

Legato NetWorker[®]
Module for Microsoft[®] SQL Server[™]

Administrator's Guide

Release 3.0
Windows NT[®] 4.0 and Windows[®] 2000 Version



© 2001, Legato Systems, Inc. All rights reserved.

Legato, the Legato logo, Legato NetWorker, ClientPak, Octopus, and SmartMedia are registered trademarks; and Celestra, Co-StandbyServer, eCluster, GEMS, SnapShotServer, 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

January 2001

22-1915-03

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
Customer Service.....	14
Customer Feedback	14
Chapter 1: Introduction.....	15
The Importance of Backing Up Microsoft SQL Server Data	15
Using Legato NetWorker to Protect SQL Server Data	15
Using NetWorker Module to Protect SQL Server Data	16
How NetWorker Module for Microsoft SQL Server Connects to NetWorker	17
Client-Server Model	17
Backing Up SQL Server Data From a NetWorker Server Client	17
Restoring SQL Server Data From a NetWorker Server Client	18
Chapter 2: NetWorker Module Backup Concepts	21
Backup Types.....	22
Scheduled Backups.....	22
Unscheduled Backups.....	23
Filegroup Backup Performance Considerations.....	24
Reduced Number of Files Needed to Restore a Damaged Disk	24
Decreased Backup Time as Compared to Full Database Backup	25

Command Line Notation and Syntax	25
Syntax for a Multiple Instance Configuration for SQL 2000.....	26
Use of MSSQL:.....	26
Backup Levels.....	27
Microsoft SQL Server 2000 Enhancements	28
Specifying Backup Levels.....	30
Choosing Backup Levels	31
Backup Level Considerations and Strategies.....	31
Backup Schedules	32
Transaction Log Maintenance.....	33
How to Prevent Log Overflow.....	34
Backup Groups	35
Backup Group Considerations.....	36
Volume Pools	36
How to Create a Volume Pool	37
Example: Specify a Volume Pool Destination	37
How to Customize a Volume Pool	38
Example: Create a Volume Pool for Each Major Backup Level.....	39
Browse and Retention Policies	40
NetWorker Backup Clients.....	41
How to Add a SQL Server Host to the NetWorker Server's Client List.....	41
Backup Operation Status	42
Chapter 3: NetWorker Module Restore Concepts.....	45
Microsoft SQL Server 2000 Recovery Models	46
FULL Recovery Model.....	47
BULK_LOGGED Recovery Model.....	47
SIMPLE Recovery Model	47
Specifying Database Recovery Models	48
Changing the Recovery Model for a Database.....	48
How to Change from FULL or BULK_LOGGED to SIMPLE	49

How to Change from BULK_LOGGED to FULL	49
How to Change from SIMPLE to FULL or BULK_LOGGED.....	49
Restore-types	50
Normal Restore-type	50
Partial Restore-type	51
Verify Only Restore-type	51
Copy Restore-type	51
Restore Modes	51
Normal Restore Mode	52
No Recovery Restore Mode.....	52
Standby Restore Mode	52
Restore Time	52
Database Backup Versions.....	53
Point-in-time Restore of a Transaction Log Backup.....	53
Microsoft SQL Server 2000 Named Log Marks	54
Specifying the Restore Time	54
NetWorker Module Database Restoration.....	55
SQL Server Master Database Maintenance	55
NetWorker Module Interaction with SQL Server Services	56
Chapter 4: NetWorker Module Usage and Concepts	57
Browsing Data Items in NetWorker User for SQL Server	57
Restore Browse Window Restrictions	58
Fake Objects	58
Display Conventions	59
Marking Items in NetWorker User for SQL Server	60
Restore Browse Window Restrictions	61
Marking Semantics and Restrictions.....	61
Chapter 5: Backing Up SQL Server Data	63
Preparing for the Backup.....	63

Creating the Backup.....	64
How to Back Up Using NetWorker User for SQL Server	65
Backup Options	66
Database Item Backup Properties	69
Target Volume Pools for Unscheduled Backups.....	70
Assigning Pools in NetWorker User for SQL Server	70
Example:	71
Specifying Matching Pools and Label Templates in NetWorker Administrator.....	71
Target Pool Specification Guidelines.....	72
Chapter 6: Restoring SQL Server Data.....	73
Preparing for the Restore	73
Performing the Restore	74
How to Restore Using NetWorker User for SQL Server	74
Database Restore Properties	78
Database Item Restore General Properties	78
Standby Undo File Selection Dialog.....	80
Database Item Restore Files Properties and Partial Restore-Type Database Item Restore General Properties	81
File Destination Dialog	84
Read File Configuration Dialog	85
Database File Relocation Restrictions.....	86
Database Item Restore Time Properties.....	87
Point-in-Time Restore Dialog	89
FileGroup Restore Properties	90
How to Ensure Proper Filegroup Restoration	92
File Restore Properties	92
Chapter 7: Special Features.....	95
Checking Database Consistency	95
Using Striping.....	96

Configuring Striping for Backups.....	97
Performance Considerations for Striping.....	97
How to Perform a Striped Backup	98
Using the RAW Save Set	99
Using Interleaving.....	99
Backing Up and Restoring Multiple Instances of SQL Server	100
Backing Up and Restoring Microsoft Cluster Server Configurations	101
How NetWorker Works with MSCS Clusters.....	101
NetWorker Module 3.0 Support for MSCS Clusters.....	102
Named Instances in Failover Cluster Configurations	102
Configuring Scheduled Backups of a Clustered SQL Server.....	103
Performing Unscheduled Operations on a Clustered SQL Server.....	103
Operations from the Command-Line	103
Operations from NetWorker User for SQL Server	104
Active/Passive Cluster Configurations.....	104
Active/Active Cluster Configurations	104
NetWorker Module Interaction with SQL Server Services	105
Chapter 8: Disaster Recovery	107
When to Reinstall Microsoft SQL Server	108
Disaster Recovery of Microsoft SQL Server Data Disk Crash or Data Corruption	109
Disaster Recovery of Microsoft SQL Server Disk Crash or Data Corruption.....	110
Disaster Recovery of Microsoft SQL Server and NetWorker Server	111
Disaster Recovery of Microsoft SQL Server (Shortcut)	112
Disaster Recovery of Microsoft SQL Server	113
Appendix A: Commands and Command Options	115
Using the nsrsqlsv, nsrsqlrc, and nwmssql Commands	115
Using the nsrsqlrc Command	116
nsrsqlrc Command Options.....	116
Using the nsrsqlsv Command.....	121

nsrsqlsv Command Options	121
Using the nwmssql Command	125
nwmssql Command Options.....	125
Glossary.....	127
Index.....	135

Preface



The *Legato NetWorker Module for Microsoft SQL Server Administrator's Guide* contains information on how to configure and manage the 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 3.0 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 SQL Server version 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 Legato NetWorker supports these technologies. In addition, users 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: restore 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 restoration of your data	 Important: You must install the NetWorker Module software in the same directory where you installed the NetWorker client binaries.

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 might 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^a	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 Inside Sales, (408) 530-3000 or <i>sales@legato.com</i>			Yes	
Legato Education Services, (408) 530-3099 or <i>training@legato.com</i>				Yes

a. For detailed information about our services, support policies, and software subscriptions, refer to the *Legato Customer Service and Support Guide* included in the media kit.

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.

Customer Service

Legato provides a variety of ways for you to contact Customer Service. To register your Legato products, go to the Legato licensing web site. To ask questions about licensing, to transfer licenses, or to learn the status of your order, contact Legato using one of the following methods.

Customer Service	Contact
Legato licensing web site	http://license.legato.com
Telephone number	(650) 812-6063 (option 3)
Fax number	(650) 812-6220
E-mail	service@legato.com
E-mail for order status	order_admin@legato.com

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 Technical Documentation page, and then look for the link to the Legato Documentation Survey.

Chapter 1: Introduction

The NetWorker Module 3.0 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 NT[®], Microsoft Windows[®] 2000, and UNIX operating systems.

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 you restore your 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 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 Legato NetWorker for storage management simplifies the process and reduces administrative burden by automating and centralizing your data storage operations.

1

With NetWorker, you can:

- Perform automated lights-out” backups during off-peak hours
- Use centralized administration to configure, monitor, and control backups from anywhere on a network
- Automate tape handling tasks using 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 NetWorker Module to Protect SQL Server Data

Entering “All” in the Save Set text box in NetWorker Administrator’s Edit Client dialog specifies an entire filesystem backup on the corresponding client host. However, a filesystem backup does not save your SQL Server data in a form that is restorable by Microsoft SQL Server. For a restorable save of SQL Server data, you must back up the data using one of the methods described in this 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 your SQL Server storage hierarchy:

- Scheduled automated backups
- Unscheduled backups that can be initiated any time
- Automatic back up 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 NetWorker Module for Microsoft SQL Server Connects to NetWorker

The SQL Server software can reside on the same system as the NetWorker server software or on a separate system. The NetWorker Module save program **nsrsqlsv** substitutes for the NetWorker server client **save** program during a backup of SQL Server data, NetWorker Module for Microsoft SQL Server must be installed on the system where SQL Server and the NetWorker client were installed. No matter what computer is the SQL Server host, that computer must be configured as a storage management client of the NetWorker server. For details about NetWorker and NetWorker Module configuration topology, refer to the *Legato NetWorker Module 3.0 for Microsoft SQL Server Installation Guide*.

Client-Server Model

NetWorker 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. You can use the NetWorker server program on Windows NT, Windows 2000, or UNIX.

The NetWorker server allows you to perform the following functions:

- Maintain the resource information
- Contact clients listed in a backup group configured on the server
- Perform manual backups when a client request is received
- Restore data upon request from the client

The online client file index and media database are maintained by the NetWorker server.

Backing Up SQL Server Data From a NetWorker Server Client


When a scheduled backup for a SQL Server client is triggered by **nsrd** on the NetWorker server, **savegrp** executes the NetWorker Module backup save. The **nsrsqlsv** program sends the backup data received from SQL Server to NetWorker through an X-Open® Backup Services Application Programming Interface (XBSA). The final results of the **savegrp** execution are sent back to the server and are included in the savegroup completion report, which is sent according to the notification set up in NetWorker Administrator for the Savegroup event.

1

NetWorker Server performs the scheduling storage management tasks, while the NetWorker Module passes the data from the SQL Server to NetWorker Server. [Figure 1 on page 18](#) shows the functional relationship between NetWorker, NetWorker Module for Microsoft SQL Server, and Microsoft SQL Server during a scheduled backup.

For background information about the NetWorker services and operations mentioned and illustrated in this chapter, see the *Legato NetWorker Administrator's Guide*.

Figure 1. NetWorker Client-Server Interaction During a Backup Initiated by NetWorker Module



Restoring SQL Server Data From a NetWorker Server Client


When a restore is initiated by **nsrsqlrc**, the NetWorker Module version of the NetWorker server client recover program, the NetWorker XBSA API translates the object names requested by the NetWorker Module into a format

understood by NetWorker and forwards them to the NetWorker server's **nsrd** service. The media service, **nsrmmd**, invokes **nsrmmdbd** to search the NetWorker server's online media database for the media containing the object(s) requested.

After the media is mounted, **nsrmmd** sends the data through the NetWorker XBSA API to **nsrsqlrc**, which restores the data to SQL Server through the LOAD command and the appropriate API (Named Pipes or Virtual Device).

Figure 2 shows how data moves from the NetWorker server to SQL Server during a NetWorker Module restore.

Figure 2. NetWorker Client-Server Interaction During a Restoration Initiated by NetWorker Module



1

Chapter 2: NetWorker Module Backup Concepts

Prior to configuring the NetWorker Administrator to work with the NetWorker Module, you should understand the differences between the different types and levels of backups that you can create, and how the type of backup created will affect your options for restoring data. Specifically, this chapter explains the following:

- Backup types: *scheduled* (automatic) and *unscheduled* (manual)
- Backup levels: full, incremental, and differential
- Client resources: backup groups, browse and retention policies
- SQL Server transaction log maintenance
- Backup schedules, volume pools, and performance considerations
- Backup status and log files



Important: It is *highly recommended* that you read this chapter and [“Chapter 3: NetWorker Module Restore Concepts” on page 45](#), prior to performing a backup or restore operation using Legato NetWorker Module. The type and level of backup you create, 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.

2

Backup Types

Using NetWorker server with Legato NetWorker Module, you can save your SQL Server data by performing one of the following types of backup:

- **Scheduled backups** — Scheduled backups are configured in NetWorker Administrator so that they can be started automatically from the NetWorker server. Scheduled backups can be configured to run at any time, using any backup level.
- **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 at any time by logging on to the NetWorker client host using an account that has SQL Server administrative privileges, and entering the **nsrsqlsv** command (with desired command options) at a command prompt. A manual backup from a command prompt can be run at the full, differential, or incremental level.
- From NetWorker User for SQL Server, the NetWorker Module graphical interface. From this interface you can run only level-full backups, as well as restore procedures.

When you run a scheduled backup from NetWorker Administrator, you can back up SQL Server data residing on any host that is a client of the NetWorker server. When you run an unscheduled backup from a NetWorker server client host, you can only back up data that is stored on that SQL Server host.



Important: It is strongly recommended that you perform regular, *scheduled* NetWorker backups of SQL Server to provide maximum protection for your data. For more information about scheduled backups, see “[Scheduled Backups](#)”).

Scheduled Backups

Only regular, *scheduled* backups ensure that *all* of your 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 SQL Server in the event of a disaster. When you perform an *unscheduled* backup of

SQL Server, the NetWorker client indexes for SQL Server and the NetWorker server's bootstrap file are not automatically backed up. Performing regular, *scheduled* NetWorker backups of SQL Server provides maximum assurance that all of your SQL Server data will be available for restoration in case of a disaster.

To customize a scheduled backup of Microsoft SQL Server data, you must configure certain resources in NetWorker Administrator, the graphical interface for NetWorker server administration:

- Backup Groups, see [“Backup Groups” on page 35](#)
- Volume Pools, see [“Volume Pools” on page 36](#)
- Backup Schedules, see [“Backup Schedules” on page 32](#)
- Browse and Retention Policies, see [“Browse and Retention Policies” on page 40](#)
- Add the SQL Server host to the NetWorker server's client list
 - If you do not create custom resource settings in NetWorker Administrator before defining your SQL Server host as a NetWorker client, NetWorker uses its default resource settings when backing up your 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, NetWorker Module supports the selection of 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 includes the database, filegroup, file, and transaction log.

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

2



Important: Because NetWorker Module can be used to manually back up specific 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 your database(s) is possible. You can *also* restore a SQL Server 7.0 or 2000 filegroup or file from a *full database* backup, if the backup was created using NetWorker Module 3.0.

When you perform an *unscheduled* backup of SQL Server, the client indexes and bootstrap file are not automatically backed up. Following a manual backup, to create client indexes and a bootstrap file, you must run the command **savegrp -O savegroup_name** from the command line on the NetWorker server.

If you never perform *scheduled* backups of SQL Server with NetWorker and NetWorker Module, the client indexes and bootstrap file are never backed up automatically. It is strongly recommended that you perform regular, *scheduled* NetWorker backups of SQL Server to provide maximum protection for your data. For more information about scheduled backups, see [“Scheduled Backups” on page 22.](#))

Filegroup Backup Performance Considerations

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 whole database. This is because, after a filegroup is restored, NetWorker Module must automatically also restore *all* transaction log backups performed since the filegroup backup. (This is a requirement imposed by Microsoft SQL Server, as mentioned in the previous section.) 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 backups.)

Reduced Number of Files Needed to Restore a Damaged Disk

SQL Server allows you to store a database in multiple filegroups and files in order to spread the database over more than one disk or disk partition. NetWorker Module allows you to back up and restore distinct filegroups and

files without having to back up entire filegroups or databases. If a disk is damaged, you might need to restore only those files on the damaged disk or disk partition.

Decreased Backup Time as Compared to 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.

2

Command Line Notation and Syntax

At least one SQL Server data item (database, filegroup, or file) must be specified for manual (unscheduled) backup or restore. More than one data object may be specified and different types of data may be combined. SQL data objects must be specified using the following syntax guidelines:

Table 1. Command Argument Syntax for SQL Server Data

SQL Server Data	Syntax for SQL Server Data Objects
All database(s) in the SQL Server storage hierarchy	<i>MSSQL:</i>
Specified database(s)	[<i>MSSQL:</i>]dbName [[<i>MSSQL:</i>]dbName [<i>MSSQL:</i>]dbName2 ...]
All filegroups in specified database(s)	[<i>MSSQL:</i>]dbName. [[<i>MSSQL:</i>]dbName. [<i>MSSQL:</i>]dbName2. ...]
Specified filegroup(s) in specified database	[<i>MSSQL:</i>]dbName.fgName [[<i>MSSQL:</i>]dbName.fgName [<i>MSSQL:</i>]dbName.fgName2 [<i>MSSQL:</i>]dbName2.fgName [<i>MSSQL:</i>]dbName2.fgName2 ...]
All files in specified filegroup(s) in specified database(s)	[<i>MSSQL:</i>]dbName.fgName. [[<i>MSSQL:</i>]dbName.fgName. [<i>MSSQL:</i>]dbName.fgName2. [<i>MSSQL:</i>]dbName2.fgName. [<i>MSSQL:</i>]dbName2.fgName2. ...]

Table 1. Command Argument Syntax for SQL Server Data (Continued)

SQL Server Data	Syntax for SQL Server Data Objects
Specified file(s) in specified filegroup(s) in specified database(s)	[MSSQL:]dbName . fgName . filename [[MSSQL:]dbName . fgName . filename [MSSQL:]dbName . fgName2 . filename [MSSQL:]dbName2 . fgName . filename [MSSQL:]dbName2 . fgName2 . filename ...]

2

Syntax for a Multiple Instance Configuration for SQL 2000

When your configuration contains multiple instances of SQL Server, the name of the instance should be specified before the data. The syntax is:

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:*

or, to restore several filegroups for *instanceTwo*, specify:

nsrsqlrc -s NetWorker_server_hostname
MSSQL\$instanceTwo: dbName . fgName
MSSQL\$instanceTwo: dbName . fgName2



Important: **nsrsqlsv** and **nsrsqlrc** only support specification of a single instance. If save sets for more than one instance are specified, the backup will fail. Restore supports mixing of instances for directed recovers; however, the same named instance can not be the source if it is not also the target. For more information about running multiple instances of SQL Server, see [“Backing Up and Restoring Multiple Instances of SQL Server” on page 100](#).

Use of MSSQL:

Specifying **MSSQL** before each individual data object name is optional and does not in any way affect the meaning of 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

```
nsrsqlsv -s NetWorker_server_hostname dbName.fgName
```

In a multiple instance configuration, **MSSQL\$** is required followed by the instance name and a colon. For example:

```
nsrsqlsv -s NetWorker_server_hostname MSSQL$instanceName:dbName.
```

2

Backup Levels

Legato NetWorker provides three main levels of backup: full, incremental, and differential, where differential is specified as any level from 1 to 9. (Another level that can be used is *skip*, which specifies that the savegroup operation is to be completed successfully, but no backup of the data is to be performed. For details about using *skip*, refer to the *Legato NetWorker Administrator's Guide*.)

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.
- Differential backup capability is new with SQL Server 7.0.) In NetWorker Module 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 database backup.

The terminology used in the Microsoft SQL Server product for the different backup levels varies from that used in NetWorker Module, as shown in [Table 2](#).

2

Table 2. NetWorker Module Backup Levels

Function	NetWorker 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 introduces enhanced support for backing up and restoring filegroups and files. In addition to creating a level full filegroup or file backup, SQL Server now supports the creation of filegroup differential and file differential backups. Performing a filegroup differential backup may actually reduce both media requirements and restore time; you can store the data 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: You can also restore a SQL Server 7.0 or 2000 filegroup or file from a *full database* backup, if the backup was created using NetWorker Module 3.0.

NetWorker Module supports the creation of 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 you have selected to backup and any SQL Server settings on those objects. The following table identifies the supported backup levels:

Table 3. Backup Levels for SQL Server Data

SQL Server Data Objects	SQL Server 7.0			SQL Server 2000		
	full	diff	incr	full	diff	incr
All database(s), including MSSQL :	yes	yes	no	yes	yes	no
Specified database(s)	yes	yes	yes ^a	yes	yes	yes ^a
All filegroups in specified database(s)	yes	no ^b	n / a	yes	yes ^a	n / a
Specified filegroup(s) in specified database	yes	no ^b	n / a	yes	yes ^a	n / a
All files in specified filegroup(s) in specified database(s)	yes	no ^b	n / a	yes	yes ^a	n / a
Specified file(s) in specified filegroup(s) in specified database(s)	yes	no ^b	n / a	yes	yes ^a	n / a

a. An incremental database, filegroup differential, or file differential backup can only be created when the “select into / bulkcopy” SQL Server database option is not set. 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 you specify a differential filegroup or file backup, NetWorker Module performs a full backup.

Specifying Backup Levels

Depending on the type of backup you are creating, the backup level can be specified in one of the following ways:

1. For scheduled backups, the backup level must be specified using NetWorker Administrator>Schedules resource>Set Level dialog. The backup level specified in this dialog box takes effect whenever that particular scheduled backup runs.
2. Unscheduled backups initiated from a Windows command prompt on the NetWorker server client host, must include the backup level as an option to the command. For example, entering the following at a command prompt would initiate an unscheduled incremental level backup:

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

For more information about **nsrsqlsv** and the backup level command options, see [“Appendix A: Commands and Command Options” on page 115.](#))

3. When an unscheduled backup is initiated from NetWorker User for SQL Server, the NetWorker Module graphical interface, only a backup level of full can be created. No interface is provided for specifying a different backup level.

[Table 4](#) summarizes where NetWorker Module backup procedures can be initiated and which backup levels are supported for each interface.

Table 4. Where to Initiate Backup Operations

Backup Type	Backup Initiated From	Backup Levels Available		
		full	incr	diff
Scheduled	NetWorker Administrator (server GUI)	Yes	Yes	Yes
Unscheduled	Command line on NetWorker server client host	Yes	Yes	Yes
	NetWorker User for SQL Server (client GUI)	Yes	No	No

Choosing Backup Levels

In addition to providing support for specifying the backup level to perform, NetWorker Module also allows you to select SQL Server data objects in various combinations to create backups of different levels:

- Full Database backup — Select one or more databases to create a level full database backup of the selected database s and their transaction log files.
- Full Filegroup or Full File backup — 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.

When performing an unscheduled level full backup of a filegroup or file, you should also perform a database incremental level backup to maintain the validity of the transaction log files.

- Database Incremental backup — Select one or more databases to create a database incremental level backup of *only* the transaction log files for the selected databases. Incremental backups are allowed when the SQL Server backup options Select Into / Bulk Copy and Truncate Log on Checkpoint are *not selected*.
- Database Differential (level 1) backup — Select one or more databases to create a database level differential backup of *only* all of the changes to the selected databases since the last full level backup was created.
- Filegroup or File Differential backup — Select one or more filegroups or one or more files to create a filegroup or file level differential backup of *only* all of the changes to the selected filegroups or files since the last full level backup was created.

Backup Level Considerations and Strategies

Because it may not be practical or efficient for you to run full backups every day, NetWorker Module with NetWorker allows you to specify other backup levels for automatic, scheduled backups. Limiting the frequency of full backups can decrease server load while nonetheless ensuring that your 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 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 will 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 transaction log backups can simplify and speed database restoration.

The following table outlines the basic differences between the backup levels:

Table 5. Backup Levels Advantages and Disadvantages

Backup Level	Advantages	Disadvantages
Full	• Fastest restore time	• Slow backup • Longest, highest load on server • High load on client and network • Uses the most volume space
Incremental	• Faster backup • Low load on server • Uses least volume space	• Slow restore • Data 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)

Backup Schedules

Legato NetWorker 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. You can customize a SQL Server backup schedule by specifying different backup levels on different days.

Choosing a backup level usually entails a trade off: the time and number of volumes required for backup versus that required for restoration. Several preconfigured backup schedules are provided, or you can create a customized schedule, which best suits your configuration requirements and resource capabilities. If your Microsoft SQL Server manages a significant amount of data, consider scheduling your backups like this:

- Schedule a full backup of each of your databases every one to two weeks.
- Schedule incremental (transaction log) backups on several successive days immediately following the previous full backup. This backs up all data that has changed since the previous incremental.

- Schedule a level 1 differential backup after several days of incremental backups. This backs up all data since the previous full backup.

Thus, a reasonable approach to backup scheduling would be similar to the following:

Table 6. Backup Scheduling Strategies

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

Another reasonable backup scheduling possibility would be something like this:

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

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

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

Transaction Log Maintenance

NetWorker Module 3.0 provides two methods for maintaining SQL Server database transaction logs:

- From NetWorker User for SQL Server, you can specify that the active portion of the transaction log should be backed up prior to restoring a database, filegroup, or file.

2

- From NetWorker User for SQL Server or from a command prompt, you can specify that the transaction log files should be truncated prior to backing up a database.

There are two other types of management, implicit and explicit:

- 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 are synonymous in terms of SQL action, resulting in the log being truncated prior to the backup. Both command options are compatible with level full and level 1 backups. If no **-T/-G** option is specified, following completion of the backup, if NetWorker Module determines that the database does not support log backups, it will tell SQL to truncate the transaction log.

How to Prevent Log Overflow

When running Windows, you can prevent your 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 percent full. An incremental backup truncates the logs and thus clears disk space.

To create this alert:

1. Create a batch file (called `sqlalert.bat` in this example) 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 you installed NetWorker Module for Microsoft SQL Server. 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 field to 10 minutes or more, enough time to run an incremental 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 field, enter the name and full path of your batch file, as in the following example:
 - e. `C:\win32app\nsr\bin\sqlalert.bat`
 - f. Select the database to monitor in the Instance field.
 - g. Click Every Time.
 - h. Click the Add button, then click the Done button.

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 115](#).

Backup Groups

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

NetWorker provides a preconfigured group named “Default.” You can modify the Default group’s attributes, but you cannot delete Default from the list of NetWorker groups. The Default group has the following attributes:

- Autostart = Disabled.
After you enable 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 = 0.
- Clones = No.

- Clone pool = Default Clone.

If you want to assign your SQL Server to another group with different attributes, you must create that group in NetWorker Administrator *before* you define the SQL Server host as a NetWorker client. For instructions on setting up a backup group, see 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. You determine which data are stored on each physical volume (for example, a magnetic tape) using the Groups field in the Pools resource in NetWorker Administrator. See [“Volume Pools” on page 36](#) for more information.

Backup Group Considerations

If you have a number of 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 your NetWorker server. You can set up any number of backup groups in NetWorker Administrator.

When you select a start time for each of your groups, 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, you must 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, see 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 your scheduled backups, you can direct related data to a particular pool according to criteria you specify in NetWorker Administrator.

Each NetWorker media volume belongs either to the preconfigured *Default* pool or to a pool you create. If you do not assign data to a specific pool of volumes, NetWorker automatically uses the Default pool and looks on the backup device for volumes with the pool label “Default.”

How to Create a Volume Pool

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

For scheduled backups only, NetWorker Module uses the pool-selection criteria you assign in NetWorker Administrator. (In NetWorker Module 3.0, any pool settings specified in the Backup Options dialog box in NetWorker User for SQL Server apply only to unscheduled backups. This was not the case in BusinessSuite Module 2.0.1 and earlier, in which pool settings in the Backup Options dialog box pertained to both scheduled and unscheduled backups.)

You can specify the pool destination for each scheduled backup according to the following criteria:

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



Example: Specify a Volume Pool Destination

This example demonstrates how to configure the pool-selection criteria to specify a pool destination.

In NetWorker Administrator:

1. Create a NetWorker pool named *SQLServerFullBackups* and select backup level **full** as a pool selection criterion.
2. Create a NetWorker pool named *SQLServerIncrBackups* and select backup level **incr** as a pool selection criterion.
3. Create a save group called *SQLGroup*.

2

4. Associate both new pools with the save group *SQLGroup*.
5. Create a NetWorker Module client named *mysqlserver*.
6. Add the NetWorker Module client *mysqlserver* to *SQLGroup*.
7. Configure the client *mysqlserver* to use the Default schedule, which performs a full backup on the first Sunday of every month and daily incremental backups in between fulls.
8. 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*, will be written to media volumes belonging to pool *SQLServerFullBackups*. When the daily incremental scheduled save of *SQLGroup* occurs, save sets for this group, including the incremental backup from *mysqlserver*, will be 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 NetWorker Administrator, the volume pool names must be specified in the NetWorker Administrator's Pools resource. A valid volume pool name is any valid NetWorker name of 1024 characters or less. At least one volume must be labeled for each pool.



Example: Create a Volume Pool for Each Major Backup Level

This example demonstrates how to configure volume pools and label templates in NetWorker Administrator to separate full, incremental, and differential (level 1) backups from one another.

1. To create a new pool, right-click the Pools icon in the server window in NetWorker Administrator and click Create. Follow the instructions in the *Legato NetWorker Administrator's Guide* and NetWorker online Help. Specify *SQLServerFullBackups*, for example, as the entry for Name.
 - a. As the choice for Groups, select a backup group that is unique for SQL Server.
 - b. Specify **full** as the selection for Levels.
 - c. In the Preferences tab of the Create Pool dialog, select **Yes** as the choice for Store Index Entries.
2. Save the Pool configuration.
3. A dialog box appears, asking you to select a label template or apply the configuration again to automatically create and select a label template. Click **OK** and reapply the configuration to automatically create a label template named *SQLServerFullBackups* and assign it to the *SQLServerFullBackups* pool configuration.
4. Repeat steps 1 through 3 in this procedure to create a pool for your Legato NetWorker Module incremental backups, selecting level *incr* instead of *full*. You might call this pool *SQLServerIncrBackups*, for example.
5. Repeat steps 1 through 3 again to create a pool for your Legato NetWorker Module differential backups, selecting level *1* instead of *full*. You might call this pool *SQLServerDiffBackups*.
6. Insert a new media volume in the backup device, or select another drive or slot with a volume loaded for labeling.
7. Select Operations>Device Operations>Label to label a media volume, selecting the label template customized for the *SQLServerFullBackups* pool. Repeat this step to label another volume for the *SQLServerIncrBackups* pool, selecting the label template customized for that pool. Repeat the step again to label a third volume for the *SQLServerDiffBackups* pool, selecting the label template customized for that pool.

You now have three different volume pools, one for each of the major backup levels. For more information about using NetWorker volume pools, see the *Legato NetWorker Administrator's Guide* and online Help.

2

Browse and Retention Policies

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

To manage the size of the online indexes, NetWorker 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 NetWorker in a format that is easily restorable.

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

You can also manage the indexes manually with the Indexes and Volumes windows in NetWorker Administrator. See the *Legato NetWorker Administrator's Guide* for a description of manual index management and index policy concepts.



Important: NetWorker 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, that 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 you want to back up must be configured as a NetWorker client using the Clients resource in NetWorker Administrator. In addition, you can configure multiple Microsoft SQL Server databases that exist on the same SQL Server host, as separate NetWorker client resources.

For each client resource, the NetWorker server performs 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, follow these steps:

1. In the Create Client dialog box in NetWorker Administrator, for the client name, 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. In the "Save set" field, specify any valid save set name(s). For example:

MSSQL:

or

MSSQL:dbname1 [MSSQL:dbname2 MSSQL:dbname3 ...]

2

(Entering only *MSSQL:*, for example, always yields a full backup of all databases on your SQL Server host.)

7. In the Remote tab, enter **nsrsqlsv** (with your desired command options) in the Backup command textbox. For information about **nsrsqlsv** options, see [“Appendix A: Commands and Command Options” on page 115](#).
8. In the Remote tab, add host names to the remote access field, to grant those hosts directed recover (restore-type copy) permission. This allows the named hosts to access the NetWorker server and receive directed recover data, if needed.
9. In the Preferences tab, enter all known aliases for the SQL Server host in the Aliases field, as in the following example:

```
mars
mars.acme.com
```

10. Save this client resource configuration.

To write the results of a scheduled backup to a log file, enter **nsrlog -f filename** (Windows NT and Windows 2000) or **/usr/ucb/logger** (UNIX) in the Notifications resource in NetWorker Administrator.

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

Backup Operation Status

NetWorker server and 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 scrolled list in the Group Detail window of NetWorker Administrator. 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.

If you are using a NetWorker server on a UNIX host, you can also send an e-mail notification of the results of a scheduled backup to the owner of a save set:

- a. Use the View Details option to edit the NetWorker Clients 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. You can edit the notification setup for this report using the Customize feature of NetWorker Administrator.

Refer to the section on setting up event notifications in the *Legato NetWorker Administrator's Guide* for details on using NetWorker Administrator to customize your notifications.

For complete details about the reports generated by NetWorker, see the *Legato NetWorker Administrator's Guide*.

2

Chapter 3: NetWorker Module Restore Concepts

NetWorker User for SQL Server can be used at any time to restore Microsoft SQL Server data backups, which were created with Legato NetWorker Module. BusinessSuite Module 2.0.1 included support for restoring SQL Server 6.5 and version 7.0 data, including the ability to verify a backup without actually performing a restore operation. NetWorker Module 3.0 provides support for SQL Server 7.0 and Microsoft SQL Server 2000. Restore capabilities have been enhanced to include new restore options such as database backup copy, restoring to a named log mark, and partial database restore.

Prior to configuring the NetWorker Administrator to work with the NetWorker Module, you should understand the differences between the different restore-types that are available, and how the type of and level of backup created will affect your options for restoring data. Specifically, this chapter examines the following:

- Microsoft SQL Server 2000 recovery models
- Restore-types: *normal*, *partial*, *verify only*, and *copy*
- Restore modes: *normal*, *no recovery*, and *standby*
- Restore time: *backup versions*, *log marks*, and *point-in-time*
- NetWorker Module database restoration and the master database



Important: It is *highly recommended* that you read this chapter and [“Chapter 2: NetWorker Module Backup Concepts” on page 21](#), prior to performing a backup or restore operation using Legato NetWorker Module. The type and level of backup you create 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.

NetWorker Module 3.0 enforces backup and restore constraints differently, depending on whether you are managing a SQL version 7.0 or version 2000 server. [Table 7](#) provides an overview of the available backup types supported by the NetWorker Module program for the recovery models available in SQL Server 2000:

Table 7. 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 8](#) provides an overview of the various backup and restore functionality supported by the NetWorker Module program for the recovery models available in each version of SQL Server:

Table 8. Supported Backup and Restore Functionality for SQL Server 2000 Recovery Models

Functionality	FULL	BULK_LOGGED	SIMPLE
Transaction Logging	Yes	Yes	Yes
Point-in-time Restore	Yes	Maybe	No
Named Log Mark Restore	Yes	Maybe	No
Bulk Copy	No	Yes	No

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

FULL Recovery Model

The FULL recovery model implements the fewest constraints on the backup and restore process, but requires the most log space of all recovery models. NetWorker Module 3.0 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 implements a few backup constraints and supports reduced log space requirements for bulk operations. NetWorker Module 3.0 enforces the following constraints differently, depending on whether you are managing a SQL Server version 7.0 or version 2000 storage hierarchy:

- All levels of database, filegroup, and file backup are supported for SQL Server version 2000, while only level full and differential database backup types are supported for SQL Server version 7.0
- All SQL Server version 2000 backup data is restorable up to the most recent transaction log, while all SQL Server version 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 version 7.0, and may be supported for SQL Server version 2000, depending on the following conditions:
 - If a bulk log change has occurred for the transaction log backup that corresponds to the current browse time, point-in-time restore is not supported. The NetWorker Module will roll forward all transactions to the end of the transaction log and leave 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 allows you to have significant control over how the storage management operations impact your system resources. NetWorker Module 3.0 enforces the following constraints on both SQL Server version 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 have been set in SQL Server, they are mapped 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 database, and therefore, the descendent filegroups and files of the database, are available for back up or restore, depends on the recovery model specified in SQL Server. NetWorker Module 3.0 modifies how items in the SQL Server storage hierarchy are displayed in the Backup and Restore Browse windows, based on the constraints enforced by the recovery model. For more information about how these constraints are visually interpreted in NetWorker User for SQL Server, see [“Fake Objects” on page 58](#).

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 your current backup and restore strategies for a database. The following sections provide instructions for preparing for the recovery model change and backup strategies for maintaining the consistency of your SQL Server storage hierarchy after the change.

How to Change from FULL or BULK_LOGGED to SIMPLE

Follow these steps when changing to the SIMPLE recovery model:

1. Prior to 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.
3. After changing the recovery model, modify your backup strategy to halt execution of level incremental (transaction log) database backups. If the change to SIMPLE is only temporary, it is not necessary to modify the backup schedule, because, for NetWorker Module version 3.0, the backup level will automatically be 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 you anticipate having to perform a point-in-time restore, it is *highly recommended* that you perform a level incremental database backup to maintain the transaction log files.
3. You do not need to modify your backup strategy.

How to Change from SIMPLE to FULL or BULK_LOGGED

Because the SIMPLE recovery model has many more constraints than either the FULL or BULK_LOGGED models, 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 your 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, 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, you will have to reinitiate the restore, 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, it is *highly recommended* that you modify your backup strategy.

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 your backup strategy to include level incremental (transaction log) database backups (*highly recommended*).

3

Restore-types

In order to restore your data, 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. 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, filegroup, or a file to the database originally backed up. The normal restore-type can restore level full, level 1 and level incr backups in the order required by SQL Server. When performing a restore the normal restore-type is the default type used by the NetWorker Module. For more information about creating backups, see [“Chapter 2: NetWorker Module Backup Concepts” on page 21](#).

Because NetWorker Module can be used to 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 your database(s) is possible. In addition, a single or multiple filegroup(s) and / or file(s) 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 for filegroups not selected are created, but will be empty.

3

Verify Only Restore-type

The verify only restore-type only verifies the backup media for the selected SQL Server 7.0 or later backups. *Selecting the verify only restore-type will NOT restore your 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 restore and share data that was previously backed up from a computer that is not currently functioning. The database backup to restore may be obtained from the local host or imported from a different NetWorker SQL Server client. In NetWorker Module 3.0, *Copy* replaces the directed recover operation, which existed in previous versions of BusinessSuite Module for Microsoft SQL Server.

Restore Modes

In order to restore a database, NetWorker Module requires that a restore mode be specified for the restore operation. A *restore mode* instructs the SQL Server on how to interact with the database after the restore operation has completed. For instance, restore modes can be used to leave the database in an

intermediate state, so that additional transaction logs may 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 during the restore, which allows database reads and writes; however, additional transaction log backups may not be applied. The normal restore mode is the default mode used by the NetWorker Module when restoring a database.

3

No Recovery Restore Mode

The no recovery restore mode activates the SQL Server NORECOVER database restore option, which 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, 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 provide the ability to restore your backups 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 you inform the NetWorker Module to discard transactions performed after a given time.

The default, or *current* restore time for each database is derived from the create time of the marked item. By default, the most recent backup is restored. If the most recent backup is level incr or 1, dependent backups are restored first. User-specified restore times allow you to restore older backup versions or point-in-time backups. 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 3.0, the global browse time *only* controls which backup data is viewable in the Restore Browse window.

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

Database Backup Versions

In the Restore Browse window, you can browse the online index as of any time in the past and mark a database backup version to restore. Use the Restore Time tab from the database restore property dialog to select any backup version for restore. The Restore Time tab makes the View>Versions option obsolete for selecting a backup version for restore. However, use View>Versions to browse and mark the filegroup or file hierarchy of an older database. The browse tree will be reconstructed with the older version of the database.

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

3

Point-in-time Restore of a Transaction Log Backup

When a level incremental database backup (transaction log backup) has been marked in the Restore Browse window, you can select to perform a point-in-time restore in the Restore Time tab of the Database Item Restore Properties dialog box. 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 Browse window will be the default row marked in the backup version table in 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 is only applicable when a transaction log backup is marked for restore.

Microsoft SQL Server 2000 Named Log Marks

Microsoft SQL Server 2000 introduces enhanced transaction log restoration, by enabling *named log marks* to be specified during transaction creation. 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 Module 3.0 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 time of the log mark is saved to the millisecond. However, the time format implemented by Legato NetWorker, 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, NetWorker Module will restore to the most recently named log mark.

For more information about transaction log marks, see the Microsoft SQL Server documentation. For more information about performing a restore using named log marks, see [“Database Item Restore Time Properties” on page 87](#).

Specifying the Restore Time

You can modify the restore time for each marked database, by changing the values on the Database Item>RightClickMenu>Properties>Restore Time tab. Unique restore times can be set on a per-item basis. For example, if two databases *foo* and *bar* are marked for restore, *foo* can have a different restore time than *bar*. Before NetWorker Module 3.0 the global browse time was used as the restore time.

For more information about modifying the restore time values, see [“Database Item Restore Time Properties” on page 87](#).

NetWorker Module Database Restoration

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

1. NetWorker Module restores the most recent full backup, then restores the most recent differential (level 1) backup (if any).
2. NetWorker Module then 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, NetWorker Module specifies the SQL Server NORECOVER 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.

This explains the appearance of the word “norecovery” in the following sample output to standard out for a database restore:

```
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 stopat='19980512 12:26:41'
Received 1.0 MB 4 file(s) from NSR server.
```

3

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. Microsoft SQL Server documentation contains detailed instructions on rebuilding the master database.
- Loading the data for the master database and SQL Server databases from the NetWorker server. See [“Disaster Recovery” on page 107](#) for more information on restoring the master database.

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 as appropriate, as follows:

1. Before the restore begins, NetWorker Module stops the SQL Server services.
2. NetWorker Module starts the SQL Server services in single-user mode.
3. NetWorker Module performs the restore.
4. After the restore finishes, NetWorker Module stops the SQL Server services.
5. NetWorker Module restarts the SQL Server services in multi-user mode.



Important: When restoring the Microsoft SQL Server *msdb* database in a non-clustered environment, *you must manually stop the SQL Agent services before initiating the restore.* This is a restriction with the current version of NetWorker Module 3.0.

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, NetWorker Module stops the SQL Server and all dependent cluster resources.
2. NetWorker Module starts the SQL Server services in single-user mode.
3. NetWorker Module performs the restore.
4. After the restore finishes, NetWorker Module stops the SQL Server services.
5. NetWorker Module restarts the SQL Server and dependent cluster resources.

Chapter 4: NetWorker Module Usage and Concepts

This chapter describes the visual appearance, interaction mode, and basic data requirements of NetWorker Module 3.0. Before attempting to back up and restore your Microsoft SQL Server version 7.0 or 2000 data, it is highly recommended that you review the information in this chapter.

Topics in this chapter include:

- Browsing and navigating the SQL Server storage hierarchy in NetWorker User for SQL Server
- Marking items in NetWorker User for SQL Server
- Notation and syntax for specifying SQL Server data items at a command prompt

Browsing Data Items in NetWorker User for SQL Server

NetWorker User for SQL Server displays the data items in the SQL Server storage hierarchy in either the Backup or the Restore Browse window, depending on which operation you have selected to perform.

The Browse window is split into two panes. The left pane consists of a tree with expand / collapse nodes, which displays the SQL Server storage hierarchy contained on the current NetWorker Module host. The levels of data within the tree can be made visible by double-clicking the expand / collapse icons. When an item is selected in the left pane, the right pane displays the descendents of the selected item.

BusinesSuite Module for Microsoft SQL Server version 2.0.1 and earlier provided support for browsing your SQL Server databases and filegroups; however, filegroups were not displayed in an hierarchical manner. Filegroups

were not children of the database, they were peers. NetWorker Module 3.0 also includes the ability to browse the filegroups and files contained in a database; however, the NetWorker User for SQL Server may display certain data items which are contained in the storage hierarchy, but 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 has been 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. For more information about this dialog, see [“Chapter 3: NetWorker Module Restore Concepts” on page 45.](#)

4

Fake Objects


In the Backup Browse window, 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 using the Fake Filegroup or Fake File conventions outlined in [Table 9 on page 59.](#) For more information about the constraints enforced by Microsoft SQL Server 2000, which determine whether an item is available for back up or restore, see [“Microsoft SQL Server 2000 Recovery Models” on page 46.](#)

Figure 3. 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 Browse window by using the fake object conventions outlined in [Table 9 on page 59](#).

Figure 4. Fake Objects in the Restore Browse Window



Display Conventions

NetWorker Module uses specific data item names, text characteristics, and data type icons to distinguish the variable qualities of SQL Server data. The following table outlines these conventions:

Table 9. SQL Server Storage Hierarchy Display Conventions

Data Item	Description	Pane	Text Quality	Icon
SQL Server	- root of the storage hierarchy - signifies all SQL Server databases on the host	Left only	- default Example: MS SQL Server	
Database	- descendent of root - signifies a database - may contain filegroups - may contain files	Left - and - Right	- default Example: Database	
Fake Database	- only visible in the Restore Browse window - signifies that the database cannot be selected for back up	Left - and - Right	- light gray - bold Example: Database	

Table 9. SQL Server Storage Hierarchy Display Conventions (Continued)

Data Item	Description	Pane	Text Quality	Icon
Filegroup	- descendent of a database - signifies a database filegroup - may contain files	Left - or - Right	- default Example: Filegroup	
Fake Filegroup	- signifies that the filegroup cannot be selected for back up	Left - or - Right	- light gray - bold Example: Filegroup	
File	- descendent of a filegroup - signifies a database file	Right only	- default Example: Filename	
Fake File	- signifies that the file cannot be selected for back up	Right only	- light gray - bold Example: Filename	

4

Marking Items in NetWorker User for SQL Server

In the Browse window, you specify items for backup or restore by marking them. To mark an item, highlight the item and click the Mark button. To deselect an item, highlight the item and click the Unmark button.

NetWorker Module 3.0 introduces support for the right click mouse button menu for data items in the browse window. Items can be marked by clicking the right mouse button to expose the right-click menu and selecting Mark. Items can be deselected by clicking the right mouse button to expose the right-click menu and selecting Unmark.



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 cues to distinguish them from normal data items. See [Table 9, "SQL Server Storage Hierarchy Display Conventions," on page 59](#) for more information.

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 has been chosen, only one *database* object may be marked. Marking the root SQL Server item is not allowed.

When the *partial* restore-type has been chosen, the subset of filegroups and files of the selected database must be marked using the Database Item Restore Properties dialog. For more information about this dialog, see ["Database Restore Properties" on page 78](#). When the *copy* restore-type has been chosen, filegroups and files of the selected database are automatically marked and restored as part of the full database restore.

4

Marking Semantics and Restrictions

In order to support the new browsing capabilities, NetWorker Module 3.0 implements 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 you are marking items, you may encounter message dialogs, which provide advice regarding the current marks and the type of operation that will be performed. These dialogs are there to assist you in creating the different types of backup and restore-types. For more information about backup types, see ["Chapter 2: NetWorker Module Backup Concepts" on page 21](#). For more information about restore-types, see ["Chapter 3: NetWorker Module Restore Concepts" on page 45](#).

The following rules identify the marking semantics and restrictions:

- It is possible to mark a single database, filegroup, or file
- It is possible to mark multiple, heterogeneous items


- It is possible to mark an item, when any of that item's siblings are already marked
- It is possible to mark all SQL Server data, by using the SQL Server root item's right click menu:

MS SQL Server>RightClickMenu>Mark all children

This rule provides the ability to easily create a *full* backup of your entire SQL Server storage hierarchy.

- It is not possible to mark an item, when any of that item's ascendants or descendants are already marked. To mark an item, the marked ascendants and descendants must first be unmarked.
- When a database is marked, all of the item's descendants are NOT automatically marked.
- When a filegroup is marked, all of the files are NOT automatically marked.
- Selecting all descendants of a database or filegroup, by marking each descendent individually, is equivalent to using the right click menu "Mark All ..." option on the root of the storage hierarchy.

Chapter 5: Backing Up SQL Server Data



This chapter describes the steps required to back up your Microsoft SQL Server data from NetWorker User for SQL Server. If you would prefer to perform the backup from a command prompt, refer to [“Appendix A: Commands and Command Options” on page 115](#) for guidelines on using the NetWorker Module commands and command options. SQL Server data cannot be *manually* backed up using NetWorker Administrator.

Preparing for the Backup


Review the following information prior to initiating the backup operation:

- Remember that the Backup Browse window displays data available for back up, based on the SQL Server database settings. For more information about the browsing items, see [“Browsing Data Items in NetWorker User for SQL Server” on page 57](#).

Creating the Backup


When you start NetWorker User for SQL Server, the main window appears:

Figure 5. Main window in NetWorker User for SQL Server



If multiple SQL Servers are active on the computer, the Select SQL Server Host dialog will appear before the main window is opened:

Figure 6. Select SQL Server Host



Select the SQL Server host that NetWorker Module should use to perform the back up operation and click OK. The main window will be opened.


How to Back Up Using NetWorker User for SQL Server

When you run a backup from NetWorker User for SQL Server, NetWorker always performs a *full* backup of your SQL Server data. To run a backup from NetWorker User for SQL Server:

1. In the NetWorker User for SQL Server window, click the Backup button or select Operation>Backup.

After you have selected to perform a backup, the Backup Browse window appears, as shown in the following figure:

Figure 7. Backup Browse Window



2. To view a list of data items available for backup, double-click the expand / collapse icon next to an item in the left pane of the Backup Browse window. The descendents of the item are displayed in the right pane of the window.
3. In the Backup window, select each item to back up, then click the Mark button or select Item>RightClickMenu>Mark. For information about marking items for backup, see [“Marking Items in NetWorker User for SQL Server” on page 60](#).
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 General Properties” on page 69](#).
5. (Optional) Specify preferences for backup operation settings by selecting Options>Backup Options and making entries as needed, then click OK. For information about setting options, see [“Backup Options” on page 66](#).
6. Click the Start button to begin the backup.

The amount of time NetWorker takes to back up your database depends on database size, network traffic, server load, and tape positioning. During the backup, NetWorker Module displays messages for each database backed up in the Backup Status window, so you can monitor the progress of your backup.

After your backup is finished, a Backup Completed message appears. Close the Backup Status window by double-clicking on the control-menu box or by selecting Close from the File menu.



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 will be displayed in the status window and the backup will wait for operator intervention.

You can cancel a backup by selecting End Backup from the File menu or by closing the Backup Status window.

Backup Options


5

You can customize the backup operation by specifying preferences in the Backup Options dialog box before you start a backup. These properties are supported for SQL Server version 7.0 and 2000.

With the Backup Window open, you can invoke the Backup Options dialog box in the following ways:

- On the toolbar, select Options>Backup Options.
- Right click the root item and 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 toolbar.

Preferences that you specify for these properties will be applied to each marked item and will be used for the current and subsequent backup operations until you close the NetWorker Module GUI. Pool settings you choose are stored in the registry and persist from one GUI session to another. When a preference is not specified, the NetWorker Module default will be used.

Figure 8. General Backup Options

The General tab of the dialog box features basic properties for configuring, performing, and storing the backup. The options on the general tab of this dialog box are as follows:

- (Optional) Compress the backup contents — (default = unchecked) Check this option to apply XBSA compression to all marked databases before writing the backup data to the storage device. In the same manual backup, it is not possible to back up certain databases with compression and others without compression.

Compressing data for a backup generates less network traffic and uses less backup media space, but takes extra time. Some hardware performs compression, which makes software compression unnecessary. Ask your administrator if you need to compress your data during backups.

- (Optional) Create a striped backup — (default = unchecked) Check this option to create a striped backup using the SQL Striped feature. When this option is checked, the Stripes option is enabled. Use the down or up arrows to specify the number of stripes you want. The maximum number of stripes allowed is the same for both SQL Server versions:
 - SQL Server 7.0 — 32 stripes
 - SQL Server 2000 — 32 stripes

For more information about using this option, see [“Using Striping” on page 96](#).


- (Optional) Use pools for media management — (default = unchecked)
Check this option to specify media volume pools for storing the backup contents. *This option only applies to manual backup operations performed from NetWorker User for SQL Server or from a command prompt.* When this option is checked, the Full backup pool and Log file backup pool filename options are enabled.

By default, NetWorker Module displays a list of valid pool names in each volume pool dropdown control. Click on the Full backup pool down arrow to select the pool to store full SQL Server backups including databases and filegroups. Click on the Log file backup pool down arrow to select the pool to store transaction log or differential SQL Server backups.

Volume pools provide the ability to segregate different kinds of backup data and allow you to direct backup data to specific devices. For more information about using this option, see [“Volume Pools” on page 36](#).

The Advanced tab of the Backup Options dialog box provides access to complex, infrequently used configuration settings for the backup.

Figure 9. Advanced Backup Options



Important: Backup optimization using the RAW save set is a new feature provided in NetWorker Module 3.0 and is *only* available for configurations running NetWorker Module 3.0 and NetWorker 6.0 or higher with a Power Edition license. In addition, the NetWorker client and the NetWorker storage node must be installed on the same computer.

The RAW save set option is available on the advanced tab:

- (Optional) Use the RAW save set — (default = unchecked) Check this option to use the RAW save set when creating a striped backup. The RAW save set may be used to lower CPU utilization when performing a striped backup. *It is highly recommended that you only use the RAW save set 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.


For more information about using the RAW save set, see [“Using the RAW Save Set” on page 99](#).

Database Item Backup Properties

These properties allow you to specify preferences for a single, *marked* database before you start a backup. The Database Item Backup Properties dialog box appears when you select a *marked* database in the Backup Browse window and select the shortcut menu>Properties, or click the Properties toolbar button. These properties are new in NetWorker Module 3.0 and are supported for SQL Server version 7.0 and 2000.

Preferences that you specify for these properties will be used for the current backup operation. When a preference is not specified, the NetWorker Module default will be used. After the backup operation has completed, and you exit from the Backup Browse window, the property values will revert back to the NetWorker Module defaults.

Figure 10. Database Item Backup General Properties



In NetWorker Module 3.0, only one database property exists:

- (Optional) Truncate the transaction log file before performing the backup — (default = unchecked) Check this option to truncate the transaction logs before backup, then back up the databases and the transaction logs. In BusinessSuite Module 2.0.1 and earlier, this property was available from the Options>Backup Options dialog box.

Target Volume Pools for Unscheduled Backups

For your manual backups, you can optionally specify target Data and Log pool names in the Backup Options dialog box in NetWorker User for SQL Server (see [Figure 8 on page 67](#)). The pool names you specify in this dialog box are stored in the operating system registry and are persistent from one GUI session to the next.

If you specify a Data pool:

- All backups you perform from NetWorker User for SQL Server are directed to this pool.
- All full (and only full) backups you initiate from a Windows command prompt are also directed to this pool. (However, when you run 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 115](#)).

Settings in the Log pool field in the Backup Options dialog box pertain only to backups you initiate from a command line. If you specify a Log pool, all incremental *and* level 1 backups you initiate from the command line are directed to that pool — unless you specify otherwise using a parameter to the `-b` command-line option.

For example, the following command entered at a Windows command prompt:

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

would override any Log pool setting specified in the Backup Options dialog box.

Assigning Pools in NetWorker User for SQL Server

To specify names for volume pools in the Backup Options dialog box, check the Use Pools for media management option (see [Figure 8 on page 67](#)). In the pool name fields, you can then specify volume pool names so that your full backups are directed to one pool and your incremental and differential backups are directed to a different pool. Legato NetWorker Module stores these values in the Windows NT and Windows 2000 registries.

**Example:**

For the name of the data pool, you might specify something like one of the following:

SQLServerFulls. You might elect to have this be the same pool to which you are also directing your SQL Server *scheduled* full backups.

SQLServerFullManuals. You could use a name like this to signify that you intend to direct all full, unscheduled backups to this pool.

SQLServerGUIFulls. You could use a name like this to signify that you will be directing only backups initiated from the GUI to this pool.

For the name of the log pool, you might choose something like **SQLServerLogs.** This could be the same pool to which your incremental and level 1 scheduled backups are directed.

Specifying Matching Pools and Label Templates in NetWorker Administrator

Before you start a manual backup, you must ensure that pools and label templates have been configured in NetWorker Administrator to match any pool names you specified in the Backup Options dialog box in NetWorker User for SQL Server. If there is not a match when the operation begins, NetWorker cannot validate the Backup Options pool selection and the backup will not succeed.

Use NetWorker Administrator to configure data and log pools and label templates as follows:

1. Create a volume pool, following the instructions provided in the *Legato NetWorker Administrator's Guide*.
2. For Levels, specify **full**.
3. Before saving the configuration, ensure that “Yes” is highlighted as the choice for Store index entries.
4. Save the pool configuration.
5. A message box appears, prompting you to either:
 - Select a new label template

- or -

- (Default) Apply the configuration, to automatically create and select a label for the volume pool
6. Click OK. The configuration will be reapplied and the corresponding label template will be automatically created.
 7. Now create a log volume pool.
 8. Specify **incr** and **1** for Levels.
 9. Repeat Steps 3 through 6 to create another label template corresponding to the second new pool.
 10. Insert a new media volume, or select another drive or slot with media loaded for labeling.
 11. Click the Label speedbar button to label the media volume, selecting the label template customized for the data pool. Repeat the process to label a media volume for the logs pool, selecting the label template customized for that pool.

Target Pool Specification Guidelines

Keep in mind the following guidelines when setting target pools:

- Pool selection criteria you customize in NetWorker Administrator pertain only to scheduled backups
- Pool settings in the Backup Options dialog box do not pertain to scheduled backups
- A Data pool setting in the Backup Options dialog box pertains only to full manual backups (from the GUI or the command line)
- A Log pool setting in the Backup Options dialog box pertains only to incr or level 1 backups from the command line
- A pool name specified on a command line (as a parameter to `nsrqlsv -b`) overrides any other target-pool settings

Chapter 6: Restoring SQL Server Data

In order to restore your SQL Server data, you must use NetWorker User for SQL Server or the command prompt features of NetWorker Module. SQL Server data cannot be restored using NetWorker Administrator.

This chapter describes the steps required to restore your Microsoft SQL Server data from NetWorker User for SQL Server. If you would prefer to perform the restore from a command prompt, refer to [“Using the nsrsqlrc Command” on page 116](#) for guidelines on using the NetWorker Module commands and command options.

Preparing for 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 you 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 in NetWorker User for SQL Server” on page 57](#).
- 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 does not allow you to restore a database while a backup of another database 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.


Performing the Restore

When you start NetWorker User for SQL Server, the main window appears, see [Figure 5 on page 64](#). If multiple SQL Servers are active on the computer, the Select SQL Server Host dialog will appear before the main window is opened, see [Figure 6 on page 64](#). Select the SQL Server host which NetWorker Module should use to perform the restore operation and click OK. The main window will be opened.

How to Restore Using NetWorker User for SQL Server

1. In the main window, initiate a restore by performing one of the following actions and selecting a restore-type, see “[Chapter 3: NetWorker Module Restore Concepts](#)” on page 45:
 - Click the Restore button and select the restore-type from the Restore-Type Selection dialog and click Continue:


Figure 11. Restore-Type Selection Dialog



- or -


- Select Operation>Restore and then choose the restore-type from the Restore-type menu:

Figure 12. Operation>Restore>Restore-Type Menu



If you select the Copy restore-type, the SQL Server Selection dialog is displayed:

Figure 13. SQL Server Selection Dialog




This dialog allows you to specify the location of the existing database backup to use for the copy operation. The options are as follows:

- The current SQL Server — (default = selected) By default, this option displays the name of the current NetWorker client host. Select this option to use a database backup from the SQL Server currently connected to by NetWorker Module.
- A different NetWorker SQL Server client — (default = not selected) By default, this option displays the name of the current NetWorker client host. Click on the down arrow to select a different SQL Server, from which to obtain the existing database backup.

For more information about selecting a restore-type, see [“Chapter 3: NetWorker Module Restore Concepts” on page 45.](#)

After you have selected the restore-type, the Restore window appears, as shown in the following figure:

Figure 14. Restore Browse Window



2. To view a list of data items available for restoring, double-click the expand / collapse icon next to an item in the left pane of the Restore Browse window. The descendents of the item are displayed in the right pane of the window.



Important: When the *partial* or *copy* restore-type has been 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. For more information, see [“Restore Browse Window Restrictions” on page 58](#).

3. In the Restore window, select each item to back up, then click the Mark button or select Item>RightClickMenu>Mark.



Important: When the *partial* restore-type has been chosen, only *one database* object may be marked. Filegroups and files of the selected database must be marked using the Database Item Restore Properties dialog. When the *copy* restore-type has been chosen, filegroups and files of the selected database are automatically marked for restoration. For more information, see [“Restore Browse Window Restrictions” on page 61](#).

4. Set the restore properties for each marked database, if any, by selecting Database Item>RightClickMenu>Properties (or select the marked database and click the properties button) and fill in the Properties dialog box. Then click OK. For information about setting Database Restore properties, see [“Database Restore Properties” on page 78](#).
5. Set the restore properties for each marked filegroup, if any, by selecting Filegroup Item>RightClickMenu>Properties, and then click OK. For information about setting Filegroup Restore properties, see [“FileGroup Restore Properties” on page 90](#).



Important: When the *partial* or *copy* restore-type has been chosen, the filegroups of the database and their associated restore properties must be specified by selecting Database Item>RightClickMenu>Properties>Files. *These properties must be set before initiating the restore operation. NetWorker Module will not allow you to start the restore before the required filegroup item properties have been set.* For information about setting Database Restore properties, see [“Database Restore Properties” on page 78](#).

6. Set the restore properties for each marked file, if any, then click OK. For information about setting File Restore properties, see [“File Restore Properties” on page 92](#).



Important: When the *partial* or *copy* restore-type has been chosen, the files of the database and their associated restore properties must be specified by selecting Database Item>RightClickMenu>Properties>Files. *These properties must be set before initiating the restore operation. NetWorker Module will not allow you to start the restore before the required file item properties have been set.* For information about setting Database Restore properties, see [“Database Restore Properties” on page 78](#).

7. Click the Start button to begin the restore.

The amount of time NetWorker Module takes to restore your databases and transaction logs depends on the amount of data, network traffic, server load, backup volume availability, and tape positioning. During this time, NetWorker Module displays messages for each restored database in the Restore Status window, so you can monitor the progress of your restore operation.

If the backup volume containing your databases is loaded at a storage node (backup device) local to the NetWorker server, the restore will proceed. If the restore does not begin soon, it is possible that either no volume or the wrong volume is mounted in the backup device. You can cancel the restore by selecting End Restore from the File menu or by closing the Restore Status window.

If the database you are restoring already exists on your disk at the target drive and pathname, and the source and target databases have the same name and logical file configuration, the target database will be overwritten by the restore operation.

After your restore is finished, a Restore Completed message appears. Close the Restore Status window by double-clicking on the control-menu box or by selecting Close from the File menu.

Database Restore Properties

These properties allow you to specify preferences for a single, *marked* database before you start a normal, partial, or copy restore-type operation. *These properties are not available for a verify only restore-type operation.* The Database Item Restore Properties dialog box appears when you select a *marked* database in the Restore Browse window and select File>Properties or Database Item>RightClickMenu>Properties.

Some of these properties are new in NetWorker Module 3.0, while others existed in BusinessSuite Module 2.0.1 and earlier, and were available from Options>Recover Options. All of the properties are context sensitive, based on which version of SQL Server you are running (7.0 or 2000) and the restore-type you have chosen.

Preferences that you specify for these properties will be used for the current restore operation. When a preference is not specified, the NetWorker Module default will be used. After the restore operation has completed, and you exit from the Restore Browse window, the property values will revert back to the NetWorker Module defaults.

6


Database Item Restore General Properties



Important: The following information applies to normal and copy restore-types *only*.

The General tab of the dialog box features basic, frequently used properties for configuring the restore. The properties are supported for SQL Server version 7.0 and 2000.

Figure 15. Database Item Restore General Properties



The options on this tab as follows:

- (Optional) Back up the active portion of the transaction log before restoring the database — (default = not checked) Check this option to back up the active portion of the transaction log prior to performing the restore.

When this option is selected, NetWorker Module will attempt a transaction log back up using the NO_TRUNCATE SQL keyword. The restore operation will proceed regardless of whether the transaction log backup succeeds or fails.

The following restrictions apply to the transaction log option:

- If the Specify a restore time checkbox on the Database Item Restore Time Properties tab is checked, an error message will be displayed and this option will be ignored. You cannot perform a transaction log backup and also restore to a specific point-in-time.
- (Required) Perform the restore using this restore mode (Normal, No Recovery, Standby) — (default = Normal) Use this option to specify which mode to use for the restore operation.
 - Normal — Select Normal to perform a standard restore operation.
 - No Recovery — Select No Recovery to activate the NORECOVER option, which places the database in an unloadable state after the restore, but still able to process additional transaction log restore operations.

- Standby — Select Standby to 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 option is enabled to specify an *undo* file for SQL Server to use when rolling back the transactions. By default, this option displays a default filename and path


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

default_dbName is the name of the database backup selected for restoration. 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 clicking on the ellipse button to open the Standby Undo File Selection dialog. The Standby undo file name text field is always disabled.

Standby Undo File Selection Dialog

The Standby Undo File Selection dialog contains properties related to specifying the location of the standby undo file.

Figure 16. Standby Undo File Selection Dialog



The options on the Standby Undo File Selection dialog are as follows:

- (Required) File location — (default = disabled) This option displays the path for the standby undo file. This option can only be modified by browsing the filesystem tree and highlighting a file. When a file is highlighted, the path for the highlighted file is displayed in the File location field.


- (Required) File name — Use this option to specify the standby undo filename by entering the filename in the field or by browsing the filesystem tree and highlighting a file. When a file is highlighted, the file name is displayed in the File name field.

Database Item Restore Files Properties and Partial Restore-Type Database Item Restore General Properties

For normal and copy restore-types, this tab is labeled *Files* and is supported for SQL Server version 7.0 and 2000. For a partial restore-type, the tab is labeled *General* and is only available for SQL Server 2000. This tab contains properties related to selecting the filegroups to restore and specifying the destination locations for the restored files.

Note: The Files tab may first invoke the Read File Configuration dialog before showing the completed Files tab. If this happens, the marked database item you selected was created by a pre-3.0 version of NetWorker Module. For details, refer to [“Read File Configuration Dialog” on page 85](#).

Figure 17. Database Item Restore Files Properties



6

The options on the Files tab of this dialog box are shown below.

Note: If the fields on this dialog are empty, you may have declined to read the file configuration information. For details, refer to [“Read File Configuration Dialog” on page 85](#).

- (Information) Database to restore — (default = disabled) This option is provided for informational purposes only and cannot be modified. This option displays the name of the database selected for restoration.
- Name for restored database — Use this option to specify the name for the restored database.
 - (Information) Normal restore-type — By default, this option displays the name of the database selected for backup and is disabled.
 - (Required) Partial and Copy restore-types — By default, this option displays a default name based on the name of the database selected for backup. To specify a different or new name, enter a name in the text field or select a name by clicking on the down arrow and highlighting a name from the list.

If necessary, file configuration will be regenerated after the Apply / OK button is depressed. If a different database name is specified, NetWorker Module automatically generates a new physical file configuration. If a new database name is specified, any changes to the file configuration are discarded.

When this option is set to the name of the database selected for backup, the Overwrite the existing database option is enabled. These two options can then be used together to specify the same name for the restored database as the existing database, when the two databases are incompatible.

- Overwrite the existing database — (default = not checked) Check this option to restore the SQL Server data to the exact location (drive and pathname) as the location on the source host from which the data was backed up.
 - (Optional) — When restoring a database backup whose files and / or metadata are the same as those of the existing database.
 - (Required) — When restoring a database backup whose files and / or metadata differ from those of the existing database. In this case, selecting this option causes the WITH REPLACE SQL keyword to be included in the restore sequence.
- Mark the filegroups to restore (Mark / Unmark) — Use this option to select / deselect the filegroups to restore.
 - (Information) Normal restore-type — (default = disabled) By default, this option displays the filegroups marked for restoration in the Restore Browse window. For more information, see [“Marking Items in NetWorker User for SQL Server” on page 60](#).
 - (Required) Partial restore-type — (default = enabled) By default, this option displays the filegroups of the database marked for restoration.

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

More than one filegroup may be marked. The PRIMARY filegroup is always marked and cannot be unmarked.

- (Information) Copy restore-type — (default = disabled) By default, this option displays all of the filegroups of the database selected for restoration. When performing a copy, individual filegroups and files of the selected database are automatically marked and restored as part of the database restore. For more information, see [“Marking Items in NetWorker User for SQL Server” on page 60](#).

The set of filegroups marked in this option, are copied into the dropdown list of the Modify the destination for the files in option.

- Modify the destination for the files in — (default = enabled) Use this option to filter which files are visible in the File / Destination List.

The dropdown combobox contains a set of different views for the database files to be restored. The following views are supported:

- All view — Display all of the files selected for restoration, including transaction log files.
- Log files view — Display *only* the transaction log files.
- Fg_Name filegroup — Display all of the files in the filegroup named Fg_Name. Only those filegroups marked for restore are included in this combobox.

To filter the view, click the down arrow button and highlight the filegroup or view for which you would like to see the files. The File Relocation List displays the source files and restored files / destinations which correspond to the selected view / filegroup.

- File / Destination List — (default = enabled) This table shows the names of files, associated with the marked database, that will be restored, and the destinations to which the files will be restored.
 - (Optional) Normal restore-type — By default, this option displays a default name and destination based on the SQL Server physical filename and logical location for the restored file.
 - (Required) Partial and Copy restore-types — By default, this option displays a default name and destination based on the SQL Server physical filename and logical location for the restored file.

The File and destination entries are not editable; they are updated based on the a dialog invoked by the Destination... button. When a file in the list is selected, the Destination... button is enabled.

To modify the destination for one file, click on the file in the list and click the Destination... button to display the File Destination dialog.

To modify the destination for multiple files, do either 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 collection of non-contiguous 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 File Destination dialog.


The following restrictions apply to the File / Destination List option:

- If necessary, file configuration will be regenerated after the Apply / OK button is depressed.
- If a different database name is specified, NetWorker Module automatically generates a new physical file configuration. If a new database name is specified, any changes to the file configuration are discarded.

File Destination Dialog

The File Destination dialog contains properties related to specifying the destination locations for the restored files.

Figure 18. File Destination Dialog




The options on the File Destination dialog are as follows:

- (Information) Source file name — (default = disabled) This option is provided for informational purposes only and cannot be modified. The name, as it exists in the backup to restore, of the file currently selected in the File / Destination list is displayed in the Source file name field. When multiple files have been selected, this option field is empty.
- (Information) Source location — (default = disabled) This option is provided for informational purposes only and cannot be modified. The filesystem location, as it exists in the backup to restore, of the file currently selected in the File / Destination list is displayed in the Source location field. When multiple files have been selected, this option field contains the filesystem location of the first selected file in the File / Destination list.
- (Required) Destination location — (default = disabled) This option displays the filesystem location for the restored file. When multiple files have been selected, the default SQL data path is opened, but not selected. This option can only be modified by browsing the filesystem tree and highlighting a directory or file. When a directory is highlighted, that path is displayed in the Destination location field. If a file is highlighted, the directory for the highlighted file is displayed.
- (Required) Destination name — By default, this option displays the name of the file currently selected in the File / Destination list. When multiple files have been selected, this option field is empty. When a file is highlighted, the filename is displayed in the Destination name field, and the path containing the file is displayed in the Destination location field. Use this option to specify the new name to use for the restored file by entering the filename in the field or by browsing the filesystem tree and highlighting a file. When a file is highlighted, the file name is displayed in the Restore as field.


Read File Configuration Dialog

Some of the data used to populate the files tab of the database restore properties dialog is obtained from new client index file-configuration meta-data objects that are created for v3.0. For backups created with a pre-v3.0 product, the file-configuration meta-data will not be present in the client index, but is available on the backup tape (or whatever media was used to store the backup). The Read File Configuration dialog box appears when you open the properties dialog for a marked database item which has no file-configuration meta-data in the client index; you get the option to load this information from the saveset media.

Figure 19. Read File Configuration Dialog

The options are:

- **Read the file configuration from save set media** (optional): If you choose to read from tape, a "Reading..." dialog appears and the tape is read in the background. When the reading is finished, the "Reading" dialog goes away and the properties dialog is invoked; the filegroup and destination table will have valid data.

Figure 20. Reading Dialog

The "Reading" dialog has a cancel button, which allows canceling the reading. If you cancel from the "Reading" dialog, the properties dialog will be invoked, but the filegroup and destination table will be empty.

- **Skip reading the file configuration** (optional): If you elect to skip reading in the "Read File Configuration" dialog, the properties dialog will be invoked, but the filegroup and destination table will be empty.

6

Database File Relocation Restrictions

The NetWorker Module contains the following restrictions regarding database file relocation:


- Only databases can be relocated. Individual filegroups, files, and transaction logs 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 support creation of a new database, the relocate procedure will fail.

Database Item Restore Time Properties

The Restore Time tab of the Database Item Restore Properties dialog box provides access to the restore time features in NetWorker Module 3.0 including restoring a specific database backup version backup version or to a specific transaction time. In addition, the ability to restore to a named log mark is *also* supported for SQL Server 2000.

Figure 21. Database Item Restore Time Properties



The “Specify a restore time” checkbox is unchecked by default. If the “Back up the active portion...” checkbox on the General tab is checked, the checkbox on this Restore Time tab will be disabled — attempting to check this box will result in an error message indicating why it cannot be checked. You cannot perform a transaction log backup and also restore to a specific time.

Checking the “Specify a restore time” box enables the other Restore-time controls:

- Using a specific backup version (Optional) — (default = greyed; if the checkbox is checked, this radio button is selected by default). Select this option to restore to a specific time. When this option is selected, the Specify a time to perform a point-in-time restore field, Point-in-time... button, and Backup Versions List are enabled.
 - Specify a time to perform a point-in-time restore (Optional) — (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 will be reinstated during the restore, and when to stop restoring transactions. If an incremental backup is marked in the Backup Versions list, this field 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, allowing you to set the time for a point-in-time backup. For details refer to [“Point-in-Time Restore Dialog” on page 89](#). When any other backup level is marked, this button displays an error message.
- Backup Versions List — (default = most recent backup version is marked) Use this list to select and mark the backup version to use for the restore. The “Specify a time...” field displays the date and time for the backup that is currently marked. *If you mark a non-incremental backup version from this list, the Point-in-time... button will be disabled. Only one backup version can be selected and / or marked.*

The Size column displays the size of the backup. The Backup 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 level incremental backup version, click on any backup in the list, and then click the Mark / Unmark button. When a backup version is marked, a checkmark is displayed next to the Backup Size in the list. If another backup version is already marked, the selected backup version will be checked and the other, marked backup version will be unchecked.

Note: You can also mark a backup version by double-clicking on a row in the list.

- Using a named log mark (Optional) — (default = not selected) Select this option to perform the restore using a named log mark. *This option is only visible in SQL Server 2000 and only enabled when log marks exist for the selected database backup.* When this option is selected, the Restore to... radio buttons and Log Marks List are enabled.
 - (Required) Restore to (the end, the beginning) of the log mark — (default = the end) Use this option to specify which type of named log mark restore you would like 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.
- Log Mark List (Required) — (default = none checked) Use this option 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 which was entered when the transaction was performed.

To select a log mark, click on the log mark in the list and click the Mark/Unmark button. When a log mark is marked, a checkmark is displayed next to the log mark name in the list. If another log mark is already marked, the selected log mark will be checked and the other, marked log mark will be unchecked.

Note: You can also mark a log mark by double-clicking on a row in the list.

Point-in-Time Restore Dialog

The Point-in-Time Restore dialog contains properties for specifying a date and time to be used for the SQL point-in-time restore. The date and time that can be specified must fall within a time frame spanned by the transaction log backup that was marked in the backup version list. The upper bound of the time frame is defined by the create time of the transaction log backup. No date and time greater than this upper bound can be specified. The lower bound of the time frame is the create time of the most recent level incr 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.

Figure 22. Point-in-time Restore Dialog

The options on the Point-in-time Restore dialog are as follows:

- (Informational) TBD — This option is provided for informational purposes only and cannot be modified.
- (Required) Date — (default = date when the marked backup version was created) Use this option 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 down 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 option 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.


FileGroup Restore Properties

These properties allow you to specify preferences for a single, *marked* filegroup before you start a normal restore-type operation. *These properties are not available for partial, copy, or verify only restore-type operations.* The Filegroup Item

Restore Properties dialog box appears when you select a *marked* filegroup from either pane of the Restore Browse window and select the shortcut menu>Properties, or click the Properties toolbar button. These properties are new in NetWorker Module 3.0 and are context sensitive, based on which version of SQL Server you are running (7.0 or 2000).

Preferences that you specify for these properties will be used for the current restore operation. When a preference is not specified, the NetWorker Module default will be used. After the restore operation has completed, and you exit from the Restore Browse window, the property values will revert back to the NetWorker Module defaults.

Figure 23. Filegroup Item Restore General Properties



The General tab of the dialog box features basic properties for configuring the restore operation. SQL Server requires both options when restoring filegroup backups. (Both options are checked, but grayed-out.)

The first option signifies that the active portion of the transaction log is backed up prior to performing a filegroup restore, so that the log can be applied to the filegroup to make it consistent with the rest of the database.

The second option forces SQL Server to ignore errors produced on behalf of nonexistent filegroups due to media failure. If there is a media failure, then the filegroup files no longer exist.

The options on this tab are as follows:

- (Information) Backup the active portion of the transaction log before restoring the filegroup — (default = checked) This option is provided for informational purposes only and cannot be modified. NetWorker Module will attempt a transaction log back up using the NO_TRUNCATE SQL keyword. The restore operation will proceed regardless of whether the transaction log backup succeeds or fails.
- (Information) Overwrite the existing filegroup with the restored file — (default = checked) This option is provided for informational purposes only and cannot be modified. NetWorker Module will specify the WITH

REPLACE SQL keyword in the restore sequence. This causes the restore operation to restore the filegroup to the exact location (drive and pathname) as the location on the source host from which the data was backed up.

How to Ensure Proper Filegroup Restoration

When restoring a file or filegroup, one must first backup the active portion of the transaction log. This is necessary because the transaction log backup must be applied to the filegroup restore to ensure the filegroup is consistent with the rest of the database. If a file or file is restored from NetWorker User, this transaction log backup will occur automatically. When using the command-line program to restore the 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 will instruct 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 due to the fact that, 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 filegroup.

1. Put the database into single-user mode.
2. Immediately perform *another* transaction log backup (which will include the backup in step 2 as a transaction).
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.


File Restore Properties

These properties allow you to specify preferences for a single, *marked* file before you start a normal restore-type operation. *These properties are not available for partial, copy, or verify only restore-type operations.* The File Item Restore Properties dialog box appears when you select a *marked* file in the

Restore Browse window and select the shortcut menu>Properties, or click the Properties toolbar button. These properties are new in NetWorker Module 3.0 and are context sensitive, based on which version of SQL Server you are running (7.0 or 2000).

Preferences that you specify for these properties will be used for the current restore operation. When a preference is not specified, the NetWorker Module default will be used. After the restore operation has completed, and you exit from the Restore Browse window, the property values will revert back to the NetWorker Module defaults.

Figure 24. File Item Restore General Properties



The General tab of the dialog box features basic properties for configuring the restore operation. SQL Server requires both options when performing a file restore. (Both options are checked, but grayed-out.)

The first option signifies that the active portion of the transaction log is backed up prior to performing a file restore, so that the log can be applied to the file to make it consistent with the rest of the database.

The second option 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.

The options on this tab are as follows:

- (Information) Backup the active portion of the transaction log before restoring the file — (default = checked) This option is provided for informational purposes only and cannot be modified. NetWorker Module will attempt a transaction log back up using the NO_TRUNCATE SQL keyword. The restore operation will proceed regardless of whether the transaction log backup succeeds or fails.
- (Information) Overwrite the existing file with the restored file — (default = checked) This option is provided for informational purposes only and cannot be modified. NetWorker Module will specify the WITH REPLACE

SQL keyword in the restore sequence. This causes the restore operation to restore the file to the exact location (drive and pathname) as the location on the source host from which the data was backed up.

Chapter 7: Special Features

This chapter provides information about several NetWorker Module features that pertain to both backup and restore operations:

- Checking database consistency
- Using striping in backup and restore operations
- Using the RAW save set to improve striped backup performance
- Using interleaving in backup and restore operations
- Backing up and restoring multiple instances of Microsoft SQL Server
- Backing up and restoring SQL Server running in a Microsoft Cluster Server environment

Checking Database Consistency

If you are running Microsoft SQL Server 7.0, you can use the NetWorker Module -j command option 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.

The option to run a DBCC is disabled by default. In NetWorker Administrator, to trigger a DBCC before each time a scheduled save is initiated, go to the appropriate client resource and in the Edit>Create Client dialog box>Remote

tab>Backup Command text box, enter `nsrsqlsv -j`. To precede a manual backup with a DBCC, log on to the client host with administrator privileges and enter `nsrsqlsv -j` at a Windows command prompt.



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

If you enable this option and the DBCC is completed successfully, NetWorker Module proceeds with a backup of the specified database(s). If the DBCC is not completed successfully, the backup is terminated and the NetWorker Module displays a message either to:

- The display indicating DBCC success or failure
- The Group Detail dialog box in NetWorker Administrator

You can also check database consistency on a restored database 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 NetWorker User for SQL Server, the client user interface.

For more information about `nsrsqlsv` and `nsrsqlrc` command options, see [“Appendix A: Commands and Command Options” on page 115](#). For more information on database consistency checking, refer to the Microsoft SQL Server documentation.

Using Striping

Legato NetWorker Module 3.0 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.

To specify striping:

- From NetWorker User for SQL Server, specify striping and the number of stripes you want in the Backup Options dialog box.
- For a scheduled backup, in NetWorker Administrator go to the appropriate client resource and in the Edit>Create Client dialog box>Remote tab>Backup Command text box append `-Sn` to the `nsrsqlsv` command, where `n` is the number of stripes you want to use.
- 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 115](#).

Configuring Striping for Backups

Before any striped backup begins, in the Set Up Server dialog box in NetWorker Administrator *you must set Parallelism to at least one more than the number of stripes you will be using*. For example, if you are using 3 stripes, specify a value of 4 or greater for the parallelism.

Performance Considerations for Striping

For improved backup performance, follow these suggestions:

- In the Create>Edit Device dialog box (in the Devices resource), set Target Sessions per device to 1. Always specify one session per device when using striping. This will yield 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 your striped backups.
- NetWorker Module automatically restores your data using the same number of stripes you specified for the data backup. For best restore performance, ensure that the same number of devices used for the backup is also available at restore time.

- NetWorker Module operation allows you to run NetWorker server on a different host from the computer that hosts NetWorker Module, the NetWorker client, and Microsoft SQL Server. *However, to optimize striping performance, it is best to have your NetWorker server on the same host as these three other components, and to have your storage node local to this host as well.*

For additional information about restoring backups, see [“Chapter 5: Backing Up SQL Server Data” on page 63.](#)

How to Perform a Striped Backup

Legato NetWorker Module supports backup striping in two ways:

1. From the SQL Server host, which is a NetWorker client, initiate a manual backup in either of the following ways:
 - a. 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.

- b. Click the Backup button in NetWorker User for SQL Server 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 8 on page 67](#)).

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 115](#). For details about using NetWorker Administrator, refer to the *Legato NetWorker Administrator’s Guide*.

2. From the NetWorker server host:

Use NetWorker Administrator to schedule a striped backup by specifying the `nsrsqlsv -S stripes_count` command in the Clients resource>Remote Tab>Backup Command text box.

Using the RAW Save Set

NetWorker Module 3.0 introduces support for use of the *RAW* save set when creating a striped backup. The RAW save set is *only* available for configurations running NetWorker Module 3.0 and NetWorker 6.0 Power Edition or higher.

When the RAW save set is used during creation of a striped backup, CPU utilization may be lowered. Performance enhancements can be achieved when a large number of output devices are directly connected to the SQL Server host.



Important: *It is highly recommended that you only use the RAW save set in high performance device configurations.*

The following restrictions apply when using the RAW save set:

- Only available for configurations running NetWorker Module 3.0 and NetWorker 6.0 Power Edition or higher.
- The NetWorker client and the NetWorker storage node must be installed on the same computer.
- RAW save set backups are not portable.
- The RAW save set cannot be used for FILE type NetWorker devices.

Using Interleaving

Legato NetWorker Module supports *interleaving* for backups of multiple clients to a single backup device.



Important: Although NetWorker Module 3.0 supports striped interleaving, interleaving stripes is not recommended. Interleaved backups can require substantial restore time. This is because, when you restore an entire interleaved backup (including all the clients backed up), the restore operation occurs for one client at a time, thus requiring multiple passes to be made over the same areas on the backup medium.

However, interleaving can be useful when you want to back up multiple clients concurrently to one device. When you restore an interleaved backup, you can restore each client's backup separately, or you can restore all the interleaved client backups at the same time.

Backing Up and Restoring Multiple Instances of SQL Server

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

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 3.0 provides support for back up and restoration of SQL Server 7.0 and SQL Server 2000 version default instances.

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

`computerName\instanceName`

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

The syntax for specifying an 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 [“Command Line Notation and Syntax” on page 25](#).

The following restrictions apply when running multiple instances:

- A maximum of 16 named SQL Server instances can be run simultaneously on a single computer.
- **nsrsqlsv** and **nsrsqlrc** only support specification of a single instance. If save sets for more than one instance are specified, the backup or restore operation will fail.
- A named instance may not be the source for a copy (directed recover) operation when the destination server is different from the source.
 - **nsrsqlrc** will fail if a directed recover from one named instance to another SQL Server instance is attempted.

- Index entries for 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 name.
 - NetWorker User for SQL Server 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.

Backing Up and Restoring Microsoft Cluster Server Configurations

This section explains the NetWorker Module 3.0 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: If you are using NetWorker Module to back up and restore SQL Server data on a node in a cluster, you must be using NetWorker Power Edition on your NetWorker server host.

For information about configuring the Legato NetWorker Module to run in a cluster environment, see the *Legato NetWorker Module 3.0 for Microsoft SQL Server Installation Guide*.

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 that are on private disks attached to that node. A client associated with a virtual server backs up only files on disks that are 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 3.0 Support for MSCS Clusters

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

Legato NetWorker Module uses the virtual server name to:

- Connect to the appropriate SQL Server
- Construct the appropriate named-pipe names to accept data from, or deliver data to, the appropriate SQL Server in the cluster
- Create entries in the NetWorker client index

Legato NetWorker Module 3.0 creates entries under the virtual server name in the NetWorker client index. Information about the SQL Server save sets of a virtual server is stored in the NetWorker client index associated with that virtual server name. NetWorker media database or index database records are indexed according to the client that performs a particular backup.

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 secondary named instances, where each secondary instance name must be unique within the cluster.

The NetWorker Module can support up to 16 named instances to act 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

- Data files installed on a clustered drive running in the same cluster group as the SQL virtual server

Configuring Scheduled Backups of a Clustered SQL Server

In NetWorker Administrator, create a client resource, specifying the name of the virtual server as the client name.



Important: With NetWorker Module 2.0, it was necessary to enter -a *virt_svr_name* in the Backup Command textbox (in the Remote Tab of the client resource). This is not necessary with NetWorker Module 2.0.1 or later. When the scheduled backup begins, the NetWorker **savegrp** command automatically passes the virtual server name to the NetWorker Module save program *nsrsqlsv*. The NetWorker Module save program automatically determines if the client name passed to it is a virtual server name; if it is, NetWorker Module uses this name to contact the virtual server and creates entries having this name in the NetWorker client index.

Performing Unscheduled Operations on a Clustered SQL Server

In a cluster environment, you can initiate an unscheduled NetWorker Module operation:

- From a Windows command prompt or from NetWorker User for SQL Server
- From either 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_svr_name -s NetWorker_server_hostname db_name
```

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

Operations from NetWorker User for SQL Server

In NetWorker User for SQL Server, the Connect to Server menu item on the Operation menu opens the Select SQL Server dialog box. Here you can select or change the SQL Server host for a backup or restore operation in a cluster environment. The Connect to Server menu item is enabled in the main NetWorker User for SQL Server window only when multiple, individual SQL Server hosts are detected. This menu item is not present when a Backup or Restore Browse window is open.

When running in a Microsoft Cluster Server (MSCS) cluster, NetWorker Module 3.0 automatically detects all active SQL Servers in the cluster, including virtual servers. This automatic detection occurs at several points:

- On NetWorker User for SQL Server startup
- Whenever you select the Connect to Server menu item
- Whenever you initiate a backup or restore

Active/Passive Cluster Configurations

When NetWorker User for SQL Server 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 NetWorker User for SQL Server on the secondary node also automatically uses the virtual server name.

Active/Active Cluster Configurations

In this situation, normally there is one SQL Server and 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, NetWorker Module automatically communicates with the virtual server running on the same physical node as the NetWorker Module software, and the Operation>Connect to Server menu item 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, on startup NetWorker User for SQL Server automatically presents the Select SQL Server dialog box containing a listbox of the SQL Server hosts. You must select a host from this dialog box, after which NetWorker Module will attempt to validate the host as a NetWorker client.

NetWorker Module Interaction with SQL Server Services

When restoring certain SQL Server system databases (namely, the *master* database or *msdb* database) in a clustering environment, Legato NetWorker Module automatically stops and then restarts the SQL Server services as appropriate, as follows:

1. Before the restore begins, Legato NetWorker Module takes the SQL Server virtual server resources offline.
2. Legato NetWorker Module starts the local SQL Server service in single-user mode.
3. Legato NetWorker Module performs the restore.
4. After the restore is completed, Legato NetWorker Module stops the local SQL Server service.
5. Legato 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 115](#).

Chapter 8: Disaster Recovery

You can readily restore lost SQL Server data if you have been using NetWorker Module for regular backups. Following are basic instructions for disaster recovery if Legato NetWorker Module binary files or Microsoft SQL Server binary files are damaged or lost. For more information about using NetWorker for disaster recovery, refer to the *Legato NetWorker Disaster Recovery Guide* and the Microsoft SQL Server documentation.



Important: Because there are so many possible variables, it is not possible to provide exact disaster recovery instructions for every situation. The examples in this chapter are designed to 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 is indicated in this document. Therefore, it is strongly recommended that you read the *Legato NetWorker Disaster Recovery Guide*, the *SQL Server Books Online* documentation, and the *Legato NetWorker Release Supplement* for the version of NetWorker that you are running, before you begin the SQL Server disaster recovery.

The following sections 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.

8

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 when the ALL backup version was created
- 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.



Important: The CD-ROM or shared network drive containing the SQL Server installation binaries is required to rebuild the master database.

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

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

1. Shutdown 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, you must copy the database files to a temporary directory on the SQL Server. If you are running SQL Server 2000 under Windows 2000, you must remove the read-only attributes from the database files 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 NetWorker Module for SQL Server.
5. Restore SQL Server master database and msdb database with the replace option. If you are running SQL Server version 7.0, you may need to restore the model and the distribution databases.
6. Restore the other SQL application databases.

8

Disaster Recovery of Microsoft SQL Server Disk Crash or Data Corruption

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

1. Shutdown 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 108](#).
3. Restart your computer.
4. Shutdown 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, you must copy the database files to a temporary directory on the SQL Server. If you are running SQL Server 2000 under Windows 2000, you must remove the read-only attributes from the database files as well. The temporary directory is also the source directory for data files 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 NetWorker Module for SQL Server.
8. Restore SQL Server master database and msdb database with the replace option. If you are running SQL Server version 7.0, you may need to restore the model and the distribution databases.
9. Restore the other SQL application databases.

Disaster Recovery of Microsoft SQL Server and NetWorker Server

If the disk(s) 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 the 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 108](#).
4. Restart your computer.
5. Shutdown 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, you must copy the database files to a temporary directory on the SQL Server. If you are running SQL Server 2000 under Windows 2000, you must remove the read-only attributes 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 NetWorker Module for SQL Server.
9. Restore SQL Server master database and msdb database with the replace option. If you are running SQL Server version 7.0, you may need to restore the model and the distribution databases.

8

10. Restore the other SQL application databases.

Disaster Recovery of Microsoft SQL Server (Shortcut)

If you can determine that the Microsoft SQL server does not need to be restored (see [“When to Reinstall Microsoft SQL Server” on page 108](#)) then the following steps can be used 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 the 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 108](#).
4. Restart your computer.
5. Shutdown 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, you must copy the database files to a temporary directory on the SQL Server. If you are running SQL Server 2000 under Windows 2000, you must remove the read-only attributes 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 NetWorker Module for SQL Server.
9. Restore SQL Server master database and msdb database with the replace option. If you are running SQL Server version 7.0, you may need to restore the model and the distribution databases.

10. Restore the other SQL application databases.



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 73](#).

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

Disaster Recovery of Microsoft 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 the failed was a NetWorker client, reinstall the NetWorker client software.
3. Reinstall the SQL Server software.
4. Start SQL Server.
5. Restore SQL Server master database and msdb database with the replace option. If you are running SQL Server version 7.0, you may need to restore the model and the distribution databases.
6. Restore the other SQL application databases.

8

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 option(s) and variables, and the appropriate syntax for specifying the SQL Server data to use for the operation.



Important: Although the command line syntax examples provided in the following sections are presented using a multiple line format, they are simply presented in this manner to improve the readability of the command. When you actually enter a command at the command prompt, the command, command options, and specified data must all be entered on a single line.

Using the nsrsqlsv, nsrsqlrc, and nwmssql Commands

Legato NetWorker Module commands perform the following functions:

- nsrsqlsv: Back up the specified SQL Server data objects. For more information, see [“Using the nsrsqlsv Command” on page 121](#).
- nsrsqlrc: Restore the specified SQL Server data objects. For more information, see [“Using the nsrsqlrc Command” on page 116](#).
- nwmssql: Invoke the client graphical user interface. For more information, see [“Using the nwmssql Command” on page 125](#).

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:

A

- The set of command options consists of lower and uppercase letters. However, 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
-1backup_level
 while the following expression is invalid because a space is not allowed between the **+** argument and **log_mark_name**:
-M + log_mark_name
- Square 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 the square brackets.

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.

nsrsqlrc Command Options

```
nsrsqlrc [-fjqV][-a virtual_server_name][-c client_host_name]
{[-d MSSQL:destination_dbName
  [(file=path,file2=path2,...)] MSSQL:source_dbName]}
[-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 ...]
[-Dn]
```

- 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 101](#).

- c** Specifies the NetWorker `client_host_name` from which the SQL Server data is to be restored.
- 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 location or another database name on the same client host. The syntax is:

```
nsrsqlrc -s NetWorker_server_hostname
-d [""]MSSQL:destination_dbName:(relocation_list) [""]
      MSSQL:source_dbName
```

`source_dbName` is the name of the SQL database to restore.
`destination_dbName` is the name of the SQL database to which the source database is to be restored. `relocation_list` specifies new locations for restored database files. The relocation list is appended to `destination_dbName` by a colon. Use of a space separator before or after the colon is not allowed.

The relocation list is composed of pairs of logical database filenames and fully qualified database filename relocation paths. Each filename and relocation path is joined by the = character and pairs are delimited in the list by commas. The entire relocation list is enclosed in parenthesis. 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:NewProjectName:
      ( MasterFile=C:\Relocation\Path\MasterFile,
        NextFile=C:\Relocation\Path\NextFile,...,
        LastFile=C:\Relocation\Path\LastFile ) MSSQL:Project
```

The set of parameters supplied to **-d** may be enclosed by double-quotes to allow for spaces in the relocation path names. 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 NetWorker from parsing the single-quote as a delimiter. For example:

```
nsrsqlrc -s NetWorker_server_hostname
-d "MSSQL:NewProjectName:
      ( 'Master File'=C:\Relocate Dir\Path\Master File,
```

A

```
NextFile=C:\Relocate Dir\Path\NextFile', ...,  
'Last=C:\Relocate Dir\Path\'Last) MSSQL:Project"
```

When no relocation list is specified, NetWorker Module reads the source database file names and location from 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. For more information about performing a copy operation, see [“Chapter 3: NetWorker Module Restore Concepts” on page 45.](#)

A

- 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.
- 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 95.](#)
- M** Performs a SQL Server data restoration of the named transaction specified in `log_mark_name` (Microsoft SQL Server 2000 only). When the `+` parameter is supplied, the data is restored to and including the named transaction. When the `-` parameter is present, the data is restored up to the named transaction. The `log_mark_name` should immediately follow the `+` or `-` symbol, using a space separator is not allowed. The default is `+`.

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
```

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 must follow all other command options and parameters on the command line. 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 [“Command Line Notation and Syntax” on page 25.](#)

When you specify **MSSQL:**, start SQL Server in single user mode to avoid failures caused by other users accessing the database(s) during restoration.

If **-M** is not specified, then normal mode is used by default.

- 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 of the SQL Server user name and password invokes the restore command 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** Use the upper-case **R** option to perform 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:

```
nsrsqlrc -s NetWorker_server_hostname
-R ["']fgName['],[']fgName2['],[']...['']["]
-d MSSQL:destination_dbName MSSQL:source_dbName
```

source_dbName is the name of the SQL database to restore.
destination_dbName is the name of the SQL database to which the source database is to be restore. **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 NetWorker 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:

A

A**-R ""**

NetWorker Module *only* restores the PRIMARY filegroup. For more information about partial restore, see [“Chapter 3: NetWorker Module Restore Concepts” on page 45](#).

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

S Use the upper-case **S** option to perform the restoration according to the specified database restore mode of normal, standby, or norecover (Microsoft SQL Server 2000 only). The syntax is:

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

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. **norecover** 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.

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

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:Project MSSQL:NewProjectName
```

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:Project MSSQL:NewProjectName
```

For more information about using the database restore modes, see [“Database Restore Properties” on page 78](#).

t The lower-case **t** option 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 of the SQL Server user name and password invokes the restore command using SQL Server integrated security.

V Use the upper-case **V** option to perform a verification of 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
```

Dn This is used to specify a debug message level.

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:

- NetWorker Administrator>Clients resource>Edit Client dialog box>Remote Tab>Backup Command text box.
(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 you are initiating a backup from the Backup Command text box in NetWorker Administrator.*)
- At a Windows command prompt.

nsrsqlsv Command Options

```
nsrsqlsv [-CGjqRTv] [-a virtual_server_name] [-b pool_name]
[-c client_host_name] [-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 ...*]

- 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 101](#).
- b** Assigns a destination pool for a save set. Specifying `-b pool_name` overrides all other pool-selection criteria provided by NetWorker or specified in the NetWorker User for SQL Server Backup Options Dialog box. The pool must be created, with a corresponding label template, before you run a command that includes the `-b` option. For more information about pool-selection criteria, see [“Volume Pools” on page 36](#).
- c** The lower-case `C` option specifies the NetWorker `client_host_name` for which the SQL Server data is to be backed up.
- C** Use the upper-case `C` option to compress the backup data before NetWorker client sends the data to NetWorker server.
- g** The lower-case `g` option specifies the save group. NetWorker server and the `savegrp` command use the `group` parameter to select the media pool.
- G** Use the upper-case `G` option to specify a `NO_LOG` transaction log backup before backing up the database. This command option is valid *only* for backups of level `full`.
- j** Performs a database consistency check before initiating the backup. For more information about this option, see [“Chapter 7: Special Features” on page 95](#).
- l** Specifies the backup level. Valid values are `full`, `diff` (equivalent to specifying level `1` 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 versions SQL Server data objects, see [“Backup Levels” on page 27](#).

A



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 of the Schedules resource in NetWorker Administrator. Do not use the `-l` option when you are initiating a backup from the Backup Command textbox in the Remote Tab of the NetWorker Administrator Edit Client dialog.

A

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 must follow all other command options and parameters on the command line. The list of data objects must follow all other command options and parameters on the command

A

line. For detailed guidelines on formatting these parameters, see [“Command Line Notation and Syntax” on page 25](#).

- 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 of the SQL Server user name and password invokes the restore command using SQL Server integrated security.
- q** Displays nsrsqlsv messages in quiet mode: only summary information and error messages are displayed.
- r** The lower-case **r** option performs a backup using the RAW save set. When performing a striped backup with five or more output devices directly connected to the SQL Server, using the RAW save set may lower CPU utilization.
- R** Uses NO_TRUNCATE option when backing up transaction logs.
- s** The lower-case **S** option specifies the NetWorker server to use for the backup operation. *Use of this command is highly recommended.*
- S** Use the upper-case **S** option to back up the specified data item(s) using *n* stripes. NetWorker Module then automatically uses the same number of stripes for the corresponding restore operation.



Important: To use backup and restore striping successfully, you must use the striping configuration described in the section on [“Using Striping” on page 96](#).

- T** The upper-case **T** option 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 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 of the SQL Server user name and password invokes the backup command using SQL Server integrated security.
- v** Use the lower-case **v** option to display **nsrsqlsv** messages in verbose mode, providing detailed information about the progress of the backup operation.

Using the nwmssql Command

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

- NetWorker Group>NetWorker Module for Microsoft SQL Server from the Windows Start menu.
- 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** The lower-case **s** option specifies the NetWorker server to use. *Use of this command is highly recommended.*
- t** The lower-case **t** option provides the ability to diagnose the current back up or restore command issued by NetWorker Module. NetWorker User for SQL Server displays the full back up or restore command in the operation status window, but does not execute the back up or restore operation.

A

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 have the capability of logging on to a domain from a workstation or a server, backing it up, and restoring the data. Backup Operators also can shut down servers or workstations.
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) backups.

backup type	The type of backup to perform using NetWorker Module. Backup type 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 NetWorker after a disk crash.
browse policy	A NetWorker policy that determines how long entries will be maintained in the online index and thus be restorable.
client	A computer that accesses the NetWorker server to back up or recover files. Clients may be workstations, PCs, or file servers.
command line	The shell prompt, where you enter commands.
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.
directed recover	In NetWorker Module 3.0, the <i>copy restore</i> type replaces the directed recover operation, which existed in previous versions of BusinessSuite Module for Microsoft SQL Server.
directive	An instruction directing NetWorker 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 object	NetWorker Module 3.0 introduces the ability to browse the filegroups and files contained in a database. To support this new behavior, NetWorker User for SQL Server may display certain data items which are contained in the storage hierarchy, but are not available for backup, these are referred to as <i>fake objects</i> .
file index	A database of information maintained by NetWorker that tracks every file or filesystem backed up.
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.
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 mark” on page 130 .

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 maintained by NetWorker that tracks every backup volume.
multiple-instance	Microsoft SQL Server 2000 introduces the ability to run multiple copies of SQL Server on a single computer. This feature is called <i>multiple instance</i> . Multiple instances allow any number of SQL Servers to run simultaneously on a single computer.
named log mark	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 to back up and recover 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 recover 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.
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 NetWorker.
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> tells you how to find a file beginning at the root directory and working down the directory tree. A <i>relative pathname</i> tells you how to find the file starting where you are now.
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.
print	Send data to a printer.
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” on page 132 .
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” on page 132 .
restore	The NetWorker Module command used to browse the SQL Server storage hierarchy and to restore files from a backup version. See “recover” on page 131 .
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 will be retained in the media database and thus be 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 NetWorker.
save set ID	An internal identification number assigned to a save set by NetWorker.
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.
stripe	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 NetWorker.
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 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 NetWorker from his or her workstation to back up and recover files.
verify only restore	Only verify the backup media for the selected SQL Server 7.0 or later backups. <i>Selecting the verify only restore-type will NOT restore your 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 NetWorker.
volume name	The name you assign 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” on page 133 .

Index

A

ad hoc backup (*See* on-demand backup)
Autostart 35
Autostart option 35, 42

B

backup
 automatic (*See* backup, scheduled)
 cancelling 66
 clone 16
 creating 64
 e-mail notification 42
 interleaving 99
 level 32
 levels
 comparison 32
 described 31
 manual (*See* backup, unscheduled)
 multiple instances of SQL 100
 on demand 22
 options 66
 Compress the backup content 67
 Create a striped backup 67, 97
 Use pools for media management 68
 Use the RAW save set 68, 99
 parallelism 97
 performance 16, 96
 preparing for 63
 scheduled 21, 33
 schedules 32
 scheduling 32
 striping 98
 time 66
 unscheduled 21, 70

 using NetWorker User for SQL Server 65

 viewing results 42
 volume 66

backup group(s) 36–37

backup level 37
 incremental 53

backups

 ad hoc 22
 lights out 16
 manual 17, 22
 on demand 22
 options 35
 scheduled 16
 Autostart option 35, 42
 unscheduled 16

bootstrap 24, 36

browse policies 40–41

BusinessSuite Module

 features highlighted 16
 restore, graphic of 19

C

cancelling a backup 66
cancelling a restore 78
client 17, 37
 resource 97
client file index 17, 24, 40–41
cluster virtual server 104
computerName 100
create time 52
customize
 backup schedule 32

D

data 15

database and log files 77
 database consistency check (DBCC) 95–96
 DBCC (*See* database consistency check)
 disaster recovery 107

E

e-mail
 "savegroup completion" notification 42
 notification(s) 43

F

fake object 59
 figures
 Backup Browse Window 65
 File Item Restore General Properties 93
 NetWorker process during BusinessSuite
 Module restore 19
 NetWorker process during NetWorker
 Module backup 18
 NetWorker User for SQL Server 64
 Operation>Restore>Restore-Type
 Menu 74
 Restore Browse Window 76
 Restore-Type Selection Dialog 74
 SQL Server Host 64
 SQL Server Selection Dialog 75
 filegroups
 backing up and restoring 92
 filesystem backup 16

G

global browse time 53

I

Installation Guide, typefaces and symbols
 explained 12
 instanceName 100
 interleaving 99

L

label template 37, 71
 legacy database options
 select into/bulk copy 48
 trunc.log.on.chkpt 48
 LOAD command 19

M

media
 database 17, 19, 40–41
 message log files 42
 Microsoft Cluster Server
 How NetWorker works with MSCS 101
 Microsoft Cluster Server (MSCS) 16, 104
 active/active cluster configurations 104
 active/passive cluster configurations 104
 virtual server 104
 MSCS (*See* Microsoft Cluster Server)
 multiple instance 100, 102

N

named instances 100–102
 Named Pipes 19
 NetWorker
 client 17, 37
 client resource 97
 client, description of 41
 features highlighted 15
 NetWorker Remote Exec Service 73
 NetWorker Module
 backup concepts 21
 backup, graphic of 18
 failover cluster support 102
 support for MSCS clusters 102
 NetWorker Module Database Restoration 55
 NetWorker User for SQL Server 57, 70

- Backup Browse window 58
 - Database Item Backup Properties 69
 - Browse window 57
 - display conventions 59
 - fake object 58
 - Fake File 58
 - Fake Filegroup 58
 - SQL Server storage hierarchy 57
 - Restore Browse window 58, 73, 93
 - Database Item Restore Properties 78, 93
 - General tab 93
 - NetWorker XBSA API 19
 - notification(s) 17
 - nsrd
 - functionality 17, 19
 - nsrmmmd
 - functionality during restore 19
 - nsrmmdbd 19
 - nsrsqlrc 18–19, 100, 103
 - nsrsqlsv 17, 70, 92, 97–98, 100, 103
- O**
- on-demand backup 22, 64–65
 - online indexes 40
 - options
 - Autostart 35
- P**
- parallelism 97
 - planning
 - using backup levels 31
- R**
- RAW save set 99
 - restrictions 99
 - recovery model
 - BULK_LOGGED 46–47
 - changing 48–50
 - FULL 46–47
 - SIMPLE 47
 - specifying 48
 - restore 19
 - backup levels 32
 - cancelling 78
 - database objects 73
 - multiple instances of SQL 100
 - point-in-time 16, 52
 - preparing for 73
 - properties 77–78
 - Backup the active portions... 93
 - Overwrite the existing file... 93
 - server data 45
 - SQL cluster resources 56
 - time 77
 - transaction log backups 55
 - uncommitted transactions, 55
 - using NetWorker User for SQL Server 74
 - restore concepts 45
 - restore mode 51
 - no recovery 45, 52
 - normal 45, 52
 - standby 45, 52
 - restore time 52
 - backup versions 45
 - database backup versions 53
 - log marks 45
 - named log mark(s) 53–54
 - point-in-time 45, 53
 - specifying 54
 - restore-type 45, 73–74
 - copy 51, 75–76
 - normal 50–51, 73
 - partial 51, 76
 - verify only 51, 73
 - retention policies 40–41
- S**
- save
 - functionality 17

save set(s) 37, 41
 savegroup completion report 17
 savegrp 17, 103
 functionality 17
 scheduled backup(s) 17
 schedules
 levels 31
 select into/bulk copy 48
 SmartMedia 16
 SQL Server
 cluster
 command-line operations 103
 operations from NetWorker User for
 SQL Server 104
 computerName 100
 database consistency check (DBCC)
 95–96
 Enterprise Manager 55
 instanceName 100
 Master Database Maintenance 55
 named instances 100–101
 services 56
 multi-user mode 56
 single-user mode 56
 virtual server 104
 SQL Server 2000 16
 multiple instance 100
 running multiple instances 100
 SQL Server Master Database Maintenance
 55
 SQL Server NORECOVER option 55
 storage management 15–16, 18
 stripes
 defined 96
 option 67
 performance considerations 97
 performing a striped backup 97–98
 specifying a value 97

T

transaction log backup 52
 transaction logs 15

trunc.log.on.chkpt 48

V

Virtual Device 19
 volume
 label templates 71
 volume pools 36, 66, 68, 70
 data pool 70–71
 log pool 70–71
 target pool 72
 volumes
 backup levels 31

X

XBSA (*See* X-Open® Backup Services
 Application Programming Interface)
 X-Open® Backup Services Application
 Programming Interface 17