of a whole instance backup or partial backup. You cannot select to restore only part of a data backup, such as only one of three databases that are included in a partial backup.

The recovery wizard enables you to start a restore or recovery immediately or schedule the operation to start at a future time.

The recovery wizard does not support the following operations, which the `nsmysqlrc` program supports:

- List image operation
- Validate operation

The following additional sources describe the wizard:

- Descriptive inline text in the wizard
- Online help in the wizard
- *NetWorker Module for Databases and Applications Release Notes*

MySQL restore and recovery with the `nsmysqlrc` program

You can run the `nsmysqlrc` program to perform the following 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.
- 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 `nsmysqlrc` 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 `nsmysqlrc` 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 `nsmysqlrc` 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 `nsmysqlrc` program restores all three databases. You cannot use `nsmysqlrc` to restore only part of a backup, such as only one of the three databases that are 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 that is 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 before 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 and later. The *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 that is 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 133 describes how to configure the MySQL 5.6 features for NMDA operations.

[Performing MySQL restores of InnoDB tables outside data directory](#) on page 234 describes the MySQL 5.6 restore procedure to a new directory structure.

MySQL replicated slave server operations

With MEB 3.11 or later, NMDA supports the backup and restore of a MySQL replicated slave server. By default, MEB 3.11 or later backs up the relay logs that are associated with the replicated slave server, along with the binary logs.

To enable the NMDA backups of a MySQL replicated slave server, you must enable the MEB log backup mechanisms. MEB must manage the backup and restore of the binary logs and relay logs in the replicated environment.

You can configure the NMDA full and incremental backups as required of the MySQL replicated slave server, but you must disable the NMDA log backups. When the MEB log backup mechanisms are enabled, NMDA relies on MEB to back up and manage the binary logs and relay logs. NMDA does not store information about the available logs in the replicated slave server backups.

[Considerations for backups of MySQL replicated slave servers](#) on page 134 describes how to configure the NMDA backups of a MySQL replicated slave server.

After you restore an NMDA MySQL backup of a replicated slave server, you can restore the binary logs by using the `mysqlbinlog` program according to the MEB documentation. [Performing](#)

[MySQL recovery of a replicated slave server](#) on page 235 describes the NMDA restores of MySQL replicated slave server backups.

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 Oracle Recovery Manager (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](#) 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 that are 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 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 324 provides a sample script to back up the archived redo log files in an Oracle Real Application Clusters (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 that are 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.

[Oracle DBA and NetWorker Backup Administrator Collaboration](#) on page 333 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.

About this task

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 that is 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 138 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.

i **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 139 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.

About this task

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 143 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 as for Oracle RAC, including the Oracle database versions, operating system versions, and NetWorker versions. The *NetWorker E-LAB Navigator* at <https://elabnavigator.emc.com/eln/elhome> describes the environment support.

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 *NetWorker Module for Databases and Applications Installation Guide* and in this administration guide.

To perform backups and restores of Exadata, 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, 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 canceled two days ago, you can use the following command:

About this task

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

To determine if a datafile requires a backup, RMAN compares the time that is specified in this command with the completion time of the most recent backup of the datafile. 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 retention policy to specify how long data is available for recovery. The NetWorker policy is 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 6](#) on page 83.

NMDA no longer supports the Oracle parameter `NSR_RETENTION_DISABLED`. The following results occurred if you used the parameter in a prior NMDA release:

- Backups remained on the backup device forever unless the DBA ran the `delete obsolete` command on a regular basis with permissions to delete the backup entries.
- Clone operations failed.

If you set the `NSR_RETENTION_DISABLED` parameter in an RMAN backup script, the NMDA Oracle backup fails with an error message. Instead of using this parameter, either you must use the `NSR_ORACLE_RETENTION` parameter if you use Oracle window-based policies or you must use the NetWorker retention policies going forward.

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 between 'sysdate-45' and
'sysdate-30' ;
```

NMDA Oracle recovery wizard

NMDA supports the Oracle recovery wizard, which is a Java application that you run from NMC to configure and run an Oracle restore and recovery. You can use the wizard to create an RMAN script for the restore and recovery of Oracle data that is backed up by NMDA. If the wizard supports the environment, use the wizard whenever possible to perform the Oracle restore and recovery operations.

You can run the recovery 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.

The NMDA Oracle recovery wizard can configure and run the following types of restore and recovery:

- Current-time restore and recovery of a whole or partial Oracle database, where a partial database can include any combination of pluggable databases, tablespaces, and datafiles.
- 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 database.

The database duplication script (created by the wizard) runs on the destination host and 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 will 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 will be created on a remote host, the wizard generates the RMAN script and stores the script in the RAP database. The recover binary on the destination host runs the RMAN script. 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 wizard enables you to start a restore or recovery immediately, or schedule the operation to start at a future time.

The recovery wizard does not support the following features:

- Cluster or Oracle RAC systems.
- RMAN automatic channels.

The following additional sources describe the wizard:

- Descriptive inline text in the wizard
- Online help in the wizard

- *NetWorker Module for Databases and Applications Release Notes*

Restores of Fujitsu NetWorker Module for Oracle backups

NMDA can restore backups that are performed with Fujitsu NetWorker Module for Oracle version 5.0 on an operating system that this NMDA release supports. The *NetWorker E-LAB Navigator* at <https://elabnavigator.emc.com/eln/elhome> 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.
- i** **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 53 provides details.

i **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 *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 141 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 backup action details 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](#) 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 *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 6](#) on page 83.

To perform save set bundling, the `nsrdasv` program connects to the Oracle database by trying 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 141 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](#) and [Example 2](#) 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 that 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 53.

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

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

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 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.

About this task

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 that is 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 that is named 12983479182:

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

The *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 retention policies between all the dependent save sets in a scheduled backup cycle, whether 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 retention policy.

[Configuring policy uniformity for scheduled Oracle backups](#) on page 142 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 56
- Archival backup through the RMAN `backup...keep` command as described in [Archival backup feature](#) on page 56
- 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 that are 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:
 1. Allocate or configure one or more channels with the `sbt_tape` device type.
 2. Back up the files with one of the following RMAN commands:


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

Note: Whether 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 that is specified in the Inactivity Timeout field of the NetWorker backup action resource.

Multiplexing for Oracle11g multisection backups is supported with AFTDs only


Due to Oracle11g and NetWorker limitations, NetWorker multiplexing for multisection backups (a new feature with Oracle11g) is supported with AFTDs only, *not* with tape or regular file type devices.

If you have tape or regular file type devices, configure a separate NetWorker device for each RMAN channel to avoid multiplexing in NetWorker.

If you multiplex the Oracle11g multisection backups on a tape or regular file type device and then experience a hang at restore time, use a single RMAN channel to restore the multisection Oracle backups.

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 *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 that is 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:

- For a scheduled backup, set the pool selection criteria in the backup configuration on the NetWorker server.
 - For a manual backup, set the NSR_DATA_VOLUME_POOL parameter in the RMAN backup script.
3. Set the retention value by performing one of the following actions:
- For a scheduled backup, set the retention period in the backup action of the policy-based backup configuration.
 - For a manual backup, 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 that is used with NMDA.

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

NMDA features specific to SAP IQ

NMDA supports specific features of SAP IQ backup and restore operations.


Files that are backed up during SAP IQ backups

Starting with NMDA 9.2, the NMDA software supports the backup of an entire SAP IQ database by default. When you configure a selective backup, NMDA backs up either all the read-write database files or only the specified read-only data objects from the SAP IQ database.

Starting with NMDA 19.1, the NMDA software supports the backup of all the SAP IQ transaction logs and log archives. After you perform the log backups, you can use NMDA to perform an SAP IQ point-in-time recovery that restores the database to a transaction log offset.

An SAP IQ database is a relational database that is optimized for use as a data warehouse, and consists of stores, log files, and server space. The database includes five types of SAP IQ stores:

- **Catalog store**—Contains metadata and information for managing the database, in the file `dbname.db`.
- **IQ main store**—Contains the database data and transaction log.
- **IQ temporary store**—Contains temporary data that is used in loading or querying the database.
- **RLV store**—Contains in-memory data.
- **IQ shared temporary store**—Contains temporary data for a multiplex environment.

 **Note:** NMDA supports only simplex environments, not multiplex environments.

These SAP IQ stores consist of various types of dbspaces and dbfiles:

- A **dbspace** is a logical name for a container of files or raw partitions called **dbfiles**.
- A **dbfile** is an operating system file or a raw partition contained within a **dbspace**.

The SAP IQ documentation provides details about SAP IQ stores, dbspaces, and dbfiles.

The following topics provide more details about which SAP IQ database objects and files are backed up by the NMDA software and at which backup level.

SAP IQ full, incremental, incremental since full, and transaction log backups

NMDA supports four backup levels for the backups of SAP IQ database data: full, incremental, incremental since full, and transaction log.

You must specify the appropriate backup level for the SAP IQ database data that you want to back up:

- **Full**—A full backup sends the command `backup database full` to SAP IQ, which backs up the entire database, including both the database data and transaction logs.
i **Note:** NMDA also supports the separate backups of the transaction logs.
- **Incremental**—An incremental backup sends the command `backup database incremental` to SAP IQ, which backs up the database data blocks that have changed since the last backup of any type.
- **Incremental since full**—An incremental since full backup (or cumulative backup) sends the command `backup database incremental since full` to SAP IQ, which backs up the database data blocks that have changed since the last full database backup.
- **Transaction log**—A transaction log backup sends the command `backup database point in time recovery logs only to 'log_backup_path/archive_file_name'` to SAP IQ. The log backup backs up all the SAP IQ transaction logs and archive logs.

SAP IQ selective and all-inclusive backups

NMDA supports two types of selective backups of SAP IQ data, which back up either read-write database files or read-only data objects from the SAP IQ database. NMDA also supports an all-inclusive backup.

An SAP IQ database includes the five types of SAP IQ stores that consist of various types of dbspaces and dbfiles. In addition to backing up the entire SAP IQ database, you can optionally specify a read-write or a read-only selective backup or an all-inclusive backup:

- **Read-write selective backup**—Backs up the complete set of all the read-write database files in the database, without any read-only dbspaces or dbfiles.
 A read-write selective backup supports all the backup levels. You can perform a full, incremental, or incremental since full backup of the read-write database files.
- **Read-only selective backup**—Backs up only the specified read-only dbspaces or read-only dbfiles or both, without any read-write database files.
 A read-only selective backup supports only the full backup level. You can perform only a full backup of the specified read-only database objects.
- **All-inclusive backup**—Backs up all the read-write database files and read-only data objects.
 An all-inclusive backup supports all the backup levels. You can perform a full, incremental, or incremental since full backup of all the read-write and read-only data objects.

SAP IQ transaction log backups

NMDA supports the backup of all the SAP IQ transaction logs and log archives as NetWorker save sets. When you run the SAP IQ transaction log backups, the NetWorker software uses a file system backup to back up the logs and log archives.

The SAP IQ transaction log backups support the following features:

- Backup of multiple log paths.
- Movement of the backed-up logs from the log location to a user-specified local archive directory.

- Retention of the backed-up logs in the local archive for a user-specified period of time.
 - ① **Note:** You can set the retention time to a value of 0 or more days. When the value is 0, the log files are retained until they are manually removed.
- Printing of the list of backed-up log files to the console.
- Logging of the retention information.

① **Note:** The SAP IQ transaction log backups do not check for any log gap or missing log files.

To enable the SAP IQ transaction log backups, you must set the mandatory parameter `NSR_BACKUP_LEVEL` to the value `txnlog` in the NMDA SAP IQ backup configuration file.

To enable the supported features of the transaction log backups, you can set other specific parameters in the NMDA configuration file, as described in [Configuring the SAP IQ transaction log backups](#) on page 145.

You can launch an SAP IQ transaction log backup manually by running the `nsrdsav -z configuration_file_path` on the NMDA client system. You can also start a log backup from the NetWorker Management Console (NMC). The log backups support scheduled backups and cloning. To enable a scheduled log backup, you must set the backup level to Logs Only in the backup action that you assign to the backup workflow in the NMC data protection policy, as described in [Configuring the backup action](#) on page 100.

An SAP IQ transaction log backup includes the following operations.

1. The NetWorker server issues a log archival SQL command to the SAP IQ server.
2. The latest transactions are recorded in the log by truncating and archiving the current transaction log and starting a new log.
3. All the log files and archive files whose modified times are earlier than the modified time of the current active log file are marked for backup.
4. When `IQ_PITR_LOG_RETENTION_PATH` is set to a valid directory pathname and `IQ_PITR_LOG_RETENTION_TIME` is set to a non-zero value, any log files that have been retained in the directory longer than the retention time are removed.
5. The NetWorker server uses a `save` command to perform a file system backup of the current log directory and all the other log paths specified in the `IQ_PIT_RESTORE_LOG_PATH` setting. Each respective save set is backed up as a transaction log save set.
6. After the backup, if `IQ_PITR_LOG_RETENTION_PATH` is set to a valid directory pathname, all the backed-up log files are moved to the specified location. Otherwise, the log files remain in the original location and are backed up again in subsequent log backups.

① **Note:** When `IQ_PITR_LOG_RETENTION_PATH` is set and the log files are moved to the specified location, the files are not backed up again in the next log backup. This setting prevents multiple recoveries of the log files and overwriting of the log files during recovery operations.

SAP IQ restores

NMDA supports the restore of SAP IQ database backups, the restore of backups of specific read-only dbspaces or dbfiles, and the restore of backups of all the read-write database files. Starting with NMDA 19.1, the NMDA software supports a point-in-time restore and recovery that uses the transaction log backups.

You can perform the following types of NMDA SAP IQ restore operations:

- Restore of the latest full all-inclusive backup to a specified backup time.
- Restore of the latest backup of a specific read-only dbspace to a specified backup time.

- Restore of the latest backup of a specific read-only dbfile to a specified backup time.
- Restore of an incremental backup or incremental since full backup to the most recent state of the database.
- Restore of the latest backup of all the read-write database files.
- Point-in-time restore and recovery of transaction log backups.

SAP IQ point-in-time recovery using log backups

NMDA supports an SAP IQ point-in-time recovery that uses the transaction log backups to restore the database to a transaction log offset. The recovery operation restores all the committed database operations prior to the specified point-in-time. The recovery first restores the database from the data backups and then restores the logs for the time period between the last full data backup and the specified point-in-time.

An SAP IQ point-in-time recovery supports the following features:

- Point-in-time recovery using a transaction log offset, not a timestamp.
 - ⓘ **Note:** The offset can be either between the previous full backup and the specified transaction log backup, or between the previous full backup and the transaction log backup just before that full backup.
- Point-in-time recovery when only one active log path is available.
- Point-in-time recovery when multiple log paths are available, after the active log path has been changed multiple times.
- Redirection of the log files to a temporary alternate directory during the point-in-time recovery.
- Clean-up of the restored logs after a successful recovery.
 - ⓘ **Note:** After a failed recovery, the log files are not cleaned up. The log file cleanup is skipped for both failed redirected recovery and failed regular recovery.
- Printing of the list of restored log files to the console.

An SAP IQ point-in-time recovery does not support specific features:

- Restore to an offset between a full backup and a cumulative or incremental backup.
- Use of the RENAME option with point-in-time recovery.
- Point-in-time recovery on a multiplex system.

To enable the SAP IQ point-in-time recovery, you must set the mandatory parameters in the NMDA SAP IQ restore configuration file:

- `IQ_PIT_RESTORE_ENABLE` must be set to `TRUE`.
- `IQ_PIT_RESTORE_LOG_PATH` must specify the pathnames of all the log files to be restored, including the active log of the database at the time to which the restore will be performed.
- `IQ_PIT_RESTORE_OFFSET` must specify the transaction log offset to use for the point-in-time restore.
- `NSR_RECOVER_TIME` must specify the transaction log save set time to which the backup data and logs will be restored.

To enable the supported features of the point-in-time recovery, you can set other specific parameters in the NMDA configuration file, as described in [Performing the SAP IQ data restore with the `nsriqrc` command](#) on page 245.

You can start an SAP IQ point-in-time recovery by running the `nsriqrc -z configuration_file_path` on the NMDA client system.

- Note:** After a successful transaction log recovery, it is recommended that you perform a full backup to maintain consistency in the logs.

NMDA features specific to Sybase

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

Sybase full, cumulative, and incremental backups

NMDA supports full, cumulative, and incremental backups of Sybase data.

These Sybase backups use specific commands to back up different types of 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.

- A cumulative Sybase backup sends the `dump database...cumulative` command to Sybase, which backs up the database data and logs that have changed since the last full database backup.

NMDA supports the cumulative backup support with Sybase ASE 15.7 SP 100 or later. [Sybase ASE 15.7 features](#) on page 64 provides more details.

- 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:
 - If `NSR_PROMOTE_FULL` is `TRUE`, the Sybase software promotes the backup to a full backup when the databases do not support incremental backups. For example, when the database data and logs are on the same device, the databases do not support incremental backups. After the backup, the Sybase software truncates the inactive portion of the database log.
 - If `NSR_PROMOTE_FULL` is `FALSE`, the Sybase software does not promote an incremental backup to a full backup if the incremental backup cannot be performed.

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

Sybase restore parameters

You set the Sybase restore parameters in the NMDA configuration file.

NMDA provides separate configuration file templates for Sybase backup and restore parameters, as described in [NMDA configuration file](#) on page 400. It is recommended that you maintain separate configuration files for the Sybase backup and restore parameters.

After setting the required restore parameters in the configuration file, you can perform a Sybase restore/recovery with the `nsrsybrc -z configuration_file_path` command. [NMDA Sybase parameters](#) on page 464 provides complete details about the Sybase restore parameters.

Restore of Sybase transaction logs only

NMDA supports the restore of one or more Sybase transaction log backups. You can update a database by restoring only the transaction logs. After initially recovering the database to a point-in-time, you can restore the transaction log backups to update the database to specific points-in-time.


For the restore of Sybase transaction logs only, you must keep the database offline during the entire restore process. You must also restore the transaction log backups in the correct order, to ensure that the database is updated properly without any missing data.

[Example 33](#) provides details about the restore of Sybase transaction logs only.

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](#).

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 same `NSR_ASE_PASSWORD` parameter in the NMDA configuration file.

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 175 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 148 and [Specifying verification of Sybase database restores](#) on page 251 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 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 [Restoring to different database types](#) on page 260.

Sybase documentation describes the supported types of databases.


Sybase ASE instance backups and restores with a large number of databases

NMDA supports a Sybase ASE instance backup and restore that includes more than 200 databases. NMDA supports the backup and restore of a Sybase ASE instance that contains many databases, up to the maximum number of databases that Sybase supports.

The `NSR_CONCURRENCY_MODE` parameter setting and the parallelism setting determine the method that NMDA uses to perform the instance backup and restore. The parallelism setting is the minimum of the `NSR_PARALLELISM`, client parallelism, and server parallelism settings.

[Best practices for Sybase backups](#) on page 146 provides more information for configuring the NMDA backups of a Sybase ASE instance that contains a large number of databases.

For example, based on the `NSR_CONCURRENCY_MODE` setting, NMDA restores either one database at a time or multiple databases concurrently. The parallelism setting determines the maximum number of databases that NMDA restores concurrently.

 **Note:** If the parallelism setting is too low for a single database to be restored, NMDA restores only one database at a time.

[Performing restores of Sybase ASE backups with a large number of databases](#) on page 254 describes the NMDA restores of the Sybase ASE instance backups with many databases.

Sybase ASE 15.7 features

NMDA supports the features of Sybase ASE release 15.7, including Sybase database compression, Sybase ASE Cluster Edition systems, concurrent full and incremental backups, cumulative backups, 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 152 describes how to configure Sybase 15.7 features for NMDA operations.

[Obtaining backup information with the `listonly` option](#) on page 252 describes how to use the `listonly` option with Sybase ASE 15.7 ESD#2 or later. You can use the `listonly` option to obtain structure information from the latest backup of a specified Sybase instance or databases, without restoring the backup.

The following topics provide information about the NMDA support of specific features.

Sybase concurrent full and incremental backups

NMDA supports concurrent full and incremental backups with Sybase ASE 15.7 ESD#2 or later. A full backup is a database backup, and an incremental backup is a transaction log backup. A full database backup can take a long time to complete. Running full and incremental backups concurrently can reduce the risk of losing large database updates due to a failure during the full backup.

During a restore, the concurrent incremental backup can only be applied to the previous full backup, not to the full backup that ran concurrently with it. The concurrent incremental backup does not truncate the transaction log.

Sybase ASE not support the following types of concurrent backups:

- Two full backups running concurrently
- Two incremental backups running concurrently
- Full and cumulative backups running concurrently
- Full and incremental backups running concurrently if the incremental backup started first

The Sybase documentation provides more details about concurrent backups.

Sybase cumulative backups and restores

NMDA supports cumulative backups with Sybase ASE 15.7 SP 100 or later. A cumulative backup makes a copy of all the pages that have changed in the database since the last full database backup. Sybase 15.7 SP 100 also supports the cumulative backups of low-durability databases, which do not support log backups. Recovery times can be reduced when cumulative backups are restored instead of log backups.

If a valid full backup of the database does not yet exist, then the initial cumulative backup of the database is promoted to a full backup after Sybase ASE generates an error message:

```
You cannot execute an incremental dump before a full database dump.
```

Note: To enable the NMDA Sybase cumulative backups and restores of a database, ensure that you use only NMDA Sybase to perform all the backups of the database. Do not use other software to back up the database.

You can perform a cumulative backup of any database except a master database or temporary database. Cumulative backups are not supported with Sybase ASE Cluster Edition.

The Sybase documentation provides more details about cumulative backups and restores.

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 176
- [Performing Sybase data restores with NetWorker User for Sybase](#) on page 262

NMDA Sybase recovery wizard

NMDA supports the Sybase recovery wizard, which is a Java application that you run from NMC to configure and run a restore and recovery of Sybase ASE data that is backed up by NMDA. If the wizard supports the environment, use the wizard whenever possible to perform the Sybase ASE restore and recovery operations.

You can run the recovery 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.

The NMDA Sybase recovery wizard can configure and run the following types of restore and recovery:

- Complete database recovery to the current time or selected backup time.
- Point-in-time recovery of a database.
- Restore and recovery of Sybase ASE data to a different database on either the same host or a different host, by using backups of the original target database.

The recovery wizard enables you to start a restore or recovery immediately or schedule the operation to start at a future time.

The recovery wizard does not support the following features:

- Sybase ASE Cluster Edition.
- Restore/recovery when you use the 32-bit NMDA Sybase binaries from the 64-bit NMDA package.

The following additional sources describe the wizard:

- Descriptive inline text in the wizard
- Online help in the wizard
- *NetWorker Module for Databases and Applications Release Notes*

NMDA features specific to Orchestrated Application Protection

NMDA 18.1 introduced the Orchestrated Application Protection feature that leverages the Data Domain BoostFS package to back up and restore the specific database types to and from a DD Boost device. The DD Boost device is configured as a NetWorker backup device.

Orchestrated Application Protection provides database backups through the NMDA program `nsroappbackup`, which leverages the native backup utilities that the database software provides. Depending on the specified backup level, you must set the Orchestrated Application Protection backup parameters in a particular section of the NMDA configuration file.

Orchestrated Application Protection provides restore operations through the NMDA program `nsroapprecover`, which receives the backed-up data through the `nsroappbackup` program and stores the data on the local disk. Then you can use the native restore utilities that the database software provides to restore and recover the database.

The current NMDA release supports the Orchestrated Application Protection feature for backups and restores of MongoDB, MySQL, and PostgreSQL database data.

The Orchestrated Application Protection feature includes the following limitations:

- Only a DD Boost device is supported as the backup device, where the data is backed up and from which the backups are restored.
- To recover cloned data that does not reside on a DD Boost device, you must clone the data to a DD Boost device and then run the recovery from the DD Boost device.
- Only nonsnapshot backups and restores are supported to and from the DD Boost device.
- The configuration wizard is not supported for the backup or restore configuration.

[NMDA features specific to MongoDB](#) on page 67 provides more information about configuring and running the MongoDB backups and restores through Orchestrated Application Protection.

[NMDA features specific to MySQL](#) on page 67 provides more information about configuring and running the MySQL backups and restores through Orchestrated Application Protection.

[NMDA features specific to PostgreSQL](#) on page 67 provides more information about configuring and running the PostgreSQL backups and restores through Orchestrated Application Protection.

NMDA features specific to MongoDB

NMDA supports the backups and restores of MongoDB database data to and from a DD Boost device through the Orchestrated Application Protection feature.

The NMDA program `nsroappbackup` works with the MongoDB utility `mongodump` to perform the database full backups. MongoDB does not support incremental database backups or transaction log backups. After you complete the required backup configuration, you can perform manual or scheduled NMDA backups of the MongoDB data.

The NMDA program `nsroapprecover` retrieves the MongoDB data backup into a user-specified directory. The directory does not initially exist on the file system, and the program creates the directory. After you run the `nsroapprecover` program, you must run the required MongoDB restore utility to complete the restore of the data backup.

[NMDA features specific to Orchestrated Application Protection](#) on page 66 provides more details about Orchestrated Application Protection.

NMDA features specific to MySQL

NMDA supports the backups and restores of MySQL database data to and from a DD Boost device through the Orchestrated Application Protection feature.

The NMDA program `nsroappbackup` works with the MySQL utility `mysqldump` to perform the database full backups. MySQL supports the transaction log (binary log), and the transaction log can be backed up separately. After you complete the required backup configuration, you can perform manual or scheduled NMDA backups of the MySQL data.

The NMDA program `nsroapprecover` retrieves the MySQL data or log backup into a user-specified directory. The directory does not initially exist on the file system, and the program creates the directory. After you run the `nsroapprecover` program to retrieve the required data and log files, you must run the required MySQL restore utility to complete the restore or recovery of the data backup.

[NMDA features specific to Orchestrated Application Protection](#) on page 66 provides more details about Orchestrated Application Protection.

NMDA features specific to PostgreSQL

NMDA supports the backups and restores of PostgreSQL database data to and from a DD Boost device through the Orchestrated Application Protection feature.

The NMDA program `nsroappbackup` works with the PostgreSQL utilities `pg_dump` and `pg_dumpall` to perform the database full backups without WAL (write ahead log) segment file archiving. The `nsroappbackup` program requires the `pg_basebackup` utility to perform the database full backups if the WAL segment file archiving is enabled. The `nsroappbackup` program also performs the transaction log backups. After you complete the required backup configuration, you can perform manual or scheduled NMDA backups of the PostgreSQL data.

The NMDA program `nsroapprecover` retrieves the PostgreSQL data backup into a user-specified directory. The directory does not initially exist on the file system, and the program creates the directory. The `nsroapprecover` program also can retrieve the archived WAL segment file to the required location. After you run the `nsroapprecover` program, you must run the required PostgreSQL restore utility to complete the restore of the data backup.

[NMDA features specific to Orchestrated Application Protection](#) on page 66 provides more details about Orchestrated Application Protection.

PostgreSQL full and transaction log backups

NMDA supports two backup levels for the PostgreSQL database backups: full backups and transaction log (WAL segment file archiving) backups. The NMDA program `nsroappbackup` performs these backups, which are enabled through the Orchestrated Application Protection configuration.

You must specify the appropriate backup level for the PostgreSQL database data that you want to back up:

- Full—A full backup in a single-instance environment must use the required PostgreSQL utility:
 - Use the `pg_dump` or `pg_dumpall` utility to back up a database when the WAL segment file archiving is disabled.
 - Use the `pg_basebackup` utility to back up a database when the WAL segment file archiving is enabled.



Note:

When the WAL segment file archiving is enabled, the full backup must leverage the `pg_basebackup` utility as required by PostgreSQL.

NMDA must use the recommended PostgreSQL methods to back up the transaction logs (WAL segment files). Simply using a copy command in the NMDA Orchestrated Application Protection backup script file or other methods for the transaction log backup can cause issues during the restore.

- Transaction log—A transaction log (WAL segment file) backup runs through the `nsroappbackup` program to the NetWorker server.



Note: The PostgreSQL WAL segment works with the PostgreSQL transaction. Only the operations inside the transactions can be logged into the WAL segment files. The WAL archiving restore requires the transaction information inside the WAL segment file to find the time point for the point-in-time restore. Without the transaction information, the WAL archiving restore and the point-in-time restore will not complete properly.

To enable the WAL segment file archiving, you must register the `nsroappbackup` program with its required command line options through the `archive_command` setting in the `postgresql.conf` file. You must specify the backup level with the `-l txnlog` command line option in the `archive_command` setting.

[PostgreSQL backup considerations](#) on page 161 provides details about how to configure and perform the NMDA PostgreSQL backup operations.

PostgreSQL restores

The NMDA program `nsroapprecover` program retrieves the PostgreSQL data backup into an empty user-specified directory. After you run the `nsroapprecover` program, you must run the required PostgreSQL restore utility to complete the recovery of the data backup.



Note: The PostgreSQL WAL segment works with the PostgreSQL transaction. Only the operations inside the transactions can be logged into the WAL segment files. The WAL archiving restore requires the transaction information inside the WAL segment file to find the time point for the point-in-time restore. Without the transaction information, the WAL archiving restore and the point-in-time restore will not complete properly.

To enable the archived WAL segment file restore, you must register the `nsroapprecover` program with its required command line options through the `restore_command` setting in the `recovery.conf` file.

You must also set the required Orchestrated Application Protection restore parameters in a restore section of the NMDA configuration file.

[Performing the restore of PostgreSQL data](#) on page 268 provides details about how to configure and perform the NMDA PostgreSQL restore operations.

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 3 NMDA components

Component name	Description
<code>nmda_application.cfg</code>	<p>Located in the directory <code>/nsr/apps/config</code> (UNIX) or <code>NetWorker_install_path\apps\config</code> (Windows). NMDA configuration file template for backup and restore parameters that apply to client-side configurations only. NMDA provides a single configuration file template for both backup and restore parameters, except for the following applications:</p> <ul style="list-style-type: none"> • For MySQL, NMDA provides both a backup configuration file template, <code>nmda_mysql_backup.cfg</code>, and a restore configuration file template, <code>nmda_mysql_restore.cfg</code>. • For SAP IQ, NMDA provides both a backup configuration file template, <code>nmda_iq_backup.cfg</code>, and a restore configuration file template, <code>nmda_iq_restore.cfg</code>. • For Sybase, NMDA provides both a backup configuration file template, <code>nmda_sybase_backup.cfg</code>, and a restore configuration file template, <code>nmda_sybase_restore.cfg</code>. <p>For Orchestrated Application Protection, NMDA provides the configuration file template, <code>nmda_oapp.cfg</code>, that can be used for both backup and restore configurations.</p>
<code>nsrdaadmin(.exe)</code>	Program that encrypts the user password for DB2, MySQL, Orchestrated Application Protection, SAP IQ, and Sybase backups.
<code>nsrdaprobe(.exe)</code>	Program that probes for the number or size of generated logs as a condition that triggers a probe-based backup.
<code>nsrdasv(.exe)</code>	<p>Main NMDA program that performs one of the following backups:</p> <ul style="list-style-type: none"> • 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 an SAP IQ server. • A manual backup of a Sybase server.
<code>nsrapprapplicationra(.exe)</code>	Program that performs operations for the NMDA configuration wizard on a DB2, Informix, Lotus, MySQL, Oracle, or Sybase host.
NMDA DB2 components:	
<code>libnsrdb2.so</code>	Located in the directory <code>/usr/lib</code> 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.

Table 3 NMDA components (continued)

Component name	Description
<code>nsrdb2cat (.exe)</code>	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.
<code>nsrdb2rlog (.exe)</code>	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:	
<code>libnsrifmx.xx libxbsa.dll</code>	Located in the directory <code>/usr/lib</code> 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:	
<code>nsrdoc1b.dll</code>	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.
<code>nsrdocrc (.exe)</code>	Program for document-level recovery of Lotus data.
<code>nsrlotus_remrecov (.bat)</code>	Script that enables NMDA to perform a directed recovery from the NetWorker User for Lotus program on Windows.
<code>nsrnotesrc (.exe)</code>	Program for recovery of Lotus database files.
<code>nwbml.dll nwbml.exe</code>	On Windows systems only. Library and program for the NetWorker User for Lotus.
NMDA MySQL components:	
<code>libmysqlapi.so.version</code>	Located in the directory <code>/usr/lib</code> on Linux only. NMDA API library that interacts with the MySQL library <code>libmysqlclient.so</code> for communicating with the MySQL server.
<code>libmysqlapiwrap.so</code>	Located in the directory <code>/usr/lib</code> on Linux only. NetWorker API library that interacts with the NMDA API library <code>libmysqlapi.so</code> .
<code>libnsrmysql.so</code> <code>nsrsbtmysql</code>	The <code>libnsrmysql.so</code> file, which is located in the directory <code>/usr/lib</code> on Linux only, is the main NMDA library that is loaded by the MySQL MEB backup or restore thread. The <code>libnsrmysql.so</code> library starts and uses <code>nsrsbtmysql</code> to perform any corresponding NetWorker operations.
<code>nsrmysqlrc</code>	Program for recovery of MySQL data.
NMDA Oracle components:	
<code>libnsrora.xx</code>	Located in the directory <code>/usr/lib</code> on UNIX only. Library (known as Media Management Library in Oracle documentation) that is loaded by an Oracle backup or restore process.
<code>nsroradrpostcmd (.bat)</code>	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.
<code>nsroraadmin (.exe)</code>	Program that is used to create resource settings in the NWORA resource file for an Oracle backup configuration.
<code>nsroraclecat (.exe)</code>	Only available on platforms that support snapshot backups. Program that is used to remove RMAN catalog entries during automatic catalog synchronization for snapshot backups of Oracle data.

Table 3 NMDA components (continued)

Component name	Description
<code>nsrorainfo(.exe)</code>	Program that determines the NetWorker volumes that are required to restore specified Oracle backup pieces from NMDA backups.
<code>nsrsbtcn.exe orasbt.dll</code>	On Windows systems only. The <code>orasbt.dll</code> file is the main NMDA library (known as Media Management Library in Oracle documentation) loaded by the Oracle backup or restore thread. The <code>orasbt.dll</code> library uses <code>nsrsbtcn.exe</code> to perform any corresponding NetWorker operations.
NMDA SAP IQ components:	
<code>libnsriqbr.so</code>	Located in the directory <code>/usr/lib</code> on Linux only. NMDA library that interacts with the SAP IQ backup and restore utilities and the NetWorker server to perform backup and restore processes for SAP IQ data.
<code>nsriqrc</code>	Program for restore of SAP IQ databases.
NMDA Sybase components:	
<code>libnsrsyb.xx</code>	Located in the directory <code>/usr/lib</code> 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.
<code>nsrsybcc(.exe)</code>	Program for database consistency check of Sybase databases.
<code>nsrsybrc(.exe)</code>	Program for recovery of Sybase databases.
<code>nwbms.dll</code> <code>nwbms.exe</code>	On Windows systems only. Library and program for the NetWorker User for Sybase.
NMDA Orchestrated Application Protection components:	
<code>nmda_oapp_mongodb.example</code>	Located in the directory <code>/nsr/apps/config</code> on Linux only. An example file to assist with MongoDB operations. The file includes examples of a backup shell script and NMDA configuration file.
<code>nmda_oapp_mysql.example</code>	Located in the directory <code>/nsr/apps/config</code> on Linux only. An example file to assist with MySQL operations. The file includes examples of a backup shell script and NMDA configuration file.
<code>nmda_oapp_postgresql.example</code>	Located in the directory <code>/nsr/apps/config</code> on Linux only. An example file to assist with PostgreSQL operations. The file includes examples of a backup shell script, NMDA configuration file, <code>archive_command</code> settings for the <code>postgresql.conf</code> file, and <code>restore_command</code> settings for the <code>recovery.conf</code> file.
<code>nsroappbackup</code>	Located in the directory <code>/usr/sbin</code> on Linux only. Program for backups that are performed through Orchestrated Application Protection.
<code>nsroapprecover</code>	Located in the directory <code>/usr/sbin</code> on Linux only. Program for restores that are performed through Orchestrated Application Protection.

NMDA backup and restore processes

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

[Snapshot Backups and Restores](#) on page 345 describes the processes in snapshot backups and restores.

Regular scheduled backup processes

A regular scheduled backup includes the following process interactions.

Procedure

1. At the scheduled backup start time, the main NetWorker server-side service, `nsrd`, starts the configured group's backup by running the `nsrworkflow` program, which runs 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, SAP IQ, 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
 - SAP IQ Open Client/Server API
 - Sybase Open Client/Server API
 - b. Each backup session that is 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 runs another `nsrdasv` process that is 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 that are started depends on the number of files for backup and the specified parallelism.
4. The NMDA shared library (for DB2, Informix, MySQL, Oracle, SAP IQ, 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, `nsrmmmd`, to store on the appropriate backup volumes.
5. The NetWorker online indexes store the tracking information:

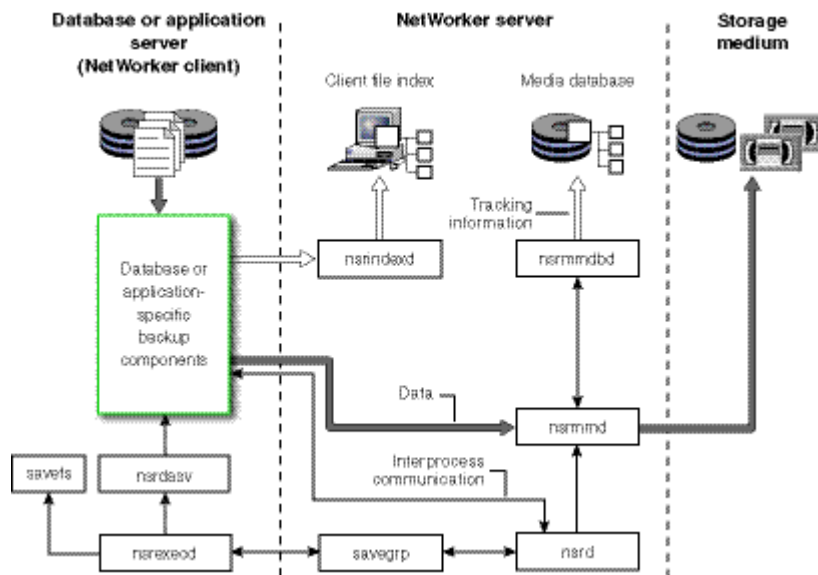
- The `nsrmmmd` service records tracking information in the NetWorker media database by using the `nsrmmdbd` service.
- The backup session sends tracking information to the NetWorker client file index by using the `nsrindexd` service.

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

Results

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.

DB2, Informix, MySQL, Oracle, SAP IQ, or Sybase restore processes

Procedure

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 `nsriqrc` command
 - 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 that are 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 that are 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.

Lotus restore processes

Procedure

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. Starts child `nsrnotesrc` processes to recover Domino data. The number of processes that are started 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 that are 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	76
• Configuring internationalization (I18N) support	77
• Verifying the database or application server configuration	81
• Verifying the basic NetWorker resource configurations	81
• Configuring scheduled backups with the wizard	86
• Configuring scheduled backups without the wizard	89
• Configuring manual backups	103
• Configuring Data Domain backups	105
• Configuring CloudBoost backups	109
• Configuring backups to AFTD storage	110
• Configuring probe-based backups	111
• Configuring parallel backups	118
• DB2 backup considerations	121
• Informix backup considerations	123
• Lotus backup considerations	124
• MySQL backup considerations	130
• Oracle backup considerations	134
• SAP IQ backup considerations	144
• Sybase backup considerations	146
• Orchestrated Application Protection backup considerations	153

Configuring NMDA backups

About this task

NOTICE Before configuring NMDA backups, determine the files to back up to prepare the system for disaster recovery. [Disaster Recovery](#) on page 273 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 336.

Procedure

1. If you have an internationalization (I18N) environment, configure I18N support according to [Configuring internationalization \(I18N\) support](#) on page 77.
2. Verify the configuration of the database server or application server according to [Verifying the database or application server configuration](#) on page 81.
3. Verify the basic NetWorker resource configurations according to [Verifying the basic NetWorker resource configurations](#) on page 81.
4. Configure the backup by using the appropriate method:
 - [Configuring scheduled backups with the wizard](#) on page 86 (server-side configuration)
 - [Configuring scheduled backups without the wizard](#) on page 89 (client-side configuration)
 - [Configuring manual backups](#) on page 103

Note:

A backup configuration that is 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 that is 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.

5. Complete any additional procedures for a specific type of configuration:
 - [Configuring Data Domain backups](#) on page 105
 - [Configuring backups to AFTD storage](#) on page 110
 - [Configuring probe-based backups](#) on page 111
 - [Configuring parallel backups](#) on page 118
6. Ensure that you meet any application-specific requirements:

- [DB2 backup considerations](#) on page 121
- [Informix backup considerations](#) on page 123
- [Lotus backup considerations](#) on page 124
- [MySQL backup considerations](#) on page 130
- [Oracle backup considerations](#) on page 134
- [SAP IQ backup considerations](#) on page 144
- [Sybase backup considerations](#) on page 146
- [Orchestrated Application Protection backup considerations](#) on page 153
 - [MongoDB backup considerations](#) on page 156
 - [MySQL backup considerations](#) on page 158
 - [PostgreSQL backup considerations](#) on page 161

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.

Note: Orchestrated Application Protection is not qualified with any specific I18N settings. However, for the database that uses the Orchestrated Application Protection feature, you can set the required I18N environment variables in the NSR_ENV_LIST parameter setting or in the backup shell script as specified by NSR_BACKUP_SCRIPT.

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`, `nsriqrc`, `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, MySQL, SAP IQ, or Sybase, you must use ASCII characters only in the password or the operation fails.
- REMOTE_RECOVCMD environment variable that is used for Lotus directed recovery, as described in [Configuring Lotus directed recovery](#) on page 215.
- Values for the parameters that are listed in the following table. [NMDA Parameters and Configuration File](#) on page 399 describes the parameters.

Table 4 NMDA parameters that support non-ASCII values

Database or application	Parameters that support non-ASCII values
All databases and applications	<ul style="list-style-type: none"> ▪ NSR_DIAGNOSTIC_DEST ▪ PRECMD ▪ POSTCMD
DB2	<ul style="list-style-type: none"> ▪ NSR_DR_FILE_LIST
Informix	<ul style="list-style-type: none"> ▪ INFORMIXSQLHOSTS (Windows only) ▪ NSR_DR_FILE_LIST ▪ ONCONFIG (Windows only)
Lotus	<ul style="list-style-type: none"> ▪ 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	<ul style="list-style-type: none"> ▪ 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	<ul style="list-style-type: none"> ▪ NSR_ORACLECAT_LOG_FILE ▪ NSR_RMAN_ARGUMENTS
SAP IQ	<ul style="list-style-type: none"> ▪ PATH
Sybase	<ul style="list-style-type: none"> ▪ 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 *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, which is 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 *NetWorker E-LAB Navigator* 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`).

About this task

To configure I18N support for Informix operations in a non-English locale, set the following environment variables in the NMDA Informix backup configuration wizard for scheduled backups. For manual backups and restores, set the environment variables manually in the environment:

- `DB_LOCALE`
- `DBLANG`
- `SERVER_LOCALE`
- `CLIENT_LOCALE`


The Informix documentation provides more details about these variables.

Configuring I18N support for Oracle operations

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

About this task

- The `format` string of an RMAN backup command
- 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 the following steps.

Procedure

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 12.1.0.2.0 - Production on Thu Dec 8 16:26:00
2016
Copyright (c) 1982, 2014, 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 SAP IQ operations

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

About this task

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

Configuring I18N support for Sybase operations

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

About this task

- 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 `nsrdsav` 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, MongoDB, MySQL, Oracle, PostgreSQL, SAP IQ, or Sybase ASE server documentation:

About this task

- 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. The following topics provide more information about these NetWorker resources.

About this task

The following storage resources must be configured for backup operations:

- Storage node that will own the backup storage devices.
- Backup storage devices.
- Label template for labeling the backup volumes.
- Media pools for sorting and storing the backup data.

Ensure that a data protection policy is also configured for a scheduled backup, to define the workflow, backup group, and other required settings. [Configuring the data protection policy with NMC](#) on page 91 provides details.

Verifying the 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.

About this task

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 5 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 provides details.

Verifying the 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, which are defined in the Privileges attribute of the User Group resource.

About this task


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

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 datazones (RDZs), you must ensure that the user has the required privileges. When you use RDZs, you can also grant the privileges by using the Users and Privileges fields of the RDZ resource. The *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 6 User group privileges for NMDA operations

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 83	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 83	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 83	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 83	Change Application Settings and any prerequisite privileges
Configuration of NetWorker resources for backup or restore	NMC user	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	Change Application Settings and any prerequisite privileges

Database-specific user definition

[Table 6](#) on page 83 lists the user group privileges that are 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 7 Database-specific user for each database or application

Database or application	Definition of database-specific user
DB2	<ul style="list-style-type: none"> 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.

Table 7 Database-specific user for each database or application (continued)

Database or application		Definition of database-specific user
Informix		OS user that runs the <code>onbar</code> or <code>onsmsync</code> command.
Lotus		Domino or Notes OS user.
MySQL		MySQL OS user.
Oracle		<ul style="list-style-type: none"> • On UNIX: <ul style="list-style-type: none"> ▪ 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. <p>Note: If Oracle ASM is used, the database-specific user also includes the OS user that runs the ASM instance.</p> <ul style="list-style-type: none"> • On Windows: User that runs the Oracle Windows services.
SAP IQ		OS user that runs the SAP IQ server.
Sybase		<ul style="list-style-type: none"> • On UNIX: OS user that runs the Sybase ASE server. • On Windows: User that runs the Sybase ASE Windows services, which is typically the System user by default.
Orchestrated Application Protection	MongoDB MySQL PostgreSQL	OS user that has the permission to use the database user to connect to the database and run the backup and restore.

Requirements for backup deletion operations

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

About this task

- 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 6](#) on page 83.

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.

Verifying the 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.

About this task

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.

Verifying the 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.

About this task

The NetWorker server configuration must include a Device resource for each physical storage device that is used for backups and restores. Also, each storage device must contain a labeled and mounted volume.

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

Verifying the 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.

About this task

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

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.

Configuring the firewall support

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

About this task

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

For Orchestrated Application Protection, refer to the *Dell EMC BoostFS Configuration Guide* for the BoostFS requirements.

Configuring scheduled backups with the wizard

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

About this task

NMDA does not support the backup configuration wizard for scheduled SAP IQ backups or for Orchestrated Application Protection, which includes scheduled MongoDB backups, scheduled MySQL backups, and scheduled PostgreSQL backups. You must configure these scheduled backups without the wizard.

The backup configuration wizard can perform the following actions:

- Configure a new group or use an existing group.
- Configure a new NetWorker Client resource and associate the Client resource with the group.
 - i** **Note:** You must configure the data protection policy, backup workflow, and backup action for the scheduled backup through NMC and associate the group with the backup workflow, as described in [Configuring the data protection policy with NMC](#) on page 91.
- Save a copy of the configuration settings to a specified file on the NMDA host for reference purposes.
- Modify a backup configuration that was created with the configuration wizard.

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

- i** **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** option 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. [Oracle DBA and NetWorker Backup Administrator Collaboration](#) on page 333 provides details.

The *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 *NetWorker E-LAB Navigator*.

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 NMC GUI and select the appropriate options to use the configuration wizard for configuring an NMDA scheduled backup.

About this task

NMDA does not support the backup configuration wizard for scheduled SAP IQ backups or for Orchestrated Application Protection, which includes scheduled MongoDB backups, scheduled MySQL backups, and scheduled PostgreSQL backups. You must configure these scheduled backups without the wizard.

Ensure that you have configured the NetWorker Device, Pool, and Label Template resources before you use the wizard. You cannot configure these resources with the wizard. The latest *NetWorker Administration Guide* provides details on the resource configurations.

Note: After you complete the backup configuration with the wizard, ensure that a data protection policy is also configured on the NetWorker server to define the backup workflow and other required settings for the scheduled backup. [Configuring the data protection policy with NMC](#) on page 91 provides details.

Procedure

1. In the NMC **Enterprise** window, right-click the NetWorker server name, and then select **Launch Application**.
The *NetWorker Administration Guide* provides details on how to access the NMC interfaces.
2. In the **Administration** window, click **Protection**.
3. Start the wizard by using the appropriate method:
 - To create a backup configuration, right-click **Clients** in the left pane, and then select **New Client Wizard**.
 - To modify a backup configuration that was created with the wizard, right-click the NMDA client in the right pane, and then select **Modify Client Wizard**.
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—When the following conditions are true, the wizard enables you to apply a password change to other NMDA backup configurations with the same credentials:

- You use the wizard to modify a configuration.
- You select to change a database user password that is used for backup.

The feature facilitates easy password changes for the same users 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—When the following conditions are true, NMDA applies a password change to all the backup configurations that use the same database and username on the host:
 - You use database authentication on the client host.
 - You select to change the password for a specific database (Oracle Net service name) and database username.

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.

When the following conditions are true, NMDA applies a password change to all the backup configurations that use the same recovery catalog database and username on the host:

- You use an Oracle recovery catalog on the client host.
- You select to change the password for a specific recovery catalog database and username.

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 screen titled **Client Configuration Summary**, click **Create** to create the configuration or click **Back** to modify the configuration.

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 336.

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).

About this task

Procedure

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 134 provides details.
2. Create an NMDA configuration file, including the required parameter settings. [Configuring the NMDA parameters without the wizard](#) on page 89 provides details.
3. Configure the data protection policy and required NetWorker resources for the scheduled backup by using the NMC program. [Configuring the data protection policy with NMC](#) on page 91 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

About this task

For the client-side configuration of a scheduled backup, set the required parameters in the NMDA configuration file. [NMDA Parameters and Configuration File](#) on page 399 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 NSR_SERVER parameter for a scheduled backup because NMDA automatically passes the server information from the NetWorker server that started the backup to the backup processes. However, you can set NSR_SERVER for a manual backup.

Table 8 Mandatory parameters for scheduled NMDA backups

Database or application	Mandatory parameters for scheduled NMDA backups
DB2	<ul style="list-style-type: none"> • DB2INSTANCE (UNIX only) • DB2_NODE_NAME • DB2PATH (Windows only) • DB2_TBS_LIST (tablespace backup only) • DB2_USER • INSTHOME (UNIX only) • USER_PSWD

Table 8 Mandatory parameters for scheduled NMDA backups (continued)

Database or application	Mandatory parameters for scheduled NMDA backups
	<p>i Note: Set the encrypted DB2 user password in the USER_PSWD parameter with the <code>nsrdaadmin -P -z configuration_file_path</code> command as described in Table 39 on page 415.</p>
Informix	<ul style="list-style-type: none"> • INFORMIXDIR • INFORMIXSQLHOSTS (UNIX) • INFORMIXSQLHOSTS (Windows, with Informix 12.10 or later) • ONCONFIG
Lotus	<ul style="list-style-type: none"> • LOTUS_USER (UNIX only) • Notes_ExecDirectory • NSR_RESOURCE_DIR (UNIX only) • PATH (UNIX only)
MySQL	<ul style="list-style-type: none"> • MYSQL_CFG_FILE • MYSQL_DATADIR • USER_PSWD <p>i Note: Set the encrypted MySQL user password in the USER_PSWD parameter with the <code>nsrdaadmin -P -z configuration_file_path</code> command as described in Table 42 on page 436.</p>
Oracle	<ul style="list-style-type: none"> • ORACLE_HOME • ORACLE_SID (specific cases only) • TNS_ADMIN (specific case only)
SAP IQ	<ul style="list-style-type: none"> • IQ_OCS_PATH • IQ_USER • PATH (specific cases only) • SYBASE • USER_PSWD (when SAP IQ server has a password) <p>i Note: Set the encrypted SAP IQ user password in the USER_PSWD parameter with the <code>nsrdaadmin -P -z configuration_file_path</code> command as described in Table 46 on page 458.</p>
Sybase	<ul style="list-style-type: none"> • 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)

Table 8 Mandatory parameters for scheduled NMDA backups (continued)

Database or application		Mandatory parameters for scheduled NMDA backups
		<p>Note: Set the encrypted Sybase user password in the USER_PSWD parameter with the <code>nsrdaadmin -P -z configuration_file_path</code> command as described in Table 47 on page 464.</p>
Orchestrated Application Protection	MongoDB MySQL PostgreSQL	<ul style="list-style-type: none"> NSR_BACKUP_NAME NSR_BACKUP_SCRIPT NSR_DATABASE_TYPE NSR_INSTANCE_NAME

Configuring the data protection policy with NMC

You must create a data protection policy for an NMDA scheduled backup. When you configure the data protection policy, you create and assign the Client resource, group, policy, workflow, and backup action that define the settings for the scheduled backup.

About this task

An NMDA scheduled backup is defined as a traditional backup action within a workflow in the NetWorker server. The Client resource that is created for the backup is assigned to a group, which is assigned to the workflow. The workflow itself is assigned to the data protection policy.

Perform the following steps to configure the data protection policy for an NMDA scheduled backup. The *NetWorker Administration Guide* and NMC online help provide details on how to use NMC to configure the required NetWorker resources and settings.

NOTICE Data protection policy configurations do not apply to client-initiated or manual NMDA backups.

Procedure

- Open the NMC Administration interface, where you perform all the policy configuration procedures:
 - In the NMC **Enterprise** window, right-click the NetWorker server name, and then select **Launch Application**.
The **Administration** window appears.
 - In the **Administration** window, click **Protection**.
The *NetWorker Administration Guide* provides details on how to access the NMC interfaces.
- Ensure that the NetWorker Client resource is configured to define the backup data that the protection policy will manage. You can configure the Client resource with or without the configuration wizard:
 - Create or modify the Client resource with the wizard according to [Configuring scheduled backups with the wizard](#) on page 86.
 - Create or modify the Client resource without the wizard according to the following topic.
- Ensure that the backup group is created for the Client resource according to [Configuring the backup group](#) on page 98.

4. Ensure that the backup policy is created for the scheduled backup according to [Configuring the backup policy](#) on page 99.
5. Ensure that the backup workflow is created for the policy according to [Configuring the backup workflow](#) on page 99.
6. Ensure that the backup action is created for the workflow according to [Configuring the backup action](#) on page 100.

Configuring the Client resource

You must customize the Client resource for an NMDA scheduled backup by setting the resource attributes.

You can create and edit a Client resource through the Client Properties dialog box in the NMC Administration interface. You must enable the NMC diagnostic mode to view the advanced attributes in the Client resource. The *NetWorker Administration Guide* provides details.

Note: It is recommended that you specify the schedule (backup level), storage node, pool, and retention period in the backup action that is created as part of the data protection policy configuration, not in the Client resource.

If a client requires different settings from other clients in the same backup group, then you can set the specific attributes in the Client resource. In the backup action that you create with the Policy Action Wizard in NMC, you can select **Client Can Override** in the **Client Override Behavior** field. Then, the attribute settings in the Client resource will override the properties set in the backup action that is created for the data protection policy.

For example, you can set the retention period to 1 month in the backup action for all clients and then set a 3-month retention in the Client resource for a specific client. Also, you can set a pool in the backup action for all clients and then set a different pool in the Client resource for a specific client.

To override the backup schedule that is specified in the policy-based configuration, you can select or customize one of the preconfigured schedules available in the NetWorker server. You can also create a custom schedule. Then you can select the new schedule in the Schedule attribute of the Client resource. The *NetWorker Administration Guide* provides details on preconfigured and custom schedules.

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
 - Run once for each backup

If NMDA retries a scheduled backup, then NMDA runs 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.

Set the Client resource attributes for an NMDA scheduled backup as shown in the following table. The table lists the basic attributes first, followed by the advanced attributes:

- You can view the basic attributes in the **Client Properties** dialog box without enabling the NMC diagnostic mode.
- To view the advanced attributes, you must enable the NMC diagnostic mode as described in the *NetWorker Administration Guide*.

Table 9 NetWorker Client resource attributes

Attribute name	NMC tab	Attribute setting
Basic attributes:		
Aliases	Globals (1 of 2)	Specify all known aliases for the system that contains the NMDA software.
Backup command	Apps & Modules	<p>For Orchestrated Application Protection backups only (including MongoDB, MySQL, and PostgreSQL backups), specify the command to be used for the backups as follows:</p> <pre>nsroappbackup -z configuration_file_path</pre> <p>where <i>configuration_file_path</i> is the pathname of the NMDA configuration file that contains the backup parameter settings.</p> <p>The PostgreSQL WAL archiving backups also require configuration settings in the <code>postgresql.conf</code> file as described in Registering the PostgreSQL archive command on page 163.</p> <p>For all other backups of other applications, specify the command to be used for the backup as follows:</p> <pre>nsrdasv(.exe) -z configuration_file_path</pre> <p>On Windows systems only, if the configuration file pathname includes any spaces, perform one of the following actions:</p> <ul style="list-style-type: none"> • Change the configuration file pathname so it does not include spaces. • Use the Windows short pathname format. For example: <pre>nsrdasv.exe -z C:\Progra~1\Legato\nsr\apps\config\config.txt</pre> • Use quotes and double backslashes in the pathname. For example: <pre>nsrdasv.exe -z "C:\\Program Files\\Legato\\nsr\\apps\\config\\config.txt"</pre>
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.
Name	General	Specify the hostname of the database server host or application server host.
Parallelism	Globals (1 of 2)	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.
Protection group list	General	Specify the NetWorker backup group to use for a scheduled backup. Configuring the backup group on page 98 describes backup groups.

Table 9 NetWorker Client resource attributes (continued)

Attribute name	NMC tab	Attribute setting	
Remote access	Globals (2 of 2)	<p>Specify the username and fully qualified hostname of a remote system that can access backups of this host, in the following format:</p> <pre>user=<i>remote_username</i>,host=<i>remote_hostname</i></pre> <p>Set this attribute in the following cases (otherwise, leave it blank):</p> <ul style="list-style-type: none"> Backup from a cluster Backup in an Oracle Data Guard environment Recovery to a host other than the one being backed up <p>On a Solaris system with Solaris zones, the Remote Access attribute must contain the hostname of the zone in which NMDA operates.</p>	
Save set	General	Type of client	Save Set attribute setting
		DB2	<p>Specify the database and node to be backed up. Use the following format for the save set name:</p> <pre>DB2:/<i>database_name</i>/<i>node_name</i></pre> <p>You can optionally specify multiple database instances on separate lines. For example:</p> <pre>DB2:/<i>DB1</i>/NODE0000 DB2:/<i>DB2</i>/NODE0000</pre> <p>where <i>DB1</i> and <i>DB2</i> are databases of the same or different DB2 instances on the same client.</p> <p>You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in NMDA configuration file syntax on page 402.</p>
		Informix	<p>Specify the Informix server instance and (optionally) dbspace to be backed up. Separate multiple dbspace names with a space:</p> <pre>INFORMIX:/<i>instance</i>[/<i>dbspace1</i> <i>dbspace2</i>]</pre> <p>To back up the entire instance:</p> <pre>INFORMIX:/<i>instance</i></pre> <p>To back up selected dbspaces:</p> <pre>INFORMIX:/<i>instance</i>/<i>dbspace1</i> <i>dbspace2</i></pre> <p>You can optionally specify multiple instances and optionally dbspaces on separate lines. For example:</p> <pre>INFORMIX:/INST1 INFORMIX:/INST2/<i>dbspace2_1</i> <i>dbspace2_2</i></pre>

Table 9 NetWorker Client resource attributes (continued)


Attribute name	NMC tab	Attribute setting	
			<p>where INST1 and INST2 are different Informix instances on the same client.</p> <p>You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in NMDA configuration file syntax on page 402.</p>
Save set	General	Type of client	Save Set attribute setting
		Lotus	<p>Specify the name for the backup save set: NOTES: <i>[description]</i></p> <p>For example, the following names are valid names:</p> <ul style="list-style-type: none"> • NOTES: • NOTES: Monday Full Backup <p>To specify the data to back up, set one of these parameters in the NMDA configuration file:</p> <ul style="list-style-type: none"> • 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. <p>NMDA Parameters and Configuration File on page 399 describes the parameters. You can optionally specify save sets for multiple Domino servers on separate lines. For example:</p> <ul style="list-style-type: none"> • NOTES:Domino_Server1 • NOTES:Domino_Server2 <p>where Domino_Server1 and Domino_Server2 are names of backups for different Domino servers on the same client.</p> <p>You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in NMDA configuration file syntax on page 402.</p>
		MySQL	<p>Specify a name starting with MYSQL:/ for the backup:</p> <p><i>MYSQL:/unique_backup_name</i></p> <p>where <i>unique_backup_name</i> is a unique name that identifies the particular MySQL backup, for example, MYSQL:/InnoDB_partial_backup.</p> <p>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 42 on page 436 describes the parameters.</p> <p> Note: You cannot use the same backup name in different NMDA backup configurations for MySQL. MySQL backup</p>



Table 9 NetWorker Client resource attributes (continued)

Attribute name	NMC tab	Attribute setting	
			<p>information in NetWorker indexes for backups without Orchestrated Application Protection on page 182 and MYSQL_BACKUP_NAME provide details on MySQL backup names and their importance for restores.</p> <p>You can optionally specify the backup names for multiple instances on separate lines. For example:</p> <pre>MYSQL:DBSID1 MYSQL:DBSID2</pre> <p>where <i>DBSID1</i> and <i>DBSID2</i> are the unique backup names for different MySQL instances on the same client.</p> <p>You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in NMDA configuration file syntax on page 402.</p>
Save set	General	Type of client	Save Set attribute setting
		Oracle	<p>Specify the complete pathname of each RMAN script to use for the scheduled backup:</p> <pre>RMAN:RMAN_script_pathname</pre> <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • RMAN:/disk/rman_scripts/archlogbkup • RMAN:/disk/rman_scripts/fullbkup <p>You can optionally specify RMAN scripts for multiple databases on separate lines. For example:</p> <ul style="list-style-type: none"> • RMAN:/disk/rman_backup_PROD • RMAN:/disk/rman_backup_SALES <p>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.</p> <p>You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in NMDA configuration file syntax on page 402.</p>

Table 9 NetWorker Client resource attributes (continued)

Attribute name	NMC tab	Attribute setting	
		Orchestrated Application Protection	Specify the required keyword to indicate a scheduled backup for Orchestrated Application Protection. Use the following keyword for the save set name: OAPP:
		SAP IQ	Specify the SAP IQ database to be backed up. Use the following format for the save set name: IQ:/ <i>database_name</i> You can optionally specify multiple databases on separate lines. For example: <ul style="list-style-type: none">• IQ:/IQ_database1• IQ:/IQ_database2 This example backs up the two databases, database1 and database2. You can specify unique parameter settings for each save set by grouping the parameter settings within braces in the configuration file as described in NMDA configuration file syntax on page 402.
		Sybase	Specify the Sybase server and optionally the Sybase databases for backup: SYBASE:/ <i>ASE_server_name</i> [/ <i>database_name</i>] To back up the entire server, specify the following setting: SYBASE:/ <i>ASE_server_name</i> You can optionally specify multiple servers and optionally databases (of the same server or different servers) on separate lines. For example: <ul style="list-style-type: none">• 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 NMDA configuration file syntax on page 402.
Advanced attributes:			
Backup target disks	Globals (2 of 2)	Specify an ordered list of AFTD and Data Domain disk devices that will receive data for this client. When you specify a value in this attribute, NetWorker ignores the values that you specify in the Storage nodes attribute. This attribute does not apply to the client resource of the NetWorker server, and applies to each instance of the client resource. You can specify devices that are local or remote to the NetWorker server.	

Table 9 NetWorker Client resource attributes (continued)

Attribute name	NMC tab	Attribute setting
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 38.
Client direct	General	Select this attribute to enable the Client Direct feature for a scheduled backup to a DD Boost device, CloudBoost device, or AFTD. If a Client Direct backup is not possible, the backup routes the data through the storage node to the NetWorker device.
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.
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: <ul style="list-style-type: none"> 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. <p> Note: To enable a restore from a Data Domain device over an FC interface, you must set this attribute to Fibre Channel.</p>
Pool	General	Specify the pool to use for a backup.
Retention policy	General	Specify the minimum length of time that the NetWorker server maintains information about the backup data in the online media database. <p> Note: If you set the NSR_SAVESET_RETENTION parameter in a manual backup as described in NMDA Parameters and Configuration File on page 399, the parameter value overrides the Retention Policy attribute setting in the Client resource.</p>
Schedule	General	For scheduled backups only. Select the NetWorker backup schedule to use for a scheduled backup. Configuring the backup action on page 100 describes backup schedules.
Scheduled backup	General	For scheduled backups only. Select the checkbox to enable the Client resource for a scheduled backup.
Storage Nodes	Globals (2 of 2)	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 that is listed in the attribute. The default storage node name, nsrserverhost, represents the NetWorker server.

Configuring the backup group

You must create a backup group, also known as a protection group, to contain the NetWorker Client resource that you configured for an NMDA scheduled backup.

About this task

A group defines the set of Client resources for the backup workflow. Client resources can belong to multiple groups at a time. However, each workflow applies to only one group, and each group

can be assigned to only one workflow. You can assign a group to a workflow either when you create or edit the group, or when you create or edit the workflow.

Note: Before you create a group, create one or more Client resources to assign to the group.

When you create a backup group, you specify the group properties and add the Client resources to the group. After you create a group, you can edit all the properties of the group except the group name and group type. To rename a group, you must delete the group and then re-create it with the new name.

You can use NMC to create and edit a group. The *NetWorker Administration Guide* provides details on how to create and edit a protection group in the NMC Administration interface. Follow the procedures for a basic client group.

Configuring the backup policy

You must create a data protection policy for an NMDA scheduled backup. A policy acts as a container for backup workflows, actions, and groups. A policy can contain one or more workflows, but you can assign a workflow to only one policy.

About this task

When you create a policy, you specify the policy name and notification settings. You cannot modify the policy name. To rename a policy, you must delete the policy and then re-create it with the new name.

You can use NMC to create and edit a data protection policy. The *NetWorker Administration Guide* provides details on how to create and edit the policy properties in the NMC Administration interface.

After you create a policy, create the workflows and actions for the policy.

Configuring the backup workflow

You must create a workflow for an NMDA scheduled backup. The backup is defined as an action within the workflow. The backup group is assigned to the workflow, and the workflow itself is assigned to the data protection policy of the backup.

About this task

The backup workflow can include one or more actions:

- The workflow can include a single action as the traditional NMDA backup, which is a scheduled backup of the save sets defined for the Client resources in the group.
- The workflow can also include a sequence of actions to be performed sequentially or concurrently.

The backup workflow can optionally include a check connectivity action before the backup action. The check connectivity action tests the connectivity between the client and NetWorker server to determine if the backup can proceed.

The backup workflow can also optionally include a clone action after the backup action. The clone action creates a cloned copy of the backup save sets. Alternatively, you can create a separate workflow for the clone action and run the clone operation separately from the backup.

The workflow defines the start time and schedule window for the action or sequence of actions, how often the workflow runs, the order of actions to perform, and the group of Client resources on which the actions occur.

Note: Before you create a backup workflow, create a policy for the workflow and create the group to assign to the workflow. You can create the actions for the workflow either when you create the workflow or separately.

You can use NMC to configure the backup workflow. The *NetWorker Administration Guide* provides details on how to create and edit the workflow in the NMC Administration interface. The guide also describes the supported workflow types and the actions that a workflow can contain.

Configuring the backup action

You must create a backup action that defines the NMDA scheduled backup in a policy-based configuration. You assign the backup action to the workflow, which is assigned to the data protection policy.

About this task

An action is the operation that occurs on the Client resources in the group assigned to the workflow. The *NetWorker Administration Guide* provides details on the supported types of actions.

For a regular (nonsnapshot) scheduled backup, you must create a traditional backup action, which is a scheduled backup of the save sets defined for the Client resources in the group.

When you create the backup action, you specify the backup level for each day of the weekly or monthly schedule. You can also specify the destination storage node, destination pool, retention period for the backup, and any advanced options as required.

If you want to override the storage node, pool, or retention period setting for a client, you can specify the corresponding attribute in the Client resource. The attribute settings in the Client resource override the settings in the backup action.

Note: Before you create the backup action, create the policy and workflow to contain the action. You can create actions for a workflow either when you create the workflow or separately. Optionally, create a check connectivity action to precede the backup action in the workflow.

You can use the Policy Action Wizard in the NMC Administration interface to configure the backup action. The *NetWorker Administration Guide* provides details on how to create and edit the action by using the Policy Action Wizard.

When you create the backup action for the NMDA scheduled backup with the Policy Action Wizard:

- From the **Action Type** list, select **Backup**.
- From the secondary action list, select **Traditional**.

Specify the backup levels in the backup action according to the information in the following table, which shows the backup levels that are supported for each type of NMDA database or application.

Table 10 Backup levels that are specified in the backup action

Database or application	Supported backup levels
DB2	<ul style="list-style-type: none"> • Full—DB2 full backup of all the database data. • Cumulative Incr—DB2 incremental backup. • Incr—DB2 delta backup. • Skip—NetWorker server skips the backup on that day. <p>To support the NetWorker backup levels from the backup action:</p> <ul style="list-style-type: none"> • Set the DB2_APPLY_NW_LEVELS parameter to TRUE in the NMDA configuration file, as described in DB2_apply_nw_levels.

Table 10 Backup levels that are specified in the backup action (continued)

Database or application	Supported backup levels
	<ul style="list-style-type: none"> Set the TRACKMOD parameter to ON with a DB2 command at the operating system command line: <pre data-bbox="480 432 1331 485">db2 update db cfg for <i>sample</i> using TRACKMOD ON</pre> where <i>sample</i> is the name of the database to be backed up.
Informix	<ul style="list-style-type: none"> Full—ON-Bar level 0 backup of all the selected dbspaces of the database instance. Cumulative Incr—ON-Bar level 1 backup of the data that has changed since the last level 0 (full) backup. Incr—ON-Bar level 2 backup of the data that has changed since the last level 1 backup. Skip—NetWorker server skips the backup on that day.
Lotus	<ul style="list-style-type: none"> Full—Lotus full backup of all the data. Incr—Lotus incremental backup of only data that has changed since the last backup. Logs Only—Lotus transaction logs only backup. Skip—NetWorker server skips the backup on that day.
MySQL	<ul style="list-style-type: none"> Full—Full backup of all the data. Cumulative Incr—Cumulative incremental backup of only the data that has changed since the last full backup. Incr—Differential incremental backup of only the data that has changed since the last full or incremental backup.
Oracle	<ul style="list-style-type: none"> Full, Cumulative Incr, or Incr—NetWorker server runs the backup script on that day. The backup level set in the RMAN backup script determines the Oracle backup level. <p data-bbox="480 1419 1331 1514">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.</p> Skip—NetWorker server skips the backup on that day.
SAP IQ	<ul style="list-style-type: none"> Full—SAP IQ full backup of all the database data. Cumulative Incr—SAP IQ "incremental since full" backup of the database block changes since the last full backup. Incr—SAP IQ incremental backup of the database block changes since the last backup of any type. Logs Only—SAP IQ transaction logs only backup. Skip—NetWorker server skips the backup on that day. <p data-bbox="432 1881 1139 1913">Note: A read-only selective backup supports a full backup only.</p>

Table 10 Backup levels that are specified in the backup action (continued)

Database or application		Supported backup levels
Sybase		<ul style="list-style-type: none"> • Full—Sybase full backup of all the data. • Cumulative Incr—Sybase cumulative backup of the database data and log changes since the last full database backup. • Incr—Sybase incremental backup of only the transaction logs generated since the last full or incremental backup. • Skip—NetWorker server skips the backup on that day. <p>Note: 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).</p>
Orchestrated Application Protection	MongoDB	<ul style="list-style-type: none"> • Full—NetWorker server runs the backup shell script that is defined in the BACKUP section and FULL subsection of the configuration file. • Skip—NetWorker server skips the backup on that day. <p>Note: MongoDB supports only a full backup.</p> <p>In the backup shell script, define the correct script behavior that is associated with the NetWorker backup level definition, to ensure the expected backup behavior in the future. In the BACKUP section of the configuration file, specify the backup shell script file pathname in the NSR_BACKUP_SCRIPT setting in the full backup level subsection. The <code>nsroappbackup</code> program uses the correct backup shell script based on the backup level.</p>
	MySQL	<ul style="list-style-type: none"> • Full—NetWorker server runs the backup shell script that is defined in the BACKUP section and FULL subsection of the configuration file. • Logs Only—NetWorker server runs the backup shell script that is defined in the BACKUP section and TXNLOG subsection of the configuration file. • Skip—NetWorker server skips the backup on that day. <p>Note: MySQL supports a full backup and transaction log backup.</p> <p>In the backup shell script, define the correct script behavior that is associated with the NetWorker backup level definition, to ensure the expected backup behavior in the future. In the BACKUP section of the configuration file, specify the different backup shell script file pathnames in the NSR_BACKUP_SCRIPT setting in the FULL and TXNLOG subsections. The <code>nsroappbackup</code> program uses the correct backup shell script based on the backup level.</p>
	PostgreSQL	<ul style="list-style-type: none"> • Full—NetWorker server runs the backup shell script that is defined in the BACKUP section and FULL subsection of the configuration file. • Skip—NetWorker server skips the backup on that day. <p>Note: PostgreSQL supports a full backup and log only backup. The PostgreSQL server must start the log only backup, so you cannot set the log only backup level in the backup action.</p> <p>In the backup shell script, define the correct script behavior that is associated with the NetWorker backup level definition, to ensure the expected backup behavior in the future. In the BACKUP section of the configuration file, specify the backup shell script file pathname in the NSR_BACKUP_SCRIPT setting in</p>

Table 10 Backup levels that are specified in the backup action (continued)

Database or application	Supported backup levels
	each required backup level subsection. The <code>nsroappbackup</code> program uses the correct backup shell script based on the backup level.

Configuring manual backups

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

Procedure

1. Ensure that you have configured the required NetWorker resources as described in [Verifying the basic NetWorker resource configurations](#) on page 81.
2. Set the required NMDA parameters:
 - For DB2, Lotus, MySQL, SAP IQ, and Sybase manual backups and for Orchestrated Application Protection, which includes MongoDB, MySQL, and PostgreSQL manual backups, 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 then set the parameters in the RMAN backup script by using the `send` command where possible. [Oracle backup considerations](#) on page 134 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, Oracle, SAP IQ, and Sybase; also, MySQL and PostgreSQL through Orchestrated Application Protection) 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.

[NMDA Parameters and Configuration File](#) on page 399 describes the parameters.

Table 11 Mandatory parameters for manual NMDA backups

Database or application	Mandatory parameters for manual NMDA backups
Lotus	<ul style="list-style-type: none"> • <code>Notes_ExecDirectory</code> • <code>NSR_RESOURCE_DIR</code> (UNIX only) • <code>PATH</code> (UNIX only)
MySQL	<ul style="list-style-type: none"> • <code>MYSQL_BACKUP_NAME</code> • <code>MYSQL_CFG_FILE</code> • <code>MYSQL_MEB_PATH</code> (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) • <code>MYSQL_SBT_LIB_PATH</code> (required only if you have 32-bit and 64-bit MEB on this same system)

Table 11 Mandatory parameters for manual NMDA backups (continued)

Database or application		Mandatory parameters for manual NMDA backups
Oracle		<ul style="list-style-type: none"> NSR_DATA_VOLUME_POOL* (for backups that generate backup copies)
SAP IQ		<ul style="list-style-type: none"> IQ_USER NSR_SAVESET_NAME SYBASE USER_PSWD (when SAP IQ server has a password) <p>Note: To set the encrypted SAP IQ user password in the USER_PSWD parameter, you must use the <code>nsrdaadmin -P -z configuration_file_path</code> command, as described in the USER_PSWD description in Table 46 on page 458.</p>
Sybase		<ul style="list-style-type: none"> NSR_BACKUP_PATHS SYBASE SYBASE_USER USER_PSWD (when Sybase server has a password) <p>Note: To set the encrypted Sybase user password in the USER_PSWD parameter, you must use the <code>nsrdaadmin -P -z configuration_file_path</code> command, as described in the USER_PSWD description in Table 47 on page 464.</p>
Orchestrated Application Protection	MongoDB	<ul style="list-style-type: none"> NSR_BACKUP_NAME
	MySQL	<ul style="list-style-type: none"> NSR_BACKUP_SCRIPT
	PostgreSQL	<ul style="list-style-type: none"> NSR_DATABASE_TYPE NSR_INSTANCE_NAME

- Perform any additional configurations that are required for the specific database backup or application backup:
 - [DB2 backup considerations](#) on page 121
 - [Informix backup considerations](#) on page 123
 - [Lotus backup considerations](#) on page 124
 - [MySQL backup considerations](#) on page 130
 - [Oracle backup considerations](#) on page 134
 - [SAP IQ backup considerations](#) on page 144
 - [Sybase backup considerations](#) on page 146
 - [Orchestrated Application Protection backup considerations](#) on page 153
 - [MongoDB backup considerations](#) on page 156
 - [MySQL backup considerations](#) on page 158
 - [PostgreSQL backup considerations](#) on page 161

Configuring Data Domain backups

About this task

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

[Deduplication backups and restores with Data Domain](#) on page 26 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 110 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 106 and [Configuring DD Boost backups to use a storage node](#) on page 108 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 105.
- Follow the best practices for Oracle backups to Data Domain in [Oracle considerations for Data Domain backups](#) on page 106.

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](#).
- 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](#).
- 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 *NetWorker E-LAB Navigator* 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.

Procedure

1. Review [Best practices for DD Boost backups](#) on page 106.
2. Configure the DD Boost device and enable the Data Domain system for NetWorker operations. The *NetWorker Data Domain Boost 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 articles:
 - *Fibre Channel Devices with Products using DD Boost in Linux/UNIX Environment* (Document ID dd95007)
 - *Fibre Channel Devices with Products using DD Boost in Windows Environment* (Document ID dd95005)

Use the document ID to search for these articles on the Support website at <https://support.emc.com>.

Note: This step is not required for backups through a storage node or for snapshot-based backups that are performed with NSM. [Table 7](#) on page 83 provides a definition of the database-specific operating system user.

4. Associate the backup with a pool that contains the DD Boost device that is configured in step 2:
 - For a scheduled backup, associate the required media pool with the backup action by using the NMC program. The *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](#) 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: The Client Direct setting in the wizard field or Client resource takes precedence over the NSR_DIRECT_ACCESS parameter setting.

Table 12 Configuring a Data Domain backup

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 that is configured with the wizard	Configuring scheduled backups with the wizard on page 86	On the Client Backup Options wizard page: <ul style="list-style-type: none"> • 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.
Scheduled backup that is configured without the wizard	Configuring scheduled backups without the wizard on page 89	In the Client resource of the NMDA host: <ul style="list-style-type: none"> • 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 103	In the NMDA configuration file: <ul style="list-style-type: none"> • Set NSR_DEVICE_INTERFACE=DATA_DOMAIN. NSR_DEVICE_INTERFACE 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.

Table 12 Configuring a Data Domain backup (continued)

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
		NSR_DATA_DOMAIN_INTERFACE 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.

Procedure

1. Follow step 1 to step 5 in [Configuring DD Boost backups to use Client Direct](#) on page 106.
2. Ensure that the Client Direct feature is disabled:
 - For a manual backup, set NSR_DIRECT_ACCESS=No in the NMDA configuration file (DB2, Lotus, MySQL, SAP IQ, Sybase), Oracle RMAN script, or environment variable (Informix). [NSR_DIRECT_ACCESS](#) describes the parameter.
 - For a scheduled backup, 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 DB2, Informix, MySQL, Oracle, or Sybase backup that backs up database data to a Data Domain device and backs up transaction logs to a traditional NetWorker device.

Procedure

1. Follow step 1 to step 5 in [Configuring DD Boost backups to use Client Direct](#) on page 106 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 12](#) on page 107.
2. Configure a pool with traditional devices for the log backups, and then 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 CloudBoost backups

About this task

This topic provides specific information about using CloudBoost with NMDA. The latest version of the *NetWorker with CloudBoost Integration Guide* provides additional information about the CloudBoost configuration and initial setup.

[Backups and restores with CloudBoost](#) on page 27 describes the NMDA support of backups with a CloudBoost device.

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

- NetWorker documentation describes the configuration of a CloudBoost backup to a CloudBoost device.
- [Configuring CloudBoost backups to use Client Direct](#) on page 109 and [Configuring CloudBoost backups to use a storage node](#) on page 110 describe the configuration of a backup to a CloudBoost device through either the Client Direct method or a storage node.

Best practices for CloudBoost backups

Before you configure a CloudBoost backup with a CloudBoost device, ensure that the NetWorker client and server releases support the CloudBoost appliance. The *NetWorker E-LAB Navigator* provides details.

Use Client Direct access for CloudBoost backups to eliminate data transfer between the NMDA client and NetWorker storage node.

CloudBoost and NetWorker documentation describes the requirements for CloudBoost backups.

Configuring CloudBoost backups to use Client Direct


You can configure a scheduled backup or manual backup to a CloudBoost device that uses the Client Direct feature on the NMDA host and passes data directly to the CloudBoost device, bypassing the storage node.

Procedure

1. Review [Best practices for CloudBoost backups](#) on page 109.
2. Configure and enable the CloudBoost appliance for NetWorker operations. The latest version of the *NetWorker with CloudBoost Integration Guide* provides details.
3. Configure and enable the CloudBoost devices for NetWorker operations. The latest version of the *NetWorker with CloudBoost Integration Guide* provides details.
4. Associate the backup with a pool that contains the CloudBoost device that is configured in steps 2 and 3:
 - For a scheduled backup, associate the required media pool and storage node with the backup action by using the NMC program. Use the CloudBoost appliance as the storage node. The *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](#) describes the parameter.
 - Specify the CloudBoost appliance as the storage node in the NetWorker Client resource attributes. [Configuring the Client resource](#) on page 92 provides details.

5. For a Client Direct backup, use the default configuration settings, except set `NSR_DIRECT_ACCESS=Yes` for a manual backup. 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:** The Client Direct setting in the wizard field or Client resource takes precedence over the `NSR_DIRECT_ACCESS` parameter setting.

Configuring CloudBoost backups to use a storage node

You can configure a scheduled backup or manual backup to a CloudBoost appliance, where the backup does not try to use the Client Direct method but instead transfers the backup data through a storage node.

About this task

Procedure

1. Follow step 1 to step 4 in [Configuring CloudBoost backups to use Client Direct](#) on page 109.
2. Ensure that the Client Direct feature is disabled:
 - For a manual backup, set `NSR_DIRECT_ACCESS=No` in the NMDA configuration file (DB2, Lotus, MySQL, SAP IQ, Sybase), Oracle RMAN script, or environment variable (Informix). [NSR_DIRECT_ACCESS](#) describes the parameter.
 - For a scheduled backup, clear the Client Direct checkbox in the wizard or disable the Client Direct attribute in the NetWorker Client resource.

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.

About this task

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.

About this task

Procedure

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 *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 *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 86
 - [Configuring scheduled backups without the wizard](#) on page 89

- [Configuring manual backups](#) on page 103
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.

About this task

Procedure

1. Follow step 1 to step 4 in [Configuring Client Direct backups to AFTD](#) on page 110.
2. Ensure that the Client Direct feature is disabled:
 - For a manual backup, set NSR_DIRECT_ACCESS=No in the NMDA configuration file (DB2, Lotus, MySQL, SAP IQ, Sybase), Oracle RMAN script, or environment variable (Informix). [NSR_DIRECT_ACCESS](#) 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

About this task

Note:

For MySQL probe-based backups, NMDA supports only user-defined probes.

NMDA supports probe-based backups of Oracle disk backups in addition to regular Oracle data.

NMDA does not support probe-based backups for snapshot backups with NSM.

Before you configure a probe-based backup, ensure that you have installed the required NetWorker releases, as described in the *NetWorker E-LAB Navigator*.

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 probe or all the probes that are associated with the backup are successful, depending on the probe success criteria that are specified in the backup configuration.
You can set the probe success criteria in the probe action.

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 probe or all 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 that were generated since the last probe-based backup.
- For Lotus, the `nsrdaprobe` program checks for the size of Domino transaction logs that were generated since the last backup of the Domino database, which is specified by `LOTUS_NSF_FILE` in the NetWorker Probe resource.
- For Sybase, the `nsrdaprobe` program checks for the number of transaction log pages that were generated since the last probe-based backup.

Note:

The `nsrdaprobe` program does not support Orchestrated Application Protection or the features that Orchestrated Application Protection supports.

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 number of log pages that were 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 number of log pages that were 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 4 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 NetWorker Probe resource and then complete the scheduled backup configuration of the Client resource and other required resources.

About this task

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

Procedure

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 13 NetWorker Probe resource attributes

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 <code>nsrdaprobe</code> to use the NMDA probe, or specify the user-created script or program.		
Command Options	Required for the <code>nsrdaprobe</code> program only. Specify a comma-separated list of parameters with their settings. Command Options settings for the nsrdaprobe program on page 116 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: <ul style="list-style-type: none"> • A backup administrator has used the wizard to configure NMDA scheduled backups of Oracle disk backups of the database. • You want <code>nsrdaprobe</code> to check if any new Oracle disk backups have been created for the database since the last successful probe-based backup. With <code>DBA_DISK_BACKUP=TRUE</code> , if the probe finds any new Oracle disk backups and no Oracle disk backup is running for the database, the probe triggers an NMDA scheduled backup of the new disk backups. Note: Do not set <code>DBA_DISK_BACKUP</code> in the same Probe resource as <code>LOG_THRESHOLD</code> .

Table 13 NetWorker Probe resource attributes (continued)

Attribute	Description		
	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 that is 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 that is required to trigger a new probe-based backup.
	LOTUS_NSF_FILE	Lotus	Recommended. Specify the full pathname of a Domino database file that is used for transaction log querying. The specified database file can be any file that is backed up by the associated scheduled backup. If you do not set this parameter, the default value that is used is <code>admin4.nsf</code> in the Lotus data directory (UNIX) or <code>C:\Program Files\IBM\Lotus\Domino\data\admin4.nsf</code> (Windows).
	NSR_DEBUG_LEVEL	All	Optional. Specify the level of debug information that is generated by the probe and written to the debug log. NSR_DEBUG_LEVEL 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 provides details.
Command Options	Parameter	Type of client	Description
	NSR_ORACLE_CONNECT_FILE	Oracle	Optional. Set if both of the following conditions are true: <ul style="list-style-type: none"> 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 that are required to connect to the Oracle database. Command Options settings for the nsrdaprobe program on page 116 provides a sample setting of this parameter.
	ORACLE_SERVICE	Oracle	Optional. Set if both of the following conditions are true:

Table 13 NetWorker Probe resource attributes (continued)


Attribute	Description		
			<ul style="list-style-type: none"> • 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 <code>nsroraadmin -r add sid=<i>Net_service_name</i> home=<i>Oracle_home</i> connect=<i>connect_filepath</i></code>. <p>Configuring the NWORA resource file with the nsroraadmin program on page 388 describes the <code>nsroraadmin</code> command.</p> <p>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.</p> <p>Command Options settings for the nsrdaprobe program on page 116 provides a sample setting of this parameter.</p>

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 scheduled backup with or without the wizard by following the instructions in [Configuring the data protection policy with NMC](#) on page 91.

You must configure the required Client resource, group, backup policy, workflow, probe action, and backup action.

- a. Configure the Client resource with or without the configuration wizard. In the Probe attribute of the Client resource, 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.

- b. Create the group to contain the Client resource by using the NMC Administration interface. The *NetWorker Administration Guide* provides details on the NMC interfaces.
- c. Create the backup policy and workflow by using the NMC Administration interface. You must assign the group to the workflow, and assign the workflow to the backup policy.
- d. Create the probe action and backup action by using the Policy Action Wizard in the NMC Administration interface:
 - When you create the probe action with the wizard, you must select **Probe** from the **Action Type** list.
 - When you create the backup action with the wizard, you must select **Backup** from the **Action Type** list, and select **Traditional** from the secondary action list.

You must assign the probe action and backup action to the workflow. In the workflow, the probe action must precede the backup action. The backup action is performed only if the probe conditions are met during the probe action.

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:
 1. 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
```

2. 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
```

3. 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 388 describes the `nsroraadmin` command.

Example 4 Multiple probes for a probe-based backup

Example 4 Multiple probes for a probe-based backup (continued)

You can configure a probe-based backup with multiple probes. If you select the checkbox **Start backup only after all probes succeed** in the probe action, then the backup starts only after all the probes succeed. Otherwise, the backup starts after any of the probes succeed.

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 hosts.

The `nsrdaprobe` program is in the `/usr/sbin` directory on the NMDA client. The `nsrdaprobe` program checks for the number of Oracle log files that are 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 complete the following steps to configure this probe-based backup with multiple probes.

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


Example 4 Multiple probes for a probe-based backup (continued)

- 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 that is 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 that are 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.

About this task

 **Note:** The Orchestrated Application Protection feature does not support parallel backups. However, the database that is supported by Orchestrated Application Protection might itself support parallel backups. In this case, the database backup utility can still work with the parallel mode.

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, Data Domain, or CloudBoost.
- 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.

i **NOTICE** You must disable NetWorker multiplexing for DB2 backups and Sybase backups. Otherwise, the restore becomes suspended or fails.

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, Data Domain, or CloudBoost 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 *NetWorker Data Domain Boost Integration Guide* and the latest version of the *NetWorker with CloudBoost Integration Guide* provide details.

To configure an NMDA parallel backup as either a manual or scheduled backup:

Procedure

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 that are used for NMDA backups. [Verifying the NetWorker Server resource](#) on page 82 and [Configuring the Client resource](#) on page 92 describe the parallelism setting.
 - b. If required, create a Device resource for each device according to [Verifying the NetWorker Device resource](#) on page 85.
 - c. For DB2 and Sybase backups, disable NetWorker multiplexing.
3. Use the instructions in the following topics to specify the number of sessions (parallelism) for a backup:
 - [Configuring DB2 parallel backups](#)
 - [Configuring Informix parallel backups](#)
 - [Configuring Lotus parallel backups](#)
 - [Configuring Oracle parallel backups](#)
 - [Configuring Sybase parallel \(multistripe\) backups](#)

i **Note:** The backup parallelism must not be greater than the total number of device sessions that are available for the backup through the media pool configuration, as determined in [step 1](#). Otherwise, the backup becomes suspended because it is waiting for an available device to be mounted for the proper pool, or the backup fails if not enough devices are available.

Configuring DB2 parallel backups

Specify that the backup will use the number of sessions equal to the number of available backup devices:

About this task

- 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.

About this task

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:

About this task

- 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.

About this task

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 118 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](#).

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

Specify the number of multistripe sessions to use in a Sybase backup:

About this task

- 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 backup topics require specific considerations:

- DB2 backup best practices
- Automatic deletion of DB2 recovery objects
- Automatic backups of DB2 transaction logs

Best practices for DB2 backups

The NMDA DB2 backup is affected by the specific settings that you use for backup devices, sessions, and the data and log backups.

Review the following best practices for DB2 backups:

- If the backup volume pool is configured with multiple backup devices, specify the number of DB2 sessions as the number of available backup devices that are used for the DB2 backup. To ensure that each session is directed to a different device, set the number of target sessions to 1 for each backup device. This setting helps to improve the backup performance.
 - ① **Note:** Ensure that the DB2 restore uses the same number of sessions as the DB2 backup.
- Set different pools for DB2 data and log operations by using different values for the `NSR_LOG_VOLUME_POOL` and `NSR_DATA_VOLUME_POOL` backup parameters. This setting helps to optimize the data throughput to the storage devices.
- Include the archive log in online database backups by using the `include logs` option in the `db2 backup database <database_name> online` command. This practice helps to restore

and roll forward the image to a consistent point-in-time during the backup if the archived log backup by `logarchmethn` is misplaced.

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:

About this task

- 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 6](#) on page 83 describes the required privileges.

You can also prune the history file and delete the backups manually. [Deleting DB2 backups](#) on page 185 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 according to the 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.

About this task

i **NOTICE** DB2 transaction log backups are manual backups, even when performed as part of a DB2 scheduled data backup. NetWorker features that are related to a data protection policy, such as the clone action and the protection period, apply only to scheduled backups and do not apply to DB2 transaction log backups.

Complete the required steps to configure the automatic backup of DB2 transaction logs when they become full.

Procedure

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`.

[NMDA Parameters and Configuration File](#) on page 399 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. If the configuration is for a high-availability environment that does not contain a common virtual hostname, set the `NSR_CLIENT` parameter to the same value on all the involved nodes.

The following sample shows the configuration file content for transaction log backups:

```
NSR_SERVER=TURBO
NSR_LOG_VOLUME_POOL=DB2INST1_Logs
```

- 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.
- 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 168 or [Performing DB2 manual backups with the DB2 GUI](#) on page 169.

 **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:

About this task

- 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 that is 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`:

About this task

- 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:

About this task

```
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

About this task

[Lotus full and incremental backups](#) on page 37 and [Lotus transaction log backups](#) on page 37 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` in the NMDA configuration file.
- For a scheduled backup configured with the NMDA configuration 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 NMDA configuration wizard, set the backup level **Logs Only** in the backup action (created with the Policy Action Wizard in NMC) to specify a transaction log backup. You assign the backup action to the backup workflow, which also contains the group of the Client resource that is configured for the NMDA Lotus transaction log backup.

The *NetWorker Administration Guide* describes how to use NMC to create a policy-based configuration for a scheduled backup, including the backup policy, workflow, group, and required NetWorker resources.

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 5](#) shows how the `NSR_RETAIN_NUM_LOGS` setting can affect log backups.

You can set the `NSR_PARALLELISM` parameter or corresponding wizard option to specify the maximum number of concurrent streams to send to the NetWorker server during a transaction logs only backup.

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 424 describes the parameters.

Example 5 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.

Example 5 Using NSR_RETAIN_NUM_LOGS to control reusable logs (continued)

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 quantity 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 that are 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.

About this task

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 that is named db.nsf for the partition:

```
NOTES:partition1_/disk2/notesdata1/db.nsf
```

Configuring Lotus DAOS backups

About this task

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.

You must perform the required steps to configure an integrated backup of Domino database files and DAOS (NLO files).

Procedure

1. Review [Best practices for Lotus DAOS backups](#) on page 127.
2. Complete [Configuring integrated Lotus DAOS backups](#) on page 128.

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.

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:

About this task

Procedure

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.

About this task

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 the following methods:


- [Wizard method \(server-side configuration\)](#) on page 128
- [Nonwizard method \(client-side configuration\)](#) on page 128

Wizard method (server-side configuration)

To configure a scheduled integrated backup with the configuration wizard:

Procedure

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 that are required to configure the Domino server backup first, followed by the DAOS backup configuration.

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.

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:

Procedure

1. Set the parameters that are 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.

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.

Results

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_RETENTION, to keep the DAOS backups for a longer time.

Note: NSR_SAVESET_RETENTION applies to manual backups only.

- 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.

[NMDA Parameters and Configuration File](#) on page 399 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 backup configuration wizard or the nonwizard method.

About this task

For either method, you can use the **Advanced Options** page in the Policy Action Wizard in NMC to optionally set **Retries** to a value greater than zero in the backup action that is created for the scheduled backup. In the corresponding workflow, you can set the **Restart Window** as required. This setting enables the automatic restart of a failed backup. Otherwise, you must restart the backup manually through the workflow restart option.

- To configure a restartable scheduled backup with the NMDA configuration wizard, complete [Configuring scheduled backups with the wizard](#) on page 86 and select the **Checkpoint Enabled** checkbox on the proper client wizard page.
- To configure a restartable scheduled backup without the NMDA configuration wizard, complete [Configuring scheduled backups without the wizard](#) on page 89 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 181 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:

- Requirements of MySQL Enterprise Backup
- Configuration parameter settings
- MySQL backup granularity
- MySQL incremental backups
- MySQL binary log backups
- MySQL backup privileges

Requirements of MySQL Enterprise Backup

The NMDA MySQL operations leverage the MySQL Enterprise Backup (MEB) to protect the MySQL database. The MySQL Enterprise Backup is part of the MySQL Enterprise Edition, which is required by the NMDA MySQL backups that are performed through the MySQL Enterprise Backup.

If you do not have the MySQL Enterprise Edition, you can choose to perform the NMDA MySQL operations through the Orchestrated Application Protection feature. [MySQL operations through Orchestrated Application Protection](#) on page 158 provides more details.

Configuration parameter settings

You must set the required configuration parameters to enable an NMDA MySQL backup:

- You must have a [client] section in the valid MySQL configuration file, `my.cnf`, that you provide for NMDA MySQL. In the [client] section, you must set the socket or port parameter

to specify how NMDA MySQL connects to the MySQL instance. This parameter setting is the same as in the [mysqld] section, which the MySQL server daemon uses. For example:

```
[client]
port = 3306
socket = /var/lib/mysql/mysql.sock
```

Note: NMDA MySQL does not support settings for multiple MySQL instances in a single MySQL configuration file. As a workaround, to specify the NMDA MySQL configuration for multiple instances, you can specify the settings for each MySQL instance in its own separate configuration file. Ensure that all the MySQL configuration files have the correct setting for the MySQL server, as set in the original configuration file. Update the [client] section in all the configuration files with the corresponding MySQL instance information.

- 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 username 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.

[NMDA Parameters and Configuration File](#) on page 399 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.

About this task

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 that is 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 that is configured with the wizard, you must specify the corresponding configuration settings in the wizard.

Table 14 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

Table 14 Configuration parameters for InnoDB partial backups (continued)

Types of InnoDB partial backups	Configuration parameters to enable the backups
Backup of specified tables when the file-per-table option is enabled	<code>MYSQL_INCLUDE = "<regular expression matching the per-table data file names>"</code>
Backup of only the tables in a specified instance	<code>MYSQL_ONLY_INNOODB_OPTIONS = NO_FRM</code>
With MEB 3.7 or later, backup of all the tables and their associated <code>.frm</code> files in a specified instance	<code>MYSQL_ONLY_INNOODB_OPTIONS = WITH_FRM_ALL</code>
With MEB 3.7 or later, backup of specified tables and their associated <code>.frm</code> files when <code>MYSQL_INCLUDE</code> is set and the file-per-table option is enabled	<code>MYSQL_ONLY_INNOODB_OPTIONS = WITH_FRM_RELATED</code>

For example, the NMDA MySQL configuration file for a partial backup contains the following parameter settings:

```
MYSQL_INCLUDE="innodb-sales*"
MYSQL_ONLY_INNOODB_OPTIONS=NO_FRM
```

With MEB 3.11, due to a limitation in MEB, a server-side scheduled partial backup (configured with the wizard) of MyISAM tables fails when only MyISAM tables are selected. A partial backup of both MyISAM and InnoDB tables succeeds with MEB 3.11.

With MEB 3.7 and 3.8.0, due to a limitation in MEB, if you set either `MYSQL_ONLY_INNOODB_OPTIONS = WITH_FRM_ALL` or `MYSQL_ONLY_INNOODB_OPTIONS = WITH_FRM_RELATED` for an InnoDB backup, you must run the InnoDB backup with the correct permissions:

- With MEB 3.7, you must run the InnoDB backup as an OS user with write permissions to the parent directory of the MySQL data directory.
- With MEB 3.8.0, you must run the InnoDB backup as an OS user with write permissions to the parent directory of the MySQL backup directory.

Considerations for MySQL incremental backups

With MEB 3.7 or later, you can set `MYSQL_INCR_OPTIONS=REDO_LOG_ONLY` to configure an incremental backup of only the redo log. 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).

About this task

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.

About this task

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.

NMDA with MEB 3.11 has a limitation with binary log backups. MEB 3.11 supports the backup of binary logs as part of a MySQL instance backup. You can run log file backups with either MEB or NMDA, but not both. If you run log file backups with both MEB and NMDA, log files can become corrupted, lost, or out of sync.

MySQL backup privileges

Before you run an NMDA MySQL backup, ensure that the backup user has the following MySQL privileges:

About this task

- `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 that are 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:

About this task

- You can store InnoDB tables in a specified directory outside the MySQL data directory. When you restore an NMDA backup of InnoDB tables that are stored outside the MySQL data directory, if any `.ibd` file from the backup 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 435 provides details about setting the parameters for MySQL backup and restore operations.

Considerations for backups of MySQL replicated slave servers

NMDA supports the backup of a MySQL replicated slave server with MEB 3.11 or later.

Ensure that you meet the following requirements for NMDA backups of a MySQL replicated slave server:

- MEB 3.11 or later is installed with a supported MySQL version.
- You have configured the NMDA full and incremental backups as required of the MySQL replicated slave server. You have disabled the NMDA log file backups.
- You have enabled the MEB log backup mechanisms. Do not use the `--use-tts`, `--skip-binlog`, or `--skip-relay-log` option because these options stop MEB from backing up the binary and relay logs.

When you perform the NMDA full or incremental backups of the MySQL replicated slave server, NMDA does not store information about available logs in the backups. After you enable the MEB log backup mechanisms, MEB is responsible for managing the backups of the binary and relay logs.

[Performing MySQL recovery of a replicated slave server](#) on page 235 describes the NMDA restores of MySQL replicated slave server backups.

Oracle backup considerations

For Oracle client-side 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 444.

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 138 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 6](#), 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

About this task

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 476 and [Backup copies](#) on page 47 provide details.

The following example provides a sample RMAN script for a manual backup.

Example 6 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.
- [Oracle RMAN Commands](#) on page 475 describes RMAN commands.
- [Verifying backup information in NetWorker indexes](#) on page 180 describes the backup information that is stored in the NetWorker indexes.

RMAN scripts for scheduled backups

About this task

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 444. 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 476 provides details.

[Common NMDA parameters](#) on page 406 and [NMDA Oracle parameters](#) on page 444 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 138 provides information about automatic channel allocation.

The following example provides a sample RMAN script for a scheduled backup.

Example 7 RMAN script for a scheduled backup

The following RMAN script is for a scheduled backup of an entire Oracle database. 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';

  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.

Example 7 RMAN script for a scheduled backup (continued)

- *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.

i **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 that are enclosed in single quotes). [Canceling scheduled backups](#) on page 166 describes how to cancel a scheduled backup.

The remainder of the preceding scheduled backup script, starting with the first `allocate channel` command, is similar to the manual backup script in [Example 6](#) except the scheduled backup script excludes the `NSR_SERVER` parameter.

i **NOTICE** Do not set the `NSR_SERVER` parameter in a scheduled RMAN backup script. NMDA sets this parameter to the value specified in the Client resource for the Oracle scheduled backup, and you cannot override this value. However, you can set `NSR_SERVER` in the scheduled backup script when you manually test-run the script as described in [Testing RMAN scripts for scheduled backups](#) on page 138.

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.

About this task

To test the RMAN script, type the following command:

```
rman cmdfile 'script_name' send '"NSR_ENV=(NSR_SERVER=NetWorker_server_name)"'
```

where:

- *script_name* is the RMAN script file pathname.
- *NetWorker_server_name* is the name of the server that starts the backup.

Configuring automatic channel allocation

i **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 list all the NMDA parameters and their requirements for Oracle operations:

- Common NMDA parameters: [Table 38](#) on page 406
- NMDA Oracle parameters: [Table 43](#) on page 445

Example 8 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 8 Using the `configure channel` command with `parms` option for automatic channels (continued)

Example 9 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_DEBUG_LEVEL` parameter separately for each channel.

You specify a debugging level 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_DEBUG_LEVEL=9)'
```

- With Oracle version 11gR2 or later:

```
configure channel 2 device type 'sbt_tape' parms
'SBT_PARMS=(NSR_DEBUG_LEVEL=9)'
```

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 then 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)

[NMDA Parameters and Configuration File](#) on page 399 describes these parameters.

Example 10 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 480 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 11 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 480 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 that is run for the manual backup is as follows:

```
connect target sys/oracle@test;
run {
  backup format '%d_%U'
  tag tag_name
```

Example 11 Using automatic channels for backup copies (continued)

```
(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:

About this task

- If you use the configuration wizard:
 - Set `NSR_BUNDLING` to `TRUE` in the **Advanced 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 that are required for save set bundling, as described in [Table 6](#) on page 83. Ensure that you have configured the corresponding User Group resource on the NetWorker server, as described in [Verifying the NetWorker User Group resource](#) on page 82.
- 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 388 describes the `nsroraadmin` command. The command sets the parameter value in the NWORA resource file, as described in [NWORA resource file](#) on page 383.

2. Ensure that you have configured the NMDA scheduled backups according to the [Configuring NMDA backups](#) on page 76.
3. Ensure that the user group privileges for the root user or administrative user on the NMDA client include the privileges that are required for save set bundling, as described in [Table 6](#) on page 83. Ensure that you have configured the corresponding User Group resource on the NetWorker server, as described in [Verifying the NetWorker User Group resource](#) on page 82.
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 the following requirements:
 - The `ORACLE_SID` parameter is set in the NMDA configuration file, as described in [NMDA Oracle parameters](#) on page 444.
 - 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 386.

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 that is specified by the parameter in the NWORA resource file.

5. In a RAC system, ensure the required settings:

- 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.
- On UNIX or Linux, create a symbolic link that points the `/nsr/apps` directory on every RAC node to a file system pathname on a shared disk, which is accessible from all the nodes.
- On Windows, set `NSR_TMPDIR` in the `nwora.res` file on every RAC node to a file system pathname on a shared disk, which is accessible from all the 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

About this task

If you enable save set bundling as described in [Configuring save set bundling for scheduled Oracle backups](#) on page 141, 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 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 388 describes the `nsroraadmin` command. The command sets the parameter value in the NWORA resource file, as described in [NWORA resource file](#) on page 383.

- Complete steps 2 to 4 under [Configuring save set bundling for scheduled Oracle backups](#) on page 141.
- In a RAC system, ensure the required settings:
 - 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.
 - On UNIX or Linux, create a symbolic link that points the `/nsr/apps` directory on every RAC node to a file system pathname on a shared disk, which is accessible from all the nodes.
 - On Windows, set `NSR_TMPDIR` in the `nwora.res` file on every RAC node to a file system pathname on a shared disk, which is accessible from all the 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:

About this task

Procedure

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 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.

Preventing possible performance issues with Oracle Exadata backups

If an Oracle Exadata database server is running on Linux and the backup data is sent over an Infiniband connection, you can prevent any backup performance issues by setting the following parameter with the `parms` option for each backup channel in the RMAN script:

```
NSR_SOCKET_BUF_UNSET=YES
```

The following RMAN script includes the parameter setting for each channel:

```
run {
  allocate channel c1 type SBT parms 'ENV=(NSR_SOCKET_BUF_UNSET=YES)';
  allocate channel c2 type SBT parms 'ENV=(NSR_SOCKET_BUF_UNSET=YES)';

  backup database;

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

The same environment variable, `NSR_SOCKET_BUF_UNSET=YES`, must also be set on the NetWorker storage node host where the backup will be sent. You can set the environment variable in the NetWorker startup script (for example, `/etc/init.d/networker`) on UNIX and Linux systems, and in the system environment variables on Microsoft Windows. The storage node must be restarted in order for the environment variables to take effect. The NetWorker documentation provides details on setting environment variables.

This environment variable setting causes the NMDA and NetWorker software to not tune the socket buffer size, which results in better performance with an Infiniband connection.

SAP IQ backup considerations

The following SAP IQ backup topics require specific considerations:

- SAP IQ interfaces file
- SAP IQ selective backups
- SAP IQ transaction log backups

Review [SAP IQ full, incremental, incremental since full, and transaction log backups](#) on page 58 for considerations for the supported levels of NMDA SAP IQ backups.

Setting up the SAP IQ interfaces file

During the NMDA SAP IQ backup operations, the Open Client Server (OCS) library communicates with the SAP IQ server. To enable this communication, the SAP IQ `interfaces` file must exist in the `$SYBASE` directory.

Before you perform any NMDA SAP IQ backups, ensure that the SAP IQ `interfaces` file exists in the `$SYBASE` directory. You can run the `dscp` utility to generate the `interfaces` file. For example:

```
~/IQ-16_0/bin64> dscp
>> open
ok
Session 1 InterfacesDriver>> add SAPIQDB
Service: [ASE]
Transport Type: [tcp]
Transport Address: 10.10.10.10 2638
Transport Type: [tcp]
Transport Address:
Security Mechanism [] :
HA Failoverserver:
Retry Count:
Retry Delay:
Added SAPIQDB
Session 1 InterfacesDriver>> list all
Distinguish name: SAPIQDB
Server Entry Version: 1
Server Name: SAPIQDB
Server Service: Adaptive Server Enterprise
Server Status: 4 (Unknown)
Server Address:
  Transport Type: tcp
  Transport Address: 10.10.10.10 2638
```

You can run the `cat interfaces` command to display the content of the resulting `interfaces` file. For example:

```
cat interfaces
```

```
SAPIQDB
  master tcp ether 10.10.10.10 2638
  query tcp ether 10.10.10.10 2638
```


Configuring the SAP IQ selective backups

NMDA supports three types of selective backups of SAP IQ data: read-write selective backups, read-only selective backups, and all-inclusive selective backups.

To configure a selective backup, you must set the required parameters in the NMDA configuration file. [NMDA Parameters and Configuration File](#) on page 399 provides details about how to set parameters in the configuration file.

Set the `IQ_SELECTIVE_TYPE` parameter to one of the following values:

- **READWRITE**—This parameter value specifies to perform a read-write selective backup, which backs up the complete set of all the read-write database files in the SAP IQ database, without any read-only dbspaces or dbfiles.

A read-write selective backup supports all the backup levels. You can perform a full, incremental, or incremental since full backup of the read-write database files.

- **READONLY**—This parameter value specifies to perform a read-only selective backup, which backs up only the specified read-only dbspaces or read-only dbfiles or both from the SAP IQ database.

You must also specify the read-only data objects to back up by setting the `IQ_READONLY_DBSPACES` or `IQ_READONLY_DBFILES` parameter or both parameters.

A read-only selective backup supports only a full level backup. If you try to perform an incremental or incremental since full backup of the read-only data objects, the backup fails with an error message.

- **ALL_INCLUSIVE**—This parameter value specifies to perform a backup of both the read-write and read-only objects from the SAP IQ database.

An all-inclusive selective backup supports all three backup levels: full, incremental, and incremental since full.

[NMDA SAP IQ parameters](#) on page 457 provides details about the NMDA SAP IQ parameters.

Configuring the SAP IQ transaction log backups

NMDA supports the log backups of all the SAP IQ transaction logs and log archives. To complete an SAP IQ transaction log backup, NetWorker software performs a file system backup of the transaction logs and log archives.

Before you perform an SAP IQ transaction log backup, you must meet the following prerequisites:

- Point-in-time logging is enabled for the SAP IQ database.
- You have confirmed the pathnames of all the transaction logs that need to be backed up.
- Each transaction log that will be backed up (except the active log) has its respective log archive at the backup time. Otherwise, a subsequent point-in-time recovery will fail.
- The SAP IQ database is up and running.
- A previous full, cumulative, or incremental backup of the database exists. Otherwise, the transaction log backup will fail.
- The required parameters are set in the NMDA configuration file.

[NMDA Parameters and Configuration File](#) on page 399 provides details about how to set the required parameters in the NMDA configuration file.

To enable the transaction log backups, set the mandatory parameters:

- **NSR_BACKUP_LEVEL**—This parameter must be set to the value `txnlog` for a transaction log backup.

- `IQ_PIT_RESTORE_LOG_PATH`—This parameter specifies a list of pathnames of the transaction log files to back up, including the pathname of the active log.

Note: NMDA does not keep track of the log files that need to be backed up. You must ensure that all the required log pathnames are included in the `IQ_PIT_RESTORE_LOG_PATH` parameter setting.

To enable specific supported features of the transaction log backups, set the appropriate optional parameters in the NMDA configuration file:

- `IQ_PITR_LOG_RETENTION_PATH`—This parameter specifies a valid pathname of an existing local archive directory where the transaction log files and log archives are moved after an SAP IQ log backup. If this parameter is not set, the log files remain in the original location.
- `IQ_PITR_LOG_RETENTION_TIME`—This parameter specifies the number of days to retain the log files in the local archive directory specified by `IQ_PITR_LOG_RETENTION_PATH`, after the files are moved to that directory. `IQ_PITR_LOG_RETENTION_PATH` must be set to a valid directory pathname.

When `IQ_PITR_LOG_RETENTION_TIME` is not set or is set to 0, the log files are retained in the local archive directory until they are manually removed.

- `NSR_LOG_VOLUME_POOL`—This parameter specifies the name of the NetWorker volume pool to store the backup save sets of the SAP IQ transaction logs.
- `NSR_DEBUG_LEVEL`—This parameter must be set to a minimum value of 3 to enable the INFO level logging of the backed-up log files.

[NMDA SAP IQ parameters](#) on page 457 provides details about the NMDA SAP IQ parameters.

Sybase backup considerations

The following Sybase backup topics require specific considerations:

- Sybase backup best practices
- Sybase roles and permissions
- Sybase database backup verification
- Sybase incremental backups to specific volume pools
- Sybase transaction log backups in an emergency
- Sybase threshold procedure
- Sybase backups on HP-UX
- Sybase ASE backups of an instance that contains many databases
- 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 112 provides details.

Review [Sybase full, cumulative, and incremental backups](#) on page 61 for considerations for the supported levels of NMDA Sybase backups.

Best practices for Sybase backups

The NMDA Sybase backup performance can be significantly affected by the `NSR_CONCURRENCY_MODE` and parallelism settings.

Set the `NSR_CONCURRENCY_MODE` parameter to the appropriate value for the preferred backup:

- Set `NSR_CONCURRENCY_MODE` to the "stripe" (default) value to specify that NMDA backs up only one database at a time. The parallelism setting determines the number of stripes that NMDA uses for the database backup.

This setting works well in most cases, although it is best used when the average size of the database to back up is 250 MB or greater.

- Set `NSR_CONCURRENCY_MODE` to the "database" value to specify that NMDA backs up multiple databases at the same time and uses one stripe per database. The parallelism setting determines the maximum number of databases that NMDA backs up concurrently.

This setting is best used when there are two or more databases and the databases are either small in size (less than 250 MB) or running incremental backups. This mode uses more resources than the stripe backup mode for the same `NSR_PARALLELISM` setting.

The `NSR_CONCURRENCY_MODE` description in [Table 47](#) on page 464 provides details about the parameter setting.

The parallelism setting that the Sybase backup uses is the minimum of the `NSR_PARALLELISM`, client parallelism, and server parallelism settings. When you set the `NSR_PARALLELISM` parameter for a backup, you must balance the performance of the backup with the increase in load on the NetWorker server and client:

- When `NSR_CONCURRENCY_MODE` is set to "stripe", a good rule of thumb for the parallelism setting is as follows:

$$\text{Parallelism} = (\text{average size of database}) / (200 \text{ MB})$$

Most systems do not have a significant performance increase when NMDA uses more than 4 to 8 stripes. It is recommended to not use striping during incremental (log) backups because these backups tend to be small in size.

- When `NSR_CONCURRENCY_MODE` is set to "database", you can set the parallelism as high as 10 in most cases. Higher values can be used, but the ASE system must be configured to handle the additional load.

In most cases, performance does not increase significantly with a parallelism setting that is greater than 8 to 10. Also, because NMDA opens and backs up multiple databases at the same time, the load on the ASE instance is considerably greater than when NMDA backs up a single database with striping.

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.

About this task

The following table lists the Sybase roles and permissions that NMDA requires for performing the Sybase administrative actions.

Table 15 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	<code>nsrdasv</code> <code>nsrsybrc</code>

Table 15 Sybase roles and permissions (continued)

Role or permission	Action	NMDA command
SA_role, DBO	Run a database consistency check	<code>nsrsybcc</code>

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](#) 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.

About this task

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:

About this task

- `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](#).

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` for a manual backup or by the pool selection in the backup action for a scheduled backup. In the preceding example, NMDA stores the incremental backup of `sybdb` as follows:

- `SYBASE:/SERVER/sybdb`—Stored in the `NSR_DATA_VOLUME_POOL` pool or in the pool that is selected in the backup action.

- SYBASE:/SERVER/sybdb.1—Stored in the NSR_LOG_VOLUME_POOL pool.

Note: If the NSR_DATA_VOLUME_POOL parameter or pool selection is not set, then NMDA backs up the metadata to the default volume pool, which is not necessarily the pool named Default.

Sybase transaction log backups in an emergency

For manual and scheduled Sybase backups, NMDA backs up the Sybase transaction log and removes the inactive portion through 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 the parameter setting can prevent you from restoring the ASE server or specific databases.

The following table describes the backup behavior for the different NSR_DUMP_LOG_OPT settings when the database and the transaction log are on the same or separate devices.

Note: Due to a Sybase limitation, NMDA does not support certain settings of the NSR_DUMP_LOG_OPT parameter 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” settings.

Table 16 Backups with NSR_DUMP_LOG_OPT

NSR_DUMP_LOG_OPT parameter setting	Database and log on separate devices	Database and log on the same device
For read/write databases:		
“no_log”	<ol style="list-style-type: none"> 1. Truncates the transaction log without logging the transaction. 2. Backs up the database. 	<ol style="list-style-type: none"> 1. Truncates the transaction log without logging the transaction. 2. Backs up the database.
“no_truncate”	Backs up the transaction log without truncating the log.	Not supported. Skips the backup.
“truncate_only”	<ol style="list-style-type: none"> 1. Truncates the transaction log. 2. Backs up the database. 	<ol style="list-style-type: none"> 1. Truncates the transaction log. 2. Backs up the database.
No setting	Backs up the transaction log.	<ol style="list-style-type: none"> 1. Backs up the database. 2. Truncates the transaction log.
For read-only databases:		
“no_log”	Backs up the database.	Backs up the database.
“no_truncate”	Not supported. Skips the backup.	Not supported. Skips the backup.
“truncate_only”	Displays an error message stating that the setting is	Displays an error message stating that the setting is

Table 16 Backups with NSR_DUMP_LOG_OPT (continued)

NSR_DUMP_LOG_OPT parameter setting	Database and log on separate devices	Database and log on the same device
	invalid for read-only databases.	invalid for read-only databases.
No setting	Backs up the database.	Backs up the database.

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.

About this task

NMDA supports the threshold procedure through the `threshold.sql` file, as described in [Sample threshold procedure](#) on page 151.

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.

i **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 restart the system and start the Sybase Backup Server. [NMDA Parameters and Configuration File](#) on page 399 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 111 describes probe-based backups.

Sample threshold procedure

Use the sample threshold procedure that is described in this topic to implement transaction log backups to free the log space. Edit the sample threshold procedure to suit the environment.

About this task

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:

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 then 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:

About this task

Procedure

1. In the shell where the Sybase backup server runs, set the variable as the Sybase user:

```
LD_PRELOAD=$SYBASE/$SYBASE_ASE/lib/libnrsyb.so
```

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.

Configuring shrink log operations during Sybase backups

Use one of the following methods to specify that NMDA perform a shrink log operation during a Sybase 15.7 database backup:

- 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 that is configured without the wizard.

[NSR_DUMP_DATA_OPT](#) provides details.

Configuring Sybase concurrent backups

You must enable the support for concurrent full and incremental backups by running the following `sp_configure` command from Sybase `isql`:

```
sp_configure 'enable concurrent dump tran', 1
```

To ensure that the backups run concurrently, you must start the incremental backup after the full backup has started. If you start the incremental backup first, the full backup waits until the incremental backup completes.

Note: If you start an incremental backup at the point when the full backup is backing up the log (the last short phase of the full backup), the incremental backup does not run concurrently but waits until the full backup completes. During a subsequent restore, the incremental backup can only be applied to that full backup.

During the concurrent backups, NMDA Sybase uses lock files that are located in the `nrsr/apps/tmp` directory. The lock files are removed after a backup completes. If lock files remain in the directory after a failed backup, the files are removed during the next backup of the database.

Configuring Sybase cumulative backups

You must enable the cumulative backup support for each database by running the following `sp_dboption` command from Sybase `isql`:

```
sp_dboption database_name, 'allow incremental dumps', true
```

You can enable Sybase cumulative and incremental backups of a database to run concurrently by running the following `sp_configure` command from Sybase `isql`:

```
sp_configure 'enable concurrent dump tran', 1
```

Note: Sybase does not support the concurrent running of full and cumulative backups of a database.

Complete the required setting to configure a Sybase cumulative backup:

- For a manual Sybase backup, set `NSR_BACKUP_LEVEL=cumulative` or `NSR_BACKUP_LEVEL=1` in the NMDA configuration file.
- For a scheduled Sybase backup, set the backup level to Cumulative Incr in the Policy Action Wizard.

After you enable cumulative backups, perform an initial full Sybase backup as a basis for any subsequent cumulative backups.

The cumulative backup level appears as level 1 in the `mminfo` command output from the NetWorker index database.

[Example 34](#) shows how to disable the restore of Sybase cumulative backups. Sybase cumulative restores are enabled by default.

Orchestrated Application Protection backup considerations

Orchestrated Application Protection provides the database/application backup operations through the NMDA program `nsroappbackup`, which leverages the native backup utilities that the database/application software provides.

Orchestrated Application Protection also leverages the Data Domain BoostFS software to back up the data to a DD Boost backup device that is configured by the NetWorker server. The backup device must be a DD Boost device, as configured by the NetWorker server through an IP configuration. Orchestrated Application Protection does not support other device types or Fiber Channel configurations.

The configuration of Orchestrated Application Protection backups requires specific considerations, based on the application and type of database to be backed up. The following topics provide details about the backup considerations for Orchestrated Application Protection.

[NMDA features specific to Orchestrated Application Protection](#) on page 66 provides information about the limitations of Orchestrated Application Protection backups.

Backup shell script settings for supported backups

For an Orchestrated Application Protection backup, you must set all the required NMDA Orchestrated Application Protection parameters in the NMDA configuration file, which includes the proper XML formats as described in [NMDA configuration file](#) on page 400.

Orchestrated Application Protection supports the backup levels `full`, `incr`, and `txnlog`. You can use a different backup shell script for each backup level that you want to run. The `nsroappbackup`

command line option `-l` indicates the backup level. The following table describes the supported command line options.

Table 18 Command line options for Orchestrated Application Protection backups

Command line options	Description	Default and valid values
<code>-l backup_level</code>	<p>Specifies the level of the backup as a full, incremental, or transaction log backup.</p> <p>Optional.</p> <p>Note: The type of database or application can affect which backup levels are supported.</p>	<ul style="list-style-type: none"> full = full backup (default). incr = incremental backup. txnlog = transaction log backup.
<code>-o pg_p_opt="%p" -o pg_f_opt="%f"</code>	<p>For a PostgreSQL WAL backup only. Specify these two options to meet the WAL log file backup requirement of the PostgreSQL server.</p> <p>Note: When both <code>-o pg_p_opt</code> and <code>-o pg_f_opt</code> options are set for a transaction log backup, the <code>NSR_BACKUP_SCRIPT</code> parameter setting is ignored in the <code><TXNLOG></code> section of the NMDA configuration file.</p> <p>Mandatory for a PostgreSQL transaction log backup.</p>	<ul style="list-style-type: none"> Undefined (default). The <code>%p</code> and <code>%f</code> options are updated by the PostgreSQL server during the backup to include the required file pathnames.
<code>-z configuration_file_path</code>	<p>Specifies the NMDA configuration file that contains the parameter settings and command line options for the backup.</p> <p>Note: The NMDA configuration file must be owned by the database user that runs the backup, or the user that is specified by <code>NSR_OS_USER</code>. The group users and other users should not have permission to access the file.</p> <p>Mandatory.</p>	<ul style="list-style-type: none"> Undefined (default). Valid complete pathname of the NMDA configuration file.

To enable each backup level, you must specify the backup shell script pathname in the `NSR_BACKUP_SCRIPT` parameter setting in each backup level subsection in the configuration file. For example, the following backup section in the configuration file includes the `NSR_BACKUP_SCRIPT` parameter setting in each backup level subsection:

```
<BACKUP>
  <FULL>
    <NSR_BACKUP_SCRIPT> full_backup_script_pathname </NSR_BACKUP_SCRIPT>
  </FULL>
  <INCR>
    <NSR_BACKUP_SCRIPT> incr_backup_script_pathname </NSR_BACKUP_SCRIPT>
  </INCR>
  <TXNLOG>
    <NSR_BACKUP_SCRIPT> txnlog_backup_script_pathname </NSR_BACKUP_SCRIPT>
```

```
</TXNLOG>
</BACKUP>
```

Note: The <FULL> subsection settings are required for all backups.

The `nsroappbackup` program performs a backup by running the backup shell script for the current backup level. If you do not want to perform a certain level of backup, you can omit the backup shell script setting in that backup level subsection.

It is recommended that the backup shell script returns the value 0 if the script completes successfully, and returns -1 if the backup fails or encounters an error. If the backup shell script returns a nonzero value, the `nsroappbackup` program assumes that the backup failed.

In the backup shell script, you must use `$OAPP_MOUNT_DIR/` as the backup target directory pathname. You can also use the `$OAPP_MOUNT_DIR` variable in the `NSR_BACKUP_SCRIPT` parameter setting in the configuration file. Orchestrated Application Protection does not automatically create a subdirectory under `$OAPP_MOUNT_DIR/`. If you want a subdirectory to be created under `$OAPP_MOUNT_DIR/`, you can create the subdirectory by a separate command in the same backup shell script. As an alternative, the backup utility can create the subdirectory.

For example, the following PostgreSQL backup shell script includes the backup utility command and a separate copy command. The script exits the processing if the executed command fails, and prints the command and returned value:

```
#!/bin/sh
( set -x; /opt/postgresql/pg94/bin/pg_basebackup --pgdata=$OAPP_MOUNT_DIR/
basebackup )
rc=$?
if [ $rc == 0 ]; then
    echo "Continued processing with the returned value $rc."
else
    echo "Exited processing with the returned value $rc."
    exit $rc
fi

( set -x; /bin/cp /home/postgres/postgresql.conf $OAPP_MOUNT_DIR/ )
rc=$?
if [ $rc == 0 ]; then
    echo "Continued processing with the returned value $rc."
else
    echo "Exited processing with the returned value $rc."
    exit $rc
fi
```

OS user that runs the `nsroappbackup` command

It is recommended that the database/application user run the `nsroappbackup` command.

The specified user must own the Orchestrated Application Protection backup script file and the configuration file. Other users or groups must not have any permissions to access these files.

You can log in as the specified user and run the following commands to verify the required file permissions and file path permissions:

```
ls -l <backup_script_file_full_pathname>
ls -l <configuration_file_full_pathname>
```

The specified user must have the read, write, and execute permissions for the backup script file, and the read and write permissions for the configuration file. The specified user can access the

backup script file and configuration file, for example by running the following commands to display the file contents:

```
cat <backup_script_file_full_pathname>
cat <configuration_file_full_pathname>
```

For a scheduled backup or when the root user runs the `nsroappbackup` command, the `NSR_OS_USER` parameter must specify the username of the database/application user. Then the `nsroappbackup` command runs the backup shell script as the specified user.

Database authentication

If you enabled the database authentication, the database backup utility might require the password to connect to the database when the backup starts. In this case, Orchestrated Application Protection supports the storage of the encrypted password in the configuration file and the password transfer to the backup utility when Orchestrated Application Protection receives the password prompt.

You can add the password in the configuration file by running the `nsrdaadmin` command. Then you can run the backup command that you plan to include in the backup script file to obtain the password prompt string, and set the password prompt string with the `USER_PSWD_PROMPT` parameter.

Save set names

Orchestrated Application Protection generates the name of the backup save set by combining into one string the values of the parameters `NSR_DATABASE_TYPE`, `NSR_INSTANCE_NAME`, `NSR_BACKUP_NAME`, and `NSR_BACKUP_LEVEL`, in that order.

For example, the following parameter settings in the configuration file will produce the save set name `mydb_myinstance_mybackup_full` for the full backup:

```
<NSR_DATABASE_TYPE> mydb </NSR_DATABASE_TYPE>
<NSR_INSTANCE_NAME> myinstance </NSR_INSTANCE_NAME>
<NSR_BACKUP_NAME> mybackup </NSR_BACKUP_NAME>
<NSR_BACKUP_LEVEL> full </NSR_BACKUP_LEVEL>
```

To obtain information about a backup save set, you can run the following command:

```
mminfo -vV -s NetWorker_server_hostname -c NetWorker_client_hostname
```

MongoDB backup considerations

Orchestrated Application Protection provides the MongoDB database backup operations through the NMDA program `nsroappbackup`, which leverages the native MongoDB backup utility.

MongoDB supports only full database backups. You must specify the full backup level when you configure a MongoDB backup. For a manual backup, use the `-l full` option. For a scheduled backup, select the Full level in the policy action.

Ensure that you review the topic [Orchestrated Application Protection backup considerations](#) on page 153. All the backup considerations from that topic also apply to the configuration of MongoDB backups.

The following MongoDB backup topics require specific considerations:

- MongoDB backup shell script
- Database authentication
- Completing the MongoDB backup configuration

MongoDB backup shell script

MongoDB provides the `mongodump` backup utility. You must include the `mongodump` utility in the backup shell script for the NMDA MongoDB backups.

For example, the following backup shell script for full MongoDB backups includes the backup utility command, `mongodump`. The script exits the processing if the executed command fails, and prints the command and returned value:

```
#!/bin/sh
( set -x; mongodump --db mydb --out $OAPP_MOUNT_DIR/backup_data )
rc=$?
if [ $rc == 0 ]; then
    echo "Continued processing with the returned value $rc."
else
    echo "Exited processing with the returned value $rc."
    exit $rc
fi
```

You must specify the backup shell script pathname in the FULL subsection of the BACKUP section in the configuration file. [Orchestrated Application Protection backup considerations](#) on page 153 provides more details.

The MongoDB server documentation provides more details about the `mongodump` backup utility.

Database authentication

If you enabled authentication, the backup utility requires the username and password to connect to the database when the backup starts. Orchestrated Application Protection supports the storage of the encrypted password in the configuration file and the password transfer to the backup utility when Orchestrated Application Protection receives the password prompt.

Add the password in the configuration file by running the `nsrdaadmin` command. Then run the backup command that you plan to include in the backup script file, to obtain the password prompt string. For example, run the following command at the command line:

```
mongodump --db admin -u admin --out $OAPP_MOUNT_DIR/backup_data
Enter password:
```

Here, "Enter password:" is the password prompt string. In the configuration file, set the `USER_PSWD_PROMPT` parameter to this password prompt string. [NMDA parameters for Orchestrated Application Protection backups](#) on page 451 provides details on setting the backup parameters in the configuration file.

Completing the MongoDB backup configuration

To enable the NMDA MongoDB backups through the Orchestrated Application Protection feature, ensure that the required environment variables and NMDA parameters are set.

Ensure that the required environment variables are set through the `NSR_ENV_LIST` parameter in the configuration file, or set the variables in the backup shell script for the MongoDB backup. The MongoDB online documentation provides more details about the required environment variables.

Ensure that all the required NMDA Orchestrated Application Protection parameters are set in the NMDA configuration file. For database operations with the Orchestrated Application Protection feature, the NMDA configuration file must include the proper XML formats as described in [NMDA configuration file](#) on page 400.

[NMDA Orchestrated Application Protection parameters](#) on page 451 provides complete details about the parameter settings in the NMDA configuration file.

MySQL backup considerations

Orchestrated Application Protection provides the MySQL database backup operations through the NMDA program `nsroappbackup`, which leverages the native MySQL backup utility.

[MySQL operations through Orchestrated Application Protection](#) on page 158 provides a comparison of differences in the NMDA MySQL operations that are performed through the MySQL Enterprise Backup and the Orchestrated Application Protection feature.

NMDA MySQL through Orchestrated Application Protection supports the full database backups and the transaction log backups:

- You must specify the full backup level when you configure a MySQL full backup. For a manual backup, use the `-l full` option. For a scheduled backup, select the Full level in the policy action.
- You must specify the Logs Only backup level when you configure a MySQL transaction log backup. For a manual backup, use the `-l txnlog` option. For a scheduled backup, select the Logs Only level in the policy action.

Ensure that you review the topic [Orchestrated Application Protection backup considerations](#) on page 153. All the backup considerations from that topic also apply to the configuration of MySQL backups.

The following MySQL backup topics require specific considerations:

- MySQL operations through Orchestrated Application Protection
- MySQL full backup shell script
- MySQL transaction log backup shell script
- Database authentication
- Completing the MySQL backup configuration

MySQL operations through Orchestrated Application Protection

NMDA supports two methods of MySQL data protection. You can perform the traditional NMDA MySQL backup and restore operations through MySQL Enterprise Backup (MEB), or you can perform the NMDA MySQL operations through Orchestrated Application Protection. The configuration requirements and supported features of the NMDA MySQL operations depend on whether you use MEB or Orchestrated Application Protection.

The following table provides a comparison of the supported features of NMDA MySQL operations that you perform through MEB and through Orchestrated Application Protection. Review the table to determine the features that match your particular requirements and whether MySQL operations through Orchestrated Application Protection provide the best solution for your environment.

Table 19 NMDA MySQL operations through MEB and Orchestrated Application Protection

Feature support	NMDA MySQL operations through MEB	NMDA MySQL operations through Orchestrated Application Protection
Backup configuration wizard	Supported.	Not supported.
Recovery configuration wizard	Supported.	Not supported.
Number of steps in recovery	One step.	Two steps:

Table 19 NMDA MySQL operations through MEB and Orchestrated Application Protection (continued)

Feature support	NMDA MySQL operations through MEB	NMDA MySQL operations through Orchestrated Application Protection
		<ol style="list-style-type: none"> 1. Retrieve the backup into a local directory. 2. Run the MySQL restore utility.
MySQL Enterprise Backup (MEB)	Required.	Not required.
Supported backup device	Any NetWorker device.	Data Domain device only.
Supported recovery device	Any NetWorker device.	Data Domain device only.
Data Domain connectivity	Fibre Channel (FC) or IP.	IP.
Backup granularity	MySQL instance, database, and table.	Depends on user's backup script.
Full backup	Supported.	Supported.
Incremental backup	Supported.	Not supported.
Transaction log backup	Supported.	Supported.
MySQL backup utility	mysqlbackup	mysqldump

MySQL full backup shell script

MySQL provides the `mysqldump` backup utility. You must include the `mysqldump` utility in the backup shell script for the NMDA MySQL backups.

For example, the following backup shell script for MySQL full backups includes the backup utility command, `mysqldump`. The script exits the processing if the executed command fails, and prints the command and the returned nonzero value:

```
#!/bin/sh
( set -x; /usr/bin/mysqldump --all-databases --result-file=$OAPP_MOUNT_DIR/
full_mysql_dump.sql)
rc=$?
if [ $rc == 0 ]; then
    echo "Continued processing with the returned value $rc."
else
    echo "Exited processing with the returned value $rc."
    exit $rc
fi
```

You must specify the backup shell script pathname in the FULL subsection of the BACKUP section in the configuration file. [Orchestrated Application Protection backup considerations](#) on page 153 provides more details.

The MySQL documentation provides more details about the `mysqldump` backup utility.

MySQL transaction log backup shell script

You can back up the MySQL transaction logs by copying the binary log files into the backup mount point.

For example, the following backup shell script for MySQL transaction log (Logs Only) backups includes the `cp` command to copy the MySQL binary logs into the backup mount point. The script exits the processing if the executed command fails, and prints the command and the returned nonzero value:

```
#!/bin/sh
( set -x; /bin/cp /var/lib/mysql/mysql-bin.* $OAPP_MOUNT_DIR )
rc=$?
if [ $rc == 0 ]; then
    echo "Continued processing with the returned value $rc."
else
    echo "Exited processing with the returned value $rc."
exit $rc
fi
```

You must specify the backup shell script pathname in the `TXNLOG` subsection of the `BACKUP` section in the configuration file. [Orchestrated Application Protection backup considerations](#) on page 153 provides more details.

Database authentication

If you enabled authentication, the backup utility requires the username and password to connect to the database when the backup starts. Orchestrated Application Protection supports the storage of the encrypted password in the configuration file and the password transfer to the backup utility when Orchestrated Application Protection receives the password prompt.

Add the password in the configuration file by running the `nsrdaadmin` command. Then run the backup command that you plan to include in the backup script file, to obtain the password prompt string. For example, run the following command at the command line:

```
mysqldump -u admin -p -all-database --result-file=$OAPP_MOUNT_DIR/
full_mysql_dump.sql
Enter password:
```

Here, "Enter password:" is the password prompt string. In the configuration file, set the `USER_PSWD_PROMPT` parameter to this password prompt string. [NMDA parameters for Orchestrated Application Protection backups](#) on page 451 provides details on setting the backup parameters in the configuration file.

Completing the MySQL backup configuration

To enable the NMDA MySQL backups through the Orchestrated Application Protection feature, ensure that the required environment variables and NMDA parameters are set.

Ensure that the required environment variables are set through the `NSR_ENV_LIST` parameter in the configuration file, or set the variables in the backup shell script for the MySQL backup. The MySQL online documentation provides more details about the required environment variables.

Ensure that all the required NMDA Orchestrated Application Protection parameters are set in the NMDA configuration file. For database operations with the Orchestrated Application Protection feature, the NMDA configuration file must include the proper XML formats as described in [NMDA configuration file](#) on page 400.

[NMDA Orchestrated Application Protection parameters](#) on page 451 provides complete details about the parameter settings in the NMDA configuration file.

PostgreSQL backup considerations

Orchestrated Application Protection provides the PostgreSQL database backup operations through the NMDA program `nsroappbackup`, which leverages the native PostgreSQL backup utilities.

The NetWorker full and transaction log backups are mapped to specific types of PostgreSQL backups, as shown in the following table. You must specify the correct NetWorker backup level when you configure a PostgreSQL backup.

Table 20 Mapping of NetWorker level backups to PostgreSQL backups

NetWorker backup level	Type of PostgreSQL backup
Full backup (full)	Dump or base backup
Transaction log backup (txnlog)	WAL archiving or archive backup

Ensure that you review the topic [Orchestrated Application Protection backup considerations](#) on page 153. All the backup considerations from that topic also apply to the configuration of PostgreSQL backups.

The following PostgreSQL backup topics require specific considerations:

- PostgreSQL backup shell script
- PostgreSQL full backup
- PostgreSQL transaction log backup
- Database authentication
- Registering the PostgreSQL archive command
- Completing the PostgreSQL backup configuration

[PostgreSQL full and transaction log backups](#) on page 68 provides considerations for the supported levels of NMDA PostgreSQL backups.

PostgreSQL backup shell script

Select the correct PostgreSQL backup utility to include in the backup shell script. PostgreSQL provides two types of backup utilities:

- The `pg_dump` and `pg_dumpall` utilities can work with the WAL archiving disabled environment. If you will not perform a WAL backup, use `pg_dump` or `pg_dumpall` as the full backup utility.
- The `pg_basebackup` utility can work with the WAL archiving enabled environment. If you will perform a WAL backup, use `pg_basebackup` as the full backup utility.

The PostgreSQL server documentation provides more details about using the backup utilities `pg_dump`, `pg_dumpall`, and `pg_basebackup`, and details about WAL archiving.

You can create the backup shell script for the full level backup. To enable the backup of the WAL segment file, PostgreSQL requires you to register the `nsroappbackup` command in the `postgresql.conf` file. [Registering the PostgreSQL archive command](#) on page 163 provides more details.

PostgreSQL full backup

Select the correct backup utility for the corresponding PostgreSQL environment setting, and create the full backup shell script file to use with that utility to perform the dump or base backup.

You must specify the backup shell script pathname in the FULL subsection of the BACKUP section in the configuration file. You must specify the full backup level by using the `-l full` option for a manual backup or by selecting the Full level in the policy action for a scheduled backup.

[Orchestrated Application Protection backup considerations](#) on page 153 provides more details.

PostgreSQL transaction log backup

To enable the PostgreSQL WAL archiving, you must register the `nsroappbackup` program with its required command line options through the `archive_command` setting in the `postgresql.conf` file.

The transaction log backup (with backup level `txnlog`) can share the same configuration file with a full backup because the parameters `NSR_DATABASE_TYPE`, `NSR_INSTANCE_NAME`, and `NSR_BACKUP_NAME` have same values for both levels of backup.

For example, the `postgresql.conf` file includes the following command settings for the transaction log backups:

```
archive_command = '/usr/sbin/nsroappbackup -o pg_p_opt="%p" -o pg_f_opt="%f" -l
txnlog -z configuration_file_path'
```

Note:

The PostgreSQL WAL segment works with the PostgreSQL transaction. Only the operations inside the transactions can be logged into the WAL segment files. The WAL archiving restore requires the transaction information inside the WAL segment file to find the time point for the point-in-time restore. Without the transaction information, the WAL archiving restore and the point-in-time restore will not complete properly.

The backup behavior of PostgreSQL WAL archiving is predefined by the `nsroappbackup` program. You do not need to specify a backup shell script for this type of backup.

The PostgreSQL server schedules the PostgreSQL WAL archiving. In the online PostgreSQL document at www.postgresql.org, the "Continuous Archiving and Point-in-Time Recovery (PITR)" section in the "Backup and Restore" chapter and the "Write Ahead Log" section in the "Server Configuration" chapter provide more details about how to configure the WAL archiving and schedule.

Database authentication

If you enabled authentication, the backup utility requires the username and password to connect to the database when the backup starts. Orchestrated Application Protection supports the storage of the encrypted password in the configuration file and the password transfer to the backup utility when Orchestrated Application Protection receives the password prompt.

Add the password in the configuration file by running the `nsrdaadmin` command. Then run the backup command that you plan to include in the backup script file, to obtain the password prompt string. For example, run the following command at the command line:

```
/usr/bin/pg_basebackup -h localhost --pgdata=/tmp/basebackup29
Password:
```

Here, "Password:" is the password prompt string. In the configuration file, set the `USER_PSWD_PROMPT` parameter to this password prompt string. [NMDA parameters for Orchestrated Application Protection backups](#) on page 451 provides details on setting the backup parameters in the configuration file.

Registering the PostgreSQL archive command

To enable the PostgreSQL WAL backups with Orchestrated Application Protection, you must register the `nsroappbackup` program with the required command line options through the `archive_command` setting in the `postgresql.conf` file.

In the `archive_command` setting, specify the `nsroappbackup` command and its command line options. For example, the `postgresql.conf` file can include the following `archive_command` setting:

```
archive_command = '/usr/sbin/nsroappbackup -o pg_p_opt="%p" -o pg_f_opt="%f" -l txnlog -z configuration_file_path'
```

Note:

- The options `-o pg_p_opt="%p" -o pg_f_opt="%f" -l txnlog` indicate that the operation is a PostgreSQL WAL backup.
- The backup and restore must use the same `NSR_DATABASE_TYPE`, `NSR_INSTANCE_NAME`, and `NSR_BACKUP_NAME` parameter settings in the configuration file so that the restore can locate the WAL segment files that belong to the backup.

Table 18 on page 154 describes the supported command line options that you can include in the `archive_command` setting.

NOTICE

The registered archive command will be run by the same user as the PostgreSQL server.

To enable WAL (write ahead log) archiving, set the following parameters in the `postgresql.conf` file:

- Set `wal_level` to `archive` (or `hot_standby`).
- Set `archive_mode` to `on`.

The PostgreSQL online documentation provides more details.

Completing the PostgreSQL backup configuration

To enable the NMDA PostgreSQL backups through the Orchestrated Application Protection feature, ensure that the required environment variables and NMDA parameters are set.

Ensure that the required environment variables are set through the `NSR_ENV_LIST` parameter in the configuration file, or set the variables in the backup shell script for the PostgreSQL backup. The PostgreSQL online documentation provides more details about the required environment variables.

Ensure that all the required NMDA Orchestrated Application Protection parameters are set in the NMDA configuration file. For database operations with the Orchestrated Application Protection feature, the NMDA configuration file must include the proper XML formats as described in [NMDA configuration file](#) on page 400.

[NMDA Orchestrated Application Protection parameters](#) on page 451 provides complete details about the parameter settings in the NMDA configuration file.

CHAPTER 3

Backup Procedures

This chapter includes the following topics:

- [Performing scheduled backups](#)..... 166
- [Performing manual backups](#)..... 168
- [Verifying backup information in NetWorker indexes](#)..... 180
- [Synchronizing backup catalogs and deleting backups](#)..... 185

Performing scheduled backups

A scheduled backup is the type of backup that the NetWorker server initiates according to a configured backup schedule.

About this task

Complete the following procedures to prepare for a scheduled backup.

Procedure

1. Ensure that the required backup configurations are in place as described in [Configuring NMDA backups](#) on page 76.
2. Run a test scheduled backup manually as described in [Testing scheduled backups](#) on page 166.

At the end of a scheduled backup, the software automatically backs up specific additional files for the application as described in [Disaster Recovery](#) on page 273.

Testing scheduled backups

Use NMC to manually test the scheduled backup of a selected backup workflow. This test backup confirms that the scheduled backup runs as expected.

Procedure

1. In the web browser, go to the URL `http://NMC_server_name:9000`, and then log in to the NMC console server.

The default NMC connection port is 9000. If you configured a different port, then use the different port number.

2. In the NMC **Enterprise** window, right-click the NetWorker server name, and then select **Launch Application** to open the **Administration** window.

The *NetWorker Administration Guide* and NMC online help describe how to use the NMC interface.

3. In the **Monitoring** window, right-click the workflow name for the backup, and then select **Start** from the list box.


The scheduled backup starts after you click **Yes** in the confirmation box.

Instead of starting only one workflow, you can start all the workflows in a particular policy by right-clicking the policy name and selecting **Start**.

Canceling scheduled backups

You can cancel a scheduled backup by using the following procedures.

About this task

 **NOTICE** If you cancel a backup, some of the backed-up data might not be recoverable. To restart a canceled scheduled backup, follow the instructions in [Restarting failed scheduled backups](#) on page 167.

Procedure

1. Open the NMC **Administration** window.
2. Select the backup workflow in the **Monitoring** window.

- For an Oracle backup, ensure that the Oracle user has the required privilege for backup deletion, as described in [Table 6](#) on page 83.

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.

- Right-click the backup workflow, and then select **Stop**.

The scheduled backup stops after you click **Yes** in the confirmation box.

Note:

When you cancel a Sybase backup from NMC, the Sybase Backup Server logs the following type of error message, which is also displayed in NMC. You can ignore this message:

```
24146:nrsdasv: Error from server SYBASE157_BS: Msg 412402, Level 2, State 1
Backup Server: 4.124.2.1: Archive API error for
device='nrsryb::master..1./nsr/apps/tmp/sybtmp_14662.txt::000': Vendor
application name=EMC, Library version=200, API routine=syb_write(), Message=
```

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:
 - Use the **Advanced Options** page in the Policy Action Wizard in NMC to set **Retries** to a value greater than zero in the backup action that is created for the scheduled backup.
 - In the corresponding workflow, set the **Restart Window** as required:
 - For a Lotus checkpoint restart (CPR) backup that restarts within the 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 canceled or failed backup from the beginning, right-click the workflow in the NMC GUI, and select **Start**.
- To manually restart a Lotus CPR backup from the point-of-failure, right-click the workflow in the NMC GUI, and select **Restart**.
- If the first backup fails in a restarted scheduled backup but the restarted backup succeeds, then information about failed save sets might remain in the **Failed** area of the NMC window. The failed save set information only indicates that the initial backup attempt failed and does not reflect the final backup result.

Monitoring scheduled backups

You can monitor a scheduled backup and the activities for specific backup policies, workflows, and actions by using the **Monitoring** window in NMC.

About this task

Note: The location of information in the **Monitoring** window might vary in different NMC releases.

The **Policies** pane and **Actions** pane in the **Monitoring** window display backup status information during and after a scheduled backup.

The *NetWorker Administration Guide* provides details about viewing scheduled backup information 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.

About this task

i **NOTICE** Ensure that you back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Perform a manual backup by using the appropriate procedure from the following table.

Table 21 Manual backup procedures

Database or application	Manual backup procedures
DB2	<ul style="list-style-type: none"> • Performing DB2 manual backups with the db2 backup command on page 168 • Performing DB2 manual backups with the DB2 GUI on page 169
Informix	<ul style="list-style-type: none"> • Performing Informix manual backups with the onbar command on page 170
Lotus	<ul style="list-style-type: none"> • Performing Lotus manual backups with the nsrdasv command on page 170 • Performing Lotus manual backups with NetWorker User for Lotus on page 171
MySQL	<ul style="list-style-type: none"> • Performing MySQL manual backups with the nsrdasv command on page 173
Oracle	<ul style="list-style-type: none"> • Performing Oracle manual backups with the rman command on page 174 • Performing Oracle manual backups with Oracle Enterprise Manager on page 174
SAP IQ	<ul style="list-style-type: none"> • Performing SAP IQ manual backups with the nsrdasv command on page 175
Sybase	<ul style="list-style-type: none"> • Performing Sybase manual backups with the nsrdasv command on page 176 • Performing Sybase manual backups with NetWorker User for Sybase on page 176 <p>i Note: Before you perform a Sybase manual backup, perform a database consistency check according to Performing Sybase database consistency checks before backups on page 175.</p>
Orchestrated Application Protection (MongoDB, MySQL, PostgreSQL)	<ul style="list-style-type: none"> • Performing Orchestrated Application Protection manual backups with the nsroappbackup command on page 178

Performing DB2 manual backups with the db2 backup command

About this task

Before you perform the manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

You can run the appropriate `db2 backup` command to perform a DB2 manual backup from the command line.

Procedure

1. Log in to the DB2 client host as the DB2 operating system user.
2. Run the `db2 backup db sample load` 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 a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Performing DB2 manual backups with the DB2 GUI

About this task

Before you perform the manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

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_path`. 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 a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Performing Informix manual backups with the onbar command

About this task

Before you perform the manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

You can run the appropriate `onbar` command to perform an Informix manual backup from the command line.

Procedure

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. [NMDA Parameters and Configuration File](#) on page 399 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 a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Performing Lotus manual backups with the nsrdasv command

About this task

Before you perform the manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

You can run the appropriate `nsrdasv` command to perform a Lotus manual backup from the command line.

Procedure

1. To specify the Lotus directories or the files to back up, set either the `NSR_BACKUP_PATHS` or `NSR_BACKUP_LOTUS_DIR` parameter in the NMDA configuration file. Do not set both parameters simultaneously. The following descriptions provide more details:
 - [NSR_BACKUP_LOTUS_DIR](#)
 - [NSR_BACKUP_PATHS](#)
2. Set the `NSR_NOTES_INI_PATH` parameter in the NMDA configuration file, as described in [NSR_NOTES_INI_PATH](#).
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 a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Performing Lotus manual backups with NetWorker User for Lotus

About this task

Before you perform a manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

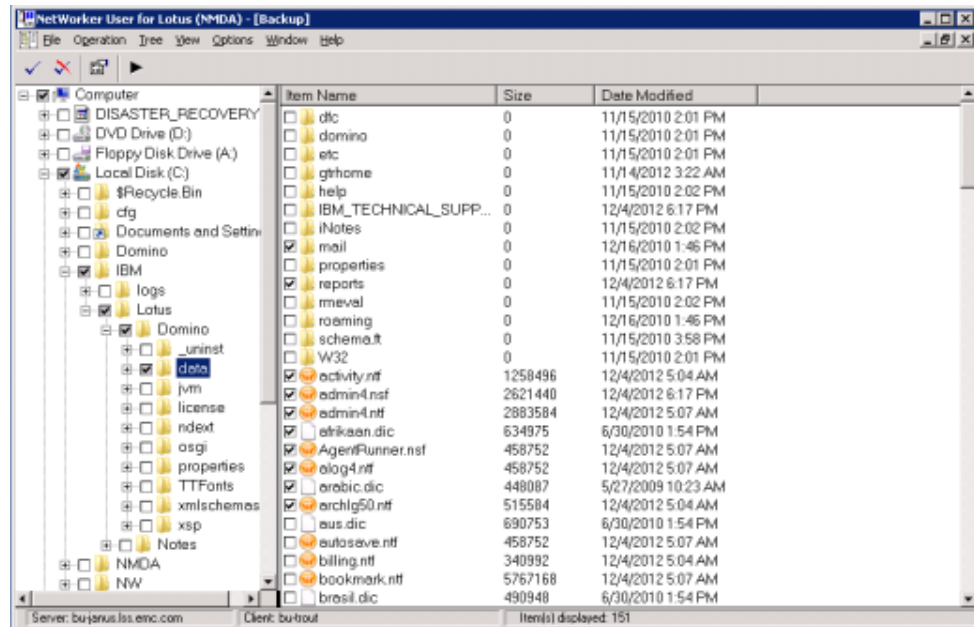
On Windows systems only, you can run the NetWorker User for Lotus GUI (`nwbml.exe`) to perform manual backups of Lotus Domino/Notes databases on the same host.

Procedure

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 the following 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



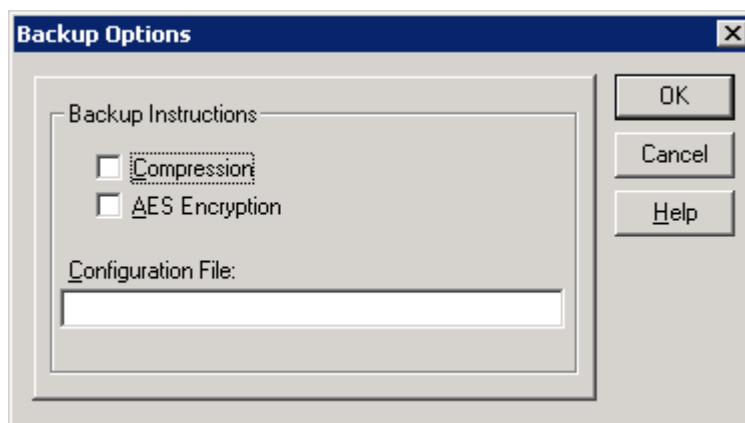
4. To view a list of files or databases available for backup, select a Lotus directory in the left pane of the **Backup** window. The Lotus directory contents appear in the right pane.
5. Select the checkbox next to each file, database, or directory to be backed up. If you select a directory, all the 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 that are specified in the **Backup Options** dialog box take precedence over the corresponding parameters that are 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 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 400 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 a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Performing MySQL manual backups with the `nsrdasv` command

About this task

Before you perform the manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

You can run the appropriate `nsrdasv` command to perform a MySQL manual backup from the command line.

Procedure

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 a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Performing Oracle manual backups with the rman command

About this task

Before you perform the manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

You can run the appropriate `rman` command to perform an Oracle manual backup from the command line.

Procedure

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 6](#) 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 a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Performing Oracle manual backups with Oracle Enterprise Manager

About this task

Before you perform the manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.


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.

Procedure

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 a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Performing SAP IQ manual backups with the nsrdasv command

About this task

Before you perform the manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

You can run the `nsrdasv` command to perform an SAP IQ manual backup from the command line.

Procedure

1. To specify the SAP IQ database to back up, set the `NSR_SAVESET_NAME` parameter in the NMDA configuration file. [Table 46](#) on page 458 provides details.
2. Log in to the SAP IQ host as the SAQ IQ operating system user.
3. Run the `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 SAP IQ manual backup.

After completing a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

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.

Procedure

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 22 Database consistency check options of the `nsrsybcc` command

The -o option	Type of database consistency check
<code>-o ckdb</code>	<code>dbcc checkdb</code>
<code>-o ckal</code>	<code>dbcc checkalloc</code>
<code>-o ckcat</code>	<code>dbcc checkcatalog</code>
<code>-o ckdbnoidx</code>	<code>dbcc checkdb (skip_ncindex)</code>
<code>-o ckstor</code>	<code>dbcc checkstorage</code>  Note: Ensure that the <code>dbccdb</code> database is set up.

The *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsrsybcc` command.

Performing Sybase manual backups with the `nsrdasv` command

About this task

Before you perform the manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

You can run the appropriate `nsrdasv` command to perform a Sybase manual backup from the command line.

Procedure

1. To specify the Sybase server or the Sybase databases to back up, set the `NSR_BACKUP_PATHS` parameter in the NMDA configuration file. [NSR BACKUP PATHS](#) 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 a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Performing Sybase manual backups with NetWorker User for Sybase

About this task

Before you perform a manual backup, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

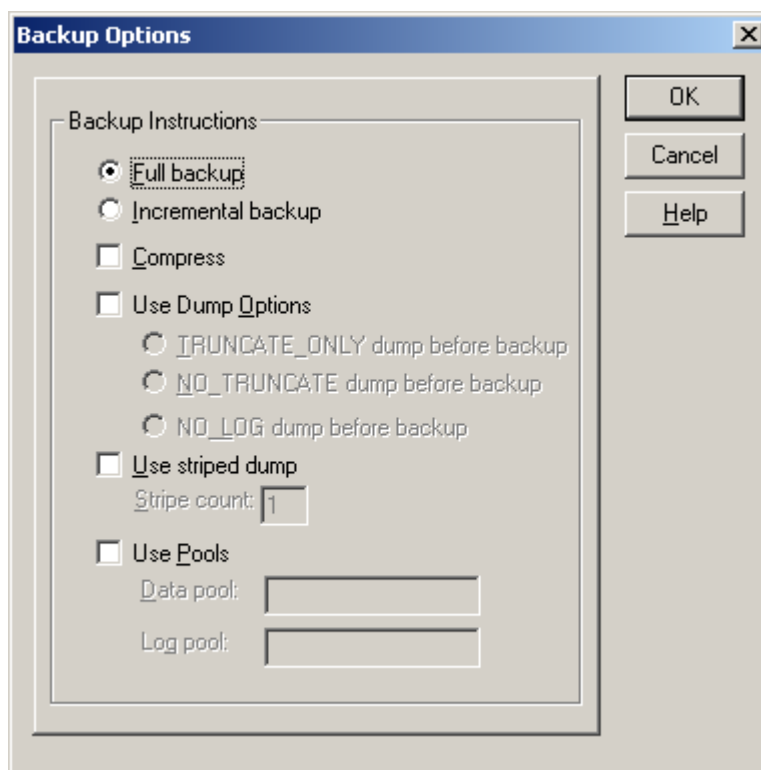
On Windows systems only, you can run the NetWorker User for Sybase GUI (`nwbms.exe`) to perform manual backups of local Sybase databases.

Procedure

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 the following 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



- In the **Backup Options** dialog box, select the required options and click **OK**.

The online help describes the options in the dialog box.

- Click **Start** in the **Backup** window.

The **Backup Status** window displays information about the backup being performed.

After completing a manual backup, back up any additional files that are required to prepare for disaster recovery as described in [Disaster Recovery](#) on page 273.

Performing Orchestrated Application Protection manual backups with the nsroappbackup command

About this task

Before you perform the manual backup of a MongoDB, MySQL, or PostgreSQL database, ensure that the required backup configurations are completed from [Configuring NMDA backups](#) on page 76.

You can run the appropriate `nsroappbackup` command to perform a database manual backup from the command line.

[Performing PostgreSQL transaction log backups with the nsroappbackup command](#) on page 179 provides additional information about performing the backup of PostgreSQL transaction logs.

Procedure

- Log in to the database server host as the database operating system user.
- Run the `nsroappbackup backup` command at the command line:

```
/usr/sbin/nsroappbackup -z configuration_file_path
```

where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the parameter settings for the manual backup.

To prepare for disaster recovery, you can include additional commands in the backup shell script to back up the disaster recovery files into the same save set with the data backup. The disaster recovery files are defined by the particular database, as described in the database documentation.

For example, the following backup shell script includes the `cp` command for the disaster recovery file backup. The script exits the processing if the executed command fails, and prints the command and returned value:

```
#!/bin/sh
( set -x; /opt/postgresql/pg94/bin/pg_basebackup --pgdata=$OAPP_MOUNT_DIR/
basebackup )
rc=$?
if [ $rc == 0 ]; then
    echo "Continued processing with the returned value $rc."
else
    echo "Exited processing with the returned value $rc."
    exit $rc
fi

( set -x; /bin/cp /home/postgres/postgresql.conf $OAPP_MOUNT_DIR/ )
rc=$?
if [ $rc == 0 ]; then
    echo "Continued processing with the returned value $rc."
else
    echo "Exited processing with the returned value $rc."
```

```
exit $rc
fi
```

Performing PostgreSQL transaction log backups with the nsroappbackup command

You can set up the backups of PostgreSQL transaction logs by registering the `nsroappbackup` command.

Procedure

1. Log in to the PostgreSQL server host as the PostgreSQL operating system user.
2. To enable the PostgreSQL transaction log backups (WAL segment file archiving), register the `nsroappbackup` command with the required command line options by completing the `archive_command` setting in the `postgresql.conf` file.

For example, complete the following setting in the `postgresql.conf` file:

```
archive_command = '/usr/sbin/nsroappbackup -o pg_p_opt="%p" -o
pg_f_opt="%f" -l txnlog -z configuration_file_path'
```

[Registering the PostgreSQL archive command](#) on page 163 provides details on the `archive_command` setting.

The PostgreSQL server runs the transaction log backup. By default, the server starts the backup when the WAL segment file is full. If you want the WAL segment file to be backed up more regularly, you can change the PostgreSQL setting to force a switch to a new WAL segment file and the backup to start based on your setting.

Canceling manual backups

You can cancel a manual backup by completing the following procedures:

About this task

NOTICE If you cancel the backup, some of the backed-up data might not be recoverable. Restart the canceled 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

About this task

Before you cancel an Oracle manual backup, ensure that the Oracle user has the required privilege for backup deletion, as described in [Table 6](#) on page 83. 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 515 or [Canceling unresponsive Oracle backups on Windows](#) on page 516.

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.

About this task

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 *NetWorker Administration Guide* provides details on how to use the NMC interfaces.

Verifying backup information in NetWorker indexes

The NetWorker server maintains information about each backup in the online indexes.

About this task

- 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, iq (for SAP IQ), 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 *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 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:

About this task

- To query the client file index, use the `nsrinfo` command, as described in [Verifying backup information in NetWorker indexes](#) on page 180.
- 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

Client	Save Set	SSID	Checkpoint ID	Save Time	Clone Retention Time	Level	Status	Size	Flags
win3e-scott-2	bootstrap	97821908		5/16/11 5:31:00 PM	5/16/12 11:59:59 PM	full	recoverable	206 KB	cr
win3e-scott-2	index:win3e-scott-2	114399121		5/16/11 5:30:57 PM	5/16/12 11:59:59 PM	full	recoverable	51 KB	cr
win3e-scott-2	NOTES:_txnlogs_1	131176284	1305581124 (003)	5/16/11 5:30:03 PM	5/16/12 11:59:59 PM	full	checkpoint, browsable	393 MB	cbk
win3e-scott-2	NOTES:_browselist	147953496	1305581124 (003)	5/16/11 5:29:57 PM	5/16/12 11:59:59 PM	full	browsable	20 KB	cb
win3e-scott-2	NOTES:_1	164730680	1305581124 (003)	5/16/11 5:29:28 PM	5/16/12 11:59:59 PM	full	checkpoint, browsable	91 MB	cbk
win3e-scott-2	NOTES:_1	181507874	1305581124 (003)	5/16/11 5:29:05 PM	5/16/12 11:59:59 PM	full	checkpoint, browsable	89 MB	cbk
win3e-scott-2	NOTES:_1	198284969	1305581124 (002)	5/16/11 5:27:06 PM	5/16/12 11:59:59 PM	full	checkpoint, aborted, suspect	133 MB	cak
win3e-scott-2	NOTES:_1	215062185	1305581124 (002)	5/16/11 5:27:05 PM	5/16/12 11:59:59 PM	full	checkpoint, aborted, suspect	107 MB	cak
win3e-scott-2	NOTES:_1	231839305	1305581124 (001)	5/16/11 5:25:28 PM	5/16/12 11:59:59 PM	full	checkpoint, aborted, suspect	46 MB	cak
win3e-scott-2	NOTES:_1	248616517	1305581124 (001)	5/16/11 5:25:24 PM	5/16/12 11:59:59 PM	full	checkpoint, aborted, suspect	56 MB	cak

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:_1	181507874	1305581124	3	1305581345	cbk
win3e-scott-2	NOTES:_1	198284969	1305581124	2	1305581226	cak
win3e-scott-2	NOTES:_1	215062185	1305581124	2	1305581225	cak
win3e-scott-2	NOTES:_1	231839305	1305581124	1	1305581128	cak
win3e-scott-2	NOTES:_1	248616517	1305581124	1	1305581124	cak

The *NetWorker Administration Guide* describes how to query information about partial save sets.

MySQL backup information in NetWorker indexes for backups without Orchestrated Application Protection

The information in this topic applies to MySQL backups that are not performed through Orchestrated Application Protection.

When NMDA performs a MySQL backup without Orchestrated Application Protection, 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 that the NetWorker server generates.

NMDA also generates an extra file that is named *backup_piece_file_name_meta*, which 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.

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
```

MongoDB backup information in NetWorker media database

When NMDA performs a MongoDB backup, the save set information is stored in the NetWorker media database.

NMDA uses a specific naming convention for the backup save set name. NMDA generates the following name for the save set file that is stored in the NetWorker media database for a MongoDB database backup:

<database_type>_<instance_name>_<backup_name>_<backup_level>

In the MongoDB save set name:

- *<database_type>* is the value MongoDB, as specified in the `NSR_DATABASE_TYPE` parameter setting.
 - [NMDA Orchestrated Application Protection parameters](#) on page 451 provides details on the parameter settings.
- *<instance_name>* is the database instance name, as specified in the `NSR_INSTANCE_NAME` parameter setting.
- *<backup_name>* is the name of the MongoDB backup image, as specified in the `NSR_BACKUP_NAME` parameter setting.
- *<backup_level>* is the MongoDB backup level, which must be full. MongoDB supports only a full database backup.

You can use the `mminfo` command to obtain the save set information from the media database about MongoDB backups.

MySQL backup information in NetWorker media database for backups with Orchestrated Application Protection

The information in this topic applies to MySQL backups that are performed through Orchestrated Application Protection.

When NMDA performs a MySQL backup through Orchestrated Application Protection, the save set information is stored in the NetWorker media database.

NMDA uses a specific naming convention for the backup save set name. NMDA generates the following name for the save set file that is stored in the NetWorker media database for a MySQL database backup:

<database_type>_<instance_name>_<backup_name>_<backup_level>

In the MySQL save set name:

- *<database_type>* is the value MySQL, as specified in the NSR_DATABASE_TYPE parameter setting.
[NMDA Orchestrated Application Protection parameters](#) on page 451 provides details on the parameter settings.
- *<instance_name>* is the database instance name, as specified in the NSR_INSTANCE_NAME parameter setting.
- *<backup_name>* is the name of the MySQL backup image, as specified in the NSR_BACKUP_NAME parameter setting.
- *<backup_level>* is the MySQL backup level, which must be full for the full level backup, and txnlog for the Logs Only level backup.

You can use the `mminfo` command to obtain the save set information from the media database about MySQL backups that are performed through Orchestrated Application Protection.

PostgreSQL backup information in NetWorker media database

When NMDA performs a PostgreSQL backup, the save set information is stored in the NetWorker media database.

NMDA uses a specific naming convention for the backup save set name. NMDA generates one of the following names for the save set file that is stored in the NetWorker media database for a PostgreSQL backup:

- For a PostgreSQL database backup:
<database_type>_<instance_name>_<backup_name>_<backup_level>
- For a PostgreSQL transaction log backup:
<database_type>_<instance_name>_<backup_name>_<backup_level>_<transaction_log_file_name>

In the PostgreSQL save set name:

- *<database_type>* is the value PostgreSQL, as specified in the NSR_DATABASE_TYPE parameter setting.
[NMDA Orchestrated Application Protection parameters](#) on page 451 provides details on the parameter settings.
- *<instance_name>* is the database instance name, as specified in the NSR_INSTANCE_NAME parameter setting.
- *<backup_name>* is the name of the PostgreSQL backup image, as specified in the NSR_BACKUP_NAME parameter setting.

- `<backup_level>` is the PostgreSQL backup level:
 - `full` specifies a full backup.
 - `incr` specifies an incremental backup.
 - `txnlog` specifies a transaction log backup.
- `<transaction_log_file_name>` is the pathname of the transaction log file, which appears only in the save set name for a transaction log backup.

You can use the `mminfo` command to obtain the save set information from the media database about PostgreSQL backups.


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 that are stored in the respective backup application catalogs. A DB2 database, Informix database, or Oracle database has its own backup catalog.

About this task

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. Also, 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 6](#) on page 83.


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, SAP IQ, and Sybase do not maintain their own backup catalogs. You can delete a Lotus backup, SAP IQ 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.

About this task

 **Note:** You cannot use the `db2 prune` command to delete snapshot backups. [Deleting DB2 snapshots](#) on page 365 provides details.

To delete DB2 backup entries on both the DB2 server and NetWorker server:

Procedure

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 that is used to store the NMDA backups in the NetWorker indexes.

The *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 `onssmsync` command.

About this task

Run the Informix `onssmsync` 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 that are associated with the Informix backup entries, ensure that the user who runs the `onssmsync` utility has the required privileges for backup deletion, as described in [Table 6](#) on page 83.

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:

About this task

- 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 47 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 [Verifying the NetWorker User Group resource](#) on page 82. 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 the following conditions:

- NetWorker client binaries are located in a nondefault directory on the Oracle Server host.

- NWORA resource file was not created.
[NSR_NWPATH](#) provides details.

CHAPTER 4

Data Restore and Recovery

This chapter includes the following topics:


- [Data restore and recovery terminology](#)..... 190
- [Performing NMDA data restore and recovery](#)..... 191
- [Performing DB2 data restore and recovery](#)..... 193
- [Performing Informix data restore and recovery](#)..... 204
- [Performing Lotus data restore and recovery](#)..... 207
- [Performing MySQL data restore and recovery](#)..... 222
- [Performing Oracle data restore and recovery](#)..... 235
- [Performing SAP IQ data restore](#)..... 244
- [Performing Sybase data restore and recovery](#)..... 251
- [Performing Orchestrated Application Protection data restore](#)..... 264
- [Canceling a data restore or recovery](#)..... 272

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.

The terms have specific meanings for the NMDA operations:

- *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, SAP IQ, Sybase—Recovery

 **Note:** You can *restore* a MongoDB backup, MySQL backup, or PostgreSQL backup that was performed through Orchestrated Application Protection.

NMDA can restore only the following data:

- Data that you backed up according to instructions in [Backup Procedures](#) on page 165.
- 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 that are used for restores

NMDA requires information from the NetWorker indexes to complete a restore operation. The NetWorker retention policies affect the available data in the NetWorker indexes and backup volumes.

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 media database entry until the retention policy specified for the client save set expires.

You define the lifecycle by setting the retention policy in the following locations:

- NetWorker Client resource for scheduled backups.
- `NSR_SAVESET_RETENTION` parameter for manual backups. If you do not set this parameter, the NetWorker server assigns the retention policy for a manual backup according to the setting in the NetWorker Client resource.

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 relabeled. When the volume is relabeled, 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.

If 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 185.

Note: For Oracle backups, the index entries that are 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 134.

The *NetWorker Administration Guide* describes how the NetWorker server uses retention policies to manage backup data and track data on volumes.

Performing NMDA data restore and recovery

To prepare for an NMDA data restore and recovery, complete the required configurations and ensure that the required backup, services, and software are available.

Before you perform a data restore and recovery, ensure that you 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 that is 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 the following attribute settings:
 - Save set: All
 - Protection group list: (blank)
 - The Remote Access attribute in the Client resource of the source client contains the following value:

```
user=database_or_application_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 articles:
 - *Fibre Channel Devices with Products using DD Boost in Linux/UNIX Environment* (Document ID dd95007)
 - *Fibre Channel Devices with Products using DD Boost in Windows Environment* (Document ID dd95005)

Use the document ID to search for these articles on the Support website at <https://support.emc.com>.

[Table 7](#) on page 83 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 [NMDA Parameters and Configuration File](#) on page 399:
 - 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 an Orchestrated Application Protection restore, including a MongoDB restore, MySQL restore, or PostgreSQL restore, in the NMDA configuration file.
 - You have set the parameters for an SAP IQ restore in the NMDA configuration file.
 - You have set the parameters for a Sybase restore in the NMDA configuration file.

The following 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 the following conditions are true:
 - NMDA backed up the data with AES encryption enabled through 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_SERVER—If the NetWorker server is not on the NMDA host, the parameter is set 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 23 Data restore procedures

Database or application	Data restore procedures
DB2	<ul style="list-style-type: none"> • <code>db2 restore</code> command • IBM Data Studio GUI or DB2 Control Center GUI • DB2 HPU utility <p>Performing DB2 data restore and recovery on page 193 provides details.</p>
Informix	<ul style="list-style-type: none"> • <code>onbar restore</code> command <p>Performing Informix data restore and recovery on page 204 provides details.</p>
Lotus	<p>Database restore and DAOS file-level restore:</p> <ul style="list-style-type: none"> • <code>nsrnotesrc</code> command • NetWorker User for Lotus GUI (Windows only) <p>Document-level recovery:</p> <ul style="list-style-type: none"> • <code>nsrdocrc</code> command

Table 23 Data restore procedures (continued)

Database or application	Data restore procedures
	<ul style="list-style-type: none"> Notes client GUI (Windows only) Performing Lotus data restore and recovery on page 207 provides details.
MySQL	<ul style="list-style-type: none"> <code>nsrmysqlrc</code> 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 222 provides details.
Oracle	<ul style="list-style-type: none"> <code>rman</code> command to run an RMAN restore script Oracle Enterprise Manager Performing Oracle data restore and recovery on page 235 provides details.
SAP IQ	<ul style="list-style-type: none"> <code>nsriqrc</code> command for restore Performing SAP IQ data restore on page 244 provides details.
Sybase	<ul style="list-style-type: none"> <code>nsrsybrc</code> command NetWorker User for Sybase GUI (Windows only) Performing Sybase data restore and recovery on page 251 provides details.
Orchestrated Application Protection (MongoDB, MySQL, PostgreSQL)	<ul style="list-style-type: none"> <code>nsroapprecovery</code> command for restore Performing Orchestrated Application Protection data restore on page 264 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.

You can also use the DB2 HPU utility to quickly unload and extract discrete data, such as table data, from an NMDA backup image and then load the data into a DB2 database.

The DB2 documentation provides details about the different commands and options that are used for restore and recovery.


Perform the DB2 data restore and recovery by using the appropriate procedures in the following topics:

- Determining how many restore devices and sessions to use
- Performing DB2 data restores with the `db2 restore` command
- Performing DB2 data recovery
- Performing DB2 data restore and recovery with the NMDA DB2 recovery wizard
- Performing DB2 data restore and recovery with the DB2 GUI
- Performing DB2 data recovery with the DB2 HPU utility
- Performing DB2 rollforward recovery after a load with `copy yes` option

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.

About this task

 **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-11et -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.

[Performing DB2 data restores to a different host](#) on page 198 describes the restore to 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

Procedure

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, 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.

Note: Restoring the most recent backup that is a tablespace backup requires a full timestamp.

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 194 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 yyyymmddhhmmss 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 yyyymmddhhmmss 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* or *pathname\nmda_db2.cfg* is the complete pathname of the NMDA configuration file.
- *yyyymmddhhmmss* 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.

Note: A restore without a timestamp always restores the latest database backup, even if there is a tablespace backup after the database backup. To restore such a tablespace backup, use the full timestamp. You can obtain the timestamp of a backup from the DB2 recover history or from the backup storage by running the

`nsrinfo` command as described in [Verifying backup information in NetWorker indexes](#) on page 180.

- *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

Procedure

1. Log in to the host that has the new DB2 instance. Log in as the DB2 operating system user for the new DB2 instance.
2. If you are restoring the most recent backup of a database, 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 194 provides details.
4. Grant the new instance, for example, `db2inst2`, permission to restore the database:
 - a. Start the NMC program, and then 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.

Ensure that you do not add the `DB2_R` setting to the Backup Options 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 yyyymmddhhmmss 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 yyyymmddhhmmss  
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 195.

Note: Ensure that the new instance has read and write permission to the script. Ensure that the source NetWorker client is available.

6. Edit the generated script, and define the following parameters:

- OPTIONS (mandatory)—Complete pathname of the NMDA configuration file.
- DBPATH ON *target_directory*—Complete pathname of the target database directory.
- ON *path_list*—Redefinition of the storage paths that are associated with a database.

Note: If the database does not exist on disk and the DBPATH ON parameter is not specified, then the first path is used as the target database directory.

- INTO—New database name, if you are redirecting the recovery to a new name.
- TAKEN AT—Timestamp of the backup to recover, *yyyymmddhhmmss*, if you are restoring the data to a point-in-time, as noted in [step 2](#).

Note: Restoring the most recent backup that is a tablespace backup requires a full timestamp.

- 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 with the new DB2 instance, type the following command to run the redirection script under the redirected new 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 restores to a different host

You can restore and recover a backup of DB2 database A to the DB2 database B on a different host.

Procedure

1. Create a NetWorker Client resource for the target host.
2. Complete the following attribute and parameter settings for the restore:
 - In the Remote Access attribute in the Client resource of the server A host, specify the username and fully qualified hostname for the remote host that has the DB2 database B. Use the following format for the Remote Access attribute value:

```
user=remote_username,host=remote_hostname
```


- Set the NSR_CLIENT parameter to the source client hostname for the server A host.
 - Set the NSR_SERVER parameter to the hostname of the NetWorker server that is used for the backup.
3. Complete the restore according to one of the following procedures:
 - [Performing DB2 data restores to the same instance](#) on page 195
 - [Performing DB2 data restores to a different instance](#) on page 196
 4. On the target host, if the restored database is in the restore pending state, apply the transaction logs according to [Performing DB2 recovery with the db2 rollforward command](#) on page 199.
 5. Complete the following changes for the DB2 database B on the target host:
 - Update the NSR_CLIENT parameter with the target client source hostname.
 - Update the NSR_SERVER parameter if the target database B uses a different NetWorker server.
 - Update the DB2 database B configuration LOGARCHOPT1 and VENDOROPT parameters with the appropriate configuration file pathnames.

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 32 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.

Procedure

1. Ensure that the database parameters LOGARCHMETH n and LOGARCHOPT n have been configured as described in [Configuring automatic backups of DB2 transaction logs](#) on page 122.
2. Apply the transaction logs by running the appropriate command:
 - Run the following command to apply all the transactions to the end of the logs:

```
db2 "rollforward db sample to end of logs and complete"
```

where *sample* is the database name.

- Run the following command to apply all the transactions to a specific point-in-time:


```
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 that are stored on the backup media, especially tape media.

Procedure

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 *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 that is 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.

Procedure

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 @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. To recover the database or tablespace to the end of the logs or to a specific point-in-time, run the `db2 recover` command with appropriate options. 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 195.

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 NMDA DB2 recovery wizard

You can use the NMDA DB2 recovery wizard to configure and run the restore and recovery of DB2 data that is backed up by NMDA.

About this task

[NMDA DB2 recovery wizard](#) on page 33 provides more details about the NMDA DB2 recovery wizard.

Before you use the NMDA DB2 recovery 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
 - Client-side configuration method without the wizard, where the value of the Save Set attribute of the Client resource has the DB2: prefix

You can use the following procedure to perform a restore and recovery with the wizard.

Procedure

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 **Protection**.
4. In the **Protection** window, click **Clients**.
5. To start the wizard, right-click the NMDA client in the right pane, and then select **Recover**.
6. On each wizard screen that appears, specify the required values for the restore and recovery 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.

i **NOTICE** The recovery wizard displays the tablespaces from the target database for tablespace-level restore selection if required.

You can select to start the restore or recovery immediately from the wizard or schedule the operation to start later.

In the wizard, you can access an existing recover configuration at a later time or you can view the recovery results. In NMC, click **Recover** on the **Administration** window toolbar to

open the **Recover** window. In the **Configured Recovers** pane, right-click the saved recover configuration and select one of the menu options:

- **New Recover**—Select this option as another way to create a new recover configuration.
- **Open Recover**—Select this option to view the recovery results.
- **Recover Again**—Select this option to access an existing recover configuration. You can make changes and save the configuration with a new name if required.

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:

About this task

- 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_path`. 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 DB2 data recovery with the DB2 HPU utility

You can use the IBM Optim High Performance Unload (HPU) utility to unload and extract a discrete volume of data from an NMDA DB2 backup image into an output file. You can then load the data from the output file into a DB2 database.

About this task

Use the following guidelines to prepare for DB2 data recovery with the HPU utility:

- Use an offline backup image when possible to help ensure the integrity and consistency of the unloaded data.
- Use an online backup image only when you are certain that no transactions took place during the most recent online backup against the objects that you will unload.
- Use a tablespace backup instead of a full database backup when possible to reduce the size of the backup image being read and enable a faster data unload.
- Include the capture of `create table` statements in the backup strategy. When a table is dropped, you must re-create the table before you can use the `load` command to recover the table.

- Determine the timestamp of the backup image, for example, 20130322093409.
- Create the control file for the HPU operation. Include a command to extract the data with the `taken at` clause as shown in the following control file that is named `sysin`:

```
GLOBAL CONNECT TO SAMPLE_DB DB2 NO;
UNLOAD TABLESPACE
QUIESCE NO
LOCK NO
USING BACKUP DATABASE SAMPLE_DB LOAD "/usr/lib/libnsrdb2.so" options
@pathname/nmda_db2.cfg taken at yyyymmddhhmmss;
SELECT CUST_ID, ADDRESS FROM CUST_TABLE;
OUTFILE("output_file")
FORMAT DEL;
```

In this sample control file:

- `pathname/nmda_db2.cfg` is the pathname of the NMDA configuration file.
- `yyymmddhhmmss` is the timestamp of the backup image.
- `output_file` is the name of the output file to which the data will be extracted.
- Do not use the HPU utility with `libnsrdb2` over a Data Domain Fibre Channel (FC) network. Due to a known limitation of the DD Boost library over FC with child processes, NMDA does not support the HPU utility operations over an FC network connection.

Run the `db2hpu` command with the control file to unload and extract the required data with the HPU utility. For example, the following command uses the sample control file named `sysin` and generates the output file that contains the extracted data:

```
db2hpu -f sysin -i db2inst1
```

Performing DB2 rollforward recovery after a load with copy yes option

After you have enabled rollforward recovery for a database and configured DB2 backups with NMDA, you can use the `load` utility with the `copy yes` option to save a copy of the DB2 database changes (during the load operation) to a specified location. A subsequent rollforward recovery of the database loads the copy of the saved changes directly into the database.

[DB2 load command with the copy yes option](#) on page 34 describes the NMDA support of the `load` command with the `copy yes` option.

Note: The rollforward recovery operation requires the saved load-copy image backup to be at the same location where it was created during the load operation. Use NetWorker and NMDA to back up both the archive logs and load-copy images to ensure that the data is secure and always available for a database recovery.

The following example shows the required configuration steps, followed by the `load` operation. The operation loads the data in delimited ASCII format (DEL) from the file `f1.dat` into table `T1` and to NetWorker through the NMDA DB2 vendor library as specified by the `copy yes` option:

```
db2 update db cfg for SAMPLE using logarchmeth1 VENDOR:/usr/lib/libnsrdb2.so
db2 update db cfg for SAMPLE using logarchopt1 @/home/db2inst1/db_nmda.config
db2 update db cfg for SAMPLE using vendoropt @/home/db2inst1/db_nmda.config
db2 "load from f1.dat of DEL insert into T1 copy yes LOAD /usr/lib/
libnsrdb2.so"
```

You can use the `nsrinfo` command to obtain information about the load copy data that is stored in the NetWorker indexes. In the following example, the `nsrinfo` command returns the line that includes `description...LOAD_COPY`:

```
nsrinfo -v -s NW_Server -n db2 -X all NW_Client

scanning client 'NW_Client' for all savetimes from the db2 namespace on
server NW_Server
version=1, DB2, objectname=/SAMPLE/NODE0000 /DB_BACKUP.20141114144126.1,
createtime=Fri Nov 14 14:41:26 2014, copytype=3 BSACopyType_BACKUP,
copyId=1415994086.1415994087, restoreOrder=1415994086.1, objectsize=0.0,
resourcetype=database, objecttype=4 BSAObjectType_DATABASE, objectstatus=2
BSAObjectStatus_ACTIVE, description=NMDA_v82:DB2_v1051:LOAD_COPY:SAMPLE:TEQ,
objectinfo=db2inst1:1, NSR size=278900
```

You can also use the `DB2 list history` command to verify the load operation:

```
db2 list history load all for db SAMPLE
```

where `SAMPLE` is the database name.

In the following example, the `restore` operation restores the last backup of the database `SAMPLE` that was performed before the `load` operation. The `rollforward` operation locates and loads the load-copy image (saved by the `load` operation) directly into the database:

```
db2 restore db SAMPLE LOAD /usr/lib/libnsrdb2.so options @/home/db2inst1/
db_nmda.config taken at 20141114140133
db2 rollforward db SAMPLE to end of logs and stop
```

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 according to the following two topics:

1. Determining the Informix restore mode
2. Performing Informix data restores with the `onbar` command

[Performing an Informix imported restore](#) on page 277 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:

About this task

- 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 `onmsync` 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.

About this task

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.ifxml
```

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 that were completed since the last dbject backup and then rolls forward the logical log files that were backed up for the dbjects. If different backup sessions are involved, the log rolls forward transactions that were made since the backup time recorded for each dbject 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.ifxml
```

 **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.ifxml
```

 **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.ifxml
```

 **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 00:00:00" -w -p
```

Performing Lotus data restore and recovery

To prepare for a Lotus restore and recovery, complete the required settings and ensure that the backup and required logs are available.

About this task

Before you perform a Lotus data restore and recovery, ensure that you 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 434 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 must 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 in the following topics:

- Performing Lotus database recovery with the `nsrnotesrc` command
- Performing Lotus database restore and recovery with the NMDA Lotus recovery wizard
- Performing Lotus database recovery with NetWorker User for Lotus
- Performing recovery of partitioned Domino servers
- Performing Lotus DAOS data recovery
- Performing Lotus document-level recovery with the `nsrdocrc` command
- Performing Lotus document-level recovery with the Notes client GUI

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 that are required for recovery in the configuration file as described in [NMDA Parameters and Configuration File](#) on page 399.

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](#).

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 12 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 434 describes how to determine the recovery point-in-time.

Example 13 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 14 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 NetWorker client:

```
NSR_BACKUP_PATHS = NOTES:
```

NOTICE Use the `NOTES:` option with caution. NMDA tries 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 15 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 18](#) 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 16 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 17 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:

Example 17 Relocation of a linked database during Lotus recovery (continued)

- 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 18 Point-in-time 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 recovery of the logged database that is 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 434 describes how to determine the recovery point-in-time.

Performing Lotus database restore and recovery with the NMDA Lotus recovery wizard

You can use the NMDA Lotus recovery wizard to configure and run the restore and recovery of Lotus Domino/Notes data that is backed up by NMDA.

About this task

[NMDA Lotus recovery wizard](#) on page 39 provides more details about the NMDA Lotus recovery wizard.

Before you use the NMDA Lotus recovery 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
 - Client-side configuration method without the wizard, where the value of the Save Set attribute of the Client resource has the NOTES: prefix

You can use the following procedure to perform a restore and recovery with the wizard.

Procedure

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 **Protection**.
4. In the **Protection** window, click **Clients**.
5. To start the wizard, right-click the NMDA client in the right pane, and then select **Recover**.
6. On each wizard screen that appears, specify the required values for the restore and recovery 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.

You can select to start the restore or recovery immediately from the wizard or schedule the operation to start later.

In the wizard, you can access an existing recover configuration at a later time or you can view the recovery results. In NMC, click **Recover** on the **Administration** window toolbar to open the **Recover** window. In the **Configured Recovers** pane, right-click the saved recover configuration and select one of the menu options:

- **New Recover**—Select this option as another way to create a new recover configuration.
- **Open Recover**—Select this option to view the recovery results.
- **Recover Again**—Select this option to access an existing recover configuration. You can make changes and save the configuration with a new name if required.

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.

About this task

Perform the Lotus data recovery by using the appropriate GUI procedure in the following two topics:

- Performing Lotus recovery of local data with the GUI
- Performing Lotus directed recovery with the GUI

Performing Lotus recovery of local data with the GUI

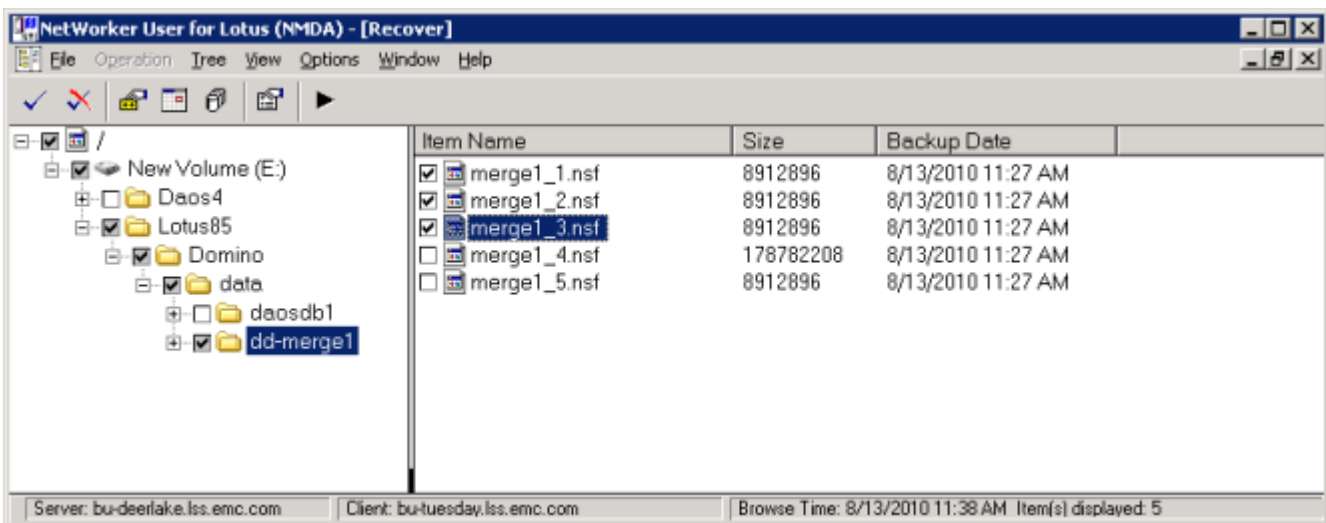
You can use the NetWorker User for Lotus GUI to recover data on the local Windows host.

Procedure

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. To refresh the list of NetWorker servers, click **Update List**.
 - c. Select or type the name of the server.
 - d. To use the server as the default NetWorker server, select **Save as Default 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. To change from the current month, click **Previous Month** or **Next 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 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.

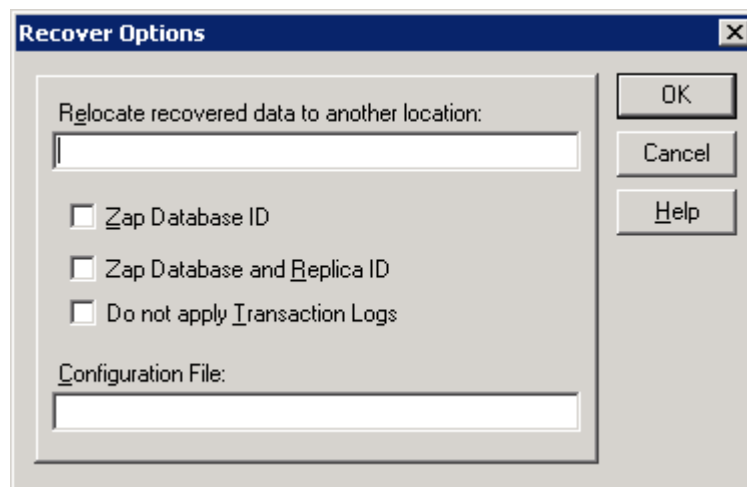
Note: Do not set NSR_BACKUP_PATHS in the NMDA configuration file. Instead, use the GUI to select the files and directories to be recovered.

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 400 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 208 includes examples of the parameters.

Table 24 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.

About this task

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 the procedures in the following two topics to perform a Lotus directed recovery with the NetWorker User for Lotus GUI:

1. Configuring Lotus directed recovery
2. Performing Lotus directed recovery

Configuring Lotus directed recovery

About this task

- 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 `nsr_lotus_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 `nsr_lotus_remrecov.bat`. On UNIX, the script file name is `nsr_lotus_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 `nsr_lotus_remrecov` script that you modified on the destination client. If you have not set the variable, the NetWorker User for Lotus uses `nsr_lotus_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=nsr_lotus_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

Procedure

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 then 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.

Steps 4 to 10 in the section [Performing Lotus recovery of local data with the GUI](#) on page 212 provide details.

- NOTICE** When you perform the directed recovery to a UNIX or Linux destination client, the total length of all the pathnames that are 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.

Procedure

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\pldata
```

Note: Do not specify only `NOTES:` for the `NSR_BACKUP_PATHS` parameter because NMDA tries to recover all the Lotus data for all the Domino servers on the host.

[NMDA Parameters and Configuration File](#) on page 399 describes the parameters in the NMDA configuration file.

Performing Lotus DAOS data recovery

To restore Domino databases and the corresponding missing NLO files, complete the following steps.

Procedure

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. To determine the missing NLO files for the restored databases, run the Domino `tell` command. 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 that is 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](#) and [NSR_RECOV_LIST_FILE](#) 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.

Results

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 the `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.

About this task

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 the following steps as the Lotus user.

Procedure

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).

[NMDA Parameters and Configuration File](#) on page 399 describes parameters in the NMDA configuration file.

[Setting the NSR_RECOVER_TIME parameter](#) on page 434 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 an empty 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 `nsrdocr` command at the command line:

```
nsrdocr(.exe) -z configuration_file_path
```

where `configuration_file_path` is the complete pathname of the NMDA configuration file.

The *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsrdocr` command.

[Example 19](#) provides an example of document-level recovery with the `nsrdocr` command.

[Performing document-level recovery of database links](#) on page 222 describes how to recover documents from a linked database.

Example 19 Recovering deleted documents for a logged database

Complete the following steps to recover deleted documents in the database `F:\Lotus\Domino\data\account.nsf` that resides on the client host `saturn`, as backed up on June 15 to the NetWorker server `mars`.

1. To obtain the backup time of the directory that contains `account.nsf`, run the `nsrinfo` command:

```
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:

About this task

- 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 *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 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

About this task

- i** **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:

Procedure

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 an empty file in the directory that is 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.

i **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 222 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 that is listed in the **NetWorker Client** text box might not be correct, which 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 that is used to back up the database if the datazone pass phrase on the NetWorker server has changed since the backup.

 **Note:** Encryption keys that are 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:

About this task

- 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 221.
- Specify the link file path with the `NSR_BACKUP_PATHS` parameter in the NMDA configuration file as described in [NSR_BACKUP_PATHS 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

To prepare for a MySQL data restore or recovery, complete the required configurations and ensure that specific directories have sufficient space to contain the backup data. Determine the binary logs that are required for the recovery.

About this task

Before you perform a MySQL data restore or recovery, ensure that you 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 MySQL recovery wizard or 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. [NMDA Parameters and Configuration File](#) on page 399 describes all the supported parameters.

- The local directory that is specified by `MYSQL_BACKUP_DIR` has enough space to contain the temporary backup files that are 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 that is 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 that is required for the data.
- You have determined the binary log range that is required for the recovery as described in [Determining the binary logs for MySQL recovery](#) on page 223. The binary logs that are 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 228.

You can run the NMDA MySQL recovery wizard or the `nsrmysqlrc` command 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 recovery wizard or `nsrmysqlrc` command must be the owner of the directory that is specified by `MYSQL_DATADIR`.

The *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsrmysqlrc` command.

Perform the MySQL data restore and recovery operations by using the appropriate procedures in the following topics:

- [Determining the binary logs for MySQL recovery](#)
- [Performing MySQL restore and recovery with the NMDA MySQL recovery wizard](#)
- [Performing MySQL recovery of whole instance backups](#)
- [Performing MySQL restores of partial backups](#)
- [Performing MySQL restores of binary log backups](#)
- [Performing MySQL redirected restores](#)
- [Performing MySQL list image operations](#)
- [Performing MySQL extract operations](#)
- [Performing MySQL extract and prepare operations](#)
- [Performing MySQL copy back operations](#)
- [Performing MySQL validate operations](#)

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.

About this task

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.

Procedure

1. To identify the binary logs that are required for the last backup, run the `nsrinfo` command. 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,
compress 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 must set `MYSQL_BINLOG` accordingly during the recovery of MySQL databases.

Performing MySQL restore and recovery with the NMDA MySQL recovery wizard

You can use the NMDA MySQL recovery wizard to configure and run the restore and recovery of MySQL data that is backed up by NMDA.

About this task

[MySQL restore and recovery with the NMDA MySQL recovery wizard](#) on page 42 provides more details about the NMDA MySQL recovery wizard.

Before you use the NMDA MySQL recovery 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.
- On the source host, you have created the NetWorker Client resource for the NMDA client by using one of the following methods:
 - Backup configuration wizard in NMDA.
 - Client-side configuration method without the wizard, where the value of the Save Set attribute of the Client resource is `MYSQL:/unique_backup_name`.

You can use the following procedure to perform a restore and recovery with the wizard.

Procedure

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 **Protection**.
4. In the **Protection** window, click **Clients**.
5. To start the wizard, right-click the NMDA client in the right pane, and then select **Recover**.
6. On each wizard screen that appears, specify the required values for the restore and recovery 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.

You can select to start the restore or recovery immediately from the wizard or schedule the operation to start later.

In the wizard, you can access an existing recover configuration at a later time or you can view the recovery results. In NMC, click **Recover** on the **Administration** window toolbar to open the **Recover** window. In the **Configured Recovers** pane, right-click the saved recover configuration and select one of the menu options:


- **New Recover**—Select this option as another way to create a new recover configuration.
- **Open Recover**—Select this option to view the recovery results.
- **Recover Again**—Select this option to access an existing recover configuration. You can make changes and save the configuration with a new name if required.

Performing MySQL recovery of whole instance backups

By default, the `nsmysqlrc` 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.

About this task

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 that is restored from the full backup before restoring the incremental backups.

For a point-in-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 that are performed before the point-in-time. The software then applies the appropriate binary log backups to recover the instance to the point-in-time.

Procedure

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.

- The `MYSQL_BINLOG` parameter specifies 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.

- The `MYSQL_DATADIR` parameter specifies the `/var/lib/mysql` directory to which a copy back operation restores the prepared backup.
- The `MYSQL_MEB_OPTIONS` parameter specifies to uncompress the compressed backup.
- The `NSR_RECOVER_TIME` parameter is set for a point-in-time recovery only.
- The `NSR_SERVER` parameter specifies the hostname of the NetWorker server host.

[NMDA Parameters and Configuration File](#) on page 399 provides details on each parameter in the NMDA configuration file.

2. To restore the whole instance backup, run the `nsrmysqlrc` command at the command line:

```
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 then 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.

About this task

The `nsrmysqlrc` program restores only the databases and tables in the backup image that is 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 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 that is 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 that were performed before the point-in-time. However, the software does not apply any binary log backups to the restored partial backup.

Procedure

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
```

In this example:

- The `MYSQL_BACKUP_NAME` and `MYSQL_BACKUP_DIR` parameters specify to extract the backup files from the backup image named `MYSQL:/mydb1` 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.
- The `MYSQL_DATADIR` parameter specifies the `/var/lib/mysql` directory to which a copy back operation restores the prepared backup.
- The `MYSQL_MEB_OPTIONS` parameter specifies to uncompress the compressed backup.
- The `NSR_RECOVER_TIME` parameter is set for a point-in-time recovery only.
- The `NSR_SERVER` parameter specifies 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.

[NMDA Parameters and Configuration File](#) on page 399 describes parameters in the NMDA configuration file.

- To restore the partial backup, run the `nsmysqlrc` command at the command line:

```
nsmysqlrc -z configuration_file_path
```

where *configuration_file_path* is the complete pathname of the NMDA configuration file that contains the required parameter settings.

- 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 that is 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 `nsmysqlrc` command to restore one or more binary logs from binary log backups.

Procedure

- 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 `nsmysqlrc`. Set `NSR_SERVER` if the NetWorker server host is different from the client host.

[NMDA Parameters and Configuration File](#) on page 399 describes parameters in the NMDA configuration file.

- To restore the binary log backup, run the `nsmysqlrc` command at the command line:

```
nsmysqlrc -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:

About this task

- 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 that is 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 284 provides details on disaster recovery.

Performing a MySQL redirected restore to a different data directory

You can use the `nsmysqlrc` command to restore backup data to a different data directory on the same host.


Procedure

- Set `MYSQL_DATADIR` to the pathname of the new data directory.
- 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.
- Use one of the following procedures to complete the restore:
 - [Performing MySQL recovery of whole instance backups](#) on page 225
 - [Performing MySQL restores of partial backups](#) on page 227

Performing a MySQL redirected restore to a different host

You can use the `nsmysqlrc` command to restore and recover a backup of MySQL server A to MySQL server B on a different host.

About this task

 **Note:** You must run the `nsmysqlrc` command on the server B host.

Procedure

- 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 that is used for the backup.

2. Use one of the following procedures to complete the restore:
 - [Performing MySQL recovery of whole instance backups](#) on page 225
 - [Performing MySQL restores of partial backups](#) on page 227

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 that is created by an NMDA MySQL backup. The list image operation does not change the backup image.

About this task

You do not need to extract files from the backup image before the list image operation. You must ensure that the device containing the backup image is mounted before you perform a list image operation.

Procedure

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.

[NMDA Parameters and Configuration File](#) on page 399 describes parameters in the NMDA configuration file.

2. To complete the list image operation, run the `nsrmysqlrc` command at the command line:

```
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.

About this task

Note: Before you perform an extract operation, you must ensure that 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.

Procedure

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.

[NMDA Parameters and Configuration File](#) on page 399 describes parameters in the NMDA configuration file.

- To complete the extract operation, run the `nsmysqlrc` command at the command line:

```
nsmysqlrc -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 `nsmysqlrc` 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.

About this task

The `nsmysqlrc` program performs the following steps to complete an extract and prepare operation:

- Locates the latest full backup and all the associated incremental backups (if any) for the backup image that is specified by the `MYSQL_BACKUP_NAME` parameter.
- Extracts the full backup and incremental backups to subdirectories under the directory that is specified by the `MYSQL_BACKUP_DIR` parameter.
- Uses an apply-log operation to prepare the full backup as required.
- 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.

Procedure

- 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.

[NMDA Parameters and Configuration File](#) on page 399 describes parameters in the NMDA configuration file.

2. To complete the extract and prepare operation, run the `nsmysqlrc` command at the command line:

```
nsmysqlrc -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 host 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 `nsmysqlrc` 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 `nsmysqlrc` program through an “extract and prepare” operation or by the MEB utilities.

About this task

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.

Procedure

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.

[NMDA Parameters and Configuration File](#) on page 399 describes parameters in the NMDA configuration file.

- To complete the copy back operation, run the `nsrmysqlrc` command at the command line:

```
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.

About this task

You must complete the required steps to restore an NMDA backup of InnoDB tables that are 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.

Procedure

- Extract and prepare the backup to a staging area. [Performing MySQL extract and prepare operations](#) on page 232 provides details.
- Edit all the `.b1` files in the backup directory so that each file includes the correct pathnames for the new directory structure.
- Copy back the data to the MySQL data directory. [Performing MySQL copy back operations](#) on page 233 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 change the backup image.

About this task

Ensure that the device containing the backup image is mounted before you perform a validate operation.

Procedure

- 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.

[NMDA Parameters and Configuration File](#) on page 399 describes parameters in the NMDA configuration file.

- To complete the validate operation, run the `nsrmysqlrc` command at the command line:

```
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 recovery of a replicated slave server

You must complete the required steps for a MySQL point-in-time recovery of a replicated slave server.

Procedure

1. Use NMDA to restore the MySQL backup of the replicated slave server.
 - Note:** MEB 3.11 or later must be used to manage the binary and relay log backups of the replicated slave server. NMDA does not store information about available logs in the replicated slave server backups.
2. Use the `mysqlbinlog` program to restore the binary logs as required. The MEB documentation provides details about how to perform the point-in-time log recovery.

Performing Oracle data restore and recovery

Perform the Oracle data restore and recovery by using the appropriate procedures in the following topics:

- Determining the volumes for Oracle restores
- Preventing restore performance degradation with Oracle 10.2 or later
- RMAN scripts for restore and recovery
- Performing Oracle restore and recovery with the NMDA Oracle recovery wizard
- Performing Oracle restores with the `rman` command
- Performing Oracle restores with the Oracle Enterprise Manager
- Performing Oracle data recovery

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 that are listed in the output of `restore...preview` are the volumes that are 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 that is used at the backup time (for example, due to NetWorker staging), the `restore...preview` operation still lists the volume name that is used at the backup time.

Example 20 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 `DBMIDData.808` and the backup of datafile 2 is on the volume `DBMIDData.802`. This output also shows that the volume `DBMIDData.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: 061dtl86_1_1 Media: DBMIDData.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: 081dtla0_1_1 Media: DBMIDData.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: 061dtl86_1_1 Media: DBMIDData.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 that are required for restoring the backup pieces.

The `nsrorainfo` command provides the following volume information for the restore of each backup piece:

Example 20 Using `restore...preview` to determine volumes for restore (continued)

- 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 that are 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 21](#) and [Example 22](#) describe how to use the `nsrorainfo` command and the command output.

Example 20 Using `restore...preview` to determine volumes for restore (continued)

The *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsrorainfo` command.

Example 21 Sample `nsrorainfo` commands for Oracle restores

Each of the following `nsrorainfo` commands provides a list of the volumes that are 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 that are 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 that are 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 22 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:

Example 22 Volume information that the `nsroraifno` command provides (continued)

- 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.

About this task

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 that is 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 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 476 provides details.

[Common NMDA parameters](#) on page 406 and [NMDA Oracle parameters](#) on page 444 describe the common parameters and Oracle parameters, respectively.

Example 23 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 243:

```
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';
```

Example 23 RMAN script for a tablespace restore (continued)

```

restore tablespace users;
recover tablespace users;
sql 'alter tablespace users online';
release channel t1;
release channel t2;
}

```

[RMAN scripts for manual backups](#) on page 135 describes the setting of NMDA parameters in an RMAN script.

Example 24 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](#) provides details.

The following RMAN script performs a restore of the database from the specified volume pool that is 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;
}

```

Performing Oracle restore and recovery with the NMDA Oracle recovery wizard

You can use the NMDA Oracle recovery wizard to create an RMAN restore/recovery script and then run the restore and recovery on the Oracle host.

About this task


[NMDA Oracle recovery wizard](#) on page 50 provides more details about the NMDA Oracle recovery wizard.

Before you use the NMDA Oracle recovery 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
- 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.

You can use the following procedure to perform a restore and recovery with the wizard.

 **Note:** For a point-in-time recovery with the wizard, the subsequent topic provides details.

Procedure

1. Start the NetWorker Management Console (NMC) 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 **Protection**.
4. In the **Protection** window, click **Clients**.
5. To start the wizard, right-click the NMDA client in the right pane, and then 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.

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 change the availability of a tablespace.

You can select to start the restore and recovery immediately from the wizard or schedule the operation to start later.

In the wizard, you can access an existing recover configuration at a later time or you can view the recovery results. In NMC, click **Recover** on the **Administration** window toolbar to open the **Recover** window. In the **Configured Recovers** pane, right-click the saved recover configuration and select one of the menu options:

- **New Recover**—Select this option as another way to create a new recover configuration.
- **Open Recover**—Select this option to view the recovery results.
- **Recover Again**—Select this option to access an existing recover configuration. You can make changes and save the configuration with a new name if required.

Performing Oracle point-in-time restore and recovery with the NMDA Oracle recovery wizard

You can use the NMDA Oracle recovery wizard to create and run the required RMAN script to perform a point-in-time restore and recovery.

About this task

[NMDA Oracle recovery wizard](#) on page 50 provides more details about the NMDA Oracle recovery wizard.

Before you use the NMDA Oracle recovery wizard, refer to the preceding topic for details on the recovery wizard requirements.

You can use the following procedure to perform a point-in-time restore and recovery with the wizard.

Procedure

1. Start the NetWorker Management Console (NMC) to access the Oracle recovery wizard, as described in the preceding topic, and click **New Recover**.
2. Select the NMDA client, select Oracle, and click **Next**.
3. Add the Oracle user details, and click **Next**.
4. Select the option of **Restore the database and logs**, and click **Next**.
5. Select the complete database to restore, and click the option.
6. Click **Yes** at the shutdown message, and start the database in the mount state.
7. Select the location to restore.
8. Under **Recovery time options**, select **Recover the database to a specific point in time** and select **Time stamp**.
9. Log in to the NetWorker server and browse for the latest and previous backups by using the following command:

```
# rmminfo -avot
```

```
cassini_c.dddefault.001 Data Domain maven 03/17/2019 08:49:01 PM 1234 MB 2123262108 cb
full RMAN:ORCL8_d2tsmitt_1_1
cassini.dddefault.001 Data Domain maven 03/17/2019 08:49:01 PM 1234 MB 2123262108 cb full
RMAN:ORCL8_d2tsmitt_1_1
cassini_c.dddefault.001 Data Domain maven 03/17/2019 08:49:27 PM 9003 MB 2106484918 cb
full RMAN:ORCL8_d3tsmiun_1_1
cassini.dddefault.001 Data Domain maven 03/17/2019 08:49:27 PM 9003 MB 2106484918 cb full
RMAN:ORCL8_d3tsmiun_1_1
cassini_c.dddefault.001 Data Domain maven 03/17/2019 08:52:13 PM 10 MB 2089707867 cb full
RMAN:ORCL8_d4tsmj3s_1_1
cassini.dddefault.001 Data Domain maven 03/17/2019 08:52:13 PM 10 MB 2089707867 cb full
RMAN:ORCL8_d4tsmj3s_1_1
cassini_c.dddefault.001 Data Domain maven 03/17/2019 08:52:17 PM 257 KB 2072930655 cb full
RMAN:ORCL8_d5tsmj41_1_1
cassini.dddefault.001 Data Domain maven 03/17/2019 08:52:17 PM 257 KB 2072930655 cb full
RMAN:ORCL8_d5tsmj41_1_1
cassini_c.dddefault.001 Data Domain maven 03/17/2019 08:52:33 PM 5 KB 2056153456 cb
full /nsr/apps/res/nwora.res
cassini.dddefault.001 Data Domain maven 03/17/2019 08:52:33 PM 5 KB 2056153456 cb
full /nsr/apps/res/nwora.res
```

10. From the command output, select the backup time that you want for the restore, add this time in the **Time stamp** option in the wizard, and click **Next**.

The wizard generates the RMAN script for the particular point-in-time recovery.

11. Name the recover configuration, and run the recovery from the wizard.

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.

About this task

If the RMAN restore script in [Example 23](#) is in the file `/disk1/scripts/restore.txt` and you configured the Net service to connect to the databases payroll and rcvcatdb, type the following 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.

About this task

i **Note:** NMDA does not support the use of the Oracle recovery 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.

i **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.

About this task

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 [Example 23](#). 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 SAP IQ data restore

Perform the SAP IQ data restore by using the procedures in the following topics:

- Prerequisites for SAP IQ data restore
- Performing the SAP IQ data restore with the `nsriqrc` command

Prerequisites for SAP IQ data restore

You must meet specific prerequisites for the NMDA SAP IQ data restore operations. NMDA can only restore the SAP IQ data that has been backed up through an NMDA SAP IQ backup.

Before you perform a restore operation, ensure that you meet the following prerequisites:

- The SAP IQ server must be running.
- The target database exists, to which data will be restored. For a database restore, the target database is at least as large as the size of the database backup.
- All-inclusive or read-write restore—The database must be shut down for the restore of an all-inclusive backup or read-write backup across all the backup levels. After the database is shut down, ensure that the `utility_db` database is running.
- Read-only dbspace restore—The database can be either running or shut down for the restore of a read-only dbspace backup. If the database is shut down, ensure that the dbspace is offline and the `utility_db` database is running.

You can restore only one read-only dbspace from a backup, which can be either an all-inclusive backup or a backup of one or more read-only dbspaces. You can restore the read-only dbspace only if the dbspace in the target database has not changed since the backup. For example, if the dbspace has changed to read-write since the backup, you cannot restore the dbspace from the backup.

- Read-only dbfile restore—The database must be running for the restore of a read-only dbfile backup.

You can restore only one read-only dbfile from a backup, which can be either an all-inclusive backup or a backup of one or more read-only dbfiles. You can restore the read-only dbfile only if the dbfile in the target database has not changed since the backup. For example, if the dbfile has changed to read-write since the backup, you cannot restore the dbfile from the backup.

- Point-in-time restore that uses the log backups—The database must be shut down. After the database is shut down, ensure that the `utility_db` database is running. The database directory and all the log directories that are listed in `IQ_PIT_RESTORE_LOG_PATH` are created by the administrator.
- The mandatory parameters are set in the NMDA SAP IQ restore configuration file:
 - `IQ_USER`, `NSR_SAVESET_NAME`, and `USER_PSWD` (when SAP IQ server has a password) must be set for all restore operations.
 - `IQ_SELECTIVE_TYPE` must be set for the restore of a selective backup, which is an all-inclusive, read-write, or read-only backup.
 - `IQ_READONLY_DBSPACES` must be set to `dbspace_name` for a restore of a single dbspace from an all-inclusive or read-only backup.
 - `IQ_READONLY_DBFILES` must be set to `dbfile_name` for a restore of a single dbfile from an all-inclusive or read-only backup.
 - `UTILDB_USER` and `UTILDB_PSWD` must be set if the database will be shut down during the restore. The `utility_db` database must also be running.

- `IQ_PIT_RESTORE_ENABLE`, `IQ_PIT_RESTORE_LOG_PATH`, `IQ_PIT_RESTORE_OFFSET`, and `NSR_RECOVER_TIME` must be set for a point-in-restore that uses the log backups.

NOTICE When the `IQ_OVERWRITE_EXISTING` parameter is not specified, you must delete or move the catalog store (`dbname.db`), IQ store files (`*.iq`), transaction logs (`*.log`), and user-defined stores. If any of these files reside in the target directory, the SAP IQ server generates an error and does not restore the backup files.

[NMDA SAP IQ parameters](#) on page 457 provides details about the NMDA SAP IQ parameters.

Performing the SAP IQ data restore with the `nsriqrc` command

Before you perform an SAP IQ restore, ensure that the restore prerequisites from the preceding topic have been met. You can perform a restore of an SAP IQ database backup or selective backup or a point-in-time restore that uses the transaction log backups, as described in the preceding topic. To perform the restore, run the `nsriqrc -z configuration_file_path` command from the command line as the operating system user that launched the SAP IQ server.

NOTICE In the `NSR_SAVESET_NAME` parameter setting for the restore, the SAP IQ database name is case-sensitive and the name must be in the same case as recorded in the corresponding backup entries in the NetWorker indexes.

When the `IQ_OVERWRITE_EXISTING` parameter is not specified, you must delete or move the catalog store (`dbname.db`), IQ store files (`*.iq`), transaction logs (`*.log`), and user-defined stores. If any of these files reside in the target directory, the SAP IQ server generates an error and does not restore the backup files.

By default, the `nsriqrc` command restores the most recent database backup. The command does not bring the database back online at the end of a restore.

An SAP IQ restore is performed to a specific directory:

- For backed-up database objects that are created with an absolute pathname, the files are restored to the original directory.
- For backed-up database objects that are created with a relative pathname:
 - If the restore uses the `utility_db` database, the files are restored to the directory where the `utility_db` database is running.
 - If the restore occurs for an active database, the database must be started from the directory where the `.db` or `.log` file resides. The backed-up files are restored to directories that are relative to the directory where the active database is running.

The *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsriqrc` command. [NMDA SAP IQ parameters](#) on page 457 provides details about the SAP IQ restore parameters in the NMDA configuration file.

After the restore operation completes, perform the following steps:

1. Check the database to ensure that the data has been restored.
2. Run a full backup of the database.

The following examples describe specific SAP IQ restore operations that NMDA supports:


- [Restore of the latest backup of an SAP IQ database](#) on page 246 describes a restore of the latest backup of a database.
- [Point-in-time restore of log backup, with offset after the previous full backup](#) on page 246 describes a point-in-time restore of a log backup, where the log offset occurs between the previous full backup and the specified transaction log backup.

- [Point-in-time restore of log backup, with offset before the previous full backup](#) on page 247 describes a point-in-time restore of a log backup, where the log offset occurs between the previous full backup and the transaction log backup just before that full backup.
- [Offline restore of read-write database files](#) on page 249 describes an offline restore of read-write database files.
- [Offline restore of a read-only dbspace](#) on page 249 describes a offline restore of a read-only dbspace.
- [Online restore of a read-only dbfile](#) on page 250 describes an online restore of a read-only dbfile.

Example 25 Restore of the latest backup of an SAP IQ database

You must set the following parameters in the NMDA configuration file for the SAP IQ database restore:

- IQ_USER—Specifies the username of the SAP IQ user account.
- NSR_SAVESET_NAME—Specifies the name of the SAP IQ database to restore.
- NSR_SERVER—Specifies the hostname of the NetWorker server.
- USER_PSWD—Specifies the encrypted password of the SAP IQ user account.

 **Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.


[NMDA SAP IQ parameters](#) on page 457 provides details about the SAP IQ restore parameters in the NMDA configuration file.

With the restore parameters set, you must run the following command to restore the database:

```
nsriqrc -z configuration_file_path
```

Example 26 Point-in-time restore of log backup, with offset after the previous full backup



For a point-in-time restore of a transaction log backup, you must specify the transaction log offset in the IQ_PIT_RESTORE_OFFSET parameter setting. In this example, the offset occurs between the previous full backup and the specified log backup.

 **Note:** If the offset occurs after the specified log backup and before the next full backup, then you must set the IQ_PIT_RESTORE_BEFORE_PREV_FULL parameter to TRUE.

You can run the `nsriqrc` command to perform the point-in-time recovery of the backup data and logs to the transaction log save set time that is specified in the NSR_RECOVER_TIME parameter setting.

You must set the following parameters in the NMDA configuration file for the point-in-time recovery:

Example 26 Point-in-time restore of log backup, with offset after the previous full backup (continued)

- `IQ_PIT_RESTORE_ENABLE`—Must be set to `TRUE` for the point-in-time restore of a log backup.
 - `IQ_PIT_RESTORE_LOG_PATH`—Specifies the pathnames of all the log files to be restored, including the active log of the database at the time to which the restore will be performed.
 - `IQ_PIT_RESTORE_OFFSET`—Specifies the transaction log offset to use for the point-in-time restore.
 - `IQ_USER`—Specifies the username of the SAP IQ user account.
 - `NSR_SAVESET_NAME`—Specifies the name of the SAP IQ database to restore.
 - `NSR_SERVER`—Specifies the hostname of the NetWorker server.
 - `NSR_RECOVER_TIME`—Specifies the transaction log save set time in either the "`MM/DD/YY HH:MM:SS`" format or the epoch time format. "`MM/DD/YY HH:MM:SS`" specifies the month, day, year, hour, minute, and seconds to which to recover the backup data and logs. The save set must contain the transaction for the offset that is specified in `IQ_PIT_RESTORE_OFFSET`.
-  **Note:** Because the NetWorker server and client can be in different time zones, set this parameter to the value of the local time on the SAP IQ server.
- `USER_PSWD`—Specifies the encrypted password of the SAP IQ user account.
-  **Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.


You can set the following optional parameters to enable specific supported features of the point-in-time restore:

- `IQ_CLEAR_LOG`
- `IQ_PIT_RESTORE_REDIRECTED_LOG_PATH`
- `NSR_DEBUG_LEVEL`
- `NSR_RECOVER_LOG_POOL`

[NMDA SAP IQ parameters](#) on page 457 provides details about the SAP IQ point-in-time restore parameters in the NMDA configuration file.

With the restore parameters set, you must run the following command to complete the point-in-time restore:

```
nsriqrc -z configuration_file_path
```

 **NOTICE** After you perform the point-in-time restore, it is recommended that you perform a full backup of the database.



Example 27 Point-in-time restore of log backup, with offset before the previous full backup

For a point-in-time restore of a transaction log backup, you must specify the transaction log offset in the `IQ_PIT_RESTORE_OFFSET` parameter setting. In this example, the offset occurs between the previous full backup and the log backup just

Example 27 Point-in-time restore of log backup, with offset before the previous full backup (continued) before that full backup, which requires the `IQ_PIT_RESTORE_BEFORE_PREV_FULL` parameter to be set to `TRUE`.

You can run the `nsriqrc` command to perform the point-in-time recovery of the backup data and logs to the transaction log save set time that is specified in the `NSR_RECOVER_TIME` parameter setting.

You must set the following parameters in the NMDA configuration file for the point-in-time recovery:

- `IQ_PIT_RESTORE_ENABLE`—Must be set to `TRUE` for the point-in-time restore of a log backup.
 - `IQ_PIT_RESTORE_LOG_PATH`—Specifies the pathnames of all the log files to be restored, including the active log of the database at the time to which the restore will be performed.
 - `IQ_PIT_RESTORE_OFFSET`—Specifies the transaction log offset to use for the point-in-time restore.
 - `IQ_PIT_RESTORE_BEFORE_PREV_FULL`—Must be set to `TRUE` when the offset is between the previous full backup and the log backup just before that full backup.
 - `IQ_USER`—Specifies the username of the SAP IQ user account.
 - `NSR_SAVESET_NAME`—Specifies the name of the SAP IQ database to restore.
 - `NSR_SERVER`—Specifies the hostname of the NetWorker server.
 - `NSR_RECOVER_TIME`—Specifies the save set time of the transaction log backup after (not before) the previous full backup, in either the "`MM/DD/YY HH:MM:SS`" format or the epoch time format. "`MM/DD/YY HH:MM:SS`" specifies the month, day, year, hour, minute, and seconds to which to recover the backup data and logs.
-  **Note:** Because the NetWorker server and client can be in different time zones, set this parameter to the value of the local time on the SAP IQ server.
- `USER_PSWD`—Specifies the encrypted password of the SAP IQ user account.
-  **Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.

You can set the following optional parameters to enable specific supported features of the point-in-time restore:

- `IQ_CLEAR_LOG`
- `IQ_PIT_RESTORE_REDIRECTED_LOG_PATH`
- `NSR_DEBUG_LEVEL`
- `NSR_RECOVER_LOG_POOL`

[NMDA SAP IQ parameters](#) on page 457 provides details about the SAP IQ point-in-time restore parameters in the NMDA configuration file.

With the restore parameters set, you must run the following command to complete the point-in-time restore:

Example 27 Point-in-time restore of log backup, with offset before the previous full backup (continued)

```
nsriqrc -z configuration_file_path
```

i **NOTICE** After you perform the point-in-time restore, it is recommended that you perform a full backup of the database.

Example 28 Offline restore of read-write database files

A restore of a read-write selective backup restores the complete set of read-write database files for the database. Before you perform the read-write restore, you must ensure that the database is shut down and the `utility_db` database is running.

You must set the following parameters in the NMDA configuration file for the read-write restore:

- `IQ_SELECTIVE_TYPE`—Set to `READWRITE` to specify the restore of a read-write selective backup.
- `IQ_USER`—Specifies the username of the SAP IQ user account.
- `NSR_SAVESET_NAME`—Specifies the name of the SAP IQ database for the restore.
- `NSR_SERVER`—Specifies the hostname of the NetWorker server.
- `USER_PSWD`—Specifies the encrypted password of the SAP IQ user account.
- `UTILDB_USER`—Specifies the username of the `utility_db` database user account.
- `UTILDB_PSWD`—Specifies the encrypted password of the `utility_db` database user account.

i **Note:** You must run the `nsrdaadmin -P iq -z configuration_file_path` command to add both the `USER_PSWD` and `UTILDB_PSWD` passwords to the configuration file.

[NMDA SAP IQ parameters](#) on page 457 provides details about the SAP IQ restore parameters in the NMDA configuration file.

With the restore parameters set, you must run the following command for the read-write restore:

```
nsriqrc -z configuration_file_path
```

Example 29 Offline restore of a read-only dbspace

When the database is shut down for a restore of a read-only dbspace, you must ensure that the dbspace is offline and the `utility_db` database is running. The operation restores a single dbspace from the specified backup.

Example 29 Offline restore of a read-only dbspace (continued)

You must set the following parameters in the NMDA configuration file for the selective restore of a read-only dbspace:

- `IQ_READONLY_DBSPACES`—Set to *dbspace_name* as the name of the single dbspace to restore.
- `IQ_SELECTIVE_TYPE`—Set to `READONLY` to specify the type of selective backup to be restored.
- `IQ_USER`—Specifies the username of the SAP IQ user account.
- `NSR_SAVESET_NAME`—Specifies the name of the SAP IQ database for the restore.
- `NSR_SERVER`—Specifies the hostname of the NetWorker server.
- `USER_PSWD`—Specifies the encrypted password of the SAP IQ user account.
- `UTILDB_USER`—Specifies the username of the `utility_db` database user account.
- `UTILDB_PSWD`—Specifies the encrypted password of the `utility_db` database user account.

Note: You must run the `nsrdaadmin -P iq -z configuration_file_path` command to add both the `USER_PSWD` and `UTILDB_PSWD` passwords to the configuration file.

[NMDA SAP IQ parameters](#) on page 457 provides details about the SAP IQ restore parameters in the NMDA configuration file.

With the restore parameters set, you must run the following command to restore the read-only dbspace:

```
nsriqrc -z configuration_file_path
```


Example 30 Online restore of a read-only dbfile

You must ensure that the database is running for the restore of a read-only dbfile. The operation restores a single dbfile from the specified backup.

You must set the following parameters in the NMDA configuration file for the selective restore of a read-only dbfile:

- `IQ_READONLY_DBFILES`—Set to *dbfile_name* as the name of the single dbfile to restore.
- `IQ_SELECTIVE_TYPE`—Set to `READONLY` to specify the type of selective backup to be restored.
- `IQ_USER`—Specifies the username of the SAP IQ user account.
- `NSR_SAVESET_NAME`—Specifies the name of the SAP IQ database for the restore.
- `NSR_SERVER`—Specifies the hostname of the NetWorker server.
- `USER_PSWD`—Specifies the encrypted password of the SAP IQ user account.

Example 30 Online restore of a read-only dbfile (continued)

 **Note:** You must run the `nsrdaadmin -P iq -z configuration_file_path` command to add the encrypted password to the configuration file.

[NMDA SAP IQ parameters](#) on page 457 provides details about the SAP IQ restore parameters in the NMDA configuration file.

With the restore parameters set, you must run the following command to restore the read-only dbfile:


```
nsriqrc -z configuration_file_path
```

Performing Sybase data restore and recovery

To prepare for a Sybase data restore or recovery, ensure that the required servers are running and the target database is ready for the restore operation.

Before you perform a Sybase data restore and recovery, ensure that you 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.

You can use the following procedures to prepare for Sybase data restore and recovery:

- [Specifying verification of Sybase database restores](#) on page 251
- [Obtaining backup information with the `listonly` option](#) on page 252
- [Performing restores of Sybase ASE backups with a large number of databases](#) on page 254

Perform the Sybase data restore and recovery by using the appropriate procedures:

- [Performing Sybase data restores with the `nsrsybrc` command](#) on page 254
- [Performing Sybase data restore and recovery with the NMDA Sybase recovery wizard](#) on page 260
- [Performing Sybase data restores with NetWorker User for Sybase](#) on page 262

 **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 the verification of a Sybase database restore at the header verification and full verification levels. NMDA also supports the verification of minimal header information without the restore of the Sybase database.

About this task

Set the `NSR_ASE_VERIFY` parameter in the NMDA configuration file to specify the restore 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 restore).
- **verifyonly**—Specifies to verify minimal header information without restoring the database.

For example, the following NSR_ASE_VERIFY setting specifies to perform a full verification of the restore:

```
NSR_ASE_VERIFY=full
```

If you do not specify a verification value, then the restore performs no verification but adds a message to the log file.

Obtaining backup information with the listonly option

You can use the `nsrsybrc` command with the `listonly` option to obtain structure information from the latest backup of a specified Sybase instance or databases, without restoring the backup. Sybase ASE 15.7 ESD#2 and later provides the `listonly` option support.

About this task

After you set the required parameters in the NMDA configuration file, you can run the `nsrsybrc -z configuration_file_path` command to display the backup structure information. The `nsrsybrc` program displays the backup information by using the `listonly` option of the Sybase `load` command.

Complete the following steps to obtain the backup information through the `listonly` option.

Procedure

1. Start an `isql` session, and then run the following `sp_configure` command to enable the Sybase resource for the `listonly` option:

```
sp_configure 'enable dump history',1
```

2. Set the required Sybase parameters in the NMDA configuration file:

- **NSR_BACKUP_PATHS**—Specify the name of the Sybase instance, a single database, or multiple databases for which the backup structure information will be displayed.
- **NSR_SERVER**—Specify the hostname of the NetWorker server that is used for the backup.
- **RECOVER_LISTONLY**—Specify `load_sql` or `create_sql` value, depending on the type of backup information to display.

Note: Instead of setting `RECOVER_LISTONLY` in the configuration file, you can specify the `-l load_sql` or `-l create_sql` option on the `nsrsybrc` command line. The latest *NetWorker Module for Databases and Applications Command Reference Guide* provides details.

- **SYBASE_USER**—Specify the username of the Sybase user account.
 - **USER_PSWD**—Specify the encrypted password of the Sybase user account.
- Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.

[NMDA Sybase parameters](#) on page 464 provides details about the Sybase parameters in the NMDA configuration file.

- To obtain the backup information, run the following command as the Sybase user:

```
nsrsybrc -z configuration_file_path
```

Example 31 Using nsrsybrc with the listonly option to display backup information

The NMDA configuration file includes the parameter setting `NSR_BACKUP_PATHS=SYBASE:/sybase160/test`, which specifies the Sybase server instance `sybase160` and the database named `test`.

With the parameter setting `RECOVER_LISTONLY=load_sql` in the configuration file, the `nsrsybrc` command displays the following sequence of SQL commands that can be used to restore the latest backup of the test database:

```
nsrsybrc -z configuration_file_path

Changed database context to 'master'.
LOAD DATABASE test FROM 'nsrsyb:..sybase160.test..7./nsr/apps/tmp/
sybtmp_20306.txt'
go
LOAD TRAN test FROM 'nsrsyb:..sybase160.test..7./nsr/apps/tmp/
sybtmp_20413.txt'
go
3:nsrsybrc: The SQL statements are generated successfully for
database 'test'. The database has not been restored.
```

With the parameter setting `RECOVER_LISTONLY=create_sql` in the configuration file, the `nsrsybrc` command displays the following sequence of commands that are obtained from the latest `dump` image in the Sybase backup history:

```
nsrsybrc -z configuration_file_path

Changed database context to 'master'.
DISK INIT
  name = 'testdata'
  , physname = '/tmp/t_dat'
  , size = '10M'
  , directio = true
go
DISK INIT
  name = 'testlog'
  , physname = '/tmp/t_log'
  , size = '10M'
  , directio = true
go
CREATE DATABASE test
  ON testdata = '10M'
LOG ON testlog = '8M'
  , testlog = '2M'
go
3:nsrsybrc: The SQL statements are generated successfully for
database 'test'. The database has not been restored.
```

Performing restores of Sybase ASE backups with a large number of databases

NMDA supports the backup and recovery of a Sybase ASE instance that contains more than 200 databases. You can set the `NSR_CONCURRENCY_MODE` parameter to specify the method that NMDA uses to restore the backup of the databases.

The parallelism setting also affects the restore of the backup, where the parallelism setting is the minimum of the `NSR_PARALLELISM`, client parallelism, and server parallelism settings.

Set the `NSR_CONCURRENCY_MODE` parameter to the appropriate value for the preferred restore of the backup:

- Set `NSR_CONCURRENCY_MODE` to the "stripe" (default) value to specify that NMDA restores only one database at a time.
- Set `NSR_CONCURRENCY` to the "database" value to specify that NMDA restores multiple databases concurrently. The parallelism setting determines the maximum number of databases that NMDA can restore concurrently.

i **Note:** If the parallelism setting is too low for a single database to be restored, NMDA restores only one database at a time.

The `NSR_CONCURRENCY_MODE` description in [Table 47](#) on page 464 provides details about the parameter setting.

In most cases, do not set `NSR_CONCURRENCY_MODE` to the "database" value for a recovery unless the ASE instance includes a very large number of databases (1000 or more). The recovery time improvements with the setting are not as significant as the backup time improvements, and the recovery is more likely to fail if the ASE system is not configured properly.

Performing Sybase data restores with the `nsrsybrc` command

After you have set the required restore parameters in the NMDA configuration file, you can run the `nsrsybrc -z configuration_file_path` command from the command line to perform a Sybase restore. You can run the command as the operating system user that launched the Sybase Backup Server, to perform a restore of the whole Sybase server, individual databases, or transaction logs.

i **NOTICE** In the `NSR_BACKUP_PATHS` parameter setting for the restore, 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 *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsrsybrc` command. [NMDA Sybase parameters](#) on page 464 describes the Sybase restore parameters that are supported in the NMDA configuration file.

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 175.
3. Run a full backup of the database.

The following examples describe specific restore and recovery operations that NMDA supports:

- [Example 31](#) describes a restore of the Sybase server and a single database or multiple databases.
- [Example 32](#) describes a Sybase point-in-time recovery.

- [Example 33](#) describes a restore of Sybase transaction logs only.
- [Example 34](#) describes a Sybase cumulative restore.


The following topics describe the other supported types of Sybase data restores:

- [Performing Sybase redirected restores](#) on page 258
- [Performing Sybase restores to a different host](#) on page 260

Example 32 Restores of a single Sybase database and multiple databases

You can set the following parameters in the NMDA configuration file for the restore of a single Sybase database or multiple databases:

- NSR_BACKUP_PATHS—Specifies the names of the Sybase server and databases to restore.
- NSR_SERVER—Specifies the hostname of the NetWorker server.
- SYBASE_USER—Specifies the username of the Sybase user account.
- USER_PSWD—Specifies the encrypted password of the Sybase user account.

 **Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.

The correct setting of the NSR_BACKUP_PATHS parameter depends on the type of Sybase restore:

- For the restore of a single database, NSR_BACKUP_PATHS must specify the names of the Sybase server and the database on the Sybase server.
- For the restore of multiple databases, NSR_BACKUP_PATHS must specify the multiple database names.
- For the restore of all databases on the Sybase server except the master database, NSR_BACKUP_PATHS must specify the Sybase server name only.

[NMDA Sybase parameters](#) on page 464 provides details about the Sybase restore parameters in the NMDA configuration file.

With the restore parameters set, you can run the following command to restore the single database or multiple databases:

```
nsrsybrc -z configuration_file_path
```

Example 33 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. To recover the data to a specific point-in-time, the `nsrsybrc` command uses the time that is specified in the RECOVER_UNTIL parameter setting in the NMDA configuration file. The operation loads the most recent full backup before the designated time and then applies transaction log backups up to the designated time.

You can set the following parameters in the NMDA configuration file for the point-in-time recovery:

Example 33 Point-in-time recovery of Sybase data (continued)

- NSR_BACKUP_PATHS—Specifies the names of the Sybase server and database to restore.
- NSR_SERVER—Specifies the hostname of the NetWorker server.
- RECOVER_UNTIL—Specifies "MM/DD/YY HH:MM:SS" for the month, day, year, hour, minute, and seconds to which to recover the data.
 - ⓘ **Note:** Because the NetWorker server and client can be in different time zones, set this parameter to the value of the local time on the Sybase server.
- SYBASE_USER—Specifies the username of the Sybase user account.
- USER_PSWD—Specifies the encrypted password of the Sybase user account.
 - ⓘ **Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.

[NMDA Sybase parameters](#) on page 464 provides details about the Sybase restore parameters in the NMDA configuration file.

With the restore parameters set, you can run the following command to complete the point-in-time recovery:

```
nsrsybrc -z configuration_file_path
```

A successful point-in-time recovery brings the database online.

- ⓘ **NOTICE** After you perform the 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.

Example 34 Restores of Sybase transaction logs only

NMDA supports the restore of only a transaction log backup, which enables you to update a database by restoring only the transaction logs. You must initially recover the database to a point-in-time and keep the database offline after the recovery. Then you can restore one or more transaction log backups to update the database to one or more specific points-in-time.

You must meet the following requirements for the restore of only Sybase transaction log backups:

- The database must be offline during the entire restore process.
- You must restore the backups in the correct order so that all the required data is restored.

Use one of the following methods to restore a transaction log backup:

- Set `RECOVER_SAVETIME=TRUE`, and set `RECOVER_UNTIL` to the exact savetime by using the date and time format as described in the preceding example.

Example 34 Restores of Sybase transaction logs only (continued)

- Run the `nsrsybrc` command with the `-t` and `-x` options. The latest *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsrsybrc` command and options.

NMDA supports the `RECOVER_SAVETIME` parameter and the `nsrsybrc -x` option.

For example, you perform a full (database) backup at the time T1 and incremental (log) backups at the times T2 and T3. Then you can perform the following sequence of recoveries.

1. To recover the database to the time T2 and keep the database offline after the recovery, set the following parameters:

```
RECOVER_UNTIL=T2
DATABASE_ONLINE=FALSE
```

The T2 value uses the data and time format, "*MM/DD/YY HH:MM:SS*".

2. To recover the database to the time T3 by restoring only the log backup, set the following parameters. The database must remain offline during the restore.

```
RECOVER_UNTIL=T3
RECOVER_SAVETIME=TRUE
```

The T3 value uses the data and time format, "*MM/DD/YY HH:MM:SS*".

When you set `RECOVER_SAVETIME=TRUE` and set `RECOVER_UNTIL` to the backup savetime, only the backup from the time T3 is restored, which is the log backup. The restore succeeds only if `RECOVER_UNTIL` is set to the exact savetime of the backup.

If a backup with the exact savetime is not found, the restore fails. If `RECOVER_UNTIL` is set to an invalid time, then the `RECOVER_SAVETIME` parameter is ignored and a regular recovery is performed to the time T3.

Example 35 Sybase cumulative restores

By default, NMDA Sybase restores all the supported types of backups, including cumulative backups. A cumulative backup includes all the database changes since the last full database backup. If required, you can restore only full and incremental backups by disabling the restore of cumulative backups.

Use one of the following methods to disable the restore of cumulative backups:

- Set `USE_CUMULATIVE=FALSE` in the NMDA configuration file.
- Run the `nsrsybrc` command with the `-C` option. The latest *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsrsybrc` command and options.

Performing Sybase redirected restores

A redirected restore loads the backup of an old database to a new database.

About this task

NMDA supports the following types of redirected restores of Sybase data:


- Restore to the same Sybase server but a different database name
- Restore to a different Sybase server but the same database name
- Restore to a different Sybase server and a different database name
- Restore to different database types

Restoring to the same Sybase server but a different database name

Before you restore data to the same Sybase server but to a different database name, you must create the new database and set the required restore parameters. You can run the `nsrsybrc` command to complete the Sybase restore.

Procedure

1. Ensure that you have created the new database with the proper device allocations.
2. Set the following Sybase restore parameters in the NMDA configuration file:
 - `NSR_BACKUP_PATHS`—Specifies the names of the Sybase server and the original database that was backed up.
 - `NSR_SERVER`—Specifies the hostname of the NetWorker server.
 - `NSR_RELOCATION_DEST`—Specifies the names of the Sybase server and the new database to which the data will be restored.
 - `SYBASE_USER`—Specifies the username of the Sybase user account.
 - `USER_PSWD`—Specifies the encrypted password of the Sybase user account.

 **Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.

[NMDA Sybase parameters](#) on page 464 provides details about the Sybase restore parameters in the NMDA configuration file.
3. To complete the restore, run the following command as the Sybase user:

```
nsrsybrc -z configuration_file_path
```


Restoring to a different Sybase server but the same database name

Before you restore data to a different Sybase server but the same database name, you must create the new database and set the required restore parameters. You can run the `nsrsybrc` command to complete the Sybase restore.

Procedure

1. Ensure that you have created the new database with the proper device allocations.
2. Set the following Sybase restore parameters in the NMDA configuration file:
 - `NSR_BACKUP_PATHS`—Specifies the names of the original Sybase server and the original database that was backed up.
 - `NSR_SERVER`—Specifies the hostname of the NetWorker server.

- NSR_RELOCATION_DEST—Specifies the names of the new Sybase server and the database to which the data will be restored.
- SYBASE_USER—Specifies the username of the Sybase user account on the new Sybase server.
- USER_PSWD—Specifies the encrypted password of the Sybase user account on the new Sybase server.

 **Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.

[NMDA Sybase parameters](#) on page 464 provides details about the Sybase restore parameters in the NMDA configuration file.

3. To complete the restore, run the following command as the Sybase user:


```
nsrsybrc -z configuration_file_path
```

Restoring to a different Sybase server and a different database name

Before you restore data to a different Sybase server and a different database name, you must create the new database and set the required restore parameters. You can run the `nsrsybrc` command to complete the Sybase restore.

Procedure

1. Ensure that you have created the new database with the same device allocations as the original database.
2. Set the following Sybase restore parameters in the NMDA configuration file:
 - NSR_BACKUP_PATHS—Specifies the names of the original Sybase server and the original database that was backed up.
 - NSR_SERVER—Specifies the hostname of the NetWorker server.
 - NSR_RELOCATION_DEST—Specifies the names of the new Sybase server and the new database to which the data will be restored.
 - SYBASE_USER—Specifies the username of the Sybase user account on the new Sybase server.
 - USER_PSWD—Specifies the encrypted password of the Sybase user account on the new Sybase server.

 **Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.

[NMDA Sybase parameters](#) on page 464 provides details about the Sybase restore parameters in the NMDA configuration file.

3. To complete the restore, run the following command as the Sybase user:

```
nsrsybrc -z configuration_file_path
```

Restoring to different database types

With Sybase ASE 15.7 or later, 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.

About this task


The restore instructions are the same as in the preceding topics for a redirected restore.


Performing Sybase restores to a different host

Before you restore Sybase data to a new destination host, you must set up the NMDA software on the new host and set the required restore parameters. You can run the `nsrsybrc` command to complete the Sybase restore.

Procedure

1. Follow the instructions in [Performing NMDA data restore and recovery](#) on page 191 to set up the NMDA software.
2. Set the following Sybase restore parameters in the NMDA configuration file:
 - `NSR_BACKUP_PATHS`—Specifies the names of the original Sybase server and the original database that was backed up.
 - `NSR_CLIENT`—Specifies the hostname of the NetWorker client where the backup was performed.
 - `NSR_SERVER`—Specifies the hostname of the NetWorker server.
 - `NSR_RELOCATION_DEST`—Specifies the names of a new Sybase server or new database to which the data will be restored.

 **Note:** You can omit the `NSR_RELOCATION_DEST` setting if the destination server name and database name are the same as on the original NetWorker client.
 - `SYBASE_USER`—Specifies the username of the Sybase user account on the destination Sybase server.
 - `USER_PSWD`—Specifies the encrypted password of the Sybase user account on the destination Sybase server.

 **Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.

[NMDA Sybase parameters](#) on page 464 provides details about the Sybase restore parameters in the NMDA configuration file.

3. To complete the restore on the destination host, run the following command as the Sybase user:

```
nsrsybrc -z configuration_file_path
```

Performing Sybase data restore and recovery with the NMDA Sybase recovery wizard

You can use the NMDA Sybase recovery wizard to configure and run the restore and recovery of Sybase ASE data that is backed up by NMDA.

About this task

[NMDA Sybase recovery wizard](#) on page 65 provides more details about the NMDA Sybase recovery wizard.

Before you use the NMDA Sybase recovery 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.
 - Client-side configuration method without the wizard, where the value of the Save Set attribute of the Client resource has the SYBASE: prefix.

You can use the following procedure to perform a restore and recovery with the wizard.

Procedure

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 **Protection**.
4. In the **Protection** window, click **Clients**.
5. To start the wizard, right-click the NMDA client in the right pane, and then select **Recover**.
6. On each wizard screen that appears, specify the required values for the restore and recovery 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.

You can select to start the restore or recovery immediately from the wizard or schedule the operation to start later.

In the wizard, you can access an existing recover configuration at a later time or you can view the recovery results. In NMC, click **Recover** on the **Administration** window toolbar to open the **Recover** window. In the **Configured Recovers** pane, right-click the saved recover configuration and select one of the menu options:

- **New Recover**—Select this option as another way to create a new recover configuration.
- **Open Recover**—Select this option to view the recovery results.
- **Recover Again**—Select this option to access an existing recover configuration. You can make changes and save the configuration with a new name if required.

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.

About this task

You must run the Sybase GUI on the destination host where the data will be restored.

Procedure

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. To refresh the list of NetWorker servers, click **Update List**.
 - c. Select or type the name of the server.
 - d. To use the server as the default NetWorker server, select **Save as Default 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. To change from the current month, click **Previous Month** or **Next 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 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 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.

Performing Orchestrated Application Protection data restore

To prepare for the restore of an Orchestrated Application Protection backup, complete the required configuration procedures and ensure that the specified directory has sufficient space to contain the restored backup data. Perform the data restore by running the `nsroapprecover` program, which retrieves the backup data into the specified directory.

The `nsroapprecover` program must retrieve the data from the DD Boost device. If the backed-up data has been cloned to conventional media, such as an AFTP or TYPE device, you must clone the requested save set to the DD Boost device before you perform the restore.

Before you perform an Orchestrated Application Protection data restore, ensure that you meet the following requirements:

- The save set to be restored is on the DD Boost device. If the save set is not on the device, clone the save set to the DD Boost device.
- You have set the required parameters in the NMDA configuration file. You can use the same NMDA configuration file for the Orchestrated Application Protection restore as was used for the Orchestrated Application Protection backup. [NMDA Orchestrated Application Protection parameters](#) on page 451 describes all the supported parameters.
- The local directory that is specified by `NSR_RELOCATION_DEST` or its parent directory has sufficient space to contain the backed-up data that is restored from the NetWorker server. You can use the output from the NetWorker `mminfo` command to determine the size of the save set to be restored.
- The directory that is specified by `NSR_RELOCATION_DEST` does not exist before the restore as the `nsroapprecover` program will create the directory. The operating system user that runs the `nsroapprecover` program has full control of that directory's parent directory.

The latest version of the *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsroapprecover` command.

The following topics provide details on the specific procedures for the restore of a MongoDB backup, MySQL backup, or PostgreSQL backup that was performed through Orchestrated Application Protection:

- [Performing the restore of MongoDB data](#) on page 264
- [Performing the restore of MySQL data](#) on page 266
- [Performing the restore of PostgreSQL data](#) on page 268

Performing the restore of MongoDB data

Perform the Orchestrated Application Protection restore of MongoDB data by using the procedures in the following topics:

- Querying an existing backup
- Completing the restore configuration
- Retrieving the backed-up MongoDB data with the `nsroapprecover` command
- Restoring the retrieved data to the MongoDB database

Querying an existing backup

You can determine the backup save set name from configuration parameter settings and then query the existing backups to determine the save time of the backup save set. You can use the query results to complete the configuration for the Orchestrated Application Protection restore.

You can use the values of the backup parameters `NSR_DATABASE_TYPE`, `NSR_INSTANCE_NAME`, and `NSR_BACKUP_NAME` to determine the save set name of the backup to be queried. Based on these parameter settings, the MongoDB save set name is as follows:

```
<NSR_DATABASE_TYPE> <NSR_INSTANCE_NAME> <NSR_BACKUP_NAME>_full
```

For example, the NMDA backup configuration file includes the following parameter settings:

```
<NSR_BACKUP_NAME> MY_TEST </NSR_BACKUP_NAME>
<NSR_DATABASE_TYPE> MongoDB </NSR_DATABASE_TYPE>
<NSR_INSTANCE_NAME> MY_DEMO </NSR_INSTANCE_NAME>
```

In this case, the backup save set name is `MongoDB_MY_DEMO_MY_TEST_full`.

When you have determined the save set name, you can run the `mminfo -vV | grep -A 2 save_set_name` command to display the backup save set information, including the save times of the save sets.

For example, the following `mminfo` command displays two corresponding save sets with the save times 1516039064 and 1516041320:

```
mminfo -vV | grep -A 2 MongoDB_MY_DEMO_MY_TEST_full
```

```
OAPPPool1.001      bu-wildcat.lss.emc.com 546 KB full
MongoDB_MY_DEMO_MY_TEST_full
4150061977 1516039064      01/15/18 12:57:44 02/15/18 02/15/18
0 558743 0 0 4180074071 558744 cr
OAPPPool1.001      bu-wildcat.lss.emc.com 546 KB full
MongoDB_MY_DEMO_MY_TEST_full
4133287017 1516041320      01/15/18 13:35:20 02/15/18 02/15/18
0 558743 0 0 4180074071 558744 cr
```

Once you have the save set name and save time, you can complete the configuration for the Orchestrated Application Protection restore.

Completing the restore configuration

To enable the NMDA Orchestrated Application Protection restore of a MongoDB backup, ensure that the required environment variables and NMDA Orchestrated Application Protection parameters are set.

The MongoDB server documentation provides more information on the required environment variables.

Ensure that all the required NMDA Orchestrated Application Protection parameters are set in the NMDA configuration file for the restore. The parameters `NSR_SAVESET_NAME` and `NSR_RELOCATION_DEST` are mandatory for a restore operation:

- For the `NSR_SAVESET_NAME` parameter setting, specify the backup save set name, as determined by the backup query results in the preceding topic.
- For the `NSR_RELOCATION_DEST` parameter setting, specify the complete pathname of the target directory that will contain the restored data. The target directory must not exist before the restore as the `nsroapprecovery` program creates the directory. Ensure that the user that

runs `nsroapprecover` has full permissions in the parent directory of the destination directory.

- For the optional `NSR_RECOVER_TIME` parameter setting, you can specify the save time of the backup save set, as determined by the backup query results in the preceding topic.

If you want to restore the most recent backup, you can omit the `NSR_RECOVER_TIME` setting.

For example, the following parameter settings in the `RECOVER` section of the NMDA configuration file include the save time of the backup save set to be restored:

```
<RECOVER>
  <NSR_SAVESET_NAME> MongoDB_MY_DEMO_MY_TEST_full </NSR_SAVESET_NAME>
  <NSR_RELOCATION_DEST> /nsr/apps/tmp/my_test </NSR_RELOCATION_DEST>
  <NSR_RECOVER_TIME> 1516039064 </NSR_RECOVER_TIME>
</RECOVER>
```

[NMDA parameters for Orchestrated Application Protection restores](#) on page 456 provides complete details on the restore parameters to set in the NMDA configuration file. For an Orchestrated Application Protection restore, the NMDA configuration file must include the proper XML formats as described in [NMDA configuration file](#) on page 400.

Retrieving the backed-up MongoDB data with the `nsroapprecover` command

Before you perform a MongoDB restore, ensure that the restore requirements from the preceding topics have been met. You can retrieve the MongoDB database backup data by running the `nsroapprecover` command from the command line as the operating system user that launched the MongoDB server.

To retrieve the MongoDB backup from NetWorker server, run the `nsroapprecover -z configuration_file_path` command.

The `nsroapprecover` program retrieves the MongoDB backup from the NetWorker server into the destination directory that is specified by the `NSR_RELOCATION_DEST` parameter. By default, unless you set `NSR_RECOVER_TIME`, the `nsroapprecover` command restores the most recent backup. The preceding topic provides more information about the restore parameter settings.

The *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsroapprecover` command.

Restoring the retrieved data to the MongoDB database

After you run the `nsroapprecover` command to retrieve the backup into the specified directory, perform the following steps:

1. Use the MongoDB restore utilities to complete the backup recovery. The MongoDB documentation provides more details on how to use the restore utilities.
2. Perform a full backup of the MongoDB database.

 **Note:** A full backup of the MongoDB database is required after a backup recovery.

Performing the restore of MySQL data

Perform the Orchestrated Application Protection restore of MySQL data by using the procedures in the following topics:

- Querying an existing backup
- Completing the restore configuration
- Retrieving the backed-up MySQL data with the `nsroapprecover` command

- Restoring the retrieved data to the MySQL database

Querying an existing backup

You can determine the backup save set name from configuration parameter settings and then query the existing backups to determine the save time of the backup save set. You can use the query results to complete the configuration for the Orchestrated Application Protection restore.

You can use the values of the backup parameters `NSR_DATABASE_TYPE`, `NSR_INSTANCE_NAME`, and `NSR_BACKUP_NAME` to determine the save set name of the backup to be queried. Based on these parameter settings, the MySQL save set names are as follows:

```
<NSR_DATABASE_TYPE> <NSR_INSTANCE_NAME> <NSR_BACKUP_NAME>_full
<NSR_DATABASE_TYPE> <NSR_INSTANCE_NAME> <NSR_BACKUP_NAME>_txnlog
```

For example, the NMDA backup configuration file includes the following parameter settings:

```
<NSR_BACKUP_NAME> MY_TEST </NSR_BACKUP_NAME>
<NSR_DATABASE_TYPE> MySQL </NSR_DATABASE_TYPE>
<NSR_INSTANCE_NAME> MY_DEMO </NSR_INSTANCE_NAME>
```

In this case, the backup save set name is `MySQL_MY_DEMO_MY_TEST_full` for the full database backup, and `MySQL_MY_DEMO_MY_TEST_txnlog` for the transaction log (Logs Only level) backup.

When you have determined the save set name, you can run the `mminfo -vV | grep -A 2 save_set_name` command to display the backup save set information, including the save times of the save sets.

For example, the following `mminfo` command displays two corresponding save sets with the save times 1516039064 and 1516041320:

```
mminfo -vV | grep -A 2 MySQL_MY_DEMO_MY_TEST_full
```

```
OAPPPool1.001      bu-wildcat.lss.emc.com 546 KB full
MySQL MY DEMO MY TEST_full
4150061977  1516039064      01/15/18 12:57:44  02/15/18 02/15/18
                0      558743      0      0 4180074071      558744 cr
OAPPPool1.001      bu-wildcat.lss.emc.com 546 KB full
MySQL MY DEMO MY TEST_full
4133287017  1516041320      01/15/18 13:35:20  02/15/18 02/15/18
                0      558743      0      0 4180074071      558744 cr
```

Once you have the save set name and save time, you can complete the configuration for the Orchestrated Application Protection restore.

Completing the restore configuration

To enable the NMDA Orchestrated Application Protection restore of a MySQL backup, ensure that the required environment variables and NMDA Orchestrated Application Protection parameters are set.

The MySQL documentation provides more information on the required environment variables.

Ensure that all the required NMDA Orchestrated Application Protection parameters are set in the NMDA configuration file for the restore. The parameters `NSR_SAVESET_NAME` and `NSR_RELOCATION_DEST` are mandatory for a restore operation:

- For the `NSR_SAVESET_NAME` parameter setting, specify the backup save set name, as determined by the backup query results in the preceding topic.

- For the `NSR_RELOCATION_DEST` parameter setting, specify the complete pathname of the target directory that will contain the restored data. The target directory must not exist before the restore as the `nsroapprecover` program creates the directory. Ensure that the user that runs `nsroapprecover` has full permissions in the parent directory of the destination directory.
- For the optional `NSR_RECOVER_TIME` parameter setting, you can specify the save time of the backup save set, as determined by the backup query results in the preceding topic.

If you want to restore the most recent backup, you can omit the `NSR_RECOVER_TIME` setting.

For example, the following parameter settings in the `RECOVER` section of the NMDA configuration file include the save time of the backup save set to be restored:

```
<RECOVER>
  <NSR_SAVESET_NAME> MySQL_MY_DEMO_MY_TEST_full </NSR_SAVESET_NAME>
  <NSR_RELOCATION_DEST> /nsr/apps/tmp/my_test </NSR_RELOCATION_DEST>
  <NSR_RECOVER_TIME> 1516039064 </NSR_RECOVER_TIME>
</RECOVER>
```

[NMDA parameters for Orchestrated Application Protection restores](#) on page 456 provides complete details on the restore parameters to set in the NMDA configuration file. For an Orchestrated Application Protection restore, the NMDA configuration file must include the proper XML formats as described in [NMDA configuration file](#) on page 400.

Retrieving the backed-up MySQL data with the `nsroapprecover` command

Before you perform a MySQL restore, ensure that the restore requirements from the preceding topics have been met. You can retrieve a MySQL database backup or transaction log backup by running the `nsroapprecover` command from the command line as the operating system user that launched the MySQL server.

To retrieve the MySQL full backup or transaction log backup from the NetWorker server, run the `nsroapprecover -z configuration_file_path` command.

The `nsroapprecover` program retrieves the MySQL backup from the NetWorker server into the destination directory that is specified by the `NSR_RELOCATION_DEST` parameter. By default, unless you set `NSR_RECOVER_TIME`, the `nsroapprecover` command restores the most recent backup. The preceding topic provides more information about the restore parameter settings.

The *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsroapprecover` command.

Restoring the retrieved data to the MySQL database

After you run the `nsroapprecover` command to retrieve the backup into the specified directory, perform the following steps:

1. Use the MySQL restore utilities to complete the backup recovery. The MySQL documentation provides more details on how to use the restore utilities.
2. Perform a full backup of the restored MySQL database before you use it.

 **Note:** A full backup of the MySQL database is required after a backup recovery.

Performing the restore of PostgreSQL data

Perform the Orchestrated Application Protection restore of PostgreSQL data by using the procedures in the following topics:

- Querying an existing backup

- Completing the restore configuration
- Registering the PostgreSQL `restore_command` for WAL archiving restore
- Retrieving the backed-up PostgreSQL data with the `nsroapprecovery` command
- Restoring the retrieved data to the PostgreSQL database

Querying an existing backup

You can determine the backup save set name from configuration parameter settings and then query the existing backups to determine the save time of the backup save set. You can use the query results to complete the configuration for the Orchestrated Application Protection restore.

You can use the values of the backup parameters `NSR_DATABASE_TYPE`, `NSR_INSTANCE_NAME`, and `NSR_BACKUP_NAME` to determine the save set name of the backup to be queried. Based on these parameter settings, the PostgreSQL save set name is as follows:

```
<NSR_DATABASE_TYPE>_<NSR_INSTANCE_NAME>_<NSR_BACKUP_NAME>_full
```

For example, the NMDA backup configuration file includes the following parameter settings:

```
<NSR_BACKUP_NAME> MY_TEST </NSR_BACKUP_NAME>
<NSR_DATABASE_TYPE> PostgreSQL </NSR_DATABASE_TYPE>
<NSR_INSTANCE_NAME> MY_DEMO </NSR_INSTANCE_NAME>
```

In this case, the backup save set name is `PostgreSQL_MY_DEMO_MY_TEST_full`.

When you have determined the save set name, you can run the `mminfo -vV | grep -A 2 save_set_name` command to display the backup save set information, including the save times of the save sets.

For example, the following `mminfo` command displays two corresponding save sets with the save times 1516039064 and 1516041320:

```
mminfo -vV | grep -A 2 PostgreSQL_MY_DEMO_MY_TEST_full
```

```
OAPPPool1.001      bu-wildcat.lss.emc.com 546 KB full
PostgreSQL_MY_DEMO_MY_TEST_full
4150061977 1516039064      01/15/18 12:57:44 02/15/18 02/15/18
0 558743 0 0 4180074071 558744 cr
OAPPPool1.001      bu-wildcat.lss.emc.com 546 KB full
PostgreSQL_MY_DEMO_MY_TEST_full
4133287017 1516041320      01/15/18 13:35:20 02/15/18 02/15/18
0 558743 0 0 4180074071 558744 cr
```

Once you have the save set name and save time, you can complete the configuration for the Orchestrated Application Protection restore.

Completing the restore configuration

To enable the NMDA Orchestrated Application Protection restore of a PostgreSQL backup, ensure that the required environment variables and NMDA Orchestrated Application Protection parameters are set.

The PostgreSQL server documentation provides more information on the required environment variables.

Ensure that all the required NMDA Orchestrated Application Protection parameters are set in the NMDA configuration file for the restore. The parameters `NSR_SAVESET_NAME` and `NSR_RELOCATION_DEST` are mandatory for a restore operation:

- For the `NSR_SAVESET_NAME` parameter setting, specify the backup save set name, as determined by the backup query results in the preceding topic.
- For the `NSR_RELOCATION_DEST` parameter setting, specify the complete pathname of the target directory that will contain the restored data. The target directory must not exist before the restore as the `nsroapprecover` program creates the directory. Ensure that the user that runs `nsroapprecover` has full permissions in the parent directory of the destination directory.
- For the optional `NSR_RECOVER_TIME` parameter setting, you can specify the save time of the backup save set, as determined by the backup query results in the preceding topic.
If you want to restore the most recent backup, you can omit the `NSR_RECOVER_TIME` setting.

For example, the following parameter settings in the `RECOVER` section of the NMDA configuration file include the save time of the backup save set to be restored:

```
<RECOVER>
  <NSR_SAVESET_NAME> PostgreSQL_MY_DEMO_MY_TEST_full </NSR_SAVESET_NAME>
  <NSR_RELOCATION_DEST> /nsr/apps/tmp/my_test </NSR_RELOCATION_DEST>
  <NSR_RECOVER_TIME> 1516039064 </NSR_RECOVER_TIME>
</RECOVER>
```

[NMDA parameters for Orchestrated Application Protection restores](#) on page 456 provides complete details on the restore parameters to set in the NMDA configuration file. For an Orchestrated Application Protection restore, the NMDA configuration file must include the proper XML formats as described in [NMDA configuration file](#) on page 400.

Registering the PostgreSQL `restore_command` for WAL archiving restore

To enable the PostgreSQL archived WAL segment file restore, you must register the `nsroapprecover` program with its required command line options through the `restore_command` setting in the `recovery.conf` file.

Note: To set all the parameters properly in the `recovery.conf` file for the particular restore and recovery, follow the details in the "Recovery Configuration" chapter in the online PostgreSQL document at www.postgresql.org.

In the `restore_command` setting, specify the `nsroapprecover` command and its command line options from the following table to provide the required restore functionality. For example, the `recovery.conf` file can include the following type of `restore_command` setting:

```
restore_command = '/usr/sbin/nsroapprecover -o pg_p_opt="%p" -o pg_f_opt="%f"
-z configuration_file_path'
```

The following table describes each `nsroapprecover` command line option that is supported in the `restore_command` setting. Each command line option starts with a dash (-).

Note: The `-z` option is mandatory for all restore operations.

Ensure that all the required Orchestrated Application Protection restore parameters are also set in the NMDA configuration file. [NMDA parameters for Orchestrated Application Protection restores](#) on page 456 provides details about the restore parameters.

Table 25 Command line options for PostgreSQL restores

Command line options	Description	Default and valid values
-l	Specifies to list the save set contents for the PostgreSQL restore. Optional.	Not applicable.
-o pg_p_opt="%p" -o pg_f_opt="%f"	For a PostgreSQL WAL restore only, specifies the target file name that is used for the PostgreSQL %p option and the source file name that is used for the %f option. The NSR_DATABASE_TYPE, NSR_INSTANCE_NAME, and NSR_BACKUP parameters are used to determine the save set. Optional.	<ul style="list-style-type: none"> • Undefined (default). • Valid target file name for the %p option, and valid source file name for the %f option.
-z configuration_file_pathname	Specifies the NMDA configuration file that contains the parameter settings and command line options for the PostgreSQL restore. Mandatory.	<ul style="list-style-type: none"> • Undefined (default). • Valid complete pathname of the NMDA configuration file.

Retrieving the backed-up PostgreSQL data with the nsroapprecover command

Before you perform a PostgreSQL restore, ensure that the restore requirements from the preceding topics have been met. You can retrieve a PostgreSQL database backup or transaction log backup by running the `nsroapprecover` command from the command line as the operating system user that launched the PostgreSQL server.

To retrieve the PostgreSQL archived WAL log files, the PostgreSQL daemon runs the `nsroapprecover` command that is registered in the `restore_command` setting in the `recovery.conf` file. [Registering the PostgreSQL restore_command for WAL archiving restore](#) on page 270 provides details.

To retrieve the PostgreSQL backup from the NetWorker server, run the `nsroapprecover -z configuration_file_path` command.

The `nsroapprecover` program retrieves the PostgreSQL backup from the NetWorker server into the destination directory that is specified by the `NSR_RELOCATION_DEST` parameter. By default, unless you set `NSR_RECOVER_TIME`, the `nsroapprecover` command restores the most recent backup. The preceding topic provides more information about the restore parameter settings.

The *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsroapprecover` command.

Restoring the retrieved data to the PostgreSQL database


After you run the `nsroapprecover -z configuration_file_path` command to retrieve the backup into the specified directory, perform the following steps:

1. Use the PostgreSQL restore utilities to complete the backup recovery. For example:

- a. Ensure that you can find the backed-up data in the directory that is specified in the restore configuration file.
- b. Shut down the PostgreSQL server.
- c. Copy the files in `recovery_location/basebackup` to the PostgreSQL data folder.
- d. Create the `recovery.conf` file and copy the file to the data folder.
- e. Start the PostgreSQL server, which causes PostgreSQL to read the `recovery.conf` file and then recover and apply the needed log files for the target time or required transaction.

The PostgreSQL documentation provides more details on how to use the restore utilities and how to restore and recover the database.


2. Perform a full backup of the PostgreSQL database.

 **Note:** A full backup of the PostgreSQL database is required after a backup recovery.

Canceling a data restore or recovery


You can cancel a data restore or recovery by using the following procedures.

About this task

 **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](#) 274
- [Preparing for DB2 disaster recovery](#)275
- [Preparing for Informix disaster recovery](#) 276
- [Preparing for Lotus disaster recovery](#) 280
- [Preparing for MySQL disaster recovery](#) 284
- [Preparing for Oracle disaster recovery](#)285
- [Preparing for Orchestrated Application Protection disaster recovery](#) 292
- [Preparing for SAP IQ disaster recovery](#)293
- [Preparing for Sybase disaster recovery](#)293

Preparing for disaster recovery

This chapter describes how to recover from a disaster. Use this chapter with [Data Restore and Recovery](#) on page 189 and the *NetWorker Server Disaster Recovery and Availability Best Practices Guide*.

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 *NetWorker Server Disaster Recovery and Availability Best Practices Guide* provides details.

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.

The operating system documentation and the database vendor documentation describe procedures for the operating system and database or application.

The *NetWorker Installation Guide* describes how to reinstall the NetWorker client.

The *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 that are described in the following topics on the new host. Also, to enable the restore of backups that are performed 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 *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.

Procedure

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.

Also, 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](#) 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 *NetWorker Administration Guide* provides details.

Performing a DB2 disaster recovery

You can use the following steps to perform a DB2 disaster recovery.

Procedure

1. Perform the required procedures to set up the system hardware, install the operating system, and install the DB2 server software. Ensure that the DB2 server version is the same as used in the NMDA DB2 backups.
2. Install the NetWorker and NMDA software according to the appropriate installation guides.
3. Run the NetWorker `recover` command to restore the DB2 configuration files and other important files that are 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_name -c original_NetWorker_client_name -a
temporary_file
```

The DB2 backup and recovery documentation describes how to recover a DB2 database. [Performing DB2 data restore and recovery](#) on page 193 provides details.

4. Recover the whole DB2 database to the required point-in-time according to [Performing DB2 data restore and recovery](#) on page 193. To perform a redirected restore to a different server than the server where the database was backed up, follow the instructions in one of the following topics:
 - [Performing DB2 data restores to a different instance](#) on page 196
 - [Performing DB2 data restores to a different host](#) on page 198

When you restore from an incremental or delta backup, ensure that you first restore the history file as described in [Delta or incremental restore failure](#) on page 493.

Note: You must use the `db2 restore` command and `db2 rollforward` command to perform the 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 32.

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:

About this task

- Informix ONCONFIG file
- ixbar file
- onconfig boot file
- sqlhosts file (UNIX only)
- sm_versions file
- Copy of Windows registry information that is stored under "HKEY_LOCAL_MACHINE \SOFTWARE\Informix" (Windows only)

[NSR_DR_BACKUP_INFO](#) provides details.

Also, 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](#) provides details.

The NetWorker `save` interface backs up the files that are 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 *NetWorker Administration Guide* provides details.

Performing an Informix imported restore


About this task

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 must have a backup of the disaster recovery files that are described in [Preparing for Informix disaster recovery](#) on page 276.

You can use the following steps to perform the Informix imported restore.

Procedure

1. Follow the instructions in [Preparing for disaster recovery](#) on page 274 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.

- i** **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 NetWorker_server_name \
-c original_NetWorker_client_name \
$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.

About this task

You can use the following steps to perform an Informix cold restore.

Procedure

1. Follow the instructions in [Preparing for disaster recovery](#) on page 274 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 that are backed up during Lotus backups](#) on page 36. 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](#) provides details.

Perform Lotus disaster recovery by using the appropriate procedures:

- [Recovering a nonlogged Lotus environment](#) on page 280
- [Recovering a logged Lotus environment](#) on page 281

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.

About this task

Note: For a partitioned Domino server, repeat these steps for each partition.

Procedure

1. Follow the instructions in [Preparing for disaster recovery](#) on page 274 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. [NMDA Parameters and Configuration File](#) on page 399 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:

About this task

- 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.

Procedure

1. Follow the instructions in [Preparing for disaster recovery](#) on page 274 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. [NMDA Parameters and Configuration File](#) on page 399 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:

About this task

- 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 that is 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 *NetWorker Administration Guide* describes the `save` command.

Performing a MySQL disaster recovery

You can use the following steps to perform a MySQL disaster recovery.

Procedure

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 that are 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_name -c original_NetWorker_client_name -d /
etc -a /etc/my.cnf
```

The *NetWorker Administration Guide* describes the `recover` command.


4. Restore any binary log backups according to [Performing MySQL restores of binary log backups](#) on page 228.
5. Recover the whole MySQL instance to the current time according to [Performing MySQL recovery of whole instance backups](#) on page 225. 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 229.

Preparing for Oracle disaster recovery

To prepare an Oracle server for disaster recovery, back up this minimum list of files:

About this task

- 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 Computer\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.

Procedure

1. Record the Oracle DBID in a text file and back up the file. [Creating the DBID text file](#) on page 286 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 286 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 287 provides details.
4. Use an RMAN backup to back up the Recovery Catalog. [Setting up RMAN backups of Recovery Catalog](#) on page 287 provides details.

Perform an Oracle disaster recovery according to [Performing an Oracle disaster recovery](#) on page 287.

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.

About this task

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 286.

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:

About this task

- 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 286
- 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 286 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 *NetWorker Administration Guide* provides details.

Sample Oracle postcommand script

The NMDA installation provides a sample postcommand script, `nsroradrpostcmd(.bat)`, 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`.

During a scheduled NMDA backup, the `nsrdasv` process passes the options `-s server_name -g group_name -y retention -b pool_name -a additional_options` to the postcommand script if the

script name begins with `nsroradr`. This action ensures that the additional files are backed up to the same devices and have the same retention policy as the Oracle datafiles that are backed up through RMAN.

Setting up RMAN backups of the database and related files

Set up an RMAN backup with NMDA to back up the following files:

About this task

- 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 286 describes the postcommand script.

Setting up RMAN backups of Recovery Catalog

About this task

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 287.

The Oracle documentation describes Recovery Catalog backups.

Performing an Oracle disaster recovery

You can use the following steps to perform an Oracle disaster recovery.

Procedure

1. Follow the instructions in [Preparing for disaster recovery](#) on page 274 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 that are 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_name -c original_NetWorker_client_name -
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 a redirect restore for disaster recovery to another host by using the RMAN command line interface (CLI), follow the procedure in [Performing a redirected restore for disaster recovery to another host using the RMAN CLI](#) on page 288.
4. 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.

Performing a redirected restore for disaster recovery to another host using the RMAN CLI

You can use the following procedure to perform a redirected restore for disaster recovery to a destination host, host 2, which is different from the source host, host 1, where the backup was originally performed. The NetWorker server is located on a separate third host, host 3.

Procedure

1. On source host 1, install Oracle in the path `/u01/app/oracle/product/11.2.0/dbhome_1` with the required SID, for example, ORADB. Install the NetWorker client and NMDA software on host 1.
2. On destination host 2, install Oracle in the same path `/u01/app/oracle/product/11.2.0/dbhome_1` with the same SID, for example, ORADB. Install the NetWorker client and NMDA software on host 2.
3. On host 3, install the NetWorker server and the NMC software.
4. Log in to NMC, create a device, and assign the device to the Data Domain Default pool or another NetWorker pool.
5. Use the client backup configuration wizard to create an NMDA Oracle client resource for source host 1.
6. Use the client backup configuration wizard to create an NMDA Oracle client resource for destination host 2.
7. Click **Client Properties** for source host 1.
8. Select the **Globals (2 of 2)** tab in the NetWorker **Client Properties** window, and add the name of the destination host 2 in the **Remote access** field.
9. Log in to source host 1 and run the following commands:

```
[root@nmdaora2 ~]# su - oracle
[oracle@nmdaora2 ~]$ . oraenv
ORACLE_SID = [oracle] ? ORADB
The Oracle base has been set to /home/oracle/app/oracle
[oracle@nmdaora2 ~]$ sqlplus / as sysdba

SQL*Plus: Release 12.1.0.1.0 Production on Thu Dec 6 23:05:19 2018
Copyright (c) 1982, 2013, Oracle. All rights reserved.

Connected to:
Oracle Database 12c Enterprise Edition Release 12.1.0.1.0 - 64bit
Production
With the Partitioning, OLAP, Advanced Analytics and Real Application
Testing options
```



```

SQL> shutdown immediate;
Database closed.
Database dismounted.
ORACLE instance shut down.
SQL> startup mount;
ORACLE instance started.

Total System Global Area 1703624704 bytes
Fixed Size                 2289160 bytes
Variable Size             1090519544 bytes
Database Buffers         603979776 bytes
Redo Buffers              6836224 bytes
Database mounted.
SQL> alter database archiveolog;
Database altered.
SQL> alter database open;
Database altered.
SQL> exit

```

10. Connect to the RMAN CLI and record the DBID of source host 1. In the following example, the DBID is 2747203147:

```

[oracle@nmdaora2 ~]$ rman target /

Recovery Manager: Release 12.1.0.1.0 - Production on Tue Dec 4 19:12:22
2018
Copyright (c) 1982, 2013, Oracle and/or its affiliates. All rights
reserved.

connected to target database: ORADB (DBID=2747203147)

```

11. In the RMAN CLI, run the RMAN backup script to back up the database on source host 1, including the control file, spfile, and archive logs.

In the RMAN backup script, specify the parameter settings for NSR_SERVER (NetWorker server IP), NSR_DATA_VOLUME_POOL (backup pool name), NSR_DEVICE_INTERFACE (device type), NSR_DEBUG_LEVEL (debug level), and NSR_DPRINTF (log print option):

```

RMAN> RUN {
ALLOCATE CHANNEL CH1 TYPE 'SBT_TAPE';
ALLOCATE CHANNEL CH2 TYPE 'SBT_TAPE';
send 'NSR_ENV=(NSR_SERVER=10.31.222.78,NSR_DATA_VOLUME_POOL="Data
Domain
Default",NSR_DEVICE_INTERFACE=DATA_DOMAIN,NSR_DEBUG_LEVEL=9,NSR_DPRINTF=
TRUE)';
BACKUP
FULL
FORMAT '%d_%U'
DATABASE
INCLUDE CURRENT CONTROLFILE SPFILE
PLUS ARCHIVELOG;
RELEASE CHANNEL CH1;
RELEASE CHANNEL CH2;
}

```

12. Log in to the destination host 2 and delete all the previously restored files by using a `clean` script or the `rm -rf` command. Ensure that you delete the `.dbf`, `.log`, and `ctl` files and

the spfile /u01/app/oracle/product/11.2.0/dbhome_1/dbs/spfileORADB.ora.
For example, run the following commands to delete the spfile:

```
oracle@linora4$ cd /u01/app/oracle/product/11.2.0/dbhome_1/dbs
rm -rf spfileORADB.ora
```

13. Log in to source host 1, list the backup through RMAN, and record the date and time of today's backup:

```
[oracle@nmdaora2 ~]$ export NLS_DATE_FORMAT='DD-MM-YYYY HH24:MI:SS'
RMAN> LIST BACKUP OF DATABASE COMPLETED AFTER '(SYSDATE-2)';
```

BS Key	Type	LV	Size	Device Type	Elapsed Time	Completion Time
72	Full		560.50M	SBT_TAPE	00:00:08	06-12-2018 18:54:31
BP Key: 72 Status: AVAILABLE Compressed: NO Tag: TAG20181206T185423						
Handle: ORADB_2bt456v_1_1 Media: blrv041c078.lss.emc.com.dddefault.001						
List of Datafiles in backup set 72						
File	LV	Type	Ckp SCN	Ckp Time		Name
3		Full	1896616	06-12-2018 18:54:23		/home/oracle/app/oracle/oradata/ORADB/datafile/o1_mf_sysaux_g0g2bvkt_.dbf
6		Full	1896616	06-12-2018 18:54:23		/home/oracle/app/oracle/oradata/ORADB/datafile/o1_mf_users_g0g2bvoq_.dbf

BS Key	Type	LV	Size	Device Type	Elapsed Time	Completion Time
73	Full		700.25M	SBT_TAPE	00:00:08	06-12-2018 18:54:31
BP Key: 73 Status: AVAILABLE Compressed: NO Tag: TAG20181206T185423						
Handle: ORADB_2atk456v_1_1 Media: blrv041c078.lss.emc.com.dddefault.001						
List of Datafiles in backup set 73						
File	LV	Type	Ckp SCN	Ckp Time		Name
1		Full	1896615	06-12-2018 18:54:23		/home/oracle/app/oracle/oradata/ORADB/datafile/o1_mf_system_g0g2bvo9_.dbf
4		Full	1896615	06-12-2018 18:54:23		/home/oracle/app/oracle/oradata/ORADB/datafile/o1_mf_undotbs1_g0g2bvlq_.dbf

14. Log in to destination host 2, connect to the database by using sql, and shut down the database:

```
[oracle@linora4 ~]$ sqlplus / as sysdba

SQL*Plus: Release 12.1.0.1.0 Production on Tue Dec 4 21:55:51 2018
Copyright (c) 1982, 2013, Oracle. All rights reserved.

Connected to:
Oracle Database 12c Enterprise Edition Release 12.1.0.1.0 - 64bit
Production
With the Partitioning, OLAP, Advanced Analytics and Real Application
Testing options

SQL> shutdown abort;
ORACLE instance shut down.
SQL> exit
Disconnected from Oracle Database 12c Enterprise Edition Release
12.1.0.1.0 - 64bit Production
With the Partitioning, OLAP, Advanced Analytics and Real Application
Testing options
```

15. On host 2, connect to the RMAN CLI and run the RMAN restore script to recover the database. Ensure that you set the DBID to the same value as recorded in step 10 on source host 1, and use the backup date and time as recorded in step 13:

```
[oracle@linora4 ~]$ rman target /

Recovery Manager: Release 12.1.0.1.0 - Production on Tue Dec 4 21:58:38
2018
Copyright (c) 1982, 2013, Oracle and/or its affiliates. All rights
reserved.

connected to target database (not started)

RMAN> SET DBID=2747203147
RUN {
startup force nomount;
ALLOCATE CHANNEL CH1 TYPE 'SBT_TAPE' PARMS
'SBT_PARMS=(NSR_SERVER=10.31.222.78,NSR_CLIENT=nmdaora2,NSR_RECOVER_POOL
=Data Domain Default)';
RESTORE SPFILE TO '/home/oracle/app/oracle/product/12.1.0/dbhome_1/dbs/
spfileORADB.ora' from autobackup MAXSEQ=01 until time =
"TO_DATE('20181206','YYYYMMDD)";
shutdown;
startup nomount;
restore controlfile from autobackup MAXSEQ=01 until time =
"TO_DATE('20181206','YYYYMMDD)";
alter database mount;

RESTORE DATABASE UNTIL TIME = "TO_DATE('06/12/2018-18:54:23', 'DD/MM/
YYYY-HH24:MI:SS)";

RECOVER DATABASE UNTIL TIME = "TO_DATE('06/12/2018-18:54:23', 'DD/MM/
YYYY-HH24:MI:SS)";
RELEASE CHANNEL CH1;
}
```

The following message appears, as expected:

```
RMAN-03002: failure of recover command at 12/06/2018 19:05:44
RMAN-06054: media recovery requesting unknown archived log for thread 1
with sequence 1 and starting SCN of 1823550
```

The recovery fails because the logs SCN is ahead, as expected. Run the following command to reset and update the logs:

```
RMAN> alter database open resetlogs;
Statement processed
```

16. On the destination host 2, open the database and validate the data.

You can reuse the script from step 15 for the database recovery by altering the specified date and time to the current date and time.

Preparing for Orchestrated Application Protection disaster recovery

To prepare a database or application environment for an Orchestrated Application Protection disaster recovery, perform regular scheduled backups of the database or application that is protected by Orchestrated Application Protection.

About this task

Refer to the database or application administration guide as provided by the vendor for details about all the files that must be backed up for a disaster recovery scenario. Add the required command in the backup shell script that is specified by `NSR_BACKUP_SCRIPT`, to copy or generate the files in the `$OAPP_MOUNT_DIR/` directory. These files will be included in the backup save set.

Procedure

1. In the backup shell script, add the command to copy the required files to the `$OAPP_MOUNT_DIR/` directory. For example:

```
cp /home/postgres/postgresql.conf $OAPP_MOUNT_DIR/
```

2. Specify the pathname of the edited backup shell script in the `NSR_BACKUP_SCRIPT` parameter setting in the NMDA configuration file.

The required files are backed up during every Orchestrated Application Protection backup.

Performing an Orchestrated Application Protection disaster recovery

You can use the following steps to perform an Orchestrated Application Protection disaster recovery.

Procedure

1. Perform the required procedures to set up the system hardware, install the operating system, and install the database or application software that was protected by Orchestrated Application Protection.

Ensure that the database or application software version is the same as used in the NMDA Orchestrated Application Protection backups.

2. Install the NetWorker and NMDA software according to the appropriate installation guides.
3. Run the Orchestrated Application Protection command `nsroapprecovery` to restore the backed-up data and the required files to a local directory. For example:

```
nsroapprecovery -z configuration_file_path
```

[Performing Orchestrated Application Protection data restore](#) on page 264 provides details on the required restore configurations and procedures.

4. Complete the database or application restore and recovery by following the instructions in the database or application administration guide.

Preparing for SAP IQ disaster recovery

To prepare an SAP IQ server environment for disaster recovery, perform regular scheduled backups of the SAP IQ server and databases.

About this task

Keep up-to-date printouts of the device allocations for the SAP IQ database.

Consider backing up the installation files by performing regular file system backups with the NetWorker software. The SAP IQ documentation provides details on the SAP IQ components to back up.

You can use the following steps to recover the SAP IQ server after a disaster.

Procedure

1. Use the printout of database device allocations to re-create the databases. The SAP IQ documentation describes the information to track for disaster recovery.
2. Recover the SAP IQ system databases and user databases. The SAP IQ documentation provides details.
3. Use the `nsriqrc` command to recover the SAP IQ data. [Performing SAP IQ data restore](#) on page 244 provides details.

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.

About this task

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 sybssystemdb on the master device.
- Back up the master database after performing actions such as initializing database devices, creating or changing databases, or adding a server login.

The Sybase documentation provides details.

You can use the following steps to recover the Sybase server after a disaster.

Procedure

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 following documents on the SAP Support portal provide details on how to restore the Sybase master database and system databases:
 - SAP Note 1611715 - SYB: How to restore an SAP ASE database server (Windows)
 - SAP Note 1618817 - SYB: How to restore an SAP ASE database server (UNIX)
3. Use the `nsrsybrc` command to recover the Sybase data. [Performing Sybase data restore and recovery](#) on page 251 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.

About this task

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.

i **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.

Procedure

1. Rebuild the lost master device by using the `datasever` 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 `syssservers` table.
4. Set the following Sybase restore parameters in the NMDA configuration file:
 - `NSR_BACKUP_PATHS`—Specifies the name of the Sybase server and master database, as the value `SYBASE:/ASE_server_name/master`.
 - i** **Note:** The `ASE_server_name` value is case-sensitive and must be in the same case as recorded in the backup entries in the NetWorker indexes.
 - `NSR_SERVER`—Specifies the hostname of the NetWorker server.
 - `SYBASE_USER`—Specifies the username of the Sybase user account.
 - `USER_PSWD`—Specifies the encrypted password of the Sybase user account.
 - i** **Note:** You must run the `nsrdaadmin -P -z configuration_file_path` command to add the encrypted password to the configuration file.

[NMDA Sybase parameters](#) on page 464 provides details about the Sybase restore parameters in the NMDA configuration file.

5. Run the following `nsrsybrc` command as the Sybase user to recover the master database from the backup:

```
nsrsybrc -z configuration_file_path
```

After the master database is loaded, the Sybase server performs postprocessing checks and validations and then shuts down.

6. Restart the Sybase server.

7. If required, recover the model database and other databases that were on the master device.
8. 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.

About this task

You can use the following steps to recover a database after the database device fails.

Procedure

1. Run the following command to perform an incremental backup of each database on the failed device that backs 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
```

[NMDA Parameters and Configuration File](#) on page 399 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 by running the `nsrsybrc` command:

```
nsrsybrc -z configuration_file_path
```

where *configuration_file_path* is the pathname of the NMDA configuration file that contains the following Sybase restore parameters:

```
NSR_BACKUP_PATHS=SYBASE:/ASE_server_name/database_name  
NSR_SERVER=NetWorker_server_name  
SYBASE_USER=Sybase_username  
USER_PSWD=encrypted_password
```

Note: In the `NSR_BACKUP_PATHS` setting, *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 *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsrsybrc` command. [NMDA Sybase parameters](#) on page 464 describes the Sybase restore parameters.

CHAPTER 6

Cluster and High-Availability (HA) Systems

This chapter includes the following topics:

• Active-passive clusters	298
• Active-active application clusters and HA systems	304
• DB2 DPF systems	306
• DB2 pureScale systems	310
• DB2 HADR systems	315
• Informix MACH systems	318
• Oracle RAC systems	320
• Sybase ASE Cluster Edition systems	325

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 (HA) solutions or to improve computer performance.

The following figures display examples of supported active-passive cluster configurations:

- Single database node with failover capability
- Multiple database nodes with failover capability
- Multiple database nodes with mutual failover capability

Each database node in the figures corresponds to one of the following items:

- DB2 node
- Informix server
- Lotus Domino server or partition
- MySQL server
- Oracle instance
- Sybase server

Note: Some cluster documents or applications documents might describe the two active-passive clusters that are configured as shown in [Figure 11](#) on page 299 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 *NetWorker E-LAB Navigator* at <https://elabnavigator.emc.com/eln/elhome> provides details about the NetWorker client supported cluster environments. NMDA supports the cluster environments that the NetWorker client supports.

The latest *NetWorker Cluster Integration Guide* provides details on how to enable the NetWorker cluster support for the NMDA client. Ensure that you specify the virtual IP address for the client name in the NetWorker configuration.

Figure 9 Single database node with failover capability


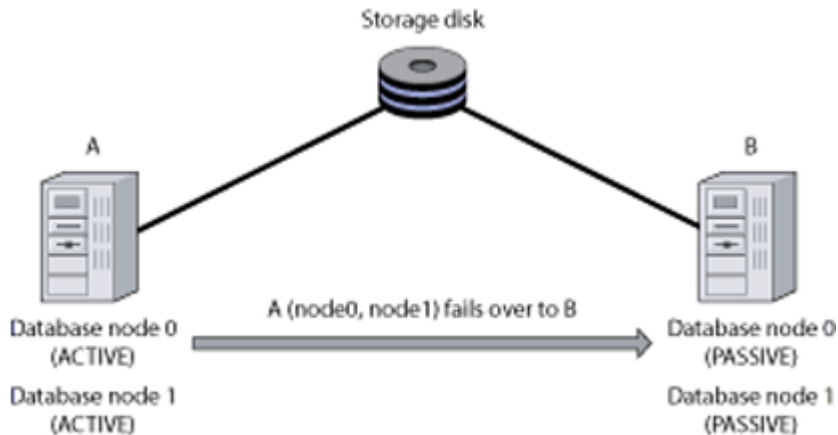
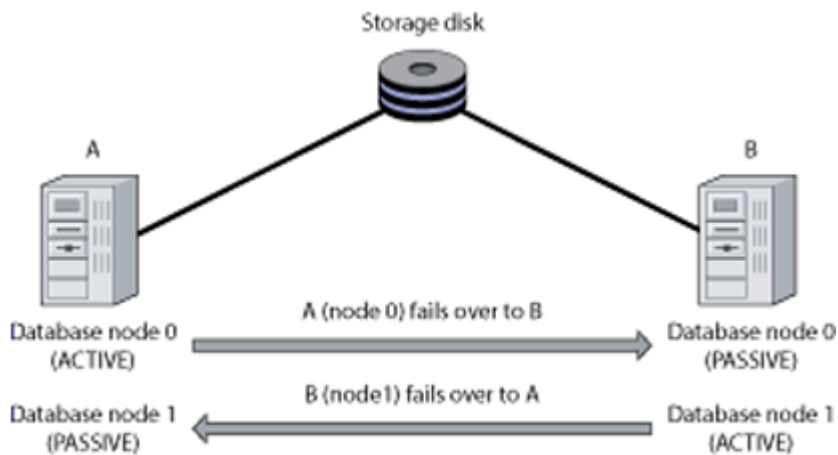


Figure 10 Multiple database nodes with failover capability**Figure 11** Multiple database nodes with mutual failover capability

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 Retries setting in the backup action of the data protection policy 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 have installed NMDA and the NetWorker software on each node of the cluster that will participate in the backups and restores. The *NetWorker Module for Databases and Applications Installation Guide* provides details.

On UNIX, ensure that the `/nsr/apps` and `/nsr/apps/tmp` directories have the `drwxrwxrwt` access permissions on all the nodes. These permissions are set by default.

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 file index has the record of backup information.
- During a restore, the name of the NetWorker client whose file index 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 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](#) describes the parameter.


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.

About this task

Ensure that you meet the cluster configuration requirements as described in [Cluster configuration requirements](#) on page 299.

Follow the procedure for [Configuring scheduled backups with the wizard](#) on page 86 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 *NetWorker Cluster Integration Guide* provides details.

Procedure

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 300 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 300 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 node 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 physical 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 backups in a cluster without the wizard

You can configure a scheduled backup in a cluster either with or without the wizard. You can configure a manual or scheduled backup without the wizard by using the following steps.

About this task

Ensure that you have installed NMDA on all the nodes as described in [Cluster configuration requirements](#) on page 299.

Complete the following steps to configure NetWorker and NMDA for the cluster. [Configuring scheduled backups without the wizard](#) on page 89 and [Configuring manual backups](#) on page 103 provide details about the resource configuration procedures and configuration file setup.

Procedure

1. Select one of the nodes from the active-passive cluster, and use the virtual hostname of that node as the client name that will store the backup data.
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.

Due to the requirement to maintain a copy of the configuration file on each physical node, it is recommended that you create the configuration file on a shared disk.

3. Set the `NSR_CLIENT` parameter to the virtual cluster hostname in the NMDA configuration file. [NSR_CLIENT setting for cluster systems](#) on page 300 provides details.

Note: For Oracle backups, you must also create an RMAN script in the shared location, if you do not store the RMAN script in the RMAN catalog. Set the `NSR_CLIENT`

parameter to the virtual cluster hostname in the script. If you perform manual backups, ensure that NSR_SERVER is also set in the script.

4. 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. [NSR_PHYSICAL_HOST_LICENSE setting for cluster systems](#) on page 300 provides details.

Note: For Oracle backups, set the NSR_PHYSICAL_HOST_LICENSE parameter as required in the RMAN script.

5. Create a generic Client resource for each physical cluster node that will run backups and restores, if a Client resource does not yet exist for a particular host.

The following example includes the attribute settings in a generic Client resource:

Note: To view some of the Client resource attributes in the **Client Properties** dialog box in the NMC Administration interface, you must enable the diagnostic mode by selecting **View > Diagnostic Mode**.

- 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)
- Protection group list: (blank)
- Save set: SKIP
- Schedule: SkipAll

6. Create the required Client resource on the NetWorker server:

- If you perform only manual backups, create a generic Client resource for the virtual hostname that is selected in step 1.
- If you perform scheduled backups, create an NMDA specific Client resource for the virtual hostname that is selected in step 1:

- The Save Set attribute contains the required value as described in [Table 9](#) on page 93. For Oracle backups, the Save Set attribute contains the pathname of the RMAN script that is created in step 3.

- 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 -c cluster_virtual_hostname

where:

- *configuration_file_path* is the pathname of the NMDA configuration file that is created in step 2.
- *cluster_virtual_hostname* is the virtual cluster hostname in the NMDA configuration file.

- If you perform Client Direct backups, the Client Direct attribute is selected as described under "Advanced attributes" in [Table 9](#) on page 93. This attribute is selected by default.
- If you use a DD Boost device, follow the relevant information in [Configuring DD Boost backups to use Client Direct](#) on page 106 or [Configuring DD Boost backups to use a storage node](#) on page 108 for the virtual cluster node.
- If you back up data to a local storage node, the Storage Nodes attribute is set to the following values in the following order:

```
curphyhost
nsrserverhost
```

[Setting up nodes to back up to a local storage node](#) on page 303 provides details.

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 -c cluster_virtual_hostname`

Protection group list: `db2group`

Remote access: `user=dbinst,host=clus_phys1`

`user=dbinst,host=clus_phys2`

`user=SYSTEM,host=clus_phys1`

`user=SYSTEM,host=clus_phys2`

Save set: `DB2:/SAMPLE/NODE0001`

Configuring probe-based backups in a cluster

Use the following procedures to configure a probe-based NMDA backup in a cluster:

About this task

- Configure a scheduled backup in the cluster according to [Configuring scheduled backups in a cluster with the wizard](#) on page 300 or [Configuring backups in a cluster without the wizard](#) on page 301.
- Configure the probe-based backup according to [Configuring probe-based backups](#) on page 111. Associate the Probe resource with the Client resource for the virtual cluster node.

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 node. You can configure a virtual cluster client to direct the backups to the storage node on the cluster node on which the backup runs (local storage node).

About this task

You can use the following steps to set up the cluster nodes to back up to a local storage node.

Procedure

1. Install the NetWorker storage node on each node that is 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 that is 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 300 or [Configuring backups in a cluster without the wizard](#) on page 301), set the Storage Nodes attribute in NMC to the following values in this order:

```
curphyhost
nsrserverhost
```

Example 35 describes how to set up three Oracle RAC nodes as storage nodes for NMDA backups.

Enabling recoveries in a cluster

For recovery in a cluster, ensure that the required configurations are in place:

About this task

- In the Client resource of the virtual cluster node, you have set the Remote Access attribute as described in [Performing NMDA data restore and recovery](#) on page 191.
- 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 300 provides details.

To run a restore and recovery, use the procedures that are described in [Data Restore and Recovery](#) on page 189.

Active-active application clusters and HA systems

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.

NMDA supports the following types of active-active application cluster software:

- 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 299. However, the Lotus partitioned Domino server is not an active-active application cluster. The NMDA documentation refers to it as an active-passive cluster.

DB2 DPF and HADR are high-availability (HA) systems that do not share common storage between the nodes. However, for NMDA configuration and licensing, these systems are treated the same as active-active application clusters.

The *NetWorker E-LAB Navigator* at <https://elabnavigator.emc.com/eln/elhome> provides details about the supported cluster environments.

The following topics describe specific types of active-active application clusters:

- [DB2 DPF systems](#) on page 306

- [DB2 pureScale systems](#) on page 310
- [DB2 HADR systems](#) on page 315
- [Informix MACH systems](#) on page 318
- [Oracle RAC systems](#) on page 320
- [Sybase ASE Cluster Edition systems](#) on page 325

Configuring an active-active application cluster

You cannot use the NMDA wizard to configure a scheduled backup in an active-active application cluster.

About this task

You must use the following steps to configure a scheduled or manual backup in an active-active application cluster.

Procedure

1. Install NMDA and the NetWorker software on each node of the cluster that will participate in the backups and restores. The *NetWorker Module for Databases and Applications Installation Guide* provides details.
2. On UNIX, ensure that the `/nsr/apps` and `/nsr/apps/tmp` directories have the `drwxrwxrwt` access permissions on all the nodes. These permissions are set by default.
3. Select one of the nodes from the active-active cluster, and use the hostname of that node as the client name that will store the backup data. If the cluster system uses virtual hostnames, use the virtual hostname of the system.

Note: You must select a single hostname to store the data for all the backups in the cluster.

4. Create the NMDA configuration file in one of the following locations:
 - On the shared disk of the system, if one exists.
 - In the same location on the local disk of each node.

Due to the overhead of maintaining a copy of the configuration file on each node, it is recommended that you create the configuration file on a shared disk, if possible.

Set the `NSR_CLIENT` parameter in the configuration file to the hostname selected in step 3. You must use the same hostname in the `NSR_CLIENT` parameter for all the backups of the cluster nodes, no matter which node is running a backup.

If you will perform manual backups, ensure that `NSR_SERVER` is also set. Depending on your licensing requirements, you might also need to set `NSR_PHYSICAL_HOST_LICENSE`. [NSR_PHYSICAL_HOST_LICENSE setting for cluster systems](#) on page 300 provides details.

Note: For Oracle RMAN operations, you must also create an RMAN script in the shared location, if you do not store the RMAN script in the RMAN catalog. Set the `NSR_CLIENT` parameter in the script to the hostname selected in step 3. If you will perform manual backups, ensure that `NSR_SERVER` is also set in the script.

5. Create a generic Client resource for each physical cluster node that will run backups and restores, if a Client resource does not yet exist for a particular host.

Note: You can skip this step for DB2 HADR systems.

6. For the node hostname selected in step 3, create the required Client resource on the NetWorker server:

Note: In an active-active cluster system that uses virtual hostnames, you must use the virtual hostname of the selected node. The use of this hostname ensures that when the backup is started or restarted after a node failure, the backup runs on the physical node that takes ownership of the specified virtual hostname.

- If you will perform scheduled backups, create an NMDA specific Client resource for the hostname of the node that is selected in step 3.

Note: For a DB2 HADR system that does not use a Virtual IP (VIP), you must create an NMDA specific Client resource for each node in the system.

- If you will perform only manual backups, create a generic Client resource for the hostname of the node that is selected in step 3.

[Configuring backups in a cluster without the wizard](#) on page 301 provides details on the Client resource configuration for both generic and NMDA specific configurations.

7. For probe-based backups, follow the instructions in [Configuring probe-based backups in a cluster](#) on page 303.

Restarting backups after a system failure

About this task

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 299.

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 that is used to start an NMDA scheduled backup, the backup might fail. If virtual hostnames are used in the cluster, the restart starts the backup on the new physical node that takes ownership of the specified virtual hostname. If virtual hostnames are not used in the cluster, you must use manual intervention to configure and restart the backup on a different node that is available.

Procedure



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:

- Database nodes reside in partitions on a single physical host
- 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

About this task

To configure database backups with NMDA in a DB2 DPF system, follow the instructions in [Configuring an active-active application cluster](#) on page 305.

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 308.
- To configure a scheduled DPF backup, follow the instructions in [Configuring scheduled DPF backups](#) on page 309.
- To restore and recover a DPF backup, follow the instructions in [Restoring and recovering DPF backups](#) on page 310.
- 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 111. 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 that are 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.

About this task

You can use the following steps to configure a manual backup of a DPF database that resides on two hosts, host1 and host2.

Procedure

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 (node where the database was created, for example, `NODE0000`).

Note: Setting `NSR_CLIENT` enables NetWorker to store backups from all the nodes under the same index as described in [Configuring an active-active application cluster](#) on page 305.

3. Use the NMC program to create a basic NetWorker Client resource for both host1 and host2. The *NetWorker Administration Guide* provides details about using the NMC program.

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.

About this task

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.

About this task

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 307 and [Figure 13](#) on page 307, 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`.

Procedure

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, for example, `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](#) provides details.
4. Create only one backup group. Use the NMC program to create a backup group, 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. For example, the attribute settings are as follows:
 - Name: `host1`
 - Save set: `DB2:/sample/NODE0000`
 - Protection group list: `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 exist. For example, the attribute settings are as follows:
 - Name: host2
 - Save set: DB2: SKIP
 - Protection group list: (blank)
 - Backup command: (blank)

The scheduled backup backs up the partitions that are specified in the DB2_PARTITION_LIST parameter.
7. Create a backup policy, a workflow, and a backup action, and associate the group with the workflow.

Restoring and recovering DPF backups

The `db2 recover` command combines the functions of the `db2 restore` command and `db2 rollforward` command.

About this task

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.

You can configure DB2 pureScale to automatically back up logs generated by each member node, to the NetWorker backup device.

NMDA supports all the features of a DB2 pureScale system except the following features:

- Automatic failover during scheduled backups
- Backup configuration with the wizard
- Snapshot-based backups

For example, after a node failure in a DB2 pureScale system, 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.

NMDA supports delta and incremental backups for the DB2 versions that support these types of backups.

i **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

About this task

Before you configure a backup or recovery in a DB2 pureScale system, ensure that you meet the requirements in [Configuring an active-active application cluster](#) on page 305.

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 *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 312.
- To configure a scheduled backup, follow the instructions in [Configuring scheduled backups in a DB2 pureScale system](#) on page 312.
- To configure a transaction log backup, follow the instructions in [Configuring automatic backups of transaction logs in a DB2 pureScale system](#) on page 313.
- For a probe-based backup, configure a scheduled backup as usual and follow the additional instructions in [Configuring probe-based backups](#) on page 111. Associate the Probe resource with the Client resource configured for the scheduled backup.

i **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 313.

Configuring manual backups in a DB2 pureScale system

About this task

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 backups in a cluster without the wizard](#) on page 301.

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.

About this task

To perform a manual backup, run the `db2 backup` command as described in [Performing DB2 manual backups with the db2 backup command](#) on page 168. For example:

Procedure

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

About this task

A scheduled backup in a DB2 pureScale system must use an NMDA configuration file. Follow the instructions in [Configuring backups in a cluster without the wizard](#) on page 301 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.

About this task

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 122.

Restoring and recovering backups in a DB2 pureScale system

About this task

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 193.

For example, you can run the following command to perform a point-in-time restore of a database that is 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 yyyyymmddhhmmss
```

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 199 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, `NODEnnnn`. 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 that are 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 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/sqlllib_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/sqlllib_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 200 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 that are used for restore and recovery.

Restoring between a DB2 pureScale instance and Enterprise Server Edition

Starting with DB2 10.5, you can restore an offline NMDA backup of a DB2 pureScale instance to DB2 Enterprise Server Edition. You can also restore an offline NMDA backup of DB2 Enterprise Server Edition to a DB2 pureScale instance.

About this task

Ensure that you follow the required steps for the particular type of NMDA restore between a DB2 pureScale instance and DB2 Enterprise Server Edition.

Restoring a backup from a DB2 pureScale instance to Enterprise Server Edition

You must follow the required steps to restore an offline NMDA backup from a DB2 pureScale instance to DB2 Enterprise Server Edition, without rollforward support through the transition. The IBM DB2 documentation provides details about the transition.

Procedure

1. In the NMDA configuration file on the target DB2 Enterprise Server, ensure that the `NSR_CLIENT` parameter setting is the identical to the `NSR_CLIENT` setting used during the backup in the DB2 pureScale environment.
2. In the NetWorker Client resource that is used in the DB2 pureScale environment, ensure that the Remote Access attribute specifies the DB2 user on the destination host.
3. On the target DB2 Enterprise Server, restore the offline backup image from the DB2 pureScale instance.
4. Complete any required changes to the restored database configuration and the NMDA DB2 configuration file according to the Enterprise Server environment:
 - Update the settings of the restored database configuration parameters, such as `logarchmeth1`, `logarchopt1`, `logarchmeth2`, `logarchopt2`, and `vendoropt`, if required for future backups and restores of the restored database the Enterprise Server. The database was restored with the original settings of database configuration parameters that are used in the pureScale environment, which might require updates for the new environment.
 - Update the settings of any required parameters in the NMDA configuration file on the Enterprise Server, such as the `NSR_CLIENT` parameter setting.
5. Perform a full offline database backup of the restored database.

Restoring a backup from Enterprise Server Edition to a DB2 pureScale instance

You must follow the required steps to restore an offline NMDA backup from DB2 Enterprise Server Edition to a DB2 pureScale instance.

Before you begin

Before you perform an NMDA backup that you will restore to a pureScale instance, you can run the `db2checkSD` command on the DB2 Enterprise Server to verify that the source database is ready for restore into a pureScale environment. With the verification complete, you can perform a full offline backup of the source database to prepare for the restore operation. The DB2 documentation provides details about these operations.

About this task

Complete the following steps for the restore of an NMDA backup from DB2 Enterprise Server Edition to a pureScale instance.

Procedure

1. In the NMDA configuration file on the DB2 pureScale instance, ensure that the `NSR_CLIENT` parameter setting is the identical to the `NSR_CLIENT` setting used during the backup on the DB2 Enterprise Server.
2. In the NetWorker Client resource that is used on the DB2 Enterprise Server, ensure that the Remote Access attribute specifies the DB2 user on the destination host.
3. On the DB2 pureScale common member (member 0), restore the offline backup image from the DB2 Enterprise Server.
4. On the DB2 pureScale instance, run the following `db2checkSD` command to perform the conversion of the database for use in the pureScale environment:

```
db2checkSD sample -l /tmp/checksd.log -u db2inst1 -p password
```

5. Complete any required changes to the restored database configuration and the NMDA DB2 configuration file according to the pureScale environment:
 - Update the settings of the restored database configuration parameters, such as `logarchmeth1`, `logarchopt1`, `logarchmeth2`, `logarchopt2`, and `vendoropt`, as required.
 - If the NMDA DB2 configuration files in `logarchopt1`, `logarchopt2`, and `vendoropt` are not located in a folder accessible to all the pureScale members, copy the files under the same pathname on each host in the pureScale environment.
 - Update the settings of any required parameters in the NMDA configuration file in the pureScale environment, such as the `NSR_CLIENT` parameter setting.
6. Perform a full offline database backup of the restored database on the pureScale member 0.

DB2 HADR systems

DB2 High Availability Disaster Recovery (HADR) is a data replication feature that provides a high-availability solution for both partial and complete site failures. HADR automatically replicates all the data changes from a primary node database to an identical database on one or more standby nodes.

In a DB2 HADR environment, each primary and standby database has its own separate storage. The primary and standby databases are connected over a standard TCP/IP network. HADR provides fast failover to a standby database if the primary database fails.

IBM recommends keeping the archived logs from both the primary and standby servers in one place within the same shared location. Setting the `logindexbuild` configuration parameter to ON ensures that a new primary server's indexes are ready for immediate use in a failover.

The process of takeover by a standby node typically takes less than a minute. During this time, all the HADR nodes are in a standby mode and backups cannot be performed on any of the nodes.

Note:

NMDA supports only HADR setups where the instance of the HADR database has the same name on all the nodes.

Database administrators must ensure that there is only one active primary node in a DB2 HADR system. A takeover by force command (used for emergencies only) turns a standby node into the primary node and does not affect the original primary node. In this configuration, which is known as a split brain, two nodes acting as independent primaries are both backed up. A database restore from such a backup might fail.

Configuring backups and restores in a DB2 HADR system

You must ensure that a DB2 HADR system is set up correctly according to the IBM DB2 documentation, including the primary server and one or more standby servers.

About this task

To configure database backups with NMDA in a DB2 HADR environment, follow the instructions in [Configuring an active-active application cluster](#) on page 305.

The following additional requirements apply to DB2 HADR systems:

- For HADR systems that use a Virtual IP (VIP) address, use the following instructions for both the manual and scheduled backup configurations:
 - **Note:** A scheduled backup also requires the setup of a NetWorker schedule.
 - Ensure that the NetWorker Client resource that is created for the configuration uses the HADR VIP address or a hostname that is associated with this VIP.
 - Ensure that the Aliases attribute in the Client resource does not contain any hostname that is not associated with the specified VIP.
 - In the Remote Access attribute in the Client resource, list the usernames and hostnames of all the HADR nodes, as described in [step 6](#).
 - In the configuration file on each node, set the `NSR_CLIENT` parameter to the same hostname associated with the VIP address.
- For HADR systems that do not use a VIP address, use the following instructions for both the manual and scheduled backup configurations:
 - Create the NMDA configuration file in the same location on each node in the HADR system.
 - In the configuration file, set the `DB2_SKIP_HADR_STANDBY=TRUE` parameter, and set the `NSR_CLIENT` parameter as described in [NSR_CLIENT setting for cluster systems](#) on page 300.
 - Create the NetWorker Client resource for each node in the HADR system.
 - For the node specified in the `NSR_CLIENT` parameter, in the Remote Access attribute in the Client resource, list the usernames and hostnames of all the other HADR nodes, as described in [step 6](#).
 - Include all the Client resources in the same backup group.

The scheduled backup starts on all the nodes but the database is backed up only on the primary node, no matter which host in the system is the primary node.

Note: If NMDA cannot verify the role as when a node is shut down for maintenance, NMDA treats the node as a stand-alone database and the node backup might fail. The group result is then failure although a good backup was actually created. To avoid this type of situation, open the Client property of this node in NMC and temporarily clear the Scheduled Backup checkbox on the General tab.

To perform a manual backup from the command line in a DB2 HADR system, run the `db2 backup` command on the primary node as described in [Performing DB2 manual backups with the db2 backup command](#) on page 168.

Configuring backups and restores with DB2 HADR in a DB2 pureScale system

Before you configure a backup or restore with DB2 HADR in a DB2 pureScale system, ensure that HADR is configured properly according to the IBM DB2 documentation in both the primary and standby DB2 pureScale environments. You should be familiar with the HADR role switch and failover procedures and the backup and recovery operations in a DB2 pureScale environment with NMDA.

About this task

The following requirements apply to the configuration of backups and restores with HADR in a DB2 pureScale environment:

- Follow the instructions in [Configuring backups and restores in a DB2 pureScale system](#) on page 311 and [Configuring backups and restores in a DB2 HADR system](#) on page 316 to install and configure NMDA on the hosts that have the running pureScale member, in both the primary and standby pureScale environments.
- Create the Client resource for each host with NMDA installed in the primary and standby pureScale environments.
- Select only one Client resource to use for the common NSR_CLIENT setting, under which all the backup save sets and index entries will be stored from both the primary and standby pureScale environments.

In the selected Client resource, set the Remote Access attribute with the usernames and hostnames from all the hosts with a DB2 pureScale member in both the primary and standby pureScale environments.

- Set NSR_CLIENT to the same client hostname in each NMDA DB2 configuration file in both the primary and standby pureScale environments. To simplify the configuration file management, use an identical configuration file in both pureScale environments.
- For a scheduled backup, select two of the Client resources, one from the primary and the other from the standby pureScale environment, and place these Client resources in the same backup group. Ensure that `DB2_SKIP_HADR_STANDBY=TRUE` is set in the configuration file, so that the backup operation is performed only in the primary pureScale environment.

After the configuration is complete, perform the manual or scheduled backup and recovery operations according to the detailed descriptions in [DB2 HADR systems](#) on page 315 and [DB2 pureScale systems](#) on page 310.

Performing DB2 HADR data recovery

You can perform a DB2 HADR system restore on a single node or multiple nodes. The DB2 documentation provides details.

About this task

Complete the following steps to restore the entire HADR system on all the nodes.

Procedure

1. Restore the database to each HADR node as if it is a stand-alone database, as described in [Performing DB2 data restore and recovery](#) on page 193.
2. Configure the HADR settings. You might need to restore the HADR settings according to the steps in [Preparing for DB2 disaster recovery](#) on page 275.
3. Start HADR on all the standby nodes. Rollforward must be run without the `complete` option. Leave the databases in the rollforward pending state as required for a standby database.
4. Start HADR on the primary node.
5. Perform a backup of the DB2 database at the primary node.

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.

i **NOTICE** NMDA operations do not support failover from primary to secondary MACH servers. [Restarting backups after a system failure](#) on page 306 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 perform 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.
 - i** **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.

About this task

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.

Procedure

1. Configure IDS to work with NMDA according to the instructions in the *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 103.
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, which is 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

i 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:

About this task

- 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 277 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 Clusters (RAC) system is an active-active application cluster environment for scalability and high availability. A node in an Oracle RAC system is a physical and virtual host with hostnames such as `host1.emc.com` and `host1-vip.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 that are 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

About this task

To configure database backups with NMDA in an Oracle RAC system, follow the instructions in [Configuring an active-active application cluster](#) on page 305.

The following additional considerations apply to an Oracle RAC system:

- 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 321 provides details.
- Follow the information in [Archived redo logs](#) on page 323 for backups and recovery of archived redo logs on the RAC system.
- For a probe-based backup, configure a RAC backup by using the preceding information and then complete the probe-based backup configuration according to [Configuring probe-based backups](#) on page 111. Associate the Probe resource with the Client resource configured for the selected virtual hostname.
- When you configure the Client resource for the virtual hostname, do not use the Oracle SCAN hostname of the cluster. Select and use one of the node's virtual hostnames in the RAC system.
- RMAN does not support Transparent Application Failover. If a node failure occurs during an Oracle backup, the backup fails but restarts based on the Retries setting in the backup action of the data protection policy for the scheduled backup. Because the Oracle RAC system uses a virtual hostname, the backup restarts on the new node to which the virtual hostname moves.

Creating RMAN backup and restore scripts

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. NMDA enables restores of the Oracle data to any physical node in the RAC, regardless of which RAC node originally performed the backup.

About this task

You can use a single RMAN backup script to run a parallel Oracle backup with NMDA on a RAC system. In the RMAN backup script that is created for running a parallel Oracle backup, allocate multiple channels for the backup:

- You can allocate or configure channels to use Oracle automatic load balancing. For example:

```
configure device type sbt parallelism 4;
```

- You can also allocate or configure channels to use a specific node. For example:

```
configure channel 1 device type sbt connect
'sys/change_on_install@node1';
configure channel 2 device type sbt connect
'sys/change_on_install@node2';
```

In the RMAN script that is used for a RAC backup or restore, ensure that the NSR_CLIENT parameter is set to the virtual hostname selected in [Configuring an active-active application cluster](#) on page 305.

To run an Oracle restore on a RAC system, none of the nodes can be open. You must mount only the node that is running the RMAN restore script.

The following example describes a sample Oracle RAC configuration, including the required RMAN backup and restore scripts.

Example 36 Configuring backups and restores for Oracle RAC

In this example, the RAC system contains three nodes with the following hostnames:

- Node A—Physical: nodeA.emc.com, virtual: nodeA-vip.lss.emc.com

Example 36 Configuring backups and restores for Oracle RAC (continued)

- Node B—Physical: nodeB.emc.com, virtual: nodeB-vip.lss.emc.com
- Node C—Physical: nodeC.emc.com, virtual: nodeC-vip.lss.emc.com

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 `nodeA-vip.lss.emc.com` 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 virtual 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.emc.com
user=database_username, host=nodeC.emc.com
```

- 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 Protection Group List attribute is set to the backup group name.
- The Save Set attribute is set to the RMAN script pathname.

[Table 9](#) on page 93 describes the Client resource attributes.

Example 36 Configuring backups and restores for Oracle RAC (continued)

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-vip.lss.emc.com)';
  backup database;
  release channel t1;
  release channel t2;
  release channel t3;
}
```

To enable restores, you must set NSR_CLIENT to the virtual hostname 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-vip.lss.emc.com)';
  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 324
- [Restoring all archived logs from each node](#) on page 324

 **Note:** The archived logs must use the same NSR_CLIENT setting that is 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-vip.emc.com, by using the following type of RMAN script:

About this task

```
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-vip.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-vip.emc.com, by using the following type of RMAN script:

About this task

```
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-vip.emc.com)';
  restore (archive log all);

  release t1;
  release t2;
}
```

Oracle RAC One Node systems

NMDA supports the Oracle RAC One Node systems. You must meet specific requirements when you relocate a node in an Oracle RAC One Node system.

Ensure that you meet the following requirements when you relocate a database instance to a target node with a different hostname from the source node. For example, the target node has a hostname that is different from the source node hostname when the target node IP address is not included in the SCAN host list:

- The NetWorker Client resource must be reconfigured to use the hostname of the target node.
- The NSR_CLIENT parameter setting must remain the same after the node relocation.

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 that are connected to that instance can fail over to any of the remaining active instances.

The database devices in the cluster (except for private devices that are 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 not assign a specific backup server to an instance. The cluster assigns the least busy backup server from a group according to availability.

i **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.

About this task

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 306 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 300. 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 [Configuring an active-active application cluster](#) on page 305.

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 111. 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` command and set the NSR_CLIENT parameter to specify the NetWorker client hostname of the node that is used to store the NetWorker indexes for the backup. The parameter setting must be the same as the NSR_CLIENT parameter setting used during a backup.

About this task

Use the NSR_RELOCATION_DEST parameter setting to specify a redirected restore. For example, run the following command to perform the redirected restore:

```
nsrsybrc -z configuration_file_path
```

where *configuration_file_path* is the pathname of the NMDA configuration file that contains the following Sybase restore parameters:

```
NSR_BACKUP_PATHS=SYBASE:/index_node_Sybase_instance_name/database_name
NSR_CLIENT=index_node_name
NSR_RELOCATION_DEST=SYBASE:/backup_server_node_Sybase_instance_name/database_name
NSR_SERVER=NetWorker_server_name
SYBASE_USER=Sybase_username
USER_PSWD=encrypted_password
```

Note: In the restore parameter settings, 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.

The *NetWorker Module for Databases and Applications Command Reference Guide* describes the `nsrsybrc` command. [NMDA Sybase parameters](#) on page 464 describes the Sybase restore parameters.

Performing backup and recovery of the quorum device

NMDA does not provide backups of the quorum device.

About this task

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.

About this task

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](#).....330
- [Multiple databases or applications with the same bitness](#).....330
- [Multiple databases or applications with 32-bit and 64-bit coexistence](#).....330

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

 **Note:** NMDA does not support 32-bit and 64-bit coexistence on Linux systems.

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.


About this task

Procedure

1. Install the NMDA software according to the instructions in the *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.
[NMDA Parameters and Configuration File](#) on page 399 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 *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:** NMDA does not support 32-bit and 64-bit coexistence on Linux systems.

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 64-bit NMDA to support the coexistence of 32-bit and 64-bit 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 26 Commands or library names for 32-bit and 64-bit coexistence

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 111
nsrdasv(.exe)	nsrdasv32(.exe)	In Backup Command attribute of the Client resource, required for a scheduled backup	Configuring the Client resource on page 92
For DB2 on 64-bit Windows only:			
libnsrdb2.dll	libnsrdb232.dll	In db2 backup command	<ul style="list-style-type: none"> • Performing DB2 manual backups with the db2 backup command on page 168 • Performing manual DPF backups on page 308
		In db2 restore command	<ul style="list-style-type: none"> • Performing DB2 data restores to the same instance on page 195 • Performing DB2 data restores to a different instance on page 196
		In db2 update db cfg command	Configuring automatic backups of DB2 transaction logs on page 122
		In DB2_VENDOR_LIB_PATH parameter setting for scheduled backups	DB2_VENDOR_LIB_PATH
For Lotus only:			
nsrdasv(.exe)	nsrdasv32(.exe)	In Lotus manual backup command	Performing Lotus manual backups with the nsrdasv command on page 170
nsrdocrc(.exe)	nsrdocrc32(.exe)	In Lotus document-level recovery command	Performing Lotus document-level recovery with the nsrdocrc command on page 218
nsrnotesrc(.exe)	nsrnotesrc32(.exe)	In Lotus database recovery command	Performing Lotus database recovery with the nsrnotesrc command on page 208
		In nsrlotus_remrecov(.bat) script for Lotus directed recovery	Performing Lotus directed recovery with the GUI on page 215
For Sybase only:			

Table 26 Commands or library names for 32-bit and 64-bit coexistence (continued)

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
nsrdasv (.exe)	nsrdasv32 (.exe)	In Sybase manual backup command	Performing Sybase manual backups with the nsrdasv command on page 176
nsrsybcc (.exe)	nsrsybcc32 (.exe)	In Sybase command for database consistency check	Performing Sybase database consistency checks before backups on page 175
nsrsybrc (.exe)	nsrsybrc32 (.exe)	In Sybase restore command	Performing Sybase data restores with the nsrsybrc command on page 254

CHAPTER 8

Oracle DBA and NetWorker Backup Administrator Collaboration

This chapter includes the following topics:

- [DBA disk backup issues in Oracle environments](#)..... 334
- [NMDA solution for DBA disk backup issues](#)..... 334
- [Configuring backups of DBA disk backups](#)..... 336
- [Reporting on backups of DBA disk backups](#)..... 340
- [Purging DBA disk backups](#)..... 342
- [Restore and recovery of DBA disk backups](#)..... 342

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 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, which is 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 when the 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 that is 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 336 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 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 54.

Note: NMDA scheduled backups of Oracle disk backups do not support the following regular NMDA Oracle features:

- Backup copies
- Oracle Data Guard
- Oracle Exadata
- Restartable backups
- Snapshot backups and restores

Configuring backups of DBA disk backups

About this task

 **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 86 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 *NetWorker Module for Databases and Applications Installation Guide*.

Procedure

1. Start the backup configuration wizard:
 - a. In the NMC **Enterprise** window, right-click the NetWorker server name and select **Launch Application**.
 The *NetWorker Administration Guide* provides details on how to access the NMC interfaces.
 - b. In the **Administration** window, click **Protection**.
 - c. Start the wizard by using the correct method:
 - To create a backup configuration, right-click **Clients** in the left pane and select **New Client Wizard**.
 - To modify a backup configuration that is created with the wizard, right-click the NMDA client in the right pane and select **Modify Client Wizard**.
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 that are 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 387.

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 338.
- 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 286. The Oracle DBA must provide the values that are 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 338.

Results

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 on an enabled FRA.
- i** **NOTICE** Ensure that the Oracle DBA does not have the Change Application Settings privilege. With this privilege, the DBA could accidentally delete some NMDA backup entries from the NetWorker index by running Oracle commands that delete obsolete Oracle disk backups.
- i** **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 336.

Table 27 Advanced NMDA parameters for backups of DBA disk backups

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.	<ul style="list-style-type: none"> • Hostname of the physical host on which the backup session runs (default). • Valid NetWorker client hostname.
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: <ul style="list-style-type: none"> • The NetWorker server is not ready because the devices are not mounted. • The <code>nsrindexd</code> service of the NetWorker server is busy due to other client sessions. Optional for a backup.	<ul style="list-style-type: none"> • 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.	<ul style="list-style-type: none"> • FALSE (default) = Enable multiplexing on the device. • TRUE = Disable multiplexing on the device. <p>i Note: Do not set this parameter to TRUE for nontape devices, such as AFTD or DD Boost devices.</p>

Table 27 Advanced NMDA parameters for backups of DBA disk backups (continued)

Parameter	Description	Default and valid values
NSR_SERVER_NIC	Specifies the name of a network interface card (NIC) on a NetWorker server. Optional for a backup.	<ul style="list-style-type: none"> 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.

About this task

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 and scheduled backup configurations.

Procedure

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 28 NetWorker Probe resource attributes

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, <code>nsrdaprobe</code> , which checks (“probes”) for the condition that triggers a probe-based backup.	
Command Options	Specify a comma-separated list of parameters with their settings.	
	Parameter	Description
	DBA_DISK_BACKUP	Set to TRUE. With <code>DBA_DISK_BACKUP=TRUE</code> , 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 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 provides details.	

2. Configure the scheduled backup by following the instructions in [Configuring the data protection policy with NMC](#) on page 91.

You must configure the required Client resource, group, backup policy, workflow, probe action, and backup action:

- a. Configure the Client resource by using the NMC Administration interface. In the Probe attribute of the Client resource, specify the name of the Probe resource from [step 1](#). You can associate a Client resource with only one probe.
- b. Create the group to contain the Client resource by using the NMC Administration interface. The *NetWorker Administration Guide* provides details on the NMC interfaces.


 **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.

- c. Create the backup policy and workflow by using the NMC Administration interface. You must assign the group to the workflow, and assign the workflow to the backup policy.
- d. Create the probe action and backup action by using the Policy Action Wizard in the NMC Administration interface:
 - When you create the probe action with the wizard, you must select **Probe** from the **Action Type** list.
 - When you create the backup action with the wizard, you must select **Backup** from the **Action Type** list, and select **Traditional** from the secondary action list.

You must assign the probe action and backup action to the workflow. In the workflow, the probe action must precede the backup action. The backup action is performed only if the probe conditions are met during the probe action.

Reporting on backups of DBA disk backups

About this task

 **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 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 in the backup. The `nsrorainfo` 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 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 *NetWorker Module for Databases and Applications Command Reference Guide* provides details about the `nsrorainfo` command.

Purging DBA disk backups

About this task

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 Change Application Settings privilege.

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.

About this task

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 that are 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, if the NMDA and NetWorker software is installed on the destination host.

Complete the following steps to enable a recovery to a remote host.

Procedure

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
```

Results

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 that is 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](#)..... 346
- [Configuring NSM snapshot backups](#)..... 351
- [Configuring NSM snapshot restore and recovery](#)..... 355
- [Setting the NSM parameters](#)..... 358
- [DB2 considerations for NSM operations](#)..... 362
- [Oracle considerations for NSM operations](#)..... 368
- [Snapshot operations with Oracle ASM](#)..... 397

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. The NSM functionality is available as part of the NetWorker extended client software.

NMDA supports NSM snapshot and ProtectPoint backups and restores of DB2 and Oracle data on the primary storage, for example, VNX Block (CLARiiON), VMAX (Symmetrix), or XtremIO. The NSM documentation provides additional details about NSM backups and restores.

NMDA supports scheduled backups for all the supported types of NSM backups.

NMDA supports manual (client-initiated) backups for only the supported types of ProtectPoint backups:

- NSM ProtectPoint backups of data on a VMAX array to a Data Domain system.
- NSM ProtectPoint with RecoverPoint backups of data on an XtremIO array to a Data Domain system.

Note:

The term *ProtectPoint* in this guide refers to both a ProtectPoint operation with NSM and a ProtectPoint with RecoverPoint operation with NSM, unless stated otherwise.

The term *snapshot* in this guide refers to both a storage-only snapshot backup with NSM and a ProtectPoint backup with NSM.

The following topics describe the software requirements and the supported types of NSM backups and restores:

- Software requirements for NSM snapshot operations
- Types of NSM backups
- Types of NSM restores
- NSM backup processes

The following topics describe how to configure and perform NSM snapshot operations:

- Configuring NSM snapshot backups
- Configuring NSM snapshot restore and recovery
- DB2 considerations for NSM operations
- Oracle considerations for NSM operations

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 29 Typical configuration for NSM operations with NMDA

Computer or device	Required software or configuration
Database server host	DB2 or Oracle server, NetWorker client, NMDA
Mount host	NetWorker client, NetWorker storage node
Storage array	Storage array that NMDA supports with NSM, for example, VMAX, VNX Block, or XtremIO

Table 29 Typical configuration for NSM operations with NMDA (continued)

Computer or device	Required software or configuration
NetWorker server host	NetWorker server, NetWorker storage node, NetWorker client
Data Domain	Required only for ProtectPoint operations, vdisk service is enabled

Types of NSM backups

NMDA supports all types of NSM backup operations as described in the following table. The following topics provide more details.

Table 30 Supported types of NSM backups

Backup type	Snapshot retention	Location of stored data
Snapshot backup	Permanent	Primary storage
Clone to NetWorker conventional media	Permanent	Secondary storage, such as AFTD or tape
ProtectPoint backup or clone to NetWorker ProtectPoint device	Permanent	Data Domain (DD) vdisk
Clone-Controlled Replication (CCR)	Permanent	Data Domain (DD, DD1, and DD2) vdisks

Snapshot backup

A snapshot backup creates a permanent point-in-time copy or snapshot of the data on the primary storage system. The snapshot is available for performing snapshot restores or for creating a clone of the snapshot to secondary storage (also known as conventional media), such as an AFTD or DD Boost device. 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 NetWorker data protection policy for a scheduled backup to control the lifecycle of the snapshot. The policy specifies the frequency of snapshot backups and the minimum amount of time that snapshot copies are retained before being recycled or released.

The NetWorker data protection policies do not apply to a manual ProtectPoint backup. The manual backup uses either user-set options or default options, such as the retention period.

Clone to conventional media

You can use NetWorker to perform a clone backup of the snapshot, which creates a cloned copy of the snapshot on secondary storage, such as an AFTD or DD Boost device. NSM can use a mount host that is different from the production host to clone or roll over the backup to secondary storage. The NSM documentation provides details about when you must use a mount host.

By using the NetWorker data protection policy, you can perform the clone backup immediately after the snapshot backup in the same backup workflow. Alternatively, you can perform the clone backup in a separate workflow.

The permanent snapshot is retained on the primary storage and is available to NMDA for performing snapshot restores for the period that is specified by the NetWorker policies.

ProtectPoint backup or clone

You can use NMDA for an NSM ProtectPoint backup to create a permanent point-in-time copy of the data from a VMAX array to a DD vdisk device. You can use NMDA for an NSM ProtectPoint with RecoverPoint backup to create a permanent point-in-time copy of the data from an XtremIO array to a DD vdisk device.

You can configure either of two types of ProtectPoint backups:

- A backup directly to the DD vdisk device.
You can then use NMDA to restore directly from the backup on the Data Domain system.
- A backup that creates a snapshot on the VMAX or XtremIO array, followed by a clone backup that backs up the snapshot to the DD vdisk device. The initial snapshot remains on the VMAX or XtremIO array, and the clone copy resides on the DD vdisk device.
You can then use NMDA to restore from either the snapshot on the array or the clone copy on the Data Domain system.

Types of NSM restores

NMDA supports all types of NSM restores as described in the following table. The following topics provide details.

Table 31 Supported types of NSM restores

Restore type	Data is retrieved from
Restore from snapshot or ProtectPoint copies (backup or clone)	Mounted snapshot from storage array or DD
Restore from clone of snapshot or ProtectPoint backup on secondary storage (conventional media)	Secondary storage
Rollback restore	Unmounted snapshot on storage array ProtectPoint backup on DD, with a rollback restore at the storage array level ProtectPoint with RecoverPoint backup on DD, with a rollback restore at the consistency group level

The `RESTORE_TYPE_ORDER` parameter setting specifies the type of NSM restore. [Configuring NSM snapshot restore and recovery](#) on page 355 provides details.

Restore from snapshot or ProtectPoint copies (backup or clone)

A snapshot restore is a file-level restore from a snapshot, performed by automatically mounting the snapshot to the mount host and copying back to the requested location. To minimize the recovery time, use the application host to which you are recovering as the mount host, if possible.

Restore from clone on secondary storage (conventional media)

A restore from a clone on a secondary storage system is the same as a traditional restore without NSM. However, this type of restore still occurs through NSM.

Rollback restore

A rollback restore is a restore of an entire snapshot backup that uses the storage array capabilities and the Data Domain capabilities for a ProtectPoint operation. You can run either a regular rollback restore or a redirected rollback restore to alternate target LUNs in the same VMAX array. A redirected rollback restore is supported with ProtectPoint for VMAX and SnapVX backups.

You run a regular rollback restore to restore the backup to the original source LUNs on the backup host. With this NMDA release, you can perform a redirected rollback restore of a ProtectPoint for VMAX database backup to relocate a database to an alternate host, configured on target LUNs in the same VMAX array.

Note:

- This NMDA release does not support a redirected rollback restore to alternate LUNs on the original backup host. During a rollback restore to the original backup host, the snapshot backup is restored to the original source LUNs. The backup file systems must exist on the source LUNs before the restore; re-create the file systems if required.
- A partitioned disk is not supported in a snapshot operation, such as a snapshot backup or rollback restore. Do not use a partitioned disk when you create the target file system. In a rollback restore, on the target devices, any extra file systems and volume management that reside on partitioned disks and are not involved in the restore must be manually cleaned up before the restore. Otherwise, the rollback restore might fail.

Use the NSM documentation to determine any limitation with the rollback on a specified storage array or in a ProtectPoint workflow.

NOTICE

A rollback restore is a destructive restore because the rollback overwrites the entire contents of a snapshot unit, such as a volume or disk.

A ProtectPoint with RecoverPoint backup and rollback restore occur at the consistency group level, regardless of which objects are included in the backup command. As a best practice for a ProtectPoint with RecoverPoint rollback restore, when you perform the backup or rollback restore, do not exclude the logs or any database files that are part of the RecoverPoint consistency group being backed up or restored. If any LUNs in the backed-up consistency group contain objects that were not included in the backup command, ensure that you manually unmount those LUNs before the rollback restore and then manually mount the LUNs back after the restore.

NSM backup processes

You can start an NSM backup by automatic or manual invocation of the NetWorker scheduled backup through the NetWorker data protection policies.

You can also use a manual backup for an NSM ProtectPoint backup from VMAX to DD, or for an NSM ProtectPoint with RecoverPoint backup from XtremIO to DD. For example, you can start the manual backup with an Oracle RMAN or DB2 backup command.


In general terms, an NSM backup includes the following processes.

1. For a scheduled backup, the NetWorker server starts the NMDA `nsrdbasv` process, which performs the DB2 or Oracle RMAN backup. For a manual backup, the user starts the DB2 or 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. Optionally for storage-only snapshots, the snapshot is mounted on the mount host for validation.
4. 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 *NetWorker Administration Guide* describes the NetWorker server and client services.


If the backup workflow includes a clone action, NSM and the NetWorker software also produce a clone copy on secondary storage as described in the NSM documentation.

The following figure illustrates the data flow of an NSM snapshot and clone operation to secondary storage in an NMDA environment. This example environment does not include support for ProtectPoint workflows with a VMAX or XtremIO array and a Data Domain system. The NSM feature provides the snapshot backup functionality as part of the NetWorker extended client software.

Figure 14 NSM snapshot and clone to secondary storage data flow with NMDA



The following figure illustrates the data flow of an NSM snapshot restore in an NMDA environment. The NetWorker storage node restores data from the snapshot target volume to the production source volume. This example environment does not include support for ProtectPoint workflows with a VMAX or XtremIO array and a Data Domain system.

Figure 15 NSM snapshot restore data flow with NMDA

Configuring NSM snapshot backups

You must configure the required NetWorker resources and parameters for NMDA snapshot backups with NSM. The NetWorker Client Configuration wizard is recommended for configuring the Client resource and parameter settings. You can also use the manual configuration method without the wizard.

About this task

The following configuration steps apply to scheduled NSM backups. If you perform only manual (client-initiated) backups, then you can use the configuration wizard or you can manually create an RMAN script or DB2 backup command with an NMDA configuration file to run on the client host.

Note:

To create the required lockbox entries for RecoverPoint and XtremIO operations, you must use the backup configuration wizard first, before you create the NetWorker client resource manually. For example, in the wizard, you can select **EMC ProtectPoint for RecoverPoint** as the type of storage, and then you must enter the RecoverPoint username and password information. The lockbox is created when the Client resource is created with the configuration wizard.

For a manual backup that you perform with the native CLI, you can specify a retention time for the backup by setting the `NSR_SAVESET_RETENTION` parameter value in UNIX time format. Use the same parameter setting in both the NMDA and NSM configurations. Set the parameter in the NMDA configuration file.

Procedure

1. Ensure that you have installed both NMDA and the NetWorker extended client software on the database host according to the instructions in the following documents:
 - *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 NetWorker extended client software on the mount host and correctly configured the Client resource for the mount host.
3. Review [Types of NSM backups](#) on page 347 to determine which type of snapshot backup to perform.
4. Ensure that partitioned disks will not be used in the snapshot backups. A partitioned disk is not supported in a snapshot operation, such as a snapshot backup or rollback restore.
5. Ensure that you have completed the basic database server configuration and NetWorker configuration according to [Configuring NMDA backups](#) on page 76.
6. If required, configure internationalization (I18N) support according to [Configuring internationalization \(I18N\) support](#) on page 353.
7. Use the configuration wizard to configure the Client resource and parameter settings for the snapshot backup. The wizard is the recommended configuration method. Follow the information in [Configuring a scheduled backup with the wizard](#) on page 87 except skip any sentence that includes a cross-reference.

On the **Select the Backup Configuration Type** page in the wizard:

- a. In the **Available Applications** table, select **DB2** or **Oracle** as the application type.
- b. Select the checkbox for **Enable NetWorker Snapshot Management on the selected application**.
- c. Click **Next**.

On the **Select the Snapshot Management Options** page, select the type of storage array or storage appliance where the snapshots will be created:

- **EMC VMAX/Symmetrix**
- **EMC ProtectPoint for VMAX3** (Select for ProtectPoint for VMAX operations)
- **EMC ProtectPoint for RecoverPoint** (Select for ProtectPoint for XtremIO operations)
- **EMC VNX/CLARiiON**

Only the arrays or ProtectPoint option that the specified application supports on the client operating system will appear as available.

The NSM documentation provides details on all the other snapshot-specific fields in the configuration wizard. Provide the required information on all the wizard pages to complete the configuration.

8. Review the considerations and perform any required procedures for DB2 or Oracle snapshot backups:
 - [DB2 considerations for NSM operations](#) on page 362
 - [Oracle considerations for NSM operations](#) on page 368
9. If you cannot use the wizard method for some reason, complete the following steps:
 - a. Configure the Client resource according to [Configuring Client resources manually for NSM backups](#) on page 354.
 - b. Set the required NSM parameters according to [Setting the NSM parameters](#) on page 358.

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](#) provides details.

10. For a scheduled snapshot backup, configure the backup group, policy, workflow, and action according to the "Data Protection Policies" chapter in the latest *NetWorker Snapshot Management Integration Guide*.

When you create the backup action for the snapshot backup with the Policy Action Wizard:

- From the **Action Type** list, select **Backup**.
- From the secondary action list, select **Snapshot**.
- From the **Minimum Retention Time** list, specify the minimum amount of time to retain the snapshot backup. After this time has elapsed, the snapshot can be recycled to release resources that are needed for new backups.

When you create an optional clone action to follow the backup action in the snapshot backup workflow, select **Clone** from the **Action Type** list in the Policy Action Wizard.

11. Test the NSM snapshot backup configuration according to [Testing an NSM scheduled backup](#) on page 354.

Configuring Pool resources for NSM backups

You must configure a pool to support NSM backups. You configure the Pool resource by using the same method as for a regular NMDA backup. However, the specified backup device must be an advanced file type.

About this task


For NSM backups, in addition to creating snapshots on the storage array (or DD in the case of ProtectPoint backups), NMDA and NSM must save the following types of data:

- Snapshot save set records.
- Any application files that cannot be backed up through a snapshot workflow. These files are backed up through a traditional workflow.

Ensure that you include the required devices in the NSM backup pool:

- For snapshot backups, include an AFTD or DD Boost device in the pool that is used for the NSM backups.
- For ProtectPoint backups, in addition to the preceding devices, include a NetWorker ProtectPoint device in the same pool. The *NetWorker Snapshot Management Integration Guide* describes how to create the NetWorker ProtectPoint device.

For a scheduled NSM backup, you specify the pool name in the **Destination Pool** field of the backup action, which you create for the snapshot backup as part of the policy-based configuration.

 **Note:** The Pool attribute setting in the Client resource can override the **Destination Pool** setting in the backup action.

The *NetWorker Snapshot Management Integration Guide* provides additional details and describes the pool configuration that is required for a snapshot clone to conventional media or to another ProtectPoint device.

If you will perform an NSM ProtectPoint backup as a manual backup, then you must specify the destination pool in the NSR_DATA_VOLUME_POOL parameter in the NMDA configuration.

Configuring internationalization (I18N) support

About this task

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 77.

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 79 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 388 describes the `nsroraadmin` command. The command sets the parameter value in the NWORA resource file.

Configuring Client resources manually for NSM backups

About this task

Configure the NetWorker Client resource manually for the database server host according to [Configuring the Client resource](#) on page 92.

The following Client resource requirements for the database server host apply to NSM snapshot backups:

- The Retention Policy attribute in the Client resource applies only to the NetWorker media database entries for NSM backups on secondary storage.
- For the Backup Command attribute, specify 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 mount host, the Remote Access attribute must include the mount hostname.
- Use the Application Information attribute in the Client resource to set the NSM parameters as described in [Setting the NSM parameters](#) on page 358.

Create a NetWorker Client resource for the mount host according to the NSM documentation.

Testing an NSM scheduled backup

About this task

 **NOTICE** You can test an NMDA scheduled backup with NSM by manual invocation of the scheduled backup workflow.

To verify the scheduled backup setup, follow the instructions for regular backups in [Testing scheduled backups](#) on page 166.

Canceling an NSM snapshot backup

About this task

Cancel an NSM snapshot backup by using the same methods as used for nonsnapshot scheduled backups. [Canceling scheduled backups](#) on page 166 provides details.

Configuring NSM snapshot restore and recovery

You can use the nonwizard configuration method and then run the NSM snapshot restore by using the application native CLI command as done for a backup. You can also configure and run an NSM snapshot restore by using the NMDA NSM recovery wizard.

Ensure that you have met the common requirements for database object access permissions in an NSM snapshot restore. The numeric user ID (UID) and group ID (GID) of the target database/instance owner must match the original UID and GID that were captured during the NSM snapshot backup.

You can restore only the files for which you have read permission, based on the files' operating system permissions at the time that the files were backed up. On a UNIX or Linux system, the read permission is associated with the numeric user ID (UID) and group ID (GID), not the username or group name. The UID and GID of the user that performs the restore must match the IDs associated with the files at backup time.

If you use the NMDA NSM recovery wizard, then you can set the NSM parameters in the **Advanced Options** table on the **Specify the Performance and Additional Options (Optional)** page in the wizard.

If you use the nonwizard configuration method, set the required parameters for the particular application:

- For a DB2 NSM restore, set the NSM parameters in a configuration file with its pathname specified in the NSR_PROXY_PFILE parameter as described in [NSR PROXY PFILE](#). Set the NMDA parameters in the NMDA configuration file as described in [NMDA configuration file](#) on page 400.
- For an Oracle NSM restore, set the parameters as described in [NSM parameter settings](#) on page 372.

Set the following parameters for an NSM snapshot restore:

- NSR_DATA_MOVER—Hostname of the mount host that is 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) restore.
 - conventional—Specifies a snapshot restore from secondary storage media.
 - rollback—Specifies a rollback restore from a snapshot copy.

The following topic provides additional configuration requirements for a rollback restore.


The default value of RESTORE_TYPE_ORDER is pit:conventional.

If you specify multiple values for RESTORE_TYPE_ORDER, the software tries each type of restore in the order that is 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 348 describes the supported restore types.

- NSR_DD_VDISK_RESTORE_DEVGRPNAME—Mandatory only for a ProtectPoint restore directly from Data Domain or a ProtectPoint with RecoverPoint restore, except for a rollback restore. Specifies the DD vdisk device group in the vdisk device pool that contains the restore LUNs to use for the restore.

- **NSR_DD_VDISK_RESTORE_POOLNAME**—Mandatory only for a ProtectPoint restore directly from Data Domain or a ProtectPoint with RecoverPoint restore, except for a rollback restore. Specifies the name of the DD vdisk device pool to use for the restore. The specified device pool must contain the restore LUNs.
 - **NSR_RECOVER_POOL**—Set this parameter to restore a ProtectPoint snapshot from a specific clone pool if there are multiple snapshot copies (clones) on Data Domain systems. [NSR_RECOVER_POOL](#) provides details.
-  **Note:** NSR_RECOVER_POOL, NSR_SNAP_TYPE, and NSR_SNAP_TECH are required for a ProtectPoint with RecoverPoint rollback restore.

The NSM documentation describes the parameters that are used to restore snapshots.

Configuration requirements for a rollback restore

The *NetWorker Snapshot Management Integration Guide* provides details on the following configuration requirements for a rollback restore of an NMDA NSM backup.

NMDA does not support partitioned disks with restored file systems in a rollback restore. On the target devices, any volume group (and its logical volumes and file systems) or file systems that are not involved in the restore on partitioned restore disks must be manually cleaned up before the restore. The extra file systems must be unmounted and removed, and the extra volume group must be removed. Otherwise, the rollback restore might fail.

For a regular rollback restore to the original source LUNs on the backup host, the file system with the same mount point as used in the backup must exist and be mounted on the host. If Logical Volume Manager (LVM) is used, then the volume group name must be the same.

For a redirected rollback restore of a ProtectPoint for VMAX backup to a different set of LUNs on an alternate host:

- NMDA software must be installed and configured properly on the host that performs the rollback restore. The CLIENT parameter must be set to the original value, as recorded in the backup.
- You must meet the common requirements for database object access permissions in an NSM snapshot restore. The numeric user ID (UID) and group ID (GID) of the target database/instance owner must match the original UID and GID that were captured during the NSM snapshot backup.
- The file system with the same mount point as used in the backup must exist and be mounted on the alternate host.
- The number of devices on which the file system resides on the alternate host must be equal to the original number of devices in the backup.
- The size of the target LUN must be equal to or greater than the size of the original LUN.
- When multiple LUNs are included in the rollback restore, the destination LUN size must be greater than or equal to the static image size.
- For an Oracle rollback restore, the Oracle-Managed Files (OMF) feature must be disabled for the Oracle database on the alternate host because the Oracle rollback restore in this release does not support the renaming of the restored files.
- If file system management is used, such as LVM or Veritas Volume Manager:
 - If a file system or volume manager exists on the backed-up devices, the file system or volume manager version on the recovery host might need to be the same as or higher than the version on the backed-up devices. The file system and volume manager documentation provides details.
 - The names of the volume group, logical volume, and physical device on the target devices do not need to match the original names, provided that no conflicts exist in the logical volume and volume group names.

- The number of file systems, volume groups, and logical volumes of the target devices do not need to match the numbers in the original devices configuration.

Note: Any extra file systems, volume groups, and logical volumes on the recovery host must be listed in the `psrollback.res` file, so that these items are skipped during the safety checks. The following topic provides details about safety checks and the `psrollback.res` file.

- Additional affected file systems that are not a part of the rollback restore will not be mounted or unmounted.

Ensure that no conflicts exist in the file system configuration that could cause issues during an LVM import. For example, ensure that no volume group or logical volume exists outside of the restore devices with the same name as used in the backup.

Configuring the safety checks for a rollback restore

During a rollback restore, NSM performs safety checks by default. The safety checks ensure that there are no files, directories, partitions, or volumes (data targets) on the rollback target LUN other than those restored with NSM. If there are additional such data targets on the target LUN that are not included in the restore session, NSM fails the rollback restore as a safety precaution to prevent the overwriting of data.

For a rollback restore of a RecoverPoint consistency group, the safety checks also ensure that all the XtremIO LUNs of the target consistency group are included in the rollback restore. If some of the LUNs do not contain data objects that are being restored, then the safety check and rollback restore both fail.

To override the safety checks, you can use the `psrollback.res` file. In the file, you must list all the files and directories to be excluded from the rollback safety checks. For a rollback restore of a RecoverPoint consistency group, you must also list in the file all the mount points of the LUNs to be excluded from the rollback safety checks.

For example, `lv011` is the logical volume at backup time, and `lv011` and `lv012` are logical volumes on the destination host. You must include `lv012` in the `psrollback.res` file to enable the rollback restore to proceed. You can also list the device name to ensure that the safety check skips all the file systems that reside on the device. To prevent `lv012` from being overwritten during the rollback restore, do not list `lv012` or the device name in the file.

NOTICE

Use the `psrollback.res` file with extreme caution to prevent possible data corruption. If you use this file to override the safety checks, the rollback restore might overwrite some database files that were not included in the restore session, such as Oracle online redo logs, which could result in data loss.

On Linux or Solaris SPARC, if a disk is configured with partitions, you can perform a rollback restore only if you list the entire disk in the `psrollback.res` file. The rollback restore then overwrites the entire disk. For example, if `/fs1` and `/fs2` are configured with partitions `/dev/sdc1` and `/dev/sdc2` respectively, then you must enable the rollback restore of `/fs1` by listing the entire disk `/dev/sdc` in `psrollback.res`. The rollback restore overwrites the entire disk `/dev/sdc`, so `/fs2` is also restored.

If a logical volume manager (LVM) controls the file system of an application host, then you must list in the `psrollback.res` file all the physical disks that belong to the LVM volume group. For example, if a volume group contains the disks `/dev/sdc` and `/dev/sdd`, and `/fs1` is the mount point of the file system, then the `psrollback.res` file must include the following lines:

```
/fs1/lost+found
/fs1/test
```

```
/dev/sdc
/dev/sdd
```

The `psrollback.res` file location is as follows:

- On UNIX systems: `/nsr/res/psrollback.res`
- On Windows systems: `C:\Program Files\EMC NetWorker\nsr\res\psrollback.res`

Example 37 Overriding the safety checks during a rollback restore

If you are restoring `/fs1/data1.df` and `/fs1/data2.df` but there are other files in the `/fs1` directory, such as the files `lost+found` and `test`, you can exclude these other files from the safety checks during a rollback restore if you do not need these files. To exclude the files, list the file pathnames in the `psrollback.res` file:

```
more /nsr/res/psrollback.res
```

```
/fs1/lost+found
/fs1/test
```

Refer to the application-specific sections that follow for any application-specific requirements for `psrollback.res` and rollback restores.

Setting the NSM parameters

The following tables describe the supported NSM parameters. The NSM documentation provides details on the NSM parameters.

About this task

If you do not use the configuration wizard, 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 scheduled backup and in the `NSR_PROXY_PFILE` configuration file for a manual backup or restore as described in [NSR_PROXY_PFILE](#).
- For Oracle NSM operations, set the parameters as described in [NSM parameter settings](#) on page 372.

The following table lists the common parameters for all NSM operations. The list is not exhaustive.

Table 32 Common NSM parameters

Parameter	Description	Default and valid values
<code>NSR_DATA_MOVER</code>	Mandatory for an NSM backup or restore that uses a mount host. Specifies the hostname of the mount host.	<ul style="list-style-type: none"> • Hostname of the local application host (default). • Valid hostname of the mount host.

Table 32 Common NSM parameters (continued)

Parameter	Description	Default and valid values
NSR_MAX_STREAMS	Optional. Specifies the maximum number of restore streams.	<ul style="list-style-type: none"> 16 (default). An integer value.
NSR_PS_SAVE_PARALLELISM	Optional. Specifies the maximum number of concurrent save streams per NSM backup. This parameter is only used when cloning the snapshot to secondary storage.	<ul style="list-style-type: none"> 16 (default). 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 <code>nmda_oracle_proxy.messages</code> under the diagnostic messages destination.	<ul style="list-style-type: none"> None (default). Full pathname of the script to run immediately after the snapshot is complete but before the backup goes to NetWorker devices. <p>If the script pathname is incorrect or the user does not have execute permissions, the backup fails before starting and generates an error.</p> <p>If the script pathname is correct but the script execution fails, the backup does not fail.</p>
RESTORE_TYPE_ORDER	Optional. Specifies the type of NSM restore to perform.  Note: If you specify multiple values, the software tries each type of restore in the order that is specified until a restore operation succeeds.	<ul style="list-style-type: none"> pit:conventional (default). One or more of the following values, with each value separated from others by a colon (:): <ul style="list-style-type: none"> - pit—Specifies a snapshot (PIT) restore - conventional—Specifies an NSM restore from secondary storage media - rollback—Specifies a rollback restore from a point-in-time copy <p>Configuring NSM snapshot restore and recovery on page 355 provides details.</p>

The following table lists the NSM parameters specifically for ProtectPoint for VMAX operations. The list is not exhaustive.

Table 33 NSM parameters for ProtectPoint for VMAX operations only

Parameter	Description	Default and valid values
NSR_DD_VDISK_RESTORE_DEVGRPNAME	Mandatory for a ProtectPoint for VMAX restore directly from Data Domain except for a rollback restore.	<ul style="list-style-type: none"> Undefined (default). Valid name of a DD vdisk device group.

Table 33 NSM parameters for ProtectPoint for VMAX operations only (continued)


Parameter	Description	Default and valid values
	Specifies the DD vdisk device group in the vdisk device pool that contains the restore LUNs to use for the restore of the ProtectPoint backup directly from Data Domain. The restore does not involve a VMAX system.	
NSR_DD_VDISK_RESTORE_POOLNAME	<p>Mandatory for a ProtectPoint for VMAX restore directly from Data Domain except for a rollback restore.</p> <p>Specifies the name of the DD vdisk device pool to use for the restore of a ProtectPoint backup directly from Data Domain. The specified device pool must contain the restore LUNs that are provided on the restore host.</p>	<ul style="list-style-type: none"> Undefined (default). Valid name of a DD vdisk device pool.

The following table lists the NSM parameters specifically for ProtectPoint with RecoverPoint (for XtremIO) operations. The list is not exhaustive.

Table 34 NSM parameters for ProtectPoint with RecoverPoint operations only

Parameter	Description	Default and valid values
NSR_DD_VDISK_RESTORE_DEVGRPNAME	<p>Optional for a ProtectPoint with RecoverPoint restore. Not used for a rollback restore.</p> <p>Specifies the DD vdisk device group in the vdisk device pool that contains the restore LUNs to use for the restore of the ProtectPoint with RecoverPoint backup.</p>	<ul style="list-style-type: none"> Undefined (default). Valid name of a DD vdisk device group.
NSR_DD_VDISK_RESTORE_POOLNAME	<p>Optional for a ProtectPoint with RecoverPoint restore. Not used for a rollback restore.</p> <p>Specifies the name of the DD vdisk device pool to use for the restore of a ProtectPoint with RecoverPoint backup. The specified device pool must contain the restore LUNs.</p>	<ul style="list-style-type: none"> Undefined (default). Valid name of a DD vdisk device pool.
NSR_SNAP_TECH	<p>Mandatory for ProtectPoint with RecoverPoint operations only.</p> <p>Specifies the RecoverPoint replication type for a backup or restore.</p>	<ul style="list-style-type: none"> Undefined (default). RP_CDP—Notifies NSM that local copies will be used to access a bookmark. RP_CRR—Notifies NSM that remote copies will be used to access a bookmark.

Table 34 NSM parameters for ProtectPoint with RecoverPoint operations only (continued)

Parameter	Description	Default and valid values
NSR_SNAP_TYPE	Mandatory for ProtectPoint with RecoverPoint operations only. Specifies the snapshot provider. Set the parameter to "protectpoint" to specify the required ProtectPoint snapshot provider.	<ul style="list-style-type: none"> Undefined (default). protectpoint.
RP_APPLIANCE_NAME	Mandatory for ProtectPoint with RecoverPoint operations. Specifies the hostname or IP address of the RecoverPoint management host.  Note: You must run the backup configuration wizard to create the mandatory lockbox that stores the password for the RecoverPoint user.	<ul style="list-style-type: none"> Undefined (default). Valid hostname or IP address of the RecoverPoint management host.

Example 38 NSM parameter settings for different types of primary storage

The NSM documentation for the specified 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
```

- ProtectPoint:

```
NSR_DATA_MOVER=datamover.emc.com
NSR_SNAP_TYPE=protectpoint
```

- ProtectPoint with RecoverPoint:

```
NSR_DATA_MOVER=datamover.emc.com
NSR_DD_VDISK_RESTORE_DEVGRPNAME=DG_a018_rp211
NSR_DD_VDISK_RESTORE_POOLNAME=ledma018_RP163_restore_1
NSR_SNAP_TECH=RP_CDP
NSR_SNAP_TYPE=protectpoint
RP_APPLIANCE_NAME=ledmd163
```

DB2 considerations for NSM operations

For NSM snapshot operations with DB2 data, the IBM DB2 software provides a feature that is 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

- The database data paths must be on snapshotable devices.
- By default, DB2 ACS includes logs in the snapshot backup, in addition to data files. In a backup that includes logs, the log directories must also be on snapshotable devices. Otherwise, the snapshot backup fails.
- To exclude logs in a snapshot backup or exclude logs in a restore of a snapshot backup that includes logs, the log directories must reside on different disk volumes than other database paths.

The DB2 ACS best practice recommends using a dedicated volume group for log paths, with the log paths contained in a snapshot volume that is separate from the database directory and database containers.

The following IBM article provides details:

http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.1.0/com.ibm.db2.luw.admin.ha.doc/doc/c0053158.html

The `DB2_ACS_LAYOUT_CHECK` parameter specifies whether to enforce the DB2 ACS best practice on the log directory layout of the database during a snapshot backup.

[DB2_ASC_LAYOUT_CHECK](#) provides details.

- 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. NMDA with NSM also does not support an incremental backup.
- For ProtectPoint with RecoverPoint, the backup objects grouping for DB2 objects is per consistency group. To enable the `exclude logs` option, the log objects must be managed in a separate consistency group from the other database objects.

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 122 describes how to configure the DB2 automatic backup of transaction logs to NetWorker devices.

After you have completed the backup configurations in [Configuring NSM snapshot backups](#) on page 351, you can run a DB2 NSM backup from the NetWorker server. You can also run a ProtectPoint NSM backup manually by using the DB2 CLI interface.

Performing manual ProtectPoint NSM backups

You can run the appropriate `db2 backup use snapshot library` command to perform a manual ProtectPoint backup of a DB2 database.

About this task

Before you run a manual ProtectPoint backup, ensure that the required backup configurations are completed as described in [Configuring NSM snapshot backups](#) on page 351.

To perform a manual ProtectPoint backup of a DB2 database, run the `db2 backup` command with the `use snapshot library` option.

For example, run the following command on Linux:

```
db2 backup db SAMPLE use snapshot library /usr/lib/libnsrdb2.so options
@pathname/nmda_db2.cfg
```

where `pathname/nmda_db2.cfg` is the complete pathname of the NMDA configuration file. The DB2 documentation provides details about the `db2 backup` command.

Restoring a DB2 NSM backup

You can use the nonwizard configuration method and then run the NSM snapshot restore by using the application native CLI command as done for a backup. You can also configure and run an NSM snapshot restore by using the NMDA NSM recovery wizard.


About this task

Use the following steps and referenced documentation to restore an NSM backup of a DB2 database by using the CLI commands.

[Configuring NSM snapshot restore and recovery](#) on page 355 provides more information about using the recovery wizard.

Procedure

1. Configure the DB2 NSM restore according to [Configuring NSM snapshot restore and recovery](#) on page 355.
2. Ensure that you meet the following requirements before you perform any restores:
 - Ensure that you meet the common requirements for database object access permissions in an NSM snapshot restore. The numeric user ID (UID) and group ID (GID) of the target database/instance owner must match the original UID and GID that were captured during the DB2 NSM backup.
 - All the file system mount points in the backup are re-created with the proper ownership and permissions.
A DB2 NSM restore does not restore the ownership and permissions of the mount points and the file system directories above them.
 - If the database contains symbolic links, then the symbolic links are re-created before you perform a restore.
A DB2 NSM backup does not back up symbolic links.
 - For a redirected rollback restore of a ProtectPoint for VMAX backup to an alternate host:
 - You have created the target file systems based on the file system configuration information in the backup. The target file system configuration must meet the requirements in [Configuration requirements for a rollback restore](#) on page 356.

 **Note:** The database does not need to exist before the rollback restore.

It is recommended that you keep the configuration information about the backed-up file system, as stored in the snapshot backup metadata, for use in re-creating the target file system. For example, with a DB2 snapshot backup in this NMDA release, you can obtain the save set ID of the snapshot save set with the `mminfo -avot | grep <backup_timestamp>` command and then extract the file system information with the `mminfo -aSvot -qssid=<save_set_ID>` command.

- You perform the rollback restore of an entire database, including the logs.
 - The database manager configuration parameter, `DFTDBPATH`, is preferably set to the database path value in the backup. This setting helps with any manual clean-up that might be required after a snapshot restore failure.
 - For a restore to a new database, as recommended by DB2, you use a consistent whole database backup. If you use an offline backup, then you can perform the rollback restore with the `without rolling forward` option.
3. To restore a database, run the `db2 restore` command with the `use snapshot library` option.

For example, run the following command 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 complete pathname of the NMDA configuration file. The DB2 documentation provides details about the `db2 restore` command.

4. 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:

About this task

- 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. The `db2 prune history` command and the automatic deletion of the recovery objects configuration do not clean up the DB2 snapshot backups.

Querying DB2 snapshots

Query of DB2 snapshots with the `db2acsutil` command produces a list of valid snapshots that are retained in the repository.

About this task

 **Note:** You cannot monitor the status of snapshots that are created with NMDA.

The following examples are snapshot queries on AIX:

```
db2acsutil load /usr/lib/libnsrdb2.so options @pathname/nmda_db2.cfg query
snapshot older than 10 days ago 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 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 that are created with NMDA supports only the `taken at yyyymmddhhmmss` option of the `db2acsutil` command. The software deletes snapshot entries from the NetWorker server indexes.

About this task

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`. [NMDA Parameters and Configuration File](#) on page 399 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 db SAMPLE taken at 20100612121212
```

where *pathname/nmda_db2.cfg* is the full pathname of the NMDA configuration file.

The DB2 documentation provides details.

In a `db2acsutil` deletion operation with a DB2 timestamp, if there is an error in searching for the snapshot save sets to delete, as when the save sets have already expired and been deleted, then the `db2acsutil` command returns an error without deleting the DB2 metadata save sets. In this case, if the `DB2_ACS_METADATA_DELETION_FORCE` parameter is set to `TRUE`, then the `db2acsutil` command deletes the corresponding metadata save sets.

Pruning DB2 snapshots

The pruning of DB2 snapshots requires the deletion of both the snapshot save sets and the corresponding metadata save sets.

About this task

For a DB2 snapshot backup, NMDA DB2 creates a metadata save set record that contains the information that is required to retrieve the actual snapshot save sets. These two types of save sets are stored separately in the NetWorker backup storage. However, the snapshot save sets information is embedded within the metadata save set entries for reference and tracking purposes.

The `db2acsutil` utility operations are based on the metadata records. For example, the `db2acsutil` utility uses the metadata records in a query for the DB2 snapshot backups available for restore.

When snapshot save sets are deleted due to the snapshot retention policy, the corresponding NMDA DB2 metadata should also be deleted. Otherwise, the output of a `db2acsutil` query operation might cause confusion.

Through the `nsrdb2cat` binary, NMDA supports the synchronous deletion (“pruning”) of the NMDA DB2 metadata save sets when the snapshot save set entries expire and are deleted from the NetWorker indexes. Before deleting the index entries, the `nsrsnapck` program runs the `nsrdb2cat` program to delete the corresponding metadata catalog entries and prune the backup entries in the DB2 recovery history.

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`, which 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. [NMDA configuration file syntax](#) on page 402 provides details.

Table 35 DB2 resource file parameters

DB2 resource file parameter	Definition	Default and valid values
DB2PATH	Mandatory only if <code>NSR_DB2CAT_MODE</code> is enabled. Specifies the location of the DB2 binary directory.	<ul style="list-style-type: none"> Undefined (default). Pathname of the DB2 binary directory. For example: <code>DB2PATH = /opt/ibm/db2/V10.5/bin</code>
NSR_DB2CAT_MODE	Mandatory. Specifies whether automatic catalog synchronization is enabled or disabled for snapshot backups. When enabled, NMDA deletes the associated NMDA DB2 metadata backups and the corresponding entries in the DB2 recovery history catalog whenever snapshot backups are expired and deleted from the NetWorker indexes.	<ul style="list-style-type: none"> undetermined (default). enabled = Snapshot backup with catalog synchronization (NMDA DB2 metadata and DB2 history pruning). disabled = Snapshot backup without catalog synchronization. <p>Note: Snapshot backups fail if this parameter value is not set to enabled or disabled.</p>
NSR_DB2CAT_SKIP_HISTORY_PRUNE	Optional. Specifies whether to skip the DB2 history pruning operation during a snapshot deletion or expiration when the snapshot is enabled for catalog synchronization. Note: The DB2 history pruning operation removes all the old backups from the DB2 recovery history, including both snapshot and non-snapshot backups and disk backups.	<ul style="list-style-type: none"> FALSE (default) = Do not skip the DB2 history pruning operation during a snapshot deletion or expiration. TRUE = Skip the DB2 history pruning operation during a snapshot deletion or expiration.
NSR_DEBUG_LEVEL	Optional. NMDA Parameters and Configuration File on page 399 provides details.	<ul style="list-style-type: none"> 0 (default) = Do not generate debug messages. 1 to 9 = Write debug messages to the debug log file (name has <code>.log</code> 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 the associated NMDA DB2 metadata backups and corresponding backup entries in the DB2 recovery history catalog are not successfully removed.	<ul style="list-style-type: none"> FALSE (default) = Remove expired NetWorker index entries only if the associated NMDA DB2 metadata backups and corresponding entries in

Table 35 DB2 resource file parameters (continued)

DB2 resource file parameter	Definition	Default and valid values
		<p>the DB2 history are successfully removed.</p> <ul style="list-style-type: none"> • TRUE = Remove all expired NetWorker index entries, even if the associated NMDA DB2 metadata backups and corresponding entries in the DB2 history are not successfully removed.

Automatic deletion of DB2 catalog entries with nsrdb2cat

For DB2 catalog synchronization, NSM runs the `nsrdb2cat` program to automatically delete the metadata save sets corresponding to an expired snapshot backup and then prune the DB2 recovery history.

The `nsrdb2cat` program runs on the DB2 server host that NMDA and NSM back up:

- The `nsrsnapck` program automatically runs the `nsrdb2cat` program.
- Only one `nsrdb2cat` program can run at a time.
- You cannot run the `nsrdb2cat` program manually.
- The default debug log is `/nsr/apps/logs/libnsrdb2_acs_cat_date_pid.log`.

Based on the snapshot save sets information that is passed from `nsrsnapck`, the `nsrdb2cat` program retrieves the corresponding metadata save sets, from which the DB2 timestamp information can also be obtained. The `nsrdb2cat` program deletes the metadata save sets and then prunes the DB2 recovery history of the database.

After deleting the metadata save sets, the `nsrdb2cat` program constructs a script to run the `db2 history prune` command with the DB2 timestamp of the expired snapshot backup. The `nsrdb2.res` file must provide the location of the `db2` binary used in the script.

The pathname of the `db2 history prune` script is `/nsr/apps/tmp/.nsrdb2pc/DB2_timestamp`. After a successful run, the script is automatically deleted.

Note: If the catalog synchronization is not enabled or the synchronous pruning of the catalog entry fails, then you must delete the leftover metadata save sets manually by running the `db2acsutil delete` utility and the `db2 prune history` command if required. NSM DB2 does not support DB2 DPF configurations.

NSM backups and restores on cluster systems

NMDA can perform NSM backups and restores of a database that is configured on an active-passive cluster system, or on active-active cluster systems only in the combinations that are listed in the *NetWorker E-LAB Navigator*.

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 [Cluster and High-Availability \(HA\) Systems](#) on page 297.

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.

About this task

You must complete the required steps to configure an NSM backup from a virtual cluster client.

Procedure

1. Install NMDA and the NetWorker extended client software 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 manually for NSM backups](#) on page 354:
 - 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 DB2 user from each physical client that can store and retrieve backups.

Oracle considerations for NSM 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. Due to NSM limitations, NMDA does not support raw devices with Oracle NSM operations on 64-bit Linux.

NMDA supports proxy backups and restores of archived redo logs. You must perform the archived logs backup separately from the database backup, for example, by using the `backup archivelog all` command.

Note: Do not use the `backup proxy database plus archivelog` command for backups with NSM.

Oracle does not support proxy backups of datafiles or archived redo logs that reside on Oracle Automated Storage, also known as Oracle Automated Storage Management (ASM).

Perform the Oracle configuration procedures according to the following topics:

- [Configuring the required Oracle settings](#) on page 369
- [Configuring the NWORA resource file](#) on page 370
- [Creating RMAN scripts for NSM snapshot backups](#) on page 370

Review the following information about performing Oracle backups and restores with NSM:

- [Checking configuration consistency](#) on page 374
- [Performing Oracle NSM backups](#) on page 375
- [Verifying Oracle NSM backup information in NetWorker indexes](#) on page 376
- [Restoring an Oracle NSM backup](#) on page 379

[NSM backups and restores on cluster systems](#) on page 394 describes Oracle operations with NSM in a cluster environment.

Configuring the required Oracle settings

About this task

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.

Settings for OCFS2 system with remote data mover

If you use an OCFS2 system and a remote data mover for NSM backups, ensure that you meet the following requirements:

- You allocate extra slots when you create an OCFS2 volume, as required to mount a device on a host outside of the RAC environment. Oracle recommends that you allocate two more slots than the number of nodes that will mount a device. For example, allocate four node slots for a two-node cluster, allocate eight slots for a six-node cluster, and so on.
- On the data mover host:
 - Ensure that the OCFS2 software is installed and running.
 - Ensure that unique OCFS2 node numbering is correctly set up.

No special configuration is required if a local data mover is used (local to one of the RAC nodes). Oracle provides support notes and documentation on how to prepare to use cloned OCFS volumes.

Windows settings for Oracle database with much read/write activity

If the Oracle database will probably have much 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 `initoracle_sid.ora`), increase the value of `LARGE_POOL_SIZE` to a large value that is appropriate for the particular system.

File types that are 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.

- ① **Note:** Due to Oracle limitations, you cannot perform NSM snapshot backups of Oracle datafiles or archived redo logs on ASM.

Control file versus recovery catalog

- ① **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.

i **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 an entry in the control file for each file that is 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

About this task

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 383. 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 382. 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

About this task

The information on RMAN backup scripts in [RMAN scripts for manual backups](#) on page 135 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 372.
- You must specify the `proxy` or `proxy only` option with each RMAN backup command.
 - i** **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.
 - i** **Note:** The NSM parameter NSR_PS_SAVE_PARALLELISM defines the NSM backup parallelism. [Table 32](#) on page 358 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 372 provides details.

Multiple channels in RMAN scripts

The allocation of multiple channels in an RMAN script does not control the degree of parallelism for snapshot operations. Oracle uses only one of the allocated channels for the backup or restore, unless you use specific backup options.

Example 39 RMAN scripts with multiple channels

Although two channels are allocated in the following RMAN script, 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';
  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 the Oracle software uses only one of the channels for the proxy backup:

```
run {
  allocate channel c1 type 'SBT_TAPE';
  allocate channel c2 type 'SBT_TAPE';
  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:

Example 39 RMAN scripts with multiple channels (continued)

```
run {
  allocate channel c1 type 'SBT_TAPE';
  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 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:

About this task

- The NMDA parameters that are described in [NMDA Parameters and Configuration File](#) on page 399.
You must set the parameters by using one of the methods in [Setting the NMDA parameters](#) on page 400.
- The NSM parameters that are described in [NSM parameter settings](#) on page 372.

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 476 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](#).

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 37](#), [Example 39](#), [Example 40](#). [Table 32](#) on page 358 provides a list of supported NSM parameters.

Example 40 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.

[Table 32](#) on page 358 provides details about 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_SNAPSHOT_POSTCMD=/space/myorcl/rman/snap.ksh) ';

  backup proxy database;
  release channel t1;
}
```

Example 41 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 that is 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) ';
```

Example 41 Setting the NSM parameter `RESTORE_TYPE_ORDER` (continued)

[Configuring NSM snapshot restore and recovery](#) on page 355 provides details.

Checking configuration consistency

During a scheduled backup, NMDA checks for consistency between the NetWorker data protection policy configuration and the RMAN backup session.

If NMDA finds a discrepancy between the data protection policy configuration and the RMAN session, warning messages appear or the backup fails as described in the following topics:

- Scheduled backups that are configured for snapshot backups
- Scheduled backups that are configured for traditional (nonsnapshot) backups

Scheduled backups that are configured for snapshot backups

Even when you have created a snapshot action for a scheduled NSM backup, RMAN might still perform 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 *NetWorker E-LAB Navigator* at <https://elabnavigator.emc.com/eln/elhome> provides details about the primary storage devices that are 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.

Scheduled backups that are configured for traditional (nonsnapshot) backups

When you have created a traditional backup action for a scheduled NSM backup, you cannot use the `proxy` option or `proxy only` option with a `backup` command in the RMAN script.

Any NSM backup that is 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.

Note:

- 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 42 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 through automatic or manual invocation of the scheduled backup workflow.

About this task

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.

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 383.

The NWORA resource file backup occurs at the backup level that is specified in configured backup action, 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 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 378 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 *NetWorker Administration Guide* provides details.

Note: The backup stores the file under the "backup" namespace, not the "oracle" namespace.

The retention policy that is applied to the NWORA resource file backup is the most conservative policy of the specified NetWorker client, not the policy that is applied to the Oracle backups. As a result, you may see a difference between the policies that are 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 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

About this task

The NetWorker server maintains information about each backup in its online indexes. [Terminology that is used in this guide](#) on page 23 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 *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 the following example.

- The second entry is for the point-in-time metadata, such as /brcmeta.1/ in the following example. 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 the following example. Only a backup to secondary storage creates this entry.

Example 43 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_0901JB_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 that is 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 that is used for the NSM backup, such as /space1/home/oracle/bp1 in the following example.

Example 44 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
```

Example 44 Oracle NSM backup entries in the media database (continued)

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 45 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 bu-bluegill

scanning client `bu-bluegill' for all savetimes from the backup
namespace
/nsr/apps/res/nwora.res, date=1396387310 Tue 01 Apr 2014 05:21:50
PM EDT
/nsr/apps/res/, date=1396387310 Tue 01 Apr 2014 05:21:50 PM EDT
/nsr/apps/, date=1396387310 Tue 01 Apr 2014 05:21:50 PM EDT
/nsr/, date=1396387310 Tue 01 Apr 2014 05:21:50 PM EDT
/, date=1396387310 Tue 01 Apr 2014 05:21:50 PM EDT
5 objects found
```

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 46 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 bu-bluegill

volume                                type            client          date
bu_jazz.lss.emc.com.001               adv_file       bu-bluegill    04/01/2014

time      size  ssid      fl  lvl  name
05:21:50 PM 3 KB  4164628438 cb  full /nsr/apps/res/
nwora.res
```

The *NetWorker Command Reference Guide* and the UNIX man pages describe these NetWorker commands.

Restoring an Oracle NSM backup

Review the information in the following topics about restores of an Oracle NSM backup:

About this task

- Creating RMAN scripts for Oracle NSM restores
- Performing Oracle NSM restores
- Relocating files during an Oracle NSM restore
- Catalog synchronization for Oracle NSM backups

[NSM backups and restores on cluster systems](#) on page 394 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.

About this task

i **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 [Data Restore and Recovery](#) on page 189.

To perform an NSM restore, you must set the appropriate parameters as described in [Setting the parameters for Oracle NSM operations](#) on page 372.

Performing Oracle NSM restores

About this task

i **NOTICE** Ensure that you meet the common requirements for database object access permissions in an NSM snapshot restore. The numeric user ID (UID) and group ID (GID) of the target database/instance owner must match the original UID and GID that were captured during the Oracle NSM backup.

The following requirements apply to Oracle NSM restores:

- You must install the NetWorker extended client software 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 380.
- 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 381.
- 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 32](#) on page 358 provides details.

Directory that is 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 that are excluded from the safety check.

Note: [Configuring NSM snapshot restore and recovery](#) on page 355 provides details about using the `psrollback.res` file with ProtectPoint rollback restores.

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.

Oracle rollback restore to a new database might fail when OMF is enabled

When the Oracle-Managed Files (OMF) database feature is enabled, a rollback restore to a new database might fail.

For example, when you perform a redirected rollback restore to alternate LUNs by using a ProtectPoint for VMAX backup of an Oracle OMF database, the restore might fail with the following error message:

```
ORA-19511: non RMAN, but media manager or vendor specific failure, error text:
A rollback is not possible when doing relocation during a restore.
Please remove 'rollback' from the RESTORE_TYPE_ORDER parameter or do not
request relocation. (114:123:2)
```

As a workaround, disable the OMF feature after you restore the spfile of the database and before you restore the control file and data files.

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 47 Symbolic link that is 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 that is specified in the `set newname` command.
- If the `set newname` command specifies a symbolic link, the base file name in the symbolic link.

Example 48 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 that is named `/tmp/rvol1`, which points to `/dev/volume_one/rvol2`.

Example 48 Relocation of a raw volume (continued)

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 that are 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 storage array does not have enough snapshot resources, 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 that are 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 in this case, 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 that are described in [Extra entries in the catalogs](#) on page 382. 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 444 provides details.
- Ensure that an `NWORA` resource file includes the required resources as described in [NWORA resource file](#) on page 383.

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 383
- [Automatic catalog synchronization with the nsroraclecat program](#) on page 391

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 388 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 384
- [NWORA SID resources](#) on page 386

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, NSR_ORACLECAT_DEBUG_FILE, NSR_ORACLECAT_LOG_FILE, NSR_ORACLECAT_MODE, NSR_REMOVE_ON_FAILURE.

The following table provides details about the required parameters.

Note: The parameter resources that are listed in the following table are the only ones supported. Do not try to add other parameter resources to the NWORA resource file.

Table 36 NWORA parameter resources

Parameter resource	Description	Default and valid values
NSR_NWPATH	<p>Specifies the directory location of the NetWorker binary <code>nsrsnapck</code>.</p> <p>Note: If you use NMDA with Sun-branded NetWorker, you must set NSR_NWPATH by using the following <code>nsroraadmin</code> command:</p> <pre>nsroraadmin -r update NSR_NWPATH=/usr/sbin/nsr</pre>	<ul style="list-style-type: none"> Directory pathname for the location of <code>nsrsnapck</code> (default). Valid directory pathname for the location of the NetWorker binary <code>nsrsnapck</code>.
NSR_ORACLECAT_DEBUG_FILE	<p>Specifies the debug file that is used by the <code>nsroraclecat</code> program. Set this parameter only for debugging the <code>nsroraclecat</code> program.</p> <p>Note: The <code>nsroraclecat</code> debug file must be in a secure location because the file includes a copy of the strings from the RMAN connection file.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the <code>nsroraclecat</code> debug file. <p>Note: If undefined, the <code>nsroraclecat</code> program does not generate debug information.</p>
NSR_ORACLECAT_LOG_FILE	<p>Specifies the operations log file that is used by the <code>nsroraclecat</code> program. The logged information includes the backup pieces that are successfully removed from the RMAN catalog, and those that failed to be removed during automatic catalog synchronization.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the <code>nsroraclecat</code> log file. <p>Note: If undefined, the <code>nsroraclecat</code> program writes the logging information to the <code>/nsr/applogs/nsroraclecat.log</code> file by default.</p>
NSR_ORACLECAT_MODE	<p>Specifies whether automatic catalog synchronization is enabled or disabled during NSM snapshot backups.</p>	<ul style="list-style-type: none"> Undetermined (default). Enabled. Disabled. <p>Note: Snapshot backups fail if the resource value is not enabled or disabled.</p>

Table 36 NWORA parameter resources (continued)

Parameter resource	Description	Default and valid values
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 77 provides details.	<ul style="list-style-type: none"> Undefined (default). Valid locale value, same as set in the NLS_LANG environment variable. <p>Note: Catalog synchronization fails if the value is not the same as the NLS_LANG variable value in a non-English environment.</p>
NSR_REMOVE_ON_FAILURE	Specifies whether the corresponding NetWorker index entries are removed when the <code>nsroraclecat</code> program fails to remove one or more RMAN catalog entries during automatic catalog synchronization. Automatic catalog synchronization with the nsroraclecat program on page 391 provides details.	<ul style="list-style-type: none"> 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 36](#) on page 384.

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 141 and [Configuring policy uniformity for scheduled Oracle backups](#) on page 142 provide details.

Depending on the command options that are 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 388 describes how to use the `nsroraadmin` command.

Example 49 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=
```

Example 49 Default NWORA parameter resources (continued)

```
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 386.

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 that are 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 391 provides details.

Note: Each NWORA SID resource must have a unique `NSR_ORACLE_SID` value.

Table 37 NWORA SID resource components

Parameter	Description	Default and valid values
<code>NSR_ORACLE_CONNECT_FILE</code>	Mandatory. Specifies the location of the file containing the connection strings that are required to create an RMAN session. The connection file on page 387 provides details.	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the RMAN connection file.
<code>NSR_ORACLE_HOME</code>	Mandatory. Specifies the home directory of the Oracle installation. The RMAN executable must be in the subdirectory <code>bin</code> in this directory.	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the Oracle home directory. <p>Note: The value must equal the values of the Oracle parameter <code>\$ORACLE_HOME</code>.</p>

Table 37 NWORA SID resource components (continued)

Parameter	Description	Default and valid values
NSR_ORACLE_LIB_PATH	Optional. Specifies the pathname of the directory containing the Oracle shared libraries on UNIX, typically <code>\$ORACLE_HOME/lib</code> .	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the Oracle shared library directory on UNIX. <p>Note: Do not set the parameter on Windows systems.</p>
NSR_ORACLE_SID	Mandatory. Specifies the SID value of the Oracle database whose RMAN catalog is to be synchronized.	<ul style="list-style-type: none"> Undefined (default). Valid SID value of the Oracle database. <p>Note: The value must be equal to the ORACLE_SID value in the NMDA configuration file that is used for the database backup. NMDA Oracle parameters on page 444 provides details.</p>
NSR_ORACLE_TNS_ADMIN	Optional. Specifies the pathname of the directory containing the Oracle Net configuration files.	<ul style="list-style-type: none"> Undefined (default). Valid pathname of Oracle network configuration directory. <p>Note: The value must be equal to the Oracle parameter <code>\$TNS_ADMIN</code> value.</p>

Using the nsroraadmin command to set SID resources

About this task

To add an NWORA SID resource to the resource file, use the `nsroraadmin -r add` command.

[Example 50](#) 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 388 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 386 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 that is required to connect to the target database.
- If you use an RMAN recovery catalog, the connection string that is 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 50 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.

About this task

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 388 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:

Procedure

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
- Adding the NWORA resources
- Updating the NWORA resources
- Deleting the NWORA SID resources

The following sources describe the `nsroraadmin` command:

- `nsroraadmin` man page on a UNIX Oracle Server that contains the NMDA software
- `nsroraadmin` entry in the *NetWorker Module for Databases and Applications Command Reference Guide*

Listing the NWORA resources

About this task

- 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

About this task

- 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 exists in the resource file, using the add keyword updates the resource value.

- To add an NWORA SID resource with the NSR_ORACLE_SID value of proddb and other specific values, use the information in the following example.

Example 51 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 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
```

Example 51 Adding an NWORA SID resource (continued)

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

About this task

- 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`:

About this task

```
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 187. The Oracle documentation describes the RMAN commands.

Review the information about automatic catalog synchronization in the following topics:

- RMAN catalog entry removals with `nsroraclecat`
- Failure of the `nsroraclecat` program

- NetWorker index entry removals with `nsrsnapck`
- Catalog synchronization after NSM Oracle backup volume is relabeled manually

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\nsracat_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 that are removed from the RMAN catalog. The program writes the information to the `nsroraclecat` log and debug files. The `NSR_ORACLECAT_LOG_FILE` and `NSR_ORACLECAT_DEBUG_FILE` parameter information in [Table 36](#) on page 384 describes 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 *NetWorker Module for Databases and Applications Command Reference Guide* at the Support website.

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 that are 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 that are described in [NetWorker index entry removals with `nsrsnapck`](#) on page 393.

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 that is 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 relabeled 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.

The following example provides sample commands for manual synchronization.

The Oracle documentation describes the RMAN `crosscheck` command.

Example 52 Using RMAN commands to synchronize the RMAN catalog entries

The following example shows the RMAN commands that are used to synchronize the RMAN catalog entries after you manually relabeled 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 that is configured on an active-passive cluster system, or on active-active cluster systems only in the combinations that are listed in the *NetWorker E-LAB Navigator*. 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 [Cluster and High-Availability \(HA\) Systems](#) on page 297.

Review the information about NSM operations on a cluster system in the following topics:

- NSM backup failover
- Configuring NSM backups from a virtual cluster client
- Configuring NSM backups from a physical cluster client
- Configuring restores from NSM backups on a cluster system

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 Retries value in the backup action resource to a nonzero value.

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 48 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.

About this task

You must complete the required steps to configure an NSM backup from a virtual cluster client.

Procedure

1. Install NMDA and the NetWorker extended client software 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 manually for NSM backups](#) on page 354:
 - 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 that are required for NSM backups as described in [Oracle considerations for NSM operations](#) on page 368:
 - To enable backup failover, specify a nonzero value in the Retries setting in the backup action that is created for the data protection policy of 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 *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 370.
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 that are created with the `-c client_name` option setting is different than without the setting.

At the start of a snapshot backup, if the storage array does not have enough snapshot resources, `nsrsnapck` might expire the oldest snapshot backup based on the **Minimum Retention Time** setting in the snapshot action. The `nsrsnapck` program also removes the NetWorker index entries of the backup based on the configuration. The `nsrsnapck` program performs the automatic expiration that is 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`.

Note:

- The host that is specified by `client_name` must have access to snapshot backups.
- NMDA and the NetWorker extended client software must be installed and configured on the host that is 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 53 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.

About this task

The following sources describe how to set up an NSM backup from a physical cluster client:

- [Configuring the required Oracle settings](#) on page 369
- [Checking configuration consistency](#) on page 374
- [Performing Oracle NSM backups](#) on page 375
- *NetWorker Administration Guide* (information 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 must 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 that are created with the `-c client_name` option setting is different than without the setting.

At the start of a snapshot backup, if the storage array does not have enough snapshot resources, `nsr_snapck` might expire the oldest snapshot backup based on the **Minimum Retention Time** setting in the snapshot action. The `nsr_snapck` program also removes the NetWorker index entries of the backup based on the configuration. The `nsr_snapck` 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`.

Note:

- The host that is specified by `client_name` must have access to snapshot backups.
- NMDA and the NetWorker extended client software must be installed and configured on the host that is specified by `client_name`.

Example 54 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.

About this task

Procedure

1. Set the parameter `NSR_CLIENT` to the correct value by using one of the methods in [Setting the NMDA parameters](#) on page 400:
 - 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.

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 346, 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.

For configuration details, refer to the manual of the corresponding replication management software.

APPENDIX A

NMDA Parameters and Configuration File

This appendix includes the following topics:

- [Setting the NMDA parameters](#)..... 400
- [NMDA configuration file](#)..... 400
- [Common NMDA parameters](#).....406
- [NMDA DB2 parameters](#).....414
- [NMDA Informix parameters](#).....421
- [NMDA Lotus parameters](#)..... 424
- [NMDA MySQL parameters](#)..... 435
- [NMDA Oracle parameters](#).....444
- [NMDA Orchestrated Application Protection parameters](#)..... 451
- [NMDA SAP IQ parameters](#).....457
- [NMDA Sybase parameters](#)..... 464

Setting the NMDA parameters

You must complete the required parameter settings for the NMDA backup and restore operations.

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 400 describes the NMDA configuration file and the exceptions to setting the parameters in the file.

[Common NMDA parameters](#) on page 406 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 applications only:

- [NMDA DB2 parameters](#) on page 414
- [NMDA Informix parameters](#) on page 421
- [NMDA Lotus parameters](#) on page 424
- [NMDA MySQL parameters](#) on page 435
- [NMDA Oracle parameters](#) on page 444
- [NMDA Orchestrated Application Protection parameters](#) on page 451
- [NMDA SAP IQ parameters](#) on page 457
- [NMDA Sybase parameters](#) on page 464

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 an Informix manual backup or restore, set parameters in the environment.
- For an Oracle backup or restore, set certain parameters in the RMAN script as described in [NMDA Oracle parameters](#) on page 444.
- For an SAP IQ backup or restore, set certain parameters in the environment as described in [NMDA SAP IQ parameters](#) on page 457.
- For a Sybase logtail only backup, as used in an up-to-the-minute recovery, run the `nsrsybrc -m logtail_bk` command as described in the *NetWorker Module for Databases and Applications Command Reference Guide*.

Use one of the following two methods to create the NMDA configuration file, depending on whether or not the configuration file will be used for database operations with the Orchestrated Application Protection feature:

- [Creating a configuration file that is not used with Orchestrated Application Protection](#) on page 401
- [Creating a configuration file that is used with Orchestrated Application Protection](#) on page 401

[NMDA features specific to Orchestrated Application Protection](#) on page 66 provides more information about the feature.

Creating a configuration file that is not used with Orchestrated Application Protection

If you will not use the Orchestrated Application Protection feature for backups and restores, you can use the following templates to create the NMDA configuration file. The NMDA software provides these configuration file templates:

- `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_iq_backup.cfg`—Template for SAP IQ backup parameters
- `nmda_iq_restore.cfg`—Template for SAP IQ restore parameters
- `nmda_sybase_backup.cfg`—Template for Sybase backup parameters
- `nmda_sybase_restore.cfg`—Template for Sybase restore parameters

The configuration file templates are located 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 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 based on the templates:

1. Copy the appropriate template file to any suitable location on the client host.
2. Customize the parameter settings in the file.
3. Ensure that the configuration file has read permissions for the group and other users.

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 writable by the administrative or application user that performs the operation.

[NMDA configuration file syntax](#) on page 402 provides details on the proper syntax to use in the NMDA configuration file.

Creating a configuration file that is used with Orchestrated Application Protection

If you will use the Orchestrated Application Protection feature for backups and restores, you can use the configuration file template to create the NMDA configuration file. The NMDA software provides the template file `/nsr/apps/config/nmda_oapp.cfg` on Linux for the Orchestrated Application Protection backup and restore parameters.

Make a copy of the template, for example, in the original directory or an alternate location. The uninstall of NMDA software removes the original template.

To create the NMDA configuration file based on the template:

1. Copy the template file to any suitable location on the client host.
2. Customize the parameter settings in the file.
3. Ensure that the configuration file has read permissions for the group and other users.

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. For example, for PostgreSQL WAL backups, you must specify the option in the `archive_command` setting in the `postgresql.conf` file, as described in [Registering the PostgreSQL archive command](#) on page 163.

i **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 writable by the application user that performs the operation.

[NMDA configuration file syntax](#) on page 402 provides details on the proper syntax to use in the NMDA configuration file.

NMDA configuration file syntax

You must use the correct syntax in the NMDA configuration file, depending on whether or not the configuration file will be used for database operations with the Orchestrated Application Protection feature.

NMDA 18.1 introduced support for Orchestrated Application Protection as described in [NMDA features specific to Orchestrated Application Protection](#) on page 66. Follow the syntax guidelines in the appropriate subtopic:

- [Configuration file syntax without Orchestrated Application Protection](#) on page 402
- [Configuration file syntax with Orchestrated Application Protection](#) on page 404

Configuration file syntax without Orchestrated Application Protection

When the NMDA configuration file will not be used with the Orchestrated Application Protection feature, ensure that the parameter settings in the file 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.
- To include any blank spaces at the start or end of a parameter value, enclose the parameter value (including the spaces) within double quote marks.
- 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
- IQ -- represents SAP IQ
- 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 128 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:, IQ:, 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, IQ, 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 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.

Configuration file syntax with Orchestrated Application Protection

When the NMDA configuration file will be used with the Orchestrated Application Protection feature, ensure that the parameter settings in the file use the required XML format tags. The NMDA configuration file must conform to the following syntax rules:

- Each parameter setting must appear on a separate line in the file.
- Parameter names and values are case-sensitive, unless specified otherwise in this appendix.
- The configuration file must begin and end with the following XML tag lines:

```
<?xml version="1.0" encoding="UTF-8"?>
<OAPP>
:
</OAPP>
```

- In each parameter setting, the parameter value must be preceded by an opening XML tag `<parameter_name>` and must be followed by a closing XML tag `</parameter_name>`. For example, the `NSR_DATABASE_TYPE` parameter setting must appear as follows, where PostgreSQL is the parameter value:

```
<NSR_DATABASE_TYPE>PostgreSQL</NSR_DATABASE_TYPE>
```

Note:

Any spaces are ignored between:

- The opening XML tag and the parameter value.
- The parameter value and the closing XML tag.

To include any blank spaces at the start or end of a parameter value, enclose the parameter value (including the spaces) within double quote marks.

- Parameter settings that apply only to backups must appear within the backup section, which begins and ends with the following XML tag lines:

```
<BACKUP>
:
</BACKUP>
```

- Parameter settings that apply only to full, incremental, or transaction log backups must appear within the corresponding subsection of the backup section. The subsection is enclosed by the `<FULL>`, `<INCR>`, or `<TXNLOG>` tags, respectively.

Parameter settings that apply to all the backup levels must appear outside of these backup level subsections.

In the following example, separate `NSR_BACKUP_SCRIPT` settings apply to each backup level, `NSR_LOG_VOLUME_POOL` applies only to transaction log backups, and the `NSR_DEBUG_LEVEL` setting applies to all backup levels:

```
<BACKUP>
  <FULL>
    <NSR_BACKUP_SCRIPT>/full_backup.sh</NSR_BACKUP_SCRIPT>
  </FULL>
  <INCR>
    <NSR_BACKUP_SCRIPT>/incr_backup.sh</NSR_BACKUP_SCRIPT>
  </INCR>
  <TXNLOG>
    <NSR_BACKUP_SCRIPT>/txnlog_backup.sh</NSR_BACKUP_SCRIPT>
    <NSR_LOG_VOLUME_POOL>log_backup_pool</NSR_LOG_VOLUME_POOL>
  </TXNLOG>
</BACKUP>
```

- Parameter settings that apply only to restores must appear within the restore section, which begins and ends with the following XML tag lines:

```
<RESTORE>
:
</RESTORE>
```

- Global parameter settings that apply to all the backup and restore operations must appear outside of the backup and restore sections.
- The following precedence rules apply to the parameter settings:
 - Global parameter settings apply to all operations, including all backups and restores.

- If a parameter is set more than once in the global section or in a backup or restore section, the last setting in the particular section takes precedence over all the previous settings of the parameter in that section.
- Use white space as preferred. NMDA ignores all the white space.

The following example shows an NMDA configuration for PostgreSQL full and transaction log backups with Orchestrated Application Protection. The global parameter settings above the <BACKUP> tag apply to all backups:

```
<?xml version="1.0" encoding="UTF-8"?>
<OAPP>
  <NSR_CLIENT>postgresql-client</NSR_CLIENT>
  <NSR_SERVER>postgresql-server</NSR_SERVER>
  <NSR_DATA_VOLUME_POOL>postgresql-backup-pool</NSR_DATA_VOLUME_POOL>
  <NSR_DATABASE_TYPE>PostgreSQL</NSR_DATABASE_TYPE>
  <NSR_INSTANCE_NAME>postgresql-64</NSR_INSTANCE_NAME>
  <NSR_BACKUP_NAME>postgresql-backup</NSR_BACKUP_NAME>
  <BACKUP>
    <FULL>
      <NSR_BACKUP_SCRIPT>/full_backup.sh</NSR_BACKUP_SCRIPT>
    </FULL>
    <TXNLOG>
      <NSR_BACKUP_SCRIPT>/txnlog_backup.sh</NSR_BACKUP_SCRIPT>
    </TXNLOG>
  </BACKUP>
</OAPP>
```

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 38 Common NMDA parameters

Parameter	Description	Default and valid values
DDBOOST_COMPRESSED_RESTORE	<p>Specifies whether to perform a compressed restore that uses the DD Boost workflow. A compressed restore uses less bandwidth by restoring the backed-up data in a compressed form from the Data Domain system to the application host.</p> <p>A compressed restore can be beneficial in a constrained bandwidth environment, but may impact the restore performance due to the usage of compression resources on the Data Domain system and application host.</p> <p>Optional for a restore through the command line or recovery wizard.</p> <p>① NOTICE SAP IQ restores and Sybase restores do not support this parameter setting.</p>	<ul style="list-style-type: none"> • FALSE (default) = Do not perform a compressed restore that uses the DD Boost workflow. • TRUE = Perform a compressed restore that uses the DD Boost workflow.

Table 38 Common NMDA parameters (continued)

Parameter	Description	Default and valid values
NSR_AES_ENCRYPTION	<p>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.</p> <p>Optional for a backup.</p> <p>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 that you used for an NMDA backup, you cannot recover the encrypted data. The <i>NetWorker Administration Guide</i> provides details on pass phrases.</p> <p>Note: Record each key (pass phrase) used for 256-bit AES encryption. You need the key to restore the backup later.</p> <p>This parameter is not supported with the Orchestrated Application Protection feature.</p>	<ul style="list-style-type: none"> 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	<p>Specifies whether the NetWorker software performs checksumming on the backup data.</p> <p>Optional for a backup.</p> <p>Note: When you restore an NMDA Sybase backup that was created with NSR_CHECKSUM set to TRUE, a <code>CRC mismatch</code> 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 <code>nmda_sybase.messages.raw</code> file and the Sybase Backup server error log.</p> <p>This parameter is not supported with the Orchestrated Application Protection feature.</p>	<ul style="list-style-type: none"> FALSE (default) = NetWorker software does not perform checksumming. TRUE = NetWorker software performs checksumming.
NSR_CLIENT	<p>Specifies the NetWorker Client resource to use for a backup or restore.</p>	<ul style="list-style-type: none"> Hostname of the physical host on which the session runs (default). Valid NetWorker client hostname.

Table 38 Common NMDA parameters (continued)

Parameter	Description	Default and valid values
	<p>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.</p> <p>Cluster and High-Availability (HA) Systems on page 297 provides details.</p> <p>For an Oracle client-side scheduled backup, set this parameter in the NMDA configuration file.</p>	
NSR_COMPRESSION	<p>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.</p> <p>Optional for a backup.</p> <p>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.</p> <p>This parameter is not supported with the Orchestrated Application Protection feature.</p>	<ul style="list-style-type: none"> FALSE (default) = NetWorker software does not perform compression. TRUE = NetWorker software performs compression.
NSR_DATA_DOMAIN_INTERFACE	<p>Specifies the network interface to use to send backup data to the DD Boost device.</p> <p>Optional for a manual backup only.</p> <p>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.</p> <p>Note: This parameter is not supported with the Orchestrated Application Protection feature.</p>	<ul style="list-style-type: none"> 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	<p>Specifies the name of the NetWorker volume pool to use for a manual backup.</p> <p>Mandatory for a DB2 or Oracle manual snapshot backup.</p>	<ul style="list-style-type: none"> Most appropriate pool as selected by the NetWorker server (default). <p>Note: The default pool is not necessarily the NetWorker pool</p>

Table 38 Common NMDA parameters (continued)

Parameter	Description	Default and valid values
	<p>Optional for a manual nonsnapshot backup with any supported application.</p> <p>Mandatory for an Oracle manual backup that uses the <code>set duplex</code> command (with <code>duplex set</code> 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 <code>parms</code> option in the RMAN script, not with the <code>send</code> command or <code>send</code> option.</p> <p>Note: This parameter overrides the volume pool that is specified in the NetWorker Client resource. This parameter is ignored by scheduled backups.</p> <p>If required, specify the associated storage node by setting the Storage Nodes attribute in the NetWorker Client resource (NMC diagnostic mode must be enabled).</p>	<p>that is named Default, although Default is often the default value of this parameter.</p> <ul style="list-style-type: none"> Valid name of a NetWorker volume pool. For a manual Oracle backup, the name must be different from the name that is used by the parameter <p>NSR_DATA_VOLUME_POOL1, NSR_DATA_VOLUME_POOL2, or NSR_DATA_VOLUME_POOL3.</p>
NSR_DEBUG_LEVEL	<p>Specifies the level of debug messages that NMDA writes to the debug log file, which is in the directory that is specified by <code>NSR_DIAGNOSTIC_DEST</code> or in the default directory, <code>/nsr/apps/logs</code> (UNIX) or <code>NetWorker_install_path\apps\logs</code> (Windows).</p> <p>Note: Use this parameter for debugging purposes with assistance from Customer Support only.</p> <p>Optional for a backup or restore.</p> <p>For an Oracle operation, set this parameter either in the configuration file or with the <code>parms</code> option in the RMAN script, not with the <code>send</code> command or <code>send</code> option.</p>	<ul style="list-style-type: none"> 0 (default) = NMDA does not generate debug messages. 1 to 9 = NMDA writes debug messages to the debug log file with a <code>.log</code> file name extension. The level of detail in the generated debug messages increases with the debug level. <p>For an SAP IQ log backup or restore, use a minimum value of 3 to enable the INFO level logging of the backed-up or restored log files.</p>
NSR_DEVICE_INTERFACE	<p>Note: This parameter is deprecated. It is supported in the current NMDA release, but will be unsupported in a future release.</p> <p>Specifies whether to always store the backup on the Data Domain device from the specified pool.</p> <p>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.</p>	<ul style="list-style-type: none"> Undefined (default) = Backup data is stored on the most appropriate device from the pool that is selected by NetWorker. DATA_DOMAIN = Backup data is always stored on the Data Domain device.

Table 38 Common NMDA parameters (continued)

Parameter	Description	Default and valid values
	<p>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.</p> <p>Note: This parameter is not supported with the Orchestrated Application Protection feature.</p>	
NSR_DIAGNOSTIC_DEST	<p>Specifies the directory location of the NMDA debug logs except the configuration wizard debug logs.</p> <p>Optional for a backup or restore.</p> <p>Note: You cannot relocate the operational and error logs (<code>nmda_*.messages.raw</code>) and the wizard debug logs.</p> <p>For an Oracle operation, set this parameter either in the configuration file or with the <code>parms</code> option in the RMAN script (not with the <code>send</code> command or <code>send</code> option).</p>	<ul style="list-style-type: none"> By default, NMDA generates the debug log files in the directory <code>/nsr/apps/logs</code> or <code>NetWorker_install_path\apps\logs</code>. Valid pathname of the directory where NMDA generates the debug logs except the wizard debug logs.
NSR_DIRECT_ACCESS	<p>Specifies the method used to perform a backup to a Data Domain device or an AFTD.</p> <p>Optional for a manual backup.</p> <p>If the target device is on a Data Domain system:</p> <ul style="list-style-type: none"> Data is deduplicated on the NMDA client when the Client Direct feature is enabled. Data is deduplicated on the storage node when the Client Direct feature is disabled. <p>This parameter is ignored during scheduled backups with NetWorker server release 8.0 or later and during restores.</p> <p>For scheduled backups, you can enable the Client Direct feature with the Client Direct setting in the wizard or the Client resource.</p> <p>Note: This parameter is not supported with the Orchestrated Application Protection feature.</p>	<ul style="list-style-type: none"> Default (default) = The backup tries to use the Client Direct feature to bypass the storage node and write data directly to the target device. If the backup cannot use this feature, then the backup uses the storage node. No = The backup does not try to use the Client Direct feature. The backup uses the storage node only. Yes = The backup tries to use the Client Direct feature. If the backup cannot use this feature, then the backup fails. <p>Note: The Yes value is deprecated. The value is supported in the current NMDA release, but will be unsupported in a future release.</p>

Table 38 Common NMDA parameters (continued)

Parameter	Description	Default and valid values
NSR_DPRINTF	<p>Specifies whether NetWorker core debug messages are written to the NMDA debug log files, as described for NSR_DEBUG_LEVEL.</p> <p>Note: NSR_DPRINTF applies only if NSR_DEBUG_LEVEL is set to a value greater than 0. Use this parameter for debugging purposes with assistance from Customer Support only.</p> <p>Optional for a backup or restore.</p> <p>For an Oracle operation, set this parameter either in the configuration file or with the <code>parms</code> option in the RMAN script, not with the <code>send</code> command or <code>send</code> option.</p>	<ul style="list-style-type: none"> 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	<p>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.</p> <p>Optional for a restore.</p> <p>If both of the following conditions are true, set this parameter to the phrase used to originally back up the data:</p> <ul style="list-style-type: none"> NMDA backed up the data with 256-bit AES encryption. The encryption phrase on the NetWorker server has changed since the backup. <p>For an Oracle restore, set this parameter with the <code>send</code> command in the RMAN script.</p> <p>Note: This parameter is not supported with the Orchestrated Application Protection feature.</p>	<ul style="list-style-type: none"> By default, an NMDA restore obtains the encryption phrase from the Datazone pass phrase attribute of the NetWorker Server resource. One or more encryption phrases to use during an NMDA restore. Each phrase must be a string that is enclosed in quotes. Separate multiple phrases with commas. <p>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:</p> <pre>NSR_ENCRYPTION_PHRASES="'key1','key2'"</pre> <ul style="list-style-type: none"> - NMDA itself supports double ("), single ('), and backward (') quotes. - Certain shells, databases, or applications might not support certain types of quotes.
NSR_MAX_START_RETRIES	<p>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.</p> <p>For example, the connection to the NetWorker server might fail for one of the following reasons:</p>	<ul style="list-style-type: none"> 4 (default). Integer number of tries to connect to the NetWorker server.

Table 38 Common NMDA parameters (continued)

Parameter	Description	Default and valid values
	<ul style="list-style-type: none"> The NetWorker server is not ready because the devices are not mounted. The <code>nsrindexd</code> service of the NetWorker server is busy due to other client sessions. <p>Optional for a backup or restore.</p>	
NSR_NWPATH	<p>Specifies the pathname of the directory that contains the NetWorker binaries. Mandatory for MySQL or Oracle backup deletions if the NetWorker client binaries are in a nondefault directory on the server host.</p> <p>Recommended for Oracle operations on Windows if multiple database instances exist on the same host.</p> <p>Recommended for an NSM snapshot backup or restore if either of the following conditions are true:</p> <ul style="list-style-type: none"> NSM (<code>nsrsnapck</code> binary) is in a nondefault location. NetWorker client software is in a nondefault location. <p>Note: You cannot use NSR_NWPATH for deduplication backups or restores with NSM.</p> <p>This parameter is not supported with the Orchestrated Application Protection feature.</p>	<ul style="list-style-type: none"> System-specific default location of the NetWorker client binaries (default). If NSM (<code>nsrsnapck</code> binary) is in a nondefault location, valid directory pathname of the NSM binary. If NetWorker client software is in a nondefault location, valid directory pathname of the NetWorker client binaries.
NSR_PHYSICAL_HOST_LICENSE	<p>Specifies whether NMDA is licensed per cluster node or per virtual cluster name in an active-passive cluster environment.</p> <p>Optional in an active-passive cluster.</p>	<ul style="list-style-type: none"> 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	<p>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 in different volume pools.</p> <p>Optional for a restore.</p> <p>Supported for the following types of restores:</p>	<ul style="list-style-type: none"> Pool that is determined by the NetWorker server (default). Valid name of a NetWorker volume pool that contains a backup or cloned backup to use for a restore.

Table 38 Common NMDA parameters (continued)

Parameter	Description	Default and valid values
	<ul style="list-style-type: none"> Restore of a nonsnapshot backup. Restore of a ProtectPoint snapshot backup on a Data Domain system. <p>Note: If required, specify the associated storage node by setting the Recover Storage Nodes attribute in the NetWorker Client resource (NMC diagnostic mode must be enabled).</p>	
NSR_SAVESET_RETENTION	<p>Specifies the retention policy of a manual backup, as the date when the save set becomes recyclable.</p> <p>Optional for a manual backup.</p> <p>Note: This parameter overrides the retention policy that is specified in the NetWorker Client resource. This parameter is ignored by scheduled backups.</p>	<ul style="list-style-type: none"> Retention policy that is specified in the NetWorker Client resource for the client (default). Valid date in <code>nsr_getdate(3)</code> format.
NSR_SERVER	<p>Specifies the hostname of the NetWorker server to perform the backup or restore.</p> <p>Mandatory for a manual backup or restore if the NetWorker server host is different from the client host.</p> <p>For an Oracle operation:</p> <ul style="list-style-type: none"> If the operation generates backup copies, set this parameter with the <code>parms</code> option in the RMAN script, not with the <code>send</code> command or <code>send</code> option. If the operation does not generate backup copies, set this parameter with the <code>send</code> command or <code>send</code> option in the RMAN script. 	<ul style="list-style-type: none"> Hostname of the NetWorker server that is detected on the client host (default). Valid hostname of a NetWorker server.
POSTCMD	<p>Specifies a postprocessing script to run after a scheduled backup:</p> <ul style="list-style-type: none"> 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: 	<ul style="list-style-type: none"> Undefined (default). Valid pathname of a postprocessing script file. The pathname must not contain any spaces. <p>For example, instead of setting POSTCMD to <code>C:\Program Files\Legato\nsr\postcmd.bat</code>, set the parameter to <code>C:\Progra~1\Legato\nsr\postcmd.bat</code>.</p>

Table 38 Common NMDA parameters (continued)

Parameter	Description	Default and valid values
	<p>#!/bin/sh</p> <p>Optional for a scheduled backup. Do not set this parameter for a manual backup.</p> <p>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.</p>	<p>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.</p> <ul style="list-style-type: none"> If the value is undefined or invalid, a postprocessing script does not run after the scheduled backup.
PRECMD	<p>Specifies a preprocessing script to run before a scheduled backup:</p> <ul style="list-style-type: none"> 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 <p>Optional for a scheduled backup. Do not set this parameter for a manual backup.</p> <p>Note: If the preprocessing script fails, NMDA does not perform the scheduled backup, an error message appears, and any postprocessing script does not run.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of a preprocessing script file. The pathname must not contain any spaces. <p>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.</p> <p>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.</p> <ul style="list-style-type: none"> If the value is undefined or invalid, a preprocessing script does not run before the scheduled backup.

NMDA DB2 parameters

You must complete the required parameter settings for the NMDA DB2 backup and restore operations.

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 that are configured without the wizard (client-side configuration)
- DB2 manual backups
- DB2 restores

[Common NMDA parameters](#) on page 406 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 122 provides details.

Table 39 NMDA DB2 parameters

Parameter	Description	Default and valid values
DB2_ACS_LAYOUT_CHECK	<p>Specifies whether to enforce the DB2 ACS best practice on the log directory layout of the database during a snapshot backup.</p> <p>The best practice requires a dedicated volume group for log paths, with the log paths contained in a different file system volume than the database directory and database containers.</p> <p>Optional for a DB2 snapshot backup.</p>	<ul style="list-style-type: none"> TRUE (default) = Enforce the DB2 ACS best practice on the log directory layout. Backups with the <code>exclude logs</code> option fail if the log paths are not in a separate file system volume. FALSE = Do not enforce the DB2 ACS best practice on the log directory layout.
DB2_ACS_METADATA_DELETION_FORCE	<p>Specifies whether to enforce the deletion of metadata save sets during a <code>db2acsutil</code> deletion operation when the snapshot save sets fail to be found during the deletion operation.</p> <p>Optional for a <code>db2acsutil</code> deletion operation.</p>	<ul style="list-style-type: none"> FALSE (default) = Do not force the deletion of the metadata save sets. TRUE = Force the deletion of the metadata save sets.
DB2_ALIAS	<p>Specifies the database alias of the DB2 database to back up. This alias is typically the same as the database name.</p> <p>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 <code>DB2:/DB_NAME/NODE.XXXX</code>, where <code>XXXX</code> is the partition number.</p> <p>Optional for a DB2 scheduled backup only.</p>	<ul style="list-style-type: none"> Undefined (default). Valid DB2 alias name.
DB2_APPLY_NW_LEVELS	<p>Specifies whether NMDA uses the backup levels that are specified in either the NMDA configuration file (for example, <code>nmda_db2.cfg</code>) or in the NetWorker Schedule resource.</p> <p>Optional for a DB2 scheduled backup only.</p>	<ul style="list-style-type: none"> FALSE (default) = NMDA uses the backup level set in the configuration file and maps the backup level to the NetWorker schedule as follows: <ul style="list-style-type: none"> DB2BACKUP_FULL maps to the full level in the schedule. DB2BACKUP_INCREMENTAL maps to the cumulative incr level in the schedule. DB2BACKUP_DELTA maps to the incr level in the schedule. TRUE = NMDA uses the backup level that is defined in the NetWorker schedule. NMDA ignores the backup

Table 39 NMDA DB2 parameters (continued)

Parameter	Description	Default and valid values
		<p>level in the configuration file. NMDA maps the levels in the NetWorker schedule to the DB2 levels as follows:</p> <ul style="list-style-type: none"> ■ The full level maps to DB2BACKUP_FULL. ■ The cumulative incr level maps to DB2BACKUP_INCREMENTAL. ■ The incr level maps to DB2BACKUP_DELTA.
DB2_BUFFER_SIZE	<p>Specifies the DB2 backup buffer size in page units. This parameter corresponds to the <code>buffer buffer_size</code> option of the DB2 <code>backup</code> command.</p> <p>Optional for a DB2 scheduled backup only.</p> <p>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.</p>	<ul style="list-style-type: none"> • Optimal value set by DB2 (default). • Number of page units for the DB2 backup buffer size. The IBM documentation provides recommended settings.
DB2INSTANCE	<p>For a DB2 backup, specifies the name (not the alias) of the DB2 instance that contains the database to be backed up.</p> <p>For a DB2 rollforward operation after a redirected recovery, specifies the name of the source DB2 instance that was used for the backup and contains the logs to be retrieved.</p> <p>Mandatory for a DB2 scheduled backup on UNIX only.</p> <p>Mandatory for a DB2 restore, recovery, and rollforward operation on all platforms. Set this parameter in the log configuration file that is specified by the DB2 database configuration parameter LOGARCHOPT1. For example, the following command specifies the file:</p> <pre style="background-color: #f0f0f0; padding: 5px;">db2 update db cfg for database using LOGARCHMETH1 vendor:c:\NetWorker\nsr\bin</pre>	<ul style="list-style-type: none"> • Undefined (default). • Valid name of the DB2 instance that contains the database, or the DB2 instance that contains the logs.

Table 39 NMDA DB2 parameters (continued)

Parameter	Description	Default and valid values
	<pre>\libnsrdb2.dll LOGARCHOPT1 @log_config_file</pre> <p>Ensure that this parameter is set correctly. The appropriate IBM DB2 documentation provides details.</p>	
DB2_NODE_NAME	<p>Specifies the alias of the DB2 instance to which the user must connect for the backup.</p> <p>Mandatory for a DB2 scheduled backup only.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid alias of the DB2 instance. If the node you are using is through a local connection, specify the instance name.
DB2_NUM_BUFFERS	<p>Specifies the number of backup buffers that are used by DB2. This parameter corresponds to the <code>with num_buffers</code> buffer option of the DB2 backup command.</p> <p>Optional for a DB2 scheduled backup only.</p> <p>DB2 automatically uses an optimal value for this parameter.</p>	<ul style="list-style-type: none"> • Optimal value set by DB2 (default). • Number of DB2 backup buffers. The IBM documentation provides more information.
DB2_OPTIONS	<p>Specifies the DB2 backup options. Optional for a DB2 scheduled backup only.</p> <p>Note: At a minimum, specify either <code>DB2BACKUP_DB</code> or <code>DB2BACKUP_TABLESPACE</code>, not both.</p>	<ul style="list-style-type: none"> • Undefined (default). • One or more of the following values (case-sensitive), with multiple values separated by commas: <ul style="list-style-type: none"> ▪ <code>DB2BACKUP_COMPRESS</code> ▪ <code>DB2BACKUP_DB</code> ▪ <code>DB2BACKUP_DEDUP_DEVICE</code> ▪ <code>DB2BACKUP_DELTA</code> ▪ <code>DB2BACKUP_EXCLUDE_LOGS</code> ▪ <code>DB2BACKUP_FULL</code> ▪ <code>DB2BACKUP_INCLUDE_LOGS</code> ▪ <code>DB2BACKUP_INCREMENTAL</code> ▪ <code>DB2BACKUP_OFFLINE</code> ▪ <code>DB2BACKUP_ONLINE</code> ▪ <code>DB2BACKUP_TABLESPACE</code>
DB2_PARALLELISM	<p>Specifies the number of tablespaces that can be read in parallel during a DB2 backup. This parameter corresponds to the <code>parallelism</code> option of the DB2 backup command.</p> <p>Optional for a DB2 scheduled backup only.</p>	<ul style="list-style-type: none"> • Optimal value set by DB2 (default). • Number of the maximum concurrent tablespaces read in parallel during a backup. The IBM documentation provides more information.

Table 39 NMDA DB2 parameters (continued)

Parameter	Description	Default and valid values
	DB2 automatically uses an optimal value for this parameter.	
DB2_PARTITION_LIST	Specifies which nodes to back up for a DPF backup. Optional for a DB2 DPF backup only.	<ul style="list-style-type: none"> 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.	<ul style="list-style-type: none"> Undefined (default). Valid path of the location of the DB2 binaries that are used for the backup.
DB2 QUIESCE	Specifies whether to quiesce the DB2 database during a backup. Optional for a DB2 scheduled backup only.	<ul style="list-style-type: none"> FALSE (default) = Do not quiesce the DB2 database during a backup. TRUE = Quiesce the DB2 database during a backup. <p>Note: You must also set the DB2_ALIAS parameter.</p>
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.	<ul style="list-style-type: none"> 1 (default). Integer number of parallel NMDA sessions.
DB2_SKIP_HADR_STANDBY	Specifies whether to skip the backup of a standby node database in an HADR system. Mandatory for a standby node configuration in an HADR system that does not use a Virtual IP configuration. Setting the TRUE value in such an HADR system prevents a possible error because DB2 does not support a standby database backup. Note: This parameter is ignored on an HADR system that uses a Virtual IP configuration and on a non-HADR database.	<ul style="list-style-type: none"> FALSE (default) = Do not skip the standby database backup. TRUE = Skip the standby database backup. Set this parameter value in an HADR system that does not use a Virtual IP configuration.
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.	<ul style="list-style-type: none"> Undefined (default). List of tablespace names, with multiple names separated by commas. For example: DB2_TBS_LIST = SYSCATSPACE, USERSPACE1

Table 39 NMDA DB2 parameters (continued)

Parameter	Description	Default and valid values
DB2_USER	<p>Specifies the name of the DB2 user that connects to the DB2 instance for the backup.</p> <p>You must specify the user password with the USER_PSWD parameter.</p> <p>Mandatory for a DB2 scheduled backup only.</p>	<ul style="list-style-type: none"> Undefined (default). Valid DB2 username.
DB2_VENDOR_LIB_PATH	<p>Specifies the complete pathname of the NMDA shared library on the DB2 host.</p> <p>The path can point to various library versions to test and evaluate hotfixes.</p> <p>Optional for a DB2 scheduled backup only.</p>	<ul style="list-style-type: none"> On UNIX systems, if you do not specify a path, NMDA assumes the default location. <p>For example, on Solaris systems: /usr/lib/libnsrdb2.so</p> <ul style="list-style-type: none"> On Windows systems, NMDA obtains the default path automatically from the registry.
INSTHOME (UNIX)	<p>Specifies the pathname of the DB2 instance home directory.</p> <p>Mandatory for a DB2 scheduled backup on UNIX only.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the DB2 instance home directory.
NSR_DR_BACKUP_INFO	<p>Specifies whether to back up additional disaster recovery information and support information along with a scheduled backup.</p> <p>If the backup of additional information fails, an error message appears but the backup report is successful.</p> <p>Optional for a DB2 scheduled backup only.</p>	<ul style="list-style-type: none"> TRUE (default) = Back up additional disaster recovery information and support information as detailed in Preparing for DB2 disaster recovery on page 275. FALSE = Do not back up additional disaster recovery information and support information. <p>Note: Set the following additional parameters:</p> <ul style="list-style-type: none"> DB2_ALIAS DB2PATH (Windows) INSTHOME (UNIX)
NSR_DR_FILE_LIST	<p>Specifies a file that contains a list of files to back up in addition to the database backup.</p> <p>NMDA backs up the files as part of the scheduled backup, before any postprocessing script defined by POSTCMD runs.</p> <p>If the backup of extra files fails, an error message appears.</p>	<ul style="list-style-type: none"> 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 <code>nmda_savelist.txt</code>, which is a file that contains this list: /space12/vendor.cfg

Table 39 NMDA DB2 parameters (continued)

Parameter	Description	Default and valid values
	Optional for a DB2 scheduled backup only.	/space12/db2inst1/sqlllib/ db2nodes.cfg
NSR_LOG_VOLUME_POOL	<p>Specifies the name of the NetWorker volume pool to use for a backup of the transaction logs.</p> <p>Mandatory for a DB2 manual snapshot backup.</p> <p>Optional for a DB2 manual nonsnapshot backup.</p> <p>Note: If required, specify the associated storage node by setting the Storage Nodes attribute in the NetWorker Client resource (NMC diagnostic mode must be enabled).</p>	<ul style="list-style-type: none"> Most appropriate pool as selected by the NetWorker server (default). Valid name of a NetWorker volume pool for the transaction logs.
NSR_PROXY_PFILE	<p>Specifies the complete pathname of a configuration file that contains the NSM parameter settings for a snapshot backup or restore.</p> <p>Mandatory for a manual backup or restore for which you must set NSM parameters.</p> <p>For a scheduled backup, you must set the NSM parameters in the Application Information attribute of the Client resource.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the configuration file. <p>Note: If undefined or an invalid pathname, the operation ignores the parameter settings in the preferred configuration file.</p>
NSR_XBSA_DEBUG	<p>Specifies whether debug messages from the NetWorker XBSA library are written to the NMDA DB2 log at the level set by NSR_DEBUG_LEVEL, described in NSR_DEBUG_LEVEL.</p> <p>Note: Use this parameter for debugging purposes with assistance from Customer Support only.</p> <p>Optional for a DB2 backup or restore.</p>	<ul style="list-style-type: none"> FALSE (default) = XBSA library debug messages are not written to the NMDA DB2 log. TRUE = XBSA library debug messages are written to the NMDA DB2 log.
SOURCE_CLIENT	<p>Specifies the source client hostname under which the archived logs backup was indexed.</p> <p>This parameter is used in a rollforward recovery of the destination database after a restore to a different client where a rollback is possible.</p> <p>Optional for a rollforward after a redirected restore.</p>	<ul style="list-style-type: none"> Undefined (default). If not specified, the NSR_CLIENT value is used. Valid hostname of the source client as used in the backup of the archived logs.

Table 39 NMDA DB2 parameters (continued)

Parameter	Description	Default and valid values
SOURCE_DBNAME	<p>Specifies the name of the source DB2 database that was originally used for the backup of the database or archived logs.</p> <p>Mandatory for a database restore to a different database and for a recovery and rollforward operation to a different database.</p>	<ul style="list-style-type: none"> Undefined (default). Valid name of the DB2 database that was originally used for the backup of the database or archived logs.
USER_PSWD	<p>Specifies the encrypted password for the DB2 user that connects to the DB2 instance, as specified by the DB2_USER parameter. The encrypted password is added as the USER_PSWD setting to the NMDA configuration file.</p> <p>Mandatory for a DB2 scheduled backup only.</p>	<ul style="list-style-type: none"> Undefined (default). Encrypted DB2 user password that you must set with the <code>nsrdaadmin -P</code> command, for example: <div style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;"> <pre>nsrdaadmin -P -z configuration_file_path</pre> </div> <p>The <i>NetWorker Module for Databases and Applications Command Reference Guide</i> describes the <code>nsrdaadmin</code> command.</p>

NMDA Informix parameters

You must complete the required parameter settings for the NMDA Informix backup and restore operations.

For Informix scheduled backups that are 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 406 describes the common parameters.

The following table describes the NMDA Informix parameters.

Table 40 NMDA Informix parameters

Parameter	Description	Default and valid values
DO_LOGFILE_BACKUPS	<p>Specifies whether to perform the logical log file backup after the dbspace backup.</p> <p>Optional for an Informix scheduled backup only.</p> <p>Note: If you set <code>DO_WHOLE_SYSTEM_BACKUP</code> to <code>TRUE</code>, NMDA ignores <code>DO_LOGFILE_BACKUPS</code>.</p>	<ul style="list-style-type: none"> <code>TRUE</code> (default) = Perform the logical log file backup after the dbspace backup. <code>FALSE</code> = Do not perform the logical log file backup after the dbspace backup. <p>Note: If you set <code>DO_LOGFILE_BACKUPS</code> to <code>TRUE</code>,</p>

Table 40 NMDA Informix parameters (continued)

Parameter	Description	Default and valid values
		<p>the <code>onbar -b -l -c</code> command runs after the backups of any Informix dbspaces.</p>
DO_WHOLE_SYSTEM_BACKUP	<p>Specifies whether to back up the logical log files during a scheduled backup.</p> <p>Optional for an Informix scheduled backup only.</p> <p>Note: If you set DO_WHOLE_SYSTEM_BACKUP to TRUE, NMDA ignores DO_LOGFILE_BACKUPS.</p>	<ul style="list-style-type: none"> TRUE (default) = Back up the logical log files during a scheduled backup. This action is equivalent to running <code>onbar</code> with the <code>-w</code> option. FALSE = Do not back up the logical log files during a scheduled backup.
INFORMIXDIR	<p>Specifies the directory pathname of the Informix RDBMS installation.</p> <p>Mandatory for an Informix scheduled backup only.</p>	<ul style="list-style-type: none"> Undefined (default). Valid directory pathname of the Informix RDBMS installation.
INFORMIXSQLHOSTS	<p>Specifies the name of the Informix SQL hosts file.</p> <p>Mandatory for an Informix scheduled backup on UNIX or (only with Informix 12.10 or later) on Windows.</p> <p>Optional for an Informix scheduled backup on Windows with Informix earlier than 12.10.</p>	<ul style="list-style-type: none"> Undefined (default). Valid name of the Informix SQL hosts file.
NSR_DR_BACKUP_INFO	<p>Specifies whether a scheduled backup backs up additional files as additional disaster recovery information and support information.</p> <p>If the backup of additional files fails, an error message appears but the backup report is successful.</p> <p>Optional for an Informix scheduled backup only.</p>	<ul style="list-style-type: none"> TRUE (default) = A scheduled backup backs up the following additional information: <ul style="list-style-type: none"> 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)

Table 40 NMDA Informix parameters (continued)

Parameter	Description	Default and valid values
		<ul style="list-style-type: none"> FALSE = A scheduled backup does not back up the additional information.
NSR_DR_FILE_LIST	<p>Specifies a file that contains a list of files to back up in addition to the database backup.</p> <p>NMDA backs up the files as part of the scheduled backup, before any postprocessing script defined by POSTCMD runs.</p> <p>If the backup of extra files fails, an error message appears but the backup report is successful.</p> <p>Optional for an Informix scheduled backup only.</p>	<ul style="list-style-type: none"> 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 <code>nmda_savelist.txt</code>, which is a file that contains a list of the extra file pathnames. <p>Note: Set NSR_DR_BACKUP_INFO to TRUE to enable this parameter setting.</p>
NSR_INACTIVITY_TIMEOUT	<p>Specifies the length of time in minutes before a timeout occurs, when the NetWorker server fails an inactive Informix manual backup.</p> <p>The preferred timeout value depends on the database size, the backup type, and the time that <code>onbar</code> needs to analyze the data. During the analysis, <code>onbar</code> does not provide any data to back up and the backup is inactive.</p> <p>The manual backup may fail if the inactive time is longer than the Inactivity Timeout setting on the NetWorker server. Increasing the timeout value can resolve this issue.</p> <p>Optional for an Informix manual backup only.</p> <p>Note: For an Informix scheduled backup, set this parameter in the Inactivity Timeout field by using the backup action wizard.</p>	<ul style="list-style-type: none"> 0 or 30 (default). Signifies no timeout or 30 minutes timeout, depending on the NetWorker server policy settings. Integer value of timeout in minutes.
NSR_LOG_VOLUME_POOL	<p>Specifies the volume pool to use for a backup of the logical logs.</p> <p>Optional for an Informix scheduled backup only.</p> <p>Note: If required, specify the associated storage node by setting the Storage Nodes attribute in the</p>	<ul style="list-style-type: none"> Most appropriate pool as selected by the NetWorker server (default). Valid name of a NetWorker volume pool for the logical logs.

Table 40 NMDA Informix parameters (continued)

Parameter	Description	Default and valid values
	NetWorker Client resource (NMC diagnostic mode must be enabled).	
ONCONFIG	Specifies the name of the Informix RDBMS configuration file. Mandatory for an Informix scheduled backup only.	<ul style="list-style-type: none"> Undefined (default). Valid name of the Informix RDBMS configuration file.

NMDA Lotus parameters

You must complete the required parameter settings for the NMDA Lotus backup and restore operations.

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 406 describes the common parameters.

The following table describes the NMDA Lotus parameters.

Table 41 NMDA Lotus parameters

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.	<ul style="list-style-type: none"> Undefined (default). Valid name of the Lotus Domino or Notes user.
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.	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the Lotus Domino or Notes directory that contains the <code>libnotes.xx</code> or <code>nnotes.dll</code> library file.
NSR_APPLY_LOGS	Specifies whether to apply the transaction logs after a Lotus backup is restored. Optional for a Lotus restore.	<ul style="list-style-type: none"> 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.	<ul style="list-style-type: none"> FALSE (default) = Lotus database or file restore occurs with user interaction.

Table 41 NMDA Lotus parameters (continued)

Parameter	Description	Default and valid values
	Optional for a Lotus restore.	<ul style="list-style-type: none"> TRUE = Lotus database or file restore occurs automatically without user interaction.
NSR_BACKUP_ALL_EXTENSIONS	<p>Specifies whether to back up all the Lotus files or only the default set of Lotus files with specific file name extensions.</p> <p>Optional for a Lotus backup.</p>	<ul style="list-style-type: none"> 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.	<ul style="list-style-type: none"> full (default) = Perform a full backup. incr = Perform an incremental backup. txnlog = Perform a backup of transaction logs only.
NSR_BACKUP_LOGS_MODE	<p>Specifies the level of transaction log backup to perform during a full backup.</p> <p>Optional for a Lotus full backup only. Ignored for an incremental backup or transaction log backup.</p> <p>Note: NMDA always backs up the transaction logs during a Lotus incremental backup unless you set NSR_INCR_BACKUP_LOGS_MODE=1.</p>	<ul style="list-style-type: none"> 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. <p>Note: 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 particular backup is guaranteed.</p>
NSR_BACKUP_LOTUS_DIR	<p>Specifies whether to back up files in the Lotus Domino or Notes data directory.</p> <ul style="list-style-type: none"> 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. <p>Optional for a Lotus backup.</p>	<ul style="list-style-type: none"> FALSE (default) = Back up the Lotus directories and files that are specified with NSR_BACKUP_PATHS. TRUE = Back up the Lotus data directory. <p>On Windows, the backup includes the notes.ini file, whether or not the file resides in the default data directory.</p> <p>Note: The NSR_BACKUP_ALL_EXTENSIONS</p>

Table 41 NMDA Lotus parameters (continued)

Parameter	Description	Default and valid values
	<p>Note: You cannot use this parameter with the NSR_BACKUP_PATHS parameter.</p>	<p>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.</p>
NSR_BACKUP_PATHS	<p>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 433 describes the use of wildcards in the pathnames.</p> <p>Optional for a Lotus backup or restore.</p> <p>Note: For a backup, you cannot use this parameter with the NSR_BACKUP_LOTUS_DIR parameter.</p> <p>For a restore, you cannot use this parameter with the NSR_RECOV_LIST_FILE parameter.</p>	<ul style="list-style-type: none"> 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 specified client. Do not use the keyword for a partitioned Domino server or multiple Domino installations on a single UNIX host.
NSR_BROWSELIST_CACHE_DEST	<p>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.</p> <p>Optional for a Lotus backup.</p>	<ul style="list-style-type: none"> By default, the browselist file is in this directory: <ul style="list-style-type: none"> - /nsr/apps/tmp (UNIX) - NetWorker_install_path\apps\tmp (Windows) Valid complete pathname of a directory on a writable device to contain the browselist file.

Table 41 NMDA Lotus parameters (continued)

Parameter	Description	Default and valid values
NSR_CATALOGFILE	<p>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.</p> <p>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.</p> <p>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.</p> <p>Optional for a Lotus backup.</p>	<ul style="list-style-type: none"> 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	<p>Specifies the comfort span value to use for an incremental backup. Lotus incremental backups with the comfort span option on page 126 describes the comfort span.</p> <p>Optional for a Lotus incremental backup.</p>	<ul style="list-style-type: none"> Undefined (default). Integer value between 65536 and 65536000, inclusive.
NSR_CROSS_MOUNT_POINTS	<p>Specifies whether NMDA crosses mount points during a Lotus backup.</p> <p>Optional for a Lotus backup.</p>	<ul style="list-style-type: none"> FALSE (default) = Lotus backup does not cross mount points. TRUE = Lotus backup crosses mount points.
NSR_DBIID	<p>Specifies that NMDA assigns either a new DBIID, or both a new DBIID and new replica ID, to a restored database.</p> <p>Optional for a Lotus restore.</p>	<ul style="list-style-type: none"> 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	<p>Specifies the complete pathname of a file that lists file pathnames to exclude from the Lotus backup.</p> <p>Optional for a Lotus backup.</p> <p>Wildcard support for Lotus operations on page 433 describes using wildcards in the pathnames.</p> <p>Note: If you set both NSR_EXCLUDE_FILE and</p>	<ul style="list-style-type: none"> If not specified, NMDA does not exclude pathnames from the backup paths that are specified by the user. Valid complete pathname of a file that lists file paths to exclude from the Lotus backup.

Table 41 NMDA Lotus parameters (continued)


Parameter	Description	Default and valid values
	<p>NSR_EXCLUDE_LIST, NMDA excludes all the files and objects that are specified through both parameters from the Lotus backup.</p> <p>NSR_EXCLUDE_FILE is a deprecated parameter that will be unsupported in a future NMDA release. Use the NSR_EXCLUDE_LIST parameter instead.</p>	
NSR_EXCLUDE_LIST	<p>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.</p> <p>Optional for a Lotus backup.</p> <p>Wildcard support for Lotus operations on page 433 describes the use of wildcards in the pathnames.</p> <p> Note: If you set both NSR_EXCLUDE_LIST and NSR_EXCLUDE_FILE, NMDA excludes all the files and objects that are specified through both parameters from the Lotus backup.</p>	<ul style="list-style-type: none"> • If not specified, NMDA does not exclude pathnames from the backup paths that are 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	<p>Specifies which of the following actions occur when Lotus link files are to be backed up or restored:</p> <ul style="list-style-type: none"> • 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. <p>Optional for a Lotus backup or restore.</p>	<ul style="list-style-type: none"> • 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	<p>Specifies the level of transaction log backup to perform during an incremental backup.</p> <p>Optional for a Lotus incremental backup only.</p> <p>Ignored for a full backup or transaction log only backup.</p>	<ul style="list-style-type: none"> • 0 = Do not process the transaction logs. • 1 (default) = Back up the transaction logs and mark the logs as reusable.
NSR_LOG_DIR	<p>Specifies the complete pathname of the log directory of a partitioned Domino server for disaster recovery only.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid complete pathname of the log directory.

Table 41 NMDA Lotus parameters (continued)


Parameter	Description	Default and valid values
	Optional for a Lotus restore.	
NSR_LOTUS_DATA_DIR	<p>Specifies the complete pathname of the directory that contains the Lotus Notes data.</p> <p>Optional for a Lotus backup. Required for a partitioned Domino server or multiple Domino installations.</p> <p> Note: This parameter does not specify that the data will be backed up.</p>	<ul style="list-style-type: none"> Undefined (default). Valid complete pathname of the directory that contains the Lotus Notes data.
NSR_MAX_TXN_LOGS	<p>Specifies the number of transaction logs stored in a single save set during a Lotus backup.</p> <p>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.</p> <p>Optional for a Lotus backup of transaction logs.</p>	<ul style="list-style-type: none"> 10 logs per save set (default). Integer number of logs per save set.
NSR_NO_NOTES_INIT	<p>Specifies whether to initialize the Notes API during a disaster recovery.</p> <p>Mandatory for disaster recovery.</p> <p>Optional for a Lotus restore.</p>	<ul style="list-style-type: none"> 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	<p>Specifies a timeout value in seconds during which NMDA retries a Lotus backup in either of the following cases:</p> <ul style="list-style-type: none"> 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. <p>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:</p> <ul style="list-style-type: none"> If NSR_SKIPDBERRORS = TRUE, NMDA skips the database and backs up the next database. 	<ul style="list-style-type: none"> 30 (default); signifies a timeout of 30 seconds. Integer value of timeout in seconds.

Table 41 NMDA Lotus parameters (continued)

Parameter	Description	Default and valid values
	<ul style="list-style-type: none"> If NSR_SKIPDBERRORS = FALSE, NMDA fails the backup with an error. Optional for a Lotus backup.	
NSR_NOTES_INI_PATH	Specifies the complete pathname of the <code>notes.ini</code> file, including the file name. Recommended for a Lotus backup or restore.	<ul style="list-style-type: none"> Undefined (default). Valid complete pathname of the <code>notes.ini</code> file.
NSR_NUMBER_LOGS	Specifies whether to restore the transaction logs during a disaster recovery only. This parameter is ignored during a regular Lotus recovery.	<ul style="list-style-type: none"> 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.	<ul style="list-style-type: none"> Value that is determined by the NetWorker server, based on the NetWorker client and server parallelisms (default). Integer number of the maximum concurrent backup or restore streams. For a scheduled backup, this parameter setting must be less than or equal to the Parallelism attribute value in the NetWorker Client resource. If multiple backups run concurrently on the client host, the total of all the NSR_PARALLELISM settings must be less than or equal to the Parallelism attribute value.
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.	<ul style="list-style-type: none"> 0 (default) = NMDA does not prefetch extra logs. NMDA restores only a log that is requested by Domino. Integer number of transaction logs that are retrieved in advance, typically the number of logs that are 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.	<ul style="list-style-type: none"> Undefined (default). n = Do not restore the current file.

Table 41 NMDA Lotus parameters (continued)


Parameter	Description	Default and valid values
	<p>The value of the parameter must be a single letter:</p> <ul style="list-style-type: none"> 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. <p>Optional for a Lotus restore.</p>	<ul style="list-style-type: none"> 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 file name, 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	<p>Specifies the complete pathname of a file that lists the files or directories to restore.</p> <p>Optional for a Lotus restore.</p> <p> Note: You cannot use this parameter with the NSR_BACKUP_PATHS parameter.</p>	<ul style="list-style-type: none"> 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	<p>Specifies additional recovery options.</p> <p>Optional for a Lotus recovery.</p>	<ul style="list-style-type: none"> Undefined (default). REMOVE_COMMON_PATH = For a restore to a new destination directory, removes the common path of the files and directories that are specified with NSR_BACKUP_PATHS or specified in a file set with NSR_RECOV_LIST_FILE.
NSR_RECOVER_TIME	<p>Specifies the point-in-time to which NMDA recovers a database.</p> <p>Optional for a Lotus restore or recovery.</p> <p>Setting the NSR_RECOVER_TIME parameter on page 434 describes the setting of this parameter.</p>	<ul style="list-style-type: none"> By default, NMDA performs the following recovery: <ul style="list-style-type: none"> 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.

Table 41 NMDA Lotus parameters (continued)

Parameter	Description	Default and valid values
		<ul style="list-style-type: none"> Valid time in <code>nsr_getdate(3)</code> format.
NSR_RELOCATION_DEST	<p>Specifies the pathname of a directory to which NMDA restores the Lotus database files.</p> <p>Optional for a Lotus restore.</p> <p>Note: You must set NSR_RELOCATION_DEST if the NSR_NUMBER_LOGS value is not zero.</p>	<ul style="list-style-type: none"> 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: <i>NSR_RELOCATION_DEST_path/original_file_path</i>. NSR_RECOVER_OPTIONS provides more information.
NSR_RESOURCE_DIR	<p>Specifies the location of the directory that contains the Lotus resource files.</p> <p>Mandatory for a Lotus backup on UNIX only.</p>	<ul style="list-style-type: none"> Undefined (default). Valid complete pathname of the Lotus directory that contains the resource files.
NSR_RETAIN_NUM_LOGS	<p>Specifies the number of archived transaction logs to retain after a Lotus backup without marking them reusable.</p> <p>Optional for a Lotus backup.</p>	<ul style="list-style-type: none"> 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 in the backup, retain all the logs without marking them reusable.
NSR_SAVESET_NAME	<p>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.</p> <p>Optional for a Lotus DAOS backup only.</p>	<ul style="list-style-type: none"> Undefined (default). Base name to use for DAOS backup save sets, for example, notes_DAOS.
NSR_SKIPDBERRORS	<p>Specifies whether NMDA continues a Notes database backup if a noncritical error occurs while backing up Notes database files, not flat files.</p> <p>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</p>	<ul style="list-style-type: none"> 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. <p>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.</p>

Table 41 NMDA Lotus parameters (continued)

Parameter	Description	Default and valid values
	<p>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.</p> <p>During the backup of multiple databases, NSR_SKIPDBERRORS enables NMDA to skip the backup of problematic databases while continuing to back up good databases.</p> <p>Optional for a Lotus backup.</p>	
NSR_VERBOSITY	<p>Specifies whether NMDA writes a list of backed-up or restored files to standard output or to the NetWorker User for Lotus GUI.</p> <p>Optional for a manual backup or restore.</p>	<ul style="list-style-type: none"> • 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	<p>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.</p> <p>Optional for a backup or restore.</p>	<ul style="list-style-type: none"> • 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	<p>Specifies the pathnames of the Domino data directory and Lotus software directory. Mandatory for a Lotus backup on UNIX only.</p> <p>Recommended for a Lotus backup on Windows and a Lotus restore.</p>	<ul style="list-style-type: none"> • 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

You must complete the required parameter settings for the NMDA MySQL backup and restore operations.

For the following MySQL backup operations, set both the common parameters and the NMDA MySQL parameters in the NMDA configuration file that is dedicated to MySQL backup parameters:

- MySQL scheduled backups that are 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 that is dedicated to MySQL restore parameters. Use a separate NMDA configuration file for the MySQL restore parameters.

[Common NMDA parameters](#) on page 406 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 42 NMDA MySQL parameters

Parameter	Description	Default and valid values
MYSQL_BACKUP_DIR	<p>Specifies the pathname of a directory to contain either the backup files that are extracted from a backup image or the binary logs that are restored from the NetWorker server. Mandatory for a MySQL restore operation except for a binary log restore or a list image operation.</p> <p>Optional for a binary log restore.</p>	<ul style="list-style-type: none"> • /nsr/apps/tmp/BACKUP (default). • Valid pathname of a directory to which the <code>nsrmysqlrc</code> user has write permissions. <p>Note: The parameter setting overrides the <code>backup_dir</code> setting in the MySQL configuration file.</p>
MYSQL_BACKUP_NAME	<p>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.</p> <p>Mandatory for a MySQL manual backup or a restore operation. Do not set this parameter for a scheduled backup.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid name for a MySQL backup: <ul style="list-style-type: none"> ■ For a MySQL backup, a backup image name that starts with <code>MYSQL:/. </code> For example: <code>MYSQL_BACKUP_NAME=MYSQL:/myinstance_whole</code> This example name is for the backup of a whole instance called <code>myinstance</code>. ■ For a MySQL restore, a backup image name or backup piece name, which is the unique save file name that is generated by NMDA for each backup. For example: <code>MYSQL_BACKUP_NAME=myinstance_whole</code> (backup image name) <code>MYSQL_BACKUP_NAME=myinstance_full_whole_1322094812</code> (backup piece name as obtained with the <code>nsrinfo</code> command) <p>Note: For an extract and prepare operation, you must specify a backup image name, not a backup piece name.</p> <p>The parameter setting overrides the <code>backup_image</code> setting in the MySQL configuration file.</p>
MYSQL_BINLOG	<p>Specifies one of the following values:</p> <ul style="list-style-type: none"> • 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. 	<ul style="list-style-type: none"> • Undefined (default). • One or two valid pathnames of binary logs: <ul style="list-style-type: none"> ■ Pathname of a single binary log to restore. For example: <code>MYSQL_BINLOG=/var/lib/mysql/bin.001</code>

Table 42 NMDA MySQL parameters (continued)

Parameter	Description	Default and valid values
	Optional for a MySQL instance recovery to the current time or a point-in-time.	<ul style="list-style-type: none"> Pathnames of the first and last logs in a log range, enclosed in square brackets and separated by a semi-colon. For example: <code>MYSQL_BINLOG=[/var/lib/mysql/bin.001; /var/lib/mysql/bin.005]</code> <p>Note: In a log range, both log pathnames must be identical except for the numeric extension in the base file name.</p>
MYSQL_CFG_FILE	<p>Specifies the pathname of the MySQL configuration file, <code>my.cnf</code>. This file is different from the NMDA MySQL configuration file that contains the <code>MYSQL_CFG_FILE</code> parameter. The MySQL configuration file provides the parameter settings for a specific server instance.</p> <p>Mandatory for a MySQL backup. Mandatory for a restore with MySQL 5.6 or later.</p> <p>Optional for a restore with MySQL 5.5.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the MySQL configuration file for a server instance. For example: <code>MYSQL_CFG_FILE=/etc/my.cnf</code>
MYSQL_COMPRESS	<p>Specifies whether the <code>mysqlbackup</code> process performs compression on the backup data.</p> <p>Optional for a MySQL backup.</p>	<ul style="list-style-type: none"> FALSE (default) = The <code>mysqlbackup</code> process does not perform compression. TRUE = The <code>mysqlbackup</code> process performs compression.
MYSQL_COMPRESS_LEVEL	<p>Specifies the level of compression that the <code>mysqlbackup</code> process performs on the backup data.</p> <p>Optional for a MySQL backup.</p>	<ul style="list-style-type: none"> 0 (default) = The <code>mysqlbackup</code> process does not perform compression. 1 to 9 = The <code>mysqlbackup</code> process performs the specified level of compression. The compression level increases with the numeric value.
MYSQL_DATABASES	<p>Specifies a list of MySQL database names or table names for backup in the form “<code>db1[.tbl1] db2[.tbl2] db3[.tbl3] ...</code>”.</p> <p>Optional for a MySQL backup.</p>	<ul style="list-style-type: none"> Undefined (default). List of database names or table names or both, enclosed in double quotation marks with multiple names separated by spaces. For example: <code>MYSQL_DATABASES="db1 db2.tbl2 db3"</code> <p>Note: The parameter setting overrides both the databases setting in the</p>

Table 42 NMDA MySQL parameters (continued)

Parameter	Description	Default and valid values
		<p>MySQL configuration file and the list of names in the file that is specified by the <code>databases_list_file</code> setting.</p>
<p><code>MYSQL_DATADIR</code></p>	<p>Specifies the pathname of a MySQL data directory for backups.</p> <p>Specifies the pathname of the original MySQL data directory or a different MySQL data directory for a recovery or copy back operation.</p> <p>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.</p> <p>Mandatory for a MySQL scheduled backup or a copy back operation.</p> <p>Optional for a MySQL recovery to the current time or a point-in-time.</p>	<ul style="list-style-type: none"> • Data directory pathname that is obtained from the <code>mysqld</code> section of the MySQL configuration file (default). • Valid pathname of the MySQL data directory. For a recovery or a copy back operation, the <code>nssrmysqlrc</code> user must have write permissions to the directory. For example: <code>MYSQL_DATADIR=/var/lib/mysql</code> <p>Note: The parameter setting overrides the <code>datadir</code> 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 <code>datadir</code> setting in the MySQL configuration file is ignored.</p>
<p><code>MYSQL_EXTRACT_PATHS</code></p>	<p>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.</p> <p>Optional for an extract operation.</p>	<ul style="list-style-type: none"> • Undefined (default). • <code>source_pathname</code> or <code>source_pathname > destination_pathname</code> where: <ul style="list-style-type: none"> ▪ <code>source_pathname</code> is the pathname of a file or directory in the backup image, relative to the root of the image. ▪ <code>destination_pathname</code> is the optional pathname on the local storage device to which the file or directory is extracted. <p>For example:</p> <pre>MYSQL_EXTRACT_PATHS=meta/backup_variables.txt > /tmp/backup_variables.txt</pre>
<p><code>MYSQL_INCLUDE</code></p>	<p>Specifies the per-table InnoDB datafiles to back up.</p> <p>Optional for a MySQL backup of InnoDB tables.</p>	<ul style="list-style-type: none"> • Undefined (default). • Regular expression that is enclosed in double quotation marks for the per-

Table 42 NMDA MySQL parameters (continued)

Parameter	Description	Default and valid values
		<p>table InnoDB datafiles to back up. For example: MYSQL_INCLUDE="db2.tbl2"</p> <p>Note: The parameter setting overrides the include setting in the MySQL configuration file.</p>
MYSQL_INCR_DIR	<p>Specifies the pathname of the MEB incremental backup directory.</p> <p>Optional for a MySQL backup.</p>	<ul style="list-style-type: none"> • /nsr/apps/tmp/BACKUP (default). • Valid pathname of the MEB incremental backup directory.
MYSQL_INCR_OPTIONS	<p>Specifies the type of incremental backup to perform for an InnoDB database with MEB version 3.7 or later.</p> <p>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.</p> <p>Optional for a MySQL incremental backup of an InnoDB database.</p>	<ul style="list-style-type: none"> • Undefined (default). • REDO_LOG_ONLY = Perform an incremental backup of only the redo log of an InnoDB database. <p>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.</p>
MYSQL_INNODB_LOG_FILE_SIZE	<p>Specifies the size of the InnoDB log file for a copy back operation.</p> <p>Mandatory for a copy back operation if the MySQL configuration file does not contain the corresponding setting.</p>	<ul style="list-style-type: none"> • Undefined (default). • Integer value in the range of 108576 to 4294967295. <p>Note: The parameter setting corresponds to the innodb_log_file_size setting in the MySQL configuration file.</p>
MYSQL_INNODB_LOG_FILES_IN_GROUP	<p>Specifies the number of InnoDB log files in the log file group for a copy back operation.</p> <p>Mandatory for a copy back operation if the MySQL configuration file does not contain the corresponding setting.</p>	<ul style="list-style-type: none"> • Undefined (default). • Integer value in the range of 2 to 100. <p>Note: The parameter setting corresponds to the innodb_log_files_in_group setting in the MySQL configuration file.</p>
MYSQL_LOG_OPTIONS	<p>Specifies options for a MySQL backup of binary logs. Based on the specified options, the backup performs one of the following operations:</p> <ul style="list-style-type: none"> • Backs up a whole instance and its binary logs. • Backs up only the binary logs for an instance. 	<ul style="list-style-type: none"> • Undefined (default). • One or more of the following options, with multiple options separated by commas: <ul style="list-style-type: none"> ▪ INCLUDE_LOGS specifies to back up the binary logs after the backup of a whole MySQL instance. If you specify both INCLUDE_LOGS and

Table 42 NMDA MySQL parameters (continued)

Parameter	Description	Default and valid values
	<p>Note: Binary logging must be enabled for the instance, and the instance must be online. Otherwise, the binary log backup fails.</p> <p>Mandatory for a MySQL backup of binary logs.</p>	<p>LOGS_ONLY_BACKUP, only the last value applies.</p> <ul style="list-style-type: none"> ■ 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. <p>Optionally, specify this option with either INCLUDE_LOGS or LOGS_ONLY_BACKUP.</p> <p>For example: MYSQL_LOG_OPTIONS=INCLUDE_LOGS, PURGE_LOGS</p>
<p>MYSQL_MEB_OPTIONS</p>	<p>Specifies additional MySQL backup or restore options. For example, you can specify the uncompress option so that a restore uncompresses a compressed backup.</p> <p>Note: You can obtain information from the backup save set metadata about whether a backup is compressed.</p> <p>With this parameter, you can specify any of the options from the mysqlbackup section of the MySQL configuration file, <code>my.cnf</code>.</p> <p>Optional for a MySQL backup or restore.</p>	<ul style="list-style-type: none"> • 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 <p>Note: The parameter setting overrides the corresponding settings in the MySQL configuration file.</p>
<p>MYSQL_MEB_PATH</p>	<p>Specifies the pathname of the installation directory of the MEB binary <code>mysqlbackup</code>.</p> <p>Mandatory for the following operations:</p> <ul style="list-style-type: none"> • 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. 	<ul style="list-style-type: none"> • <code>/opt/mysql/meb/bin</code> (default). • Valid pathname of the installation directory of the MEB binary. For example: MYSQL_MEB_PATH=/opt/mysql/meb-3.7.1/bin
<p>MYSQL_ONLY_INNO_DB_OPTIONS</p>	<p>Specifies the type of MySQL backup to perform when the backup contains only InnoDB databases or tables.</p> <p>Due to a limitation with MEB 3.7 or later, when you specify the WITH_FRM_ALL or</p>	<ul style="list-style-type: none"> • Undefined (default). • One of the following options for a backup of InnoDB databases or tables:

Table 42 NMDA MySQL parameters (continued)

Parameter	Description	Default and valid values
	<p>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.</p> <p>Optional for a MySQL backup that contains only InnoDB databases or tables.</p>	<ul style="list-style-type: none"> ■ NO_FRM specifies to exclude <code>.frm</code> files from the backup. ■ WITH_FRM_ALL (MEB 3.7 or later only) specifies to include <code>.frm</code> files for all the InnoDB tables in the instance. ■ WITH_FRM_RELATED (MEB 3.7 or later only) specifies to include <code>.frm</code> files for only the tables that are in the partial backup that is specified with <code>MYSQL_INCLUDE</code>. <p>For example: <code>MYSQL_ONLY_INNODB_OPTIONS=WITH_FRM_ALL</code></p> <p>Note: The parameter setting overrides these settings in the MySQL configuration file: <code>only-innodb</code>, <code>only-innodb-with-frm=all</code>, <code>only-innodb-with-frm=related</code>.</p>
<p><code>MYSQL_RESTORE_OPERATION</code></p>	<p>Specifies the restore operation to perform.</p> <p>Mandatory for a binary log restore, copy back, extract, extract and prepare, list image, or validate operation.</p>	<ul style="list-style-type: none"> • Undefined (default). • One of the following values to specify a restore operation: <ul style="list-style-type: none"> ■ <code>binlog_restore</code> specifies to perform a binary log restore. ■ <code>copy_back</code> specifies to perform a copy back of prepared backup data to a specified directory. ■ <code>extract</code> specifies to extract data from a backup image. ■ <code>extract_and_prepare</code> specifies to extract data from a backup image and prepare the data as required (preparation is only required for InnoDB data). ■ <code>list_image</code> specifies to list files from a backup image. ■ <code>validate</code> specifies to validate the integrity of a backup image, only with MEB 3.7 or later. <p>For example: <code>MYSQL_RESTORE_OPERATION=extract_and_prepare</code></p>

Table 42 NMDA MySQL parameters (continued)

Parameter	Description	Default and valid values
MYSQL_SBT_LIB_PATH	<p>Specifies the pathname of the NMDA SBT library that is used for MySQL.</p> <p>Mandatory when you have 32-bit and 64-bit MEB on the same system.</p>	<ul style="list-style-type: none"> • /usr/lib/libnsrmysql.so (default). • Valid pathname of the NMDA SBT library for MySQL, libnsrmysql.so. <p>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:</p> <p>MYSQL_SBT_LIB_PATH=/usr/lib/libnsrmysql32.so</p> <p>Note: The parameter setting overrides the sbt_lib_path setting in the MySQL configuration file.</p>
MYSQL_SRC_ENTRY	<p>Specifies the pathname of a single file or directory to be listed from a backup image.</p> <p>Optional for a list image operation.</p>	<ul style="list-style-type: none"> • 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	<p>Specifies the name of the MySQL backup user that connects to the MySQL instance.</p> <p>Note: You can specify the user password with the USER_PSWD parameter.</p> <p>Mandatory for a MySQL backup if the MySQL configuration file does not contain the username.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid username of the MySQL backup user. <p>Note: The parameter setting overrides the user setting in the MySQL configuration file.</p>
NSR_BACKUP_LEVEL	<p>Specifies the NetWorker backup level to use for a MySQL manual backup.</p> <p>Optional for a MySQL manual backup. Do not set this parameter for a scheduled backup.</p>	<ul style="list-style-type: none"> • 0 or FULL (default) = Perform a full backup. • 1 = Perform a cumulative incremental backup, which backs up all the data that is changed since the last full backup. • INCR = Perform a differential incremental backup, which backs up the data that is changed since the last full or incremental backup. <p>Note: A setting of 0, 1, or FULL overrides the MYSQL_INCR_OPTIONS parameter.</p>

Table 42 NMDA MySQL parameters (continued)

Parameter	Description	Default and valid values
NSR_DR_BACKUP_INFO	<p>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 that is specified by <code>MYSQL_CFG_FILE</code>.</p> <p>If the backup of additional information fails, an error message appears but the backup report is successful.</p> <p>Optional for a MySQL scheduled backup only.</p>	<ul style="list-style-type: none"> 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	<p>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 <code>POSTCMD</code> runs.</p> <p>If the backup of extra files fails, an error message appears but the backup report is successful.</p> <p>Optional for a MySQL scheduled backup only.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of a file that contains a list of additional files to back up during a scheduled backup. For example, <code>NSR_DR_FILE_LIST</code> is set to the pathname of a <code>nmda_savelist.txt</code> file that contains a list of file pathnames for backup. <p>Note: Set <code>NSR_DR_BACKUP_INFO=TRUE</code> to enable this parameter setting.</p>
NSR_LOG_VOLUME_POOL	<p>Specifies the NetWorker volume pool to use for a backup of the MySQL binary logs.</p> <p>Optional for a MySQL binary log backup.</p> <p>Note: If required, specify the associated storage node by setting the Storage Nodes attribute in the NetWorker Client resource (NMC diagnostic mode must be enabled).</p>	<ul style="list-style-type: none"> Most appropriate pool as selected by the NetWorker server (default). Valid name of a NetWorker volume pool for the MySQL binary logs.
NSR_RECOV_INTERACT	<p>Specifies whether to disable the prompt that the <code>nsmysqlrc</code> program issues for shutting down the MySQL database server prior to a copy back of backup data to the data directory.</p> <p>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.</p>	<ul style="list-style-type: none"> Undefined (default). <code>n</code> = The <code>nsmysqlrc</code> program does not issue the prompt to shut down the database server prior to a copy back operation to the data directory. For example: <code>NSR_RECOV_INTERACT=n</code>
NSR_RECOVER_TIME	<p>Specifies the point-in-time to which the <code>nsmysqlrc</code> program recovers the MySQL backup data.</p>	<ul style="list-style-type: none"> By default, the <code>nsmysqlrc</code> program recovers the backup data to the most recent available backup.

Table 42 NMDA MySQL parameters (continued)

Parameter	Description	Default and valid values
	Mandatory for a MySQL point-in-time recovery.	<ul style="list-style-type: none"> Valid date in <code>nsr_getdate(3)</code> format.
USER_PSWD	<p>Specifies the encrypted password for the MySQL backup user that connects to the MySQL instance. The encrypted password is added as the USER_PSWD setting to the NMDA MySQL configuration file.</p> <p>Note: You can specify the name of the MySQL backup user with the MYSQL_USER parameter.</p> <p>Mandatory for a MySQL backup if the MySQL configuration file does not contain the password.</p>	<ul style="list-style-type: none"> Undefined (default). Encrypted MySQL backup user password that you must set with the <code>nsrdaadmin -P</code> command, for example: <pre>nsrdaadmin -P -z configuration_file_path</pre> <p>The <i>NetWorker Module for Databases and Applications Command Reference Guide</i> describes the <code>nsrdaadmin</code> command.</p> <p>Note: The parameter setting overrides the password setting in the MySQL configuration file.</p>

NMDA Oracle parameters

You must complete the required parameter settings for the NMDA Oracle backup and restore operations.

For Oracle client-side scheduled backups, you must set the following parameters in the NMDA configuration file as described in [NMDA configuration file](#) on page 400:

- 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 can set all the other NMDA parameters for Oracle operations in the RMAN script. Alternatively, you can set any of these other parameters in an NMDA configuration file and specify the configuration file pathname with the CONFIG_FILE setting in the RMAN script. [CONFIG_FILE](#) provides details.

When you set an NMDA parameter in the RMAN script, use one of the following methods, unless specified otherwise in [Table 38](#) on page 406 (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 46 describes automatic channels. [Backup copies](#) on page 47 describes backup copies.

On Windows, do not set a parameter with the `parms 'ENV=(...)'` option. If you set a parameter with `parms 'ENV=(...)'` on Windows, 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 43 NMDA Oracle parameters

Parameter	Description	Default and valid values
CONFIG_FILE	<p>Specifies the complete pathname of an NMDA configuration file that contains the settings of NMDA Oracle parameters that you could also set in an RMAN script. Only specific parameters can be set in an RMAN script, as described at the start of this topic.</p> <p>Optional for an Oracle backup or restore.</p> <p>You must specify the CONFIG_FILE setting in the RMAN script. For example:</p> <pre>run { allocate channel t1 device type sbt; send 'NSR_ENV=(CONFIG_FILE=/ home/oradb/config/nmda.cfg)'; backup database plus archivelog;</pre>	<ul style="list-style-type: none"> • Undefined (default). • Valid pathname of the NMDA configuration file.

Table 43 NMDA Oracle parameters (continued)

Parameter	Description	Default and valid values
	<pre data-bbox="416 348 863 411"> release channel t1; } </pre> <p data-bbox="416 443 863 600">You can include the same CONFIG_FILE setting in different RMAN scripts, which enables you to share the parameters from the configuration file between those RMAN scripts.</p> <p data-bbox="416 621 863 709">NMDA configuration file on page 400 provides details on how to set parameters in an NMDA configuration file.</p>	
NSR_DATA_VOLUME_POOL 1	<p data-bbox="416 747 863 810">Specifies the name of the volume pool to use for a duplexed Oracle backup.</p> <p data-bbox="416 831 863 1020">Mandatory for an Oracle manual backup that uses the <code>set duplex</code> command (with <code>duplex set</code> to 2, 3, or 4) or other RMAN commands to generate two or more backup copies.</p> <p data-bbox="416 1041 863 1136">Set this parameter with the <code>parms</code> option in the RMAN script (not with the <code>send</code> command or <code>send</code> option).</p>	<ul data-bbox="885 747 1316 982" style="list-style-type: none"> • Undefined (default). • Valid NetWorker pool name that is different from the name that is used by the parameter <code>NSR_DATA_VOLUME_POOL</code>, <code>NSR_DATA_VOLUME_POOL2</code>, or <code>NSR_DATA_VOLUME_POOL3</code>.
NSR_DATA_VOLUME_POOL 2	<p data-bbox="416 1167 863 1230">Specifies the name of the volume pool to use for a duplexed Oracle backup.</p> <p data-bbox="416 1251 863 1440">Mandatory for an Oracle manual backup that uses the <code>set duplex</code> command (with <code>duplex set</code> to 3 or 4) or other RMAN commands to generate three or more backup copies.</p> <p data-bbox="416 1461 863 1556">Set this parameter with the <code>parms</code> option in the RMAN script (not with the <code>send</code> command or <code>send</code> option).</p>	<ul data-bbox="885 1167 1316 1402" style="list-style-type: none"> • Undefined (default). • Valid NetWorker pool name that is different from the name that is used by the parameter <code>NSR_DATA_VOLUME_POOL</code>, <code>NSR_DATA_VOLUME_POOL1</code>, or <code>NSR_DATA_VOLUME_POOL3</code>.
NSR_DATA_VOLUME_POOL 3	<p data-bbox="416 1587 863 1650">Specifies the name of the volume pool to use for a duplexed Oracle backup.</p> <p data-bbox="416 1671 863 1860">Mandatory for an Oracle manual backup that uses the <code>set duplex</code> command (with <code>duplex set</code> to 4) or other RMAN commands to generate four backup copies.</p> <p data-bbox="416 1881 863 1955">Set this parameter with the <code>parms</code> option in the RMAN script (not with the <code>send</code> command or <code>send</code> option).</p>	<ul data-bbox="885 1587 1316 1822" style="list-style-type: none"> • Undefined (default). • Valid NetWorker pool name that is different from the name that is used by the parameter <code>NSR_DATA_VOLUME_POOL</code>, <code>NSR_DATA_VOLUME_POOL1</code>, or <code>NSR_DATA_VOLUME_POOL2</code>.

Table 43 NMDA Oracle parameters (continued)

Parameter	Description	Default and valid values
NSR_ENV_LIST	<p>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 that are supported by RMAN.</p> <p>Optional for an Oracle scheduled backup.</p> <p>Set this parameter in the configuration file only.</p>	<ul style="list-style-type: none"> Undefined (default). A list of one or more environment variable settings that are supported by RMAN. Each setting includes the variable name, an equal sign, and the variable value. Multiple settings are separated by commas, and the entire list is enclosed in double quotes, such as "name1=value1, name2=value2". <p>Note: Do not include an equal sign (=) or comma (,) or single quote (') within a variable value, such as value1 or value2.</p> <p>For example, the following parameter specifies the NLS_DATE_FORMAT environment variable setting: NSR_ENV_LIST="NLS_DATE_FORMAT=dd-mm-yyyy hh.mi.ss"</p>
NSR_LOG_VOLUME_POOL	<p>Specifies the name of the NetWorker volume pool to use for a backup of the archived redo logs.</p> <p>Mandatory for an Oracle manual snapshot backup.</p> <p>Optional for an Oracle manual nonsnapshot backup.</p> <p>Note: Set both NSR_DATA_VOLUME_POOL and NSR_LOG_VOLUME_POOL to send the database files and log files to separate volume pools. If required, specify the associated storage node by setting the Storage Nodes attribute in the NetWorker Client resource (NMC diagnostic mode must be enabled).</p>	<ul style="list-style-type: none"> Most appropriate pool as selected by the NetWorker server (default). Valid name of a NetWorker volume pool for the archived redo logs.
NSR_MMDB_RETRY_TIME	<p>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 tries.</p> <p>Optional for an Oracle backup or restore.</p>	<ul style="list-style-type: none"> 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.

Table 43 NMDA Oracle parameters (continued)

Parameter	Description	Default and valid values
NSR_NO_MULTIPLEX	<p>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.</p> <p>Optional for an Oracle backup.</p> <p>To optimize restore operations, RMAN requires Oracle backups not to be multiplexed.</p> <p>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.</p>	<ul style="list-style-type: none"> • FALSE (default) = Enable multiplexing on the device that the RMAN channel uses. • TRUE = Disable multiplexing on the device that the RMAN channel uses. <p>Note: Do not set this parameter to TRUE if you use a random access NetWorker device, such as an advanced file device.</p>
NSR_ORACLE_RETENTION	<p>Specifies whether to use the Oracle RMAN retention policy to manage the backup lifecycle.</p> <p>Optional for an Oracle backup.</p> <p>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.</p> <p>Note: NMDA does not support this parameter for a scheduled backup of Oracle disk backups.</p>	<ul style="list-style-type: none"> • FALSE (default) = Enable the NetWorker retention policy. Use the policy 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 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	<p>Specifies the complete pathname of a configuration file that contains the NSM parameter settings for a snapshot backup or restore.</p> <p>Mandatory for a manual backup or restore if a configuration file contains the NSM parameters settings for the operation. Supported for an NSM snapshot backup or restore only.</p> <p>For a scheduled backup, you must set the NSM parameters in the Application Information attribute of the Client resource.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid pathname of the configuration file. <p>Note: If undefined or an invalid pathname, the operation ignores the parameter settings in the preferred configuration file.</p>
NSR_RMAN_ARGUMENTS	<p>Specifies any valid combination of options for the RMAN executable, <code>rman (.exe)</code>. The Oracle Recovery Manager</p>	<ul style="list-style-type: none"> • Undefined (default).

Table 43 NMDA Oracle parameters (continued)

Parameter	Description	Default and valid values
	<p>documentation describes the valid options.</p> <p>Optional for an Oracle scheduled backup.</p> <p>Set this parameter in the configuration file only.</p> <p>Note: NMDA does not support this parameter for a scheduled backup of Oracle disk backups.</p>	<ul style="list-style-type: none"> String that contains any valid combination of options for the RMAN executable, <code>rman(.exe)</code>. For example, set the parameter to append RMAN output to the message log file <code>/nsr/apps/logs/msglog.log</code> if you do not use a Recovery Catalog: <code>NSR_RMAN_ARGUMENTS="nocatalog log /nsr/apps/logs/msglog.log append"</code>
NSR_RMAN_OUTPUT	<p>Specifies options to control how the RMAN output is redirected.</p> <p>Optional for an Oracle scheduled backup.</p> <p>Set this parameter in the configuration file only.</p>	<ul style="list-style-type: none"> 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 <code>NSR_RMAN_ARGUMENTS</code> specifies to append RMAN output to a message log file with a date stamp in the log file name: <code>NSR_RMAN_ARGUMENTS="nocatalog log /nsr/apps/logs/dbid_arch_msglog_`date +%y%m%d_%H`.log append"</code> <code>NSR_RMAN_OUTPUT=SHELL</code>
NSR_SERVER_NIC	<p>Specifies the name of a network interface card (NIC) on a NetWorker server.</p> <p>Optional for an Oracle backup or restore.</p> <p>When you set this parameter with the RMAN send command for an allocated channel, the parameter value overrides the <code>NSR_SERVER</code> setting for that channel only.</p> <p>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 <code>NSR_SERVER</code> value for a scheduled backup.</p>	<ul style="list-style-type: none"> Undefined (default). Valid name of a NetWorker server NIC.
NSR_VOLUMES_INFO	<p>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 <code>restore...preview</code> command or <code>restore...preview recall</code> command</p>	<ul style="list-style-type: none"> FALSE (default) = NMDA does not obtain the latest NetWorker volume information. TRUE = NMDA obtains the latest NetWorker volume information.

Table 43 NMDA Oracle parameters (continued)

Parameter	Description	Default and valid values
	Recommended for NMDA Oracle restore preview functionality.	
ORACLE_HOME	<p>Specifies the home directory pathname of the Oracle Server installation.</p> <p>Mandatory for an Oracle scheduled backup.</p> <p>Set this parameter in the configuration file only.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid pathname of the home directory of the Oracle Server installation.
ORACLE_SID	<p>Specifies the system identifier (SID) value of the Oracle database to be backed up.</p> <p>Mandatory for an Oracle scheduled backup in the following cases:</p> <ul style="list-style-type: none"> • You store the <code>connect target</code> and <code>connect rcvcat</code> commands for the scheduled backup in a separate file, and you run the <code>connect</code> commands in the RMAN script by using the <code>@</code> command. • You enable save set bundling for the scheduled backup. • You perform an NSM snapshot backup with catalog synchronization enabled. Snapshot Backups and Restores on page 345 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. <p>Set this parameter in the configuration file only.</p>	<ul style="list-style-type: none"> • 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: <code>ORACLE_SID=orcl10</code>
ORACLE_USER (UNIX)	<p>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.</p> <p>Optional for an Oracle scheduled backup in a client-side configuration (configured with the NMC method, not the wizard) on UNIX systems only.</p>	<ul style="list-style-type: none"> • 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 43 NMDA Oracle parameters (continued)

Parameter	Description	Default and valid values
	<p>Set this parameter in the configuration file only.</p> <p>Note: The following types of NMDA backups do not support the use of ORACLE_USER for performing an Oracle backup through operating system authentication:</p> <ul style="list-style-type: none"> - Scheduled backup on Windows - Probe-based backup 	
TNS_ADMIN	<p>Specifies the directory pathname of the Oracle Net configuration files.</p> <p>Mandatory for an Oracle scheduled backup if the Oracle Net configuration files reside in a directory other than the default <code>\$ORACLE_HOME/network/admin</code> directory.</p> <p>Set this parameter in the configuration file only.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid pathname of the directory that contains the Oracle Net configuration files.

NMDA Orchestrated Application Protection parameters

You must complete the required parameter settings for the NMDA backup and restore operations that are performed through the Orchestrated Application Protection feature.

Refer to the appropriate topic for details on how to set the NMDA parameters for backup or restore operations with the particular database or application:

- [NMDA parameters for Orchestrated Application Protection backups](#) on page 451
- [NMDA parameters for Orchestrated Application Protection restores](#) on page 456

NMDA parameters for Orchestrated Application Protection backups

You must complete the required parameter settings for the NMDA backup operations that are performed through the Orchestrated Application Protection feature.

Set the parameters in an NMDA configuration file that you create with the required XML tags, as described in [NMDA configuration file](#) on page 400. Set the backup parameters in the backup parameter section of the configuration file, which starts with a `<BACKUP>` tag and ends with a `</BACKUP>` tag.

The following table describes the NMDA parameters for Orchestrated Application Protection backups.

Table 44 NMDA parameters for Orchestrated Application Protection backups

Parameters	Description	Default and valid values
NSR_BACKUP_NAME	Specifies a logical name for the backup.	<ul style="list-style-type: none"> • Undefined (default). • Name of the backup.

Table 44 NMDA parameters for Orchestrated Application Protection backups (continued)






Parameters	Description	Default and valid values
	<p> Note: If you change this parameter setting after several backups, run a full backup immediately.</p> <p>Mandatory.</p>	<p> Note: The backup name cannot include a space.</p>
NSR_BACKUP_SCRIPT	<p>Specifies the complete pathname of the backup shell script to use for the backup. Based on the backup level, this parameter must be set in the <FULL>, <INCR>, or <TXNLOG> subsection of the configuration file.</p> <p>The backup shell script must include one or more pathnames of target files or directories for the backup, and each pathname must start with \$OAPP_MOUNT_DIR/.</p> <p>If you include one or more directory pathnames in the script, ensure that the native backup utility can create these directories if required. The documentation of the Orchestrated Application Protection protected database provides information on the native backup utility.</p> <p> Note: The backup script file must be owned by the database user that runs the backup, or the user that is specified by NSR_OS_USER. The group users and other users should not have permission to access the file.</p> <p> NOTICE The backup shell script must return a nonzero value when the backup fails. Verify that this is the case before you run any backups.</p> <p>Mandatory.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid complete pathname of the backup shell script, enclosed within double quotes. For example: <NSR_BACKUP_SCRIPT>/backup.sh</NSR_BACKUP_SCRIPT>
NSR_DATABASE_TYPE	<p>Specifies the database type for the backup.</p> <p>Mandatory.</p>	<ul style="list-style-type: none"> • Undefined (default). • Database type, such as MongoDB, MySQL, or PostgreSQL. <p> Note: This parameter value cannot include a space.</p>
NSR_DEBUG_BOOSTFS	<p>Specifies whether to generate the BoostFS binary debug log file, which is in the directory that is specified by NSR_DIAGNOSTIC_DEST or in the default directory /nsr/app/logs.</p>	<ul style="list-style-type: none"> • FALSE (default) = NMDA does not generate the BoostFS debug log file. • TRUE = NMDA generates the BoostFS debug log file with a .log file name extension.

Table 44 NMDA parameters for Orchestrated Application Protection backups (continued)

Parameters	Description	Default and valid values
	<p>Use this parameter for debugging purposes with assistance from Customer Support only.</p> <p>Optional.</p>	
NSR_ENV_LIST	<p>Specifies a list of additional environment variables to set before starting the backup command during a scheduled backup or before running a manual backup with the <code>nsroappbackup</code> command. The database administration documentation describes the environment variables that are required by the backup command.</p> <p>Ensure that you set <code>LD_LIBRARY_PATH</code> or the equivalent environment variable value with <code>NSR_ENV_LIST</code> if the database backup utility requires the setting.</p> <p>Set this parameter in the configuration file only.</p> <p>Mandatory if <code>LD_LIBRARY_PATH</code> or equivalent must be set.</p>	<ul style="list-style-type: none"> Undefined (default). A list of one or more environment variable settings. Each setting includes the variable name, an equal sign, and the variable value. Multiple settings are separated by commas, and the entire list is enclosed within double quotes, such as <code>"name1=value1,name2=value2"</code>. <p>Note: Do not include an equal sign (=) or comma (,) within a variable value, such as <code>value1</code> or <code>value2</code>.</p>
NSR_FORCED_VOLLOC	<p>Specifies the host to use as the storage node for the backup. The NetWorker <code>nsr_storage_node(5)</code> man page provides more details.</p> <p>Optional.</p>	<ul style="list-style-type: none"> Hostname of the NetWorker server (default). Valid hostname of the storage node host.
NSR_INSTANCE_NAME	<p>Specifies a name to describe the database instance that is protected by an Orchestrated Application Protection backup.</p> <p>Note: If you change this parameter setting after several backups, run a full backup immediately.</p> <p>Mandatory.</p>	<ul style="list-style-type: none"> Undefined (default). Name or nickname of the database instance that is protected by Orchestrated Application Protection. <p>Note: This parameter value cannot include a space.</p>
NSR_LOG_VOLUME_POOL	<p>For a transaction log backup only, specifies the NetWorker volume pool to use for the log backup.</p> <p>Note: If this parameter is not set in the <code><TXNLOG></code> section of the configuration file, the Default pool is used.</p> <p>Optional.</p>	<ul style="list-style-type: none"> Default pool (default). Valid name of a NetWorker volume pool for the transaction log backup.

Table 44 NMDA parameters for Orchestrated Application Protection backups (continued)

Parameters	Description	Default and valid values
NSR_OS_USER	<p>Specifies the username of the operating system user that will run the backup shell script that is specified by the NSR_BACKUP_SCRIPT setting.</p> <p>Note: This user must own the configuration file and the backup script file. The group users and other users should not have permission to access those files.</p> <p>Mandatory for a scheduled backup.</p>	<ul style="list-style-type: none"> Undefined (default). Valid username of the operating system user to run the backup shell script.
NSR_PROMOTE_FULL	<p>Specifies whether to automatically promote an incremental (incr) backup or transaction log (txnlog) backup to a full backup when the full backup is missing.</p> <p>Note: This parameter only applies when a corresponding full backup does not exist at the time that you run an incremental or transaction log backup.</p> <p>Optional.</p>	<ul style="list-style-type: none"> 0 (default) = Do not automatically promote an incremental or transaction log backup to a full backup when a full backup does not exist. 1 = Automatically promote an incremental or transaction log backup to a full backup when a full backup does not exist. <p>Note: It is recommended that you run the required incremental or transaction log backup after the full backup has completed successfully.</p>
POSTCMD	<p>Specifies a postprocessing script to run after a scheduled backup:</p> <ul style="list-style-type: none"> The postprocessing is run by the user that is specified by NSR_OS_USER or (when NSR_OS_USER is not specified) the user that runs the nsroappbackup command. 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 <p>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. This parameter can be set in all the <FULL>, <INCR>, and <TXNLOG> sections of the configuration file.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of a postprocessing script file. The pathname must not contain any spaces. If the value is undefined or invalid, a postprocessing script does not run after the scheduled backup.

Table 44 NMDA parameters for Orchestrated Application Protection backups (continued)


Parameters	Description	Default and valid values
	Optional for a scheduled backup. Do not set this parameter for a manual backup.	
PRECMD	<p>Specifies a preprocessing script to run before a scheduled backup:</p> <ul style="list-style-type: none"> The preprocessing is run by the user that is specified by NSR_OS_USER or (when NSR_OS_USER is not specified) the user that runs the <code>nsroappbackup</code> command. 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: <code>#!/bin/sh</code> <p> Note: If the preprocessing script fails, NMDA does not perform the scheduled backup, an error message appears, and any postprocessing script does not run. This parameter can be set in all the <FULL>, <INCR>, and <TXNLOG> sections of the configuration file.</p> <p>Optional for a scheduled backup. Do not set this parameter for a manual backup.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of a preprocessing script file. The pathname must not contain any spaces. If the value is undefined or invalid, a preprocessing script does not run before the scheduled backup.
USER_PSWD	<p>Specifies the encrypted password if the database backup utility requires the password to be inputted through the command line. The encrypted password is added as the USER_PSWD setting to the NMDA configuration file.</p> <p>Mandatory when the backup requires a password.</p>	<ul style="list-style-type: none"> Undefined (default). Encrypted password that you must set with the <code>nsrdaadmin -P</code> command, for example: <pre>nsrdaadmin -P -z configuration_file_path</pre> <p>The <i>NetWorker Module for Databases and Applications Command Reference Guide</i> describes the <code>nsrdaadmin</code> command.</p>
USER_PSWD_PROMPT	Specifies the password prompt that the database backup utility displays when a password must be inputted through the command line. This parameter setting	<ul style="list-style-type: none"> Undefined (default).

Table 44 NMDA parameters for Orchestrated Application Protection backups (continued)

Parameters	Description	Default and valid values
	<p>enables the <code>nsroappbackup</code> program to pass the encrypted password that is set in <code>USER_PSWD</code> to the backup utility.</p> <p>Mandatory when the backup requires a password.</p>	<ul style="list-style-type: none"> Password prompt that the database backup utility displays when the backup requires a password. Each database or application requires a particular setting of the password prompt. To determine the required prompt, you can manually run the database backup command (that you include in the backup script file) from the command line. <p>For example, the PostgreSQL database backup utility requires the following setting:</p> <pre><USER_PSWD_PROMPT>Password: </USER_PSWD_PROMPT></pre>

NMDA parameters for Orchestrated Application Protection restores

You must complete the required parameter settings for the NMDA restore operations that are performed through the Orchestrated Application Protection feature.

Set the parameters in an NMDA configuration file that you create with the required XML tags, as described in [NMDA configuration file](#) on page 400. Set the restore parameters in the restore parameter section of the configuration file, which starts with a `<RESTORE>` tag and ends with a `</RESTORE>` tag.

The following table describes the NMDA parameters for Orchestrated Application Protection restores.

Table 45 NMDA parameters for Orchestrated Application Protection restores

Parameters	Description	Default and valid values
<code>NSR_BACKUP_NAME</code>	<p>Specifies the name of the data backup for the restore. The backup name is used to locate the backup save set.</p> <p>Mandatory for a PostgreSQL transaction log restore.</p>	<ul style="list-style-type: none"> Undefined (default). Name of the data backup.
<code>NSR_DATABASE_TYPE</code>	<p>Specifies the same value as used for the backup.</p> <p>Mandatory for a PostgreSQL transaction log restore.</p>	<ul style="list-style-type: none"> Undefined (default). Same value as used for backup.
<code>NSR_FORCED_VOLLOC</code>	<p>Specifies the hostname of the storage node to use for the recovery. The NetWorker <code>nsr_storage_node(5)</code> man page provides more details.</p>	<ul style="list-style-type: none"> Hostname of the NetWorker server (default). Valid hostname of the storage node host.

Table 45 NMDA parameters for Orchestrated Application Protection restores (continued)

Parameters	Description	Default and valid values
	Optional.	
NSR_INSTANCE_NAME	Specifies the same value as used for the backup. Mandatory for a PostgreSQL transaction log restore.	<ul style="list-style-type: none"> Undefined (default). Same value as used for backup.
NSR_RECOVER_TIME	Specifies to recover the files as of the specified date, which includes the time. Optional.	<ul style="list-style-type: none"> Undefined (default). Valid date in <code>nsr_getdate(3)</code> format.
NSR_RELOCATION_DEST	Specifies the destination directory where the recovered files will be relocated. The directory must not exist initially, as the <code>nsroapprecovery</code> program creates the directory. Ensure that the user that runs <code>nsroapprecovery</code> has full permissions in the parent directory of the destination directory. Mandatory.	<ul style="list-style-type: none"> Undefined (default). Valid complete pathname of the destination directory where the recovered files will be relocated.
NSR_SAVESET_NAME	Specifies the name of the backup save set to be restored. The specified name must work with NSR_RECOVER_TIME, if set. If NSR_RECOVER_TIME is not set, the latest save set with the specified name is restored. Mandatory for all types of restores except a PostgreSQL transaction log restore.	<ul style="list-style-type: none"> Undefined (default). Valid save set name for the backup to be restored. <p>Note: The save set name is case-sensitive and must be in the same case as recorded in the corresponding backup entry in the NetWorker media database.</p>

NMDA SAP IQ parameters

You must complete the required parameter settings for the NMDA SAP IQ backup and restore operations.

For the following SAP IQ operations, set both the common parameters and the NMDA SAP IQ parameters in the NMDA configuration file:

- SAP IQ scheduled backups that are configured without the wizard (client-side configuration)
- SAP IQ manual backups
- SAP IQ restores

[Common NMDA parameters](#) on page 406 describes the common parameters.

The following table describes the NMDA SAP IQ parameters.

Note: The parameter names are case-sensitive. The parameter values are case-insensitive unless stated otherwise.

Table 46 NMDA SAP IQ parameters

Parameter	Description	Default and valid values
IQ_CLEAR_LOG	<p>Specifies whether to clear the logs during an SAP IQ point-in-time restore.</p> <p>Optional for an SAP IQ point-in-time restore.</p>	<ul style="list-style-type: none"> FALSE (default) = Do not clear the logs during a point-in-time restore. TRUE = Clear the logs during a point-in-time restore.
IQ_OCS_PATH	<p>Specifies the pathname of the directory where the Open Client Server (OCS) library is installed, for example, /sap/OCS-15_0.</p> <p>Mandatory for an SAP IQ scheduled backup.</p> <p>Optional for an SAP IQ manual backup.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the directory where the OCS library is installed.
IQ_OVERWRITE_EXISTING	<p>Specifies whether to overwrite the existing database files during an SAP IQ restore.</p> <p>Optional for an SAP IQ restore.</p>	<ul style="list-style-type: none"> FALSE (default) = Do not overwrite the existing database files during a restore. TRUE = Overwrite the existing database files during a restore.
IQ_PIT_RESTORE_BEFORE_PREV_FULL	<p>Specifies whether the transaction log offset occurs between the previous full backup and the log backup before that full backup.</p> <p>Mandatory for an SAP IQ point-in-time restore when the offset is between the previous full backup and the log backup before that full backup.</p>	<ul style="list-style-type: none"> FALSE (default) = Transaction log offset does not occur between the previous full backup and the log backup before that full backup. TRUE = Transaction log offset occurs between the previous full backup and the log backup before that full backup. <p>Note: When this parameter is set to TRUE, NSR_RECOVER_TIME must be set to the save set time of the transaction log backup after the previous full backup.</p>
IQ_PIT_RESTORE_ENABLE	<p>Specifies whether to perform an SAP IQ point-in-time restore.</p> <p>Mandatory for an SAP IQ point-in-time restore.</p>	<ul style="list-style-type: none"> FALSE (default) = Do not perform a point-in-time restore. TRUE = Perform a point-in-time restore.
IQ_PIT_RESTORE_LOG_PATH	<p>Specifies a list of pathnames of the transaction log files to include in an SAP IQ log backup or point-in-time restore:</p> <ul style="list-style-type: none"> List of log files to back up during an SAP IQ log backup, including the active log file. List of log files to restore during an SAP IQ point-in-time restore. You must include the active log file of the database at the time to which the restore is to be performed. 	<ul style="list-style-type: none"> Undefined (default). One or more transaction log file pathnames for a log backup or point-in-time restore. Separate the multiple log file pathnames with a comma: IQ_PIT_RESTORE_LOG_PATH= <i>log_file_pathname1</i>, <i>log_file_pathname2...</i>

Table 46 NMDA SAP IQ parameters (continued)

Parameter	Description	Default and valid values
<p>IQ_PIT_RESTORE_OFFSET</p>	<p>Mandatory for an SAP IQ log backup or point-in-time restore.</p> <p>Specifies a transaction log offset for an SAP IQ point-in-restore. The database is restored to the specified offset during the point-in-time restore operation.</p> <p>Note: The transaction log save set (with the NSR_RECOVER_TIME save time) must include the transaction for the specified offset.</p> <p>Mandatory for an SAP IQ point-in-time restore.</p>	<ul style="list-style-type: none"> Undefined (default). Valid value of the transaction log offset for the point-in-time restore.
<p>IQ_PIT_RESTORE_REDIRECTED_LOG_PATH</p>	<p>Specifies an alternate directory location to which the backed-up log files are restored before an SAP IQ point-in-time restore is performed.</p> <p>Optional for an SAP IQ point-in-time restore.</p>	<ul style="list-style-type: none"> Undefined (default). Note: If this parameter is not set, a temporary directory is created inside the first path in IQ_PIT_RESTORE_LOG_PATH and the backed-up log files are restored to this temporary directory prior to a point-in-time restore. Valid pathname of an existing directory where the backed-up log files are restored prior to a point-in-time restore.
<p>IQ_PITR_LOG_RETENTION_PATH</p>	<p>Specifies a local archive directory where the transaction log files and log archives are moved after an SAP IQ log backup.</p> <p>Optional for an SAP IQ log backup.</p>	<ul style="list-style-type: none"> Undefined (default). Note: If this parameter is not set, the log files remain in the original location after a log backup. Valid pathname of an existing directory (local archive) where the transaction log files and log archives are moved after a log backup.
<p>IQ_PITR_LOG_RETENTION_TIME</p>	<p>Specifies how long to retain the log files in the local archive directory specified by IQ_PITR_LOG_RETENTION_PATH, after the files are moved to that directory.</p> <p>IQ_PITR_LOG_RETENTION_PATH must be set to a valid directory pathname.</p> <p>Optional for an SAP IQ log backup.</p>	<ul style="list-style-type: none"> 0 (default). Note: When this parameter is not set or is set to 0, the log files are retained in the local archive directory until they are manually removed. Integer value of 0 or more, as the number of days to retain the log files in the directory.

Table 46 NMDA SAP IQ parameters (continued)

Parameter	Description	Default and valid values
<p>IQ_READONLY_DBFILES</p>	<p>When IQ_SELECTIVE_TYPE is set to READONLY, specifies one of the following values:</p> <ul style="list-style-type: none"> • For a full backup, a list of read-only dbfiles to back up. <p>Incremental and incremental since full backups are not supported for read-only data objects.</p> <ul style="list-style-type: none"> • For a restore, the single read-only dbfile to restore. <p>Mandatory for an SAP IQ selective backup or restore of read-only dbfiles.</p> <p>This parameter is ignored when IQ_SELECTIVE_TYPE is not set to READONLY.</p>	<ul style="list-style-type: none"> • Undefined (default). • One or more read-only dbfiles to back up. Separate the multiple dbfile names with a comma: IQ_READONLY_DBFILES=<i>dbfile_name1</i> [, <i>dbfile_name2...</i>] • A single read-only dbfile to restore.
<p>IQ_READONLY_DBSPACES</p>	<p>When IQ_SELECTIVE_TYPE is set to READONLY, specifies one of the following values:</p> <ul style="list-style-type: none"> • For a full backup, a list of read-only dbspaces to back up. <p>Incremental and incremental since full backups are not supported for read-only data objects.</p> <ul style="list-style-type: none"> • For a restore, the single read-only dbspace to restore. <p>Mandatory for an SAP IQ selective backup or restore of read-only dbspaces.</p> <p>This parameter is ignored when IQ_SELECTIVE_TYPE is not set to READONLY.</p>	<ul style="list-style-type: none"> • Undefined (default). • One or more read-only dbspaces to back up. Separate the multiple dbspace names with a comma: IQ_READONLY_DBSPACES= <i>dbspace_name1</i>, <i>dbspace_name2...</i>] • A single read-only dbspace to restore.
<p>IQ_SELECTIVE_TYPE</p>	<p>Specifies the type of selective backup or restore to perform.</p> <p>Mandatory for an SAP IQ selective backup or restore of only the read-write database files or specified read-only data objects.</p>	<ul style="list-style-type: none"> • ALL_INCLUSIVE (default). • One of the following values: <ul style="list-style-type: none"> ▪ ALL_INCLUSIVE = Perform a backup or restore of both the read-only and read-write data objects. All backup levels are supported for an all-inclusive backup. ▪ READONLY = Perform a backup or restore of the specified read-only dbspaces or dbfiles or both. IQ_READONLY_DBFILES or

Table 46 NMDA SAP IQ parameters (continued)

Parameter	Description	Default and valid values
		<p>IQ_READONLY_DBSPACES or both must also be set. Only a full backup is supported for a read-only backup.</p> <ul style="list-style-type: none"> ▪ READWRITE = Perform a backup or restore of all the read-write database files. The backup skips any read-only dbspaces and dbfiles. All backup levels are supported for a read-write backup.
IQ_USER	<p>Specifies the name of the SAP IQ user that runs the backup or restore and connects to the SAP IQ server. You must specify the password of the user by setting the USER_PSWD parameter.</p> <p>Mandatory for an SAP IQ backup or restore.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid SAP IQ username.
NSR_BACKUP_LEVEL	<p>Specifies the level of SAP IQ manual backup to perform.</p> <p>Optional for an SAP IQ manual backup.</p> <p>Note: Do not set this parameter for a scheduled backup.</p> <p>For a manual backup of read-only data objects, the backup level must be full; otherwise, the backup fails.</p>	<ul style="list-style-type: none"> • full (default) = Perform a full backup, which backs up the database. • cumulative = Perform an "incremental since full" backup, which backs up the data blocks that have changed since the last full backup. • incr = Perform an incremental backup, which backs up the data blocks that have changed since the last backup of any type. • txnlog = Perform a log backup that backs up the active log and the transaction log files that are specified by IQ_PIT_RESTORE_LOG_PATH.
NSR_DEBUG_LEVEL	<p>Specifies the level of debug messages that NMDA writes to the debug log file, which is in the directory that is specified by NSR_DIAGNOSTIC_DEST or in the default directory, /nsr/apps/logs (UNIX) or NetWorker_install_path\apps\logs (Windows).</p> <p>Note: Use this parameter for debugging purposes with assistance from Customer Support only.</p> <p>Optional for a backup or restore.</p>	<ul style="list-style-type: none"> • 0 (default) = NMDA does not generate debug messages. • 1 to 9 = NMDA writes debug messages to the debug log file with a .log file name extension. The level of detail in the generated debug messages increases with the debug level. <p>For an SAP IQ log backup or restore, use a minimum value of 3 to enable the INFO level logging of the backed-up or restored log files.</p>
NSR_LOCALE	<p>Specifies the locale to use to connect to the SAP IQ server for a client-side</p>	<ul style="list-style-type: none"> • Undefined (default).

Table 46 NMDA SAP IQ parameters (continued)

Parameter	Description	Default and valid values
	<p>scheduled backup or a manual backup configuration.</p> <p>Optional for an SAP IQ backup.</p>	<ul style="list-style-type: none"> Valid locale identifier.
NSR_LOG_VOLUME_POOL	<p>Specifies the name of the NetWorker volume pool to use for a backup of the SAP IQ transaction logs.</p> <p>Optional for an SAP IQ log backup.</p>	<ul style="list-style-type: none"> Most appropriate pool as selected by the NetWorker server (default). Valid name of a NetWorker volume pool for the transaction logs.
NSR_RECOVER_LOG_POOL	<p>Specifies the name of the NetWorker volume pool from which to restore the transaction log save sets during an SAP IQ point-in-time restore.</p> <p>Optional for an SAP IQ point-in-time restore.</p>	<ul style="list-style-type: none"> Most appropriate pool as selected by the NetWorker server (default). Valid name of a NetWorker volume pool for recovery of the transaction log save sets.
NSR_RECOVER_TIME	<p>Specifies the time for either an SAP IQ regular restore or point-in-time restore:</p> <ul style="list-style-type: none"> For an SAP IQ regular restore, the time to which the <code>nsriqrc</code> program restores the backup data. For an SAP IQ point-in-time restore, the transaction log save set time to which the NetWorker <code>recover</code> program restores the backup data and logs. The selected save set must have the transaction for the offset specified in <code>IQ_PIT_RESTORE_OFFSET</code>. <p>Optional for a regular SAP IQ restore.</p> <p>Mandatory for an SAP IQ point-in-time restore.</p>	<ul style="list-style-type: none"> Current time (default). Valid date and time in the <code>nsr_getdate(3)</code> format.
NSR_SAVESET_NAME	<p>Specifies the name of the save set to back up or restore.</p> <p>Mandatory for an SAP IQ manual backup or restore.</p>	<ul style="list-style-type: none"> Undefined (default). Valid save set name for the manual backup or restore in the following format: <code>NSR_SAVESET_NAME=IQ:/<i>database_name</i></code>
PATH	<p>Specifies the pathname of the directory that contains the NetWorker binaries on the Linux system.</p> <p>Set on Linux only if you relocated the NetWorker client binaries.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathname of the directory where the NetWorker binaries are installed.

Table 46 NMDA SAP IQ parameters (continued)

Parameter	Description	Default and valid values
SYBASE	<p>Specifies the pathname of the directory where the SAP IQ software is installed.</p> <p>Mandatory for an SAP IQ backup or restore. Set the parameter with the required method:</p> <ul style="list-style-type: none"> For a backup, set the parameter in the configuration file. For a restore, set the parameter in the environment. 	<ul style="list-style-type: none"> Undefined (default). Valid directory pathname for the SAP IQ software installation.
USER_PSWD	<p>Specifies the encrypted password for the SAP IQ user that connects to the SAP IQ server, as specified by the IQ_USER parameter.</p> <p>The encrypted password is added as the USER_PSWD setting to the NMDA configuration file.</p> <p>Mandatory for an SAP IQ backup or restore if the SAP IQ server has a password.</p> <p>Note: For an offline restore, you must set both the UTILDB_PSWD and USER_PSWD passwords by running the <code>nsrdaadmin -P iq -z configuration_file_path</code> command. This command prompts first for the UTILDB_PSWD password and then for the USER_PSWD password.</p>	<ul style="list-style-type: none"> Undefined (default). Encrypted SAP IQ user password, which you must set by running the appropriate <code>nsrdaadmin</code> command: <ul style="list-style-type: none"> For all operations except an offline restore, run the following command: <pre>nsrdaadmin -P -z configuration_file_path</pre> For an offline restore, run the following command: <pre>nsrdaadmin -P iq -z configuration_file_path</pre> <p>The <i>NetWorker Module for Databases and Applications Command Reference Guide</i> describes the <code>nsrdaadmin</code> command.</p>
UTILDB_USER	<p>Specifies the name of the SAP IQ user for login to the <code>utility_db</code> database. You must specify the password of the user by setting the UTILDB_PSWD parameter.</p> <p>Mandatory for an offline SAP IQ restore, which accesses the <code>utility_db</code> database.</p>	<ul style="list-style-type: none"> Undefined (default). Valid name of the SAP IQ user for login to the <code>utility_db</code> database.
UTILDB_PSWD	<p>Specifies the encrypted password for the SAP IQ user that performs an offline restore, as specified by the UTILDB_USER parameter.</p> <p>The encrypted password is added as the UTILDB_PSWD setting to the NMDA configuration file.</p>	<ul style="list-style-type: none"> Undefined (default).

Table 46 NMDA SAP IQ parameters (continued)

Parameter	Description	Default and valid values
	<p>Mandatory for an offline SAP IQ restore, which accesses the <code>utility_db</code> database.</p> <p>Note: For an offline restore, you must set the <code>UTILDB_PSWD</code> password by running the <code>nsrdaadmin -P iq -z configuration_file_path</code> command, which also prompts you to set the <code>USER_PSWD</code> password as a second step.</p>	<ul style="list-style-type: none"> Encrypted SAP IQ user password, which you must set by running the <code>nsrdaadmin</code> command. For example: <pre>nsrdaadmin -P iq -z configuration_file_path</pre> <p>The <i>NetWorker Module for Databases and Applications Command Reference Guide</i> describes the <code>nsrdaadmin</code> command.</p>

NMDA Sybase parameters

You must complete the required parameter settings for the NMDA Sybase backup and restore operations.

For the following Sybase operations, set both the common parameters and the NMDA Sybase parameters in the NMDA configuration file:

- Sybase scheduled backups that are configured without the wizard (client-side configuration)
- Sybase manual backups
- Sybase restores and recovery

[Common NMDA parameters](#) on page 406 describes the common parameters.

The following table describes the NMDA Sybase parameters.

Table 47 NMDA Sybase parameters

Parameter	Description	Default and valid values
<code>DATABASE_ONLINE</code>	<p>Specifies whether to bring the Sybase databases back online after a recovery.</p> <p>Optional for a Sybase restore/recovery.</p>	<ul style="list-style-type: none"> TRUE (default) = Bring the Sybase databases online after recovery. FALSE = Do not bring the Sybase databases online after recovery.
<code>DBCCOPT</code>	<p>Specifies one or more <code>-o</code> options to pass to the <code>nsrsybcc</code> command, which performs a database consistency check for the Sybase backup.</p> <p>Optional for a Sybase backup.</p> <p>The <i>NetWorker Module for Databases and Applications Command Reference Guide</i> describes the <code>nsrsybcc</code> command.</p> <p>Note: If you set <code>USE_CONSISTENCY_CHECK</code> to TRUE and you do not set <code>DBCCOPT</code>,</p>	<ul style="list-style-type: none"> Undefined (default). One or more of the following options of the <code>nsrsybcc</code> command, with a single option or multiple space-separated options enclosed in double quotes: <ul style="list-style-type: none"> <code>-o ckdb</code> <code>-o ckal</code> <code>-o ckat</code>

Table 47 NMDA Sybase parameters (continued)

Parameter	Description	Default and valid values
	<p><code>nsrsybcc</code> performs all the possible checks before a backup.</p>	<p>-o ckdbnoidx</p> <p>-o ckstor</p> <p>For example:</p> <p>DBCCOPT="-o ckcat -o ckal -o ckdb"</p>
LD_LIBRARY_PATH	<p>Specifies the directory pathname of the Open Client Server (OCS) library.</p> <p>Mandatory for a Sybase scheduled backup on the following platforms:</p> <ul style="list-style-type: none"> HP-UX Itanium Linux AMD64/EM64T Solaris SPARC <p>Optional for a Sybase manual backup on these platforms.</p> <p>For a Sybase manual backup, you can alternately set this parameter as an environment variable.</p>	<ul style="list-style-type: none"> 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	<p>Specifies the directory pathname of the Open Client Server (OCS) library.</p> <p>Mandatory for a Sybase scheduled backup on Solaris AMD64/EM64T.</p> <p>Optional for a Sybase manual backup on Solaris AMD64/EM64T.</p> <p>For a Sybase manual backup, you can alternately set this parameter as an environment variable.</p>	<ul style="list-style-type: none"> 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	<p>Specifies the directory pathname of the Open Client Server (OCS) library.</p> <p>Mandatory for a Sybase scheduled backup on AIX.</p> <p>Optional for a Sybase manual backup on AIX.</p> <p>For a Sybase manual backup, you can alternately set this parameter as an environment variable.</p>	<ul style="list-style-type: none"> 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	<p>Specifies an unencrypted password to use for a password-protected Sybase backup or the restore/recovery of such a backup. The password is added to the Sybase <code>dump</code> command for password-protecting the</p>	<ul style="list-style-type: none"> Undefined (default). Unencrypted password, from 6 to 30 characters in length, to add with the

Table 47 NMDA Sybase parameters (continued)

Parameter	Description	Default and valid values
	<p>backup, or added to the Sybase <code>load</code> command for restore/recovery of the backup.</p> <p>Optional for a Sybase backup or restore/recovery.</p> <p>Note: The NMDA configuration file contains the unencrypted password value. This parameter is different than the <code>USER_PSWD</code> parameter, which specifies an encrypted password for the Sybase user.</p>	<p><code>passwd=</code> clause to the Sybase <code>dump</code> or <code>load</code> command.</p>
NSR_ASE_VERIFY	<p>Specifies one of the following options for backup or restore verification:</p> <ul style="list-style-type: none"> • full • header • verifyonly <p>Optional for a Sybase backup or restore/recovery.</p>	<ul style="list-style-type: none"> • Undefined (default). • One of the following values for backup or restore verification: <ul style="list-style-type: none"> ▪ full = Verify both the header information and rows structure for full verification of a backup. ▪ header = Verify the page header information only for a backup. ▪ verifyonly = Verify the backed-up databases by performing a consistency check with the <code>nsrsybcc</code> command, without performing a restore/recovery. <p>For example, the following setting specifies a full verification of the backup:</p> <p>NSR_ASE_VERIFY=full</p>
NSR_BACKUP_LEVEL	<p>Specifies the level of Sybase manual backup to perform.</p> <p>Optional for a Sybase manual backup.</p> <p>Note: Do not set this parameter for a scheduled backup.</p>	<ul style="list-style-type: none"> • full (default) = Perform a full backup, which backs up the database. • cumulative or 1 = Perform a cumulative backup, which backs up the database data and log changes since the last full backup. <p>NMDA supports this backup level with Sybase ASE 15.7 SP 100 or later.</p> <ul style="list-style-type: none"> • incr = Perform an incremental level backup, which backs up the transaction logs. <p>Note: A whole instance incremental backup skips the backup of any database that does not support</p>

Table 47 NMDA Sybase parameters (continued)

Parameter	Description	Default and valid values
		<p>incremental backups (for example, when the database data and transaction logs are on the same device).</p>
NSR_BACKUP_PATHS	<p>Specifies the backup or restore/recovery of either the entire Sybase server or one or more Sybase databases.</p> <p>Mandatory for a Sybase manual backup or restore/recovery.</p> <p>Note: Do not specify both a server name and a list of databases. The master database must be restored separately before the server instance is restored.</p>	<ul style="list-style-type: none"> Undefined (default). Valid pathnames in either of the following forms, with multiple database names separated by a comma: SYBASE:/ASE_server_name (backs up or restores the entire server) SYBASE:/ASE_server_name/database_name1,SYBASE:/ASE_server_name/database_name2...]
NSR_CONCURRENCY_MODE	<p>Specifies the method that NMDA uses to back up or restore a Sybase ASE instance that contains many databases.</p> <p>The behavior of the instance backup or restore depends on the parallelism setting, which is the minimum of the NSR_PARALLELISM, client parallelism, and server parallelism settings.</p> <p>Optional for a Sybase ASE instance backup or restore that contains many databases.</p>	<ul style="list-style-type: none"> stripe (default) = During an instance backup, back up only one database at a time. The parallelism setting determines the number of stripes that NMDA uses for the database backup. During an instance restore, restore only one database at a time. database = During an instance backup, back up multiple databases at the same time and use one stripe per database. The parallelism setting determines the maximum number of databases that NMDA backs up concurrently. During an instance restore, restore multiple databases concurrently. The parallelism setting determines the maximum number of databases that NMDA restores concurrently. <p>Note: It is recommended that you use the "database" value only when an instance has a moderate number of databases (such as 50) of roughly similar sizes.</p> <p>If the parallelism setting is too low for a single database to be restored, NMDA restores only one database at a time.</p>

Table 47 NMDA Sybase parameters (continued)



Parameter	Description	Default and valid values
NSR_DUMP_DATA_OPT	<p>Specifies options for the backup of a Sybase database.</p> <p>Optional for a Sybase backup.</p>	<ul style="list-style-type: none"> 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	<p>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.</p> <p>Optional for a Sybase backup. Sybase transaction log backups in an emergency on page 149 provides details on using the options for transaction log backups.</p> <p> Note: Set this parameter only when there is an emergency during the transaction log backup.</p>	<ul style="list-style-type: none"> Undefined (default). One of the following values: <ul style="list-style-type: none"> “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. <p>For example, the following setting specifies to back up the transaction log without truncation: NSR_DUMP_LOG_OPT=“no_truncate”</p>
NSR_EXCLUDE_FILE	<p>Specifies the complete pathname of a file that lists databases to exclude from a backup of a Sybase server.</p> <p>Optional for a Sybase backup of a server.</p> <p> Note: Do not specify this parameter for a Sybase backup of one or more databases.</p>	<ul style="list-style-type: none"> 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
NSR_LOCALE	<p>Specifies the locale to use to connect to the Sybase server for a Sybase client-side scheduled backup or a manual backup configuration.</p> <p>Optional for a Sybase client-side scheduled backup or manual backup.</p>	<ul style="list-style-type: none"> Undefined (default). Valid locale identifier.

Table 47 NMDA Sybase parameters (continued)

Parameter	Description	Default and valid values
	<p>i Note: This parameter is not supported for a server-side scheduled backup or probe-based backup.</p>	
NSR_LOG_VOLUME_POOL	<p>Specifies the volume pool to use for a backup of the transaction logs.</p> <p>Optional for a Sybase backup.</p> <p>i 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 148 provides details. If required, specify the associated storage node by setting the Storage Nodes attribute in the NetWorker Client resource (NMC diagnostic mode must be enabled).</p>	<ul style="list-style-type: none"> • Most appropriate pool as selected by the NetWorker server (default). • Valid name of a NetWorker volume pool for the transaction logs.
NSR_PARALLELISM	<p>Specifies the maximum number of NetWorker sessions that can be opened at a time. Each stripe of a database requires a NetWorker session to be opened for a backup or restore.</p> <p>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.</p> <p>Optional for a Sybase backup or restore.</p> <p>i Note: During a restore, the <code>nsrsybrc</code> program can open more than the NSR_PARALLELISM number of NetWorker sessions, although the program cannot restore databases concurrently during this time.</p>	<ul style="list-style-type: none"> • 1 (default). • Integer value of 1 or greater. <p>For a scheduled backup, this parameter setting must be less than or equal to the Parallelism attribute value in the NetWorker Client resource. If multiple backups run concurrently on the client host, the total of all the NSR_PARALLELISM settings must be less than or equal to the Parallelism attribute value.</p>
NSR_PROMOTE_FULL	<p>Specifies whether to promote an incremental backup to a full backup when an incremental backup cannot be performed.</p> <p>i Note: This parameter only applies when the incremental backup is not a whole instance incremental backup. A whole instance incremental backup ignores this parameter and skips the backup of any database that does not support incremental backups.</p> <p>Optional for a Sybase backup.</p>	<ul style="list-style-type: none"> • 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 47 NMDA Sybase parameters (continued)

Parameter	Description	Default and valid values
NSR_RELOCATION_DEST	<p>Specifies an alternate Sybase server name and database name to use for the restore of database files during a redirected restore/recovery.</p> <p>The correct parameter setting depends on the NSR_BACKUP_PATHS setting:</p> <ul style="list-style-type: none"> • If NSR_BACKUP_PATHS specifies a server instance, then the NSR_RELOCATION_DEST value must also be a server instance. • If NSR_BACKUP_PATHS specifies multiple databases, then all the databases will be restored to the same database specified by NSR_RELOCATION_DEST. <p>Optional for a Sybase redirected restore/recovery.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid pathname in the following format: SYBASE:/ASE_server_name[/database_name] <p>Note: This parameter must not specify more than one pathname.</p>
NSR_XBSA_DEBUG	<p>Specifies whether debug messages from the NetWorker XBSA library are written to the NMDA Sybase log at the level set by NSR_DEBUG_LEVEL, described in NSR_DEBUG_LEVEL.</p> <p>Note: Use this parameter for debugging purposes with assistance from Customer Support only.</p> <p>Optional for a Sybase backup or restore/recovery.</p>	<ul style="list-style-type: none"> • FALSE (default) = XBSA library debug messages are not written to the NMDA Sybase log. • TRUE = XBSA library debug messages are written to the NMDA Sybase log.
PATH	<p>Specifies one of the following values:</p> <ul style="list-style-type: none"> • 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. <p>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.</p> <p>Set on UNIX only if you relocated the NetWorker client binaries.</p>	<ul style="list-style-type: none"> • Undefined (default). • Valid directory pathnames as follows: <ul style="list-style-type: none"> ▪ 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. <p>For example, set the parameter as follows to add the OCS library pathnames on Windows:</p> <pre>PATH=%PATH%;%SYBASE%\%SYBASE_OCS%\bin;%SYBASE%\%SYBASE_OCS%\dll</pre>
RECOVER_LISTONLY	<p>Specifies to display the structure information from the latest backup of a</p>	<ul style="list-style-type: none"> • Undefined (default).

Table 47 NMDA Sybase parameters (continued)


Parameter	Description	Default and valid values
	<p>specified Sybase instance or databases, without restoring the backup. The <code>nsrsybrc</code> program displays the information by using the <code>listonly</code> option of the Sybase <code>load</code> command.</p> <p>The <code>NSR_BACKUP_PATHS</code> parameter must specify the Sybase instance, database, or multiple databases for which the structure information will be displayed.</p> <p> Note: This parameter is supported with Sybase ASE 15.7 ESD#2 or later.</p> <p>Optional to prepare for a Sybase restore/recovery.</p>	<ul style="list-style-type: none"> • <code>load_sql</code> = Displays a sequence of <code>load database</code> or <code>load transaction</code> SQL commands, which can be used to restore the latest Sybase backup. • <code>create_sql</code> = Displays a sequence of <code>disk init</code>, <code>sp_cacheconfig</code>, <code>create database</code>, or <code>alter database</code> commands that are obtained from the latest <code>dump</code> image in the Sybase backup history.
RECOVER_SAVETIME	<p>Specifies whether to restore a single backup with the exact savetime in the <code>RECOVER_UNTIL</code> setting.</p> <p>This parameter enables you to restore a logs only backup, which you use to roll forward the offline database to the savetime that is specified by <code>RECOVER_UNTIL</code>.</p> <p>If a backup with the exact savetime in the <code>RECOVER_UNTIL</code> setting is not found, the restore fails. If <code>RECOVER_UNTIL</code> is set to an invalid time, then the <code>RECOVER_SAVETIME</code> parameter is ignored and a regular recovery is performed.</p>	<ul style="list-style-type: none"> • <code>FALSE</code> (default) = Perform a regular restore and recovery of Sybase data and logs as specified by the parameter settings. Do not restore only a single backup with the savetime as specified by <code>RECOVER_UNTIL</code>. • <code>TRUE</code> = Restore only a single backup with the savetime as specified by <code>RECOVER_UNTIL</code>.
RECOVER_UNTIL	<p>Specifies one of the following values for a Sybase recovery:</p> <ul style="list-style-type: none"> • The "until" date and time to which the backup is recovered. The backup is recovered as close as possible to this point-in-time, but no later. • Either the <code>up_to_min</code> or <code>up_to_min_noskip</code> value for an up-to-the-minute recovery as described in Sybase up-to-the-minute recovery on page 62. <p>Optional for a Sybase restore/recovery.</p>	<ul style="list-style-type: none"> • Date and time of most recent available backup (default). • Valid backup date and time in <code>nsr_getdate(3)</code> format. • <code>up_to_min</code> = Performs a logtail backup, then restores the latest full backup and all corresponding incremental backups, including the logtail backup, to the specified databases. Skips all the restores if the logtail backup fails. • <code>up_to_min_noskip</code> = Performs a logtail backup, then restores the latest full backup and all corresponding incremental backups, including the logtail backup, to the specified

Table 47 NMDA Sybase parameters (continued)

Parameter	Description	Default and valid values
		databases. Performs all the restores even if the logtail backup fails.
STRIPE_DEVICE_DEST	Specifies a shorter alternative directory pathname to use in the temporary full pathname for the device during a Sybase restore. Specify the shorter alternative if the device name is longer than the maximum Sybase limit, for example, 127 bytes. Optional for a Sybase restore/recovery.	<ul style="list-style-type: none"> Undefined (default). Valid full directory pathname to use in the temporary device pathname for the Sybase restore. For example: STRIPE_DEVICE_DEST=C:\tmp
SYBASE	<p>Specifies the pathname of the directory where you installed the Sybase ASE software.</p> <p>Mandatory for a Sybase backup or restore/recovery. Set the parameter with the required method:</p> <ul style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> Undefined (default). Valid directory pathname for the Sybase ASE software installation.
SYBASE_USER	<p>Specifies the name of the Sybase user that connects to the Sybase server for the backup or restore. You must specify the password of the user by setting the USER_PSWD parameter.</p> <p>Mandatory for a Sybase backup or restore/recovery.</p>	<ul style="list-style-type: none"> Undefined (default). Valid Sybase username.
USE_CONSISTENCY_CHECK	<p>Specifies whether the <code>nsrsybcc</code> command performs a database consistency check before a backup occurs or after a restore/recovery completes.</p> <p>Optional for a Sybase scheduled backup or restore/recovery.</p>	<ul style="list-style-type: none"> FALSE (default) = Do not run the database consistency check command, <code>nsrsybcc</code>, before the backup or after the restore/recovery. TRUE = Run the database consistency check command, <code>nsrsybcc</code>, before the backup or after the restore/recovery. The specific checks that are performed depend on the DBCCOPT setting.
USE_CUMULATIVE	Specifies whether the <code>nsrsybrc</code> command restores Sybase cumulative backups in addition to full backups and incremental (transaction log) backups.	<ul style="list-style-type: none"> TRUE (default) = Restore any available cumulative backup in addition


Table 47 NMDA Sybase parameters (continued)

Parameter	Description	Default and valid values
	Optional for a Sybase restore/recovery.	<p>to full backups and incremental backups.</p> <ul style="list-style-type: none"> FALSE = Restore only full backups and incremental backups. Do not restore any cumulative backups.
USER_PSWD	<p>Specifies the encrypted password for the Sybase user that connects to the Sybase server, as specified by the SYBASE_USER parameter.</p> <p>The encrypted password is added as the USER_PSWD setting to the NMDA configuration file.</p> <p>Mandatory for a Sybase backup or restore/recovery if the Sybase server has a password.</p>	<ul style="list-style-type: none"> Undefined (default). Encrypted Sybase user password, which you must set by using the <code>nsrdaadmin -P</code> command. For example: <pre data-bbox="927 741 1334 821" style="background-color: #f0f0f0; padding: 5px;">nsrdaadmin -P -z configuration_file_path</pre> <p>The <i>NetWorker Module for Databases and Applications Command Reference Guide</i> describes the <code>nsrdaadmin</code> command.</p>

APPENDIX B

Oracle RMAN Commands

This appendix includes the following topics:

 **Note:** These topics describe Oracle RMAN commands that you can use in the RMAN scripts for Oracle backups and restores.

- [The pool option of the backup command](#)..... 476
- [The send command](#)..... 476

The pool option of the backup command

i **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 517 describes this error message.

Use the appropriate method to specify the NetWorker volume pool that NMDA will use:

- For a scheduled Oracle backup, specify the pool in the configured backup action or Client resource.
- For a manual Oracle backup, set the `NSR_DATA_VOLUME_POOL` parameter in the RMAN script.

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 that are described in [NMDA Oracle parameters](#) on page 444. 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 476
- [Two ways to run the send command](#) on page 478
- [Precedence rules](#) on page 480

i **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 477
- [The send command options](#) on page 477

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

You must run the `send` command with the correct options and option values to ensure the proper settings of the NMDA parameters.

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 that are allocated for NMDA.
- `send channel` sets the parameters for the specified channels only.

NOTICE You can use the `device_type` option 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 479 provides details.

Example 55 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 preceding RMAN script, listing the values for options that the `send` command uses.

Table 48 Option values in the send command

Option value	Describes
<i>device_specifier</i>	The device type as specified in an <code>allocate channel</code> command in the RMAN script. For a backup tape device, use SBT or SBT_TAPE.
<i>channel_id</i>	The channel identifier as specified in an <code>allocate channel</code> command in the RMAN script. In the example, the identifier is t1.
<i>name1</i>	The first NMDA parameter name. In the example, the first parameter name is NSR_SERVER.
<i>value1</i>	The value that is assigned to the first parameter. In the example, the first value is server1.
<i>name2</i>	The second NMDA parameter name. In the example, the second parameter name is NSR_DATA_VOLUME_POOL.
<i>value2</i>	The value that is 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 479.
- In the `run` job of the RMAN script, as described in [The send command in the RMAN script](#) on page 480.

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 that are listed in [The send command string](#) on page 477, 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 477 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 trying 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 480 provides details.

Example 56 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 that are listed in [The send command string](#) on page 477.
- 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 477.
- Specify the correct option values in the `send` command, as described in [The send command options](#) on page 477.
- 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 in [Example 6](#). The single `send` command sets the parameters for both channels.

Precedence rules

Parameters are set for channels that are 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 57 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](#).....484
- [NMDA error messages](#).....487
- [DB2 troubleshooting tips and error messages](#)..... 492
- [Informix troubleshooting tips and error messages](#)..... 502
- [Lotus troubleshooting tips and error messages](#)..... 504
- [MySQL troubleshooting tips and error messages](#).....510
- [Oracle troubleshooting tips and error messages](#)..... 513
- [Orchestrated Application Protection troubleshooting tips](#)..... 531
- [SAP IQ troubleshooting tips and error messages](#).....532
- [Sybase troubleshooting tips and error messages](#).....533

General troubleshooting tips

Use the following list to troubleshoot basic problems in running backup and restore operations with NMDA.

About this task

To set up an NMDA backup and restore system, perform the following steps.

Procedure

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:

- *NetWorker Module for Databases and Applications Installation Guide*
- *NetWorker E-LAB Navigator*
- *NetWorker Installation Guide*
- *NetWorker Administration Guide*
- [Backup Configuration](#) on page 75

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:

- a. Open Windows Explorer, and locate the `nsrdasv.exe` file in the `NetWorker_install_path\bin` directory.
- b. Right-click the file icon, and select **Properties**.
- c. 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, type the correct 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 168.
4. Ensure that you can perform a scheduled backup with NMDA according to the instructions in [Performing scheduled backups](#) on page 166.

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
- NSR_DIAGNOSTIC_DEST
- NSR_DPRINTF

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`

i 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 the following file name format:

```
binaryname[_app].date.time.pid[_threadid].log
```

where *app* indicates the type of application:

- db2
- informix
- IQ
- lotus
- mysql
- oracle
- sybase

When the NSR_DPRINTF parameter is set to TRUE, NetWorker core debug messages are written to the NMDA debug log files. NSR_DPRINTF applies only if NSR_DEBUG_LEVEL is set to a value greater than 0.

The NetWorker server also writes diagnostic information from a manual or scheduled backup to specific log files on the NetWorker server. The *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, complete the following steps:

1. Use the NMC program to ensure that the Lockbox resource is created for the particular client and the Users attribute contains the superuser of the client. The *NetWorker Administration Guide* describes the lockbox password management.
2. Delete the NSR Peer Information resource for the NetWorker client from the NSRLA database on the NetWorker server. This deletion causes the NetWorker server to create a new NSR Peer Information resource for the client. The topic "Deleting the NSR Peer Information resource" in the *NetWorker Security Configuration Guide* provides details.

To check for invalid NSR Peer Information resources, you can run the following command:

```
nsradmin -s server_name -p nsrexecd -C "NSR peer information"
```

To automatically correct the resource errors, you can run the `nsradmin` command with the `-y` option. The *NetWorker Command Reference Guide* provides details on the command and options.

Authentication error, unable to obtain the user credentials from the lockbox

After you upgrade or reinstall the NetWorker server, you might encounter an authentication error. This type of error causes a backup to fail because the authentication information in the NSRLA database on the NetWorker server does not match the client credentials.

The following error messages might appear:

```
Authentication error; why = Server rejected credential
Unable to obtain the user credentials from the lockbox due to the following
error: 'error_message'
Unable to retrieve the password using lockbox_name for client client_name.
Failed to retrieve password from lockbox.
Password could not be retrieved from lockbox.
```

As a solution, delete the NSR Peer Information resource for the NetWorker client from the NSRLA database on the NetWorker server. This deletion causes the NetWorker server to create a new NSR Peer Information resource for the client. The topic "Deleting the NSR Peer Information resource" in the *NetWorker Security Configuration Guide* provides more details.

To check for invalid NSR Peer Information resources, you can run the following command:

```
nsradmin -s server_name -p nsrexecd -C "NSR peer information"
```

To automatically correct the resource errors, you can run the `nsradmin` command with the `-y` option. The *NetWorker Command Reference Guide* provides details on the command and options.

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 the `NSR_MAX_START_RETRIES` parameter to an appropriate value as described in [NSR_MAX_START_RETRIES](#).

After you set the `NSR_MAX_START_RETRIES` parameter, if 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](#) 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 in the backups and restores. The error log file has the following file name for all applications:

```
nmda_app.messages.raw
```

where *app* indicates the type of application:

- db2
- informix
- lotus
- mysql
- oapp
- oracle
- sybase

The oapp value indicates Orchestrated Application Protection, as supported for MongoDB, MySQL, and PostgreSQL databases.

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 *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 that are 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 49 NetWorker XBSA error messages

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	An API call sequence occurred that does not conform to the XBSA Data Movement API State Diagram document.

Table 49 NetWorker XBSA error messages (continued)

Error message	Description
<p>The sequence of API calls is incorrect. Must call item1 before item2</p>	
<p>BSA_RC_BAD_HANDLE The handle used to associate this call with a previous BSAInit() call is invalid because explanation</p>	<p>The value that was passed into the function for bsaHandle contained a NULL pointer.</p>
<p>BSA_RC_BAD_PARAMETER received parameter parm with value value, which is invalid</p>	<p>The software received an invalid parameter.</p>
<p>BSA_RC_BUFFER_TOO_SMALL Buffer is too small to hold the object entry to be returned. n bytes required for the object entry</p>	<p>The buffer is too small to hold the object entry.</p>
<p>BSA_RC_DESCRIPTION_TOO_LONG The description field contained too many characters (n >= n)</p>	<p>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.</p>
<p>BSA_RC_INVALID_COPYTYPE the copyType field contained an unrecognized value of n</p>	<p>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.</p>
<p>BSA_RC_INVALID_DATABLOCK the dataBlock parameter contained inconsistent values: bufferLength: n, bufferPtr: n, numBytes: n</p>	<p>The fields of a supplied DataBlock parameter are not internally consistent. This can occur under one of the following conditions:</p> <ul style="list-style-type: none"> • 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.
<p>BSA_RC_INVALID_KEYWORD an entry in the environment structure is invalid (variable=value)</p>	<p>One of the environment strings that was 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:</p> <ul style="list-style-type: none"> • 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.

Table 49 NetWorker XBSA error messages (continued)

Error message	Description
	<ul style="list-style-type: none"> The VALUE string was invalid. The VALUE string could not be validated, for example, as in the case of a hostname string that the <code>gethostbyname()</code> function could not find.
<p>BSA_RC_INVALID_OBJECTSTATUS</p> <p>the object Status field contained an unrecognized value of n</p>	<p>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.</p>
<p>BSA_RC_INVALID_OBJECTTYPE</p> <p>the objectType is invalid (n)</p>	<p>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.</p>
<p>BSA_RC_INVALID_TIME</p> <p>a time field contained an unrecognized value of n</p>	<p>The software received an invalid time value.</p>
<p>BSA_RC_MATCH_EXISTS</p> <p>object matching the specified predicate already exists</p>	<p>The object exists in the NetWorker server that the NetWorker XBSA session uses, and the requested operation cannot be completed.</p>
<p>BSA_RC_MORE_DATA</p> <p>more data is available. Data can be obtained through BSAGetData() or BSAGetNextQueryObject()</p>	<p>This error has two meanings in the XBSA Data Movement API:</p> <ul style="list-style-type: none"> 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 and BSA_RC_NO_MORE_DATA. 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.
<p>BSA_RC_NO_MATCH</p> <p>The ResourceType predicate value of D does not match the reference value of L</p>	<p>The client index and media database are out of synch. To resynchronize the client index and media database, run the <code>nsrck -X</code> command. Alternatively, wait for NetWorker to run <code>nsrck</code> automatically.</p>
<p>BSA_RC_NO_MATCH</p> <p>The variable predicate value of value does not match the reference value of variable</p>	<p>No objects matching the specified QueryDescriptor were found in the NetWorker server that the NetWorker XBSA session uses.</p>
<p>BSA_RC_NO_MORE_DATA</p>	<p>This error has two meanings in the XBSA Data Movement API:</p>

Table 49 NetWorker XBSA error messages (continued)

Error message	Description
there is no more data for the current object	<ul style="list-style-type: none"> Object Data Retrieval—This error is used when all the data for an object being retrieved from a NetWorker server was placed into the specified DataBlock parameter for a function call. Query Result Retrieval—This error is used when the last (or only) object matching a query is returned to the caller.
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 .
BSA_RC_NULL_OBJECTNAME an object name is required	The ObjectName parameter that was 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 object Info field contained too many characters	The ObjectInfo parameter that was passed into the function, either directly or in one of the following data structures, was found to have more than BSA_MAX_OBINFO 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 path Name field contained too many characters	The string pathname contains more than BSA_MAX_PATHNAME characters in an ObjectName structure.
BSA_RC_RESOURCECTYPE_TOO_LONG The resourceType field contained too many characters (n >= n)	The string resourceType contains more than BSA_MAX_RESOURCECTYPE 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.

DB2 rollforward might fail on Windows with DB2 9.7 due to the logarchopt2 value

On Windows with DB2 9.7, due to a DB2 limitation, a DB2 rollforward operation might crash the DB2 instance if the log query or retrieval uses logarchmeth2 with the following settings:

- The logarchmeth2 setting specifies the vendor archive method.
- The logarchopt2 setting exceeds 50 characters in length.

The rollforward operation failure produces the following type of error message:

```
C:\Program Files\IBM\SQLLIB\BIN> db2 rollforward db sample to end of logs
```

```
SQL1224N The database manager is not able to accept new requests, has terminated all requests in progress, or has terminated the specified request because of an error or a forced interrupt. SQLSTATE=55032
```

In this case, restart the DB2 instance and reconfigure logarchopt2 to point to a value of 30 characters or less, as required by the IBM standard. Then restart the rollforward operation.

Debug log issue with load command and copy yes option

Due to a DB2 limitation, the NMDA debug log contains the erroneous message, “num_sessions = 0”, when you run the `load` command with the `copy yes` option.

This issue is resolved in FixPack 5 for DB2 version 10.1 and 10.5.

Deletion of DB2 snapshot save sets might fail with an error

In a `db2acsutil` deletion operation with a DB2 timestamp, if there is an error in searching for the snapshot save sets to delete, as when the save sets have already expired and been deleted, then the `db2acsutil` command returns an error.

If the `DB2_ACS_METADATA_DELETION_FORCE` parameter is not set to `TRUE`, the `db2acsutil` command also does not delete the corresponding metadata save sets. To enforce the metadata cleanup in such a case, you can set `DB2_ACS_METADATA_DELETION_FORCE` to `TRUE` and rerun the `db2acsutil` operation. The operation will still return an error, but then running a query with the DB2 timestamp will produce the expected result of no matched records.

Delta or incremental backup failure in a DB2 pureScale system

This topic applies only if the DB2 version does not support delta and incremental backups in a DB2 pureScale system.

If you try 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 that does not support this type of backup, then 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.
```

DB2 supports delta and incremental backups in a pureScale system starting with DB2 v10.5 fp4. You must perform a DB2 full backup in a pureScale system that does not support delta and incremental backups.

Delta or incremental restore failure

A restore from a DB2 incremental or delta backup might fail with an “out of order image” error. For example, the restore might fail with the following error message:

```
SQL2572N  Attempted an incremental restore of an out of order image.
```

This error might occur when the DB2 history file is damaged or missing, or when the backup is being restored to a new instance or host that does not yet have the correct history file.

You can run the following command to check if the history file is correct and complete:

```
db2 list history backup all for sample
```

where *sample* is the name of the database to be restored.

If the history file is damaged or missing, run the appropriate command to restore the history file:

- On UNIX:

```
db2 restore database sample history file load /usr/lib/libnsrdb2.so
options @pathname/nmda_db2.cfg
```

- On Windows:

```
db2 restore database sample history file load NetWorker_install_dir\nsr
\bin\libnsrdb2.dll options @pathname\nmda_db2.cfg taken at yyyymmddhhmmss
```

where:

- *sample* is the name of the database to be restored.
- *pathname*/nmda_db2.cfg or *pathname*\nmda_db2.cfg is the complete pathname of the NMDA configuration file.
- *yyyymmddhhmmss* is the date and time of the backup.

After you restore the history file, you can perform other restore operations.

Log retrieval issue for re-created database

Indexing of the archived log backups is based on the DB2 instance, the database name, and the log chain and sequence numbers. A database might have archived logs with the same index name from different lifespans, where a lifespan extends from when the database was created to when it was dropped. Not removing the old log backups can cause issues during a log retrieval, as in a rollforward operation or an online backup with the INCLUDE LOGS option.

The log retrieval operation might fail with the following error messages in the db2diag.log file:

```
Database ID does not match. Extent does not belong to this database.
Database ID does not match. Extent probably for another database.
```

To prevent this issue, remove the old log backups before you perform an archived log backup for a newly created database. Using a different log volume pool for the backup might not prevent the issue because the backup index search is not based on the pool setting.

The IBM documentation describes a similar requirement and similar issues:

- Requirement for an archived log path of a database:
The archived log path must not contain log files that do not belong to the current database. If the archived log path was previously used for a database of the same name, these log files must be removed before using the current archived log path. The following IBM article provides details:
http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.1.0/com.ibm.db2.luw.admin.config.doc/doc/r0011448.html?cp=SSEPGG_10.1.0%2F2-2-4-8-69
- Issue with a DB2 DISK archived log backup:
The following IBM article describes an issue with a DISK archived log backup that uses a previously used location:
<http://www-01.ibm.com/support/docview.wss?uid=swg21450419>

Pruning of DB2 log history after a rollforward might produce errors

If you perform a DB2 restore without the `replace history file` option and then perform a rollforward operation, the log history file might contain duplicated entries for the logs that are archived after the backup and used in the rollforward. A DB2 history pruning operation with `auto_del_rec_obj` set to `on` might then produce misleading errors. The errors occur when the DB2 software requests the removal of the same archived log backup multiple times due to the duplicated entries in the history.

Note: This issue does not occur when you perform the DB2 restore with the `replace history file` option.

For example, the `db2 prune history` operation produces the following message in the `db2diag.log` file:

```
2015-05-04-13.39.24.676564-240 E684889E515          LEVEL: Info
PID       : 1027                TID    : 47814394505536PROC : db2sysc 0
INSTANCE: db2inst1            NODE   : 000                DB     : TEST
APPHDL   : 0-5718             APPID: *LOCAL.db2inst1.150504173845
AUTHID   : DB2INST1
EDUID    : 2878                EDUNAME: db2agent (TEST) 0
FUNCTION: DB2 UDB, database utilities, sqluhDeletionReport, probe:381
MESSAGE  : ADM8506I Successfully deleted the following database logs "3 - 5"
in log chain "1".

2015-05-04-13.39.24.676841-240 E685405E502          LEVEL: Error
PID       : 1027                TID    : 47814394505536PROC : db2sysc 0
INSTANCE: db2inst1            NODE   : 000                DB     : TEST
APPHDL   : 0-5718             APPID: *LOCAL.db2inst1.150504173845
AUTHID   : DB2INST1
EDUID    : 2878                EDUNAME: db2agent (TEST) 0
FUNCTION: DB2 UDB, database utilities, sqluhDeletionReport, probe:387
MESSAGE  : ADM8509N Unable to delete the database logs "3 - 4" in log chain
"1".
```

To work around this issue, perform the following steps:

1. To verify that the log entries are removed, check the `db2diag.log` file and the backup index.
2. Clean up the DB2 history file by running a `db2 prune history...` with `force` option command without the `and delete` option.

Pruning of log backups fails due to the `logarchoptn` value

Due to a DB2 limitation, when the `logarchoptn` value exceeds 30 characters, the value is not recorded correctly in the DB2 recover history. In such cases, the pruning operation might encounter an issue. The output of the `db2 prune history and delete` command might indicate success but the deletion of log backups might have an issue. Errors might appear in the `db2diag.log` file, and the log entries might remain in both the DB2 database history and backup storage.

In this case, you can perform the following workaround.

1. To adjust the length of the parameter setting, run the following command:

```
db2 update history file EID entry-eid with comment
'new_location_of_logarchopt1'
```

For example, run the following command:

```
db2 update history EID 10 with comment '@C:\tmp\other.cfg'
```


2. Rerun the `db2 prune history and delete` command.

Sometimes, there might be multiple entries of the same archived log in the DB2 history as shown by the `db2 list history archive log` command. However, there might be only one corresponding entry in the backup storage. During the pruning of the log backup, DB2 might send more than one deletion request. In such cases, the entries are removed from the backup storage but might still appear in the DB2 archived log history after the prune operation. You can remove the entries from the DB2 history by performing a manual prune operation without the `delete` option.

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.

About this task

 **Note:** The NetWorker man pages, the *NetWorker Command Reference Guide*, and the *NetWorker Administration Guide* describe the NetWorker commands to remove a failed backup.

To remove a failed backup:

Procedure

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 that are required to perform the operation, as described in [Table 6](#) on page 83.

SOURCE_CLIENT prevents issue in rollforward with archived log backup

A DB2 archived log backup might take place in the same timeframe as a log restore during a rollforward operation. For example, to complete a rollforward, uncommitted transactions are rolled back. The rollback action starts the log archiving. In this case, both the log restore and log archiving use the same configuration file.

A redirect recovery to a different client should use two client parameters that refer to different hosts: the source client that has the log to roll forward and the target client that has the log to be backed up. If the recovery uses only one client parameter that points to the source client, the concurrent log backup is saved incorrectly under the source client name. Then a rollforward operation (when needed) of the source database might fail with the following error when the wrong log is retrieved:

```
Database ID does not match. Extent probably for another database.
```

The SOURCE_CLIENT parameter was introduced to prevent this issue by pointing to the source client that has the log backups that are used to roll forward. NSR_CLIENT is used to point to the target host under which the archived log backups of the rollback will be stored. If SOURCE_CLIENT is not specified, NSR_CLIENT is used for both the archived log backup and restore.


Wizard configuration or scheduled backup on AIX might fail due to DB2 10.5 and 10.1 issues

On AIX, a DB2 wizard configuration or scheduled backup might fail when NMDA tries to attach to the DB2 instance or connect to the database.

The problem occurs with certain Fix Packs of DB2 versions 10.5 and 10.1, and is due to a DB2 known issue as described at the following IBM reference site:

<http://www-01.ibm.com/support/docview.wss?crawler=1&uid=swg1IT05079>

```
IT05079: APPLICATION RETURNS -1042. DB2DIAG.LOG SHOWS "INVALID DATA VALUE:
ICCLIB.C:989"
Local fix:
* set ICC_IGNORE_FIPS=YES (please note: GSKit to run in non-FIPS mode)
```

 **Note:** The DB2 issue is expected to be fixed in the next DB2 Fix Pack release.

The failure produces the following type of stack trace in the db2diag.log file:

```

FUNCTION: DB2 Common, Cryptography, cryptContextRealInit, probe:60
MESSAGE : ECF=0x90000403=-1879047165=ECF_CRYPT_UNEXPECTED_ERROR
          Unexpected cryptographic error
DATA #1 : Hex integer, 4 bytes
0x00000002
DATA #2 : Hex integer, 4 bytes
0x00000002
DATA #3 : Hex integer, 4 bytes
0x00000000
DATA #4 : String, 32 bytes
Invalid data value: icclib.c:990
CALLSTCK: (Static functions may not be resolved correctly, as they are
resolved
to the nearest symbol)
 [0] 0x09000000180329F4 pdOSSeLoggingCallback + 0x184
 [1] 0x0900000001736588 oss_log__FP9OSSLogFacUiN32U1N26iPPc + 0x1C8
 [2] 0x0900000001736380 ossLog + 0xA0
 [3] 0x09000000171D8244 cryptLogICCErrWithStatus + 0xF8
 [4] 0x09000000171D97E8 cryptContextRealInit + 0x1454
 [5] 0x09000000171D7920 cryptContextCheckAndInit + 0x68
 [6] 0x09000000171C4B04 cryptDHInit + 0x1C4
 [7] 0x09000000171C4190 sqlxSlcServerEncryptAccsec + 0x1FC8
 [8] 0x09000000171C4CE0 sqlxAppAuthenticate__FP14db2UCinterface + 0x10C
 [9] 0x09000000180BCD1C sqlxAppAuthenticate__FP14db2UCinterface + 0x1314
 [10] 0x09000000180BB0A8 sqljrDrdaArAttach__FP14db2UCinterface + 0x84
 [11] 0x09000000180BAECC sqlUCdrdaARinit__FP14db2UCconHandle + 0x3D4
 [12] 0x090000001806FC60 sqlUCappAttach + 0xAE4
 [13] 0x0900000018077E4C sqleatin__FPcN41sP5sqlca + 0x238
 [14] 0x0900000018077390 sqleatcp_api + 0x1A4
 [15] 0x0900000018076FE4 sqleatin_api + 0x60
 [16] 0x0000000010047F924 InstAttach + 0xE4
 [17] 0x00000000100461714 validate_db2_credentials + 0x1D4
 [18] 0x00000000100463AA4 db2ra_verify_credentials + 0x64

```

To apply the workaround that is recommended by DB2, specify the environment setting “ICC_IGNORE_FIPS=yes” by setting the internal NMDA DB2 parameter DB2CFG_OPT_INTERNAL as follows:

DB2CFG_OPT_INTERNAL="ICC_IGNORE_FIPS=yes"

To apply the workaround in the NMDA DB2 backup configuration wizard:

1. Add “ICC_IGNORE_FIPS=yes” in the **Connection parameters** field on the **Specify the DB2 Information** page.
2. Add the parameter setting DB2CFG_OPT_INTERNAL="ICC_IGNORE_FIPS=yes" in the **Advanced Options** table to apply the workaround in the modification workflow.

To apply the workaround in the NMDA DB2 recovery wizard:

1. Add “ICC_IGNORE_FIPS=yes” in the **Connection parameters** field on the **Specify the DB2 Information** page.
2. Add the parameter setting DB2CFG_OPT_INTERNAL="ICC_IGNORE_FIPS=yes" in the **Advanced Options** table to apply the workaround in the modification workflow when you select the **Recover Again** option on the **Recover** window.

The specified environment setting will be used in the connection to DB2 during the NMDA operations. The specified setting will also be added in the wizard **Advanced Options** table, but you can remove the setting from the table if you do not want it to be used in the backup or recovery.

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 50 Path and suffix for the load libnsrdb2 command

Operating system	Path with suffix
AIX, Linux, Solaris, HP-UX Itanium	<code>/usr/lib/libnsrdb2.so</code>
Microsoft Windows	<code>NetWorker_install_path\bin \libnsrdb2.dll</code>

DB2 SQL 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 488 describes the error messages that are generated by the NetWorker XBSA interface.

Table 51 DB2 SQL error messages

Error message	Description
SQL1024N A database connection does not exist. SQLSTATE=08003	The command, for example, <code>db2 load ... copy yes</code> , requires a connection to the database. Issue a <code>CONNECT</code> command to connect to the database, and then reissue the command.
SQL1116N A connection to or activation of database "DB_NAME" failed because the database is in BACKUP PENDING state. SQLSTATE=57019	The <code>load</code> command with the <code>copy no</code> option was issued on a rollforward recoverable database. Back up the database to resolve this issue.
SQL1268N A rollforward recovery stopped due to error "SQL1042" while retrieving log file <i>logfile</i> for database <i>db</i> on node "0".	The <code>NSR_ENCRYPTION_PHRASES</code> parameter does not contain the datazone pass phrase that was used to back up the transaction logs. Set the <code>NSR_ENCRYPTION_PHRASES</code> parameter to the proper phrase. NSR_ENCRYPTION_PHRASES provides details.
SQL1776N	DB2 does not support a backup on an HADR standby node. Set the following parameter in the NMDA configuration file:

Table 51 DB2 SQL error messages (continued)

Error message	Description
The command cannot be issued on an HADR database. Reason code: "1".	DB2_SKIP_HADR_STANDBY = TRUE
SQL2025N An I/O error "3" occurred on media "VENDOR".	<p>The client is not registered on the NetWorker server to which the NMDA software is backing up. Create a valid client on the NetWorker server.</p> <p>Test to ensure that the connection between the client and server is valid:</p> <pre data-bbox="651 646 1329 695">save -s servername/testfile</pre>
SQL2025N An I/O error "3" occurred on media "VENDOR".	<p>The NSR_CLIENT parameter is set to an invalid client name while running a backup.</p> <p>Set the NSR_CLIENT parameter to the name of the client from which the backup is running. NSR_CLIENT provides details.</p>
SQL2025N An I/O error "3" occurred on media "VENDOR".	<p>In the vendor configuration file, the NSR_DATA_VOLUME_POOL parameter is set to a pool name that does not exist during a backup.</p> <p>If possible, remove the DB2 vendor.cfg file. NMDA DB2 parameters on page 414 provides details. Otherwise, do as follows:</p> <ul style="list-style-type: none"> • 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.
SQL2025N An I/O error "3" occurred on media "VENDOR".	<p>There is no NMDA license on the server. Each client requires a separate client license.</p> <p>Obtain a valid NMDA license.</p>
SQL2025N An I/O error "25" occurred on media "VENDOR".	<p>The NSR_ENCRYPTION_PHRASES parameter does not contain the datazone pass phrase that was used to back up the database.</p> <p>Set the NSR_ENCRYPTION_PHRASES parameter to the proper phrase. NSR_ENCRYPTION_PHRASES provides details.</p>
SQL2025N An I/O error "25" occurred on media "VENDOR".	<p>The user does not have the restore privilege on the NetWorker server. Add the "recover local data" privilege for the user.</p>
SQL2025N An I/O error occurred. Error code: "25". Media on which this error occurred: "VENDOR".	<p>The DB2 host has been renamed or has multiple names.</p>

Table 51 DB2 SQL error messages (continued)


Error message	Description
	Ensure that all the hostnames or IP addresses assigned to the host are included in the Aliases attribute in the NetWorker Client resource.
<p>SQL2036N</p> <p>The path for the file, named pipe, or device "FileName.dat" is not valid.</p>	<p>The file pathname or permissions of the data file are incorrect.</p> <p>Ensure the correct file pathname and permissions for the data file.</p>
<p>SQL2062N</p> <p>An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "0".</p>	<p>The file pathname or permissions are incorrect for the configuration file that is specified in the <code>vendoropt</code> option.</p> <p>Ensure the correct file pathname and permissions for the configuration file.</p> <p> Note: You must specify the full pathname of the configuration file.</p>
<p>SQL2062N</p> <p>An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "0".</p>	<p>Permissions or ownership of debug files are incorrect for the database instance.</p> <p>Ensure that each database instance has a unique debug file name.</p>
<p>SQL2062N</p> <p>An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "0".</p>	<p>The options file pathname is wrong or a relative pathname is used for the options file in the DB2 backup command.</p> <p>Ensure the correct file pathname of the options file. You must specify the full pathname in the backup command.</p>
<p>SQL2062N</p> <p>An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "4".</p> <p>Table 50 on page 498 provides the correct libnsrdb2 path and suffix information.</p>	<p>The NSR_CLIENT parameter value is an incorrect client name for the restore.</p> <p>Set the NSR_CLIENT parameter to the name of the client from which the restore is running.</p> <p>NSR_CLIENT provides details.</p>
<p>SQL2062N</p> <p>An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "4".</p> <p>Table 50 on page 498 provides the correct libnsrdb2 path and suffix information.</p>	<p>You have not registered the client on the NetWorker server to which the module is restoring:</p> <ol style="list-style-type: none"> 1. Create a valid client on the NetWorker server. 2. Test to ensure that the connection between the client and server is valid: <pre style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;">save -s servername/testfile</pre>
<p>SQL2062N</p> <p>An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "11".</p>	<p>You did not specify a valid timestamp for the object being restored. Specify a valid timestamp.</p> <p>Or:</p>

Table 51 DB2 SQL error messages (continued)

Error message	Description
	<p>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:</p> <pre data-bbox="651 464 1153 491">DB2_R=database_name: db2inst1:db2inst2:</pre> <p>Performing DB2 data restores with the db2 restore command on page 194 provides details.</p>
<p>SQL2062N</p> <p>An error occurred while accessing media "/usr/lib/libnsrdb2.so". Reason code: "11". Table 50 on page 498 provides the correct libnsrdb2 path and suffix information.</p>	<p>The NSR_SERVER parameter value is an invalid server name. For example, the server might not exist and you cannot ping it.</p> <p>Set the NSR_SERVER parameter to a valid NetWorker server that has the DB2 server defined as a client. NSR_SERVER provides details.</p>
<p>SQL2062N</p> <p>An error occurred while accessing media "libnsrdb2.so". Reason code: "25".</p>	<p>The backup of the <i>database_name</i> database failed due to a backup request failure. The BRC API call pb_open failed. Only serverless backups support a deduplication backup.</p> <p>Perform a serverless type of backup for the NSM deduplication backup of the database.</p>
<p>SQL2071N</p> <p>An error occurred while accessing the shared library "c:\progra~1\legato\nsr\bin\libnsrdb2.dll". Reason code: "1". or An error occurred while accessing the shared library "/usr/lib/libnsrdb2.so". Reason code: "1".</p>	<p>The missing NMDA DB2 library, libnsrdb2.xx, is not in the correct place as indicated or is not accessible.</p> <p>Reinstall NMDA or place the libnsrdb2.xx library into the correct location as indicated.</p>
<p>SQL2071N</p> <p>An error occurred while accessing the shared library "/usr/lib/libnsrdb2.so". Reason code: "2". Table 50 on page 498 provides the correct libnsrdb2 path and suffix information.</p>	<p>An error message occurs when you use the following software:</p> <ul style="list-style-type: none"> • 32-bit NMDA to back up a 64-bit database • 64-bit NMDA to back up a 32-bit database <p>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.</p>
<p>SQL2079N</p> <p>On a Windows system, the backup failed for DB2 version 9.x.</p> <p>An error was reported by the shared library: "c:\progra~1\legato\nsr\bin\libnsrdb2.dll". Return code: "30".</p>	<p>The stack size is insufficient for the db2syscs.exe file. Increase the stack size as follows:</p> <ol style="list-style-type: none"> 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: <pre data-bbox="692 1759 1225 1814">C:\Program Files\IBM\SQLLIB\BIN> ..\misc\db2hdr db2syscs.exe /s 1024,32</pre> 3. Start the database engine with the db2start command.
<p>SQL2079N</p>	<p>The VENDOROPT parameter is null.</p>

Table 51 DB2 SQL error messages (continued)

Error message	Description
An error was reported by the shared library during an attempted recovery: "/usr/lib/libnsrdb2.so". Return code "30".	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).
SQL3522N The load operation failed because the COPY YES parameter was specified but the database does not use recoverable logging.	Perform one of the following actions to resolve this issue: <ul style="list-style-type: none"> • Enable the database for rollforward recovery, and then run the load command with the copy yes option. • Run the load command again without specifying the copy yes option.

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 that is 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 488 describes the error messages that are generated by the NetWorker XBSA interface.

Timeout might occur for an NMDA Informix manual onbar backup after upgrade to NetWorker 9.x

After you upgrade from NetWorker 8.x to 9.x, an NMDA Informix manual `onbar` backup might fail when the inactive time is longer than the Inactivity Timeout setting on the NetWorker server.

An increase in the timeout value can resolve this issue. Refer to the `NSR_INACTIVITY_TIMEOUT` parameter details in [Table 40](#) on page 421.

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 that is 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 time stamped 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 and recovery of Domino server fails with secured console

If a Lotus Domino server is configured with a secured (password-protected) console, NMDA cannot back up or recover the Lotus databases.

Under these conditions, backups and recoveries fail and report the following error:

```
Notes Library initialization failed, error = 417
```

For backups and recoveries to succeed, you must remove the password protection for the Domino server console by using the Domino Administrator program or the `Set Secure` command. The Domino Administrator online help provides instructions on removing password protection.

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 that are specified on the operating system level.
- On UNIX and Linux, the Domino server ignores these settings in the locale that would affect the format that is 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.

Child process termination for canceled Lotus backups

If a Lotus manual or scheduled backup is canceled while child processes are waiting for writable volumes on the NetWorker server, the child processes might not be terminated.

To terminate the child processes, make available the writable volumes on the NetWorker server, as requested by the Lotus backup. When the volumes become available, the backup resumes and immediately fails, and the processes are then terminated.

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
```

Prevent memory allocation failure with 32-bit Lotus Domino restores on AIX

If an NMDA Lotus restore with 32-bit Lotus Domino on AIX involves multiple parallel processes (for example, NSR_PARALLELISM is set to a value greater than 1) and produces a memory allocation failure, set the environment variable value EXTSHM=ON and rerun the restore, based on IBM recommendations.

For example, with 32-bit Domino on AIX:

- If an NMDA Lotus restore fails with a memory allocation error, set EXTSHM=ON on the command line where you run the `nsrnotesrc` program.
- If an NMDA Lotus directed recovery through the NetWorker User for Lotus program fails with a memory allocation error, uncomment the following line in the `nsrlotus_remrecov` script:

```
EXTSHM=ON; export EXTSHM
```

Recovery report from NetWorker User for Lotus contains unexpected characters

When NetWorker User for Lotus is used to recover a Domino database that has non-ASCII characters in its name, the file name might be displayed incorrectly in the GUI.

Note:

On Windows, if the non-ASCII characters in the `nsrnotesrc` command line output are not readable, ensure that you use the proper code page for the language. For example, if the language is Spanish, type the following command to set the proper code page before you use the `nsrnotesrc` command for recovery:

```
chcp 1251
```

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](#) 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 `nsrdocr` 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*, and *error_number* are internal numbers representing an error type or error code to report to Technical Support, reached through Customer Support at the Support website.

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 488 describes the error messages that are generated by the NetWorker XBSA interface. To ensure that XBSA error and debug messages are in the NMDA Lotus debug log, set the parameter NSR_XBSA_DEBUG=TRUE as described in [NSR_XBSA_DEBUG](#).

The following table lists error messages that are 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 52 Error messages from Lotus backups and restores

Error message	Description
Error messages common to both Lotus backups and restores:	
A NW server with a NSR client resource for <i>hostname</i> 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.
The call to NotesInitExtended failed: <i>Notes_API_error_number</i> .	NMDA could not initialize the Notes session. To resolve the error, locate the error message for the particular <i>Notes_API_error_number</i> and perform the appropriate corrective

Table 52 Error messages from Lotus backups and restores (continued)

Error message	Description
	action. If the error number is 421, set the NSR_NOTES_INI_PATH parameter.
The LoadLibrary() call failed.	<p>NMDA could not find or load the libnotes.xx or nnotes.dll library file.</p> <p>To resolve the error, ensure that the Notes_ExecDirectory parameter is set to the directory path containing the Lotus library.</p>
Error messages from Lotus backups:	
Could not open file: <i>file_name</i> .	<p>NMDA could not open the file for the backup.</p> <p>To resolve the error, ensure that the file has the correct permissions.</p>
Directory link points to its own directory.	<p>NMDA detected a circular link.</p> <p>To resolve the error, delete the directory link file or correct the path set in the file.</p>
Duplicate entry: <i>file_pathname</i> . Ignoring.	<p>NMDA did not add another entry for the file to the backup list because the file was already on the list.</p> <p>This is an informational message only, and does not require a resolution.</p>
Either nothing was specified to backup or system is down.	<p>NMDA could not derive a list of files to back up.</p> <p>To resolve the error, ensure that you configure and perform the manual backup or scheduled backup according to the instructions in this administration guide.</p>
Error finding last backup instance for <i>file_name</i> ; error = <i>error_description</i> .	<p>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.</p> <p>To resolve the error, ensure that the NetWorker server is up and accessible to the Lotus user on the client computer.</p>
Failed to build exclude list, ignoring exclude list entries.	<p>NMDA could not correctly process the values were specified with the NSR_EXCLUDE_FILE or NSR_EXCLUDE_LIST parameter, so nothing was excluded from the backup.</p> <p>To resolve the error, ensure that the parameters specify the correct values.</p>
Failed to open directory link file: <i>file_name</i> , errno = <i>error_number</i> .	<p>NMDA could not open the file.</p> <p>To resolve the error, ensure that the file has the correct pathname or permissions or both.</p>
Failed to open exclude list file, errno = <i>error_number</i> .	<p>The exclude list file did not exist or did not have the correct access permissions.</p>

Table 52 Error messages from Lotus backups and restores (continued)

Error message	Description
	<p>To resolve the error, ensure that you meet the following requirements:</p> <ul style="list-style-type: none"> • The NSR_EXCLUDE file parameter specifies the correct pathname for the exclude list file. • The exclude list file has the correct permissions.
<i>pathname</i> is not a directory.	<p>The <i>pathname</i> in the Domino directory link (.dir) file was not a directory.</p> <p>To correct the error, delete the .dir file from the directory if it is not used.</p>
The password file lookup for the user <i>Notes_user</i> failed. The user name may be invalid.	<p>The LOTUS_USER parameter specified an invalid Notes user.</p> <p>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.</p>
The ReadFile() call failed.	<p>NMDA could not read the contents of a file to be backed up.</p> <p>To resolve the error, ensure that the file is a readable file.</p>
Error messages from Lotus restores:	
Full backup not found, cannot recover without full backup <i>file_name</i> .	<p>NMDA could not find a full backup for the file in the NetWorker index.</p> <p>To resolve the error, ensure that the path for the <i>file_name</i> matches the case-sensitive path in the NetWorker index.</p>
No objects found for recover.	<p>Either nothing was specified for recovery or the specified items were removed from the recovery list.</p> <p>To resolve the error, perform one of the following actions:</p> <ul style="list-style-type: none"> • 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.	<p>You specified the NSR_NO_NOTES_INIT parameter for a database recovery.</p> <p>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.</p>
The log directory may only be <i>number</i> characters in length.	<p>The NSR_LOG_DIR parameter specified a log directory pathname that is too long.</p> <p>To resolve the error, specify a different (shorter) pathname with the NSR_LOG_DIR parameter.</p>

Table 52 Error messages from Lotus backups and restores (continued)

Error message	Description
The recover time '-t' was either not specified or contained an invalid value.	<p>During a document-level recovery, you did not specify the mandatory <code>-t</code> option correctly with the <code>nsrdocrc</code> command.</p> <p>To resolve the error, specify the <code>-t time</code> option correctly by using the <code>nsr_getdate()</code> format for the time. To recover the backup to the current time, use the <code>-t now</code> option.</p>

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 435 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_INNOODB_OPTIONS=WITH_FRM_ALL`
 - `MYSQL_ONLY_INNOODB_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 that was 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 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*, and *error_number* are internal numbers representing an error type or error code to report to Technical Support, reached through Customer Support at the Support website.

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 that are generated during MySQL backup and restore operations, in alphabetical order.

Table 53 Error messages from MySQL backups and restores

Error message	Description
<p>mysqlbackup: ERROR: Log scan was only able to reach 4845560832, but a checkpoint was at 4845561254.</p>	<p>A MySQL incremental backup failed because the database server overwrote part of the circular InnoDB log file before the <code>ibbackup</code> process was able to read the log file. To resolve the error, perform one of the following actions:</p> <ul style="list-style-type: none"> • 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.
<p>Error: operation failed. global_error_count: 2 backup_config.error_count: 0 92824:nsrdasv: The program 'mysqlbackup' terminated with the error code 11.</p>	<p>A MySQL restore of a database failed due to one of the following reasons:</p> <ul style="list-style-type: none"> • The MySQL database server was not shut down before the restore. • The default InnoDB log size was too small. <p>In either case, to resolve the error, complete the following steps:</p> <ol style="list-style-type: none"> 1. At the <code>mysql</code> command prompt, set <code>innodb_fast_shutdown=0</code>. This setting ensures that all the data is flushed and committed to the database and not the logs. For example: <pre data-bbox="683 1394 1327 1451" style="background-color: #f0f0f0; padding: 5px;">mysql> SET GLOBAL innodb_fast_shutdown=0</pre> 2. Move or delete the <code>ib_logfile</code> file from the <code>mysql</code> directory. For example: <pre data-bbox="683 1566 1327 1623" style="background-color: #f0f0f0; padding: 5px;">/var/lib/mysql/mv ib_logfile ./tmp/</pre> 3. Only if the InnoDB log size was too small, update the <code>innodb_log_file_size</code> value in the MySQL configuration file. 4. Restart the database. The InnoDB engine creates a set of logs. 5. Perform a MySQL full backup with NMDA.
<p>Since the <code>MYSQL_BACKUP_NAME</code> parameter setting, '<i>backup_name</i>', is not an image name, the <code>MYSQL_RESTORE_OPERATION</code> parameter must be set to 'extract'.</p>	<p>A MySQL extract and prepare operation failed because the parameter <code>MYSQL_BACKUP_NAME</code> was set to a backup piece name instead of a image name. With <code>MYSQL_BACKUP_NAME</code> set to the backup piece name, you can perform an extract</p>

Table 53 Error messages from MySQL backups and restores (continued)

Error message	Description
	operation by setting <code>MYSQL_RESTORE_OPERATION</code> to the value "extract". To perform an extract and prepare operation, change the <code>MYSQL_BACKUP_NAME</code> setting to a valid image name.
The first specified binary log does not match the binary log, ' <i>binary_log_name</i> ', for the last restored backup.	A MySQL current time restore or PIT restore failed because the first binary log specified by <code>MYSQL_BINLOG</code> does not correspond to the binary log in use when the last backup piece was created. Use the <code>nsrinfo -v</code> option to find out which binary log was in use for the backup.
The <code>nsrmysqlrc</code> program could not find a save set for recovery.	<p>A MySQL current time restore or PIT restore failed due to one of the following reasons:</p> <ul style="list-style-type: none"> • <code>MYSQL_BACKUP_NAME</code>, <code>NSR_SERVER</code>, or <code>NSR_RECOVER_TIME</code> is set to an incorrect value. <code>NSR_RECOVER_TIME</code> is used only for a PIT restore. • The backup that is 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)	<p>NMDA could not access the MySQL server during an offline MySQL backup. To resolve the error, complete the following steps:</p> <ol style="list-style-type: none"> 1. Add <code>no_connection</code> in the <code>[mysqlbackup]</code> section of the MySQL configuration file, <code>my.cnf</code>. 2. Specify <code>innodb_log_files_in_group</code> with the correct number of logs (2 is the default number) in the <code>[mysqld]</code> 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 484, 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 *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. [NMDA Parameters and Configuration File](#) on page 399 describes how to set the parameters. [RMAN scripts for manual backups](#) on page 135 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 485 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 514.

- 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 135.

- 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 that are described in this topic.

RMAN stores information and RMAN error messages in the log file that is 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 444 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` in the directory that is 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 *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 54 Trace option values and conditions traced

Trace value	Conditions traced
0 (default)	All error conditions.
1	<ul style="list-style-type: none"> • All error conditions. • Entry and exit for each System Backup to Tape (SBT) function (the NMDA implementation of the Oracle SBT interface).
2	<ul style="list-style-type: none"> • 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.

About this task

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.

Procedure

1. Include the `set command id to 'xxx'` command in the RMAN backup script that is used for the Oracle backup. Otherwise, the query in the next step will fail. [Example 7](#) provides a sample script with the command.
2. To determine the Oracle process ID that corresponds to each RMAN channel, run the following query in the Oracle `sqlplus` program:

```
select spid, client_info from v$process p, v$session s where
p.addr=s.paddr and client_info like '%id=%';
```

3. To cancel the backup process, type the `kill` command:

```
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.

About this task

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.

Oracle rollback restore to a new database might fail when OMF is enabled

When the Oracle-Managed Files (OMF) database feature is enabled, a rollback restore to a new database might fail.

For example, when you perform a redirected rollback restore to alternate LUNs by using a ProtectPoint for VMAX backup of an Oracle OMF database, the restore might fail with the following error message:

```
ORA-19511: non RMAN, but media manager or vendor specific failure, error text:
A rollback is not possible when doing relocation during a restore.
Please remove 'rollback' from the RESTORE_TYPE_ORDER parameter or do not
request relocation. (114:123:2)
```

As a workaround, disable the OMF feature after you restore the spfile of the database and before you restore the control file and data files.

Rollback restore fails for Oracle NSM backup in FRA

If you create a copy of an Oracle datafile or archived log to the FRA, and then perform a successful NSM snapshot backup of that copy in the recovery area, a rollback restore of the backup fails to restore the Oracle datafile or archived log to its original location outside the recovery area. The rollback restore produces the following RMAN error:

```
ORA-27037: unable to obtain file status
```

Note: The FRA is a directory where RMAN places all the disk type backups that it creates (backups that are not created with NMDA). The NMDA software can back up the backups from the recovery area.

For example, you create a datafile copy and an NSM snapshot backup as follows:

1. Use the `backup as copy datafile` command to create a datafile copy of the file `/fs1/test1.dbf` to the recovery area at `/fs2`. The datafile copy is created as `/fs2/datafile/o1_mf_test1_2nnplt1z_.dbf`.
2. Use the `backup proxy recovery area` command to perform an Oracle NSM backup of the datafile copy in the recovery area.

Then if you perform a rollback restore of this snapshot backup, the restore occurs to /fs2. The restore produces the RMAN error because RMAN expects the backed-up file to be restored to /fs1, the original file location.

As a workaround, use the additional command `restore...from datafilecopy` to restore from the datafile copy in the recovery area to the original datafile location.

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 517
- [Error messages from the nsroraadmin program](#) on page 527
- [Error messages from the nsrorainfo program](#) on page 530

Error messages from Oracle backups and restores

The following table lists error messages that are 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*, and *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 55 Error messages from Oracle backups and restores

Error message	Description
Error messages from nonproxy backups and restores:	
A connection to NW server ' <i>server</i> ' could not be established because ' <i>reason</i> '.	NMDA could not connect to the NetWorker client file index due to the specified 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 ' <i>file_name</i> ' (<i>errno</i>)	NMDA failed to create the lock file that was required for an index deletion operation. Report the error number (<i>errno</i>) to Technical Support.
Could not decode the 'sf_check' value: xdrs = <i>Oxvalue</i>	This error is an internal XDR error due to a network read or write operation. Report the error to Technical Support.

Table 55 Error messages from Oracle backups and restores (continued)

Error message	Description
Could not decode the 'sf_magic' value: xdrs = 0x <i>value</i>	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 <i>value</i>	This error is an internal XDR error due to a network read or write operation. Report the error to Technical Support.
Could not locate the LNM save file ' <i>backup_piece_name</i> ' on server ' <i>server</i> '.	NMDA could not locate an index record for the backup piece. The index record is probably missing. Use the <code>mminfo</code> and <code>nsrinfo</code> commands to verify the status of the index record.
Could not locate the LNM save time ' <i>save_time</i> ' on server ' <i>server</i> '.	NMDA could not locate an index record for the save time in the client file index. The index record is probably missing. Use the <code>mminfo</code> and <code>nsrinfo</code> commands to verify the status of the index record.
Could not lock ' <i>file_name</i> ' for index deletion. There were <i>number</i> attempts. (<i>errno</i>)	NMDA was able to create the lock file that was required for an index deletion operation, but could not lock the file after the specified number of tries. Report the error number (<i>errno</i>) to Technical Support.
Error in mmdb lookup by time: <i>reason</i>	A lookup in the media database failed for the specified reason. Use the <code>mminfo</code> 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 <code>nsrindexd</code> . This issue was probably due to the NetWorker services being shut down. Restart the NetWorker services on the server, as required.
Exceeded the number of retries for <code>nsr_init()</code> . The NetWorker server may be down or unreachable.	After a maximum of five tries, NMDA failed to call the NetWorker core function, <code>nsr_init()</code> . This issue was probably due to the NetWorker services being shut down. Restart the NetWorker services on the server, as required.
Exceeded the number of retries for <code>nsr_start()</code> . The NetWorker server may be down or unreachable.	After a maximum of five tries, NMDA failed to call the NetWorker core function, <code>nsr_start()</code> . This issue was probably due to the NetWorker services being shut down. Restart the NetWorker services on the server, as required.
Invalid retention policy. Value Ignored.	The <code>NSR_SAVESET_RETENTION</code> parameter had an invalid time value. Ensure that this parameter in the RMAN script has a valid value in the NetWorker date format.
Invalid KEY word	The syntax of the string in the RMAN <code>send</code> command was incorrect. The send command on page 476 provides the correct <code>send</code> command syntax.
Invalid retention policy: <i>retention_time</i> . Value Ignored.	The <code>NSR_SAVESET_RETENTION</code> parameter had an invalid time value, <i>retention_time</i> . Ensure that the parameter

Table 55 Error messages from Oracle backups and restores (continued)

Error message	Description
	NSR_SAVESET_RETENTION in the RMAN script has a valid value in the NetWorker date format.
NSR_DATA_VOLUME_POOL <i>n</i> 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, <i>n</i> 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 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 Backup Configuration on page 75 or Snapshot Backups and Restores on page 345.
Oracle pools are not supported	NMDA does not support Oracle pools. NMDA supports NetWorker pools only. Remove the <code>pool</code> option of the <code>backup</code> command in the RMAN script or set the <code>pool</code> option to zero. The pool option of the backup command on page 476 provides details.
<i>'string'</i> should be in format: KEY=(<i>xxxxx</i>)	The syntax of the string in the RMAN <code>send</code> command was incorrect. The send command on page 476 provides the correct <code>send</code> 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: <i>backup_piece_name</i>	NMDA could not complete the backup because the backup piece name already existed in the NetWorker client file index. Change the <code>format</code> 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: <i>reason</i>	A call of the NetWorker core function, nsr_init(), failed due to the specified 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: <i>reason</i>	A call of the NetWorker core function, nsr_start(), failed due to the specified 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 (<i>errno</i>).	NMDA could not obtain the current time due to an operating system error. Report the operating system error (<i>errno</i>) 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: <i>reason</i>	During a restore, a call of the NetWorker core function, mm_retrieve(), failed due to the specified reason. Take the

Table 55 Error messages from Oracle backups and restores (continued)

Error message	Description
	corrective action suggested by the error message. If required, report the error to Technical Support.
The function <code>nsr_bind_recov_mm()</code> failed with the error: <i>reason</i>	During a restore, a call of the NetWorker core function, <code>nsr_bind_recov_mm()</code> , failed due to the specified reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The function <code>nsr_end()</code> failed with the error message: <i>reason</i>	A call of the NetWorker core function, <code>nsr_end()</code> , failed due to the specified reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The function <code>nsr_rstart()</code> failed with the error: <i>reason</i>	During a restore, a call of the NetWorker core function, <code>nsr_rstart()</code> , failed due to the specified reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The function <code>sbtinit2()</code> has already been called.	This error is an internal error due to Oracle calling the function <code>sbtinit2()</code> twice. Report the error to Technical Support.
The functions <code>sbtinit()</code> or <code>sbtinit2()</code> 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: <code>cfx_name(backup_piece_name)</code> <code>save_time(save_time)</code>	An index lookup located the entry in the client file index but not in the media database. Restart the NetWorker services, and use the <code>mminfo</code> and <code>nsrinfo</code> commands to verify the backup information in the indexes. Run the <code>nsrck</code> program to resolve any corruption of the indexes.
The lookup of ' <i>backup_piece_name</i> ' on server ' <i>server</i> ' failed - ' <i>reason</i> '	NMDA could not locate <i>backup_piece_name</i> in the indexes due to the <i>reason</i> . The indexes might be corrupted. Run the <code>nsrck</code> 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 <code>NSR_CLIENT</code> to the NetWorker client name by using the <code>send</code> command.
The name of the NSR server could not be determined.	NMDA could not determine the name of the NetWorker server. Set the parameter <code>NSR_SERVER</code> to the NetWorker server name by using the <code>send</code> command.
The NSR client name could not be determined.	NMDA could not determine the name of the NetWorker client. Set the parameter <code>NSR_CLIENT</code> to the NetWorker client name by using the <code>send</code> command.
The NSR server name could not be determined.	NMDA could not determine the name of the NetWorker server. Set the parameter <code>NSR_SERVER</code> to the NetWorker server name by using the <code>send</code> command.
The <code>NSR_CLIENT</code> parameter was not set.	NMDA could not determine the name of the NetWorker client. Set the parameter <code>NSR_CLIENT</code> to the NetWorker client name by using the <code>send</code> command.
The <code>NSR_SERVER</code> parameter was not set.	NMDA could not determine the name of the NetWorker server. Set the parameter <code>NSR_SERVER</code> to the NetWorker server name by using the <code>send</code> command.

Table 55 Error messages from Oracle backups and restores (continued)

Error message	Description
The NW authentication for client ' <i>client</i> ' was refused by server ' <i>server</i> ' because ' <i>reason</i> '.	NMDA could not obtain the required authentication to connect to the NetWorker client file index due to the specified 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 <code>send</code> 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 <code>send</code> command.
The NWORA file ID could not be XDR'd. xdrm: <i>Ox value</i> NWORA fid: <i>Ox value</i> ssid: <i>Ox value</i> ssoff: <i>Ox value</i>	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 ' <i>save_time1</i> '. The save time queried was ' <i>save_time2</i> '.	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 <code>nsrck</code> program to resolve any corruption of the indexes.
The removal of SSID ' <i>save_set_id</i> ' failed with error: <i>reason</i>	An index deletion operation failed for the specified reason. Use the <code>mminfo</code> and <code>nsrinfo</code> commands to verify the status of the index record. If required, report the error to Technical Support.
The <code>savefile_fini()</code> call failed. <i>reason</i>	During a restore, a call of the NetWorker core function, <code>savefile_fini()</code> , failed due to the specified reason. Take the corrective action suggested by the error message. If required, report the error to Technical Support.
The <code>sfhead</code> 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 retention time is not in the future: current time: <i>current_time</i> retention: <i>retention_time</i>	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: <i>Ox value</i> ua: <i>Ox value</i>	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 ' <i>file_name</i> ' to raw device ' <i>device_name</i> '.	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 ' <i>device_name</i> ' to file ' <i>file_name</i> '.	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.

Table 55 Error messages from Oracle backups and restores (continued)

Error message	Description
Cannot back up object <i>object_name</i> with proxy copy.	The RMAN <code>backup</code> command included the <code>proxy only</code> option, but the object <i>object_name</i> did not reside on a primary storage device that NSM supports. When the <code>backup</code> command includes the <code>proxy only</code> option, ensure that the object <i>object_name</i> resides on a primary storage device that NSM supports.
Could not find the <code>nsrsnapck</code> binary.	During an index removal for a proxy backup, NMDA could not locate the <code>nsrsnapck</code> binary, which is probably in a nondefault location. Ensure that the parameter <code>NSR_NWPATH</code> is set correctly.
Could not lstat - <i>file_name</i>	The <code>lstat()</code> system call failed. The file <i>file_name</i> either did not exist or had invalid permissions. Ensure that the file is an existing file with valid permissions.
Could not lstat secondary link - <i>file_name</i>	The <code>lstat()</code> system call failed. The file <i>file_name</i> 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 <code>NSR_ORACLECAT_MODE</code> from NWORA resource file.	The error was due to one of the following conditions: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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 383 provides details.
Could not read link - <i>pathname</i>	A proxy backup failed due to the <i>pathname</i> that was an invalid symbolic link. Before a proxy backup, ensure that any symbolic link is a valid link.
Error creating staging directory ' <i>directory</i> '.	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, <code>.nworapc</code> . 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.
<i>nsrsnapck_binary_name</i> process failed with error - <i>reason</i>	During an index removal for a proxy backup, the <code>nsrsnapck</code> binary failed. The binary name is <code>nsrsnapck</code> on UNIX systems and <code>nsrsnapck.exe</code> on Windows systems. Report the error to Technical Support.

Table 55 Error messages from Oracle backups and restores (continued)

Error message	Description
Path <i>pathname</i> is too long.	A proxy backup failed because the specified pathname exceeded the limit of 1,024 bytes. Ensure that any pathname in a proxy backup does not exceed 1,024 bytes.
pb_init() failed with (<i>reason</i>): invalid BRC API 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 <i>NetWorker E-LAB Navigator</i> at https://elabnavigator.emc.com/eln/elhome describes the supported systems.
The BRC API did not return an error string for the SBTPC object: <i>object_name</i>	An unknown error occurred during a BRC API function call by NSM. Report the error to Technical Support.
The BRC status of logical object ' <i>file_name</i> ' was failure: <i>file_status</i>	NSM reported a failure during a proxy backup of the file <i>file_name</i> . Report the error to Technical Support.
The call to pb_environment() failed with error: <i>reason</i>	During a proxy operation, a pb_environment() function call failed due to the specified 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: <i>reason</i>	During a proxy operation, a pb_open() function call failed due to the specified 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: <i>reason</i>	During a proxy operation, a pb_prepare() function call failed due to the specified 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 ' <i>object_name</i> ' with the error: <i>reason</i>	During a proxy operation, a pb_status() function call failed due to the specified 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 ' <i>object_name</i> ' failed with error: <i>reason</i>	During a proxy operation, a pb_status() function call failed due to the specified 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: <i>file_name</i>	The operating system <i>file_name</i> 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 - <i>file_name</i>	The file <i>file_name</i> 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.
The destination does not have the same terminating name as the source ' <i>device_name</i> '.	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/rdisk/c1t2d0s2. Perform a proxy restore of the raw device to a location with the same basename as the backed-up source.

Table 55 Error messages from Oracle backups and restores (continued)

Error message	Description
The file being recovered could not be found in its staging location: <i>file_name</i>	During a proxy restore, an error occurred at the point where the file <i>file_name</i> was to be moved from the staging directory <code>.nworapc</code> to the destination directory. Ensure that there are no permission problems or other problems with the destination directory and the staging directory <code>.nworapc</code> , and then restart the proxy restore. If the error recurs, report the error to Technical Support.
The file ' <i>file_name</i> ' cannot be removed from the staging directory (<i>errno</i>).	During a proxy restore of the file <i>file_name</i> , a file with the same name was found in the <code>.nworapc</code> subdirectory, probably left there by a previous failed restore. The <i>errno</i> is the error number from the failure to remove the existing file. Remove the file <i>file_name</i> from the <code>.nworapc</code> subdirectory, and restart the proxy restore.
The NMDA BRCAPI version <i>version</i> is outside the range supported by the BRC service: <i>earliest_version - latest_version</i>	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 <i>NetWorker E-LAB Navigator</i> at https://elabnavigator.emc.com/eln/elhome describes the supported releases.
The NWORA resource file does not exist. Please create it with <code>nsroraadmin</code> .	A proxy backup failed because the NWORA resource file did not exist. Create the NWORA resource file by using the <code>nsroraadmin</code> command, and restart the proxy backup. NWORA resource file on page 383 provides details.
The NWORA resource lock file does not exist. Please create it by running ' <code>nsroraadmin -r list</code> '	A proxy backup failed because the NWORA resource lock file did not exist. Create the NWORA resource lock file by using the <code>nsroraadmin -r list</code> command, and restart the proxy backup. NWORA resource file on page 383 provides details.
The NWORA resource <code>NSR_ORACLECAT_MODE</code> is in the 'undetermined' state.	In the NWORA resource file, <code>NSR_ORACLECAT_MODE</code> was set to the default value of undetermined. Set the value of <code>NSR_ORACLECAT_MODE</code> to either enabled or disabled (as required) by using the <code>nsroraadmin</code> command.
The object ' <i>file_name</i> ' is not a file.	A proxy backup failed because the file <i>file_name</i> 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 <code>ORACLE_SID</code> must be set when performing proxy copy backups.	During a scheduled proxy backup, the parameter <code>ORACLE_SID</code> was not set in the configuration file. In the configuration file, set the parameter <code>ORACLE_SID</code> to the SID value of the Oracle database.
The OS file name has been specified multiple times by Oracle: <i>file_name</i>	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: <i>file_name</i>	The configuration file specified by the parameter <code>NSR_PROXY_PFILE</code> 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 <code>NSR_PROXY_PFILE</code> is a valid pathname of the configuration file.

Table 55 Error messages from Oracle backups and restores (continued)

Error message	Description
The pb_cancel() call for object ' <i>object_name</i> ' returned the error message: <i>error</i>	The pb_cancel() function call failed during a proxy operation. Report the error to Technical Support.
The pb_inquiry() call failed for object ' <i>object_name</i> ': <i>error</i>	The pb_inquiry() function call failed during a proxy operation. Report the error to Technical Support.
The pb_inquiry() for object ' <i>object_name</i> ' failed because: <i>error</i>	The pb_inquiry() function call failed during a proxy operation. Report the error to Technical Support.
The pb_inquiry() of object ' <i>object_name</i> ' returned error: <i>error</i>	The pb_inquiry() function call failed during a proxy operation. Report the error to Technical Support.
The pb_restore() for object ' <i>object_name</i> ' failed with error: <i>error</i>	The pb_restore() function call failed during a proxy operation. Report the error to Technical Support.
The pb_save() of object ' <i>object_name</i> ' returned error: <i>error</i>	The pb_save() function call failed during a proxy operation. Report the error to Technical Support.
The pb_snapshot() call for object ' <i>object_name</i> ' failed with error: <i>error</i>	The pb_snapshot() function call failed during a proxy operation. Report the error to Technical Support.
The restore destination path is not valid: <i>file_name</i>	During a proxy restore operation, NMDA found the specified restore destination, <i>file_name</i> , to be invalid. Ensure that the specified restore destination is a valid pathname.
The restore operation for the file failed for an unknown reason: <i>file_name</i>	During a proxy restore, an error occurred at the point where the file <i>file_name</i> was to be moved from the staging directory <code>.nworapc</code> to the destination directory. Ensure that there are no permission problems or other problems with the destination directory and the staging directory <code>.nworapc</code> , 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: <i>object_name</i>	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object <i>object_name</i> . Report the error to Technical Support.
The SBTPC object is not in the SBTPCSTATUS_NOTREADY state: <i>object_name</i>	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object <i>object_name</i> . Report the error to Technical Support.
The SBTPC object ' <i>object_name</i> ' failed with the error message: <i>reason</i>	The proxy backup or restore of a file failed during an NSM operation, for the specified reason. Report the error to Technical Support.
The SBTPC object ' <i>object_name</i> ' is entering the SBTPCSTART backup state but its BRC type is: <i>type</i>	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object <i>object_name</i> . Report the error to Technical Support.
The SBTPC object ' <i>object_name</i> ' is entering the SBTPCSTART restore state but its BRC type is: <i>type</i>	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object <i>object_name</i> . Report the error to Technical Support.

Table 55 Error messages from Oracle backups and restores (continued)

Error message	Description
The SBTPC object ' <i>object_name</i> ' is entering the SBTPCSTART state but its status is: <i>status</i>	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object <i>object_name</i> . Report the error to Technical Support.
The SBTPC object ' <i>object_name</i> ' is leaving the BRC prepare state but its status is: <i>status</i>	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object <i>object_name</i> . Report the error to Technical Support.
The SBTPC object ' <i>object_name</i> ' is leaving the BRC save state but its status is: <i>status</i>	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object <i>object_name</i> . Report the error to Technical Support.
The SBTPC object ' <i>object_name</i> ' is leaving the BRC snapshot state but its status is: <i>status</i>	During a proxy operation, NMDA and NSM became unsynchronized as to the status of the object <i>object_name</i> . Report the error to Technical Support.
The SBTPC object ' <i>object_name</i> ' 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 ' <i>directory</i> ' has invalid permissions (<i>errno</i>).	During a proxy restore, NMDA was unable to write to the staging directory, <i>directory</i> . The <i>errno</i> 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: <i>backup_piece_name</i>	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):	
<i>client</i> : WARNING! The NWORA resource file 'save' process output error messages. <i>client</i> : Please check the save log file for more information: <i>log_file</i>	The NWORA resource file could not be backed up after a successful RMAN backup. Analyze the <i>log_file</i> 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 <i>client_name</i> 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.

Table 55 Error messages from Oracle backups and restores (continued)

Error message	Description
The temporary file ' <i>rman_script_path</i> ' could not be created (errno).	The scheduled backup binary, <i>nsrdasv</i> , could not create the file <i>rman_script_path</i> to write the RMAN script that is 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 <i>rman_script_path</i> file.

Error messages from the nsroraadmin program

The following table lists error messages that are 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 56 Error messages from the nsroraadmin program

Error message	Description
Command line arguments are not understood.	The <i>nsroraadmin</i> command included one or more invalid options. Use the <i>nsroraadmin</i> command with the correct options. Configuring the NWORA resource file with the nsroraadmin program on page 388 provides details.
Could not create the NWORA resource file (<i>errno</i>)	The <i>nsroraadmin</i> command could not create the NWORA resource file, possibly due invalid permissions. Ensure that valid permissions exist to enable the <i>nsroraadmin</i> command to create the NWORA resource file. NWORA resource file on page 383 provides details.
Could not create the NWORA resource lock file (<i>errno</i>)	The <i>nsroraadmin</i> command could not obtain the required lock file in the <i>/nsr/tmp</i> or <i>NetWorker_install_path\tmp</i> directory. The command needs the lock file to access the NWORA resource file. Report the error to Technical Support.
Could not open resource file ' <i>file_name</i> ' (<i>errno</i>).	The <i>nsroraadmin</i> 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 <i>nsroraadmin</i> command, or modify the file permissions.
No command line parameters are set.	The <i>nsroraadmin</i> command did not include the correct options. Use the <i>nsroraadmin</i> command with the correct options.
NSR_ORACLECAT_MODE can only be set to 'enabled', 'disabled' or 'undetermined'.	The <i>nsroraadmin</i> command included a NSR_ORACLE_CAT_MODE parameter resource setting other than enabled, disabled, or undetermined. In the <i>nsroraadmin</i> 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 388 provides details.

Table 56 Error messages from the nsroraadmin program (continued)

Error message	Description
NSR_REMOVE_ON_FAILURE can only be set to 'TRUE' or 'FALSE'.	The <code>nsroraadmin</code> command included a <code>NSR_REMOVE_ON_FAILURE</code> parameter resource setting other than TRUE or FALSE. In the <code>nsroraadmin</code> command, set the <code>NSR_REMOVE_ON_FAILURE</code> parameter resource to either TRUE or FALSE only.
NWORA parameter resources must be specified in the 'ResourceName ResourceValue' format.	The <code>nsroraadmin</code> command did not include the name and value of the NWORA parameter resource in the correct format. In the <code>nsroraadmin</code> 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 <code>nsroraadmin</code> command with the <code>-r delete</code> option did not include the SID value of an Oracle database. In the <code>nsroraadmin</code> command with the <code>-r delete</code> option, specify the correct SID value.
The '-r' flag cannot be set multiple times.	The <code>nsroraadmin</code> command contained more than one <code>-r</code> option. Use the <code>nsroraadmin</code> command with only one <code>-r</code> option.
The '-r' option requires an NWORA resource specification.	The <code>nsroraadmin</code> command with the <code>-r</code> option did not include the required resource specification. In the <code>nsroraadmin</code> command with the <code>-r</code> option, specify the required resource name and resource value.
The '-r' option requires either an 'add', 'update', 'list' or 'delete' option.	In the <code>nsroraadmin</code> command, the <code>-r</code> option did not include one of the required keywords: add, update, list, or delete. In the <code>nsroraadmin</code> command, include one of the required keywords with the <code>-r</code> option. Configuring the NWORA resource file with the nsroraadmin program on page 388 provides details.
The first NWORA resource is not a header (<i>errno</i>).	The NWORA resource file is probably corrupted. Verify the contents of the NWORA resource file. If required, repair the resource file by using the <code>nsroraadmin</code> command.
The NWORA resource file does not contain the NSR_NWPATH resource.	The NWORA resource file does not contain the mandatory <code>NSR_NWPATH</code> parameter resource. The file might be corrupted. Verify the contents of the NWORA resource file. If required, repair the resource file by using the <code>nsroraadmin</code> command.
The NWORA resource file does not exist.	The NWORA resource file does not exist. Create the NWORA resource file by using the <code>nsroraadmin</code> command.
The NWORA resource named ' <i>resource_name</i> ' is not found.	The <code>nsroraadmin</code> command specified the name of a resource that does not exist in the NWORA resource file. In the <code>nsroraadmin</code> 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 <code>nsroraadmin</code> command.
The NWORA resource parameter list contains the invalid element ' <i>resource_name</i> '.	The NWORA resource file contains an invalid resource name. The file might be corrupted. NMDA does not support manual editing

Table 56 Error messages from the nsroraadmin program (continued)

Error message	Description
	of the file. Repair the NWORA resource file by using the <code>nsroraadmin</code> command. Configuring the NWORA resource file with the nsroraadmin program on page 388 provides details.
The NWORA resource parameter list for a SID requires the <i>item1</i> , <i>item2</i> and <i>item3</i> information.	The <code>nsroraadmin</code> command for creating or updating an NWORA SID resource did not include the required items. In the <code>nsroraadmin</code> command for creating or updating an NWORA SID resource, include the required items.
The NWORA resource ' <i>resource_name</i> ' is not a SID resource.	The <code>nsroraadmin</code> command with the <code>-r delete</code> option did not include a valid name of an NWORA SID resource. In the <code>nsroraadmin</code> command with the <code>-r delete</code> option, specify a valid name of an NWORA SID resource.
The NWORA resource specified is not supported: <i>resource_name</i> = <i>resource_value</i>	The <code>nsroraadmin</code> command included an invalid name or invalid value for an NWORA parameter resource. In the <code>nsroraadmin</code> command, specify a valid name and valid value for an NWORA parameter resource. NWORA parameter resources on page 384 provides details.
The NWORA SID resource for ' <i>sid_value</i> ' already exists.	The <code>nsroraadmin</code> command tried to add an NWORA SID resource that already existed. In the <code>nsroraadmin</code> command, specify the values for a new NWORA SID resource.
The SID token 'connect' is an empty string.	The <code>nsroraadmin</code> command did not include the required pathname of the RMAN connection file with the connect keyword. In the <code>nsroraadmin</code> 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 388 provides details.
The SID token 'home' is an empty string.	The <code>nsroraadmin</code> command did not include the required pathname of the Oracle home directory with the home keyword. In the <code>nsroraadmin</code> command, specify a valid pathname of the Oracle home directory with the home keyword.
The SID token ' <i>ORACLE_SID</i> ' is invalid.	The <code>nsroraadmin</code> command with the <code>sid</code> keyword included an invalid SID value for the Oracle database. In the <code>nsroraadmin</code> command, specify a valid SID value with the <code>sid</code> keyword.
The SID token 'sid' is an empty string.	The <code>nsroraadmin</code> command did not include the required SID value of the Oracle database with the <code>sid</code> keyword. In the <code>nsroraadmin</code> command, specify a valid SID value with the <code>sid</code> keyword.
The tokens 'sid', 'home' and 'connect' must be set when adding a SID.	The <code>nsroraadmin</code> command to add an NWORA SID resource did not include the settings of the mandatory <code>sid</code> , <code>home</code> , and <code>connect</code> keywords. In the <code>nsroraadmin</code> command to add an NWORA SID resource, include the settings of the <code>sid</code> ,

Table 56 Error messages from the nsroraadmin program (continued)

Error message	Description
	home, and connect keywords.
The value of the NWORA resource is missing.	The <code>nsroraadmin</code> command with the <code>-r update</code> option did not include the NWORA resource value with the resource name. In the <code>nsroraadmin</code> command with the <code>-r update</code> option, specify the NWORA resource value with the resource name.
Unrecognized argument ' <i>option</i> '.	The <code>nsroraadmin</code> command included the unrecognized option <i>option</i> . Use the <code>nsroraadmin</code> command with the correct options. Configuring the NWORA resource file with the nsroraadmin program on page 388 provides details.
You must be the super-user to update the NWORA resource file.	The wrong user typed the <code>nsroraadmin</code> command. Type the <code>nsroraadmin</code> 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 that are 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 57 Error messages from the nsrorainfo program

Error message	Description
A connection to NW server ' <i>server</i> ' could not be established because ' <i>reason</i> '.	NMDA could not connect to the NetWorker client file index due to the specified 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 ' <i>backup_piece_name</i> ' on server ' <i>server</i> '.	NMDA could not locate an index record for the backup piece. The index record is probably missing. Use the <code>mminfo</code> and <code>nsrinfo</code> commands to verify the status of the index record.
Could not locate the LNM save time ' <i>save_time</i> ' on server ' <i>server</i> '.	NMDA could not locate an index record for the save time in the client file index. The index record is probably missing. Use the <code>mminfo</code> and <code>nsrinfo</code> commands to verify the status of the index record.
Error in mmdb lookup by time: <i>reason</i>	A lookup in the media database failed for the specified reason. Use the <code>mminfo</code> command to verify the status of the media database record. Take the corrective action suggested by the error message.

Table 57 Error messages from the nsrorainfo program (continued)

Error message	Description
Exceeded the number of retries. The NetWorker server may be down or unreachable.	NMDA could not contact the NetWorker index service <code>nsrindexd</code> . This issue was probably due to the NetWorker services being shut down. Restart the NetWorker services on the server, as required.
The file ' <i>file_name</i> ' could not be opened.	The file that was specified with the <code>-f</code> option of the <code>nsrorainfo</code> command could not be accessed. Ensure that the specified file exists, and then type the <code>nsrorainfo</code> command again with the <code>-f</code> option.
The file name provided is NULL.	In the <code>nsrorainfo</code> command, the <code>-f</code> option did not include the required file name. In the <code>nsrorainfo</code> command, include the required file name with the <code>-f</code> option.
The index entry failed the cross check: <code>cfx_name(backup_piece_name)</code> <code>save_time(save_time)</code>	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 <code>mminfo</code> and <code>nsrinfo</code> commands to verify the backup information in the indexes. Run the <code>nsrck</code> program to resolve any corruption of the indexes.
The lookup of 'backup_piece_name' on server ' <i>server</i> ' failed - ' <i>reason</i> '	NMDA could not locate <i>backup_piece_name</i> in the indexes due to the <i>reason</i> . The indexes might be corrupted. Run the <code>nsrck</code> program to resolve any corruption of the indexes.
The NW authentication for client ' <i>client</i> ' was refused by server ' <i>server</i> ' because ' <i>reason</i> '.	NMDA could not obtain the required authentication to connect to the NetWorker client file index due to the specified reason. The client might not be configured as a client on the server. Take the corrective action suggested by the error message.
The record obtained has the wrong save time ' <i>save_time1</i> '. The save time queried was ' <i>save_time2</i> '.	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 <code>nsrck</code> program to resolve any corruption of the indexes.

Orchestrated Application Protection troubleshooting tips

If you encounter issues with the Orchestrated Application Protection operations, you can troubleshoot the issues by setting specific parameters that generate the appropriate level of debugging information.

If you encounter issues with an Orchestrated Application Protection backup or restore operation, set the `NSR_DEBUG_LEVEL` and `NSR_DPRINTF` parameters to the following values in the NMDA configuration file:

```
<NSR_DEBUG_LEVEL> 9 </NSR_DEBUG_LEVEL>
<NSR_DPRINTF> TRUE </NSR_DPRINTF>
```

If you encounter issues with Data Domain BoostFS, set the `NSR_DEBUG_LEVEL`, `NSR_DPRINTF`, and `NSR_DEBUG_BOOSTFS` parameters to the following values in the NMDA configuration file:

```
<NSR_DEBUG_LEVEL> 9 </NSR_DEBUG_LEVEL>
<NSR_DPRINTF> TRUE </NSR_DPRINTF>
<NSR_DEBUG_BOOSTFS> TRUE </NSR_DEBUG_BOOSTFS>
```

[NMDA Orchestrated Application Protection parameters](#) on page 451 provides more details about the supported configuration parameters.

The following topics provide additional troubleshooting tips for Orchestrated Application Protection issues.

Orchestrated Application Protection backup might fail on Linux with a fusermount error

With some Linux distributions and releases, an Orchestrated Application Protection backup might fail with the following type of error message:

```
The backup command '/usr/bin/backup.sh' did not complete successfully: The
command '/opt/emc/boostfs/bin/boostfs' did not complete successfully. Error:
fusermount: /nsr/apps/tmp/d631ab4b_318204_3494 not mounted. Return code: 0.
```

In this example message, `/usr/bin/backup.sh` is the pathname of the user-defined backup script file and `d631ab4b_318204_3494` is a temporary directory. These names will be different in the error message that is displayed in each particular environment.

The error occurs when the `fusermount` binary considers the mount point to be unmounted but the mount point remains in the mounted state.

For assistance with this issue, contact Data Domain Technical Support and reference the Data Domain bug number 206596.

Orchestrated Application Protection backup might become suspended on Linux with database password authentication

During an Orchestrated Application Protection backup, the password prompt appears on the command console. For example, you might see one of the following password prompts on the command console:

```
Password:
```

or

```
Enter Password:
```

The database requires the password authentication to enable the backup. When the `USER_PSWD` and `USER_PSWD_PROMPT` parameter settings are missing in the NMDA configuration file, the backup becomes suspended when the password prompt appears.

In this case, ensure that the `USER_PSWD` and `USER_PSWD_PROMPT` parameters are set in the NMDA configuration file. If these parameters are already set, ensure that the `USER_PSWD_PROMPT` parameter setting matches the password prompt on the console.

[Database authentication](#) on page 156 provides more details.

SAP IQ troubleshooting tips and error messages

An NMDA SAP IQ operation can encounter one of the following issues. If you encounter any of these issues, use the recommended resolution.


 **Note:** The SAP IQ server generates log messages in the log files `$$SYBASE/IQ-16_0/logfiles/*.srvlog`.

Table 58 SAP IQ issues and resolutions

SAP IQ issues	Resolutions
Unable to connect to the database or <code>utility_db</code> database.	Ensure that the <code>interfaces</code> file under the OCS library installation directory contains the correct port number for the corresponding database.
Unable to restore multiple read-only dbspaces or dbfiles.	Restore only one read-only dbspace or dbfile at a time. NMDA does not support the restore of multiple dbspaces or dbfiles.
Unable to restore a read-only dbfile when the database is shut down.	Ensure that the database is running when you perform the restore of the read-only dbfile.
Unable to restore a read-only dbspace when the database is shut down and the dbspace is online.	Ensure that the database is running for the restore. Alternatively, change the dbfile to the offline state before you run the restore.

Sybase troubleshooting tips and error messages

The following topics provide NMDA Sybase troubleshooting tips and error messages.

Sybase backup failure due to `.lock` files

An NMDA Sybase backup or restore interruption due to a failure, termination, or other reason might generate `.lock` files in the `/nsr/apps/tmp` directory. For example, the failure of an NMDA Sybase backup generates the following `.lock` files:

```
/nsr/apps/tmp/testdb1.lock
/nsr/apps/tmp/testdb2.lock
/nsr/apps/tmp/testdb3.lock
```

The `.lock` files might cause a subsequent NMDA Sybase backup to terminate with the following type of error message:

```
97301:nsrdasv: Another Sybase backup is running. Exiting the backup.
```

To prevent or resolve this issue, remove all the `.lock` files from the `/nsr/apps/tmp` directory.

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 488 describes the error messages that are 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 534
- [Sybase backup server and libnsrsyb error messages](#) on page 538

Error messages from Sybase consistency checks, backups, and restores

The following table lists error messages that are 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 59 Error messages from Sybase consistency checks, backups, and restores

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 <i>Sybase_server</i> : 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 <i>nsrdsav</i> , <i>nsrsybrc</i> , and <i>nsrsybcc</i> commands each operate on a database (or for <i>nsrsybrc</i> and <i>nsrsybcc</i> , 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 <i>-s NetWorker_server</i> option to specify the NetWorker server to receive the command.
non fatal internal error from server <i>server_name</i> : 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 <i>nsrsybcc</i> 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:/ <i>ASE_server_name</i>) or a list of databases (SYBASE:/ <i>ASE_server_name/database_name1</i> SYBASE:/ <i>ASE_server_name/database_name2</i>) 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 <i>nsrdsav</i> , <i>nsrsybcc</i> , or <i>nsrsybrc</i> program can operate on a single Sybase server because the user ID and password that are 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 <i>database_name</i> has a length greater than the maximum of 30	The database name that was supplied at the command line was longer than 30 characters. The maximum database name length is 30 characters.

Table 59 Error messages from Sybase consistency checks, backups, and restores (continued)

Error message	Description
The instance name was not provided in the command line <code>command_line_value</code> . The command line has the form <code>SYBASE:/instance_name[/database_name]</code> .	You specified 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 <code>nsrdasv</code> , <code>nsrsybrc</code> , and <code>libnsrsyb</code> . 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 <code>\$USER</code> .
Error message from Sybase consistency checks:	
invalid check option <code>-o</code> value was supplied	The database consistency check option was not valid. The <i>NetWorker Module for Databases and Applications Command Reference Guide</i> or the <code>nsrsybcc</code> man page provides a list of supported options.
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 that was supplied to the <code>nsrdasv</code> command.
cannot find database <code>database_name</code> in instance <code>server_name</code>	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 <code>process_number</code> running command PRECMD or POSTCMD completed with a result of <code>n</code>	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 and POSTCMD provide details.
the exit status of process <code>process_number</code> 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 <code>NSR_BACKUP_LEVEL</code> specified a level other than full, incremental, or skip.
the <code>NSR_BACKUP_PATHS</code> parameter was not set in the configuration file	The configuration files does not include the required setting of the <code>NSR_BACKUP_PATHS</code> parameter.
The Sybase user name was not set. Please specify the Sybase user in the configuration by setting the <code>SYBASE_USER</code> parameter	Supply a username for Sybase login through the <code>SYBASE_USER</code> parameter setting in the configuration file.
unable to create directory entries	NMDA could not create the directory entries. Check the <code>xbsa.messages</code> file for the specific reason for the entry creation failure.

Table 59 Error messages from Sybase consistency checks, backups, and restores (continued)

Error message	Description
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 <i>database_name</i> in instance <i>server_name</i>	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 <i>database_name</i> in instance <i>server_name</i>	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 <i>database_name</i>	The command to dump the transaction log with the <code>no_truncate</code> 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 that are specified in \$PATH.
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 <i>database_name</i>	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 <i>database_name</i> with the <code>no_log</code> 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 " <i>database_name</i> " because it does not exist in " <i>instance_name</i> "	The <code>nsrsybrc</code> command could not find a backup of the database for recovery. Run the <code>nsrinfo</code> command to see if a backup exists, and ensure that the user ID used for the <code>nsrsybrc</code> command matches the object owner that is displayed. To prevent this problem, run the Sybase Backup server and the <code>nsrsybrc</code> command and <code>nsrdasv</code> command from the same user ID.
cannot restore to the destination database <i>database_name</i> because it does not exist in the instance <i>server_name</i>	The database to which the <code>nsrsybrc</code> command is recovering data does not exist. Create the database and retry the <code>nsrsybrc</code> 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.

Table 59 Error messages from Sybase consistency checks, backups, and restores (continued)

Error message	Description
if the destination is an instance, the source must be an instance, too	<p>You used the <code>-d <i>destination</i></code> 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:</p> <pre data-bbox="639 457 1318 562"> nrsybrc -U sa -P xxx -d SYBASE:/ destination_server/destination_database SYBASE:/source_server/source_database </pre> <p>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.</p>
if the source is an instance, the destination must be an instance, too	<p>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:</p> <pre data-bbox="639 877 1318 940"> nrsybrc -U sa -P xxx -d SYBASE:/ destination_server SYBASE:/source_server </pre> <p>Note: The Sybase server names are case-sensitive and must be in the same case as recorded in the NetWorker backup indexes.</p>
internal error. Full backup expected but not found.	<p>A full backup was found, but was then no longer available before the <code>nrsybrc</code> 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.</p>
invalid time specification: time value	<p>The <code>-t <i>time</i></code> option supplied with the <code>nrsybrc</code> command was not valid. This option must be in the <code>nsr_getdate</code> form. The <code>nsr_getdate</code> man page provides details.</p>
no NetWorker server was specified	<p>The NetWorker server was not specified or could not be found. Use the <code>-s <i>server_name</i></code> option to specify the NetWorker server to receive the command.</p>
Recover option validation error.	<p>You must specify either the entire server (<code>SYBASE:/<i>ASE_server_name</i></code>) or a list of databases (<code>SYBASE:/<i>ASE_server_name/database_name1</i>SYBASE:/<i>ASE_server_name/database_name2</i></code>) at the command line. You cannot specify both a server name and a list of databases simultaneously.</p>
the command line did not specify a database or an instance to restore	<p>You must specify the name of the database or Sybase server to be recovered with the <code>nrsybrc</code> command.</p>
there are no databases to restore in instance <i>server_name</i>	<p>NMDA could not find any databases in the directory entry for the Sybase server database.</p>
there is no backup of the instance for the time supplied	<p>NMDA could not find a backup for the Sybase server name. Ensure that you run the <code>nrsybrc</code> command with the same user</p>

Table 59 Error messages from Sybase consistency checks, backups, and restores (continued)

Error message	Description
	ID as used to run the <code>nsrdasv</code> 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 <code>database_name</code> in instance <code>server_name</code> 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 <code>nsrinfo</code> command to determine when the last full backup occurred.
unable to query backup	There was an error querying the backup from the server. Check the <code>xbsa.messages</code> 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 that are logged in the Sybase Backup server error log. The Sybase documentation describes other Sybase Backup server errors.

Table 60 Sybase backup server and libnsrsyb error messages

Error message	Description
libnsrsyb opened with an unknown mode: internal error	The <code>libnsrsyb</code> 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 <code>ddddddd</code> has non digits in it. Timestamps are composed of digits in the form <code>YYYYMMDDhhmmsslll</code> .	<p>The timestamp that was supplied for the <code>load</code> command from the <code>isql</code> command line has a timestamp with an incorrect format. The timestamp must have the format <code>YYYYMMDDhhmmsslll</code>, where:</p> <ul style="list-style-type: none"> • <code>YYYY</code> indicates the year. • <code>MM</code> indicates the month. • <code>DD</code> indicates the day. • <code>hh</code> indicates the hour. • <code>mm</code> indicates the minutes. • <code>ss</code> indicates the seconds. • <code>lll</code> indicates the milliseconds. The millisecond position is optional. Alternatively, you can enter <code>000</code> for the milliseconds.
time stamps are not valid for dump command	The <code>isql</code> command line specified a timestamp for a <code>dump</code> command. Timestamps are not valid with the <code>dump</code> 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 <code>xbsa.messages</code> file for specific details.

Table 60 Sybase backup server and libnsrsyb error messages (continued)

Error message	Description
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 <code>xbsa.messages</code> file for specific details.
unable to create environment variables	The resources that are required to create the internal environment variable array were not available. This might be due to access problems in the <code>/nsr/tmp</code> 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 <code>/nsr/applogs/xbsa.messages</code> 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 <code>xbsa.messages</code> file for the error text. If the debug level is not 2, change the setting to 2 and retry the operation. Check the <code>xbsa.messages</code> 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 <code>xbsa.messages</code> 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 <code>xbsa.messages</code> 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 <code>nsrinfo</code> to see which save sets are available.	The item to be loaded could not be found. Use the <code>nsrinfo</code> 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 that was used was the current time, which means that no backup exists. Use the <code>nsrinfo</code> 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 <code>isql</code> 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 <code>xbsa.messages</code> 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 <code>xbsa.messages</code> file for specific details.
unknown backup type supplied	The backup type that was supplied from the NetWorker server was not a database or a transaction log.

GLOSSARY

This glossary contains the definitions of terms that are found in this manual. Most of the terms are specific to NMDA. For terms specific to the NetWorker software, refer to the latest *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 network computers, and who adds users and defines user privileges.
- attribute** Name or value property of a resource.
- autochanger** See [library](#).

B

- backup**
1. Duplicate of database or application data, or an entire computer system, stored separately from the original, which can be used to recover the original if it is lost or damaged.
 2. Operation that saves data to a volume for use as a backup.
- 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** A volume used to store backup data. NetWorker backup data cannot be stored on an archive volume or a clone volume.
- blob space** (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. The bootstrap consists of three components that reside on the NetWorker server: the media database, the resource database, and a server index.

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 NWORA resource file .
client	Host on a network, such as a computer, workstation, or application server whose data can be backed up and restored with the backup server software.
Client Direct	Feature that enables clients to deduplicate backup data and send it 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 is purged from the index after the browse time of each backup expires.
client-initiated backup	See manual backup .
client-side configuration	Backup configuration that is created without the configuration wizard. The configuration is performed by using the NetWorker Management Console and the configuration files or parameters that are stored on the client host, as compared to a server-side configuration. See server-side configuration .
clone	<ol style="list-style-type: none"> 1. Duplicate copy of backed-up data, which is indexed and tracked by the NetWorker server. Single save sets or entire volumes can be cloned. 2. Type of mirror that is specific to a storage array.
clone volume	Exact duplicate of a backup or archive volume. NetWorker software can index and track four types of volumes (backup, archive, backup clone, and archive clone). Save sets of these different types may not be intermixed on one volume. Clone volumes may be used in exactly the same way as the original backup or archive volume.
CloudBoost device	Logical storage device that is created on a CloudBoost appliance and is used to store NetWorker backups. Each device appears as a folder on the CloudBoost appliance and appears with a storage volume name in NMC.
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	Wizard that is integrated with the NetWorker Management Console GUI, which you can use to configure scheduled backups and also configure and run data restore and recovery. The NetWorker modules support the recovery wizard for specific applications.
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.
database	<ol style="list-style-type: none"> 1. Collection of data arranged for ease and speed of update, search, and retrieval by computer software. 2. Instance of a database management system (DBMS), which in a simple case might be a single file containing many records, each of which contains the same set of fields.
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. A Data Domain device is also known as a DD Boost device.
DBA (database administrator)	Person who is typically responsible for installing, configuring, and maintaining database systems.
dbobject	Informix database object, a term that can refer to a blob space, db space, 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 db spaces.
DD Boost	Optimized library and communication framework with a special Data Domain API that allows the backup 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 data is returned 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 database files are restored in a directed recovery.
device	<ol style="list-style-type: none"> 1. Storage folder or storage unit that can contain a backup volume. A device can be a tape device, optical drive, autochanger, or disk connected to the server or storage node. 2. General term that refers to storage hardware. 3. Access path to the physical drive, when dynamic drive sharing (DDS) is enabled.
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.
direct file access (DFA)	See Client Direct .
disaster recovery	Restore and recovery of data and business operations in the event of hardware failure or software corruption.

distributed segment processing (DSP) Part of the DD Boost interface, which enables data deduplication to be performed on a host before the data is sent to the Data Domain system for storage.

drive Hardware device through which media can be read or written to. **See** [device](#).

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.

F

firewall Security software designed to prevent unauthorized access to or from a private network.

full backup Type of backup that backs up all data objects or files, including the transaction logs contained in databases, regardless of when they last changed. **See** [level](#).

G

group One or more client computers that are configured to perform a backup together, 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 that the application services continue despite a hardware or software failure. Each cluster node has its own IP address with private resources or disks that are available only to that computer.

host Computer on a network.

hot backup **See** [online backup](#) .

I

incremental backup **See** [level](#).

instance Combination of processes that runs each time a database starts up.

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 a backup application.

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 or files, regardless of when they last changed.
- An incremental backup backs up only data objects or files that have changed since the previous backup.

library Hardware device that contains one or more removable media drives, as well as 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.

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 a disk file system or magnetic tape, to which backup data is written. See [volume](#).

media index Database that contains indexed entries of storage volume location and the life cycle status of all data and volumes managed by the NetWorker server. Also known as media database.

media pool See [pool](#).

mount host Host used in NSM backups that is separate from the database server host, with access to the primary storage unit. During a backup to conventional storage, either the database server host or mount host backs up a point-in-time copy (snapshot) from the primary storage to conventional storage.

multiplex To simultaneously write data from more than one save set to the same storage device.

N

NetWorker Management Console (NMC) Software program that is 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.
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 storage nodes on the same network.
NetWorker Snapshot Management (NSM)	Technology that provides point-in-time snapshot copies of data. NetWorker software backs up data from the snapshot. This allows applications to continue to write data during the backup operation, and ensures that open files are not omitted.
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 shut down and unavailable to users. Also known as a cold backup.
online backup	Backup of database objects performed while the corresponding database or instance is running and available to users. Also known as a hot backup.
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: <ul style="list-style-type: none"> • 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	Feature that enables a maximum number of concurrent streams of data during backup or restore operations. For example, parallelism values can be set for the NetWorker server, clients, pools, and groups.

pathname	Set of instructions to the operating system for accessing a file: <ul style="list-style-type: none"> • An absolute pathname indicates how to find a file by starting from the root directory and working down the directory tree. • A relative pathname indicates how to find a file by 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	Backup client that is bound to a physical host in the cluster and 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 that contains an image of the data as it appeared at a specific point in time. A PIT copy is also called a snapshot or shadow copy.
policy	Set of defined rules for client backups that can be applied to multiple groups. Groups have dataset, schedule, browse, and retention policies.
policy uniformity	Consistency of the 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	<ol style="list-style-type: none"> 1. NetWorker sorting feature that assigns specific backup data to be stored on specified media volumes. 2. Collection of NetWorker backup volumes to which specific data has been backed up.
primary storage	Server storage subsystem, such as a disk array, that contains application data and any persistent snapshots of data.
probe-based backup	Type of scheduled backup, also known as an event-based backup, where the NetWorker server initiates the backup only when specified conditions are met, as determined by one or more probe settings.
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 snapshot backup .
proxy restore	Restore of Oracle data from a proxy backup through NSM. See snapshot restore . See rollback restore .
Q	
quiesce	State in which all writes to a disk are stopped and the file system cache is flushed. Quiescing the database prior to creating the snapshot provides a transactionally consistent image that can be remounted.

R

recover	To restore data files from backup storage to a client and apply transaction (redo) logs to the data to make it consistent with a given point-in-time.
recyclable save set	Save set whose browse and retention policies have expired. Recyclable save sets are removed from the media database.
recyclable volume	Storage volume whose data has exceeded both its browse and retention policies and is now available to be relabeled and reused.
regular backup or restore	NMDA backup or restore that does not use snapshot technologies through NSM.
remote device	<ol style="list-style-type: none"> 1. Storage device that is attached to a storage node that is separate from the NetWorker server. 2. Storage device at an offsite location that stores a copy of data from a primary storage device for disaster recovery.
resource	Software component whose configurable attributes define the operational properties of the NetWorker server or its clients. Clients, devices, schedules, groups, and policies are all NetWorker resources.
resource database	NetWorker database of information about each configured resource.
restore	To retrieve individual data files from backup media and copy the files to a client without applying transaction logs.
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 period of time that backup data is retained on a storage volume and available for recovery. After this time is exceeded, the data is eligible to be overwritten.
RMAN	See Oracle Recovery Manager (RMAN) .
rollback restore	Process by which a snapshot is restored to its source or alternate location by using the capability of the storage array. A rollback restore destroys existing data on the target location.
roll forward	To apply transactional logs to a recovered database to restore it 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.

S

save	NetWorker command that backs up client files to backup media volumes and makes data entries in the online index.
-------------	--

save set	<ol style="list-style-type: none"> 1. Group of tiles or a file system copied to storage media by a backup or snapshot rollover operation. 2. NetWorker media database record for a specific backup or rollover.
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 ID (ssid)	Internal identification number assigned to a save set.
save stream	Data and save set information that is 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 that is 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	A storage library attached to the NetWorker server or storage node, used to store traditional or snapshot backups. A NetWorker server Device resource must be configured for each secondary storage device. See primary storage .
server-side configuration	Backup configuration that is performed through the configuration wizard, with settings saved on the NetWorker server, as compared to client-side configuration. See client-side configuration .
shared disk	Storage disk that is connected to multiple nodes in a cluster.
snapshot	Point-in-time, read-only copy of specific data files, volumes, or file systems on an application host. Operations on the application host are momentarily suspended while the snapshot is created on a proxy host. Also called a PiT copy, image, or shadow copy.
snapshot backup	Snapshot created on a storage array as a backup. Previously called instant backup.
snapshot policy	Sets of rules that control the life cycle of snapshots. These rule specify the frequency of snapshot creation, how long snapshots are retained, and which snapshots will be backed up to conventional storage media.
snapshot restore	Restore from a snapshot backup. Previously called instant restore.
snapshot save set	Group of files or other data included in a single snapshot. Previously called a snapshot.
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	Computer that manages physically attached storage devices or libraries, whose backup operations are administered from the controlling NetWorker server. Typically a “remote” storage node that resides on a host other than the NetWorker server.

T

- tablespace** 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 list of changed files in a database, stored in a log file to execute quick restore and rollback transactions.

U

- unscheduled backup** See [manual backup](#).

V

- virtual cluster client** NetWorker client that is not permanently bound to one physical host but is managed by a cluster manager. It is also referred to as a logical cluster client or a virtual client.
- volume**
1. Unit of physical storage medium, such as a disk or magnetic tape, to which backup data is written.
 2. Identifiable unit of data storage that may reside on one or more computer disks.
- volume name** Name that you assign to a backup volume when it is labeled.