



LEGATO®

Keeping the World's Business-Critical Information Available

Legato NetWorker®

Module for Microsoft® SQL Server™

Release 3.1
Windows® Version

Administrator's Guide

© 2002, Legato Systems, Inc. All rights reserved.

Legato, the Legato logo, Legato NetWorker, ClientPak, Octopus, GEMS, AlphaStor, Celestra, and SmartMedia are registered trademarks; and Co-StandbyServer, eCluster, SnapShotServer, RepliStor, and StandbyServer are trademarks of Legato Systems, Inc. in the United States and/or other countries.

All other trademarks are the trademarks of their respective owners.

Legato NetWorker Module for Microsoft SQL Server Administrator's Guide, Release 3.1, Windows Version

January 2002

01-8561-3.1

THIS PRODUCT CONTAINS CERTAIN COMPUTER PROGRAMS AND OTHER PROPRIETARY MATERIAL, THE USE OF WHICH IS SUBJECT TO THIS END-USER LICENSE AGREEMENT (THE "AGREEMENT"). DO NOT PROCEED WITH THE INSTALLATION OF THIS PRODUCT UNTIL YOU (LATER DEFINED AS LICENSEE) HAVE READ THIS AGREEMENT AND AGREE TO BE BOUND BY AND BECOME A PARTY TO THIS AGREEMENT. BY PROCEEDING WITH THE INSTALLATION OF THIS PRODUCT (OR AUTHORIZING ANY OTHER PERSON TO DO SO), YOU AND YOUR COMPANY ACCEPT THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, YOU MAY EXIT THE INSTALLATION PROCESS AND RETURN THIS PRODUCT, ALL MEDIA AND DOCUMENTATION, AND PROOF OF PAYMENT, TO THE PLACE YOU OBTAINED THEM FOR A FULL REFUND WITHIN TEN (10) DAYS OF FIRST ACQUIRING THIS PRODUCT. WRITTEN APPROVAL IS NOT A PREREQUISITE TO THE VALIDITY OR ENFORCEABILITY OF THIS AGREEMENT AND NO SOLICITATION OF ANY SUCH WRITTEN APPROVAL BY OR ON BEHALF OF LEGATO SHALL BE CONSTRUED AS AN INFERENCE TO THE CONTRARY. IF YOU HAVE ORDERED THIS PRODUCT, LEGATO'S ACCEPTANCE IS EXPRESSLY CONDITIONAL ON YOUR ASSENT TO THESE TERMS TO THE EXCLUSION OF ALL OTHER TERMS; IF THESE TERMS ARE CONSIDERED AN OFFER BY LEGATO, ACCEPTANCE IS EXPRESSLY LIMITED TO THESE TERMS.

1. DEFINITIONS

1.1 Authorization Code: means the code provided to Licensee by Legato for permanent authorization to use the Software. The Authorization Code is provided to Licensee once the Enabler Code is registered with Legato.

1.2 Documentation: means any user reference materials on any media, provided by Legato for use with the Software.

1.3 Enabler Code: means the code provided by Legato for activation of the Software.

1.4 Licensee: means the person or entity acquiring this License or for whom this License was acquired.

1.5 Software means the object code copy of the software program provided to You in association with this Agreement, together with the associated original electronic media and all accompanying manuals and other documentation, and together with all enhancements, upgrades, and extensions thereto that may be provided by Legato to You from time to time.

2. OWNERSHIP AND ADMINISTRATION OF SOFTWARE

2.1 Ownership and Title. As between the parties, Legato, and its licensors, own and shall retain all right, title, and interest in and to: (i) the Software including all intellectual property rights embodied therein; (ii) all of the service marks, trademarks, trade names, or any other designations associated with the Software; and (iii) all copyrights, patent rights, trade secret rights, and other proprietary rights relating to the Software. Licensee further acknowledges and agrees that it shall have no rights with respect to any of the foregoing rights other than the rights expressly set forth in this Agreement.

2.2 Software Activation. Legato employs Enabler Codes and Authorization Codes that enable the use of the Software. The Software is shipped in a "Media Kit" which consists of object code software on CD-ROM and an Enabler Code for initial activation of the Software or the Software and Enabler Code may be delivered electronically. Once Legato receives confirmation from Licensee that the Enabler Code is installed and is provided with the host ID information, Legato will provide the Authorization Code to Licensee. Legato administers the generation and distribution of Enabler and Authorization Codes, which administration may be modified by Legato from time to time.

2.3 Administration of Software. Legato may include on the media with the Software additional computer programs which are not currently licensed for use by Licensee and to which the Enabler Code or Authorization code will not permit access. Inclusion of such additional computer programs in no way implies a license from Legato and access or use of such programs is strictly prohibited unless Licensee procures the right to use any such program and the applicable Enabler Code is provided thereto.

3. LICENSE GRANT

3.1 Grant. Legato grants to Licensee a non-exclusive, nontransferable, non-sublicensable, perpetual, unless terminated in accordance with the provisions of this Agreement, license (the "License") to (i) use the Software installed in accordance with the Documentation and only on the licensed computer solely for its own internal operations; and (ii) move the Software temporarily in case of computer system malfunction. The License granted under this Agreement does not constitute a sale of the Software or any portion or copy of it. Licensee may not use the Software on more than one computer system unless otherwise specifically authorized by an explicit Software product, or additional licenses for additional computers are purchased. Rights not expressly granted are reserved by Legato. Where the Software is provided to Licensee at no charge for evaluation purposes only, the License granted is limited to a continuous thirty (30) day period, commencing with the acceptance of this Agreement (the "Evaluation Period"). At the conclusion of the Evaluation Period, Licensee agrees to destroy the Software and certify its destruction to Legato, in writing, within ten (10) days, or shall return the Software to Legato or purchase a perpetual license.

3.2 Copies. Licensee may make copies of the Software provided that any such copy is: (i) created as an essential step in utilization of the Software on the licensed computer and is used in no other manner; or (ii) used for archival purposes to back-up the licensed computers. All trademark and copyright notices must be reproduced and included on such copies. Licensee may not make any other copies of the Software.

3.3 Restrictions on use. Licensee shall not, and shall not aid, abet, or permit any third party to: (i) decompile, disassemble, or otherwise reverse engineer or attempt to reconstruct or discover any source code or underlying ideas or algorithms of the Software by any means whatsoever; (ii) remove any identification, copyright, or other notices from the Software; (iii) provide, lease, lend, use for timesharing or service bureau purposes, or otherwise use or allow others to use the Software or any portion thereof to or for the benefit of third parties; (iv) create a derivative work of any part of the Software; or (v) develop methods to enable unauthorized parties to use the Software. If EC law is applicable, the restrictions in Section 3.3 (i) are limited so that they prohibit such activity only to the maximum extent such activity may be prohibited without violating the EC Directive on the legal protection of computer programs. Notwithstanding the foregoing, prior to decompiling, disassembling, or otherwise reverse engineering any of the Software, Licensee shall request Legato in writing, to provide Licensee with such information or assistance and Licensee shall refrain from decompiling, disassembling, or otherwise reverse engineering any of the Software unless Legato cannot or has not complied with such request in a commercially reasonable amount of time.

3.4 Purchase Orders. Nothing contained in any purchase order, acknowledgment, or invoice shall in any way modify the terms or add any additional terms or conditions to this Agreement.

3.5 Updates. This section applies if the Software acquired is an update to the original Software (the "Update"). An Update does not constitute a legally licensed copy of the Software unless purchased as an Update to a previous version of the same Software. The Update may only be used in accordance with the provisions of this Agreement. The Update, together with the original Software, constitute one (1) legally licensed copy of the Software.

3.6 Evaluation License. This Section applies if the Software is being used for an initial thirty (30) day evaluation period. The license is valid only for a period of thirty (30) days from the delivery of the Software, and is designed to allow Licensee the right to evaluate the Software during such period. In the event that Licensee desires to enter into a longer-term license agreement with Legato, Licensee shall obtain an appropriate Enabler and Authorization Code in accordance with Section 2.2 above, upon payment of applicable fees, which authorizes use of the Software after such evaluation period, but only subject to all of the terms and conditions of this Agreement. In the event Licensee determines not to enter into a licensing transaction with Legato at the end of such thirty (30) day evaluation period, then Licensee's rights under this Agreement shall terminate automatically and Licensee shall promptly return to Legato or destroy all copies of the Software and so certify to Legato.

4. MAINTENANCE AND SUPPORT

4.1 Legato has no obligation to provide support, maintenance, upgrades, modifications, or new releases under this Agreement. Legato may provide such services under separate agreement.

5. LIMITED WARRANTY

5.1 Media and Documentation. Legato warrants that if the media or documentation are damaged or physically defective at the time of delivery of the first copy of the Software to Licensee and if defective or damaged product is returned to Legato (postage prepaid) within thirty (30) days thereafter, then Legato will provide Licensee with replacements at no cost.

5.2 Software. Subject to the conditions and limitations of liability stated herein, Legato warrants for a period of thirty (30) days from the delivery of the first copy of the Software to Licensee that the Software, as delivered, will materially conform to Legato's then current published Documentation for the Software. This warranty covers only problems reported to Legato during the warranty period.

5.3 Remedies. The remedies available to Licensee hereunder for any such Software which does not perform as set out herein shall be either repair or replacement, or, if such remedy is not practicable in Legato's opinion, refund of the license fees paid by Licensee upon a return of all copies of the Software to Legato. In the event of a refund this Agreement shall terminate immediately without notice.

6. TERM AND TERMINATION

6.1 Term. The term of this Agreement is perpetual unless terminated in accordance with its provisions.

6.2 Termination. Legato may terminate this Agreement, without notice, upon Licensee's breach of any of the provisions hereof.

6.3 Effect of Termination. Upon termination of this Agreement, Licensee agrees to cease all use of the Software and to return to Legato or destroy the Software and all Documentation and related materials in Licensee's possession, and so certify to Legato. Except for the License granted herein and as expressly provided herein, the terms of this Agreement shall survive termination.

7. DISCLAIMER AND LIMITATIONS

7.1 Warranty Disclaimer. EXCEPT FOR THE LIMITED WARRANTY PROVIDED IN SECTION 5 ABOVE, LEGATO MAKES NO WARRANTIES WITH RESPECT TO ANY SOFTWARE AND DISCLAIMS ALL STATUTORY OR IMPLIED WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING OR USAGE OF TRADE AND ANY WARRANTIES OF NONINFRINGEMENT. ALL SOFTWARE IS PROVIDED "AS IS" AND LEGATO DOES NOT WARRANT THAT THE SOFTWARE WILL MEET ANY REQUIREMENTS OR THAT THE OPERATION OF SOFTWARE WILL BE UNINTERRUPTED OR ERROR FREE. ANY LIABILITY OF LEGATO WITH RESPECT TO THE SOFTWARE OR THE PERFORMANCE THEREOF UNDER ANY WARRANTY, NEGLIGENCE, STRICT LIABILITY OR OTHER THEORY WILL BE LIMITED EXCLUSIVELY TO THE REMEDIES SPECIFIED IN SECTION 5.3 ABOVE. Some jurisdictions do not allow the exclusion of implied warranties or limitations on how long an implied warranty may last, so the above limitations may not be applicable. This warranty provides specific legal rights. Other rights vary from state to state. Licensee acknowledges and agrees that Legato is not responsible for and will have no liability for hardware, software, or other items or services provided by persons other than Legato.

8. LIMITATION OF LIABILITY

8.1 Limitation of Liability. NOTWITHSTANDING ANYTHING ELSE IN THIS AGREEMENT OR OTHERWISE, AND EXCEPT FOR BODILY INJURY, LEGATO WILL IN NO EVENT BE LIABLE OR RESPONSIBLE FOR OR OBLIGATED WITH RESPECT TO THE SUBJECT MATTER OF THIS AGREEMENT UNDER ANY CONTRACT, NEGLIGENCE, STRICT LIABILITY, OR OTHER LEGAL CONTRACTUAL OR EQUITABLE THEORY FOR: (I) ANY INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND WHETHER OR NOT ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH DAMAGES; OR (II) DAMAGES FOR LOST PROFITS OR LOST DATA; OR (III) COST OF PROCUREMENT OF SUBSTITUTE GOODS, TECHNOLOGY, SERVICES, OR RIGHTS; OR (IV) FOR ANY MATTER BEYOND ITS REASONABLE CONTROL. NOTWITHSTANDING ANYTHING IN THIS AGREEMENT TO THE CONTRARY OR THE FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITATION OF LIABILITY OR LIMITED REMEDY, LEGATO'S ENTIRE AGGREGATE LIABILITY ARISING FROM OR RELATING TO THIS AGREEMENT OR THE SUBJECT MATTER HEREOF, UNDER ANY LEGAL THEORY (WHETHER IN CONTRACT, TORT, INDEMNITY OR OTHERWISE) SHALL BE LIMITED TO THE AMOUNTS RECEIVED BY LEGATO FOR THE SOFTWARE DURING THE SIX (6) MONTH PERIOD PRIOR TO THE DATE THE CLAIM AROSE FOR THE PARTICULAR LEGATO SOFTWARE THAT CAUSED THE LIABILITY.

9. MISCELLANEOUS

9.1 Governing Law. This Agreement shall be governed by the laws of the State of California, as applied to agreements entered into and to be performed entirely within California between California residents, without regard to the principles of conflict of laws or the United Nations Convention on Contracts for the International Sale of Goods.

9.2 Government Restricted Rights. This provision applies to Software acquired directly or indirectly by or on behalf of any government. The Software is a commercial software product, licensed on the open market at market prices, and was developed entirely at private expense and without the use of any government funds. Any use, modification, reproduction, release, performance, display, or disclosure of the Software by any government shall be governed solely by the terms of this Agreement and shall be prohibited except to the extent expressly permitted by the terms of this Agreement, and no license to the Software is granted to any government requiring different terms.

9.3 U.S. Government Users. As defined in FAR section 2.101, DFAR section 252.227-7014(a)(1), DFAR section 252.227-7014(a)(5), or otherwise, all software and accompanying documentation provided in connection with this Agreement are "commercial items," "commercial computer software" and/or "commercial computer software documentation." Consistent with DFAR section 227.7202 and FAR section 12.212, any use, modification, reproduction, release, performance, display, disclosure, or distribution thereof by or for the U.S. Government shall be governed solely by the terms of this Agreement and shall be prohibited except to the extent expressly permitted by the terms of this Agreement. Licensee will ensure that each copy used or possessed by or for the government is labeled to reflect the foregoing.

9.4 Export and Import Controls. Regardless of any disclosure made by Licensee to Legato of an ultimate destination of the Products, Licensee will not directly or indirectly export or transfer any portion of the Software, or any system containing a portion of the Software, to anyone outside the United States (including further export if Licensee took delivery outside the U.S.) without first complying with any export or import controls that may be imposed on the Software by the U.S. Government or any country or organization of nations within whose jurisdiction Licensee operates or does business. Licensee shall at all times strictly comply with all such laws, regulations, and orders, and agrees to commit no act which, directly or indirectly, would violate any such law, regulation or order.

9.5 Assignment. This Agreement may not be assigned or transferred by Licensee without the prior written consent of Legato which shall not be unreasonably withheld. Legato may assign or otherwise transfer any or all of its rights and obligations under this Agreement upon notice to Licensee.

9.6 Sole Remedy and Allocation of Risk. Licensee's sole and exclusive remedy is set forth in this Agreement. This Agreement defines a mutually agreed-upon allocation of risk, and the License price reflects such allocation of risk.

9.7 Equitable Relief. The parties agree that a breach of this Agreement adversely affecting Legato's intellectual property rights in the Software may cause irreparable injury to Legato for which monetary damages may not be an adequate remedy and Legato shall be entitled to equitable relief in addition to any remedies it may have hereunder or at law.

9.8 No Waiver. Failure by either party to enforce any provision of this Agreement will not be deemed a waiver of future enforcement of that or any other provision, nor will any single or partial exercise of any right or power hereunder preclude further exercise of any other right hereunder.

9.9 Severability. If for any reason a court of competent jurisdiction finds any provision of this Agreement, or portion thereof, to be unenforceable, that provision of the Agreement will be enforced to the maximum extent permissible so as to effect the intent of the parties, and the remainder of this Agreement will continue in full force and effect.

10. ENTIRE AGREEMENT

10.1 This Agreement sets forth the entire understanding and agreement between the parties and may be amended only in a writing signed by authorized representatives of both parties. No vendor, distributor, dealer, retailer, sales person, or other person is authorized by Legato to modify this Agreement or to make any warranty, representation, or promise which is different than, or in addition to, the warranties, representations, or promises made in this Agreement.

Contents

Preface	11
Audience.....	11
Product Documentation	11
Conventions.....	12
Information and Services.....	13
General Information	13
Technical Support	13
Licensing and Registration.....	14
Customer Feedback	14
Chapter 1: Introduction.....	15
The Importance of Backing Up Microsoft SQL Server Data	15
Using Legato NetWorker Software to Protect SQL Server Data.....	16
Using the NetWorker Module to Protect SQL Server Data	16
How the NetWorker Module Works with SQL Server Data	17
Security Requirements.....	18
Client/Server Model.....	19
SQL Server Data Backups	21
SQL Server Data Restoration	22
Chapter 2: NetWorker Module Backup.....	25
Backup Types.....	26
Scheduled Backups	27
Unscheduled Backups	27
Considerations for Filegroup Backup Performance	28
Reduced Number of Files Needed to Restore a Damaged Disk.....	29
Filegroup Backup Versus Full Database Backup.....	29

Backup Levels.....	29
Microsoft SQL Server 2000 Enhancements	30
How to Specify Backup Levels	31
Combining Data Objects to Create Backup Levels	33
Promoting Backup Levels	34
Backup Strategies.....	35
Strategies for Scheduling Backups.....	36
Transaction Log Maintenance.....	37
How to Prevent Log Overflow	38
Backup Groups	40
Backup Group Considerations.....	41
Volume Pools	41
How to Create a Volume Pool	41
How to Customize a Volume Pool	43
Browse and Retention Policies	44
NetWorker Backup Clients	45
How to Add a SQL Server Host to the NetWorker Server's Client List.....	45
Backup Operation Status	47
Chapter 3: NetWorker Module Restore	49
Microsoft SQL Server 2000 Recovery Models	50
Full Recovery Model	51
Bulk_Logged Recovery Model.....	51
Simple Recovery Model.....	52
Specifying Database Recovery Models	52
Changing the Recovery Model for a Database.....	53
How to Change from Full or Bulk_Logged to Simple	53
How to Change from Bulk_Logged to Full	53
How to Change from Simple to Full or Bulk_Logged	54
Restore Types.....	54
Normal Restore Type	55

Partial Restore Type	55
Verify-Only Restore Type	55
Copy Restore Type	56
Restore Modes	56
Normal Restore Mode	56
No Recovery Restore Mode	56
Standby Restore Mode	56
Restore Time	57
Database Backup Versions	57
Point-in-Time Restore of a Transaction Log Backup	58
Microsoft SQL Server 2000 Named Log Marks	58
Specifying the Restore Time	59
NetWorker Module Database Restoration	59
SQL Server Master Database Maintenance	60
NetWorker Module Interaction with SQL Server Services	60
Chapter 4: Using the NetWorker User for SQL Server Program	63
Browsing Data Items	63
Restore Browse Window Restrictions	64
Fake Objects	64
Display Conventions	65
Marking Items	66
Restore Browse Window Restrictions	68
Marking Semantics and Restrictions	68
Chapter 5: Backing Up SQL Server Data	71
Preparing for the Backup	71
Creating the Backup	72
How to Back Up Using NetWorker User for SQL Server	73
Backup Options	75

How to Specify General Backup Options	76
How to Specify Advanced Backup Options	78
Database Item Backup Properties	79
Target Volume Pools for Unscheduled Backups	80
Assigning Pools in NetWorker User for SQL Server	81
Example:	82
Chapter 6: Restoring SQL Server Data	83
Before the Restore	83
Specifying Database Restore Properties	84
How to Specify Database Item Restore General Properties	85
How to Specify Standby Undo File Properties	87
How to Specify Database Item Files Properties	87
Modifying a File Destination	91
How to Specify the Restored File Destination	92
Read File Configuration Properties	93
How to Specify Read File Configuration Properties	93
Database File Relocation Restrictions	94
How to Specify Database Item Restore Time Properties	95
How to Specify Point-in-Time Restore Properties	97
Backing Up the Active Portion of the Transaction Log	99
Filegroup and File Restore Properties	100
How to View the Restore Properties	100
Performing the Restore	102
Chapter 7: Special Features	109
Database Consistency Check	109
How to Perform a Database Consistency Check	110
Striping	111
How to Specify Striping	111
Configuring Striping for Backups	112

Performance Considerations for Striping.....	112
How to Perform a Striped Backup	113
Performing a Striped Recovery	113
Optimal Striped Recovery Versus Fail-Safe Striped Recovery.....	114
Performing an Optimal Striped Recovery (Best-Case Scenario).....	114
Performing a Fail-Safe Striped Recovery (Worst-Case Scenario).....	115
How to Perform a Fail-Safe Striped Recovery on SQL Server 7.0	115
How to Perform a Fail-Safe Striped Recovery on SQL Server 2000	116
Diagnosing Errors During Backup and Recovery	118
RAW Save Set.....	118
Overview	119
Key Benefits	119
Supported NetWorker Server Features.....	120
Supported Configurations	122
RAW Save Set Restrictions	122
Recommended Configuration	123
Performing a RAW Save Set Backup or Restore.....	123
RAW Save Set Notifications.....	124
Interleaving	124
Backing Up and Restoring Multiple Nonclustered Instances of SQL Server	125
Backing Up and Restoring Microsoft Cluster Server Configurations	126
How NetWorker Works with MSCS Clusters.....	127
NetWorker Module 3.1 Support for MSCS Clusters.....	128
Named Instances in Failover Cluster Configurations	128
Configuring Scheduled Backups of a Clustered SQL Server.....	128
Performing Manual Operations on a Clustered SQL Server	129
Operations from the Command Line	129
Operations from the NetWorker User for SQL Server Program	129
Active/Passive Cluster Configurations.....	130
Active/Active Cluster Configurations	130
NetWorker Module Interaction with SQL Server Services	131

Chapter 8: Disaster Recovery	133
When to Reinstall Microsoft SQL Server.....	136
Disaster Recovery of SQL Server Primary Data Disk Crash or Data Corruption.....	136
Disaster Recovery of SQL Server Binary Disk Crash or Data Corruption.....	137
Disaster Recovery of SQL Server and NetWorker Server	139
Disaster Recovery of SQL Server (Shortcut)	140
Disaster Recovery of SQL Server	142
Appendix A: Commands and Command Options.....	143
NetWorker Module Commands.....	143
Using the nsrsqlrc Command.....	144
The nsrsqlrc Command Options.....	145
Using the nsrsqlsv Command	151
The nsrsqlsv Command Options	152
Using the nwmssql Command	156
nwmssql Command Options.....	156
Backup and Restore Command Syntax for SQL Server Data	156
Syntax for a Named Instance Configuration for SQL 2000	158
Syntax for Names Containing a Period, Backslash, or Colon	159
Glossary.....	163
Index.....	173

Preface

The *Legato NetWorker Module for Microsoft SQL Server Administrator's Guide* contains information on how to configure and manage the Legato NetWorker[®] Module for Microsoft SQL Server software.

You *must* install the NetWorker Module software on your server and clients to use the information presented in this guide. If you have not yet installed the software, refer to the *Legato NetWorker Module for Microsoft SQL Server Installation Guide* for installation instructions.

Audience

The information in this guide is intended for system administrators who are responsible for performing backup and restoration of data from Microsoft[®] SQL Server[™] versions 7.0 and 2000. Administrators should have a working knowledge of the technical concepts related to data storage, data retrieval, and networking, and be familiar with how the Legato NetWorker product supports these technologies. In addition, you may find it helpful to have access to the Microsoft SQL Server, Windows[®], and Legato NetWorker documentation sets when performing backup and restore procedures.


Product Documentation

Legato offers an extensive archive of product documentation at its web site www.legato.com. Most of the documents are in Adobe[®] Acrobat[®] Portable Document Format (PDF), and can be viewed by downloading and installing the Adobe Acrobat Reader. The Reader is available in the */viewers/acroread* directory on the Legato Documentation Suite CD-ROM, or directly from

Adobe at *www.adobe.com*. To install and use the Reader on the platform of your choice, refer to the instructions in the CD-ROM's */viewers/acroread/readme.txt* file or at the Adobe web site.

Conventions

This document uses the following typographic conventions and symbols to make information easier to access and understand.

Convention	Indicates	Example
boldface	Names of DOS or UNIX line commands, daemons, options, programs, or scripts	The nsradmin command starts the command line version of the administration program.
<i>italic in text</i>	Pathnames, filenames, computer names, new terms defined in the Glossary or within the chapter, or emphasized words	Displayed messages are also written to <i>/nsr/logs/daemon.log</i> .
<i>italic in command line</i>	A variable that you need to provide in the command line	nwadmin -s <i>server-name</i>
fixed-width	Examples and information displayed on the screen	media waiting: recover waiting for 8mm 5GB tape volume name
fixed-width, boldface	Commands and options that you must type exactly as shown	nsr_shutdown -a
Menu_Name> Command	A path or an order to follow for making selections in the GUI	Volume>Change Mode>Appendable
Important:	Information that you must read and follow to ensure successful backup and recovery of your data	 Important: Use the no_verify flag with extreme caution.

Information and Services

Legato offers a variety of methods, including electronic, telephone, and fax support for you to obtain company, product, and technical information.

General Information

The Legato web site provides most of the information that you need. If you want technical bulletins and binary patches, you can also access the Legato FTP site. If you have specific sales or training needs, e-mail or call Legato.

Legato Service or Resource	Technical Bulletins	Binary Patches	Company & Product Information	Training Programs
<i>www.legato.com</i>	Yes	Yes	Yes	Yes
<i>ftp.legato.com</i> (log in as <i>anonymous</i>)	Yes	Yes		
Legato Sales (650) 210-7000 (option 1) <i>sales@legato.com</i>			Yes	
Legato Education Services (650) 842-9357 <i>training@legato.com</i>				Yes

Technical Support

The Support section of the Legato web site provides contact information, software patches, technical documentation, and information about available support programs.

- If you have an active support agreement, you may access TechDialog, the Legato integrated product knowledge base. You may also contact Legato Technical Support for help with Legato software issues.
- If you do not have an active support agreement, contact LegatoCare to purchase annual Software Update Subscriptions, Legato Technical Support services, or per-update/per-incident support.

Licensing and Registration

To license and register your Legato products, go to the Legato licensing web site. To change contact information, transfer licenses, or ask questions about licensing, contact Legato using one of the following methods.

Licensing and Registration	Contact
Legato licensing web site	http://license.legato.com
Telephone number	(650) 812 6000 (option 3, option 2) ^a +31 23 554 8881 ^b
Fax number	(650) 745-1477 ^a +31 23 554 8808 ^b
E-mail	licensing@legato.com ^a licensingemea@legato.com ^b

a. Contact information for Americas, Asia, and Pacific.

b. Contact information for Europe, Middle East, and Africa.

Customer Feedback

Legato welcomes your comments and suggestions about software features, the installation procedure, and documentation. Please send any suggestions and comments to feedback@legato.com. You will receive a notice confirming receipt of your e-mail. Although we cannot respond personally to every request, we consider all your comments and suggestions during product design.

Help us improve our documentation and be eligible to win a prize by completing a brief survey. Visit the Legato web site at www.legato.com, go to the Documentation page under the Support section, and look for the survey link.

Chapter 1: Introduction



The NetWorker Module for Microsoft SQL Server software provides services that enable the NetWorker Server backup and restore functionality for Microsoft SQL Server data. The NetWorker Module software integrates this capability into the comprehensive network-wide data protection solution that NetWorker software provides on the Microsoft Windows and UNIX[®] operating systems.

This section gives an overview of how to use NetWorker and NetWorker Module software together to back up and restore SQL Server data.

The Importance of Backing Up Microsoft SQL Server Data

In a client/server environment, data can be lost due to hardware failures, software bugs, procedural flaws, and user errors. A viable database backup strategy can help restore data in the event of such a disaster.

There is considerable risk in failing to perform regular backups of databases and frequent backups of transaction logs. If all transaction logs are lost, a database can be restored only to the time of its last full backup. Without backups or transaction logs, the database cannot be restored at all.

Without a centralized storage management solution, backups of mission-critical data must be implemented at the local system level. Local backups do not provide a cohesive policy for managing and protecting vital, enterprise-wide data.

Using Legato NetWorker Software to Protect SQL Server Data

Legato NetWorker software is a high-capacity, easy-to-use management solution for network data storage that provides data backup and restoration for heterogeneous networks of servers and clients. Using this product for storage management simplifies the process and reduces administrative burden by automating and centralizing data storage operations.

The NetWorker product can:

- Perform automated backups during off-peak hours.
- Use centralized administration to configure, monitor, and control backups from anywhere on a network.
- Automate tape handling tasks using Legato GEMS[®] SmartMedia[®], the intelligent media manager that Legato provides.
- Increase backup performance by simultaneously sending data from multiple clients to the storage management server.
- Use concurrent device support to direct data streams to multiple backup devices for even greater speed.
- Clone backup media for local and off-site storage strategies.

Using the NetWorker Module to Protect SQL Server Data

A filesystem backup does not save SQL Server data in a form that Microsoft SQL Server can restore. For a restorable save of SQL Server data, it must be backed up using one of the methods described in this *Administrator's Guide*.

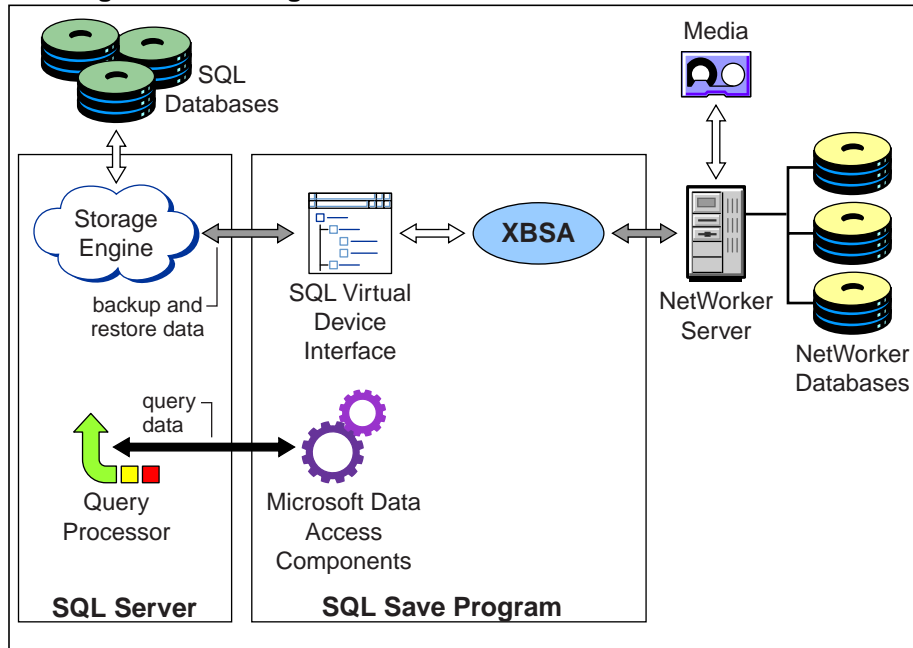
The NetWorker Module software is an add-on that provides automated storage management for Microsoft SQL Server databases. This module provides the following capabilities for managing the SQL Server storage hierarchy:

- Scheduled automated backups
- Unscheduled backups that can be initiated any time
- Automatic backup of the active portion of a transaction log
- SQL Server data restoration from various levels and types of backups
- Data retrieval options, including point-in-time, copy, and partial restore
- Support for Microsoft Cluster Server (MSCS) cluster configurations
- Support for SQL Server 2000 multiple instance configurations

How the NetWorker Module Works with SQL Server Data

The SQL Server software can reside on the same system as the NetWorker server software or on a separate system. However, because the NetWorker Module save program substitutes for the NetWorker server client save program during SQL Server data backup, the NetWorker Module software must be installed on the system where SQL Server and the NetWorker client are installed. No matter which computer acts as the SQL Server host, that computer must be configured as a storage management client of the NetWorker server. [Figure 1](#) shows how Legato NetWorker software works with SQL Server. For details about NetWorker software and NetWorker Module configuration, refer to the *Legato NetWorker Module for Microsoft SQL Server Installation Guide*.

Figure 1. How Legato NetWorker Software Works with SQL Server



Security Requirements

Using the NetWorker Module to back up Microsoft SQL Server databases requires that the proper privileges be granted to the NetWorker Module process. Microsoft SQL Server imposes the following requirements on third-party backup products such as the NetWorker Module:

- The logon account that the third-party backup process uses to connect to SQL Server must be a member of the SQL Server *sysadmin* role in order to issue the T-SQL BACKUP query.
- The Windows logon account under which the third-party backup process is running must be a member of the SQL Server *sysadmin* role in order to open a shared memory handle when initializing the Microsoft Virtual Device Interface (VDI).

Microsoft Windows imposes the following requirement on third-party products such as the NetWorker Module:

- The Windows logon account under which the third-party backup process is running must be a member of the Microsoft Cluster Server (MSCS) Administrators group in order to open a cluster handle when accessing resources managed by MSCS.

NetWorker Module imposes the following requirements:

- When using the **save** and **restore** commands, the Windows logon account must be a member of the SQL Server *sysadmin* role.
- NetWorker User for SQL Server must be a member of the local Backup Operators group.
- NetWorker User for SQL Server must be a member of the local Administrators group.

Microsoft SQL Server provides two authentication modes: Windows Authentication Mode and Mixed Mode (Windows Authentication and SQL Server Authentication). Regardless of the authentication mode, you must add the Windows logon account to the SQL Server *sysadmin* role under which the NetWorker Module will run.

When running in Windows Authentication Mode, the Windows logon account that the NetWorker Module process uses must be a member of the SQL Server *sysadmin* role. The administrator and BUILTIN\administrator accounts are automatically members of the *sysadmin* role.

When running in mixed mode and using a SQL Server logon account to connect to Microsoft SQL Server, then the SQL Server logon account must be a member of the *sysadmin* role. For more information on how to add members to a SQL Server role, refer to the Microsoft SQL Server Books Online.

To ensure that scheduled backups run with the proper security context, the Client resource must be configured for the SQL Server host using the NetWorker Administrator program. For instructions on how to configure Client resources, refer to "Chapter 1: Installing NetWorker Module for Microsoft SQL Server" in the *Legato NetWorker Module for Microsoft SQL Server Installation Guide*.

When creating the Client resource for the SQL Server host, specify the Windows logon account name and the logon account password for the Remote User and Password attributes (under the Remote tab). For the SQL Server logon account, enter the appropriate command using the **-U** and **-P** options in the Backup Command attribute text box. Be aware that this SQL Server logon account password is visible in plain text and may present a security risk. For existing Client resources, specify both the Windows logon account information and the SQL Server logon account in the same manner in the Edit Client dialog box.

Client/Server Model

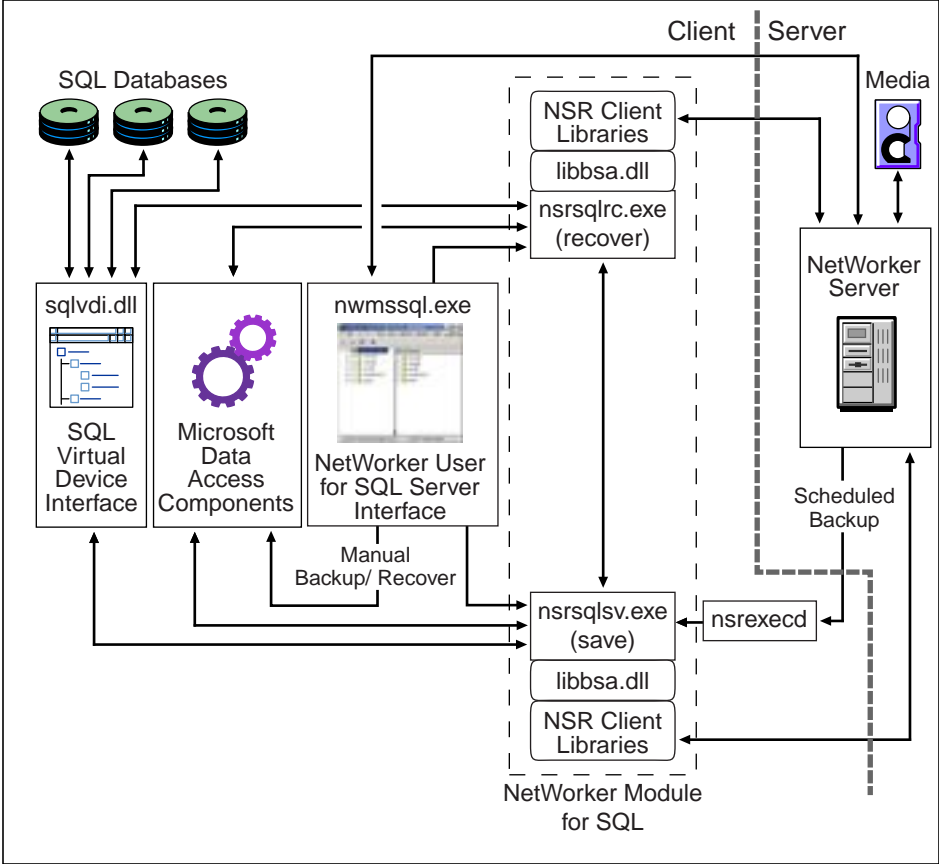
NetWorker software uses a client/server model to provide storage management services. At least one computer on the network is designated to serve as a NetWorker server host. Computers with data to back up are configured as clients of the NetWorker server. The NetWorker server program can be used on Windows NT, Windows 2000, or UNIX environments.

The NetWorker server performs the following functions:

- Maintains the resource information
- Contacts clients listed in a backup group configured on the server
- Runs manual backups when a client request is received
- Restores data upon request from the client

In addition, the NetWorker server maintains the online client file index and media database. See [Figure 2 on page 20](#) for details on how the NetWorker Module software works with SQL Server.

Figure 2. How the NetWorker Module Software Works with SQL Server



SQL Server Data Backups

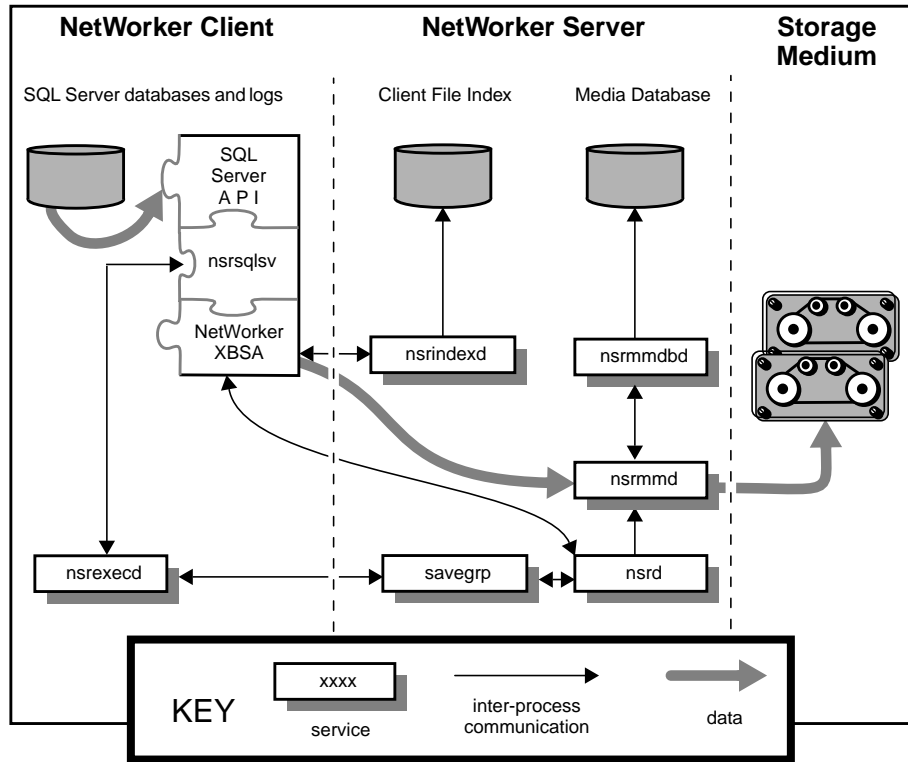
Figure 3 on page 22 shows the functional relationship between the NetWorker server, the NetWorker Module, and the SQL Server products during a scheduled backup. In general, the NetWorker server performs the scheduling of storage management tasks, while the NetWorker Module passes the data from the SQL Server to the NetWorker server.

When a scheduled backup for a SQL Server client occurs, the following events take place:

1. The backup is triggered by the **nsrd** function on the NetWorker server.
2. The NetWorker **savegrp** function executes the NetWorker Module backup save.
3. The NetWorker Module **nsrsqlsv** function sends the backup data received from the SQL Server to the NetWorker server through an X/Open[®] Backup Services Application Programming Interface (XBSA).
4. The results of the **savegrp** are sent back to the NetWorker server and are included in the savegroup completion report. This report is sent according to the notification set up in the NetWorker Administrator program for the Savegroup event.

For information about the NetWorker services and operations described in this chapter, refer to the *Legato NetWorker Administrator's Guide*.

Figure 3. NetWorker Client/Server Interaction During a Backup Initiated by the NetWorker Module



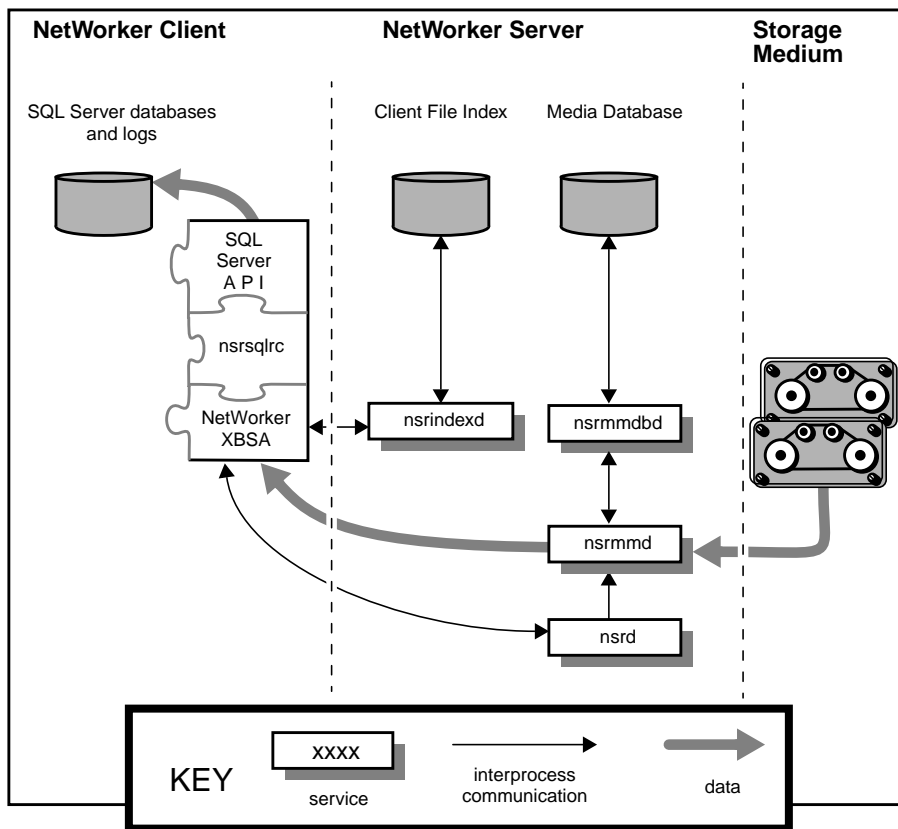
SQL Server Data Restoration

Figure 4 on page 23 shows the functional relationship between the NetWorker server, the NetWorker Module, and the SQL Server products during a restore. When the NetWorker Module software initiates a restore, the following events occur:

1. The NetWorker Module **nsrsqlrc** function initiates the restore.
2. The NetWorker software translates the object names requested into a NetWorker-compatible format and forwards them to the NetWorker **nsrd** service.

3. The NetWorker media service, **nsrmmd**, invokes **nsrmmdbd** to search the online media database for the media containing the objects requested.
4. After the media is mounted, **nsrmmd** sends the data to **nsrsqlrc**, which restores the data to SQL Server using the **restore** command and the SQL Server VDI.

Figure 4. NetWorker Client/Server Interaction During a Restore Initiated by NetWorker Module



1

Chapter 2: NetWorker Module Backup

Before configuring the NetWorker Administrator program to work with the NetWorker Module, it is important to understand the different types and levels of backups that can be created, and how the type of backup created affects the options for restoring data. This chapter includes the following sections:

- ["Backup Types" on page 26](#)
- ["Considerations for Filegroup Backup Performance" on page 28](#)
- ["Backup Levels" on page 29](#)
- ["Backup Strategies" on page 35](#)
- ["Transaction Log Maintenance" on page 37](#)
- ["Backup Groups" on page 40](#)
- ["Volume Pools" on page 41](#)
- ["Browse and Retention Policies" on page 44](#)
- ["NetWorker Backup Clients" on page 45](#)
- ["Backup Operation Status" on page 47](#)



Important: It is *highly recommended* that you read this chapter and ["Chapter 3: NetWorker Module Restore" on page 49](#), before performing a backup or restore operation using this NetWorker Module. The type and level of backup created impacts the type of restore that can be performed. For example, regularly scheduled, full-level backups are required to enable recovery from a disaster situation.

Backup Types

Using the NetWorker server with this NetWorker Module, save SQL Server data by performing one of the following types of backups:

- **Scheduled (automatic) backups** — Scheduled backups are configured in the NetWorker Administrator program to start automatically from the NetWorker server. Scheduled backups can be configured to run at any time, using any backup level.

When running a scheduled backup from NetWorker Administrator, SQL Server data residing on any host that is a client of the NetWorker server can be backed up.

- **Unscheduled (manual) backups** — This type of backup is immediate and is sometimes referred to as “on-demand” or “ad hoc”.

Unscheduled backups can be run at any time:

- From a Windows command prompt. A user or administrator can initiate an unscheduled SQL Server backup by logging on to the NetWorker client host with an account that has SQL Server administrative privileges. A manual backup can be run from a command prompt at the full, differential, or incremental (transaction log) level.
- From the NetWorker User for SQL Server program, the NetWorker Module graphical interface. Only full level backups restore procedures can be run from this interface.

When an unscheduled backup is run from a NetWorker server that is a client host, only data stored on that SQL Server host can be backed up.



Important: Performing regular, *scheduled* NetWorker backups of SQL Server databases is strongly recommended for maximum data protection.

The following sections describe the backup types in greater detail and the considerations for filegroup performance of large databases with minimal transaction activity.

Scheduled Backups

Only regular, scheduled backups ensure that *all* of SQL Server data is automatically saved, including the NetWorker server's client indexes and bootstrap file. The indexes and bootstrap file are vital for restoring data to the SQL Server in the event of a disaster.

To customize a scheduled backup of Microsoft SQL Server data:

- Configure the following resources in the NetWorker Administrator program:
 - Backup groups; see ["Backup Groups" on page 40](#)
 - Volume pools; see ["Volume Pools" on page 41](#)
 - Backup schedules; see ["Strategies for Scheduling Backups" on page 36](#)
 - Browse and retention policies; see ["Browse and Retention Policies" on page 44](#)
- Configure a NetWorker client. For information on configuring NetWorker clients, refer to the *Legato NetWorker Administrator's Guide*.

Note: If custom resource settings are not created in the NetWorker Administrator program before defining the SQL Server host as a NetWorker client, the NetWorker server uses its default resource settings when backing up data. For more information on default resource settings, refer to the *Legato NetWorker Administrator's Guide*.

Unscheduled Backups

When performing an unscheduled backup of SQL Server data, the NetWorker Module you can select the following combinations of data objects:

- The entire SQL Server storage hierarchy
- One or more entire *databases*
- One or more *filegroups* in a SQL Server 7.0 or later database
- One or more *files* in a SQL Server 7.0 or later filegroup
- A heterogeneous collection of databases, filegroups, and files
- Transaction log backups

The storage hierarchy is defined as the database storage components exposed to third-party backup vendors by the SQL Server Storage Engine. The storage components include the database, filegroup, file, and transaction log.

To create a backup, select the objects to back up and specify the level of backup to perform. If a backup level is not specified, a level full backup is created.



Important: When using the NetWorker Module to manually back up specific files and filegroups from a SQL Server 7.0 or later database, it is possible to also restore from a media failure that affects only a subset of the data in the databases. A SQL Server 7.0 or 2000 filegroup or file can *also* be restored from a *full database* backup, if the backup was created using release 3.0 or later of the NetWorker Module.

When performing an unscheduled backup of the SQL Server, the client indexes and bootstrap file are not automatically backed up. See "[Target Volume Pools for Unscheduled Backups](#)" on page 80 for information on creating client indexes and a bootstrap file following a manual backup.

Considerations for Filegroup Backup Performance

For large databases with little transaction activity, filegroup backups and restores can provide advantages, such as increasing the speed of restoration. However, in some cases the total time required for restoration can take as long as restoring the entire database. This is because after a filegroup is restored, the NetWorker Module must automatically also restore *all* transaction log backups performed since the filegroup backup. (This is a Microsoft SQL Server requirement.) If the transaction logs are spread out over a number of days, the logs could be stored on multiple backup media volumes. The time required to change media volumes to restore each transaction log might be much greater than the time it would take to restore the original full database from a single media volume. Of course, this problem might also be encountered any time a database is to be restored to a given point in time after a series of incremental (transaction log) backups.

Reduced Number of Files Needed to Restore a Damaged Disk

With SQL Server a database can be stored in multiple filegroups and files in order to spread the database over more than one disk or disk partition. With the NetWorker Module, distinct filegroups and files can be backed up and restored without having to restore entire databases. If a disk is damaged, only those files on the damaged disk or disk partition may need to be restored.

Filegroup Backup Versus Full Database Backup

A filegroup backup might also be used if there is not enough time to create a backup of the entire database. For example, if a database backup takes three hours to create but only two hours per night are available for backups, some of the files can be backed up one night and the rest of the files can be backed up the following night.

Backup Levels

The NetWorker product provides three main levels of backup: full, incremental, and differential, where differential is specified as any level from 1 to 9.

The NetWorker Module also provides full, incremental, and differential backups.

- An incremental backup, done from the command line, corresponds to a Microsoft SQL Server transaction log backup. A log file backup by itself cannot be used to restore a database. A log file is used after a database restore to recover the database to the point of the original failure.
- A differential backup, specified as any level from 1 to 9, is done from the command line, and makes a copy of all the pages in a database modified after the last full database backup.

Table 1 shows how the terminology for backup levels used in the Microsoft SQL Server product differs from the terminology used in the NetWorker Module product.

Table 1. NetWorker Module Backup Levels

Function	NetWorker Module Term	SQL Server Term
Backs up an entire database, filegroup, or file.	Full database, filegroup, or file backup (full)	Database, filegroup, or file backup
Backs up all transaction log changes since the most recent full, differential, or transaction log backup.	Incremental database backup (incr)	Transaction log (also called <i>xlog</i>) backup
Backs up all database changes, including filegroups and files, since the last full backup.	Database, filegroup, or file differential backup (diff)	Differential backup

Microsoft SQL Server 2000 Enhancements

Microsoft SQL Server 2000 provides enhanced support for backing up and restoring filegroups and files. In addition to creating a level full filegroup or file backup, SQL Server 2000 supports the creation of filegroup differential and file differential backups. Performing a filegroup differential backup may actually reduce both media requirements and restore time; the data can be stored across more than one disk or disk partition, and restore time may be reduced because a differential can substitute for any log backups performed between the full and differential backups (the full backup must be performed first).



Important: If a backup was created using NetWorker Module, release 3.0 or later, a SQL Server 7.0 or 2000 filegroup or file can also be restored from a *full database* backup.

The NetWorker Module allows you to specify these new backup levels, along with database full, database differential, and database incremental. The availability of a backup level depends on the type of data selected for backup and any SQL Server settings on those objects. [Table 2](#) identifies the supported backup levels.

Table 2. Backup Levels for SQL Server Data

SQL Server Data Objects	SQL Server 7.0			SQL Server 2000		
	full	diff	incr	full	diff	incr
All databases, including MSSQL:	yes	yes	yes ^{c,d}	yes	yes ^c	yes ^{c,d}
Specified databases	yes	yes ^c	yes ^{a,c}	yes	yes ^c	yes ^{a,c}
All filegroups in specified databases	yes ^{a,c}	no ^{a,b,c}	n/a	yes ^{a,c}	yes ^{a,c}	n/a
Specified filegroups in specified database	yes ^{a,c}	no ^{a,b,c}	n/a	yes ^{a,c}	yes ^{a,c}	n/a
Specified files in specified filegroups in specified databases	yes ^{a,c}	no ^{a,b,c}	n/a	yes ^{a,c}	yes ^{a,c}	n/a

a. An incremental database, filegroup, or file backup can only be created when the SQL Server database options are properly configured. For more information, refer to the Microsoft SQL Server 7.0 or 2000 documentation.

b. SQL Server 7.0 does not allow differential (level 1) backups of filegroups or files. If a differential filegroup or file backup is specified, NetWorker Module performs a full backup.

c. Individual items are subject to promotion. For more information, see "[Promoting Backup Levels](#)" on page 34.

d. New in NetWorker Module for Microsoft SQL Server release 3.1.

How to Specify Backup Levels

Depending on the type of backup being performed, specify the backup level in *one* of the following ways:

- For a scheduled backup, set the level in the Set Level dialog box in the Manage Schedules resource of the NetWorker Administrator program. The backup level specified in this dialog box takes effect whenever that scheduled backup runs.

- For an unscheduled backup, enter the NetWorker **nsrsqlsv** command from a Windows command prompt on the NetWorker server. The backup level can be included as an option to the command. If the backup level is not indicated, a full backup is performed.

For example, to initiate an unscheduled incremental (transaction log) level backup, enter the following:

```
nsrsqlsv [-s NetWorker_server_hostname] -l incr
```

For more information about **nsrsqlsv** and backup level options, see ["Appendix A: Commands and Command Options" on page 143](#).

Note: An unscheduled backup started from the NetWorker User for SQL Server program can only be performed at level full.

[Table 3](#) summarizes where the NetWorker Module backup procedures can be initiated and which backup levels are supported for each interface.

Table 3. Where to Initiate Backup Operations

Backup Type	Backup Initiated From	Backup Levels Available		
		full	incr	diff
Scheduled	NetWorker Administrator program on the server	Yes	Yes	Yes
Unscheduled	Command line on the NetWorker server that is a client host	Yes	Yes	Yes
	NetWorker User for SQL Server program on the client	Yes	No ^a	No

a. When restoring a database filegroup or file from NetWorker User for SQL Server, you can request that the active portion of the transaction log be backed up prior to the restore. For more information, see [step 2](#) under ["How to Specify Database Item Files Properties" on page 87](#).

Combining Data Objects to Create Backup Levels

In addition to backup levels, the NetWorker Module allows you to select SQL Server data objects in various combinations to create backups of different levels, as shown in [Table 4](#).

Table 4. Creating Additional Backup Levels with Data Objects

Type of Backup	Procedure
Full database backup	Select one or more databases to create a level full database backup of the selected databases and their transaction log files.
Full filegroup or full file backup	<p>Select one or more filegroups, or one or more files, to create a level full filegroup or file backup of the selected filegroups or files, but not their transaction logs.</p> <p>When performing an unscheduled level full backup of a filegroup or file, also perform a database incremental level backup to maintain the validity of the transaction log</p>
Database incremental (transaction log) backup	<p>Select one or more databases to create a database incremental level backup of <i>only</i> the transaction log for the selected databases.</p> <p>The SQL database must be preconfigured to allow transaction log backups.</p>
Database differential (level 1) backup	Select one or more databases to create a database level differential backup of <i>only</i> the changes made to the selected databases since the last full level backup was created.
Filegroup or file differential backup	For SQL Server 2000 only, select one or more filegroups or one or more files to create a filegroup or file level differential backup of <i>only</i> the changes made to the selected filegroups or files since the last full level backup was created.

Promoting Backup Levels

Microsoft SQL Server Best Practices guidelines indicate that a full database backup should be the first step in implementing a recovery strategy for a database. In adhering to these guidelines, the NetWorker Module now supports backup level *promotion*. Table 5 explains what prompts a promotion.

Table 5. Backup Level Promotion Process

Item	If Requested Level is...	Level is Promoted to...	Reason
Database	differential	full	<ul style="list-style-type: none"> Database full backup does not exist
Database	incremental	full	<ul style="list-style-type: none"> Database full backup does not exist. Database does not support incremental (transaction log backups)^a Database is currently in emergency mode.^a
Filegroup/File	full	database full	<ul style="list-style-type: none"> Database full backup does not exist.
Filegroup/File	differential	database full	<ul style="list-style-type: none"> Database full backup does not exist.
Filegroup/File, when connected to SQL Server version 7.0	differential	full	<ul style="list-style-type: none"> SQL Server version 7.0 does not support filegroup/file differential backups.
Filegroup/File	incremental	full	<ul style="list-style-type: none"> Filegroup/file incremental backups are not supported.

a. See the Microsoft SQL Server Books Online for more information.



Important: A database restore using the NetWorker Module will fail if:

- A full backup of the database was performed using the NetWorker Module.
- The backup was removed from the NetWorker Server by removing or recycling the volume.
- An incremental backup was then performed.

The restore fails because the NetWorker Module cannot find a full backup.

The NetWorker Module checks SQL Server to determine if a full database backup has been performed; it does not check to see if a full backup still exists on the NetWorker Server.

Backup Strategies

Legato NetWorker software allows you to set dates, times, and backup levels for automatic backups using a backup schedule. A backup schedule defines the backup level to be performed on each day of a backup cycle. A SQL Server backup schedule can be customized by specifying different backup levels on different days.

Because it may not be practical or efficient to run full backups every day, other backup levels can be specified in the NetWorker Module for automatic, scheduled backups. Limiting the frequency of full backups can decrease server load while ensuring data is protected. Consider the following when selecting backup levels:

- Full backups take more time than differential backups, while differential backups take more time than incremental (transaction log) backups. However, restoring data immediately following a differential backup is generally faster than restoring data following a number of successive transaction log backups.
- If you have only a stand-alone storage device and a full backup does not fit on a single media volume, an operator must monitor the backup so the volume can be changed at the appropriate time.
- An incremental backup saves only transactions that have occurred since the most recent full or transaction log backup. For this reason, using incremental (transaction log) backups can simplify and speed database restoration.

Table 6 outlines the basic differences between the backup levels:

Table 6. Backup Level Advantages and Disadvantages

Backup Level	Advantages	Disadvantages
Full	Fastest restore time	<ul style="list-style-type: none">• Slow backup• Longest, highest load on server• High load on client and network• Uses the most volume space
Incremental (transaction log)	<ul style="list-style-type: none">• Faster backup• Low load on server• Uses least volume space• Point-in-time restore	<ul style="list-style-type: none">• Slow restore• Data can spread across multiple volumes• Multiple transaction logs can spread across multiple volumes
Differential	Faster than a full backup and captures all changes since last full	Generally more time-consuming than a transaction log backup (depending on the backup schedule strategy)

Strategies for Scheduling Backups

Choosing a backup level usually entails trading off the time and number of volumes required for backup versus that required for restoration. Several preconfigured backup schedules are provided, or a customized schedule can be created to suit configuration requirements and resource capabilities. If the SQL Server manages a significant amount of data, schedule a backup of the databases every one to two weeks, as shown in Table 7.

Table 7. Full Backup Every One to Two Weeks

Fri	Sat	Sun	Mon	Tues	Wed	Thurs
full	incr	incr	incr	incr	diff	incr
incr	incr	incr	diff	incr	incr	incr
full	Repeat					

Another reasonable backup scheduling possibility would be to schedule incremental (transaction log) backups on several successive days immediately following the previous full backup, as shown in Table 8. This backs up all data that has changed since the previous incremental (transaction log) backup.

Table 8. Incremental (Transaction Log) Backup after a Full Backup

Fri	Sat	Sun	Mon	Tues	Wed	Thurs
full	incr	incr	incr	diff	incr	incr
Repeat						

A level 1 differential backup can also be scheduled after several days of incremental (transaction log) backups. This backs up all data since the previous full backup.

Note: Use the Group Control feature in the NetWorker Administrator program to start a scheduled backup at a time not specified in the schedule. For instructions on using this feature, refer to the section on monitoring and controlling backups in the *Legato NetWorker Administrator's Guide*.

For further details and examples of planning backup strategies and creating schedules, refer to the section on backup levels in the *Legato NetWorker Administrator's Guide*.

Transaction Log Maintenance

The NetWorker Module provides implicit and explicit methods for managing SQL Server database transaction logs:

- Implicit management is done using log backups to manage log space. This can occur when a backup schedule is implemented that includes **incr** (transaction log level) backups. It can also occur for manual saves at the command line when the **nsrsqlsv.exe** command option **-l incr** is specified.
- Explicit management is done by specifying the **nsrsqlsv.exe** command on the command line, with or without the **-T** option (Truncate Only) or the **-G** option (No Log). Both command options result in the log being truncated prior to the backup, and both options are compatible with level full and level 1 (differential) backups. Following completion of the backup, if the NetWorker Module determines that the database does not support log backups, and no **-T/-G** option is specified, it tells the SQL server to truncate the transaction log.

How to Prevent Log Overflow

In Windows, prevent database logs from overflowing available log space by creating an alert in the SQL Server Performance Monitor that forces an incremental backup when the database's log space reaches a certain capacity (for example, 80% full). An incremental (transaction log) backup truncates the logs and clears disk space.

To create this alert in Windows NT 4.0 and SQL Server 7.0:

1. Create a batch file (for example, *sqlalert.bat*) that contains the backup command and any appropriate command flags, for example:

```
installpath\nsr\bin\nsrsqlsv.exe  
-s NetWorker_server_hostname -l incr -b pool_name  
MSSQL:db_Name  
PAUSE
```

In this command, replace *installpath* with the pathname where the NetWorker Module software is installed. The **pause** command is optional.

2. In the Microsoft SQL Server program group, select SQL Performance Monitor.
3. From the View menu, select Alert.
4. From the Options menu, select Alert. Change the Periodic Update text box to 10 minutes or more, enough time to run an incremental (transaction log) backup.
5. From the Edit menu, select Add to Alert.
6. In the Add to Alert dialog box:
 - a. Select SQL Server: Databases from the Object menu.
 - b. Select Percent Log Used (%) from the Counter menu.
 - c. In the Alert If group box, click the Over button and enter 80.
 - d. In the Run Program On Alert text box, enter the name and full path of the batch file, for example:
C:\win32app\nsr\bin\sqlalert.bat
 - e. Select the database to monitor in the Instance text box.
 - f. In the Run Program On Alert group box, select Every Time.
 - g. Click Add, then click Done.

To create this alert in Windows 2000 and SQL Server 2000:

1. Create a batch file (for example, *sqlalert.bat*) that contains the backup command and any appropriate command flags, for example:

```
installpath\nsr\bin\nsrsqlsv.exe
-s NetWorker_server_hostname -l incr -b pool_name
MSSQL:db_Name
PAUSE
```

In this command, replace *installpath* with the pathname where the NetWorker Module software is installed. The **pause** command is optional.

2. In the Microsoft SQL Server program group, select Profiler.
3. From the Tools menu, select Performance Monitor.
4. In the left pane, double-click Performance Logs and Alerts.
5. Click Alerts. Any current alerts appear in the right pane.
6. Right-click Alerts and select New Alert Settings from the shortcut menu.
7. In New Alert Settings dialog box, type the name of the new alert in the Name text box, then click OK.
8. Under the General tab, click Add.
9. Set the options in the Select Counters dialog box as follows:
 - a. From the Performance Object list, select SQL Server: Databases.
 - b. Under the Select Counter From List option, select Percent Log Used from the list.
 - c. Under the Select Instances From List option, select the relevant database from the list.
 - d. Click Add, then click Close.
10. From the Alert When The Value Is list, select Over.
11. In the Limit text box, type 80.
12. In the Interval text box, enter 10 (or more)—enough time in minutes to run an incremental (transaction log) backup.
13. From the Units list, select Minutes.
14. Under the Action tab, click Run This Program, then type the full path of the batch file. For example:


```
C:\Program Files\nsr\bin\sqlalert.bat
```
15. Click OK.

For the complete syntax of Legato NetWorker Module backup and restore commands, and an explanation of the command options, see ["Appendix A: Commands and Command Options" on page 143](#).

Backup Groups

2

A *backup group* is a group of one or more NetWorker clients that start their backups at the same specified time. It is possible to assign one or more SQL Server hosts to a NetWorker backup group. NetWorker groups can distribute backups to alleviate network traffic and to schedule backups for a time of day when performance demands on the database and NetWorker servers are lowest.

NetWorker software provides a preconfigured group named Default. The Default group's attributes can be modified, but it cannot be deleted from the list of NetWorker groups. The Default group has the following attributes:

- Autostart = Disabled
After enabling the autostart feature, clients in a backup group start their scheduled backups at a designated time. Be sure to enable the Autostart option for the configured groups; otherwise, the scheduled backups cannot take place.
- Start time = 3:33
- Client retries = 1
- Clones = No
- Clone pool = Default Clone

To assign the SQL Server to another group with different attributes, create that group in the NetWorker Administrator program *before* defining the SQL Server host as a NetWorker client. For instructions on setting up a backup group, refer to the section on groups and backup schedules in the *Legato NetWorker Administrator's Guide*.

Data backed up for a NetWorker group can be written to a specific set of media, or allowed to mix with backup data from other NetWorker groups. Determine which data is stored on each physical volume (for example, a magnetic tape) using the Groups attribute in the Pools resource in the NetWorker Administrator program. For more information, see ["Volume Pools" on page 41](#).

Backup Group Considerations

To back up several large SQL Server databases, consider creating a separate backup group with a different start time for each database. This can help reduce network traffic or load on the NetWorker server. Any number of backup groups can be set up in the NetWorker Administrator program.

When selecting a start time for each group, be sure to schedule the groups far enough apart for one group to complete its backup before the next group starts, and choose times when the network is usually less busy, for example, nights and weekends.

To send a copy of the server's bootstrap report to a specific printer, enter the name of the designated printer in the printer attribute for the group. To locate the printer attribute for the group, use the Expert Mode (for Windows NT and Windows 2000) or View Details (for UNIX) option.

For details on configuring backup groups and using the Expert Mode and View Details options, refer to the *Legato NetWorker Administrator's Guide*.

Volume Pools

A *volume pool* (also known as a *pool*) is a collection of labeled media volumes to which data from a NetWorker backup may be directed. To logically and systematically track, organize, and sort scheduled backups, direct related data to a particular pool according to criteria specified in the NetWorker Administrator program.

Each NetWorker media volume belongs either to the preconfigured Default pool or to a custom created pool. If data is not assigned to a specific pool of volumes, NetWorker software automatically uses the Default pool and searches the backup device for volumes with the pool label Default.

How to Create a Volume Pool

Each volume pool created must be associated with a label template, which provides an automated method of identifying media assigned to the pool. NetWorker software writes backup data to labeled volumes associated with a specified volume pool. Besides providing the ability to segregate different kinds of backup data, pools can direct backup data to specific devices.

For scheduled backups only, the NetWorker Module uses the pool-selection criteria assigned in the NetWorker Administrator program. In release 3.0 and later of the NetWorker Module, pool settings specified in the Backup Options dialog box of the NetWorker User for SQL Server program apply only to

unscheduled backups. This was not the case in NetWorker BusinessSuite® Module 2.0.1 and earlier, in which pool settings in the Backup Options dialog box pertained to both scheduled and unscheduled backups.

Specify the pool destination for each scheduled backup according to the following criteria:

- Object type
- Backup group
- Device
- NetWorker client
- Save set
- Backup level

To configure the pool-selection criteria to specify a pool destination:

1. Start the NetWorker Administrator program.
2. Create a NetWorker pool (named *SQLServerFullBackups*, for example) and select backup level **full** as a pool selection criterion.
3. Create a NetWorker pool (named *SQLServerIncrBackups*, for example) and select backup level **incr** as a pool selection criterion.
4. Create a save group called *SQLGroup*, for example.
5. Associate the pools created in steps 2 and 3 with the save group *SQLGroup*.
6. Create a NetWorker Module client named *mysqlserver*, for example.
7. Add the NetWorker Module client *mysqlserver* to *SQLGroup*.
8. Configure the client *mysqlserver* to use the Default schedule, which performs a full backup on the first Sunday of every month and daily incremental (transaction log) backups in between fulls.
9. Enable the save group *SQLGroup*.

When the full scheduled save of *SQLGroup* occurs, save sets for this group, including the full SQL Server database backup from *mysqlserver*, are written to media volumes belonging to pool *SQLServerFullBackups*. When the daily incremental scheduled saves of *SQLGroup* occur, save sets for this group, including the incremental (transaction log) backup from *mysqlserver*, are written to volumes belonging to pool *SQLServerIncrBackups*.

For more information, refer to the *Legato NetWorker Administrator's Guide*.

How to Customize a Volume Pool

To customize volume pools for the NetWorker Module in the NetWorker Administrator program, the volume pool names must be specified in the NetWorker Administrator's Pool resource. A valid volume pool name is any valid NetWorker name of 64 characters or less. At least one volume must be labeled for each pool.

The following procedure uses an example that demonstrates how to configure volume pools and label templates in the NetWorker Administrator program to separate full, incremental (transaction log), and differential (level 1) backups from one another.

1. Start the NetWorker Administrator program. Click Manage Pools.
2. To create a new pool, right-click the Pools icon and select Create to open the Create Pool dialog box.
3. For the Name attribute, specify *SQLServerFullBackups*, for example.
4. For the Groups attribute, select a backup group that is unique for SQL Server.
5. For the Levels attribute, specify *full*.
6. Under the Preferences tab of the Create Pool dialog box, select Yes for the Store Index Entries attribute.
7. Click OK to save the Pool configuration.

A dialog box opens in which to select a label template or to apply the configuration again to automatically create and select a label template.

8. Click OK and reapply the configuration to automatically create a label template named *SQLServerFullBackups* and assign it to the *SQLServerFullBackups* pool configuration.
9. Repeat steps 2 through 7 to create a pool for the Legato NetWorker Module incremental (transaction log) backups, selecting level *incr* instead of *full*. Name this pool *SQLServerIncrBackups*, for example.
10. Repeat steps 2 through 7 again to create a pool for the Legato NetWorker Module differential backups, selecting level *1* instead of *full*. Name this pool *SQLServerDiffBackups*, for example.
11. Insert a new media volume in the backup device, or select another drive or slot with a volume loaded for labeling.

12. From the Operations menu, select Device Operations>Label and complete the following:
 - a. Label a media volume for full backups by selecting the label template customized for the *SQLServerFullBackups* pool.
 - b. Label a second media volume for the *SQLServerIncrBackups* pool by selecting the label template customized for that pool.
 - c. Label a third media volume for the *SQLServerDiffBackups* pool by selecting the label template customized for that pool.

There are now three different volume pools, one for each of the major backup levels. For more information about using NetWorker volume pools, refer to the *Legato NetWorker Administrator's Guide* and the Online Help.

Browse and Retention Policies

The *browse policy* is the amount of time the NetWorker server maintains entries for a client's backup data in the online index; preset values in the NetWorker software are decade, year, quarter, month, and week (the default is month). The *retention policy* is the amount of time the NetWorker server keeps the media that contains the backup data; preset values in the NetWorker software are decade, year, quarter, month, and week (the default is year). Retention policies are usually set to a longer value than browse policies.

To manage the size of the online indexes, the NetWorker software uses browse policies for the client index entries and retention policies for the media database entries. These policies provide a mechanism for specifying how long save sets should be maintained by the NetWorker server in a format that is easily restorable.

Choose one of the preconfigured policies provided with the NetWorker software or create custom policies. A policy can be used as either a browse or a retention policy. Specifying shorter time periods for the browse and retention policies allows the NetWorker server to recycle expired save sets, thereby increasing the amount of available space within the indexes.

The indexes can also be managed manually through the Indexes and Volumes windows in the NetWorker Administrator program. For more information about manual index management and index policy concepts, refer to the *Legato NetWorker Administrator's Guide*.



Important: The NetWorker software does not allow a browse policy to exceed its retention policy. The entry for a save set must be removed from the file index before the save set can be removed or marked recyclable in the media index.

NetWorker Backup Clients

A NetWorker backup client is a resource configured on the NetWorker server. This resource provides the server with the following information about a particular client, including:

- The client data to back up
- The backup schedule for the client
- The browse policy for the backup data
- The retention policy for the backed-up data

Each SQL Server host to be backed up must be configured as a NetWorker client using the Client resource in the NetWorker Administrator program. In addition, multiple Microsoft SQL Server databases that exist on the same SQL Server host can be configured as separate NetWorker Client resources.

For each Client resource, the NetWorker server does the following:

- Maintains the Client resource information, including entries in the online client file index and media database
- Contacts the clients listed in a backup group configured on the server
- Performs the scheduled backups when a client request is received
- Restores the data upon a request from the client

How to Add a SQL Server Host to the NetWorker Server's Client List

To add a SQL Server host to the NetWorker server's list of backup clients:

1. In the Create Client dialog box in the NetWorker Administrator program, enter the name of the SQL Server host.
2. Choose a schedule.
3. Select a backup group.
4. Select a browse policy.

5. Select a retention policy.
6. For the Save Set attribute, specify any valid save set names. For example:
`MSSQL:`
`or`
`MSSQL:dbname1 [MSSQL:dbname2 MSSQL:dbname3 ...]`
(Entering only `MSSQL;` for example, always yields a backup of all databases on the SQL Server host.)
7. Under the Remote tab, enter **nsrsqlsv** (with any desired command options) in the Backup Command attribute. For information about **nsrsqlsv** options, see ["Appendix A: Commands and Command Options" on page 143](#).
8. Under the Remote tab, add host names to the Remote Access attribute, to grant those hosts directed recover (copy restore type) permission. This allows the named hosts to access the NetWorker server and receive directed recover data, if needed.
9. On the Remote tab of each virtual server Client resource in the cluster, in the Remote User and Password fields, add the user name and password, respectively, for a Windows user account that has SQL Server administrator privileges. This allows the NetWorker Module to back up the SQL Server virtual server.
10. Under the Preferences tab, enter all known aliases for the SQL Server host in the Aliases attribute, as in the following example:
`mars`
`mars.acme.com`
11. Click OK to save this Client resource configuration.
12. To write the results of a scheduled backup to a log file, enter the following for the Notification resource in the NetWorker Administrator program:
`nsrlog -f filename` (Windows NT and Windows 2000)
`or`
`/usr/ucb/logger` (UNIX)

After configuring the SQL Server host as a NetWorker client and enabling the Autostart option, the SQL Server data is backed up automatically according to the schedule set. For more information about configuring and monitoring clients, refer to the *Legato NetWorker Administrator's Guide*.

Backup Operation Status

The NetWorker server and the NetWorker Module provide the following status information regarding backup operations:

- All messages related to the operation are provided in three lists: pending save sets, completed save sets, and failed save sets. These lists are viewable via:
 - A list in the Group Detail window of the NetWorker Administrator program. You have the option of printing a columnar version of the details displayed in the Group Detail window.
 - The NetWorker message log files.
- A savegroup *completion* notice upon completion of a scheduled backup. When using a NetWorker server on a UNIX host, an e-mail notification of the results of a scheduled backup can be sent to the owner of a save set:
 - a. Use the View Details option to edit the NetWorker Client resource for the NetWorker Module host.
 - b. Edit the Owner Notification attribute and enter a notification command directed to the login ID of the owner. Use the Customize feature of the NetWorker Administrator program to edit the notification setup for this report.

For more information about using the NetWorker software to set up event notifications and generate reports, refer to the *Legato NetWorker Administrator's Guide*

Chapter 3: NetWorker Module Restore

The NetWorker User for SQL Server program can be used at any time to restore Microsoft SQL Server data backups that were created with the Legato NetWorker Module. Release 3.1 of the NetWorker Module provides support for Microsoft SQL Server 7.0 and SQL Server 2000.

Before configuring the NetWorker Administrator program to work with the NetWorker Module, understand the different restore types that are available, and how the type of and level of backup created affects the options for restoring data. This chapter includes the following sections:

- ["Microsoft SQL Server 2000 Recovery Models" on page 50](#)
- ["Specifying Database Recovery Models" on page 52](#)
- ["Restore Types" on page 54](#)
- ["Restore Modes" on page 56](#)
- ["Restore Time" on page 57](#)
- ["NetWorker Module Database Restoration" on page 59](#)



Important: It is *highly recommended* that you read this chapter and ["Chapter 2: NetWorker Module Backup" on page 25](#), before performing a backup or restore operation using this Legato NetWorker Module. The type and level of backup created impacts the type of restore that can be performed. For example, regularly scheduled, level full backups are required to enable recovery from a disaster situation.

Microsoft SQL Server 2000 Recovery Models

Microsoft SQL Server 2000 introduces recovery models to control the work-loss exposure when restoring a database. Recovery models represent a trade-off between performance, log-space requirements, and work-loss exposure. SQL Server 2000 supports three recovery models: full, bulk_logged, and simple.

Release 3.1 of the NetWorker Module enforces different backup and restore constraints, depending on whether you are managing a SQL version 7.0 or version 2000 server. [Table 9](#) provides an overview of the available backup types the NetWorker Module program supports for the recovery models available in SQL Server 2000:

Table 9. Supported Backup Types for SQL Server 2000 Recovery Models

Backup Type	Full	Bulk_Logged	Simple
Database	All levels	All levels	Full and Differential
Filegroup	All levels	All levels	None
File	All levels	All levels	None

[Table 10](#) provides an overview of the various backup and restore functionality that the NetWorker Module program supports for the recovery models available in each version of SQL Server:

Table 10. Supported Backup and Restore Functionality for SQL Server 2000 Recovery Models

Functionality	Full	Bulk_Logged	Simple
Transaction Log Backup	Yes	Yes	No
Point-in-time Restore	Yes	Maybe ^a	No
Named Log Mark Restore	Yes	Maybe ^a	No

a. Point-in-time restore is not supported if bulk changes are present in the transaction log that contains the time or day marker.

For more information about setting recovery models for SQL Server data, refer to the Microsoft SQL Server 2000 documentation.

Full Recovery Model

The full recovery model imposes the fewest constraints on the backup and restore process, but requires the most log space of all recovery models. Release 3.1 of the NetWorker Module enforces the following constraints on SQL Server 7.0 and SQL Server 2000 data items:

- All levels of database, filegroup, and file backup are supported.
- All backup data is restorable up to the most recent transaction log.
- Point-in-time and named log mark restores (2000 only) are supported.

Bulk_Logged Recovery Model

The bulk_logged recovery model also imposes a few backup constraints and supports reduced log space requirements for bulk operations. NetWorker Module release 3.1 enforces the following constraints differently, depending on whether you are managing a SQL Server 7.0 or 2000 storage hierarchy:

- All levels of database, filegroup, and file backup are supported for SQL Server 2000, while only level full and differential database backup types are supported for SQL Server 7.0.
- All SQL Server 2000 backup data is restorable up to the most recent transaction log, while SQL Server 7.0 backup data is only restorable up to the last level full or differential database backup.
- Point-in-time and named log mark restores are *not* supported for SQL Server 7.0.
- A point-in-time restore is not supported for SQL Server 2000 if a bulk log change has occurred for the transaction log backup that corresponds to the current browse time. In addition, a point-in-time restore is not supported if bulk changes are present in the transaction log that contains the time or day marker. The NetWorker Module rolls forward all transactions to the end of the transaction log and leaves the database in an operational state.

Simple Recovery Model

The simple recovery model provides the fastest performance by minimally logging operations, thereby reducing log space, but not supporting transaction log backups. The simple recovery model implements the most constraints on the backup and restore process. This recovery model provides significant control over how storage management operations impact system resources. NetWorker Module release 3.1 enforces the following constraints on both SQL Server 7.0 and 2000 data items:

- Only level full and differential database backup types are supported.
- Backup data is restorable up to the last level full or differential database backup.
- Point-in-time and named log mark restores are *not* supported.

Specifying Database Recovery Models

SQL Server 2000 allows a different recovery model to be specified for each database and provides support for the legacy database options **trunc.log.on.chkpt** and **select into/bulk copy**. Depending on how the legacy options are set in SQL Server, they map to the new recovery models as follows:

		select into/bulk copy	
		True	False
trunc.log.on.chkpt	True	Simple	Simple
	False	Bulk_Logged	Full

Whether or not a transaction log, and the descendent filegroups and files of the database, are available for backup or restore, depends on the recovery model specified in SQL Server. NetWorker Module release 3.1 modifies how items in the SQL Server storage hierarchy are displayed in the Backup and Restore windows, based on the constraints enforced by the recovery model. For more information about how these constraints are visually interpreted in the NetWorker User for SQL Server program, see ["Fake Objects" on page 64](#).

Changing the Recovery Model for a Database

SQL Server allows the recovery model for a database to be changed. However, the NetWorker Module does not enforce the constraints of the new recovery model after a change. Therefore, changing the recovery model may impact the current backup and restore strategies for a database. The following sections provide instructions for preparing for the recovery model change, as well as backup strategies for maintaining the consistency of the SQL Server storage hierarchy after the change.

How to Change from Full or Bulk_Logged to Simple

To change to the simple recovery model:

1. Before changing the recovery model, it is *highly recommended* that you perform a level incremental database backup to maintain the transaction log files.
2. In SQL Server 2000, change the recovery model to simple.

After changing the recovery model, modify backup strategy to halt execution of level incremental database backups. If the change to simple is only temporary, it is not necessary to modify the backup schedule, because, for NetWorker Module release 3.1, the backup level is automatically promoted to full if the recovery model (or legacy database status bit) does not allow the specified level.

How to Change from Bulk_Logged to Full

Follow these steps when changing from bulk_logged to the simple recovery model:

1. In SQL Server 2000, change the recovery model to full.
2. If performing a point-in-time restore is anticipated, performing a level incremental database backup is *highly recommended* to maintain the transaction log files.
3. It is not necessary to modify the backup strategy.

How to Change from Simple to Full or Bulk_Logged

Because the simple recovery model has many more constraints than the full and bulk_logged models, the NetWorker Module may allow certain scenarios to occur, although under normal conditions these situations would be flagged as warnings or failures by SQL Server. The following scenarios highlight the importance of modifying the backup strategy after changing from simple to a full or bulk_logged recovery model:

- Level incremental backups can be created if the most recent level full or differential database backup was created when the database recovery model was simple. SQL Server allows the transaction log backup to occur, but displays a warning.
- You cannot restore a level incremental database backup if the level full or differential database backup being restored was created when the database recovery model was simple. This is consistent with the SQL Server **norecover** behavior.



Important: In this scenario, the restore will have to be reinitiated using the most recent level full database backup, and all database transactions performed after the backup is created will be lost.

After changing from simple to a full or bulk_logged recovery model, modifying the backup strategy is *highly recommended*.

1. In SQL Server 2000, change the recovery model to full or bulk_logged.
2. After changing the recovery model, perform a level full or differential database backup.
3. Modify the backup strategy to include level incremental database backups (*highly recommended*).

Restore Types

In order to restore data, the NetWorker Module requires that a **restore type** be specified. The restore type is based on the level of backup created, as well as the set of data needed to restore from a backup. The restore type must be specified prior to browsing and selecting objects for restoration. The NetWorker Module supports four restore types: normal, partial, verify-only, and copy (which replaces the directed recover operation).

Normal Restore Type

The normal restore type restores the entire set of data associated with one or more SQL Server 7.0 or later backups, including full, incremental, and differential backups. The normal restore type supports recovery of a database, a filegroup, or a file to the database originally backed up. The normal restore type can restore level full, level 1 (differential), and level incremental backups in the order required by SQL Server. When performing a restore, the NetWorker Module uses the normal restore type as the default. For more information about creating backups, see ["Chapter 2: NetWorker Module Backup" on page 25](#).

Because the NetWorker Module can back up *only* specified files and filegroups of a SQL Server 7.0 or later database, restoring from a media failure that affects only a subset of the data in the databases is possible. In addition, a single filegroup, or multiple filegroups and/or files, can be restored from a full database backup.

Partial Restore Type

The partial restore type restores a *portion* of the filegroups associated with a single SQL Server 2000 database backup (Microsoft SQL Server 2000 only). When a partial database restore is performed, the Primary filegroup and associated files are always restored, in addition to the files specified for restore. The Primary filegroup contains information necessary for restoring the database to the proper structure. Files or filegroups not selected are created, but are empty. Note that only a single item can be marked for this operation. In addition, a copy of a system database can be marked, but it cannot be overwritten.

Verify-Only Restore Type

The verify-only restore type only verifies the backup media for the selected SQL Server 7.0 or later backups.

Note: Selecting the verify-only restore type does *not* restore the SQL Server data. In addition, when verify-only is specified, item-level properties for database, filegroup, and file objects are not available.

Copy Restore Type

The copy restore type creates a copy of a database by restoring a SQL Server 7.0 or later database backup to a new location or to a new database name. The copy restore type makes it easy to replicate a database that was previously backed up. The database backup to restore may be obtained from the local host or imported from a different NetWorker SQL Server client. Note that you can only mark a single item for this operation. In addition, you can copy a system database, but you cannot overwrite it.

Restore Modes

3

To restore a database, NetWorker Module requires that a *restore mode* be specified. A restore mode instructs the SQL Server how to interact with the database after the restore operation completes. For instance, restore modes can leave the database in an intermediate state, so that additional transaction logs can be applied. Restore modes correspond to SQL Server database restore options and include: normal, no recovery, and standby.

Normal Restore Mode

The normal restore mode instructs the SQL Server to leave the database in an operational state after the restore completes, which allows database reads and writes. The normal restore mode is the default mode the NetWorker Module uses when restoring a database.

No Recovery Restore Mode

The no recovery restore mode activates the SQL Server NORECOVERY database restore option for the last stage restored. This places the database in an unloadable state after the restore, but still able to process additional transaction log restore operations.

Standby Restore Mode

The standby restore mode activates the SQL Server STANDBY database restore option for the last stage restored, which forces the database to be in a read-only state between transaction log restore operations. The standby restore mode provides an *undo* file for SQL Server to use when rolling back the transactions.

Restore Time

Microsoft SQL Server and the NetWorker Module allow backups to be restored to a specific *restore time*. The restore time controls which backup data should be reinstated when a database is restored. The restore time may also control which portions of a level incremental backup are to be restored, when the NetWorker Module is instructed to discard transactions performed after a given time.

The default, or *current*, restore time for each database comes from the create time of the marked item. By default, the most recent backup is restored. If the most recent backup is level incremental or 1, dependent backups are restored first. User-specified restore times can restore older backup versions or perform point-in-time restore operations. For instance, a point-in-time restore may be specified by using a restore time that is earlier than the create time of the transaction log backup, but later than the create time of the previous backup.

In previous versions of BusinessSuite Module for Microsoft SQL Server, the global browse time controlled which backup versions were available for restoration. In NetWorker Module 3.1, the global browse time *only* controls which backup data is viewable in the Restore window.

The NetWorker Module provides three methods for restoring to a specific time: database backup versions, point-in-time restore of a transaction log (level incremental) backup, and restoring to a named log mark.

Database Backup Versions

In the NetWorker User for SQL Server Restore window, you can browse the online index and mark a database backup version to restore. Use the Restore Time tab from the Database Restore Properties dialog box to select any backup version for restore.

The global browse time can be modified from the View>Change Browse Time menu to display backup versions for a different restore time. However, the browse time cannot change to a point earlier than the creation time of the first SQL Server backup because the NetWorker file index would be empty. If the browse time changes to some point in the future, the databases appear as they currently exist in the SQL Server storage hierarchy.

Point-in-Time Restore of a Transaction Log Backup

When a transaction log backup is marked under the Restore Time tab in the Database Restore Properties dialog box, you can select a point-in-time restore. The Restore Time properties support selection of a backup version and modification of the restore date and time. The backup level marked in the Restore window is the default row marked in the Backup Version table under the Restore Time tab. You can mark a different backup version from the Restore Time version table. When a point-in-time restore is performed, the restore procedure only reinstates transactions from the backup version that occurred before the specified restore date and time.

Note: Point-in-time restore only applies when a transaction log backup is marked for restore.

3

Microsoft SQL Server 2000 Named Log Marks

Microsoft SQL Server 2000 introduces enhanced point-in-time restoration by allowing *named log marks* to be specified during transaction creation. Database applications create named log marks when transactions are performed. The marks enable access to specific transaction points in a database transaction log backup. NetWorker Module 3.1 provides support for restoring to the beginning or end of a named log mark during database restoration. Restoring data using named log marks is an improvement over point-in-time restore, because the time associated with restoring to a specific transaction can be more accurately determined.



Important: When a named log mark is created in the SQL Server *msdb* database, the log mark time is saved to the millisecond. However, the NetWorker software's time format, which is used to specify point-in-time restore, only supports granularity to the second. If named log marks with duplicate names are created within a millisecond of each other, the NetWorker Module restores to the most recently named log mark.

For more information about transaction log marks, refer to the Microsoft SQL Server documentation. For more information about performing a restore using name log marks, see ["How to Specify Database Item Restore General Properties" on page 85](#).

Specifying the Restore Time

Unique restore times can be set on a per-item basis. For example, if two databases, *earth* and *moon*, are marked for restore, *earth* can have a different restore time than *moon*.

Note: Before NetWorker Module release 3.0, the global browse time was used as the restore time.

To modify the restore time for each marked database:

1. Start the NetWorker User for SQL Server program.
2. In the Restore browse window, right-click a database item and select Properties from the shortcut menu.
3. Under the Restore Time tab, specify a restore time and click OK.

For more information about modifying the restore time values, see ["How to Specify Database Item Restore Time Properties" on page 95](#).

NetWorker Module Database Restoration

When you start a restore, the NetWorker Module automatically implements the following process:

1. The NetWorker Module restores the most recent full backup, then restores the most recent differential (level 1) backup (if any).
2. The NetWorker Module restores all transaction log backups that ran after the most recent differential backup (or that ran after the last full backup, if there was no differential backup). To correctly handle restoration of uncommitted transactions, the NetWorker Module specifies the SQL Server NORECOVERY option for all intermediate transaction logs.
3. Restore of the final transaction log specifies the restore mode if a restore mode of STANDBY or NORECOVERY was selected.

For example, if you selected a restore mode of NORECOVERY, that specification appears in the output for a database restore as follows:

```
C:> nsrsqlrc -s NetWorker_server_host my_database
nsrsqlrc: Restoring database my_database...
nsrsqlrc: RESTORE database my_database FROM
virtual_device='BSMSQL' WITH norecovery, stats
nsrsqlrc: RESTORE database my_database from
virtual_device='BSMSQL' WITH norecovery (differential)
nsrsqlrc: RESTORE transaction my_database FROM
```

```
virtual_device='BSMSQL' WITH norecovery
nsrsqlrc: RESTORE transaction my_database FROM
virtual_device='BSMSQL' WITH norecovery
Received 1.0 MB 4 files from NSR server.
```

SQL Server Master Database Maintenance

The *master* database contains information about all SQL Server databases on the SQL Server host. The *master* database can be restored in one of the following ways:

- Using the SQL Server Enterprise Manager to rebuild the *master* database. For detailed instructions on rebuilding the *master* database, refer to the Microsoft SQL Server documentation.
- Restoring the data for the *master* database and other SQL Server databases from the NetWorker server. For more information on restoring the *master* database, see [“Disaster Recovery of SQL Server” on page 142](#).

3

NetWorker Module Interaction with SQL Server Services

When restoring the Microsoft SQL Server *master* database, the NetWorker Module **restore** program automatically stops and restarts the SQL Server services appropriately, as follows:

1. Before the restore begins, the NetWorker Module stops the SQL Server service and other dependent services.
2. The NetWorker Module starts the SQL Server service in single-user mode.
3. The NetWorker Module performs the restore.
4. After the restore finishes, the NetWorker Module waits for the SQL Server service to shut down.
5. The NetWorker Module restarts the SQL Server service in production mode.

When restoring the Microsoft SQL Server *master* database running in a cluster configuration, the NetWorker Module **restore** program controls the SQL cluster resources, as follows:

1. Before the restore begins, the NetWorker Module stops the SQL Server and all dependent cluster resources, including the SQL Server Agent.
2. The NetWorker Module starts the SQL Server service in single-user mode.
3. The NetWorker Module performs the restore.

4. After the restore finishes, the NetWorker Module waits for the SQL Server service to shut down.
5. The NetWorker Module restarts the SQL Server and dependent cluster resources.

When restoring the Microsoft SQL Server *msdb* database, the NetWorker Module **restore** program automatically stops and restarts the SQL Server services appropriately, as follows:

1. Before the restore begins, the NetWorker Module stops the SQL Server Agent.
2. The NetWorker Module performs the restore.
3. After the restore finishes, the NetWorker Module restarts the SQL Server Agent.

When restoring the Microsoft SQL Server *master* database *along with the msdb* database, the NetWorker Module **restore** program automatically stops and restarts the SQL Server services appropriately, as follows:

1. Before the restore begins, the NetWorker Module stops the SQL Server service and other dependent services, including the SQL Server Agent.
2. The NetWorker Module starts the SQL Server service in single-user mode.
3. The NetWorker Module restores the *master* database.
4. After the restore finishes, the NetWorker Module waits for the SQL Server service to shut down.
5. The NetWorker Module restarts the SQL Server service in production mode.
6. The NetWorker Module restores the *msdb* database.
7. The NetWorker Module restarts the dependent services, including the SQL Server Agent.

Chapter 4: Using the NetWorker User for SQL Server Program

This chapter describes the visual appearance, interaction mode, and basic data requirements of NetWorker User for SQL Server release 3.1. Before attempting to back up and restore Microsoft SQL Server 7.0 or 2000 data, reviewing the information in this chapter is highly recommended. This chapter includes the following sections:

- ["Browsing Data Items" on page 63](#)
- ["Marking Items" on page 66](#)

Browsing Data Items

The NetWorker User for SQL Server program displays the data items in the SQL Server storage hierarchy in either the Backup or the Restore window, depending on which operation is performed.

The Backup and Restore browse windows are split into two panes. The left pane consists of expandable trees that display the SQL Server storage hierarchy contained on the current NetWorker Module host. To view the levels of data, double-click the data items to expand the tree.

NetWorker Module 3.1 allows you to browse the filegroups and files contained in a database; however, the program may display certain data items that are contained in the storage hierarchy, but which are not available for backup or restore operations. These are referred to as *fake objects*.

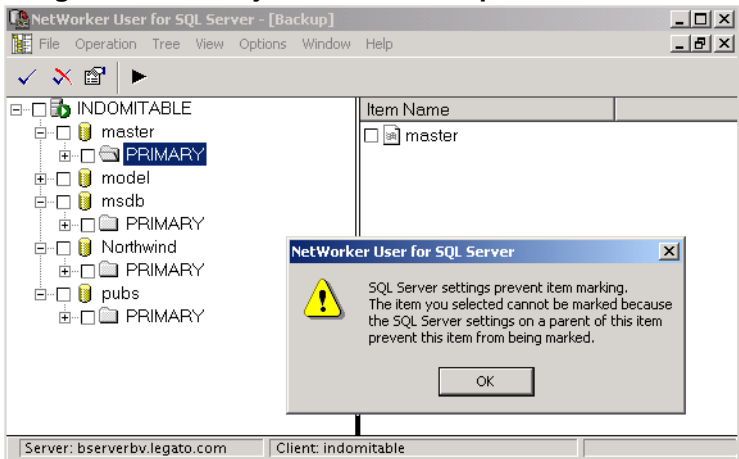
Restore Browse Window Restrictions

The Restore browse window displays backups available for restoring, based on the selected restore type. The normal and verify restore types do not restrict the contents of the browse window, but the partial and copy restore types only display restorable *database* objects. When the partial or copy restore type is chosen, filegroups and files of the selected database are not visible in the browse window, but are available in the Database Item Restore Properties dialog box. For more information about this dialog box, see ["Chapter 3: NetWorker Module Restore" on page 49.](#)

Fake Objects

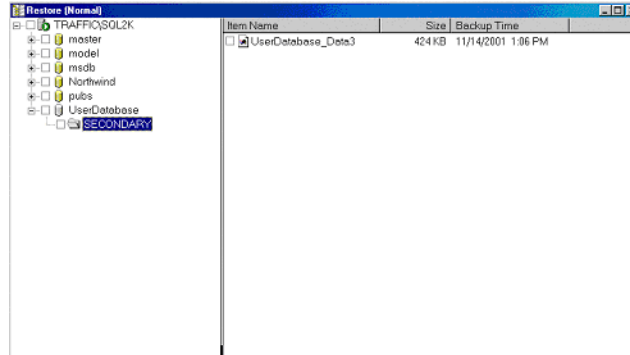
When a filegroup or file exists in the SQL Server storage hierarchy, but cannot be backed up due to SQL Server settings on the database, the item is displayed in the Backup browse window, shown in [Figure 5](#), using the Fake Filegroup or Fake File conventions outlined in [Table 11 on page 65](#). In addition, if you try to mark a fake object, an error dialog box appears, as shown in [Figure 5](#). For more information about the constraints Microsoft SQL Server 2000 enforces that determine whether an item is available for back up or restore, see ["Microsoft SQL Server 2000 Recovery Models" on page 50.](#)

Figure 5. Fake Objects in the Backup Browse Window



When one or more of the parent items do not have backup items available, then the item is displayed in the Restore window, shown in [Figure 6](#), using the fake object conventions outlined in [Table 11](#).

Figure 6. Fake Objects in the Restore Browse Window






Display Conventions

The NetWorker Module uses specific data item names, text characteristics, and data type icons to distinguish the variable qualities of SQL Server data. [Table 11](#) outlines these conventions:

Table 11. SQL Server Storage Hierarchy Display Conventions (Part 1 of 2)

Data Item	Description	Pane	Icon
SQL Server	<ul style="list-style-type: none"> • Root of the storage hierarchy • Signifies all SQL Server databases on the host 	Left only	
Database	<ul style="list-style-type: none"> • Descendent of root • Signifies a database • May contain filegroups 	Left and right	
Fake Database	<ul style="list-style-type: none"> • Only visible in the Restore window • Signifies that the database cannot be selected for backup 	Left and right	
Filegroup	<ul style="list-style-type: none"> • Descendent of a database • Signifies a database filegroup • May contain files 	Left or right	

Table 11. SQL Server Storage Hierarchy Display Conventions (Part 2 of 2)

Data Item	Description	Pane	Icon
Fake Filegroup	Signifies that the filegroup cannot be selected for backup or restore	Left or right	
File	<ul style="list-style-type: none"> • Descendent of a filegroup • Signifies a database file 	Right only	
Fake File	Signifies that the file cannot be selected for backup or restore	Right only	

Marking Items

Mark items for backup or restore in the browse window. The NetWorker Module provides several methods for marking and unmarking items:

- Click the check box to the left of the item you want to mark. To unmark it, click the check box again.
- Click the item you want to mark to select it, then click the Mark button in the task bar. To unmark it, click the Unmark button in the task bar.
- Right-click the item you want to mark and select Mark from the shortcut menu. To unmark it, select Unmark from the shortcut menu.
- Click the item you want to mark to select it, then choose Mark from the File menu. To unmark it, choose Unmark from the File menu.

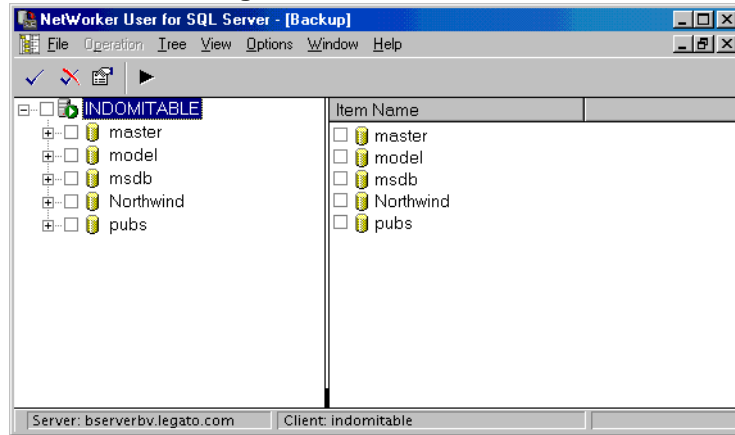


Important: Some SQL Server data items are visible in the browse window, but cannot be marked. For example, when certain database options are set, SQL Server does not allow a filegroup or file backup to be performed. These items are displayed using visual and behavioral cues to distinguish them from normal data items. See [Table 11 on page 65](#) for more information.

NetWorker Module release 3.1 provides enhanced marking indicators that help determine the state of each item in the browse tree:

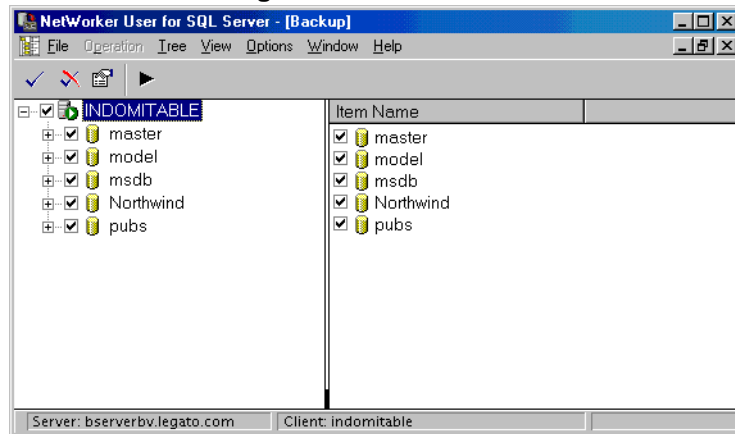
- **Unmarked** — An unmarked item is one that is not selected for backup or restore. An empty check box appears to the left of an item to indicate it is unmarked, as shown in [Figure 7](#).

Figure 7. Unmarked Items



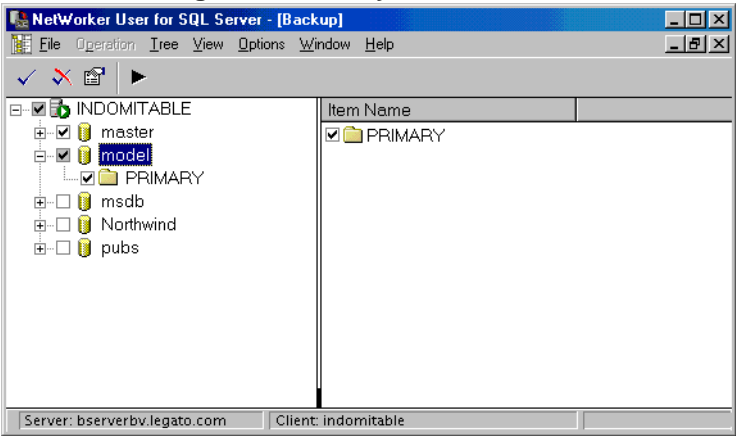
- **Marked** — A marked item is one that is selected for backup or restore. A check mark appears in the check box to the left of each marked item, as shown in [Figure 8](#).

Figure 8. Marked Items



- Partially marked — A partially marked item is one that has marked descendents, but the item itself is not explicitly marked. A partially marked item is not backed up or restored. A check mark appears in a gray check box to the left of each partially marked item, as shown in [Figure 9](#).

Figure 9. Partially Marked Items



4

Restore Browse Window Restrictions

In the Restore window, the rules for marking an item are based on the selected restore type. The normal and verify restore types do not restrict marking in any way; all restorable objects (database, filegroup, file) are markable. When the partial or copy restore type is chosen, only one *database* object may be marked. Marking the root SQL Server item is not allowed.

When the *partial* restore type is chosen, the subset of filegroups and files of the selected database must be marked using the Database Item Restore Properties dialog box. (For more information about this dialog box, see ["Specifying Database Restore Properties" on page 84](#).) When the *copy* restore type is chosen, filegroups and files of the selected database are automatically marked and restored as part of the full database restore.

Marking Semantics and Restrictions

To support the new browsing capabilities, NetWorker Module 3.1 imposes certain semantics and restrictions regarding how items may be marked. Whether an item can be marked is based on the mark status of that item's ascendants and descendents.

While marking items, message dialog boxes may appear that provide guidance on the current marks and the type of operation that will be performed. These dialog boxes are helpful in creating the different backup and restore types. For more information about backup types, see ["Chapter 2: NetWorker Module Backup" on page 25](#). For more information about restore types, see ["Chapter 3: NetWorker Module Restore" on page 49](#).

The NetWorker Module marking capabilities are as follows:

- You can mark a single database, filegroup, or file.
- You can mark multiple, heterogeneous items.
- You can mark an item when any of that item's siblings are already marked.
- You can mark all SQL Server data by right-clicking the root item and selecting Mark All Children from the shortcut menu. This allows you to easily create a *full* backup of the entire SQL Server storage hierarchy.
- You can unmark all databases from the SQL Server root.
- You cannot mark an item when any of that item's ascendants or descendents are already marked, except in the SQL Server root. To mark an item, the marked ascendants and descendents must first be unmarked.
- When a database is marked, all of the item's descendents are *not* automatically marked.
- When a filegroup is marked, all of the files are *not* automatically marked.
- Rather than mark each descendent of a database or each filegroup individually to select them all, you can right-click the root item and select "Mark All..." from the shortcut menu.

The SQL Server root has special marking properties. This includes the ability to mark and unmark all databases. When you right-click the SQL Server root, the shortcut menu provides the following options:

- Mark All Databases — Use this option to mark all databases. This can be done at any time, if the following is true:
 - All databases are not already marked. The menu item is dimmed this case.
 - A filegroup or file is not already marked. The root mark is not allowed in this case.
- Unmark All Databases — Use this option to unmark all marked databases.

Chapter 5: Backing Up SQL Server Data

This chapter describes the steps required to back up Microsoft SQL Server data using the NetWorker User for SQL Server program. For guidelines on using the NetWorker Module commands and command options to perform the backup from the command prompt, see ["Appendix A: Commands and Command Options" on page 143](#). SQL Server data cannot be *manually* backed up using the NetWorker Administrator program. This chapter includes the following sections:

- ["Preparing for the Backup" on page 71](#)
- ["Creating the Backup" on page 72](#)
- ["Target Volume Pools for Unscheduled Backups" on page 80](#)

Preparing for the Backup

The Backup browse window displays data available for backup based on the SQL Server database settings. SQL database items that cannot be backed up are not displayed in the Backup browse window. This includes, but is not limited to, databases in the following state:

- Standby
- Offline
- Not recovered
- Loading
- Prerecovery
- Single user with active user connections


For more information about browsing items, see ["Browsing Data Items" on page 63](#).

Creating the Backup

1. Start the NetWorker User for SQL Server program.

If multiple instances of SQL Server are active on the computer, the Select SQL Server Instance dialog box, shown in [Figure 10](#), opens before the main window opens; otherwise, the main window, shown in [Figure 11](#), opens.


Figure 10. Select SQL Server Instance




2. Select the SQL Server instance that the NetWorker Module should use to perform the backup operation, and click OK.

The main window opens, as shown in [Figure 11](#).

Figure 11. Main Window in NetWorker User for SQL Server



- (Optional) Click the Select NetWorker Server button on the task bar.
The Change Server dialog box opens, as shown in [Figure 12](#).

Figure 12. Change Server Dialog Box

- (Optional) Select a NetWorker server from the list, then click OK.

How to Back Up Using NetWorker User for SQL Server

When performing a backup using the NetWorker User for SQL Server program, the NetWorker server always performs a *full* backup of the SQL Server data.

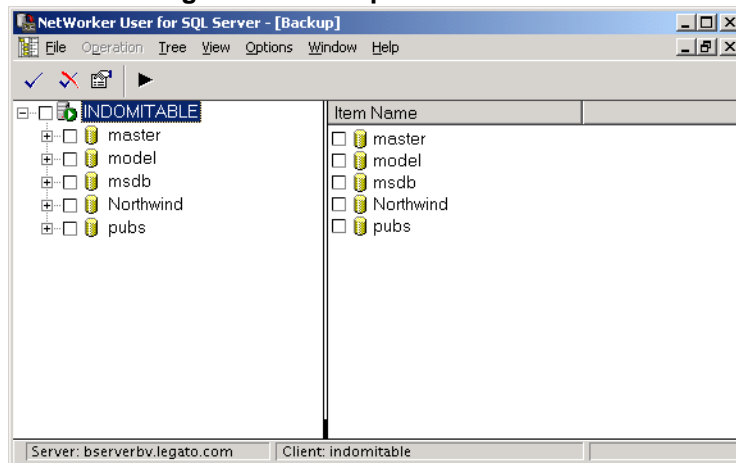
To run a backup from the NetWorker User for SQL Server program:

- In the NetWorker User for SQL Server window, click the Backup button (or select Backup from the Operation menu).

The Backup browse window opens, as shown in [Figure 13 on page 74](#).



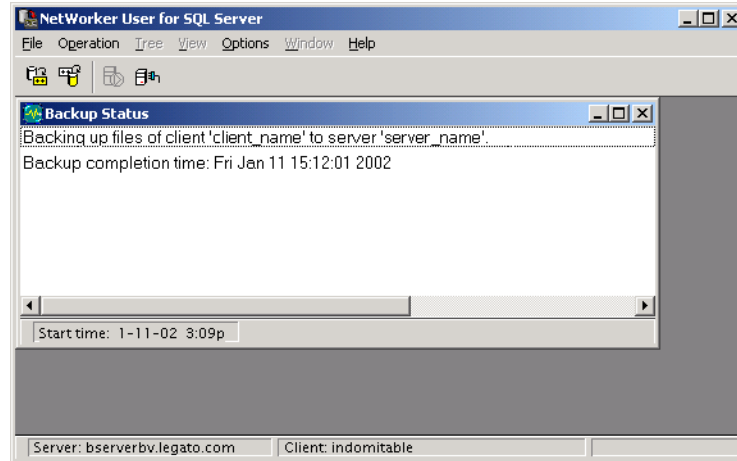
Important: SQL database items that cannot be backed up are not displayed in the Backup browse window. For more information, see ["Preparing for the Backup" on page 71](#).

Figure 13. Backup Browse Window

2. To view a list of data items available for backup, click the expand/collapse icon next to an item in the left pane of the Backup browse window. Alternatively, double-click an item name in the left pane to view the list. The descendants of the item appear in the right pane of the window.
3. Mark each item to back up. For information about marking items for backup, see ["Marking Items" on page 66](#).
4. (Optional) Set the backup properties for each marked database, if any, then click OK. For information about setting backup properties, see ["Database Item Backup Properties" on page 79](#).
5. (Optional) Select Options>Backup Options and specify preferences for backup operation settings. Then click OK. For information about setting options, see ["Database Item Backup Properties" on page 79](#).
6. Click the Start button to begin the backup.

The amount of time to back up a database depends on database size, network traffic, server load, and tape positioning. To monitor the progress of each database as it is backed up, the NetWorker Module displays messages in the Backup Status window, as shown in [Figure 14 on page 75](#).

Figure 14. Backup Status Window



After the backup is finished, a Backup Completed message appears. Close the Backup Status window.



Important: A backup cannot run if there is no media volume mounted at the storage node (backup device). Therefore, it is best to have a labeled media volume mounted in the storage node at all times. The volume should be labeled for the volume pool where the backup is to be directed. If there is no volume in the backup device when a backup is triggered, no messages appear in the status window and the backup waits for operator intervention.

Note: To cancel a backup, select File>End Backup or close the Backup Status window.

Backup Options

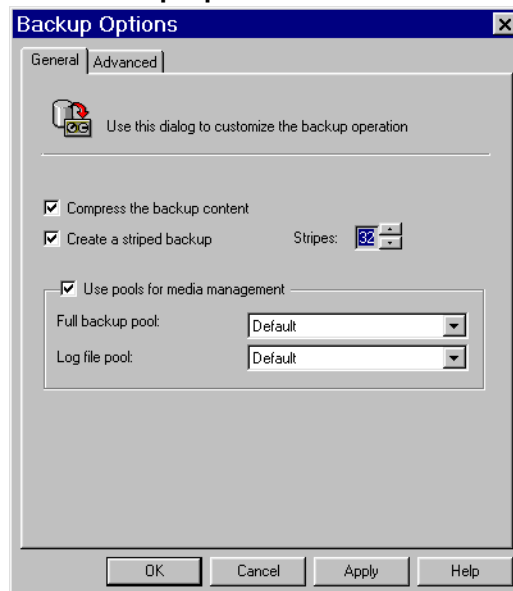
To customize the backup operation, specify preferences in the Backup Options dialog box before starting the backup. These properties are supported for SQL Server 7.0 and 2000.

Preferences specified for these properties are applied to each marked item and are used for the current and subsequent backup operations until the NetWorker Module program is closed. Pool settings chosen are stored in the registry and persist from one session to another. When a preference is not specified, the NetWorker Module defaults are used.

How to Specify General Backup Options

1. With the Backup window open, invoke the Backup Options dialog box, shown in [Figure 15](#), in one the following ways:
 - Select Options>Backup Options.
 - Select Backup Options from the shortcut menu. The root item does *not need to be marked* in order to access these properties.
 - Select the root item and then click the properties button on the task bar.

Figure 15. Backup Options Window — General Tab



2. Select the following options as desired:

- **Compress The Backup Content** (default = unchecked) — Applies XBSA compression to all marked databases before writing the backup data to the storage device. In the same manual backup, certain databases cannot be backed up with compression and others without. Compressing data for a backup generates less network traffic and uses less backup media space, but consumes additional CPU resources. Most modern tape devices perform compression, which makes software compression unnecessary.
- **Create A Striped Backup** (default = unchecked) — Creates a striped backup using the SQL Striped feature. When this option is checked, the Stripes option is enabled. Use the arrows to specify the number of stripes you want. The maximum number of stripes the NetWorker Module software supports (32) is the same for both SQL Server versions.

For more information about using this option, see ["Striping" on page 111](#).

- **Use Pools For Media Management** (default = unchecked) — Specifies media volume pools for storing the backup contents. When this option is checked, the Full Backup Pool and the Log File Pool filename options are enabled.

Volume pools allow you to segregate different kinds of backup data and direct backup data to specific devices. For more information about using this option, see ["Volume Pools" on page 41](#).

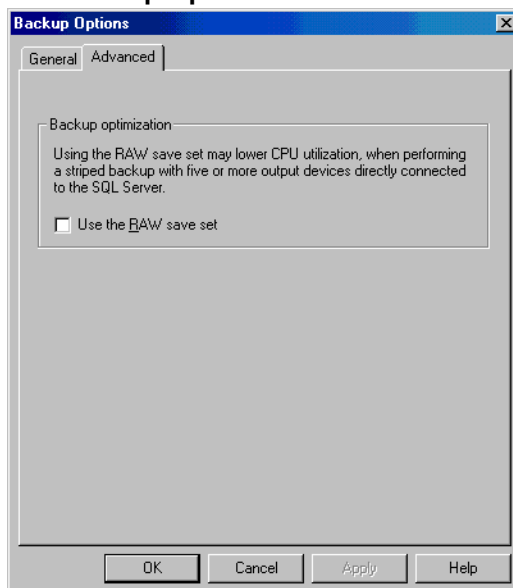
From the Full Backup Pool list, select the pool to store full SQL Server backups, including databases and filegroups. From the Log File Pool list, select the pool to store transaction log or differential SQL Server backups.

Note: This option only applies to manual backup operations performed from the NetWorker User for SQL Server program or from a command prompt.

How to Specify Advanced Backup Options

The Advanced tab of the Backup Options dialog box, shown in [Figure 16](#), provides access to complex, infrequently used configuration settings for the backup.

Figure 16. Backup Options Window — Advanced Tab



5



Important: Backup optimization using the RAW save set is a NetWorker Module release 3.1 feature and is *only* available for configurations running NetWorker Module release 3.1 or higher and NetWorker Server software release 6.0 or higher with a Power Edition license. In addition, the NetWorker Server software and NetWorker client, or the NetWorker client and the NetWorker storage node, must be installed on the same computer as the SQL Server being backed up.

1. Select the following option as desired:

Use The RAW Save Set (default = unchecked) — Specifies to use the RAW save set when creating a striped backup. The RAW save set can be used to lower CPU utilization when performing a striped backup.

It is highly recommended using the RAW save set only in high-performance device configurations. Performance enhancements can be achieved when a large number of output devices are directly connected to the SQL Server host.

2. Close the Backup window when finished.

For more information about using the RAW save set, see ["RAW Save Set" on page 118](#).

Database Item Backup Properties

Use these properties to specify preferences for a single, *marked* database before starting the backup operation. When a preference is not specified, the NetWorker Module default is used. After the backup operation is complete, and the Backup browse window is closed, the property values revert back to the NetWorker Module defaults.

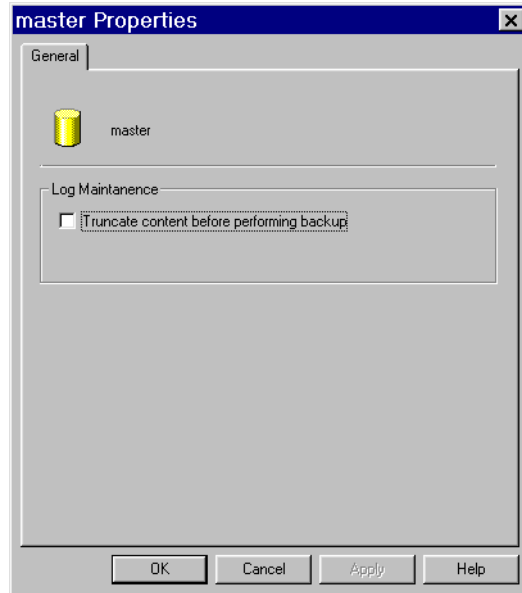
To specify these properties:

1. Right-click a *marked* database in the Backup browse window and select Properties from the shortcut menu, or click the Properties task bar button.

The database Properties dialog box opens, as shown in [Figure 17 on page 80](#).

In NetWorker Module release 3.1, only one database property exists:

- Truncate Content Before Performing Backup (default = unchecked) — Truncates the transaction logs before backup, then backs up the databases and the transaction logs.

Figure 17. Database Backup Properties

2. Select the Truncate Content Before Performing Backup option if desired, then click OK.

5

Target Volume Pools for Unscheduled Backups

For manual backups, specify target data and log pool names in the Backup Options dialog box in the NetWorker User for SQL Server program (see [Figure 15 on page 76](#)). The pool names specified in this dialog box are stored in the operating system registry and are persistent.

If a data pool is specified:

- All backups performed from the NetWorker User for SQL Server program are directed to this pool.
- All full (and only full) backups initiated from a Windows command prompt are also directed to this pool. (However, when running a backup from a command prompt, you can override the data pool setting by specifying a different pool using the **-b** option. For more information about command line options, see "[Appendix A: Commands and Command Options](#)" on page 143).

Settings in the Backup Options>Use Pools For Media Management field pertain only to backups initiated from a command line. If a log pool is specified, all incremental *and* level 1 (differential) backups initiated from the command line are directed to that pool, unless specified otherwise using the **-b** option.

For example, the following command overrides any log pool setting specified in the Backup Options dialog box:

```
nsrsqlsv -l incr -b SQLServerLogs my_dbase
```

Note: When performing an unscheduled backup of SQL Server, the client indexes and bootstrap file are not automatically backed up. To create client indexes and a bootstrap file following a manual backup, run the following command from the command line on the NetWorker server:

```
savegrp -G savegroup_name
```

Assigning Pools in NetWorker User for SQL Server

To specify names for volume pools in the Backup Options dialog box:

1. Check Use Pools For Media Management (see [Figure 15 on page 76](#)).
2. In the Full Backup Pool and Log File Pool fields, specify volume pool names so that the full backups are directed to one pool and the incremental and level 1 (differential) backups are directed to a different pool. Legato NetWorker Module stores these values in the Windows registry.



Example:

You might specify the following data pools:

SQLServerFulls — This name might signify the pool for SQL Server *scheduled* full backups.

SQLServerFullManuals — This name might signify the pool for all full, unscheduled backups.

SQLServerGUIFulls — This name might signify the pool for backups initiated from the NetWorker Module.

In addition, you might specify the following log pool:

SQLServerLogs — This name might signify the pool for incremental and level 1 (differential) scheduled backups.

Chapter 6: Restoring SQL Server Data

This chapter describes the steps required to restore Microsoft SQL Server data from the NetWorker User for SQL Server program. For guidelines on using the NetWorker Module commands and command options to perform the restore from a command prompt, see ["Using the nsrsqlrc Command" on page 144](#). This chapter includes the following sections:

- ["Before the Restore" on page 83](#)
- ["Specifying Database Restore Properties" on page 84](#)
- ["Backing Up the Active Portion of the Transaction Log" on page 99](#)
- ["Filegroup and File Restore Properties" on page 100](#)
- ["Performing the Restore" on page 102](#)

Note: In order to restore SQL Server data, use the NetWorker User for SQL Server program or the command prompt features of the NetWorker Module. SQL Server data cannot be restored using the NetWorker Administrator software.

Before the Restore

Review the following prerequisites prior to initiating the restore operation:

- Remember that the Restore browse window displays backups available for restoring based on the restore type selected. The normal and verify-only restore types do not restrict the contents of the browse window, but the partial and copy restore types only display restorable *database* objects. For more information about the browsing items, see ["Browsing Data Items" on page 63](#).
- Ensure that the NetWorker server software is running on the appropriate host and the NetWorker Remote Exec Service is started on the SQL Server host.

- If a backup of another database is in progress, wait for it to finish. Microsoft SQL Server will not restore a database while another database backup is in progress.
- Ensure that all database users are logged off the database. A restore will fail if other users try to use the database while the restore operation is in progress.

Specifying Database Restore Properties

Use the database restore properties to specify preferences for a single, *marked* database before starting a normal, partial, or copy restore type operation.

Note: These properties are *not* available for a verify-only restore type operation.

All of the properties are context sensitive, based on the version of SQL Server running (7.0 or 2000) and the restore type chosen.

Preferences specified for these properties are used for the current restore operation. When a preference is not set, the NetWorker Module uses the default. After the restore operation is complete, and the Restore window is closed, the property values revert back to the NetWorker Module defaults.



Important: The SQL Server databases must be restored in the correct order.

How to Specify Database Item Restore General Properties



Important: The information in this section applies to normal and copy restore types *only*.

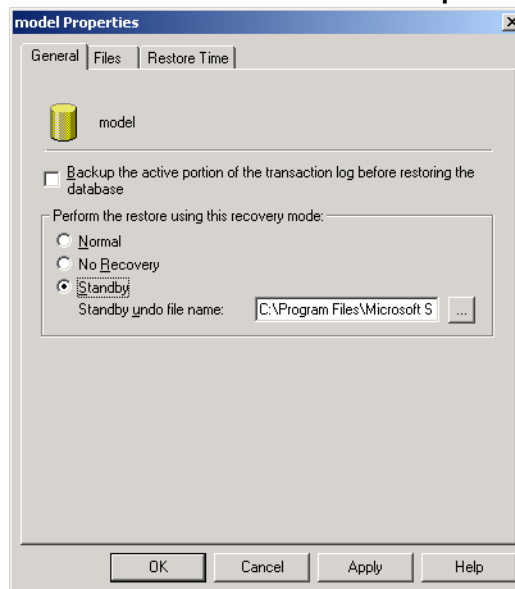
To set database restore properties:

1. In the Restore window, mark a database, then right-click the database and select Properties from the shortcut menu.

The Database Item Restore Properties dialog box opens, as shown in [Figure 18](#).

Note that these properties are supported for SQL Server 7.0 and 2000.

Figure 18. Database Item Restore General Properties Dialog Box



2. Select the following attributes as needed:

- Back Up The Active Portion Of The Transaction Log Before Restoring The Database (default = not checked) — Backs up the active portion of the transaction log prior to performing the restore.

When this attribute is selected, the NetWorker Module attempts a transaction log backup using the NO_TRUNCATE SQL keyword. The restore operation proceeds regardless of whether the transaction log backup succeeds or fails.

Note: If the Specify A Restore Time box under the Restore Time tab is checked, the latest transactions captured in the log tail backup are not restored. The log tail backup proceeds, but a message regarding this restriction appears.

- (Required) Perform The Restore Using This Recovery Mode (default = Normal) — Specify which of the following modes to use for the restore operation:
 - Normal — Restore the database to an operational state.
 - No Recovery — Activate the NORECOVERY option, which places the database in an unloadable state after the restore, but still able to process additional transaction log restore operations.
 - Standby — Activate the STANDBY option, which forces the database to be in a read-only state between transaction log restore operations.

When STANDBY is selected, the Standby Undo File Name attribute is enabled to specify an *undo* file for SQL Server to use when rolling back the transactions. By default, this attribute displays a default filename and path:

`%DriveLetter:\<default_path>\<default_dbName>undo.ldf.`

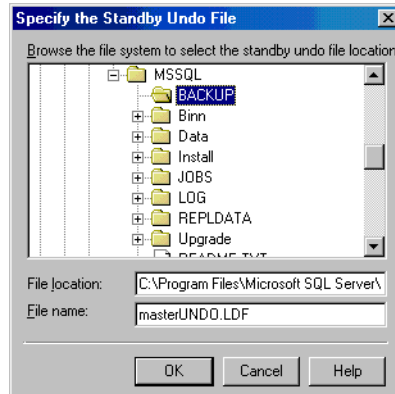
where *default_dbName* is the name of the database backup selected for restoration and *default_path* is the default SQL Server data path obtained from the SQL Server registry. Specify the name and path for the Standby Undo File by typing a valid name and path in the text field, or by clicking the ellipses button to open the Specify the Standby Undo File dialog box.

How to Specify Standby Undo File Properties

Use the Specify the Standby Undo File dialog box to specify the location of the standby undo file as follows:

1. From the Database Item Restore Properties dialog box, click on the ellipses button to open the Specify the Standby Undo File dialog box, shown in [Figure 19](#).

Figure 19. Specify the Standby Undo File Dialog Box



2. Specify attributes as follows:
 - (Required) File Location (default = SQL default data path) — Displays the path for the standby undo file. Modify this attribute by typing a path in the File Location text box, or by browsing the filesystem tree and highlighting a file. When a file in the browse tree is highlighted, the file's path appears in the File Location field.
 - (Required) File Name — Specify the standby undo filename by entering the filename in the text box or by browsing the filesystem tree and highlighting an existing file. When a file is highlighted, the filename appears in the File Name field.

How to Specify Database Item Files Properties

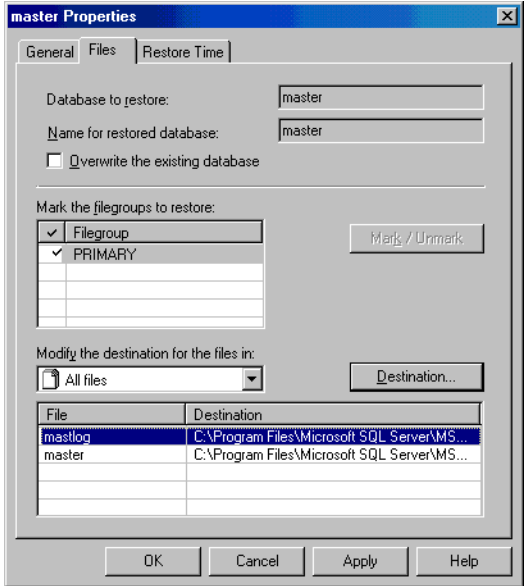
The Database Item Restore Properties dialog box contains a tab in which to select the filegroups to restore and specify the destination locations for the restored files. For normal and copy restore types, this tab is labeled Files and is supported for SQL Server 7.0 and 2000. For a partial restore type, the tab is labeled General and is only available for SQL Server 2000.

Note: The Files tab may first invoke the Read File Configuration dialog box before showing the completed Files tab. This means the marked database item selected was created by a release of NetWorker Module earlier than 3.0, or the most recent backup is a log tail backup for a database that was corrupt. For details, see ["Read File Configuration Properties" on page 93](#).

To specify database item Files Properties:

1. Select the Files tab. The Files Properties dialog box appears, as shown in [Figure 20](#).

Figure 20. Files Properties Dialog Box



2. Specify attributes as follows:

Note: If the text boxes in this dialog box are empty, review the file configuration information. For details, see ["Read File Configuration Properties" on page 93](#).

- (Informational) Database To Restore (default = disabled) — Displays the name of the database selected for restoration. This attribute is informational only and cannot be modified.

- Name For Restored Database — Specify the name for the restored database.
 - If performing a normal restore — By default, this text box displays the name of the database selected for backup and is disabled.
 - (Required) If performing a partial or copy restore — The NetWorker Module software generates the default name by appending *CopyOf* or *PartOf* to the source database name. To specify a different name, enter a new name in the text field or select a name from the list. The name must comply with SQL Server naming conventions.

File configuration is regenerated after clicking the Apply or OK button. If a different database name is specified, the NetWorker Module automatically generates a new physical file configuration. *If a new database name is specified, any changes to the file configuration are discarded.*



Important: When this attribute is set to the name of an existing database, the Overwrite The Existing Database attribute is enabled when changes are applied. These two attributes can then be used together to specify the same name for the restored database as an existing database, when the two databases are incompatible.

- Overwrite The Existing Database (default = not checked) — Instructs the SQL Server to create the specified database and its related files, even if another database already exists with the same name. In such a case, the existing database is deleted.
 - (Optional) — When restoring a database backup whose database ID is the same as that of the existing database, and the files are intact.
 - (Required) — When restoring a database backup whose database ID differs from that of the existing database, or whose files have been damaged. In this case, selecting this attribute causes the WITH REPLACE SQL keyword to be included in the restore sequence. Normally a restore will refuse to overwrite preexisting files, but the WITH REPLACE keyword specifies that it is okay for each file to be restored over existing files of the same name and location.

- Mark The Filegroups To Restore (Mark/Unmark) — Select or unselect the filegroups to restore.
 - (Informational) If performing a normal restore or copy restore (default = disabled) — Displays the filegroups of the database selected for restoration. For more information, see ["Marking Items" on page 66](#).
 - (Required) If performing a partial restore (default = enabled) — By default, this attribute displays the filegroups of the database marked for restoration.

To select or unselect a filegroup for restoration, highlight the filegroup in the list and click the Mark/Unmark button. When a filegroup is marked, a check mark is displayed next to the filegroup name in the list.

You can mark more than one filegroup. The PRIMARY filegroup is always marked and cannot be unmarked. SQL Server requires that the PRIMARY filegroup be included in a partial restore.

Note: The set of filegroups marked in this attribute is copied into the list of the Modify The Destination For The Files In attribute.

- Modify The Destination For The Files In (default = enabled) — Use this attribute to filter which files are visible in the File and Destination lists.

The list contains a set of different views for the database files to be restored. Icons also help identify the filter type. The following views are supported:

- All Files — Displays all of the files for the database, including transaction log files.
- All Log Files — Displays *only* the transaction log files.
- All Data Files — Displays *only* the data files.
- Filegroup Name — Displays only the data files for a specific filegroup.
- Drive Letter — Displays all files located on a given drive at the time the backup occurred, even if those files have since been relocated to a different drive.

To filter the view, click the arrow button and highlight the filegroup or view for which you would like to see the files. The File and Destination lists display the source files and restored files and destinations that correspond to the selected view or filegroup.

- File and Destination lists (default = enabled) — These list the names of files associated with the marked database to be restored, and the destinations to which the files will be restored. The File column is the SQL Server logical filename, and the Destination column is the physical filename and its location.
 - If performing a normal restore — Displays the current name and destination based on the SQL Server physical filename and logical location for the restored file.
 - (Required) If performing a partial or copy restore — Displays a default name and destination based on the SQL Server physical filename and logical location for the restored file.

The File and Destination lists are not editable in this table; however, the Destination entry can be modified by clicking the Destination button. For more information on modifying the destination, see ["Modifying a File Destination"](#) below.

Modifying a File Destination

To modify the destination for one file, do *one* of the following:

- Double-click a file in the list to display the Specify the File Destination dialog box, as shown in [Figure 21 on page 92](#). Then follow the instructions in ["How to Specify the Restored File Destination" on page 92](#).
- Click a file in the list, then click the Destination button to display the Specify the File Destination dialog box. Then follow the instructions in ["How to Specify the Restored File Destination" on page 92](#).

To modify the destination for multiple files, do *one* of the following:

- To select a contiguous range of files, click on the first file to be modified, hold down [Shift], and then click on the last file to be modified.
- To select a series of noncontiguous files, click on the first file to be modified, hold down [Ctrl], and then click each subsequent file until you have selected all of the files to modify.
- To modify the destination for selected files, click the Destination button to display the Specify the File Destination dialog box. Then follow the instructions in ["How to Specify the Restored File Destination" on page 92](#).

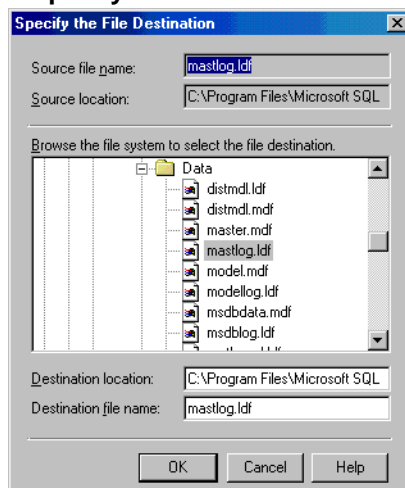
The following restrictions apply to the File and Destination lists:

- If necessary, file configuration is regenerated after clicking the Apply or OK button.
- If a different database name is specified, the NetWorker Module automatically generates a new physical file configuration when the change is applied. *If a new database name is specified, any applied changes to the file configuration are discarded.*

How to Specify the Restored File Destination

Specify the destination locations for the restored files in the Specify the File Destination dialog box, shown in [Figure 21](#).

Figure 21. Specify the File Destination Dialog Box



1. Specify attributes as follows:

- (Informational) Source File Name (default = disabled) — Displays information only and cannot be modified. The name, as it exists in the backup to restore the file currently selected in the File and Destination lists, is displayed in the Source File Name text box. When multiple files are selected, this text box is empty.
- (Informational) Source Location (default = disabled) — Displays information only and cannot be modified. The filesystem location, as it exists in the backup to restore the file currently selected in the File and Destination lists, is displayed in the Source Location text box. When multiple files are selected, this text box contains the filesystem location of the first selected file in the File and Destination lists.

- (Required) Destination Location (default = disabled) — This attribute displays the filesystem location for the restored file. When multiple files are selected, the default SQL data path is opened, but not selected.

This attribute can be modified by typing a path in the Destination Location text box, or by browsing the filesystem tree and highlighting a directory or file. When a directory is highlighted, that path is displayed in the Destination Location text box. If a file is highlighted, the directory for the highlighted file is displayed.

- (Required) Destination File Name — By default, this attribute displays the name of the file currently selected in the File and Destination lists. When multiple files are selected, this text box is empty. When a file is highlighted, the file name is displayed in the Destination File Name text box, and the path containing the file is displayed in the Destination Location text box.

Use this attribute to specify the new name to use for the restored file by entering the file name in the text box or by browsing the filesystem tree and highlighting a file. When a file is highlighted, the file name is displayed in the Destination File Name text box.

Read File Configuration Properties

Some of the data used to populate the attributes under the Files tab of the Database Item Restore Properties dialog box is obtained from new file-configuration meta data objects created in the client file index. For backups created with a version earlier than release 3.1, the file-configuration meta data is not be present in the client file index, but is available in the save set media. Also note that NetWorker Module release 3.0 log-tail backups may not create meta data.

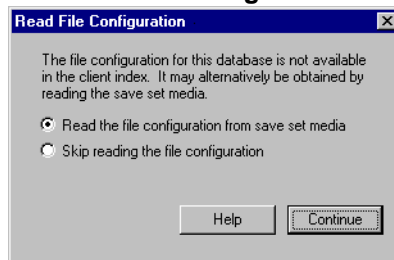
How to Specify Read File Configuration Properties

To specify Read File Configuration properties:

1. Open the Properties dialog box for a marked database item that has no file-configuration meta data in the client file index. You may load this information from the save set media.

The Read File Configuration dialog box opens, as shown in [Figure 22](#).


Figure 22. Read File Configuration Dialog Box



2. Specify attributes as follows:

- **Read The File Configuration From Save Set Media** — If you choose to read from tape, a Reading dialog box opens, as shown in [Figure 23](#), and the save set media is read in the background. When the reading is finished, the Properties dialog box is invoked; the filegroup and destination table has valid data.

Figure 23. Reading Dialog Box



To cancel the reading, select the Cancel button. If you cancel from the Reading dialog box, the Properties dialog box is invoked, but the filegroup and destination table is empty.

- **Skip Reading The File Configuration** — If you elect to skip reading in the Read File Configuration dialog box, the Properties dialog box opens, but the filegroup and destination table is empty.

Database File Relocation Restrictions

The NetWorker Module imposes the following restrictions on database file relocation:

- Only database backups can be relocated. Individual file and filegroup backups cannot be relocated without relocating the database that contains those files.

- If the configuration of a database has changed since the most recent level full database backup was created, you cannot relocate the database. Configuration changes include adding or deleting filegroups, files, or transaction log files.
- System databases may not be the destination database of a relocate operation.
- If the file system relocation destination does not have sufficient space to create a new database, the relocate procedure fails.

How to Specify Database Item Restore Time Properties

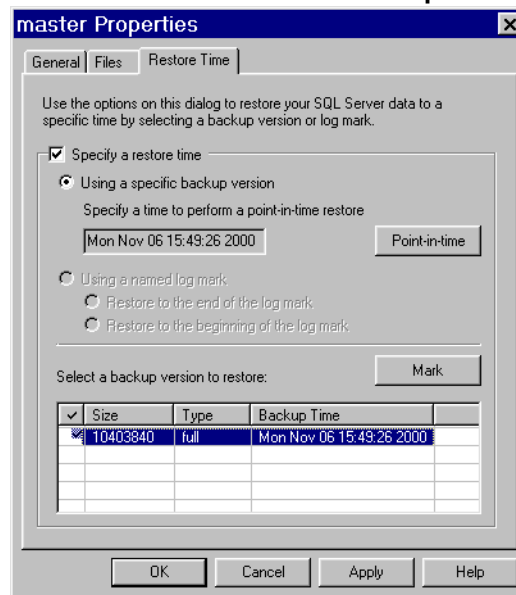
Specify restore time properties under the Restore Time tab, including restoring a specific database backup version or a specific transaction time. The ability to restore to a named log mark is *also* supported for SQL Server 2000.

To specify database item Restore Time properties:

1. In the Database Item Restore Properties dialog box, select the Restore Time tab.

The Restore Time dialog box opens, as shown in [Figure 24](#).

Figure 24. Database Item Restore Time Properties Dialog Box



Note: The Specify A Restore Time check box is unchecked by default. If the Back Up The Active Portion check box under the General tab is checked, a warning message appears indicating the transactions in the log-tail backup will not be present in the restored database.

2. Select the Specify A Restore Time check box to enable the following restore time controls:

- Using A Specific Backup Version (default = grayed-out; if the Specify A Restore Time check box is checked, this radio button is selected by default) — Restore to a specific time. When this attribute is selected, the Specify A Time To Perform A Point-In-Time Restore text box, Point-In-Time button, and Select A Backup Version To Restore list are enabled.
 - Specify A Time To Perform A Point-In-Time Restore (default = the restore time for the backup version selected in the Backup Versions list) — The restore time indicates what data from the marked backup version is reinstated during the restore, and when to stop restoring transactions. If an incremental backup is marked in the Backup Versions list, this text box can be modified by clicking the Point-In-Time button.
 - Point-In-Time button — When a level incremental backup is marked in the Backup Versions list, the Point-In-Time button is enabled for setting the time for a point-in-time backup. For details, see ["How to Specify Point-in-Time Restore Properties" on page 97](#). When any other backup level is marked, this button displays an error message.
- Backup Versions List — (default = most recent backup version is marked) Select and mark the backup version to use for the restore. The Specify A Time text box displays the date and time for the backup that is currently marked. *If you mark a nonincremental backup version from this list, the Point-In-Time button is disabled. Only one backup version can be selected and/or marked.*

The Size column displays the size of the backup. The Type column contains the backup type, including “full” for level full backups, “incr” for level incremental backups, and “1” for level differential backups. The Backup Time column displays the date and time, in granularity of seconds, on which the backup was created.

To select a backup version, double-click any backup in the list, or click any backup in the list and then click the Mark/Unmark button. When a backup version is marked, a check mark appears next to the

backup Size in the list. If another backup version is already marked, the selected backup version is checked and the other backup version is unchecked.

- Using A Named Log Mark (default = not selected) — Perform the restore using a named log mark. *This attribute is only visible in SQL Server 2000 and only enabled when log marks exist for the selected database backup.* When this attribute is selected, the Restore To The End Of The Log Mark and the Restore To The Beginning Of The Log Mark radio buttons are enabled.
- (Required) Restore To (The End, The Beginning) Of The Log Mark (default = the end) — Use this attribute to specify which type of named log mark restore to perform:
 - Restore To The End Of The Log Mark — Restore the backup and stop immediately after the named log mark. This type of restore includes the named transaction in the restore.
 - Restore To The Beginning Of The Log Mark — Restore the backup and stop immediately before the named log mark. This excludes the named transaction.
- (Required) Log Mark List (default = none checked) — Use this attribute to specify the log mark to use for the restore. Only one log mark may be selected.

The Log Mark column displays the name of the transaction log mark. The Date column displays the date and time, in granularity of milliseconds, on which the named transaction log mark was created. The Description column contains any information about the log mark that was entered when the transaction was performed.

To select a log mark, double-click any log mark in the list, or click any log mark in the list and then click the Mark/Unmark button. A check mark appears next to the log mark name in the list. If another log mark is already marked, the selected log mark is checked and the other log mark is unchecked.

How to Specify Point-in-Time Restore Properties

Specify a date and time for the SQL point-in-time restore in the Point-in-Time Restore dialog box. Specify a date and time that falls within a time frame spanned by the transaction log backup that was marked in the Backup Versions list. The create time of the transaction log backup defines the upper boundary of the time frame. No date and time greater than this upper bound can be specified. The lower boundary of the timeframe is the create time of the

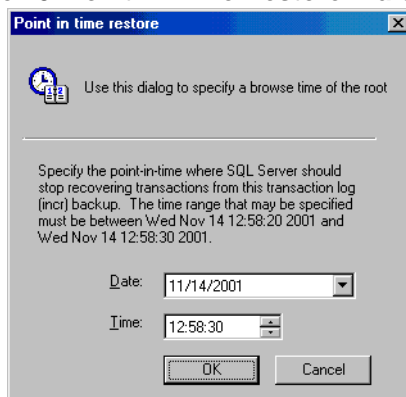
most recent level incremental or level full backup created prior to the selected backup. No date and time less than or equal to this lower bound can be specified.

To specify point-in-time restore properties:

1. On the Restore Time tab, click the Specify A Restore Time check box, then click Using A Specific Backup Version. Click the Point-In-Time button.

The Point-In-Time Restore dialog box opens, as shown in [Figure 25](#).

Figure 25. Point-in-Time Restore Dialog Box



2. Specify the following attributes as needed:

- (Required) Date (default = date when the marked backup version was created) — Use this attribute to specify the restore time date for the marked backup version. The syntax for this field is mm/dd/yy.

To modify the date, enter a date using the appropriate syntax or click the arrow to display the Calendar. In the Calendar, click a date; use the Previous Month and Next Month arrow buttons to change from the current month.

The restore time date must be in the range of transaction dates included in the marked backup version.

- (Required) Time (default = time when the marked backup version was created) — Use this attribute to specify the restore time for the marked backup version. The syntax for this field is hh:mm:ss.

To modify the time, enter a time using the appropriate syntax or use the scroll arrow buttons to change the time.

The restore time must be in the range of transaction times included in the marked backup version.

Backing Up the Active Portion of the Transaction Log

When restoring a file or filegroup, first back up the active portion of the transaction log. This is necessary because the transaction log backup must be applied to the file or filegroup restore to ensure the file or filegroup is consistent with the rest of the database. If a file or filegroup is restored from the NetWorker User for SQL Server program, this transaction log backup occurs automatically. When using the command line to restore the file or filegroup, the transaction log backup must be explicitly performed using the command line **save** program as follows:

```
nsrsqlsv [-s NetWorker_server_hostname] -l incr -R db_name
```

This command instructs the SQL Server to perform a transaction log backup of the database *db_name*. The **-R** flag is required so the transaction log is not truncated after the backup completes.

When restoring a file or filegroup for a Microsoft SQL Server 7.0 database, an additional transaction log backup may be required. This is because, for Microsoft SQL Server 7.0, performing the first transaction log backup may result in additional transactions on the database.

If, after performing the first transaction log backup, the **restore** command fails, use the following procedure to restore the file or filegroup:

1. Put the database into single-user mode.
2. Immediately perform *another* transaction log backup.
3. Restore the desired file or filegroup.
4. Take the database out of single-user mode.

For more information about filegroup backup and restore, refer to the Microsoft SQL Server 7.0 documentation.

Filegroup and File Restore Properties

The filegroup and file restore properties in the Properties dialog box are provided for informational purposes only. SQL Server dictates the settings for these properties, therefore, they are not configurable.

The properties in this dialog box apply to a single, *marked* filegroup or file for a normal restore type operation.

Note: These properties are not available for partial, copy, or verify-only restore type operations.



Important: When restoring a filegroup or file, first back up the active portion of the transaction log. This is necessary because the transaction log backup must be applied to the restore to ensure the filegroup or file is consistent with the rest of the database. For more information, see ["Backing Up the Active Portion of the Transaction Log" on page 99](#).

How to View the Restore Properties

To view the restore properties for a filegroup or file:

1. Right-click a *marked* filegroup or file from either pane of the Restore window.
2. Select Properties from the shortcut menu, or click the Properties button on the task bar.

The Properties dialog box opens, as shown in [Figure 26 on page 101](#) for a filegroup and in [Figure 27 on page 101](#) for a file. These properties were added in NetWorker Module release 3.0 and are version sensitive, based on whether you are running SQL Server 7.0 or 2000.

Figure 26. Filegroup Restore Properties Dialog Box

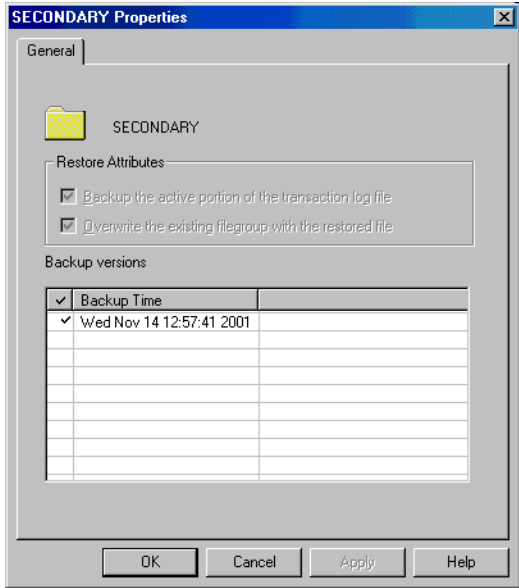
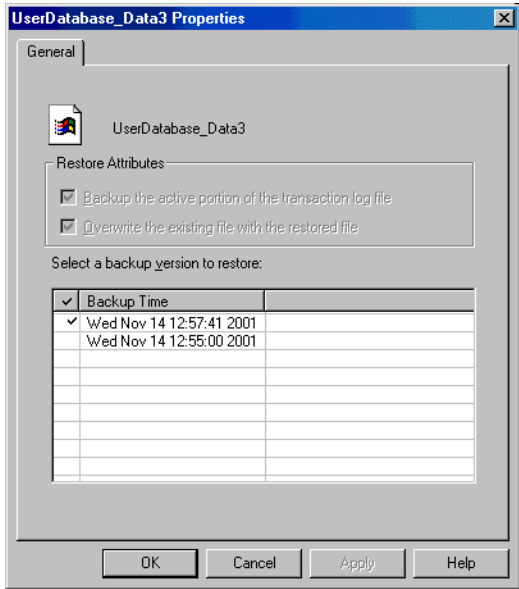


Figure 27. File Restore Properties Dialog Box



The following attributes appear in the Properties dialog box:

- (Informational) Backup The Active Portion Of The Transaction Log File (default = checked) — Signifies that the active portion of the transaction log is backed up prior to performing a restore, so that the log can be applied to the filegroup or file to make it consistent with the rest of the database. It is required by SQL Server when restoring damaged or lost data files. The NetWorker Module attempts a transaction log backup using the NO_TRUNCATE SQL keyword. The restore operation proceeds regardless of whether the transaction log backup succeeds or fails.
- (Informational) Overwrite The Existing Filegroup/File With The Restored File (default = checked) — Forces SQL Server to ignore errors produced on behalf of nonexistent files due to media failure. If there is a media failure, then the files no longer exist. NetWorker Module specifies the WITH REPLACE SQL keyword in the restore sequence. This causes the restore operation to restore the filegroup or file to the exact location (drive and pathname) as the location on the source host from which the data was backed up.
- (Informational) Backup Versions List — (default = most recent backup version is marked) Lists the date and time of all the backups available for restoration.

Performing the Restore

To perform the restore:

1. Start the NetWorker User for SQL Server program. The main window opens, as shown in [Figure 11 on page 72](#).


If multiple SQL Servers are active on the computer, the Select SQL Server Instance dialog box opens before the main window opens, as shown in [Figure 10 on page 72](#).

2. Select the SQL Server host that the NetWorker Module should use to perform the restore operation and click OK. The main window opens.
3. (Optional) Click the Select NetWorker Server button on the task bar.

The Change Server dialog box opens, as shown in [Figure 12 on page 73](#).

4. (Optional) Select a NetWorker server from the list, then click OK.
5. In the main window, click the Restore button on the task bar.


The Restore Operation dialog box appears, as shown in [Figure 28 on page 103](#).

Figure 28. Restore Operation Dialog Box

You can also select Operation>Restore, and select Normal, Copy, Verify Only, or Partial from the menu.

6. Select a restore type from the Restore Operation dialog box and click Continue.

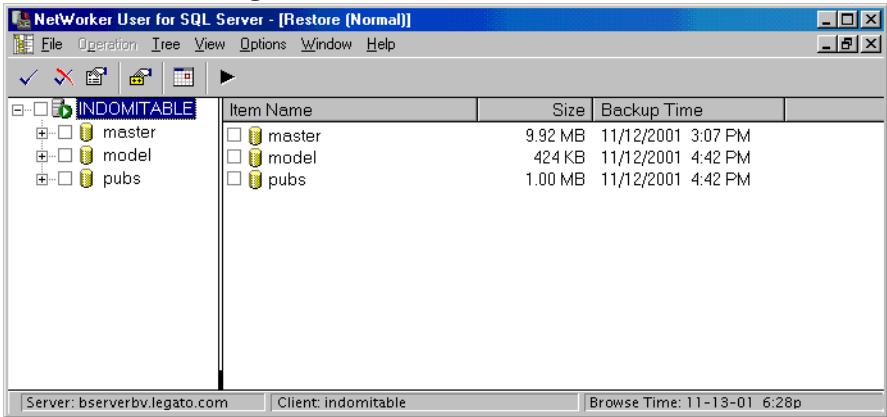
If you Copy restore type is selected, the SQL Server Selection dialog box appears, as shown in [Figure 29](#). In this dialog box, select the NetWorker client from which SQL Server backups are imported for restore to the local SQL Server.

Figure 29. SQL Server Selection Dialog Box

7. Select one of the following attributes to specify the location of the existing database backup to use for the copy operation:
 - The Current SQL Server (default = selected) — By default, this attribute displays the name of the current NetWorker client host. Select this attribute to use a database backup from the SQL Server currently connected to the NetWorker Module.
 - A Different NetWorker SQL Server Client (default = not selected) — By default, this attribute displays the name of the current NetWorker client host. Click the arrow to select a different SQL Server from which to obtain the existing database backup.
8. Click Continue.

The Restore window opens, as shown in [Figure 30](#).

Figure 30. Restore Browse Window



9. To view a list of data items available for restoring, expand any item in the left pane of the Restore window. The descendents of the item are displayed in the right window pane.

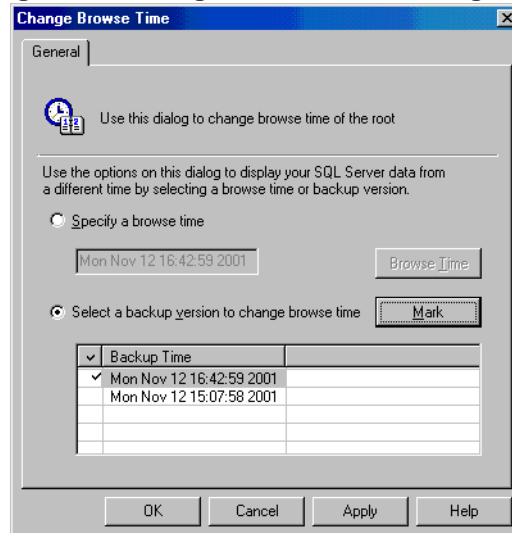


Important: If the *partial* or *copy* restore type is selected, filegroups and files of the selected database are not visible in the browse window. This is because these operations require a database backup. The filegroups and files are available in the Database Item Restore Properties dialog box. For more information, see "[Restore Browse Window Restrictions](#)" on [page 64](#).

10. To change the browse time, select View>Change Browse Time, or click the Change Browse Time button on the task bar.

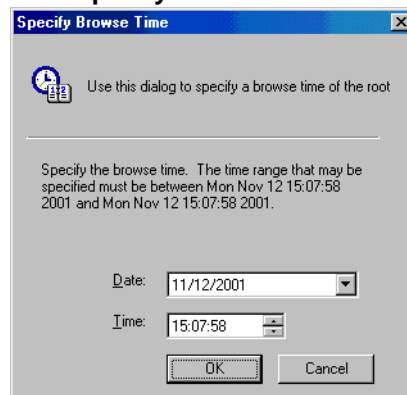
The Change Browse Time dialog box opens, as shown in [Figure 31](#).

Figure 31. Change Browse Time Dialog Box



11. Select one of the following attributes to change the browse time:
 - Specify A Browse Time (default = selected) — By default, this attribute displays the current browse time. Select this attribute to enable the Browse Time button, then click the Browse Time button to invoke the Specify Browse Time dialog box, as shown in [Figure 32](#).

Figure 32. Specify Browse Time Dialog Box



Select a date and time from the list, then click OK. The dialog box closes and the Specify A Browse Time attribute in the Change Browse Time dialog box is updated with the date and time selected. The backup version list is also updated to reflect the new setting.

- Select A Backup Version To Change Browse Time (default = disabled) — By default, this attribute marks the current browse time in the backup version list. Select this attribute to enable the backup version list and Mark button. The backup version list displays the known backup version times, which correspond to save times of the root browse directory. The root browse directory is created anytime a SQL Server backup is performed. To select a new browse time, double-click a backup time in the list, or select a backup time in the list and then click the Mark button. The Specify A Browse Time is updated to reflect the new setting. Click OK to save the new browse time and close the dialog box.



Important: When you change the browse time, all items marked in the Restore window prior to the change are discarded, and any database, filegroup, or file property settings are lost.

12. In the Restore window, mark the item to restore by clicking the check box to the left of the item. For more information, see ["Marking Items" on page 66](#)



Important: If you select the *partial* or *copy* restore type, you can mark only *one* database object for partial restore. For more information, see ["Restore Browse Window Restrictions" on page 68](#).

13. Set the restore properties for each marked database, if any, by right-clicking the database item and selecting Properties from the shortcut menu (or select the marked database and click the Properties button on the task bar). Fill in the Properties dialog box, then click OK. For information about setting database restore properties, see ["Specifying Database Restore Properties" on page 84](#).



Important: When selecting a *partial* or *copy* restore type, specify the restore properties. To do this, right-click the database item and select Properties from the shortcut menu. Then specify the restore properties under the Files tab.

These properties *must* be set before starting the restore operation, or the NetWorker Module does not allow the restore to start.

14. View the restore properties for each marked filegroup or file, if any, by right-clicking the filegroup or file, selecting Properties from the shortcut menu, and then clicking OK. For information about viewing restore properties, see ["Filegroup and File Restore Properties" on page 100](#).
15. Click the Start button to begin the restore.

Note: To cancel the restore, select File>End Restore or close the Restore Status window.

The amount of time to restore databases depends on the amount of data, network traffic, server load, backup volume availability, and tape positioning. While the restore is in progress, the NetWorker Module displays messages in the Restore Status window for each restored database, for monitoring the progress of the restore operation.

If the backup volume containing the databases is loaded at a storage node (backup device) local to the NetWorker server, the restore proceeds. If the restore does not begin, it is possible that either the wrong volume or no volume is mounted in the backup device.

When restoring an incompatible database using the name of an existing database, or when restoring from a media failure where one or more database files were lost, the Overwrite The Existing Database attribute must be selected under the Files tab. To get to the Files tab, right-click the database item and select Properties from the shortcut menu.

After the restore is finished, the restore completion time appears in the Restore Status window.

Chapter 7: Special Features

This chapter provides information about several NetWorker Module features that pertain to backup and restore operations. This chapter includes the following sections:

- ["Database Consistency Check" on page 109](#)
- ["Striping" on page 111](#)
- ["Diagnosing Errors During Backup and Recovery" on page 118](#)
- ["RAW Save Set" on page 118](#)
- ["Interleaving" on page 124](#)
- ["Backing Up and Restoring Multiple Nonclustered Instances of SQL Server" on page 125](#)
- ["Backing Up and Restoring Microsoft Cluster Server Configurations" on page 126](#)

Database Consistency Check

When running Microsoft SQL Server 7.0, use the NetWorker Module `-j` command to request that SQL Server run a comprehensive database consistency check (DBCC) before a backup or after a restore. The DBCC includes the following automatic checks:

- `DBCC CHECKDB database_name`
- `DBCC CHECKALLOC database_name`
- `DBCC TEXTALL database_name`
- `DBCC CHECKCATALOG database_name`

In order for the entire DBCC to complete successfully, each of these tests must succeed.



Important: The option to run a DBCC is disabled by default. To trigger a DBCC before a scheduled save is initiated, see ["How to Perform a Database Consistency Check"](#) below.

If this option is enabled and the DBCC is completed successfully, the NetWorker Module proceeds with a backup of the specified databases. If the DBCC does not complete successfully, the backup terminates and the NetWorker Module displays a message to either:

- The display, indicating DBCC success or failure.
- The Group Detail dialog box in the NetWorker Administrator program.

Database consistency on a restored database can also be checked by initiating a restore from the command line on the restore host using the **-j** command line option with **nsrsqlrc**. This command invokes the DBCC on the database after the restore operation finishes.

You cannot check database consistency from the NetWorker User for SQL Server program, the client user interface.

For more information about **nsrsqlsv** and **nsrsqlrc** command options, see ["Appendix A: Commands and Command Options"](#) on page 143. For more information on database consistency checking, refer to the Microsoft SQL Server documentation.

How to Perform a Database Consistency Check

To perform a DBCC from the NetWorker Administrator program prior to a scheduled backup:

1. Go to the appropriate Client resource.
2. In the Edit Client or Create Client dialog box, select the Remote tab.
3. In the Backup Command attribute, enter **nsrsqlsv -j**.

To perform a DBCC from the Windows command prompt prior to a manual backup:

- Log on to the client host with administrator privileges and enter **nsrsqlsv -j** at a Windows command prompt.

Note: For large databases, running a database consistency check can take several hours. Thus, for production systems, the **-j** switch should be used with caution.

Striping

Legato NetWorker Module release 3.1 supports the use of multiple *stripes* for backup and restoration of Microsoft SQL Server 7.0 and 2000. Stripes are one or more streams of data that may be extracted, in parallel, from a database, and written in parallel to multiple media devices, such as tape drives. With the NetWorker Module, striping can yield a significant performance advantage when a large amount of data is backed up and restored using multiple tape drives.

How to Specify Striping

To specify striping:

1. From the NetWorker User for SQL Server program, specify striping and the number of stripes you want in the Backup Options dialog box.
2. For a scheduled backup, in the NetWorker Administrator program go to the appropriate Client resource.
3. In the Edit Client dialog box, select the Remote tab.
4. In Backup Command attribute, append **-Sn** to the **nsrsqlsv** command, where *n* is the number of stripes you want to use.
5. From the Windows command prompt on the client host, append **-Sn** to the **nsrsqlsv** command, again where *n* is the number of stripes you want to use. For example:

```
nsrsqlsv -s NetWorker_server_name -S3 db_name
```

Note the *lowercase* **-s** to specify the NetWorker server name and the *uppercase* **-S** to specify the number of stripes to use. For more information about using NetWorker Module command options, see ["Appendix A: Commands and Command Options" on page 143](#).

Configuring Striping for Backups

Before any striped backup begins, *the Parallelism attribute must be set to at least one more than the number of stripes you will be using*. Set this attribute in the Set Up Server dialog box in the NetWorker Administrator program. For example, if you are using three stripes, specify a value of four or greater for the parallelism.

Performance Considerations for Striping

For improved backup performance, follow these suggestions:

- In the Create Device or Edit Device dialog box (in the Device resource), configure target sessions for the NetWorker devices performing the striped backup. Always specify one session per device when using striping. This yields the best backup performance on the client host by eliminating interleaving.
- Allot a separate backup device (such as a tape drive) for each stripe in the backup or restore operation. To help ensure that this recommendation is always satisfied, consider using a different NetWorker server for backup and restore operations not related to Microsoft SQL Server, or dedicating a storage node exclusively to the striped backups.
- The NetWorker Module automatically restores data using the same number of stripes specified for the data backup. For best restore performance, ensure that the same number of devices used for the backup are also available at restore time.
- The NetWorker Module operation allows you to run the NetWorker server on a different host from the computer that hosts the NetWorker Module, the NetWorker client, and Microsoft SQL Server. *However, to optimize striping performance, it is best to have the storage node on the same host as these three other components.*

For additional information about restoring backups, see ["Chapter 5: Backing Up SQL Server Data" on page 71](#).

How to Perform a Striped Backup

Legato NetWorker Module supports manual and scheduled backup striping.

To perform a *manual* striped backup from the SQL Server host, which is a NetWorker client, use *one* of the following methods:

- From the Windows command prompt, specify the **-S** option with the **nsrsqlsv** command. For example:

```
nsrsqlsv -s NetWorker_server_hostname -Sn db_name
```

where *n* specifies the number of stripes you want to use.

Specify the number of stripes immediately following **-S**, as for instance **-Sn** in the preceding example. (In these commands, note the lowercase **-s** to specify the NetWorker server name and the uppercase **-S** to specify the number of stripes to use.)

For complete **nsrsqlsv** and **nsrsqlrc** command option syntax, see ["Appendix A: Commands and Command Options" on page 143](#). For details about using the NetWorker Administrator program, refer to the *Legato NetWorker Administrator's Guide*.

- Click the Backup button in the NetWorker User for SQL Server program and then select Backup Options from the Options menu. In the Backup Options dialog box, check Create A Striped Backup and use the down or up arrows to specify the number of stripes you want (see [Figure 15 on page 76](#)).

To perform a *scheduled* striped backup from the NetWorker server host:

1. In the NetWorker Administrator program, go to the appropriate Client resource.
2. In the Edit Client dialog box, select the Remote tab.
3. In Backup Command attribute, enter **nsrsqlsv -S stripes_count**.

Performing a Striped Recovery

NetWorker Module release 3.1 now performs striped recovery optimizations. These optimizations are designed to allow striped recovers to proceed as quickly as possible. This feature, however, represents a fundamental change in the striped recovery strategy, and impacts the ability to restore a striped recovery under certain situations.

Optimal Striped Recovery Versus Fail-Safe Striped Recovery

NetWorker Module release 3.1 now assumes a best-case scenario for striped recovery. Because of this change in product philosophy, the worst-case scenario cited in ["Performing a Fail-Safe Striped Recovery \(Worst-Case Scenario\)" on page 115](#) causes striped recovery to deadlock by default. The deadlock occurs because Microsoft SQL Server is waiting for the NetWorker server to mount a volume, but the NetWorker server has no tape sessions available to grant the volume mount request. The NetWorker Module does not detect this deadlock, and stops responding until the NetWorker server cancels the inactive tape session. This may take as long as 30 minutes. [Table 12](#) provides guidelines for determining when to activate the fail-safe striped recovery algorithm.

Table 12. Guidelines for Fail-Safe Striped Recovery

If...	Then...
The backup was interleaved	Use the NetWorker Administrator program to determine if a second mount request has occurred for a volume already assigned tape sessions.
Fewer devices are available	Use the NetWorker Administrator program to determine if a device has been deactivated for some reason, or is busy performing an operation for another client.

Performing an Optimal Striped Recovery (Best-Case Scenario)

Striped recovery assumes, by default, that enough NetWorker tape sessions are available to allow the striped recovery to proceed without blocking. This is the best-case scenario. The NetWorker Module then assumes the following:

- Each stripe was written to a different volume during backup. Interleaving was not used during the striped backup.
- The same number of devices used during backup is available during restore. No device failures have occurred since the backup.
- All devices used for backup are currently available. No other client sessions are currently assigned to the devices.

If the configuration does *not* meet these requirements, see "[Performing a Fail-Safe Striped Recovery \(Worst-Case Scenario\)](#)" below for more information.

If the configuration *does* meet all of the above requirements, then the NetWorker Module optimally restores the striped recovery. The NetWorker Module begins delivering backup data to the SQL Server as soon as each tape session is established. This method provides the highest performance and lowest restore time.

Previous versions of the NetWorker Module allowed a striped recovery even if one or more of the requirements above were not met. In doing so, the NetWorker Module was required to load each volume containing a stripe member before sending any backup data to the SQL Server. This increased the restore time due to the overhead of loading and reading tapes.

Performing a Fail-Safe Striped Recovery (Worst-Case Scenario)

If a striped backup is interleaved (i.e., multiple stripes were written to the same volume), or if fewer tape sessions are available at restore time (due to a failed or busy device), then a fail-safe striped recovery is necessary.

Enabling this option causes the NetWorker Module to predetermine the number of tape sessions available prior to performing the striped recovery. As a result, the NetWorker Module contacts the NetWorker server with a list of striped save sets and asks how many tape sessions can be simultaneously assigned for the save sets. If the NetWorker server reports a number of sessions that is smaller than the number of striped save sets, then the NetWorker Module uses a fail-safe striped recovery algorithm. The algorithm selected depends on the SQL Server version. In addition, for SQL Server 2000, the fail-safe algorithm may drastically impact striped recovery.

How to Perform a Fail-Safe Striped Recovery on SQL Server 7.0

When enabling a fail-safe striped recovery to restore a SQL Server 7.0 client, the striped recovery requires no additional configuration settings.

To enable a fail-safe striped recovery on SQL Server 7.0:

1. In the Restore window of NetWorker User for SQL Server, right-click the root item in the left pane and select Restore Options.
2. Click the Detect Available Tape Sessions Prior To Restoring A Striped Backup attribute.
3. Start the recovery from NetWorker User for SQL Server or from the command line.

4. If the volumes containing the striped recovery are not managed by an autochanger, then monitor the media wait events in the NetWorker Administrator program. When a media wait event occurs, load the appropriate volume.

Note: You may be required to load a volume more than once due to the fail-safe striped recovery algorithm implemented for SQL Server 7.0. To prevent the restore from deadlocking, the NetWorker Module reads a required amount of information for each stripe into memory before proceeding.

5. (Optional) Once the recovery is complete, disable the fail-safe striped recovery attribute in NetWorker User for SQL Server by following steps 1 and 2 above to uncheck the option.

Because this setting is maintained in the Windows registry, disabling the option allows the next striped recovery to proceed at maximum performance. When performing a fail-safe striped recovery because of a permanent anomaly in the configuration, fixing the situation is recommended.

How to Perform a Fail-Safe Striped Recovery on SQL Server 2000

When enabling a fail-safe striped recovery to restore a SQL Server 2000 client, additional configuration settings are recommended to allow the restore to proceed at maximum throughput.

NetWorker Module release 3.1 takes advantage of a SQL Server 2000 VDI feature called *removable pipes*. This feature allows third-party backup vendors to restore a striped backup from fewer devices. However, there is currently no way for the NetWorker Module to accurately determine how many tape sessions the NetWorker server can assign. The NetWorker server does not currently support striped recoveries by default, and indicates to the NetWorker Module that only one tape session is available to restore the striped save sets during the predetection phase. Because of the way SQL Server 2000 VDI removable pipes function, the third-party backup vendor can only start as many stripe restore threads as there are tape sessions available; otherwise, the restore will deadlock.

Due to the removable pipes requirements and the NetWorker server limitation, the NetWorker Module will, by default, restore one stripe save set at a time for SQL Server 2000 when the fail-safe striped recovery option is enabled. Thus, if data was backed up with two stripes, it takes twice as long to restore. It is possible, however, to enable striped recoveries on the NetWorker server. This is not recommended as a permanent solution, because it may

cause file-system restores to fail. However, it can be enabled temporarily to restore the SQL Server 2000 striped backup to achieve the maximum performance during the restore.

To enable a fail-safe striped recovery on SQL Server 2000:

1. In the Restore window of NetWorker User for SQL Server, right-click the root item in the left pane and select Restore Options.
2. Click the Detect Available Tape Sessions Prior To Restoring A Striped Backup attribute.
3. (Optional) Enable striped recovers on the NetWorker server by creating the file *striped_recovery* in the *~\nsr\debug* directory on the NetWorker server.

Note: It is not necessary to restart the NetWorker services to activate this setting.

4. Start the recovery from NetWorker User for SQL Server or from the command line.
5. If the volumes containing the striped recovery are not managed by an autochanger, then monitor the media wait events in the NetWorker Administrator program. When a media wait event occurs, load the appropriate volume.
6. (Recommended) Once the recovery is complete, disable the fail-safe, striped recovery attribute in NetWorker User for SQL Server by following steps 1 and 2 above to uncheck the option.

Because this setting is maintained in the Windows registry, disabling the option allows the next striped recovery to proceed at maximum performance. When performing a fail-safe striped recovery because of a permanent anomaly in the configuration, fixing the situation is recommended.

7. (Recommended) Once the recovery is complete, disable striped recovers on the NetWorker server by deleting the file *striped_recovery* in the *~\nsr\debug* directory on the NetWorker server.

Diagnosing Errors During Backup and Recovery

NetWorker Module release 3.1 provides enhancements for diagnosing errors that occur during backup and recovery. These enhancements come in the form of an application-specific log that provides the following:

- Software configuration information
- Operation parameters
- Operation status and error messages

An application log is generated for both **nsrsqlsv.exe** and **nsrsqlrc.exe** programs. The logs are written into the *nsr\applogs* folder on the SQL Server host. The logs are cumulative and are appended each time the **Save** or **Recover** program runs. The log space management is rudimentary—the log is truncated when disk space is exhausted.

Log filenames have the following conventions:

- Each program has its own log file:
 - *nsrsqlsv.log*
 - *nsrsqlrc.log*
- Each named instance has a unique log file for both **Save** and **Recover**:
 - *nsrsqlsv.<Instance_name>.log*
 - *nsrsqlrc.<Instance_name>.log*
- Each virtual server has a unique log file for both **Save** and **Recover**:
 - *nsrsqlsv.<Virtual_server>.log*
 - *nsrsqlrc.<Virtual_server>.log*

Note: The logging capabilities of **nsrsqlsv.exe** and **nsrsqlrc.exe** are not cluster-aware. For both clustered and nonclustered configurations, the logs are stored on a local disk.

RAW Save Set

NetWorker Module release 3.0 introduced a performance-enhancing feature called RAW save set. This feature is designed to reduce CPU utilization for enterprise-class configurations. RAW save set works for configurations where the backup storage devices are locally attached to the SQL Server host. This includes configurations where either the NetWorker server or a NetWorker storage node is running on the same system as the NetWorker Module.

Overview

RAW save set takes advantage of the fact that the backup storage devices are locally attached. When RAW save set is enabled, the NetWorker Module bypasses NetWorker tape services provided by the media management daemon (MMD), and instead performs native writes and reads to the backup storage devices. This significantly reduces CPU utilization because of the following:

- Memory copies associated with moving backup data from SQL Server to the NetWorker MMD are eliminated. The NetWorker Module writes backup data received from SQL Server directly to the backup storage device. The shared memory the SQL Server provides is used to perform the backup storage I/O operation.
- Process-context switching that occurs when NetWorker MMD manages the backup storage traffic is eliminated. Again, the NetWorker Module writes backup data received from SQL Server directly to the backup storage device. The backup storage I/O operation occurs in the NetWorker Module process context, not in the NetWorker MMD process context.

RAW save set uses new tape services provided by NetWorker MMD. These tape services allow a backup process to acquire control of the backup storage device from NetWorker MMD controlling the device. Once the NetWorker Module has acquired the device, the NetWorker Module opens a handle for the device and performs I/O operations using the handle. When SQL Server has delivered all of the backup data, the NetWorker Module closes the device handle and hands the device back to the NetWorker MMD.

Key Benefits

RAW save set decreases CPU utilization. Unofficial performance results observe that CPU utilization during backup is significantly reduced. The backup throughput with RAW save set is slightly lower than with a normal backup. However, the CPU utilization should outweigh this slight reduction in backup throughput for configurations that anticipate a high transaction load on the SQL Server host while the backup is in progress. The additional CPU cycles made available through the use of RAW save set are now available to support additional transaction load.

Supported NetWorker Server Features

RAW save set currently provides limited NetWorker server functionality. [Table 13](#) lists NetWorker features and the expected RAW save set behavior for each one.

Table 13. RAW Save Set Interaction with NetWorker Server (Part 1 of 2)

NetWorker Server Feature	Description	Expected RAW Save Set Behavior
Archiving	File system feature.	Not currently supported.
Cloning	Save set-level replication following scheduled save.	Not currently supported.
Device configuration	Configuration of the tape device for optimal performance.	<ul style="list-style-type: none"> • The NetWorker Module relies on the MMD to perform the device configuration. • The NetWorker Module always uses a device block size of 64KB.
Directed recovery	Restoration of a SQL Server database from a different NetWorker client.	Not supported if the SQL Server clients do not both reside on the same physical computer.
Dynamic drive sharing (DDS)	Sharing of backup storage among hosts on a SAN.	Not currently supported.
Media errors	Handling of errors such as end-of-tape and tape cleaning.	<ul style="list-style-type: none"> • The NetWorker Module assumes the device is already in a mode compatible with device cleaning. • The NetWorker Module detects the end of the media. • The NetWorker Module detects limited tape errors.
Media pool configuration	Configuration of the NetWorker server to direct backup traffic to specific volumes.	Ensure the volumes are used exclusively by RAW save set. Do this using media pools.

Table 13. RAW Save Set Interaction with NetWorker Server (Part 2 of 2)

NetWorker Server Feature	Description	Expected RAW Save Set Behavior
Media notifications	Satisfying a save or restore tape session request for a volume.	<ul style="list-style-type: none"> • Server should load correct volume required for RAW save set. • Server should successfully recycle the volume with RAW save set.
Removable Storage Manager (RSM)	Windows 2000 feature.	Not currently supported.
Save set cutoff	Limit on maximum save set size.	RAW save set ignores the cutoff size.
Scanning	Reconstruction of media and file index entries from tape structures.	Not currently supported. ^a
Spanning	Crossing of a media boundary by a save set.	<ul style="list-style-type: none"> • The NetWorker Module detects the end of the media and requests a new volume. • RAW save set media spanning is supported.
Staging	Writing a backup to a local disk and then moving it to tape.	Staging is unsupported because file type devices may not be enabled on the NetWorker server.
Tape Broker	Assigning tape sessions to the appropriate device.	<ul style="list-style-type: none"> • Tape broker prevents other sessions from being assigned once a RAW save set is active. • Tape broker does not currently prevent a RAW session from being assigned when other sessions are currently active.

a. To recover an index associated with multiple volumes containing RAW save sets, use the **nmrecov.exe** command.

Supported Configurations

RAW save set support is limited to the following configurations:

- NetWorker Server 6.0 or higher — Media management services required for RAW save set to function properly are not available in prior NetWorker server releases.
- NetWorker Server Power Edition or higher — RAW save set also functions when the NetWorker server is in evaluation mode, so customers can evaluate how this feature affects CPU performance.
- File type devices must be disabled on the NetWorker server — RAW save set does not function properly with file type devices. Because the NetWorker Module cannot prevent the NetWorker server from assigning tape sessions to file type devices, the NetWorker Module chooses not to support RAW save set on configurations with file type devices enabled.

In addition to the above configuration requirements, the following requirement must also be met:

- Backup storage devices must be locally attached — RAW save set is only supported when the storage device is locally attached, because the NetWorker Module must open a handle to the backup storage device. If the NetWorker server assigns a backup device to the NetWorker Module that is not locally attached, *the backup or restore operation fails*.

RAW Save Set Restrictions

Because RAW save set backups are closely tied to the configuration on which they are performed, restoration of these backups is limited to that particular configuration; backups are not portable. This restricts copy restore (directed recover) operations. In general, you cannot perform a copy restore of a RAW save set to a different SQL Server host. Copy restore is only possible when the SQL Server instance to which you want to restore the backup is running on the same computer as the SQL Server instance for which the backup was performed.

Recommended Configuration

To deploy RAW save set in your environment, the following configuration steps are recommended:

1. Locally attach the backup storage device to the SQL Server host.
2. Install and configure either the NetWorker server or the NetWorker storage node. To consolidate the index and policy management onto one NetWorker server, then installation and configuration of the NetWorker storage node is recommended.
3. Configure a volume pool dedicated to RAW save set backups. Ensure that the pool selection criteria for this pool excludes all other client data from being saved on volumes labeled for this pool.
4. Label volumes in the locally attached backup storage for the RAW save set volume pool.

Performing a RAW Save Set Backup or Restore

Once properly configured, you can perform a RAW save set backup or restore.

To perform a RAW save set *scheduled* backup from the NetWorker Administrator program:

1. Log on with administrator privileges to the NetWorker server host.
2. Start the NetWorker Administrator program. Under the Configure tab, click Manage Clients.
3. In the Clients window, right-click the name of a client and select Edit. The Edit Client dialog box appears.
4. In the Edit Client dialog box, click the Remote tab.
5. In the Backup Command attribute, add the `-r` parameter to the command, then click OK.

To perform a RAW save set *manual* backup from NetWorker User for SQL Server:

1. In the Backup window of NetWorker User for SQL Server, right-click the root item in the left pane and select Backup Options.
2. Under the Advanced tab, click Use The RAW Save Set, then click OK.

Note: This setting is not persistent across NetWorker User for SQL Server program invocations.

To perform a RAW save set *manual* backup from the command line **save** program:

- From the command line **save** program, **nsrsqlsv.exe**, specify the **-r** flag, in addition to any other parameters.

To perform a RAW save set restore:

- No special action is required. If a backup was performed using RAW save set, the NetWorker Module automatically performs the restore using RAW save set.

RAW Save Set Notifications

The NetWorker Module provides the following notifications for RAW save set:

- If all configuration requirements are met (see ["Supported NetWorker Server Features" on page 120](#)), the following message appears:
RAW save set is enabled
- If the NetWorker server version is not sufficient to support RAW save set, the following message appears:
Server version does not support raw save set.
- If the NetWorker server licensing is not sufficient to support RAW save set, the following message appears:
Server is not properly licensed to support raw save set.
- If the NetWorker server device configuration is not sufficient to support RAW save set, the following message appears:
Server is not properly configured to support raw save set.

Interleaving

Legato NetWorker Module supports *interleaving* for backing up multiple clients to a single backup device.



Important: Although NetWorker Module release 3.1 supports striped interleaving, interleaving stripes is not recommended. Interleaved backups can require substantial restore time. This is because, when restoring an entire interleaved backup, the one stripe at a time is restored, thus requiring multiple passes over the same areas on the backup medium.

However, interleaving is useful when backing up multiple clients concurrently to one device. When restoring an interleaved backup, restore each client's backup separately.

Backing Up and Restoring Multiple Nonclustered Instances of SQL Server

Microsoft SQL Server 2000 allows you to run multiple copies of SQL Server on a single computer. This feature is called *multiple instance*. Microsoft allows up to 16 named SQL Server instances to run simultaneously.

The first installation of SQL Server on a computer is called the *default instance*. The name of the default instance is the network name for the local computer. NetWorker Module release 3.1 provides support for backup and restoration of SQL Server 7.0 and SQL Server 2000 default instances.

Additional installations of SQL Server on a computer are called *named instances*. Each named instance has a unique instance name of the form:

computerName\instanceName

where *computerName* is the network name of the computer, and *instanceName* is the logical name of the named instance.

The syntax for specifying a nonclustered instance of SQL Server at a command prompt is:

`MSSQL$instanceName[:dbName ...] [.fgName ...] [.fileName . . .]`

For detailed guidelines on specifying the save sets for an instance, see ["Backup and Restore Command Syntax for SQL Server Data" on page 156](#).

The following restrictions apply when running multiple instances:

- For Microsoft SQL Server 2000, a maximum of 16 named SQL Server instances can be run simultaneously on a single computer.
- The **nsrsqlsv** and **nsrsqlrc** commands only support specification of one instance at a time. If save sets for more than one instance are specified, the backup or restore operation fails.

- A *nonclustered* named instance may not be the source for a copy (directed recover) operation when the destination server is different from the source. This restriction does not affect *clustered* named instances.
 - The **nsrsqlrc** command fails if a directed recover from one nonclustered named instance to another SQL Server instance is attempted.
 - The NetWorker User for SQL Server program prevents the selection of a named instance as the source for a directed recover operation; named instances do not appear in the NetWorker Client list.
- Index entries for nonclustered named instances are created using the local host on which the instance is running. To differentiate backups for the default instance and named instances, the index name has been extended to logically partition the index for a client.

Backing Up and Restoring Microsoft Cluster Server Configurations

This section explains the NetWorker Module release 3.1 backup and restore procedure for Microsoft Cluster Server (MSCS) clusters, and how this procedure relates to MSCS cluster support provided for NetWorker filesystem clients.



Important: You must have NetWorker Power Edition or higher installed on the NetWorker server host in order to use the NetWorker Module to back up and restore SQL Server data on a node in a cluster.

For more detailed information about configuring the NetWorker Module to run in a cluster environment, refer to the *Legato NetWorker Module for Microsoft SQL Server Installation Guide*.



Important: If the NetWorker Module is backing up or restoring a SQL Server virtual server, then the Windows account used to execute the NetWorker Module **save** or **restore** commands *must* be an MSCS administrator member. To determine which accounts have MSCS administrator privileges, refer to the MSCS Online Help. If the Windows account does not have MSCS administrator privileges, the SQL Server modules cannot communicate with MSCS and the various MSCS cluster resources, including the SQL Server virtual servers.

How NetWorker Works with MSCS Clusters

An MSCS cluster appears as a set of nodes and virtual servers. On Windows NT 4.0 and Windows 2000, MSCS allows a cluster to contain two nodes. On Windows 2000 Data Center, the maximum number of supported nodes is four.

Each node is a physical computer with its own IP address and network name. A cluster may be configured to contain any number of virtual servers. Each virtual server has its own IP address and network name. Each virtual server also owns a subset of shared cluster disks and is responsible for starting cluster applications that can fail over from one cluster node to another.

The NetWorker filesystem client software treats each physical node and each virtual server in a cluster as a separate NetWorker client. Each NetWorker client has its own client index on a NetWorker server, and each client is responsible for backing up its own files. A client associated with a physical node in the cluster backs up files on private disks attached to that node. A client associated with a virtual server backs up only files on disks in the cluster resource group belonging to that virtual server.



Important: In order to back up all of the cluster's data, both the NetWorker Module software and the NetWorker client software must be installed on a private disk on each physical node in the cluster.

NetWorker Module 3.1 Support for MSCS Clusters

The NetWorker Module can back up or restore data exported by SQL Server running as a virtual server in an MSCS cluster.

The NetWorker Module uses the virtual server name to:

- Connect to the appropriate SQL Server.
- Initialize the SQL Server VDI to accept data from, or deliver data to, the appropriate SQL Server in the cluster.
- Create entries in the NetWorker client file index.

The NetWorker media database or client file index are indexed according to the client that performs a particular backup. NetWorker Module release 3.1 creates entries under the virtual server name in the NetWorker client file index. Information about the SQL Server save sets of a virtual server is stored in the NetWorker client file index associated with that virtual server name.

Named Instances in Failover Cluster Configurations

The NetWorker Module provides failover cluster support by utilizing the multiple instance features provided in SQL Server 2000. In a failover configuration, virtual servers run as either the default instance or as named instances. One default instance of a virtual server may be installed. Additional virtual servers may be installed as named instances, where each instance name must be unique within the cluster.

The NetWorker Module supports up to 16 named instances as virtual servers in a cluster configuration, where each named instance virtual server has the following qualities:

- A unique IP address, network name, and instance name
- Datafiles installed on a clustered drive running in the same cluster group as the SQL virtual server

Configuring Scheduled Backups of a Clustered SQL Server

To configure a scheduled backup of a clustered SQL Server:

1. In the NetWorker Administrator program, create a Client resource for each SQL Server virtual server in the cluster. For the Name attribute, specify the name of the virtual server. For detailed instructions on creating a Client resource, refer to the *Legato NetWorker Module for Microsoft SQL Server Installation Guide*.

2. Under the Remote tab, in the Remote Access attribute, grant access to all physical nodes in the cluster by adding entries similar to the following:

`administrator@cluster_computer_name`

This allows the NetWorker Module server access to the cluster node to authenticate the computer (on which the virtual server is running) as a NetWorker Module client before any backup or restore operation begins. Do this for each virtual server Client resource in the cluster.

3. For the Remote User and Password attributes, add the user name and password, respectively, for a Windows user account that has SQL Server administrator privileges and NT administrator privileges. This allows the NetWorker Module to back up the SQL Server virtual server. Do this for each virtual server Client resource in the cluster.

Performing Manual Operations on a Clustered SQL Server

In a cluster environment, a manual (unscheduled) NetWorker Module operation can be initiated from:

- A Windows command prompt
- The NetWorker User for SQL Server program
- Any computer in the cluster, regardless of whether a failover has occurred, and regardless of which computer the virtual server is running on at the time

Operations from the Command Line

To perform a command line operation with SQL Server running in a cluster, use the **-a** command option with the **nsrsqlsv** or **nsrsqlrc** command, for example:

```
nsrsqlsv -a virt_srvr_name -s NetWorker_server_hostname db_name
```

Specifying **nsrsqlsv -a virt_srvr_name** causes the NetWorker Module software to contact the SQL Server virtual server and to create save set entries under *virt_srvr_name* in the NetWorker client index.

Operations from the NetWorker User for SQL Server Program

In the NetWorker User for SQL Server program, the Select SQL Instance command on the Operation menu opens the Select SQL Instance dialog box. Here you can select or change the SQL Server host for a backup or restore operation in a cluster environment. The Select SQL Instance command is

enabled in the main NetWorker User for SQL Server window only when multiple, individual SQL Server instances are detected. This command is not present when a Backup or Restore window is open.

When running in an MSCS cluster, NetWorker Module release 3.1 automatically detects all active SQL Servers in the cluster, including virtual servers. This automatic detection occurs whenever:

- The NetWorker User for SQL Server program starts up
- The Select SQL Instance menu item is selected
- A backup or restore is initiated

Active/Passive Cluster Configurations

When the NetWorker User for SQL Server program starts up on the primary node, the NetWorker Module automatically uses the cluster virtual server as the client name for reading or writing to the NetWorker media database and client index. If SQL Server fails over to the secondary cluster node, opening the NetWorker User for SQL Server program on the secondary node also automatically uses the virtual server name.

Active/Active Cluster Configurations

This situation occurs when there is one SQL Server virtual server running on each physical node in the cluster. In this case, if no failover has occurred and each virtual server is running on a different physical node, the NetWorker Module automatically communicates with the virtual server running on the same physical node as the NetWorker Module software, and the Operation>Select SQL Instance command is disabled.

However, if one of the physical cluster nodes goes offline (for example, if a failover occurs) and both SQL Server virtual servers are then running on the same physical cluster node, when the NetWorker User for SQL Server program starts up, the Select SQL Server dialog box appears with a list of the SQL Server instances. Select an instance from this dialog box. The NetWorker Module then attempts to validate the instance as a NetWorker client.

NetWorker Module Interaction with SQL Server Services

When restoring the *master* database in a clustering environment, the NetWorker Module automatically stops and then restarts the SQL Server services as appropriate, as follows:

1. Before the restore begins, the NetWorker Module takes the SQL Server virtual server resources offline.
2. The NetWorker Module starts the local SQL Server service in single-user mode.
3. The NetWorker Module performs the restore.
4. After the restore is complete, the NetWorker Module stops the local SQL Server service.
5. The NetWorker Module brings the SQL Server virtual server resources back online.

For information about using the NetWorker Module commands and command options with MSCS clusters, such as **-a**, see ["Appendix A: Commands and Command Options" on page 143](#).

Chapter 8: Disaster Recovery

Using the NetWorker Module for regular backups provides for easy restoration of lost SQL Server data. Lost SQL Server binaries can also be restored, in certain configurations, when using the NetWorker client for regular filesystem backups of the ALL save set. This chapter includes the following sections:

- ["When to Reinstall Microsoft SQL Server" on page 136](#)
- ["Disaster Recovery of SQL Server Primary Data Disk Crash or Data Corruption" on page 136](#)
- ["Disaster Recovery of SQL Server Binary Disk Crash or Data Corruption" on page 137](#)
- ["Disaster Recovery of SQL Server and NetWorker Server" on page 139](#)
- ["Disaster Recovery of SQL Server \(Shortcut\)" on page 140](#)
- ["Disaster Recovery of SQL Server" on page 142](#)



Important: The SQL Server databases must be restored in the correct order. Follow the database restore steps in ["Chapter 6: Restoring SQL Server Data" on page 83](#).



Important: Because of the many variables, it is not possible to provide exact disaster recovery instructions for every situation. The examples in this chapter provide general principles and procedures for data restoration.

Depending upon the severity of the disaster, restoration of the operating system, NetWorker server, or NetWorker client may be more intensive than indicated in this document. Therefore, before beginning a SQL Server disaster recovery, review the *Legato NetWorker Disaster Recovery Guide*, the *SQL Server Books Online* documentation, and the appropriate version of the *Legato NetWorker Release Supplement* for your system.

The NetWorker Module provides the following features for disaster recovery:

- System database restore automation — certain system databases require SQL service control. This includes the *master* database and the *msdb* database. NetWorker Module automates the control of these services as follows:
 - For the *master* database, the SQL Server service restarts in single-user mode as required by Microsoft SQL Server.
 - For the *msdb* database, the SQL Agent shuts down to close connections to *msdb*.

For more information on how the NetWorker Module controls SQL Server services, see ["NetWorker Module Interaction with SQL Server Services" on page 60](#).

- Database restore order — when restoring a complete backup of all databases, or when restoring certain system databases, the restore must occur in a specific order. When system databases are present in the restore list, the NetWorker Module ensures the restore order follows Microsoft SQL Server procedures as follows:
 - The *master* database is always restored first. This ensures meta data present in the *master* database is correct for all subsequent restored databases.
 - The *msdb* database is always restored after *master* and before all other databases. This ensures scheduling and other system data present in the *msdb* database is correct for all subsequent restored databases.
 - The model database is always restored after *master* and *msdb*, and before all other databases. This ensures the database configuration is correct for all subsequent restored databases.

- **SQL Server Startup Complete detection** — when the SQL Server service starts, it launches a database startup process. The SQL Server service allows user connections while the startup process is running. However, if the startup process is interrupted by a database RESTORE query, any database that hasn't yet started up is marked as suspect. When the interruption occurs, subsequent restore of the *msdb* database will fail. Subsequent restore of any user database may also fail if the Overwrite attribute is not specified for the restore.

NetWorker Module will wait for SQL Server to complete the database startup process for all databases prior to performing a restore. NetWorker Module will also wait prior to initial login and after restore of the *master* database. This wait process allows proper restore of the *msdb* database and user databases following a SQL Server startup. It is recommended that **-f** always be specified for restores that follow a SQL Server service startup.

- **Disaster recovery settings** — the NetWorker User for SQL Server program can now specify that all databases be overwritten when performing a restore. This is necessary when restoring from a disk crash. See the Files Properties dialog box in [Figure 20 on page 88](#).



Important: The NetWorker Module does not handle all dependent services. When restoring application services, such as Meta Data Services or Replication services, and the databases on which these services depend, these services must be manually shut down.

In general, the NetWorker Module does not ensure that all connections to a database are closed prior to restoring a database as required by Microsoft SQL Server. Open connections to all databases being restored must be manually terminated. For more information, refer to the *SQL Server Books Online* documentation.

Following are basic instructions for disaster recovery if NetWorker Module binary files or Microsoft SQL Server binary files are damaged or lost. The instructions explain how to recover from a disaster on a NetWorker Module client host running Windows NT 4.0 or Windows 2000, and SQL Server 7.0 or SQL Server 2000. For more information about using NetWorker software for disaster recovery, refer to the *Legato NetWorker Disaster Recovery Guide* and the Microsoft SQL Server documentation.

When to Reinstall Microsoft SQL Server

When performing a disaster recovery of Microsoft SQL Server using NetWorker Module 2.0.1 or later on a NetWorker 5.7 or later client host running Microsoft Windows NT 4.0 or Windows 2000, it is not necessary to reinstall SQL Server, if the following conditions are true:

- A recent NetWorker backup of the ALL save set exists
- SQL Server was installed prior to the creation of the ALL backup version
- SQL backups for all of the SQL Server databases exist

If SQL Server is already installed when the ALL backup is performed, critical state information for SQL Server is backed up as part of the ALL save set. When the ALL save set contains this information, and backups for all of the SQL Server databases exist, the SQL Server can be reinstated by simply restoring the ALL save set and the SQL databases.

Disaster Recovery of SQL Server Primary Data Disk Crash or Data Corruption

If the primary disk containing critical SQL Server data is damaged, perform the following steps:

1. Shut down SQL Server.



Important: The Rebuild Master utility requires SQL system database files located in the ...*Data* directory of the SQL Server installation CD-ROM or shared network drive.

For SQL Server 2000, the database files must be copied to a temporary directory on the SQL Server. When running SQL Server 2000 under Windows 2000, the read-only attributes from the database files must be removed as well. The temporary directory is also the source directory for data files requested by the Rebuild Master utility.

2. Run the Rebuild Master utility, **rebuildm**, located in the SQL ...*Binn* directory.
3. Restart SQL Server.
4. Start the NetWorker Module for SQL Server program.

5. Use the **replace** option to restore the SQL Server *master* database and the *msdb* database. If running SQL Server 7.0, the model and the distribution databases may need to be restored.
6. Restore the other SQL application databases.

Steps 5 and 6 can also be performed as a single operation from the command line. At a DOS prompt, run the following command:

```
nsrsqlrc -s ... -f MSSQL:
```

Note: Steps 5 and 6 can be performed as a single operation from the NetWorker User for SQL Server program:

1. From the Operation menu, select Restore>Normal. The Restore (Normal) window opens.
2. Right-click the SQL Server root and select Mark All Databases.
3. From the Options menu, select Restore Options. The Restore Options dialog box appears.
4. Click Automatically Overwrite Any Database Marked For Restore, and click OK.
5. Click the Start button.

Disaster Recovery of SQL Server Binary Disk Crash or Data Corruption

If the disk containing SQL Server binaries is damaged, perform the following steps:

1. Shut down SQL Server.
2. Restore the ALL save set. For more information about the ALL save set, see ["When to Reinstall Microsoft SQL Server" on page 136](#).
3. Restart the computer.
4. Shut down SQL Server.



Important: The Rebuild Master utility requires SQL system database files located in the ...*Data* directory of the SQL Server installation CD-ROM or shared network drive.

For SQL Server 2000, the database files must be copied to a temporary directory on the SQL Server. If running SQL Server 2000 under Windows 2000, the read-only attributes must be removed from the database files as well. The temporary directory is also the source directory for datafiles requested by the Rebuild Master utility.

5. Run the Rebuild Master utility, **rebuildm**, located in the SQL ...*Binn* directory.
6. Restart SQL Server.
7. Start the NetWorker Module for SQL Server program.
8. Use the **replace** option to restore the SQL Server *master* database and the *msdb* database. If running SQL Server 7.0, the model and the distribution databases may need to be restored.
9. Restore the other SQL application databases.

Steps 8 and 9 can also be performed as a single operation from the command line. At a DOS prompt, run the following command:

```
nsrsqlrc -s ... -f MSSQL:
```

Note: Steps 8 and 9 can be performed as a single operation from the NetWorker User for SQL Server program:

1. From the Operation menu, select Restore>Normal. The Restore (Normal) window opens.
2. Right-click the SQL Server root and select Mark All Databases.
3. From the Options menu, select Restore Options. The Restore Options dialog box appears.
4. Click Automatically Overwrite Any Database Marked For Restore, then click OK.
5. Click the Start button.

Disaster Recovery of SQL Server and NetWorker Server

If the disks containing the SQL Server binaries and the NetWorker online indexes (the *nsr* filesystem) and/or the SQL Server data are damaged, then perform the following steps:

1. Restore the NetWorker binaries and online indexes. For details, refer to the *Legato NetWorker Disaster Recovery Guide*.
2. Reinstall the NetWorker software.
 - If the computer that failed was the NetWorker server host, reinstall the NetWorker server software.
 - If the computer that failed was a NetWorker client, reinstall the NetWorker client software.
3. Restore the ALL save set. For more information about the ALL save set, see ["When to Reinstall Microsoft SQL Server" on page 136](#).
4. Restart the computer.
5. Shut down SQL Server.



Important: The Rebuild Master utility requires SQL system database files located in the ...*Data* directory of the SQL Server installation CD-ROM or shared network drive.

For SQL Server 2000, the database files must be copied to a temporary directory on the SQL Server. If running SQL Server 2000 under Windows 2000, the read-only attributes must be removed from the database files as well. The temporary directory is also the source directory for data files requested by the Rebuild Master utility.

6. Run the Rebuild Master utility, **rebuildm**, located in the SQL ...*Binn* directory.
7. Restart SQL Server.
8. Start the NetWorker Module for SQL Server program.

9. Use the **replace** option to restore the SQL Server *master* database and the *msdb* database. If running SQL Server 7.0, the model and the distribution databases may need to be restored.

10. Restore the other SQL application databases.

Steps 9 and 10 can also be performed as a single operation from the command line. At a DOS prompt, run the following command:

```
nsrsqlrc -s ... -f MSSQL:
```

Note: Steps 9 and 10 can be performed as a single operation from the NetWorker User for SQL Server program:

1. From the Operation menu, select Restore>Normal. The Restore (Normal) window opens.
2. Right-click the SQL Server root and select Mark All Databases.
3. From the Options menu, select Restore Options. The Restore Options dialog box appears.
4. Click Automatically Overwrite Any Database Marked For Restore, then click OK.
5. Click the Start button.

Disaster Recovery of SQL Server (Shortcut)

If the Microsoft SQL server does not need to be restored (to determine this, see ["When to Reinstall Microsoft SQL Server" on page 136](#)), use the following steps for SQL Server disaster recovery:

1. Reinstall the operating system.
2. Reinstall the NetWorker software.
 - If the computer that failed was the NetWorker server host, reinstall the NetWorker server software.
 - If the computer that failed was a NetWorker client, reinstall the NetWorker client software.
3. Restore the ALL save set. For more information about the ALL save set, see ["When to Reinstall Microsoft SQL Server" on page 136](#).
4. Restart the computer.
5. Shut down SQL Server.



Important: The Rebuild Master utility requires SQL system database files located in the ...*Data* directory of the SQL Server installation CD-ROM or shared network drive.

For SQL Server 2000, the database files must be copied to a temporary directory on the SQL Server. If running SQL Server 2000 under Windows 2000, the read-only attributes must be removed from the database files as well. The temporary directory is also the source directory for data files requested by the Rebuild Master utility.

6. Run the Rebuild Master utility, **rebuildm**, located in the SQL ...*Binn* directory.
7. Restart SQL Server.
8. Start the NetWorker Module for SQL Server program.
9. Use the **replace** option to restore the SQL Server *master* database and the *msdb* database. If running SQL Server 7.0, the model and the distribution databases may need to be restored.
10. Restore the other SQL application databases.

Steps 9 and 10 can also be performed as a single operation from the command line. At a DOS prompt, run the following command:

```
nsrsqlrc -s ... -f MSSQL:
```

Note: Steps 9 and 10 can be performed as a single operation from NetWorker User for SQL Server:

1. From the Operation menu, select Restore>Normal. The Restore (Normal) window opens.
2. Right-click the SQL Server root and select Mark All Databases.
3. From the Options menu, select Restore Options. The Restore Options dialog box appears.
4. Click Automatically Overwrite Any Database Marked For Restore, then click OK.
5. Click the Start button.

SQL Server should now be restored to the most recent backup.

Disaster Recovery of SQL Server

To recover the SQL Server:

1. Reinstall the operating system.
2. Reinstall the NetWorker software.
 - If the computer that failed was the NetWorker server host, reinstall the NetWorker server software.
 - If the computer that failed was a NetWorker client, reinstall the NetWorker client software.
3. Reinstall the SQL Server software.
4. Start SQL Server.
5. Use the **replace** option to restore the SQL Server *master* database and the *msdb* database. If running SQL Server 7.0, the model and the distribution databases may need to be restored.
6. Restore the other SQL application databases.

Steps 5 and 6 can also be performed as a single operation from the command line. At a DOS prompt, run the following command:

```
nsrsqlrc -s ... -f MSSQL:
```

Note: Steps 5 and 6 can be performed as a single operation from the NetWorker User for SQL Server program:

1. From the Operation menu, select Restore>Normal. The Restore (Normal) window opens.
2. Right-click the SQL Server root and select Mark All Databases.
3. From the Options menu, select Restore Options. The Restore Options dialog box appears.
4. Click Automatically Overwrite Any Database Marked For Restore, then click OK.
5. Click the Start button.

Appendix A: Commands and Command Options

This appendix provides syntax, notation, and related information for NetWorker Module for Microsoft SQL Server commands and command options. Many command line examples are provided, where each example consists of a command, command options and variables, and the appropriate syntax for specifying the SQL Server data to use for the operation.



Important: The command line syntax examples provided in the following sections are presented using a multiple-line format in order to improve the readability of the command. When actually entering a command at the command prompt, the command, command options, and specified data must all be entered on a single line.

NetWorker Module Commands

Legato NetWorker Module commands perform the following functions:

- **nsrsqlsv** — Backs up the specified SQL Server data objects. For more information, see ["Using the nsrsqlsv Command" on page 151](#).
- **nsrsqlrc** — Restores the specified SQL Server data objects. For more information, see ["Using the nsrsqlrc Command" on page 144](#).
- **nwmssql** — Invokes the client graphical user interface. For more information, see ["Using the nwmssql Command" on page 156](#).

A

Descriptions, command options, and usage scenarios for each command are defined in the following sections. However, when initiating the commands, keep in mind these notes regarding syntax:

- Case is very important when specifying command-line flags. Each command option is either lowercase or uppercase and, frequently, both the upper and lowercase of a letter are included in the set of command options. For example, **-r** initiates use of the RAW save set, while **-R** sets the NO_TRUNCATE option when backing up transaction logs.
- Depending on the command option, the space separator between an option and its corresponding argument can be optional, required, or not allowed. For example, the following expressions are both valid:

-1 *backup_level*

-1*backup_level*

while the following expression is invalid because a space is not allowed between the **+** argument and *log_mark_name*:

-M + *log_mark_name*

- Brackets ([]) are used to denote the optional portions of a command (for example, command options and corresponding arguments, if any). When initiating an actual backup or restore operation, do not include brackets.
- Data items must follow all other command options and parameters on the command line.

An application log is generated for both **nsrsqlsv.exe** and **nsrsqlrc.exe** programs. The logs are written into the *nsr\applogs* folder on the SQL Server host. The logs are cumulative and are appended each time the **save** or **recover** program runs. The log space management is rudimentary—the log is truncated when disk space is exhausted. For more information about logging, see ["Diagnosing Errors During Backup and Recovery" on page 118](#).

Using the nsrsqlrc Command

The **nsrsqlrc** command restores specified SQL Server data (databases, filegroups, and files) from the NetWorker system. To initiate a restore operation, specify **nsrsqlrc** and any command options at a Windows command prompt.

The nsrsqlrc Command Options

Example:

```
nsrsqlrc [-fjqV] [-a virtual_server_name] [-c client_host_name]
        [-Dn]
        { [-d MSSQL:destination_dbName]
          [-C file=path,file2=path2,...] }
        [-M [+|-] log_mark_name] [-R fgName1,fgName2,...]
        [-t date] -s NetWorker_server_hostname
        [-S normal|standby:undo_file|norecover]
        [[-U username] [-P password]]
        [MSSQL: dbName dbName.fgName dbName.fgName.filename ...]
```

Command options defined:

- a** Specifies the *virtual_server_name*, when SQL Server is configured to run as a clustered service in an MSCS cluster. For more information, see ["Backing Up and Restoring Microsoft Cluster Server Configurations" on page 126](#).
- c** Specifies the NetWorker *client_host_name* from which the SQL Server data is to be restored.
- C** Uppercase **-C** copies the database being restored to either the same SQL Server or a different SQL Server. It can be used for normal, copy, and partial restores. Use *relocation_list* to specify new locations for restored database files. The relocation list is composed of pairs of logical database filenames and fully qualified database filename relocation paths. You can specify the relocation list only when restoring a database. Each filename and relocation path is joined by the equal sign (=), and pairs are delimited in the list by commas. The syntax is:

```
[ " ] [ \ ] file [ ' ] = [ \ ] path [ ' ] , [ \ ] file2 [ ' ] = [ \ ] path2 [ ' ] , ... [ " ]
```

For example, to copy a database named *Project* from a client host named *ClientHost1* to a different location on *ClientHost1*, specify the relocation list for the database files, but do not include the client host name command option:

```
nsrsqlrc -s NetWorker_server_hostname
-d MSSQL:CopyOfProject
-C Project_Data=C:\Relocation\Path\Project_Data.mdf,
  Project_Data2=C:\Relocation\Path\Project_Data2.ndf,...,
  Project_Log=C:\Relocation\Path\Project_Log.ldf MSSQL:Project
```

The relocation list may be enclosed by double quotes to allow for spaces in the relocation elements and pathnames. A logical filename or relocation path may be enclosed by single quotes to also allow for spaces. If a filename or path contains a single quote, precede the single quote with another single quote to prevent the NetWorker Module software from parsing the single quote as a delimiter. For example:

```
nsrsqlrc -s NetWorker_server_hostname
-d MSSQL:CopyOfProject
-c "'File'='C:\Relocate Dir\Path\',
    ='C:\Relocate Dir\Path\'',...,
    ''='C:\Relocate Dir\Path\''' MSSQL:Project
```

When no relocation list is specified, the NetWorker Module reads the source database filenames and location from the client index meta data or the backup. This information is used to generate a default list by relocating all files to the default SQL data path for the target SQL server. The filenames are guaranteed to be unique, but sufficient disk space is not ensured. For more information about performing a copy operation, see ["Chapter 3: NetWorker Module Restore" on page 49](#).

- d** Performs a **copy** operation (Microsoft SQL Server 7.0 and higher). The **copy** operation restores SQL Server data from a client host to another database name on the same client host. The syntax is:

```
nsrsqlrc -s NetWorker_server_hostname -c client_host_name
-d MSSQL:destination_dbName MSSQL:source_dbName
```

where *source_dbName* is the name of the SQL database to restore, and *destination_dbName* is the name of the SQL database to which the source database is to be restored.

- Dn** Uppercase **-D** specifies a debug message level. This causes the NetWorker Module to display all debug messages defined at or below the given level. [Table 14](#) defines debug message levels displayed based on the value for *n*.

Table 14. Debug Message Levels Defined (Part 1 of 2)

Debug Level	Messages Displayed
0	No debug messages enabled
1 – 5	<ul style="list-style-type: none">• Additional information on operation status• Additional information on fatal error messages

Table 14. Debug Message Levels Defined (Part 2 of 2)

Debug Level	Messages Displayed
6, 7	Undefined; will default to level 5
8, 9	<ul style="list-style-type: none"> • Additional information on cluster-related operation status • Additional information on cluster-related fatal error messages
Greater than 9	Undefined; will default to level 9

Note: When **-C**, **-M**, **-R**, or **-d** are used, the list of data objects can include only *one* database. The list of data objects must follow all other command options and parameters on the command line. For detailed guidelines on formatting these parameters, see ["Backup and Restore Command Syntax for SQL Server Data" on page 156](#).

- f** Performs a REPLACE restore operation. This option restores a source database to an existing, incompatible database of the same name on the target host, replacing the target database with the source. This option is also used to restore damaged files.
- j** Performs a database consistency check between the SQL Server data backup and the restored SQL Server data. For more information about this option, see ["Chapter 7: Special Features" on page 109](#).

A

- M** Performs a SQL Server data restoration of the named transaction specified in *log_mark_name* (Microsoft SQL Server 2000 only). When the mark name is prefixed with a plus sign (+), the data is restored to and includes the named transaction. When the mark name is prefixed with a minus sign (-), the data is restored up to the named transaction. The *log_mark_name* should immediately follow the plus or minus symbol; using a space separator is not allowed. The default is the plus sign.

For example, to restore the SQL data to and including the named transaction *transaction_mark*, enter the following command:

```
nsrsqlrc -s NetWorker_server_hostname  
-M +transaction_mark MSSQL:dbName
```

To restore the SQL data *only* to the named transaction *transaction_mark*, enter the following command:

```
nsrsqlrc -s NetWorker_server_hostname  
-M -transaction_mark MSSQL:dbName
```

Only one SQL Server database may be specified, and the database must follow all other command options and parameters on the command line.

- P** Specifies the Microsoft SQL Server user password. When the **-U** user name command option is specified, the password command option must also be provided, as follows:

```
nsrsqlrc -s NetWorker_server_hostname -U username  
-P password MSSQL:
```

Use the SQL Server user name and password to log onto SQL Server using SQL Server integrated security.

- q** Displays **nsrsqlrc** messages in *quiet mode*, providing minimal information about the progress of the restore operation, including error messages.

- R** Uppercase **-R** performs a partial database restore of the specified filegroups (Microsoft SQL Server 2000 only). The partial database restore operation restores specific filegroups from a single full SQL Server database backup. Supply the filegroups to the **-R** command option in a list, with items separated by commas. The **-C** parameter may be used, and should specify all files for the database. The **-d** parameter is also required:

```
nsrsqlrc -s NetWorker_server_hostname
```

```
-R [ " ] [ \ ]fgName[ ' ], [ \ ]fgName2[ \ ], [ \ ]... [ \ ] [ " ]
```

```
-C Project_Data=C:\Relocation\Path\Project_Data.mdf,  
    Project_Data2=C:\Relocation\Path\Project_Data2.ndf,...,  
    Project_Log=C:\Relocation\Path\Project_Log.ldf MSSQL:Project
```

```
-d MSSQL:PartOfProject MSSQL:Project
```

where *Project* is the name of the SQL database to restore, *PartOfProject* is the name of the SQL database to which the source database is to be restored, and *fgName,...* are the names of the filegroups to restore.

The set of parameters supplied to **-R** may be enclosed by double quotes to allow for spaces between the filegroup names. Each filegroup name may be enclosed by single-quotes to also allow for spaces. If a filegroup name contains a single quote, precede the single quote with another single quote to prevent the NetWorker software from parsing the single quote as a delimiter.

For example, to allow for the space in *Filegroup A*, the space after the first comma, and the single quote in *Filegroup A'*, use the following syntax:

```
-R " "Filegroup A ', 'Filegroup A ' ' "
```

When an empty relocation list is supplied:

```
-R ""
```

The NetWorker Module *only* restores the PRIMARY filegroup. For more information about partial restore, see ["Chapter 3: NetWorker Module Restore" on page 49](#).

- s** Lowercase **-s** specifies the NetWorker server to use for the restore operation. *Use of this command option is highly recommended.*

- S** Uppercase **-S** performs the restoration according to the specified database restore mode of normal, standby, or no recovery (Microsoft SQL Server 2000 only). The syntax is:

```
nsrsqlrc -s NetWorker_server_hostname [-d destination_dbName]
-S normal | [ "standby:undo_file" ] | norecover
MSSQL:source_dbName
```

where *source_dbName* is the name of the SQL database to restore, and *destination_dbName* is the name of the SQL database to which the source database is to be restored.

normal mode restores the database in normal, operational mode

standby mode activates the STANDBY option, which forces the database to be in a read-only state between transaction log restore operations.

no recovery mode activates the NORECOVER option, which places the database in an unloadable state after the restore, but still able to process additional transaction log restore operations.

For example, to restore a database named *Project* in normal, operational mode to a new database named *NewProjectName*, enter the following command:

```
nsrsqlrc -s NetWorker_server_hostname -S normal
-d MSSQL:NewProjectName MSSQL:Project
```

To restore the database in standby mode, the **standby** parameter must be immediately followed by a colon, and the standby undo file location and filename must be specified. When a space occurs in the file name or location, the set of parameters supplied to **-S** may be enclosed by double quotes as follows:

```
nsrsqlsv -s NetWorker_server_hostname
-S "standby:C:\temp\undo filename"
-d MSSQL:NewProjectName MSSQL:Project
```

For more information about using the database restore modes, see ["Specifying Database Restore Properties" on page 84](#).

- t** Lowercase **-t** performs a restoration of SQL Server data as of the specified *date*. When the date of a backup version occurs before or is equivalent to the *date*, the backup version is restored. Follow the **nsr_getdate** command syntax guidelines when formatting the *date*.

- U** Specifies the Microsoft SQL Server user name. When the user name command option is specified, the **-P** password command option must also be provided, as follows:

```
nsrsqlrc -s NetWorker_server_hostname -U username  
-P password MSSQL:
```

Use the SQL Server user name and password to log onto SQL Server using SQL Server integrated security.

- V** Uppercase **-V** verifies the SQL Server database selected for restoration (Microsoft SQL Server 2000 only). The **-V** command option only verifies that the selected database backup is suitable for restoring, *the backup is not restored*. The syntax is:

```
nsrsqlrc -s NetWorker_server_hostname -V MSSQL:dbName
```

A

Using the nsrsqlsv Command

The **nsrsqlsv** command is used to back up SQL Server data objects, which consist of databases, filegroups, and files. To initiate a backup operation, specify **nsrsqlsv** and its command options at one of the following locations:

- In the NetWorker Administrator program:
 1. Under the appropriate server, go to the Configure tab and click Manage Clients.
 2. Right-click the client you want to back up and select Edit. The Edit Client dialog box appears.
 3. Under the Remote tab, in the Backup Command text box enter **nsrsqlsv** and its command options.
(Exception: The **-b** and **-l** command options are valid *only* for manual backups initiated from a Windows command prompt on a client host. *Do not use either of these options when initiating a scheduled save in the NetWorker Administrator program.*)
 4. Click OK.
- At a Windows command prompt.

The nsrsqlsv Command Options

Example:

```
nsrsqlsv [-CGjqRtV] [-a virtual_server_name] [-b pool_name]
[-c client_host_name] [-Dn] [-g group] [-l backup_level]
[-N name] -s NetWorker_server_hostname
[-S stripes_count]
[[-U username] [-P password]]
[MSSQL: dbName dbName.fgName dbName.fgName.filename ...]
```

Command options defined:

- a** Specifies the *virtual_server_name* when SQL Server is configured to run in an MSCS cluster. For more information, see ["Backing Up and Restoring Microsoft Cluster Server Configurations"](#) on page 126.
- b** Assigns a destination pool for a save set. Specifying **-b pool_name** overrides all other pool-selection criteria provided by the NetWorker software or specified in the NetWorker User for SQL Server Backup Options Dialog box. The pool must be created, with a corresponding label template, before running a command that includes the **-b** option. For more information about pool-selection criteria, see ["Volume Pools"](#) on page 41.
- c** Lowercase **-c** specifies the NetWorker *client_host_name* for which the SQL Server data is to be backed up.
- C** Uppercase **-C** specifies compression of the backup data before the NetWorker client sends the data to the NetWorker server.
- Dn** Uppercase **-D** specifies a debug message level. This causes the NetWorker Module to display all debug messages defined at or below the given level. [Table 15](#) defines debug message levels displayed based on the value for *n*.

Table 15. Debug Message Levels Defined (Part 1 of 2)

Debug Level	Messages Displayed
0	No debug messages enabled
1 – 5	<ul style="list-style-type: none">• Additional information on operation status• Additional information on fatal error messages
6, 7	Undefined; will default to level 5

Table 15. Debug Message Levels Defined (Part 2 of 2)

Debug Level	Messages Displayed
8, 9	<ul style="list-style-type: none"> • Additional information on cluster-related operation status • Additional information on cluster-related fatal error messages
Greater than 9	Undefined; will default to level 9

Note: When **-C**, **-M**, **-R**, or **-d** are used, the list of data objects can include only *one* database. The list of data objects must follow all other command options and parameters on the command line. For detailed guidelines on formatting these parameters, see ["Backup and Restore Command Syntax for SQL Server Data" on page 156](#).

- g** Lowercase **-g** specifies the save group. The NetWorker server and the **savegrp** command use the **group** parameter to select the media pool.
- G** Uppercase **-G** specifies a NO_LOG transaction log backup before backing up the database. This command option is valid *only* for level full backups.
- j** Performs a database consistency check before initiating the backup. For more information about this option, see ["Chapter 7: Special Features" on page 109](#).
- l** Specifies the backup level. Valid values are *full*, *level 1* (equivalent to specifying *diff* in BusinessSuite Module 2.0 and 2.0.1), and *incr* (equivalent to specifying *xlog* in BusinessSuite Module 2.0 and 2.0.1). For more information about which backup levels are supported for various SQL Server data objects, see ["Combining Data Objects to Create Backup Levels" on page 33](#).



Important: The **-l** option is valid only for manual backups initiated from a Windows command prompt on a client host. For scheduled backups, set the backup level in the Set Level dialog box of the Schedule resource in the NetWorker Administrator program. Do not use the **-l** option when initiating a backup in the NetWorker Administrator program from the Backup Command attribute of the Edit Client dialog box (under the Remote Tab).

Following are sample display outputs for each of the three different levels of backup:

- Full backup of the database *my_database* to a specified NetWorker server host:

```
nsrsqlsv -s NetWorker_server_host my_database
nsrsqlsv: Backing up my_database...
nsrsqlsv: BACKUP database my_database TO
virtual_device='BSMSQL' WITH stats
nsrsqlsv: my_database: level=full, 33 MB 00:00:05
1 file(s)
```

- Differential backup:

```
nsrsqlsv -s NetWorker_server_host -l diff my_database
nsrsqlsv: Backing up my_database...
nsrsqlsv: BACKUP database my_database TO
virtual_device='BSMSQL' WITH differential,
stats
nsrsqlsv: my_database level=diff, 719 KB
00:00:05      1 file(s)
```

- Incremental backup:

```
nsrsqlsv -s NetWorker_server_host -l incr my_database
nsrsqlsv: Backing up my_database...
nsrsqlsv: BACKUP log my_database TO
virtual_device = 'BSMSQL'
nsrsqlsv: my_database level=incr, 61 MB
00:00:05      1 file(s)
MSSQL: dbName dbName.fgName dbName.fgName.filename ...
```

At least one SQL Server data item (database, filegroup, or file) must be specified, and the data items and list of data objects must follow all other command options and parameters on the command line. For detailed guidelines on formatting these parameters, see ["Backup and Restore Command Syntax for SQL Server Data" on page 156](#).

- N** Specifies the symbolic name of the save set. By default, the most common prefix of the path arguments is used as the save set name.

- P** Specifies the Microsoft SQL Server user password. When the **-U** user name command option is specified, the password command option must also be provided, as follows:

```
nsrsqlrc -s NetWorker_server_hostname -U username  
-P password MSSQL:
```

Use the SQL Server user name and password to log onto SQL Server using SQL Server integrated security.

- q** Displays **nsrsqlsv** messages in quiet mode; only summary information and error messages are displayed.
- r** Lowercase **-r** performs a backup using RAW save set. When performing a striped backup with five or more output devices directly connected to the SQL Server, using RAW save set may lower CPU utilization. For more information, see ["RAW Save Set" on page 118](#).
- R** Uppercase **-R** uses the NO_TRUNCATE option when backing up transaction logs.
- s** Lowercase **-s** specifies the NetWorker server to use for the backup operation. *Use of this command option is highly recommended.*
- S** Uppercase **-S** backs up the specified data items using *n* stripes.



Important: To use backup and restore striping successfully, see the striping configuration described in ["Striping" on page 111](#).

- T** Uppercase **-T** performs a TRUNCATE_ONLY transaction log backup before backing up the database; valid for full backups only.
- U** Specifies the Microsoft SQL Server user name. When this command option is specified, the **-P** password command option must also be provided, as follows:

```
nsrsqlrc -s NetWorker_server_hostname -U username  
-P password MSSQL:
```

Use the SQL Server user name and password to log onto SQL Server using SQL Server integrated security.

- v** Lowercase **-v** displays **nsrsqlsv** messages in verbose mode, providing detailed information about the progress of the backup operation.

Using the nwmssql Command

The **nwmssql** command invokes the NetWorker User for SQL Server program, the NetWorker Module's client graphical user interface. To invoke the client graphical user interface, specify **nwmssql** at one of the following locations:

- From the Windows Start menu, go to NetWorker Group>NetWorker Module for Microsoft SQL Server.
- As a shortcut on the Windows desktop.

To create a desktop shortcut, go to the `<install_path>\nsr\bin` directory and drag the `nwmssql.exe` file to your desktop while pressing the [Ctrl] key.

nwmssql Command Options

nwmssql [-t] -s NetWorker_server_hostname

- s** Lowercase **-s** specifies the NetWorker server to use. *Use of this command is highly recommended.*
- t** Lowercase **-t** diagnoses the current backup or restore command issued by the NetWorker Module. The NetWorker User for SQL Server program displays the full backup or restore command in the operation status window, but does not execute the backup or restore operation.

Backup and Restore Command Syntax for SQL Server Data

With the standard Networker backup and restore commands (**nsrsqlsv** and **nsrsqlrc**), use the additional command syntax shown in [Table 16 on page 157](#) to back up or restore SQL Server data.

Enter the Networker commands with the SQL Server data syntax for either scheduled or unscheduled backups as follows:

- Scheduled backup — In the Networker Administrator program, enter the command in the Save Set attribute of the Create Client or Edit Client dialog box.
- Unscheduled backup — Enter the command at the Windows command prompt on the NetWorker server or from the NetWorker User for SQL Server program.

Note that at least one SQL Server data item (database, filegroup, or file) must be specified for an unscheduled backup or restore.

You can specify more than one data object and combine different types of data. SQL data objects must be specified per the syntax shown in [Table 16](#).

Table 16. Command Syntax for SQL Server Data

SQL Server Data	Syntax for SQL Server Data Objects
All databases in the SQL Server storage hierarchy (optional)	MSSQL:
Specified databases	[MSSQL:] <i>dbName</i> [[MSSQL:] <i>dbName</i> [MSSQL:] <i>dbName2 ...</i>]
All filegroups in specified databases	[MSSQL:] <i>dbName.</i> [[MSSQL:] <i>dbName.</i> [MSSQL:] <i>dbName2 ...</i>]
Specified filegroups in specified database	[MSSQL:] <i>dbName.fgName</i> [[MSSQL:] <i>dbName.fgName</i> [MSSQL:] <i>dbName.fgName2</i> [MSSQL:] <i>dbName2.fgName</i> [MSSQL:] <i>dbName2.fgName2 ...</i>]
Specified files in specified filegroups in specified databases	[MSSQL:] <i>dbName.fgName.filename</i> [[MSSQL:] <i>dbName.fgName.filename</i> [MSSQL:] <i>dbName.fgName2.filename</i> [MSSQL:] <i>dbName2.fgName.filename</i> [MSSQL:] <i>dbName2.fgName2.filename ...</i>]

Specifying **MSSQL** before each individual data object name is optional and does not affect the expression or the resulting operation. However, when **MSSQL** is specified, it must be followed by a colon (:). For example, the following two commands are equivalent:

```
nsrsqlsv -s NetWorker_server_hostname dbName.fgName
```

```
nsrsqlsv -s NetWorker_server_hostname MSSQL:dbName.fgName
```

In a nonclustered named instance configuration, **MSSQL\$** is required, followed by the instance name and a colon. For example:

```
nsrsqlsv -s NetWorker_server_hostname MSSQL$instanceName:dbName.fgName
```

Syntax for a Named Instance Configuration for SQL 2000

When the configuration contains nonclustered named instances of SQL Server, the name of the instance should be specified before the data, as follows:

```
MSSQL$instanceName[ :dbName ... ] [ .fgName ... ] [ .fileName ... ]
```

For example, to back up all of the databases for *instanceOne*, enter the following:

```
nsrsqlsv -s NetWorker_server_hostname MSSQL$instanceOne:
```

To restore several filegroups for *instanceTwo*, specify:

```
nsrsqlrc -s NetWorker_server_hostname  
MSSQL$instanceTwo:dbName.fgName  
MSSQL$instanceTwo:dbName.fgName2
```

Clustered named instances should *not* use this syntax. Clustered instance SQL Server Virtual Server names should instead be passed using the **-a** or **-c** parameters. For example:

```
nsrsqlsv -s NetWorker_server_hostname -a SQL_virtual_server_DNS_name MSSQL:
```

```
nsrsqlsv -s NetWorker_server_hostname -c SQL_virtual_server_DNS_name MSSQL:
```

where *SQL_virtual_server_DNS_name* is the Domain Name System (DNS) name for the SQL virtual server. A Client resource should be created under this name.

For scheduled saves of a SQL virtual server client, it is not necessary to specify **-a** or **-c** parameters with the SQL virtual server name. The **savegrp** process automatically specifies the virtual server name to the **nsrsqlsv** process using the **-m** parameter.



Important: The **nsrsqlsv** and **nsrsqlrc** commands only support specification of a single instance. If save sets for more than one instance are specified, the backup will fail. The **restore** command supports mixing of instances for directed recovers; however, a nonclustered named instance cannot be the source. For more information about running multiple instances of SQL Server, see ["Backing Up and Restoring Multiple Nonclustered Instances of SQL Server" on page 125](#).

Syntax for Names Containing a Period, Backslash, or Colon

NetWorker Module release 3.1 introduces a new command line syntax that allows you to back up and restore databases, filegroups, and filenames containing a period, backslash, or colon. By typing a backslash (\) before the period or backslash, the **save** and **restore** commands interpret the period or backslash as a literal character.

Table 17 below, Table 18 on page 160, Table 19 on page 161, and Table 20 on page 161 show the syntaxes for databases, filegroups, and filenames containing a period, backslash, colon, or any combination of the three.

Keep the following notes in mind regarding the information in the tables:

- The item command line syntax shown in the right-hand column is used to specify the items for the **save** or **recover** commands.
- The notation **MSSQL:** remains optional only for the **save** command.
- A single period continues to delimit SQL identifiers.
- The syntax also supports named instances.
- In each case, the backslash period (\.) character sequence replaces each literal period in the SQL identifier.
- In each case, the double backslash (\\) character sequence replaces each literal backslash in the SQL identifier.

Table 17. Command Syntax for Names Containing a Period (Part 1 of 2)

Name As Visible from SQL Utilities	Equivalent Command Line Syntax
The SQL database named <i>MyDatabase.COM</i>	MyDatabase\. COM MSSQL:MyDatabase\. COM MSSQL\$MyInstance:MyDatabase\. COM
The SQL filegroup named <i>MyFileGroup.2</i> for the SQL database named <i>MyDatabase.COM</i>	MyDatabase\. COM. MyFileGroup\. 2 MSSQL:MyDatabase\. COM. MyFileGroup\. 2 MSSQL\$MyInstance:MyDatabase\. COM. MyFileGroup\. 2

Table 17. Command Syntax for Names Containing a Period (Part 2 of 2)

Name As Visible from SQL Utilities	Equivalent Command Line Syntax
The SQL file named <i>MyFile.2</i> , which is a member of the SQL filegroup named <i>MyFileGroup.2</i> for the SQL database named <i>MyDatabase.COM</i>	MyDatabase\COM.MyFileGroup\2.MyFile\2 MSSQL:MyDatabase\COM.MyFileGroup\2.MyFile\2 MSSQL\$MyInstance:MyDatabase\COM.MyFileGroup\2.MyFile\2

Table 18. Command Syntax for Names Containing a Backslash

Name As Visible from SQL Utilities	Equivalent Command Line Syntax
The SQL database named <i>MyDatabase\COM</i>	MyDatabase\\COM MSSQL:MyDatabase\\COM MSSQL\$MyInstance:MyDatabase\\COM
The SQL filegroup named <i>MyFileGroup\2</i> for the SQL database named <i>MyDatabase\COM</i>	MyDatabase\\COM.MyFileGroup\2 MSSQL:MyDatabase\\COM.MyFileGroup\2 MSSQL\$MyInstance:MyDatabase\\COM.MyFileGroup\2
The SQL file named <i>MyFile\2</i> , which is a member of the SQL filegroup named <i>MyFileGroup\2</i> for the SQL database named <i>MyDatabase\COM</i>	MyDatabase\\COM.MyFileGroup\2.MyFile\2 MSSQL:MyDatabase\\COM.MyFileGroup\2.MyFile\2 MSSQL\$MyInstance:MyDatabase\COM.MyFileGroup\2.MyFile\2

Table 19. Command Syntax for Names Containing a Colon

Name As Visible from SQL Utilities	Equivalent Command Line Syntax
The SQL database named <i>MyDatabase:COM</i>	<code>MyDatabase:COM</code> <code>MSSQL:MyDatabase:COM</code> <code>MSSQL\$MyInstance:MyDatabase:COM</code>
The SQL filegroup named <i>MyFileGroup:2</i> for the SQL database named <i>MyDatabase:COM</i>	<code>MyDatabase:COM.MyFileGroup:2</code> <code>MSSQL:MyDatabase:COM.MyFileGroup:2</code> <code>MSSQL\$MyInstance:MyDatabase:COM.MyFileGroup:2</code>
The SQL file named <i>MyFile:2</i> , which is a member of the SQL filegroup named <i>MyFileGroup:2</i> for the SQL database named <i>MyDatabase:COM</i>	<code>MyDatabase:COM.MyFileGroup:2.MyFile:2</code> <code>MSSQL:MyDatabase:COM.MyFileGroup:2.MyFile:2</code> <code>MSSQL\$MyInstance:MyDatabase:COM.MyFileGroup:2.MyFile:2</code>

Table 20. Command Syntax for Names Containing Periods, Backslashes, and Colons

The SQL filegroup named <i>My/FileGroup.2</i> for the SQL database named <i>My:Database.COM</i>	<code>My:Database\ .COM.My\\FileGroup\ .2</code> <code>MSSQL: My:Database \.COM.My\\FileGroup\ .2</code> <code>MSSQL\$MyInstance: My:Database\ .COM.My\\FileGroup\ .2</code>

A

Glossary

This glossary contains terms and definitions found in this manual. Most of the terms are specific to NetWorker Module for Microsoft SQL Server products.

1-9	Intermediate backup levels. Each number represents a backup level. Lower levels back up more files.
Administrators group	Windows NT user group whose members have all the rights and abilities of users in other groups, plus the ability to create and manage all the users and groups in the domain. Only members of the Administrators group can modify Windows NT OS files, maintain the built-in groups, and grant additional rights to groups.
Application Specific Module (ASM)	A program that, when used in a directive, specifies the way a set of files or directories is to be backed up and recovered.
ASM	See Application Specific Module (ASM) .
Backup Operators group	A Windows NT group whose members can log on to a domain from a workstation or a server, back it up, and restore the data. Backup Operators also can shut down servers or workstations.
backup group	A group of one or more NetWorker clients that start their backups at the same time.
backup level	The level or amount of data the NetWorker Module should provide in the backup. Supported backup levels include: full, incremental, and differential (level 1).

backup type	The type of backup to perform using the NetWorker Module. Backup types include: scheduled and unscheduled (manual).
backup volume	Backup media, such as magnetic tape or optical disk.
bootstrap	Information that includes the server index, media index, and configuration files needed for recovering the NetWorker server after a disk crash.
browse policy	A NetWorker policy that determines how long entries are maintained in the online index, and thus restorable.
client	A computer that accesses the NetWorker server to back up or recover files. Clients may be computers or file servers.
client file index	A database of information the NetWorker server maintains that tracks every file or filesystem backed up.
cluster	1. A group of independent network sectors that operate and appear to clients as if they were a single unit. 2. A group of disk sectors. The operating system assigns a unique number to each cluster and then keeps track of files according to which clusters they use.
command line	The shell prompt, where commands are entered.
compression	Compress the contents of a database before backup to generate less network traffic and use less backup media space. Compression can increase the backup time.
copy restore	Create a <i>copy</i> of a database by restoring a SQL Server 7.0 or later database backup to a new location or to a new database name.
device	The backup device (tape drive, optical drive, or autochanger) connected to the NetWorker server; used for backing up and recovering client files.
default instance	In a Microsoft SQL Server 2000 multiple instance configuration, the first installation of SQL Server on a computer is called the <i>default instance</i> . The name of the default instance is the network name for the local computer.

differential	A backup level that corresponds to a NetWorker Module level (1-9) backup. All of the pages in a database that were modified after the last database backup are saved.
directed recover	The <i>copy restore</i> type replaces the directed recover operation, which existed in versions of the NetWorker Module earlier than release 3.0.
directive	An instruction directing the NetWorker server to take special actions on a given set of files.
enabler codes	Special codes provided by Legato that allow you to run your NetWorker software product.
fake objects	Data items contained in the storage hierarchy that are not available for backup. To be able to browse the filegroups and files contained in a database, the NetWorker User for SQL Server program may display these fake objects.
file index	See client file index .
fileserver	A computer with disks that provides services to other computers on the network.
filesystem	1. A file tree on a specific disk partition or other mount point. 2. The entire set of all files. 3. A method of storing files.
full (f)	A backup level that corresponds to a Microsoft SQL Server database backup. The entire database is saved, which includes both data pages and transaction log pages.
group	A client or group of clients that starts backing up its files at a designated time.
heterogeneous networks	Networks with systems of different platforms that interact meaningfully across the network.
incremental (i)	A backup level that corresponds to a Microsoft SQL Server transaction log backup. Only the log file is saved.

instance	A copy of SQL Server running on a computer. A computer can run multiple instances of SQL Server 2000. A computer can run only one instance of SQL Server version 7.0 or earlier, although in some cases it may also be running multiple instances of SQL Server 2000.
interoperability	The ability of software and hardware on multiple computers from multiple vendors to communicate meaningfully.
level (1-9)	A backup level that corresponds to a Microsoft SQL Server differential backup. All of the pages in a database that were modified after the last database backup are saved.
log mark	See named log marks .
master database	The SQL Server master database contains information about all SQL Server databases on the SQL Server host.
media index	A database of information the NetWorker server maintains that tracks every backup volume.
media management daemon (MMD)	A media management service, also known as nsrmmd , that provides device support, including labeling and multiplexing writes (interleaved data) to media during backups, demultiplexing reads (unweaving the data) from media during recovery, and generating requests for volume mounting and unmounting when appropriate. An MMD service is created for each enabled device known to the NetWorker Server. The MMD service runs on the computer where the device is located.
MMD	See media management daemon (MMD) .
multiple instance	Microsoft SQL Server 2000 feature that allows multiple copies of SQL Server to run on a single computer.
named instance	An installation of SQL Server 2000 that is given a name to differentiate it from other named instances and from the default instance on the same computer. A named instance is identified by the computer name and instance name.

named log marks	Named log marks are created by database applications when transactions are performed. The marks enable access to specific transaction points in a database transaction log backup.
NetWorker	A Legato network-based software product for backing up and recovering filesystems.
NetWorker client	A computer that can access the backup and recover services from a NetWorker server.
NetWorker server	The computer on a network running the NetWorker software, containing the online indexes, and providing backup and recovery services to the clients on the same network.
no recovery	Equivalent to the SQL Server NORECOVER option, which places a database in an unloadable state after a restore, but enables the database to process additional transaction log restore operations.
notice	A response to a NetWorker event.
nsrhost	The logical <i>hostname</i> of the computer that is the NetWorker server.
nsrmmmd	See media management daemon (MMD) .
online indexes	The databases located on the server that contain all the information pertaining to the client backups and backup volumes.
operator	The person who monitors the server status, loads backup volumes into the server devices, and otherwise executes day-to-day tasks using the NetWorker software.
override	A backup level that takes place instead of the scheduled one.

partial restore	Only restore a <i>portion</i> of the filegroups and/or files associated with a single SQL Server 2000 database backup (Microsoft SQL Server 2000 only). When a partial database restore is performed, the PRIMARY filegroup and associated files are always restored. The PRIMARY filegroup contains information necessary for restoring the database to the proper structure.
pathname	Instructions for accessing a file. An <i>absolute pathname</i> indicates how to find a file beginning at the root directory and working down the directory tree. A <i>relative pathname</i> indicates how to find a file from the current location.
point-in-time	Restore SQL Server data to a specific point in time, such as a named log mark or transaction time within a backup version.
preconfigured	Existing selections or configurations for different NetWorker features.
PRIMARY	The SQL Server PRIMARY filegroup contains information necessary for restoring a database to the proper structure. When a partial database restore is performed, the PRIMARY filegroup and associated files are always restored.
promotion	When the server performs a backup at a higher level than originally requested. For example, the server performs a level full backup when a level differential backup was requested.
recover	The NetWorker command used to browse the server index and to recover files from a backup volume to a client's disk. See restore .
recovery model	The Microsoft SQL Server 2000 recovery model represents the trade-offs made when deciding which databases to back up and how often, and the impact back up and restore time have on system performance. Recovery models include: full, bulk_logged, and simple.

Registry	A database of configuration information central to Windows NT operations. The overall effect centralizes all Windows NT settings and provides security and control over system, security, and user account settings.
relocation list	The relocation list is used during a normal or copy restore type operation to specify where to locate the restored files. The list is composed of pairs of logical database filenames and fully qualified database filename relocation paths.
relocation path	See relocation list .
restore	The NetWorker Module command used to browse the SQL Server storage hierarchy and to restore files from a backup version. See recover .
restore mode	The restore mode instructs the NetWorker Module on how to interact with a database after a restore operation has completed. Restore modes correspond to SQL Server database restore options and include: normal, no recovery, and standby.
restore time	The restore time controls which backup data should be reinstated when a database is restored; may also control which portions of a level incremental backup are to be restored, when you inform the NetWorker Module to discard transactions performed after a given time.
restore type	The restore type is based on the level and type of backup created, as well as the set of data needed to restore from a backup. The restore type must be specified prior to browsing and selecting objects for restoration. Restore types include normal, copy, partial, and verify only.
retention policy	A NetWorker policy that determines how long entries are retained in the media database, and thus restorable.
root item	The top level data object in a SQL Server storage hierarchy.
save	The NetWorker command that backs up client files to backup volumes and makes data entries in the online index.

save set	A set of files or a filesystem backed up onto backup media using the NetWorker software.
save set ID	An internal identification number assigned to a save set by the NetWorker software.
server	The computer on a network running the NetWorker software, containing the online indexes and providing backup and recover services to the clients on a network.
stand-alone device	A backup device that contains a single drive for backing up data. Stand-alone devices cannot store or automatically load backup volumes.
standby	Equivalent to the SQL Server STANDBY option, which forces a database to be in a read-only state between transaction log restore operations.
stripes	One or more streams of data that may be extracted, in parallel, from a database, and written in parallel to multiple media devices, such as tape drives.
system administrator	The person normally responsible for installing, configuring, and maintaining the NetWorker server and software.
transaction log	A SQL Server transaction log contains named transactions or listings of changed files of a SQL Server database. Transaction logs can be truncated prior to a full database backup or backed up separately by performing an incremental level backup.
truncate	Equivalent to the SQL Server TRUNCATE_ONLY option, which causes the transaction log files to be truncated before creating a backup.
user	A person who can use the NetWorker software from his or her computer to back up and recover files.
verify-only restore	Only verifies the backup media for the selected SQL Server 7.0 or later backups. <i>Selecting the verify-only restore type will NOT restore the SQL Server data.</i> In addition, when verify-only is specified, item-level properties for database, filegroup, and file objects are not available.

virtual server	In a Microsoft Cluster Server configuration, SQL Servers appear as a set of two nodes and virtual servers. Each node is a physical computer with its own IP address and network name, and the virtual servers have their own IP addresses and network names. Each virtual server also owns a subset of shared cluster disks and is responsible for starting cluster applications that can fail over from one cluster node to another.
volume	Backup media, such as magnetic tape or optical disk.
volume ID	The internal identification assigned to a backup volume by the NetWorker software.
volume name	The name assigned to a backup volume when it is labeled.
volume pool	A feature that allows you to sort backup data to selected volumes. A volume pool contains a collection of backup volumes to which specific data has been backed up.
xlog	See transaction log .

A

ad hoc backup (*See* on-demand backup)
application log 118
Autostart 40
Autostart option 40, 46

B

backup
 canceling 75
 client 45
 clone 16
 combining SQL Server data objects 33
 creating 72
 database properties, specifying 79
 e-mail notification 47
 functionality supported for SQL Server
 2000 50
 group 40–41
 interleaving 124
 levels
 comparison 36
 defined 29
 described 35
 specifying 31
 SQL Server terminology 30
 supported 31
 multiple instances of SQL 125
 on demand 26
 operation status 47
 options 75
 advanced 78
 creating striped backup 77, 111
 general 76
 RAW save set 78
 parallelism 112
 performance 16, 111

 preparing for 71
 procedures, where to initiate 32
 restrictions 158
 scheduled 37
 scheduled events during 21
 schedules 36
 scheduling 35
 strategies 35–36
 striping 113
 time 74
 types 26
 types supported for SQL Server 2000 50
 unscheduled 80
 using NetWorker User for SQL
 Server 73
 viewing results 47
 volume 75
backup groups 41–42
backup level 42
backups
 ad hoc 26
 manual 26
 on demand 26
 options 40
 scheduled 16
 Autostart option 40, 46
 unscheduled 16
binary disk crash 137
bootstrap 41
browse policies 44–45
browse time, changing 105
BusinesSuite Module
 features highlighted 16

C

canceling a backup 75
canceling a restore 107
client 17, 19, 42

- resource 110–111, 113
- client file index 19, 44–45
- client server model 19
- cluster virtual server 130
- clustered SQL Server
 - configuring scheduled backups of 128
 - performing unscheduled operations on 129
- create time 57
- customize
 - backup schedule 35

D

- data 15
- data pools
 - example 82
- database and log files 107
- database consistency check (DBCC) 109
- database file relocation restrictions 94
- DBCC (*See* database consistency check)
- default instance 125
- disaster recovery
 - basic instructions 135
 - features 134
 - SQL Server 142
- display conventions 65

E

- e-mail
 - "savegroup completion" notification 47
 - notification(s) 47
- errors, diagnosing 118

F

- failover cluster support 128
- fake objects 65
- Figures
 - Backup Status Window 75
- figures
 - Backup Browse Window 74

- Backup Options Advanced Tab 78
- Backup Options General Tab 76
- Change Browse Time Dialog Box 105
- Database Item Restore Time Properties 95
- Fake Objects in the Backup Browse Window 64
- Fake Objects in the Restore Browse Window 65
- File Destination Dialog Box 92
- General Properties Dialog Box 101
- Main Window in NetWorker User for SQL Server 72
- Marked Items 67
- Master Backup Properties 80
- NetWorker Client Server Interaction
 - During a Backup Initiated by NetWorker Module 22
- NetWorker Client Server Interaction
 - During a Restoration Initiated by NetWorker Module 23
- NetWorker Client/Server Interaction
 - During a Backup Initiated by the NetWorker Module 22
- NetWorker Client/Server Interaction
 - During a Restore Initiated by NetWorker Module 23
- NetWorker User for SQL Server 72
- Partially Marked Items 68
- Point-in-Time Restore Dialog Box 98
- Read File Configuration Dialog Box 94
- Reading Dialog Box 94
- Restore Browse Window 104
- Restore Type Selection Dialog Box 103
- Select SQL Server Instance 72
- Specify Browse Time Dialog Box 105
- SQL Server Host 72
- SQL Server Selection Dialog Box 103
- Standby Undo File Selection Dialog Box 87
- Unmarked Items 67
- filegroups
 - backing up and restoring 102

restore properties, specifying 100
 specifying for restore 88
 files
 backing up and restoring 102
 restore properties, specifying 100

G

global browse time 57

H

host, adding to NetWorker server client
 list 45

I

interleaving 124

L

label template 41
 legacy database options
 select into/bulk copy 52
 trunc.log.on.chkpt 52
 log filenames 118
 log pool
 example 82
 logging
 scheduled backup to file 46

M

marking items 66
 capabilities 69
 indicators 67
 Mark All Databases option 69
 partially 68
 restrictions 66, 68
 semantics 68
 SQL Server root 69
 media

database 19, 23, 44–45
 message log files 47
 Microsoft Cluster Server (MSCS) 16
 active/active cluster
 configurations 130
 active/passive cluster
 configurations 130
 how NetWorker works with 127
 virtual server 130
 MSCS (*See* Microsoft Cluster Server)
 multiple instance 128
 multiple nonclustered instance 125

N

named instances 125–126, 128, 158
 named log mark 95
 defined 58
 option 97
 specifying restore time properties 95
 NetWorker
 client 17, 42
 client resource 110–111, 113
 client, description of 45
 features highlighted 16
 NetWorker Remote Exec Service 83
 NetWorker Module
 backup concepts 25
 failover cluster support 128
 how it works with SQL Server data 17
 interaction with SQL Server services 60
 support for MSCS clusters 128
 NetWorker Module database restoration 59
 NetWorker Module, interaction with 60
 NetWorker User for SQL Server 63, 80
 NORECOVERY option 59
 notifications 21
 nsrd
 functionality 21–22
 nsrlog command 46
 nsrmmd
 functionality during restore 23
 nsrmmdbd 23

nsrsqlrc 22–23, 125, 129
nsrsqlsv 21, 125

O

on-demand backup 26, 72–73
online indexes 44
options
 Autostart 40
 NORECOVERY 59
 RAW save set 79
 STANDBY 56
 Unmark All Databases 69

P

parallelism 112
partially marking items 68
planning
 using backup levels 35
point-in-time backup 96
primary data disk crash 136

R

RAW save set
 backup
 from command line save
 program 124
 from NetWorker User for SQL
 Server 123
 scheduled 123
 notifications 124
 options 79
 overview 119
 recommended configuration 123
 restore 124
 restrictions 78
 supported configurations 122
 supported NetWorker server
 features 120
 when to use 79

Read File Configuration properties,
 specifying 93
Rebuild Master utility 136
recovery model
 bulk_logged 50–51
 changing 53–54
 full 50
 simple 52
 specifying 52
restore 22
 backup levels 36
 canceling 107
 database objects 83
 file destination 92
 functionality supported for SQL Server
 2000 50
 how to 102
 multiple instances of SQL 125
 named log mark 95
 overwriting existing database 89
 point-in-time 16, 57
 properties 84, 106
 restrictions 158
 scheduled events during 22
 server data 49
 SQL cluster resources 60
 time 95, 107
 transaction log backups 59
 uncommitted transactions 59
 using NetWorker User for SQL
 Server 102
restore concepts 49
restore mode
 no recovery 49, 56
 normal 49, 56
 standby 49, 56
restore time 57
 backup versions 49
 database backup versions 57
 log marks 49
 named log marks 58
 point-in-time 49, 57
 specifying 59

restore type 83
 copy 56
 normal 55
 partial 55
 verify only 55

restore-type
 copy 103–104
 normal 83
 partial 104
 verify only 83

retention policies 44–46

S

save

 functionality 17

save sets 42, 46

savegroup completion report 21

savegrp 21

 functionality 21

scheduled backups 21

schedules

 levels 35

select into/bulk copy 52

SmartMedia 16

SQL Server

 binary disk crash 137

 cluster

 command line operations 129

 operations from NetWorker User for
 SQL Server 129

 database consistency check
 (DBCC) 109

 disaster recovery 142

 Enterprise Manager 60

 host, adding to NetWorker server client
 list 45

 master database maintenance 60

 named instances 125–126

 primary data disk crash 136

 reinstall, when to 136

 root, marking 69

 services

 multi-user mode 60

 NetWorker Module interaction 60

 single-user mode 60–61

 virtual server 130

SQL Server 2000 16

 enhancements for backup and
 restore 30

 multiple nonclustered instance 125

 named log mark 95

 recovery models 50

 running multiple nonclustered
 instances 125

SQL Server master database

 maintenance 60

SQL Server NORECOVERY option 59

STANDBY option 56

Standby Undo File Selection properties,
 specifying 87

storage management 15–16, 21

stripes

 defined 111

 option 77

 performance considerations 112

 performing striped backup 112–113

 performing striped recovery 113

 fail-safe 115

 Microsoft SQL Server 2000 116

 Microsoft SQL Server 7.0 115

 optimal 114

 optimal versus fail-safe 114

 specifying a value 112

syntax

 command arguments for SQL Server
 data 157

 database, filegroup, and filename
 containing periods 159

T

tables

 Backup Levels Advantages and
 Disadvantages 36

 Backup Levels for SQL Server Data 31

- Command Argument Syntax for SQL Server Data 157
- Command Syntax for Names Containing a Period Character 159
- Creating Additional Backup Levels with Data Objects 33
- Full Backup Every One to Two Weeks 36
- Guidelines for Fail-Safe Striped Recovery 114
- Incremental (Transaction Log) Backup After a Full Backup 37
- NetWorker Module Backup Levels 30
- RAW Save Set Supported NetWorker Server Features 120
- SQL Server Storage Hierarchy Display Conventions 65
- Supported Backup and Restore Functionality for SQL Server 2000 Recovery Models 50
- Supported Backup Types for SQL Server 2000 Recovery Models 50
- Where to Initiate Backup Operations 32
- transaction log
 - backup 57
 - maintenance 37
 - overflow prevention 38
- trunc.log.on.chkpt 52

U

- unmarking items 66
 - indicators 67
- Unmark All Databases option 69

V

- VDI (*See* Virtual Device Interface)
- Virtual Device Interface 23
- virtual server 128
- volume pools 41, 75, 80
 - data pool 80
 - destination 42

- selection criteria 42
- volumes
 - backup levels 35

X

- XBSA (*See* X-Open® Backup Services Application Programming Interface)
- X-Open® Backup Services Application Programming Interface 21