

Legato NetWorker[®] Module for Sybase[®]

Administrator's Guide

**Releases 1.1 and 2.0
Windows NT[®] and UNIX[®] Versions**



© 2000, 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 Sybase Administrator's Guide

August 2000

22-2307-01

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
About This Guide	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
Importance of Backing Up Mission-Critical Data	15
Sybase Server	15
Legato NetWorker	16
Legato NetWorker Module for Sybase	16
How NetWorker Connects to the Database Server	16
Chapter 2: Scheduled Backups	19
Required Sybase Roles and Permissions	20
Using a Backup Group	20
What Is a Backup Group?	21
Suggestions for Setting Up NetWorker Module for Sybase Groups	21
Using NetWorker Backup Schedules	22
What Is a Full NetWorker Backup?	23
What Is an Incremental NetWorker Backup?	23
When are Incremental Backups Not Allowed?	23

Using Volume Pools.....	23
What Is a Volume Pool?	23
Customizing Volume Pools.....	24
Example: Customize Media Organization	25
Creating a Volume Label Template.....	25
How to Create a Volume Pool.....	26
Using NetWorker Policies	26
Configuring a Database Server as a NetWorker Backup Client.....	27
What Is a NetWorker Client?	27
Creating a NetWorker Client.....	27
Specifying a Save Set.....	28
How to Send an E-Mail Notification after a Scheduled Backup	29
Viewing the Results of a Backup.....	29
Using the nsrinfo Command	30
Customizing the nsrsyb Script.....	30
How to Customize the nsrsyb Script.....	30
Using Pre- and Post-Processing Commands	30
PRECMD.....	31
USE_CONSISTENCY_CHECK	31
POSTCMD	31
Changing the SYBASE Variable.....	31
Changing the DBCCOPT Variable	31
Changing the BACKUP_OPT Variable	32
Changing the NetWorker XBSA Environment Variables	32
Chapter 3: On-Demand Backups.....	33
On-Demand Backups.....	33
How to Backup Sybase Databases	33
Monitoring Backups from NetWorker User for Sybase	34
NetWorker User for Sybase Backup Options	34
Backup Type	34

Data Compression	34
Dump Options.....	34
NetWorker Pool Names.....	34
On-Demand Backups from the Command Line.....	35
Running nsrsybsv at the Command Line	35
Running nsrsybcc at the Command Line	36
On-Demand Backups and NetWorker Indexes	36
Manual Backup of NetWorker Indexes for Use during Disaster Recovery	37
Example of an On-Demand Backup from the Command Line	37
Running the Dump Command at the isql Command Line.....	38
Dump Database Syntax.....	38
Dump Transaction-Log Syntax	38
Threshold Procedures	38
Backing Up Transaction Logs Using Thresholds	39
Chapter 4: Recovering Data	41
Information About Recovering Data	41
How to Perform a Recovery from NetWorker User for Sybase	42
NetWorker User for Sybase Recovery Options.....	43
How to Relocate Recovered Data	43
How to Determine Required Volumes	43
How to Change the Browse Time	43
Recovering Data with nsrsybrc	44
Running nsrsybrc at the Command Line.....	44
Types of Recoveries nsrsybrc Can Perform	45
Point-in-Time Recoveries	45
Relocated Recoveries	46
Imported Recoveries	46
Combining Relocated and Imported Recoveries	47
Running the Load Command at the isql Command Line	47
Disaster Recovery	48

How to Recover a Database Server After a Disk Crash	48
How to Recover a NetWorker Server and Database Server After a Disk Crash	49
How to Recover the Master Database.....	49
How to Recover Databases Not on the Master Device	50
Appendix A: XBSA Environment Variables.....	53
NetWorker XBSA	53
Changing NetWorker XBSA Environment Variables	53
Default Values and Valid Options.....	54
NSR_BACKUP_LEVEL.....	55
NSR_CLIENT	55
NSR_COMPRESSION.....	56
NSR_DATA_VOLUME_POOL.....	56
NSR_DEBUG_FILE	57
NSR_DEBUG_LEVEL.....	57
NSR_GROUP	58
NSR_LOG_VOLUME_POOL.....	58
NSR_NO_BUSY_ERRORS	59
NSR_SAVESET_NAME.....	59
NSR_SERVER	60
Appendix B: Error Messages.....	61
NetWorker Error Messages.....	61
nsrck	62
nsrexecd	63
nsrindexd	64
nsrmmdbd.....	65
nsrsyb	66
nsrsybcc	66
nsrsybrc	69
nsrsybsv	75

savegrp	82
Sybase Backup Server and libbms Messages	90
NetWorker XBSA and libbms Error Messages	93
Appendix C: Command Line Utilities	103
Syntax for the NetWorker Module Command Line Utilities	103
Conventions	103
Sybase Usernames and Passwords	103
nsrsybsv	104
nsrsybcc	105
nsrsybrc	106
Backup Options	106
Full Backups on Separate Devices	107
Full Backups on the Same Device	107
Incremental Backups on Separate Devices	108
Incremental Backups on the Same Device	108
Syntax for Dump and Load Commands from the isql Command Line	109
Dump Database Syntax	109
Dump Transaction Log Syntax	109
Finding the Timestamp for a Save Set	109
Load Database Syntax	110
Load Transaction Log Syntax	111
Glossary	113
Index	117

Preface



About This Guide

The *Legato NetWorker Module for Sybase Administrator's Guide* contains information on how to configure and manage the Legato NetWorker Module for Sybase software.

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

Audience

The information in this guide is intended for system administrators who are responsible for installing software and maintaining the servers and clients on a network. Operators who monitor the daily backups may also find this manual useful.

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 */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 */acroread/readme.txt* file or at the Adobe web site.

Conventions

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

Table 1. Document Conventions

Convention	Indicates	Example
boldface	Names of DOS or UNIX line commands, daemons, options, programs, or scripts	The nsradmin command starts the command line version of the administration program.
<i>italic in text</i>	Pathnames, filenames, computer names, new terms defined in the Glossary or within the chapter, or emphasized words	Displayed messages are also written to <i>/nsr/logs/daemon.log</i> .
<i>italic in command line</i>	A variable that you need to provide in the command line	nwadmin -s <i>server-name</i>
fixed-width	Examples and information displayed on the screen	media waiting: recover waiting for 8mm 5GB tape volume name
fixed-width, boldface	Commands and options that you must type exactly as shown	nsr_shutdown -a
Menu_Name>Command>Command	A path or an order to follow for making selections in the GUI.	Volume>Change Mode>Appendable
Important:	Information that you must read and follow to ensure successful backup and recovery of your data	 Important: 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 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	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
Note: For detailed information about our services, support policies, and software subscriptions, refer to the <i>Legato Customer Service and Support Guide</i> 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	<i>http://license.legato.com</i>
Telephone number	(650) 812-6063 (option 3)
Fax number	(650) 812-6220
E-mail	<i>service@legato.com</i>
E-mail for order status	<i>order_admin@legato.com</i>

Customer Feedback

The Legato NetWorker Module for Sybase team 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 Manuals page, and then look for the link to the Legato Documentation Survey.

Chapter 1: Introduction

Legato NetWorker Module for Sybase works with Sybase-SQL Server or Sybase-Adaptive Server Enterprise and Legato NetWorker to provide reliable, high-performance data protection for local or distributed Sybase databases. The NetWorker Module for Sybase integrates backup and recovery procedures for your databases into the comprehensive network-wide data protection facilities that NetWorker provides.

Note: Hereafter in this manual, Sybase-SQL Server and Sybase-Adaptive Server Enterprise will be referred to as Sybase server.

Importance of Backing Up Mission-Critical Data

Although the reliability of computer equipment has improved in recent years, hardware failures still occur, sometimes with catastrophic results.

In a client-server environment, you can lose data not only after hardware failures, but also because of user errors. Software bugs, procedural flaws, and simple user error are common culprits, requiring database media recoveries. A viable backup strategy can help you recover from these potentially disastrous situations.

Many database administrators do not recognize the danger of failing to make regular backups of databases and transaction logs. If you lose all transaction logs, you can still recover a database to the time of its last full backup. Without backups or transaction logs, you cannot recover the database at all.

Sybase Server

The Sybase database server works with the Sybase Backup Server to provide:

- Online, backups and recoveries of databases and transaction logs.
- Automated transaction log backups using a stored threshold procedure, or on-demand transaction log backups.

Legato NetWorker

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

With NetWorker, you can:

- Perform automated backups during off-peak hours.
- Use centralized administration to configure, monitor, and control backups from anywhere on a network.
- Automate tape-handling tasks using NetWorker SmartMedia™ software, which supports a wide variety of 4-mm Digital Audio Tape (DAT), 8-mm, and high-end devices, as well as bar-code label recognition and cleaning cartridge support.
- Increase backup performance by simultaneously sending data from multiple clients to the backup server.

Legato NetWorker Module for Sybase

Legato NetWorker Module for Sybase is an add-on module for NetWorker that provides automated backup media management and scheduling for the database server. The NetWorker Module for Sybase provides:

- True “after-hours” database storage management through automated scheduling, autochanger support, and electronic tape labeling and tracking.
- Support for local backup or distributed backup to a centralized backup server.
- Automated, scheduled backups of databases and transaction logs.
- High performance through support for high-speed devices, such as Digital Linear Tape (DLT) drives.

How NetWorker Connects to the Database Server


NetWorker and the NetWorker Module for Sybase provide storage management services for Sybase databases through an X-Open® Backup Services Application Programming Interface (XBSA API) connection to the Sybase Backup Server Archive API. The **libbms** shared library accepts API calls from the Sybase Backup Server, translates them into XBSA calls, and

sends them to NetWorker. You use the NetWorker administration program to set up backup schedules, label backup volumes, and configure the system running Sybase as a storage management client of the NetWorker server.

When **nsrd** triggers a scheduled backup for a database server instance on the NetWorker server, **savegrp** executes the **nsrsyb** script instead of performing a standard **save**. The **nsrsyb** script interacts with Sybase Open Client/C to coordinate a backup of the specified Sybase database.

NetWorker takes care of the scheduling and storage management tasks, while **libbms** takes care of passing the data from Sybase to NetWorker. Figure 1 shows the functional relationship between NetWorker, NetWorker Module for Sybase, and Sybase during a scheduled backup.

Figure 1. Data Movement through the Sybase Server during Backups




The Sybase server can exist on the same system as the NetWorker server software, or it can exist on a separate system. Because **nsrsyb** is substituted for the client program **save** during a backup, the NetWorker Module for Sybase

software must be installed on the system where you installed Sybase server. No matter where the system running the database server software resides, the system is considered a storage management client of the NetWorker server.

When a recovery request is initiated, the NetWorker **libbms** shared library translates the object names requested by the **nsrsybrc** command or the load command into a format understood by NetWorker and forwards it to the NetWorker server's **nsrd** daemon. The **libbms** shared library also sends a request to the **nsrindexd** index daemon. The **nsrindexd** daemon checks the client file index to ensure that an entry for the requested objects exists and passes this information back to **libbms**.

The **libbms** shared library then requests the data from the **nsrmmmd** media daemon. The **nsrmmmd** media daemon searches the NetWorker server's media database for the media containing the objects requested and recovers the data to the database server. Figure 2 shows how data moves from the NetWorker server to the database server during a recovery session.

Figure 2. Data Movement during a Recovery Initiated by nsrsybrc



Chapter 2: Scheduled Backups

NetWorker can back up your Sybase databases and transaction logs according to a schedule you set up. You add your database server to the NetWorker server's list of clients, and then specify which data you want backed up and the level of backup you want to run on specific days.

This chapter discusses the following topics:

- Ensuring that you have the required Sybase roles and privileges for backing up and recovering databases
- Using backup groups for your database server instances
- Using NetWorker backup schedules to automate your backups
- Customizing volume pools to track, organize, and sort your database and transaction log backups
- Using NetWorker policies to manage the life cycle of your data
- Configuring your database server as a NetWorker client
- Customizing the backup script to suit your environment

If you want to begin using NetWorker Module with its default configurations, see the following sections:

- [“Suggestions for Setting Up NetWorker Module for Sybase Groups” on page 21](#)
- [“Using NetWorker Backup Schedules” on page 22](#)
- [“Configuring a Database Server as a NetWorker Backup Client” on page 27](#)

Be sure to read the *Legato NetWorker Administrator's Guide*, shipped with your NetWorker server software, for complete information about using the NetWorker administrative program's graphical interface.

Required Sybase Roles and Permissions

Because the NetWorker Module for Sybase relies on the administrator’s ability to dump and load databases when performing backup and recovery operations, the administrator must have the appropriate Sybase roles and permissions. [Table 2](#) lists the Sybase roles and permissions required for performing NetWorker Module administrative actions.

Table 2. Sybase Roles and Permissions

Role/Permission	Action	NetWorker Module Command
SA_role or create database privileges	Create a database	n/a
SA_role, DBO (database ownership), or OPER_role	Dump and load databases	nsrsybsv; nsrsybrc
SA_role, DBO	Run a database consistency check	nsrsybcc

Note: During a scheduled NetWorker backup, a database consistency check (**nsrsybcc**) runs by default. The Sybase OPER_role does not have permission to run a database consistency check; therefore, a user with an OPER_role cannot run a scheduled NetWorker backup of the database server without disabling the consistency check first.

Using a Backup Group

You can assign one or more server instances to a group of NetWorker clients. NetWorker groups enable you to distribute backups to alleviate network traffic and to schedule backups for a time of day when performance demands on your database and NetWorker servers are lower.

What Is a Backup Group?

A *backup group* is a NetWorker *resource* that allows NetWorker clients assigned to a group to start backing up files at a designated time. You can have data backed up for a NetWorker group written to a specific set of media, or you can allow it to mix with backup data from other NetWorker groups.

NetWorker provides a preconfigured group named “Default,” which ships with preconfigured settings:

- Autostart = Disabled
- Start time = 3:33
- Client retries = 0
- Clones = No
- Clone pool = Default Clone

You can modify the Default group’s attributes, but you cannot delete Default from the list of NetWorker groups.

Suggestions for Setting Up NetWorker Module for Sybase Groups

If you have a large number of server instances, consider creating backup groups with different start times to help reduce network traffic. You can have any number of backup groups configured on your NetWorker server.

When you select the start times for each group, be sure to schedule them far enough apart so that one group completes its backup before the next group starts.

For instructions that explain how to set up a backup group, refer to the *Legato NetWorker Administrator’s Guide* shipped with your NetWorker server software.



Important: Be sure to enable the Autostart option for the group you configure; otherwise, the group’s scheduled backup will not take place.

Using NetWorker Backup Schedules

NetWorker uses backup schedules to automate the level of a scheduled backup. The NetWorker server’s administration program provides a graphical calendar that accepts entries designating the level of backup that should occur on a given day of the week.

NetWorker provides several preconfigured schedules. You can select the preconfigured Default schedule or create your own. The NetWorker Module translates valid NetWorker levels to Sybase **dump** commands as shown in [Table 3](#). Create a schedule for your database server backups using the information provided and the instructions for setting up backup schedules in the *Legato NetWorker Administrator’s Guide*.

Table 3. NetWorker Module Backup Level Translations

NetWorker	Sybase Server	Backup Operation
full	dump database	Full backup of a database and its transaction logs
incremental	dump transaction log	Backup of all changes to the database since the last dump
1–9	not valid	Fails with an error message
skip	skip	Skips the schedule backup

If your database server manages a significant amount of data, consider scheduling a full backup every one to two weeks and incremental backups on the other days to protect data changes between full backups.

You can start a scheduled backup at any time by using the “Group Control” feature in the NetWorker administration program. For instructions about monitoring and controlling backups, refer to the *Legato NetWorker Administrator’s Guide*.



Important: If you have scheduled an incremental backup, but a full backup has never been completed, NetWorker will perform a full backup instead of incremental.

What Is a Full NetWorker Backup?

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

For more information, see [“Full Backups on Separate Devices” on page 107](#) and [“Full Backups on the Same Device” on page 107](#).

What Is an Incremental NetWorker Backup?

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

For more information, see [“Incremental Backups on Separate Devices” on page 108](#) and [“Incremental Backups on the Same Device” on page 108](#).

When are Incremental Backups Not Allowed?

Incremental backups are not allowed under the following conditions:

- The database and transaction log are on the same device.
- The **select into/bulk copy** option is selected.
- The **truncate log on checkpoint** option is selected.
- A full backup has never been performed.

Using Volume Pools

NetWorker provides a means for directing your backups to specific sets of storage media. Volume pools enable you to establish a logical and systematic method for tracking, organizing, and sorting your backup data. Volume pools always have a label template associated with them to allow an automated method of identifying media assigned to a pool.

What Is a Volume Pool?

A *volume pool* is an assigned collection of backup volumes containing specific data sorted during a NetWorker backup. All NetWorker volumes belong to a pool, either one preconfigured by NetWorker or one you create.

Each pool has a Pool type designation that indicates whether the volume contains data that has been archived, backed up, or migrated.

Note: For the NetWorker Module for Sybase, the only valid Pool types are “Backup” and “Backup clone.”

Volume pools provide the ability to segregate backed-up data, such as database and transaction log backups, from regular filesystem backups. Additionally, you can separate full (database and transaction logs) backups and incremental (transaction log) backups. Pools also enable you to direct backup data to specific devices.

NetWorker uses the choices you select to sort backed-up data to specific backup volumes labeled for the pool. You can sort NetWorker filesystem backup data by several categories:

- Backup group
- NetWorker client
- Save set
- Backup level

For more information about using NetWorker volume pools, refer to the *Legato NetWorker Administrator's Guide*.

Customizing Volume Pools

By default, NetWorker Module directs your database and transaction log backups to the Default volume pool. To separate your Sybase backups from your regular filesystem backups, create a customized volume pool for your Sybase data. You might want to organize backup data even further, for example, by department or type of database maintained.

To ensure that NetWorker Module directs your Sybase backups to your custom pool, you must edit the `NSR_DATA_VOLUME_POOL` and `NSR_LOG_VOLUME_POOL` environment variables in the **nsrsyb** script to include your custom pool name. NetWorker Module uses the XBSA environment variables `NSR_DATA_VOLUME_POOL` for full backups and `NSR_LOG_VOLUME_POOL` for incremental backups to determine where to send Sybase backup data. The **nsrsyb** entry in the client resource relays the information needed to direct full and incremental backups to the correct backup media.

Alternatively, you can set the **-b poolname** option in the `BACKUP_OPT` variable in the **nsrsyb** script to direct both full and incremental backups to the same volume pool. The **-b poolname** option takes precedence over settings in the `NSR_DATA_VOLUME_POOL` and `NSR_LOG_VOLUME_POOL` environment variables.

You can direct your database backups and transaction log backups to different pools, but keep in mind that this process requires separate volumes from each pool to complete a backup or recovery operation. Sending your database and transaction logs to the same volume pool is the preferred method.



Example: Customize Media Organization

To customize media organization for NetWorker Module:

1. Create a label template for your Sybase backups.
 2. Create a volume pool for your Sybase backups.
 3. In the backup script, modify the values assigned to the NSR_DATA_VOLUME_POOL and NSR_LOG_VOLUME_POOL environment variables to reflect the custom pool name. Alternatively, you can add the **-b poolname** option to the BACKUP_OPT environment variable to send full and incremental backups to the same pool.
-

For additional details, see [“Customizing the nsrsyb Script” on page 30](#).

Creating a Volume Label Template

NetWorker generates labels for backup volumes according to the rules of a label template configured on the NetWorker server. To add a label template resource, you create a name for the label template, specify the fields to use in the label, provide the alphabetic or numeric range for the volumes, and select a separator to use between the fields. The order in which you enter the fields determines the order of the fields in the label templates. The first field you enter is the first field NetWorker uses in the label template.

The label template’s name cannot include any of the characters reserved for use as field separators. NetWorker uses the period, dash, underscore, and colon as field separators. NetWorker displays an invalid character message if you try to save a configuration that contains a field separator in the label template name.

When you save the configuration, NetWorker displays the label to apply to the next volume in the pool you associate with the label. The label template is also added to the choices available for label template selection for NetWorker volume pools.

For complete instructions about using NetWorker label templates and labeling and mounting backup volumes, refer to the *Legato NetWorker Administrator’s Guide* shipped with your NetWorker server software.

How to Create a Volume Pool

To create a customized pool for database and transaction log backups:

1. Open the Pools resource on the NetWorker server.
2. Create a name for the new volume pool.
3. Select Backup for the pool type.
4. Select the label template you created or select a preconfigured template.
5. Select the group to which the client resource belongs.
6. Select a device from the choices displayed.
7. Select the full and incremental levels from the choices displayed. This step ensures that the database and transaction log backups go to the same volume pool.
8. Save the settings.

When you save the configuration, NetWorker adds the pool to its list of volume pool resources. NetWorker also adds the new pool to the choices available for labeling volumes. When you select a pool for labeling volumes, NetWorker uses the label template you created and assigned to the pool to display the next label available in the series.

Using NetWorker Policies

NetWorker uses browse policies for the *online client indexes* and retention policies for the *media database* entries to manage and reduce the size of the online indexes. You can choose one of the preconfigured policies provided with NetWorker or create policies of your own. You can use a policy as either a browse policy or a retention policy.

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

You should keep copies of your transaction log backups until the associated database backup has exceeded its browse policy.



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

Configuring a Database Server as a NetWorker Backup Client

NetWorker uses a client-server model to provide storage management services. You must designate at least one computer on the network as the NetWorker server. You configure computers with data to back up as clients of the NetWorker server. If you have multiple instances of the database server on one computer, you can configure a separate NetWorker client resource for each database instance.

You configure NetWorker clients using the Clients resource provided by the NetWorker administration program. NetWorker maintains the resource information, contacts clients listed in a backup group configured on the server, performs on-demand backups when it receives a client request, and recovers data upon request from the client. The NetWorker server maintains the online file index and media database.

What Is a NetWorker Client?

A NetWorker client is a resource configured on the NetWorker server. The client resource provides the server with information about the data to back up for a client, how long to maintain entries for the data in the online index for recovery (browse policy), and how long to keep the media containing the client's backed-up data (retention policy).

Creating a NetWorker Client

To configure a Sybase server as a NetWorker client:

1. Create a client resource on the NetWorker server.
2. Select the Default schedule or a customized backup schedule that uses the valid backup levels from [Table 3 on page 22](#).
3. Select a browse policy.
4. Select a retention policy.
5. Select a group.
6. Specify a *save set*. For details on specifying a database server instance or an individual database, see [“Specifying a Save Set” on page 28](#).
7. Enter your Sybase user name in the Remote user text box and your Sybase password in the Password text box.
8. Enter the name of a customized version of the backup script in the Backup command text box. If you want to use backup options during a scheduled backup, you must edit the BACKUP_OPT variable in the **nsrsyb** script to include the backup options.

9. Leave the text boxes for directives, remote access, and archive users blank. The **nsrsyb** script controls data compression.
10. Enter all known aliases for the computer running your database server in the Aliases text box. For example:

venus
venus.acme.com

11. Save the client resource.

For details about using the NetWorker administrative graphical interface to configure a client, refer to the *Legato NetWorker Administrator's Guide* shipped with your NetWorker server software.

Specifying a Save Set

If you want to back up the entire database server, delete the choice All and enter the database server name. The entry shown in the example is the equivalent of performing a **dump database** at the command line for each database in the database server:

SYBASE: /Sybase_Server_name



Important: NetWorker Module does not back up **tempdb** because this database contains only temporary data.

If you want to back up a specific database for the database server instance, include the database name in the save set string. You can specify more than one database by making a separate save set entry for each database. The entry shown in the example is the equivalent of performing a **dump database database_name** at the command line:

SYBASE: /Sybase_Server_name/database_name

You can set up multiple instances of Sybase that exist on the same computer as separate NetWorker client/save set combinations by adding a new save set to the save set list with the *instance name* of the other database server:

SYBASE: /second_Sybase_Server_name

For further details about the **dump database** and **dump transaction log** commands, refer to the *Sybase SQL Server System Administration Guide* or *Sybase Adaptive Server Enterprise System Administration Guide* included with your database server software.

How to Send an E-Mail Notification after a Scheduled Backup

This is for UNIX NetWorker servers only.

To send an e-mail notification of the results of a scheduled backup to the owner of a save set:

1. Open the NetWorker client resource for the database server instance.
2. Select the View Details option.
3. In the Owner Notification attribute, enter a notification command directed to the login ID of the owner appropriate for your platform:

Solaris example:

```
/usr/ucb/mail -s "jupiter-mars-venus backup" jdoe
```

AIX and HP-UX example:

```
/usr/bin/mail -s "jupiter-mars-venus backup" jdoe
```

Viewing the Results of a Backup

NetWorker provides several reports about the results of a backup:

- An e-mail “savegroup completion” notice upon completion of a scheduled backup.
- A series of error messages written to the NetWorker message log files. For more information about NetWorker and NetWorker XBSA error messages, refer to [“Appendix B: Error Messages” on page 61](#).
- A scrolling list of messages displayed in the main window of the NetWorker administrative graphical interface.
- A scrolling list of messages displayed in the Group Control window of the NetWorker administrative graphical interface.
- A printout of the NetWorker server’s bootstrap file for the backup session.

To see more detailed information than the completion report displays, refer to the `/nsr/logs/messages` file.

The *Legato NetWorker Administrator’s Guide* provides complete details about the reports generated by NetWorker. For suggestions about using these reports as a part of your disaster recovery plan, refer to the *Legato NetWorker Disaster Recovery Guide*.

Using the nsrinfo Command

To view the NetWorker server's index entries for databases backed up for a database server instance, use the **nsrinfo** command. To see the results of a backup operation, enter the following command at the command line:

```
nsrinfo -X all -n sybase client_name
```

For complete information about using the **nsrinfo** command, refer to the *Legato NetWorker Administrator's Guide* and **nsrinfo** man page or Command Reference page.

Customizing the nsrsyb Script

When you install the NetWorker Module for Sybase on a Sybase server, the installation places a working copy of the **nsrsyb** script or batch file in the same directory as the NetWorker executables. You can modify options in the **nsrsyb** script, or you can have customized scripts control backups of different instances of the database server.



Important: The **nsrsyb** script must reside in the same directory as the NetWorker **save** command. If you create a modified version of the script, keep the modified version in the same directory.

How to Customize the nsrsyb Script

To create a customized version of the backup script:

1. Copy *nsrsyb.sh* or *nsrsyb.bat*, to a new file. The name of your copied file must begin with either *nsr* or *save*, for example, *nsr_pubs2*.
2. Open the copied script to modify it.
3. Save your changes.
4. Enter the name of the customized backup script in the Backup command text box in the client resource for the database server.

Using Pre- and Post-Processing Commands

The **nsrsyb** script contains several variables that allow you to use Sybase Server commands to perform tasks before and after a scheduled backup.

PRECMD

The PRECMD variable enables you to customize your backup with a command that will execute before the backup starts. For example, you could run a script that puts the database server in single-user mode. If the script returns a nonzero value, the backup session fails and an error message indicates the failure. The POSTCMD, however, is still executed, even if the PRECMD fails.

USE_CONSISTENCY_CHECK

The USE_CONSISTENCY_CHECK variable executes the **nsrsybcc** database consistency check program before a backup occurs. It creates a PRECMD with the appropriate arguments for **nsrsybcc**. If you set the USE_CONSISTENCY_CHECK variable and a PRECMD is already set, the commands in the PRECMD variable override the commands in the USE_CONSISTENCY_CHECK variable, and **nsrsybcc** might not run.

POSTCMD

The POSTCMD variable enables you to customize your backup with a command that will execute after the backup is completed. For example, you could run a script that brings the database server into multi user mode. If the script returns a nonzero value, an error message indicates the failure. Operator intervention might be required to return the database to a ready state.

Changing the SYBASE Variable

The Sybase variable points to the directory where Sybase is installed. You can set the Sybase variable in the **nsrsyb** script to include the complete directory path to Sybase, for example:

```
SYBASE=/space2/sybase11
```

or

```
SET SYBASE=C:\Sybase
```

Changing the DBCCOPT Variable

The DBCCOPT variable enables you to set any of the command line options for the **nsrsybcc** program within the **nsrsyb** script. When the **nsrsyb** script runs, it uses the options you set when it runs the database consistency check program. For the complete syntax of **nsrsybcc** and an explanation of the options, see [“Appendix C: Command Line Utilities” on page 103](#).

Changing the BACKUP_OPT Variable

The BACKUP_OPT variable enables you to set any of the command line options for the **nsrsybsv** program within the **nsrsyb** script. These options are passed to **nsrsybsv** during scheduled backups. For example, you can set options to indicate the debug level (-D).

If you created a customized volume pool for your Sybase backups, you can use the **-b pool_name** option to direct your database and transaction log backups to the same volume pool.

For the complete syntax of **nsrsybsv** and an explanation of the options, see [“Appendix C: Command Line Utilities” on page 103](#).

Changing the NetWorker XBSA Environment Variables

The **nsrsyb** script contains several NetWorker XBSA environment variable settings shared by NetWorker and the **libbms** shared library for backup tasks. The script is already configured to the default settings shown in [“Appendix A: XBSA Environment Variables” on page 53](#). You can modify NetWorker XBSA settings shown in the script to any of the valid values shown in Appendix A.

Chapter 3: On-Demand Backups

This chapter discusses running on-demand backups and database consistency checks from the NetWorker User for Sybase (Windows NT) or the command line. The example commands in this chapter are in C-shell format.

On-Demand Backups

You can initiate on-demand backups of your database server from the NetWorker User for Sybase program or from the command line with the **nsrsybsv** program. You can also initiate on-demand database consistency checks from the command line with the **nsrsybcc** program.

Before you perform an on-demand backup or database consistency check, you must set the Sybase variable in your environment to include the complete path to the directory where Sybase is installed.

How to Backup Sybase Databases

To perform a backup of a Sybase server from the NetWorker User for Sybase:

1. Make sure that the NetWorker server and the Sybase server are running.
Also make sure that the NetWorker Remote Exec services are running on the computer where the Sybase server is installed.
2. Start the NetWorker User for Sybase program and select Backup.
3. In the Sybase Server Login Dialog Box, enter the following:
 - The name of a Sybase server instance
 - Sybase user id for that instance
 - Sybase password for the above user id
4. Select either a single database or a server instance to back up, and select Mark. To unmark an item, select the item and select Unmark.

5. Optional: Select which backup options to use, see [“NetWorker User for Sybase Backup Options” on page 34](#).
6. Click Start to run the backup.

Monitoring Backups from NetWorker User for Sybase

A NetWorker Module for Sybase manual backup can be monitored from the NetWorker User for Sybase. When the backup is complete, the NetWorker Module for Sybase produces a backup session report.

NetWorker User for Sybase Backup Options

Several backup options are available when using the NetWorker User for Sybase. These options are described in the remainder of this section.

Backup Type

The NetWorker Module for Sybase can perform full or incremental backups. To specify what type of backup to perform, simply click on full or incremental.

Data Compression

NetWorker client systems can compress data during backup before data is moved over the network or written to tape. Compressing data may significantly accelerate the backup process. Compression is turned off by default.

Dump Options

The behavior of these backup options is affected by the backup level (full or incremental) and by whether the database and transaction logs are on the same or separate devices. For more information, see [“Backup Options” on page 106](#).

- TRUNCATE_ONLY - Clicking this check box is equivalent to using the -T option with the **nsrsybsv** command.
- NO_TRUNCATE - Clicking this check box is equivalent to using the -R option with the **nsrsybsv** command.
- NO_LOG - Clicking this check box is equivalent to using the -G option with the **nsrsybsv** command.

NetWorker Pool Names

For full backups, you can enter the name of a data pool to which you want to backup.

1. Select the Full backup option.
2. Click the Use Pools check box.

3. Enter the name you want in the Data pool text box, for example Sybase_Data.

For incremental backups, you can enter the name of a log pool to which you want to backup.

1. Select the Incremental backup option.
2. Click the Use Pools check box.
3. Enter the name you want in the Log pool text box, for example Sybase_Logs.

On-Demand Backups from the Command Line

The following sections describe how to perform on-demand backups from the command line.

Sybase recommends that you perform a database consistency check before running an On-Demand backup. For information about running a database consistency check, see [“Running nsrsybcc at the Command Line” on page 36](#).

Running nsrsybsv at the Command Line

Occasionally, you might want to initiate a database backup from the command line. You can use the **nsrsybsv** program to dump a database or transaction log from the command line.

By default, backups initiated with the **nsrsybsv** program write data to the Default pool. If you created a customized volume pool for your Sybase data, use the **-b pool_name** option with the **nsrsybsv** program at the command line

See [“Appendix C: Command Line Utilities” on page 103](#) for details about the options and correct syntax for the **nsrsybsv** program.

Alternatively, you can direct your Sybase data to a specific volume pool by setting the XBSA volume pool environment variables to match the customized volume pool name.

You can specify several different options for backing up the transaction logs when you use the **nsrsybsv** program. The backup option flags for transaction logs that you can specify are `truncate_only (-T)`, `with no_log (-G)`, and `with no_truncate (-R)`. For details, see [“Backup Options” on page 106](#).

Running nsrsybcc at the Command Line

NetWorker Module for Sybase ships with a database consistency check program, **nsrsybcc**, that you can run from the command line. For specific details, refer to the **nsrsybcc** man page or **nsrsybcc** Command Reference page.

For faster consistency checks for backups, you can configure **nsrsybcc** not to check non clustered indexes by setting the **-o ckdbnoidx** option. For a complete list of the options and correct syntax for the **nsrsybcc** program, see [“Appendix C: Command Line Utilities” on page 103](#).

For SQL Server, the default settings for **nsrsybcc** check the catalog, the device allocations, and the database (while skipping non clustered indexes). For Adaptive Server Enterprise, if the **dbccdb** database is set up, the default settings for **nsrsybcc** check the storage. If the **dbccdb** database is not set up, the default settings are the same as for SQL Server.

Occasionally, the **nsrsybcc** program might return an error caused by database locking problems. For complete details on database consistency checking and locked databases, refer to the *Sybase Administration Guide*.

On-Demand Backups and NetWorker Indexes

During each scheduled NetWorker backup, the NetWorker server’s **savegrp** program invokes a level 9 backup of the client indexes and the NetWorker server’s bootstrap file. The **savegrp** program is only invoked during a *scheduled* backup.

When you initiate an on-demand backup from the command line using the **nsrsybsv** program, the NetWorker client indexes for the database server and the NetWorker server’s bootstrap are not backed up. If you *never* perform scheduled NetWorker backups, the NetWorker client indexes for the database server and the NetWorker server’s bootstrap file are *never* backed up for use in the event of a catastrophic disaster.



Important: The client index for the database server and the NetWorker server’s bootstrap file are vital for recovering data to the database server in the event of a disaster. Therefore, perform regular, scheduled backups of the NetWorker server to provide maximum protection for your critical data.

Manual Backup of NetWorker Indexes for Use during Disaster Recovery

Occasionally, you might need to perform an on-demand backup of your database server using the **nsrsybsv** program. After performing an on-demand backup, back up the client file index and bootstrap by invoking the following command from the command line:

```
savegrp -0 -l full -P printer_name -c client_name
```

This ensures that you have the necessary information to recover from a disaster.

Refer to the **savegrp** man page or the **savegrp** Command Reference page for a complete description of the **savegrp** command and command options.

Example of an On-Demand Backup from the Command Line

The following is an example on-demand backup from a UNIX command line:

1. Designate a specific NetWorker server.

```
setenv NSR_SERVER NetWorker_server_name
```

2. Set the SYBASE variable to the Sybase directory.

```
setenv SYBASE /space2/sybase
```

3. For Sybase 12.0 or later, run the Sybase.sh or Sybase.csh script.

```
source $SYBASE/SYBASE.sh
```

4. Back up a specific database and transaction logs on the database server.

```
nsrsybsv -U userid -P password -l full  
SYBASE:/server_name/db_name
```

5. Save the NetWorker server index and bootstrap save set.

```
savegrp -0 -l full -P printer_name -c client_name group_name
```

If you want to back up all of the databases in the database server, use the command:

```
nsrsybsv -U userid -P password -l full SYBASE:/server_name
```

Running the Dump Command at the isql Command Line

You can run the Sybase **dump** command from the **isql** command line.

The **dump** command uses the environment variables that are set for the Sybase Backup Server. If you made any changes to the XBSA environment variables, such as setting `NSR_SERVER`, `NSR_DATA_VOLUME_POOL`, or `NSR_LOG_VOLUME_POOL`, make sure that these variables are also set for the Sybase Backup Server. You cannot specify a NetWorker server or volume pool with the **dump** command at the **isql** command line. For details on setting the environment variables for the Sybase Backup Server, refer to your *Sybase Administration Guide*.

Use the NetWorker Module **nsrsybsv** and **nsrsybrc** programs rather than the Sybase **dump** and **load** commands to backup and recover NetWorker save sets.

Dump Database Syntax

The syntax for each database to be dumped is (including quotes):

```
dump database database_name to "bms::"
```

Dump Transaction-Log Syntax

The syntax for each transaction log to be dumped is (including quotes):

```
dump transaction database_name to "bms::"
```

For more details concerning the appropriate syntax for the **dump** and **load** commands, see ["Syntax for Dump and Load Commands from the isql Command Line" on page 109](#).

Threshold Procedures

NetWorker Module ships with a sample threshold procedure that you can use to implement transaction log backups to free log space. The sample threshold procedure is configured to perform the following tasks:

- When the threshold is reached, NetWorker Module dumps the transaction log.
- If dump transaction log is not allowed, the NetWorker Module performs a full database dump and truncates the transaction log.

If NetWorker Module cannot perform a full database dump, you must either add space to the transaction log or abort any processes that were suspended when the threshold was crossed. Refer to the *Sybase Administration Guide* for details.

You can edit the sample threshold procedure to suit your environment. [Table 4](#) lists the default location for the sample threshold procedure:

Table 4. Threshold Procedure Location

Platform	Location
Windows NT	<i>c:\Program Files\nsr\bin\threshold.sql</i>
Solaris	<i>/nsr/examples/sybase/threshold.sql</i>
AIX	<i>/nsr/examples/threshold.sql</i>
HP-UX	• <i>/opt/networker/examples/threshold.sql</i>

Follow the instructions in the *Sybase Administration Guide* for using the **sp_addthreshold** command to add the NetWorker Module threshold procedure to your database server and for details on managing free space with thresholds.

Backing Up Transaction Logs Using Thresholds

To ensure that your transaction logs are backed up regularly to free log space, use a threshold procedure that is set lower than the last-chance threshold. You can use the sample threshold procedure included with NetWorker Module or create one of your own. Refer to the *Sybase Administration Guide* for instructions on determining the appropriate value for the last-chance threshold.

When the database and transaction logs are on the same device, tripping the threshold causes a full database backup (**dump database**).

When the database and transaction logs are on separate devices, tripping the threshold causes an incremental database backup (**dump transaction log**). Sybase recommends keeping your user databases and transaction logs on separate devices.

Chapter 4: Recovering Data

NetWorker Module for Sybase provides a means for recovering data backed up by NetWorker to your database server in the event of data corruption or a disk crash. NetWorker Module enables you to perform on-demand recovers of database server instances and individual databases using the NetWorker User for Sybase (Windows NT) or the **nsrsybrc** program from the command line.

Use the information in this chapter in conjunction with the *Sybase Administration Guide*, the *Legato NetWorker Disaster Recovery Guide*, and the **nsr_crash** NetWorker man page or the **nsr_crash** Command Reference page. The command examples in this chapter are in C-shell format.

Information About Recovering Data

During a scheduled or on-demand backup, the NetWorker server makes an entry in an online client file index and records the location of the data in an online media database. These entries provide recovery information needed for every database server object backed up. The client index entry is maintained in the index until the browse policy configured for the client's save set expires. When the retention policy configured for the client's save set expires, the save set's status changes from Recoverable to Recyclable in the media database.

When all the save sets on the storage media change status to Recyclable, the media mode changes status to Recyclable, and the media is eligible for relabeling. The save set entries, however, remain in the media database until the media is relabeled. Until the media is relabeled, the data is still recoverable, using the NetWorker **scanner** command.



Important: NetWorker Module does not support recovery by save set. This means that you cannot use the NetWorker **recover** program to recover Sybase data that has passed its browse policy. You must recover the save set entries in the client index with the **scanner -i** command then use the **nsrsybrc** program to recover your data.

The **scanner** man page and the **scanner** Command Reference page provide complete details about recovering save sets whose client index entries have expired.

You can view the online client file index entries in the NetWorker administrator program. The listing displays the save set ID assigned during a backup session, the number of files and their size, and the date and level of the backup session.

NetWorker sends a record of the server's *bootstrap* generated during the backup to your default printer, so you have a printed record of the dates, locations, and save set ID numbers for the server's online indexes needed for recovering data. Keep the bootstrap printouts in a file for quick reference in the event of a disaster, such as a disk crash or server malfunction.

4

How to Perform a Recovery from NetWorker User for Sybase

To perform a recovery using the NetWorker User for Sybase:

1. Ensure that the NetWorker server services are running.
2. Start the NetWorker User for Sybase and select Recover.
3. In the Sybase Server Dialog Box, enter the name of the Sybase server instance and the host name of the Sybase server computer that you wish to recover from. Click OK.
4. Select a file to recover, and select Mark.
To unmark an item, select the item and select Unmark.
5. Optional: Select which recover options to use, see [“NetWorker User for Sybase Recovery Options” on page 43](#).
6. Select Start to begin the recovery process.
7. Enter a Sybase user id and password for the server instance to which you are recovering. Click OK.

NetWorker User for Sybase Recovery Options

The following options are available when you perform a recovery with the NetWorker User for Sybase.

- [“How to Relocate Recovered Data” on page 43](#)
- [“How to Determine Required Volumes” on page 43](#)
- [“How to Change the Browse Time” on page 43](#)

How to Relocate Recovered Data

When you recover Sybase data from the NetWorker server, NetWorker copies the data from the volume to a location on your system. You can tell NetWorker to place the data back in a different database. To relocate recovered data:

1. From the Options menu, select Recover Options.
2. In the Server name text box, enter the name of the Sybase server instance to which you want to recover.
3. In the Database name text box, enter the name of a database to which you want to recover. The database name you enter must exist.
4. Select OK.

How to Determine Required Volumes

To determine the volumes required for a recovery:

1. Mark the database file(s) to recover.
2. From the View menu, select Required Volumes.

The backup volumes required for the recovery are displayed in the Required Volumes dialog box.

How to Change the Browse Time

From the Recover window, you can browse the entries for each backed-up database file listed in the client file index from any time in the past. The Change Browse Time option in the View menu allows you to view the entries for your backed-up files from a specific point in time. To change the Browse Time:

1. Select Change Browse Time from the View menu to open the Change Browse Time dialog box.
2. Select a day from the calendar to enter a new date.
3. Select Previous Month or Next Month to change from the current month.

4. Enter a time to browse, where:
 - Time is based on a 12-hour clock.
 - a represents a.m.
 - p represents p.m.

After you change the browse time, the Recover window displays the entries for the specified time. The new browse time is displayed in the status bar at the bottom of the Recover window. All databases you mark are recovered from a backup made at the selected time.

Recovering Data with nsrsybrc

You can use the **nsrsybrc** program to recover data backed up from a scheduled or on-demand NetWorker backup. The **nsrsybrc** program enables you to recover databases and roll transactions forward by recovering the transaction logs. You cannot recover database backups with the NetWorker **recover** program.

The recovered database overwrites any existing tables in the original database. If you want to recover data to a new database, you must first create the new database with the **for load** option. For instructions about creating a new database, refer to your *Sybase Administration Guide*.

Before you use the command line to initiate the **nsrsybrc** command, you must set the Sybase variable in your environment to include the complete path to the Sybase directory that contains the interfaces file. For details, see [“Changing the SYBASE Variable” on page 31](#).



Important: The **nsrsybrc** program does not create the database. If necessary, you must re-create the database and then use the **nsrsybrc** program. Sybase recommends that you maintain hard-copy records of your database definitions and device allocations so that you can easily recreate them.

Running nsrsybrc at the Command Line

The **nsrsybrc** program accepts command line options. If you enter a command line option that differs in value from one of the settings in the NetWorker XBSA environment variables, the command line option overrides the environment variable settings. For a list of the options and syntax that the **nsrsybrc** program accepts, see [“Appendix C: Command Line Utilities” on page 103](#).

Types of Recoveries nsrsybrc Can Perform

The **nsrsybrc** program supports point-in-time, relocated, and imported recoveries of databases from the backup media managed by NetWorker.

Point-in-Time Recoveries

For Sybase-ASE 11.5 or later, you can restore up to a specific point-in-time in the middle of an incremental backup. The **nsrsybrc** program uses the time supplied in the **-t** option to restore data to that specific time in a transaction log backup, if possible. This feature allows you to restore data to a point in the middle of a transaction log instead of restoring the entire transaction log. The database is brought online after the restore takes place.

```
nsrsybrc -U userid -P password -s networker_server  
-t "MM/DD/YY HH:MM:SS" SYBASE:/server_name[/db_name]
```

During a point-in-time restore, the **nsrsybrc** program restores a database up to a specific point in time. It loads the most recent full backup before the designated time and then applies any transaction log backup made before that time. The **nsrsybrc** program accepts a wide variety of time designations. For details, refer to the **nsrsybrc** man page or **nsrsybrc** Command Reference page.

Note: Setting the specific NetWorker server can help reduce the recovery time. Otherwise, NetWorker must search the network for the correct server.



Important: After performing a point-in-time recovery, the Adaptive Server restarts the databases's log sequence. This means that until you perform another full backup, you cannot perform another recover.

Relocated Recoveries

Relocated recoveries enable you to load the backup of an old database to a new database. The **nsrsybrc** program uses the **-d** option to specify where to recover the relocated data.

The following scenarios exist for relocated recoveries:

- Recovering to the same server instance, but to a different database name
nsrsybrc -U userid -P password -s networker_server
-d SYBASE:/server_name/new_db_name
SYBASE:/server_name/old_db_name
- Recovering to a different server instance, but to the same database name
nsrsybrc -U userid -P password -s networker_server
-d SYBASE:/new_server_name/db_name
SYBASE:/old_server_name/db_name
- Recovering to a different server instance and to a different database name
nsrsybrc -U userid -P password -s networker_server
-d SYBASE:/new_server_name/new_db_name
SYBASE:/old_server_name/old_db_name

Before you perform an relocated recovery:

- Be sure that your new database is already created and has the same device allocations as the original database.
- Refer to your *Sybase Administration Guide* for instructions about setting up a new database.
- Consider creating the new database with the “for load” option. This can help save you time when you create the new database.

Imported Recoveries

Imported recoveries enable you to recover a backup from a Sybase server on another NetWorker client. This can be done by using the **-c** option in the command line or setting **NSR_CLIENT** to that client.

```
nsrsybrc -U userid -P password -s networker_server  
-c networker_client SYBASE:/server_name[/db_name]
```

Note: Since the **-d** option is not used, the destination server name and database name used are the same as the original.

Combining Relocated and Imported Recoveries

Relocated and imported recoveries can be combined together, to relocate recovered data from a different client computer.

The following example recovers a database from a Sybase server instance on another host to a different Sybase server instance, which is on the computer you are running the command from.

```
nsrsybr -U userid -P password -s networker_server
      -c networker_client -d SYBASE:/new_server_name/db_name
SYBASE:/old_server_name/db_name
```

Running the Load Command at the isql Command Line

You can run the Sybase **load** command from the **isql** command line to recover a database or a transaction log.

The **load** command uses the environment variables that are set for the Sybase Backup Server. If you made any changes to the XBSA environment variables, such as setting **NSR_SERVER**, or **NSR_CLIENT**, make sure that these variables are also set for the Sybase Backup Server.

After you load a database from the **isql** command line, you must bring the database back online, with the following command:

```
online database database_name
```

Use the **nsrsybr** program rather than the **load** command to recover your data. When you use the **nsrsybr** program, the database is brought back online automatically.

To load a database or transaction log from the **isql** command line, you might need the timestamp for each database or transaction log. If you do not include a timestamp, NetWorker uses the most recent backup. Indicating the timestamp when you load a transaction log is useful if you have more than one transaction log.

You can find the timestamp for a save set using either of the following methods:

- Enter the following command to get a list of all the Sybase save sets for the NetWorker client:
nsrinfo -X all -n sybase *networker_client*
- Use the Save Set Recover window in the NetWorker Administration program to select the save set you want to recover. The date and time are displayed in the Instances window. Then enter the **load** command at the **isql** command line.



Important: You cannot use the Save Set Recover window to recover Sybase data. Use the **nsrsybrc** program to recover your databases and transaction logs. For details, see [“Recovering Data with nsrsybrc” on page 44](#).

The syntax for loading the most recent database backup is:

```
load database database_name from “bms:”
```

The syntax for loading the most recent transaction log is:

```
load transaction database_name from “bms:”
```

For more details concerning the appropriate syntax for the dump and load commands, see [“Syntax for Dump and Load Commands from the isql Command Line” on page 109](#).

Disaster Recovery

4

If you have a disaster recovery plan in place and use NetWorker Module for Sybase to maintain regular backups of your SQL Server or Sybase-ASE instances and the transaction logs associated with them, you are well-equipped to recover critical data quickly.

The information in this section assumes that you have read and are familiar with the procedures in your database server documentation and the information in the *Legato NetWorker Disaster Recovery Guide*.

How to Recover a Database Server After a Disk Crash

To recover a primary disk containing critical Sybase server and NetWorker client binaries:

1. Reinstall the NetWorker client binaries, the NetWorker Module for Sybase software, and the database server software, if needed. If you perform regular NetWorker backups of your system binaries, you can use NetWorker to recover the them.
2. Use your printout of database device allocations to re-create the databases. For details about what information you should track for disaster recovery, refer to your *Sybase Administration Guide*.
3. Follow the instructions in your *Sybase Administration Guide* for recovering the system databases. For an overview of the tasks you should perform, see [“Recover the Master Database” on page 49](#).
4. Use the **nsrsybrc** program to recover the data.

How to Recover a NetWorker Server and Database Server After a Disk Crash

To recover both the NetWorker server's primary disk containing the online indexes (the */nsr* filesystem) and the primary disk for the database server:

1. Reinstall the NetWorker server binaries, if needed.
2. Find the latest bootstrap printout for the NetWorker server and follow the procedure in "Recovering from a Disk Crash" in the *Legato NetWorker Disaster Recovery Guide* to recover the server's online indexes.



Important: Do not attempt to recover the NetWorker server's online client file index or media database to a different directory than the one from which they were backed up. After you recover the indexes to their original location, you can safely move them to another directory. For details about moving the indexes, refer to the *Legato NetWorker Disaster Recovery Guide*.

3. Reinstall the NetWorker Module for Sybase software and the database server, if needed.
4. Follow the instructions in your *Sybase Administration Guide* for recovering the system databases. For an overview of the tasks you should perform, see ["Recover the Master Database" on page 49](#).
5. Use the **nsrsybrc** program to recover the data.

How to Recover the Master Database

In the event of a disaster, you will need to recover the master database. The master database contains data about all the databases in the database server. This section provides an overview of the tasks you need to perform to recover the master database. For specific instructions refer to the *Sybase Troubleshooting and Error Messages Guide* and the *Sybase Administration Guide* for details.



Important: If you recover the master database to a different SQL Server or Sybase-ASE, all the device allocations are copied to the new database server. If you recover the master database to another database server on the same computer as the original, they both attempt to use the same database files. Refer to your Sybase documentation for instructions on recovering the master database to a different database server.

To recover the master database, follow the steps in the section “Master Database is Lost and Valid Dump Exists” in “Chapter 1: System Database Recovery” of the *Sybase Troubleshooting Guide*.

If the valid dump uses the server’s default sort order, perform the following tasks. If not, refer to the *Sybase Troubleshooting Guide* for the section on “Valid Dump with Non-Default Sort Order.”

To recover the master database:

1. Rebuild the lost master device using “buildmaster.”
2. Start the database server in single-user mode. This mode is also called master-recover mode.
3. Re-create any extra device allocations for master (db_id 1) so that they match the space allocations for the master device before the disaster. Refer to the *Sybase Administration Guide* for specific instructions.
4. Ensure that the database server has the correct name for the Sybase Backup Server in the **sys.servers** table.
5. Load the master database from backup with:

```
# nsrsybrc -U userid -P password  
SYBASE:/server_name/master
```

This command loads the master database and shuts down the database server.

6. While the database server is shut down, manually reestablish the number of devices configuration parameter, if necessary.
7. Restart the database server.
8. If you made changes to the system catalog information for the master database since the last dump, recover the changes now.
9. Load the model database.
10. Drop and re-create any other database that is on the master device.

How to Recover Databases Not on the Master Device

This section overviews the tasks you should perform to recover databases that either are not on the master device or are partially on the master device. Refer to the *Sybase Administration Guide* for specific instructions on recovering databases.

To recover databases not on the master device:

1. Obtain a current log dump of every database on each device.

2. Examine the space usage of every database on each device, for example from the **isql** command line:

```
select * from sysdatabases  
select * from sysusages  
select * from sysdevices
```

3. Once you have this information for all databases on each device, drop each database, except for the databases entirely on the master device. For example, enter the following command to drop each database from the **isql** command line:

```
dbcc dbrepair (database_name, dropdb)
```

4. Manually delete the rows in the **sysdevices** file that refer to the missing devices. This step drops the failed devices.
5. Restart the database server in single-user mode.
6. Initialize new devices.
7. Re-create the databases, one at a time.
8. Restart the database server. Be sure that the database is *not* in single-user mode.
9. Load the most recent database backup into each database with:

```
# nsrsybrc -U userid -P password  
SYBASE:/server_name/database_name
```

This command recovers the last full database backup of the specific database and applies all of the associated transaction log backups in the order that they were created.

Appendix A: XBSA Environment Variables

This appendix describes how to change NetWorker XBSA environment variables and lists their default values and valid options.

NetWorker XBSA

NetWorker XBSA allows for configuration of environment options to activate certain features of NetWorker not directly supported by X/Open specifications. NetWorker XBSA and the **libbms** shared library enable the Sybase Backup Server API and NetWorker to interact during backups and recoveries.

Changing NetWorker XBSA Environment Variables

To change a value for a NetWorker XBSA environment variable that does not appear in the default **nsrsyb** script:

1. On Windows NT, open the *nsrsyb.bat* file located in the */Program Files/nsr/bin/* directory.

On UNIX, open the *nsrsyb.sh* template file:

- For Solaris and AIX platforms, the default location of the template file is */nsr/examples/sybase/nsrsyb.sh*
 - For HP-UX platforms, the default location of the template file is */opt/networker/examples/nsrsyb.sh*
2. Add the NetWorker XBSA environment variable to the script and assign a valid value to the variable.
 3. On Unix, add the NetWorker XBSA environment variable to the list of variables exported.

```
export_environment_variables()
{
  If ["$NSR_VOLUME_POOL"!="custom_pool_name"]
  then
    export NSR_VOLUME_POOL
  fi
  .
  .
  .
export new_NSR_variable
}
```

4. Save the edited script with a new filename beginning with either *nsr* or *save*, for example, *nsr_pubs2*.
5. Edit the NetWorker client resource configured for the database server instance and enter the filename of the customized **nsrsyb** script in the Backup command field.



Important: If you choose to use a threshold procedure for transaction log backups or the **isql** command line, the environment variables set in the shell that launches the Sybase Backup Server override the default settings for the NetWorker XBSA environment variables.

A

Default Values and Valid Options

This section contains tables for each of the NetWorker XBSA environment variables you can modify by customizing the **nsrsyb** script. Most of the environment variables described appear in the default **nsrsyb** script, while others are set in the NetWorker XBSA shared libraries included with NetWorker Module for Sybase.

NSR_BACKUP_LEVEL

Definition	The NSR_BACKUP_LEVEL environment variable indicates the NetWorker backup level to use for the XBSA session.
Default Value	full
Possible Values	Although NetWorker supports backups of levels full , 1 through 9 , incremental , and skip , NetWorker Module for Sybase supports only backups of levels full , incremental , and skip . If you assign a level other than full , incremental , or skip , the backup fails and displays the message “Only level full or incremental backups allowed.” If you specify the skip level for a save set, the savegroup operation is completed successfully but does not perform a backup of the data.

NSR_CLIENT

Definition	The NSR_CLIENT environment variable indicates the NetWorker client resource to use for the XBSA session.
Default Value	The host from which the XBSA session is initiated, as indicated by getlocalhost() is the default value for NSR_CLIENT.
Possible Values	Because the client name is an arbitrary string, the value for NSR_CLIENT is not checked directly. An incorrect value might cause an authentication or system error in NetWorker.

NSR_COMPRESSION

Definition	The NSR_COMPRESSION environment variable indicates whether to compress the backup data as it is sent to the NetWorker server.
Default Value	FALSE
Possible Values	Setting NSR_COMPRESSION to a value of TRUE means that the standard compression technique for XBSA for NetWorker is performed on the data backed up. Setting NSR_COMPRESSION to a value of FALSE means that compression is not performed. Note: Compressing data from the database server might accelerate backups, as long as the database server can send data to the backup server fast enough to keep the tape drive streaming. Data compression during backup impacts CPU usage on the backup server, but reduces the amount of data sent to the NetWorker server.

NSR_DATA_VOLUME_POOL

Definition	The NSR_DATA_VOLUME_POOL environment variable indicates the volume pool to which data files should be backed up.
Default Value	XBSA does not set a pool by default. If no pool is specified, the NetWorker server selects the pool based on its pool resources configuration.
Possible Values	You can use any valid NetWorker pool name of 1024 characters or less as the NSR_DATA_VOLUME_POOL value. You set this value in the NetWorker pools resource and can explicitly assign it in a shell script.

NSR_DEBUG_FILE

Definition	The NSR_DEBUG_FILE environment variable indicates the full <i>pathname</i> and <i>filename</i> to which NetWorker XBSA messages should be written. Message logs for XBSA are separated from regular NetWorker messages. NetWorker XBSA error messages are indicated by the prefix BSA.
Default Value	/nsr/applogs/xbsa.messages
Possible Values	Any valid <i>pathname</i> and <i>filename</i> is acceptable. If NetWorker cannot open the file specified, it writes a BSA_RC_INVALID_KEYWORD error message to one of the following locations: • The alternate <i>/nsr/applogs</i> messages directory created during installation • The directory assigned to the TMPDIR environment variable • The <i>/tmp</i> directory, if TMPDIR is not set

NSR_DEBUG_LEVEL

Definition	The NSR_DEBUG_LEVEL environment variable enables you to set the level of NetWorker XBSA error report messages sent to the <i>xbsa.messages</i> log file.
Default Value	The default value is 2 , which means that critical error messages and all network (RPC) errors are written to <i>xbsa.messages</i> .
Possible Values	Any integer from 0 to 9 is valid. Higher values generate more detailed reports: • A value of 0 means that no error messages are written to <i>xbsa.messages</i>. • A value of 1 means that only critical error messages are written to <i>xbsa.messages</i>. • A value of 2 means that all network (RPC) errors are written to <i>xbsa.messages</i>. • A value of 3 means that all NetWorker XBSA informational messages are written to <i>xbsa.messages</i>. • You can override the default value in the nsrsybsv, nsrsybrc, and nsrsybcc programs with the -D option.

NSR_GROUP

Definition	The NSR_GROUP environment variable indicates the group configuration to use for a backup session.
Default Value	None
Possible Values	You can use any valid NetWorker group name of 1024 characters or fewer as the NSR_GROUP value. Invalid group names might cause authentication or system errors in another routine. A NetWorker group configuration acts as an alarm clock, notifying the NetWorker server that a group of clients has a backup scheduled to occur at the time designated in the group's Start Time attribute.

NSR_LOG_VOLUME_POOL

Definition	The NSR_LOG_VOLUME_POOL environment variable indicates the volume pool to which transaction logs should be backed up.
Default Value	NetWorker XBSA does not set a pool by default. If no pool is specified, the NetWorker server selects the pool based on its pool resources configuration.
Possible Values	You can use any valid NetWorker pool name of 1024 characters or fewer as the value for NSR_LOG_VOLUME_POOL. You set this value in the NetWorker pools resource.

NSR_NO_BUSY_ERRORS

Definition	The NSR_NO_BUSY_ERRORS environment variable indicates whether the savegroup should wait for a busy NetWorker server or fail immediately.
Default Value	The default value is FALSE , which means the savegroup will wait for the NetWorker server to accept the connection.
Possible Values	Setting NSR_NO_BUSY_ERRORS to a value of TRUE causes the backup to fail immediately when the NetWorker server is busy. A network error message describing the reason for the failure is written to the <i>xbsa.messages</i> file. If NSR_DEBUG_LEVEL is set to 1 and the NetWorker server is busy, the backup process stops and the error message "BSA_RC_ABORT_SYSTYEM_ERROR" is written to the <i>xbsa.messages</i> file.

NSR_SAVESET_NAME

Definition	The NSR_SAVESET_NAME environment variable indicates the save set name NetWorker XBSA should use for a save session.
Default Value	NetWorker XBSA uses the format "SYBASE:/server_name/database_name," where <i>server_name</i> is the name of the database server instance. Do not change this value.
Possible Values	You can use any valid NetWorker save set name as the value for NSR_SAVESET_NAME. Only the first 63 characters are meaningful to the NetWorker server's media database. For dbspaces, use: SYBASE : /<i>server_name</i>/<i>database_name</i>

NSR_SERVER

Definition	The NSR_SERVER environment variable indicates the <i>hostname</i> of the server that NetWorker XBSA should use for a save session.
Default Value	The default value is the most appropriate server, based on the index name and client name for the session. See also NSR_CLIENT.
Possible Values	The server name defined by the NSR_SERVER environment variable is checked using gethostbyname() . If this routine call fails, the NetWorker XBSA error code "BSA_RC_INVALID_KEYWORD" is returned.

Appendix B: Error Messages

This appendix lists error messages you might encounter while using NetWorker Module for Sybase and provides suggestions to resolve the problems. The error messages are organized into these categories:

- [“NetWorker Error Messages”](#)
- [“Sybase Backup Server and libbms Messages” on page 90](#)
- [“NetWorker XBSA and libbms Error Messages” on page 93](#)

NetWorker Error Messages

NetWorker error messages are displayed in the NetWorker Administrator window. The display lists the messages encountered during the previous 24 hours. The messages are also written to */nsr/messages/daemon.log*.

NetWorker error messages appear in the format:

day hh:mm:ss daemon_or_program_name: *message*

NetWorker XBSA and **libbms** error messages are written to the */nsr/applogs/xbsa.message* file. The **libbms** error messages are also reported to the Sybase Backup Server, which prints them to the *stdout* file and logs them in the Sybase Backup Server error log.

The messages are organized alphabetically by NetWorker daemon name and program name to make them easier to match to the NetWorker message displayed.

nsrck

cannot lock flag file for *clientname*: *reason*

The flag file signifying the end of the first part of index compression is already in use by another instance of the **nsrck** program, or by the **nsrindexd** daemon. Because disaster ensues if two processes access the same index at the same time, **nsrck** refuses to act on the named file.

checking index for *clientname*

Informative message that the files associated with the named client are being inspected.

completed checking *count* clients

Displayed as the program finishes, provided some form of checking was accomplished.

compressing index for *clientname*

Displayed when the **-x** or **-C** option has taken effect.

cross-checking index for *clientname*

Displayed when the **-X** option is in effect.

more space needed to compress *clientname* index, *size required*

Indication that the **nsrck** program cannot find enough disk space to hold the *db.CMP* temporary file. Free some disk space on any local filesystem and retry the command. You can use the **df(8)** command to see how much free space is available on any filesystem.

rolling forward index compression for *clientname*

Indicates that after a reboot, if index compression completed its first copy, the compression is being rolled forward.

WARNING no valid savetimes - cross-check not performed for *clientname*

Indicates that during a cross-check, no save sets were found for this client. Because this situation can occur during disaster recovery, **nsrck** avoids deleting the entire contents client index and instead does nothing.

nsrexecd

/path/nsrexecd: Can't make pipe
/path/nsrexecd: Can't fork
fork: No more processes

The specified client-side resource has been exceeded. There are too many other services running on the client while **savegrp** is running. Inspect the client and determine why it has run out of resources. You might need to reboot the client. You should also consider rescheduling any jobs automatically started on the client—for example, through **cron(8)**—that run while **savegrp** is running.

/path/nsrexecd: Couldn't look up address for your host
/path/nsrexecd: Host address mismatch for server

The **nsrexecd** daemon on the client managed to look up the server in the client's host table, but the address listed there did not match the address of the server. Every interface of the server must have a unique name listed in the host table (possibly with nonunique aliases or CNAMEs), and each unique name must be listed as a valid server to **nsrexecd**.

/path/nsrexecd: Host server cannot request command execution
/path/nsrexecd: Your host cannot request command execution

The server is not listed in **nsrexecd**'s list of valid servers on the specified client. The list of valid servers is either on the **nsrexecd** command line (with one or more **-s** server options to **nsrexecd**), or in a file (with the **-f** file option to **nsrexecd**). It might also be the case that the server is not listed in one or more of */etc/hosts*, NIS, or DNS on the client, in which case **nsrexecd** cannot validate the server until the client's host-naming configuration is fixed.

/path/nsrexecd: Invalid authenticator
/path/nsrexecd: Invalid command

The **savegrp** program did not follow its protocol correctly. These two messages should never occur in a savegroup completion message.

/path/nsrexecd: Permission denied
Permission denied

Messages generated by **nsrexecd** and **rshd**, respectively. In either case, the server does not have permission to execute commands on the client. In the case of the first message, make sure that the server is listed as a valid server on the client (see “/path/nsrexecd: Host server cannot request command execution” for details). In the case of the second message, which does not mention **nsrexecd**, make sure that “servername” is listed in the client’s *.rhosts* file (or, if you set the remote user attribute for this client, the *.rhosts* file in the home directory for that user on the client).

socket: protocol failure in circuit setup

The client does not seem to support the TCP/IP protocol stack, or has not used a privileged port for setting up its connection. The latter could occur if you use **nsrexecd** but did not start it as *root* on the specified client. The **nsrexecd** daemon must run as *root* on each client.

nsrindexd

lock on *filename* acquired

Informative message that eventually follows the “waiting for lock” message.

waiting for lock on *filename*

Indicates that another program is accessing the same file that the **nsrindexd** daemon requires. The daemon waits for the advisory lock to be cleared.

nsrmmdbd

A copy of this process is already running

Indicates that another copy of nsrmmdbd is currently running and has exclusive access to the media database. Only one nsrmmdbd process should be running on a given machine at a time. This can happen if the previous nsrmmdbd was not properly killed off. Use nsr_shutdown or ps and kill(1) to identify and kill off all the NetWorker daemons before restarting nsrd(8) again.
--

Cannot open lock file

An internal error. Check the permissions on the <i>/nsr/tmp</i> and <i>/nsr/mm</i> directories.

media db is cross checking the save sets
--

Displayed each time the daemon is restarted. Upon startup, the daemon checks its records before providing its service.
--

media db is open for business

Displayed after any of the other media database messages are displayed to indicate that the service is available again.

media db is recovering, this might take a while

Displayed when the nsrmmdbd daemon is reloading its database. The service is unavailable while the data is being reloaded.

media db is saving its data, this might take a while
--

Displayed when the daemon is dumping its records to a temporary file when the database is being backed up. The service is unavailable while the database is dumping.
--

nsrsyb

nsrsybsv returned status of <i>value</i> nsrsyb exiting
Indicates that an error occurred when nsrsybsv was executed. Check the nwadmin display or the <i>/nsr/applogs/xbsa.messages</i> file on the NetWorker server to determine which error caused nsrsybsv to fail.

nsrsybcc

<i>CS-LIBRARY</i> or <i>CT-LIBRARY</i> error: <i>error_message</i> . Operating system error number(<i>n</i>): <i>error_message</i> .
Displayed when an error occurs in the Sybase Open Client library layer. The operating system error message part is displayed only if an operating system error occurred. You might see these error messages when you recover the master database because this operation shuts down the database server, but they are not normal during other operations. The error message text describes the specific problem.
the command line did not specify any databases or instances
The nsrsybsv , nsrsybrc , and nsrsybcc programs each operate on a database (or, for nsrsybrc and nsrsybcc , a list of databases). No databases or instances were specified on the command line.
the command line may specify the entire instance or a list of individual databases, but not both
You can specify either the entire instance (SYBASE:/ <i>server_name</i>) or a list of databases (SYBASE:/ <i>server_name/database_name1</i> SYBASE:/ <i>server_name/database_name2</i>) at the command line. An instance name and a list of databases cannot be specified at the same time.

The command line specifies more than one Sybase instance. Only a single instance may be supplied with each command line.

Each invocation of the **nsrsybsv**, **nsrsybcc**, or **nsrsybrc** programs can operate on a single Sybase server because the user ID and password supplied are unlikely to be the same over multiple servers. Retry the command and run it once for each Sybase server.

the database name *database_name* has a length greater than the maximum of 32

Indicates that the database name supplied at the command line was longer than 32 characters. The maximum database name length is 32 characters.

error from server *server_name*: Msg *number*, Level *number*, State *number*

The Sybase server returned an error. Check the error message that follows this message to determine the reason for the error.

The instance name was not provided in the command line *command_line_value*. The command line has the form SYBASE:/instance_name[/database_name].

The database to be processed was specified as "SYBASE:", but the instance name was not supplied.

invalid check option *-o value* was supplied

The database consistency check option that was supplied is not valid. See ["Appendix C: Command Line Utilities" on page 103](#) or refer to the **nsrsybcc** man page for a list of supported options.

no NetWorker server was specified

This message indicates that no NetWorker server was specified or could be found. You can specify the NetWorker server to which you want to issue the command with the **-s server_name** option.

non fatal internal error from server *server_name*: Msg *number*, Level *number*, State *number*

The Sybase server returned a nonfatal error. This error does not stop the operation, but you should examine the message to ensure that the error does not lead to future problems.

path needs to begin with SYBASE:. The command line has the form SYBASE:/instance_name[/database_name]

The database name option for the **nsrsybcc** program did not begin with the characters "SYBASE:". All Sybase server save sets must begin with this name.

SQL Server *server_name* version is too old. It must be 11.0 or later, and it is *version_number*.

NetWorker Module for Sybase is supported on SQL Server 11.x or later and Adaptive Server Enterprise 11.5 or later.

User name and password were not found in the client *client_name* resource in group *group_name*. Edit the client resource for saveset *save_set_name* to include these.

The NetWorker client resource for this save set and group did not have the Sybase user ID and password set. This means that the scheduled backup could not log in to the Sybase server.

user name is required and was not supplied

You must supply a user name for Sybase login. This user name can be queried from the client resource in the NetWorker server, entered on the command line, or obtained from the environment variable, \$USER.

nsrsybrc

CS-LIBRARY or *CT-LIBRARY* error: *error_message*.
Operating system error number(*n*): *error_message*.

An error occurred in the Sybase Open Client library layer. The operating system error message part is displayed only if an operating system error occurred. You might see these error messages when you recover the master database because this operation shuts down the database server, but they are not normal during other operations. The error message text describes the specific problem.

cannot restore to the destination database
database_name because it does not exist in the instance
server_name

The database to which **nsrsybrc** is recovering does not exist. Create the database and try the **nsrsybrc** program again.

the command line did not specify a database or an
instance to restore

You must supply the name of the database or Sybase server instance you want to recover when you use the **nsrsybrc** program from the command line.

the command line did not specify any databases or
instances

The **nsrsybsv**, **nsrsybrc**, and **nsrsybcc** programs each operate on a database (or, for **nsrsybrc** and **nsrsybcc**, a list of databases). No databases or instances were specified on the command line.

the command line may specify the entire instance or a
list of individual databases, but not both

You can specify either the entire instance (SYBASE:/*server_name*) or a list of databases (SYBASE:/*server_name/database_name1* SYBASE:/*server_name/database_name2*) at the command line. Both an instance name and a list of databases cannot be specified at the same time.

The command line specifies more than one Sybase instance. Only a single instance may be supplied with each command line.

Each invocation of the **nsrsybsv**, **nsrsybcc**, or **nsrsybrc** programs can operate on a single Sybase server because the user ID and password supplied are unlikely to be the same over multiple servers. Retry the command and run it once for each Sybase server.

the database name *database_name* has a length greater than the maximum of 32

The maximum database name length is 32 characters. This error message indicates that the database name supplied at the command line was longer than 32 characters.

error from server *server_name*: Msg *number*, Level *number*, State *number*

The Sybase server returned an error. You should check the error message that follows this message to determine the reason for the error.

If master is being restored, no others can be restored in the same session. The database must be in master recover mode to recover master, and this precludes restoring any other database.

A list of databases to recover was specified, and the master database was listed along with others. Recovering the master database shuts down the Sybase server, which makes recovering other databases impossible.

if the destination is an instance, the source must be an instance, too

The **-d destination** option was used to specify a server instance, but the item to be recovered is a single database. You should retry the command and specify the destination database, for example,
nsrsybrc -U sa -P xxx -d
SYBASE: /destination_server/destination_database
SYBASE: /source_server/source_database.

if the source is an instance, the destination must be an instance, too

The object to be recovered is an entire Sybase server instance, but the destination specified to recover the instance to is a database name. You should retry the command and specify the destination as an instance, for example,

```
nsrsybrc -U sa -P xxx -d  
SYBASE:/destination_server SYBASE:/source_server.
```

The instance name was not provided in the command line *command_line_value*. The command line has the form SYBASE:/instance_name[/database_name].

The database to be processed was specified as "SYBASE:", but the instance name was not supplied.

internal error. Full backup expected but not found.

This error occurs only when a full backup is found but is then no longer available before **nsrsybrc** recovers the database. For example, if you manually relabel the volume containing the full backup at the same time you try to recover the incremental backup that depends on that full backup, this error occurs.

invalid time specification: *time value*

The **-t time** option supplied on the command line to the **nsrsybrc** program was not valid. This option should be supplied in the **nsr_getdate** form. Refer to the **nsr_getdate** man page for details.

no backup was found for database *database_name*

The **nsrsybrc** program could not find a backup of the database specified for recovery. Run the **nsrinfo** program to see if a backup exists, and make sure that the user ID used for **nsrsybrc** matches the object owner that is displayed. You should run the Sybase Backup Server and the **nsrsybrc** and **nsrsybsv** programs from the same user ID to avoid this problem.

no NetWorker server was specified

This message indicates that no NetWorker server was specified or could be found. You can specify the NetWorker server to which you want to issue the command with the **-s *server_name*** option.

Non fatal internal error from server *server_name*: Msg *number*, Level *number*, State *number*

The Sybase server returned a nonfatal error. This error does not stop the operation, but you should examine the message to ensure that the error does not lead to future problems.

path needs to begin with SYBASE:. The command line has the form SYBASE:/instance_name[/database_name]

The **-d destination** option or the database name option for the **nsrsybrc** program did not begin with the characters "SYBASE:". All Sybase server save sets must begin with this name.

SQL Server *server_name* version is too old. It must be 11.0 or later, and it is *version_number*.

NetWorker Module for Sybase is supported on SQL Server 11.x or later and Adaptive Server Enterprise 11.5 or later.

Sybase server version *version_number* does not support the "checkstorage" option. Versions 11.5 and later support it

The database consistency check checkstorage option only works with Adaptive Server Enterprise 11.5 and later. SQL Server 11.x does not support this option.

there are no databases to restore in instance *server_name*

There were no databases found in the directory entry for the Sybase server database.

there is no backup of the instance for the time supplied

No backup could be found for the Sybase server name supplied. Make sure that you run the **nsrsybrc** program with the same user ID that you used to run the **nsrsybsv** program. Otherwise, make sure that the time you use is correct. If you do not enter a time, the current time is used.

there is no full backup of database *database_name* in instance *server_name* for the time supplied

Backups of this database exist, but there was not a full backup available for the time requested. Try an earlier time, or run the **nsrinfo** program to determine when the last full backup occurred. For example, if the full backup has passed its browse policy, the full backup might be listed in the media database but not in the client index. In this situation you must re-create the entry in the client index with the **scanner -i** command, and then recover the database with the **nsrsybrc** program.

Unable to close temporary file that has environment variables. Check for disk full or privilege errors in /nsr/tmpdir.

The temporary file used to pass environment variables between **nsrsybsv** and **libbms** could not be closed. The permissions might be incorrect, or the disk might have insufficient space to write the file. Redirect the **nsrsybsv** program to create a temporary directory in a different place by setting the NSR_TMPDIR variable.

unable to create directory /nsr/tmpdir

The access privileges for the user running **nsrsybsv** are insufficient to create the /nsr/tmpdir directory. Re-create the directory manually or change the permissions so that this directory can be created by this user. Redirect the **nsrsybsv** program to create a temporary directory in a different place by setting the NSR_TMPDIR variable.

unable to open temporary file to pass environment variables

The temporary file used to pass environment variables between **nsrsybsv**, **nsrsybrc**, and **libbms** could not be opened. Check for file access or disk problems. Redirect the **nsrsybsv** program to create a temporary directory in a different place by setting the NSR_TEMPDIR variable.

unable to query backup

There was an error querying the backup from the server. Check the */nsr/applogs/xbsa.messages* file for the specific error text.

unable to write environment variables to the temporary file

The system could not write to the temporary file used to pass environment variables between **nsrsybsv** and **libbms**. Check for file access or disk problems.

User name and password were not found in the client *client_name* resource in group *group_name*. Edit the client resource for saveset *save_set_name* to include these.

The NetWorker client resource for this save set and group did not have the Sybase user ID and password set. This means that the scheduled backup could not log in to Sybase server.

user name is required and was not supplied

You must supply a user name for Sybase login. This user name can be queried from the client resource in the NetWorker server, entered on the command line, or obtained from the environment variable, \$USER.

nsrsybsv

Both the truncate only (-T) and the no truncate (-R) options were specified. Only one may be specified for a given backup.

The -T and -R options, when supplied together, ask for opposite results. Select one or the other option and retry the operation.

cannot find database *database_name* in instance *server_name*

The database to be backed up does not exist in the Sybase Server.

the command line did not specify a database or an instance to back up

The command line is missing the SYBASE:/*server_name/database_name* value that specifies the database to back up.

the command line did not specify any databases or instances

The **nsrsybsv**, **nsrsybrc**, and **nsrsybcc** programs each operate on a database (or, for **nsrsybrc** and **nsrsybcc**, a list of databases). No databases or instances were specified on the command line.

the command line may specify the entire instance or a list of individual databases, but not both

You can specify either the entire instance (SYBASE:/*server_name*) or a list of databases (SYBASE:/*server_name/database_name1* SYBASE:/*server_name/database_name2*) at the command line. An instance name and a list of databases cannot be specified at the same time.

The command line specifies more than one Sybase instance. Only a single instance may be supplied with each command line.

Each invocation of the **nsrsybsv**, **nsrsybcc**, or **nsrsybrc** programs can operate on a single Sybase server because the user ID and password supplied are unlikely to be the same over multiple servers. Retry the command and run it once for each Sybase server.

CS-LIBRARY or *CT-LIBRARY* error: *error_message*.
Operating system error number(*n*): *error_message*.

An error occurred in the Sybase Open Client library layer. The operating system error message part is displayed only if an operating system error occurred. You might see these error messages when you recover the master database because this operation shuts down the database server, but they are not normal during other operations. The error message text describes the specific problem.

the database name *database_name* has a length greater than the maximum of 32

The maximum database name length is 32 characters. This error message indicates that the database name supplied at the command line was longer than 32 characters.

environment variable NSR_BACKUP_LEVEL has an invalid value of *value*, defaulting to a full backup

The environment variable NSR_BACKUP_LEVEL specified a level other than full, incremental or skip.

error from server *server_name*: Msg *number*, Level *number*, State *number*

The Sybase server returned an error. You should check the error message that follows this message to determine the reason for the error.

the exit status of process *process_number* could not be determined

The PRECMD or POSTCMD that was executed did not exit, but the process no longer exists.

a full database backup is required and will be done before the transaction log backup

The incremental backup failed because a full backup must first be performed. A full backup will be performed and then the transaction log backup will be retried.

The -G option is not supported where the transaction log is on a separate device. This option will be ignored.

The -G option is not supported for a full backup of a database when the transaction log is on a separate device.

The -G option is not valid for an incremental backup where the log is on a separate device. It will be ignored. To truncate the transaction log without logging it, use the same command with a level of full instead of incremental.

The -G option is not supported for incremental backups when the transaction log is on a separate device.

The instance name was not provided in the command line *command_line_value*. The command line has the form SYBASE:/instance_name[/database_name].

The database to be processed was specified as "SYBASE:", but the instance name was not supplied.

An invalid backup level was supplied. Valid backup levels are full, incremental, and skip

The backup level supplied to **nsrsybsv** is not permitted.

no command to execute in *PRECMD* or *POSTCMD*

The *PRECMD* or *POSTCMD* environment variable is set, but it has no value. You should either unset the environment variable or enter a value to execute in the environment variable.

no NetWorker server was specified

This message indicates that no NetWorker server was specified or could be found. You can specify the NetWorker server to which you want to issue the command with the **-s** *server_name* option.

Non fatal internal error from server *server_name*: Msg *number*, Level *number*, State *number*

The Sybase server returned a nonfatal error. This error does not stop the operation, but you should examine the message to ensure that the error does not lead to future problems.

only one database or instance may be specified

More than one database or instance was supplied on the command line to the **nsrsybsv** program. The **nsrsybsv** program only supports a single instance (SYBASE:/*server_name*) or database (SYBASE:/*server_name/database_name*) per invocation.

path needs to begin with SYBASE:. The command line has the form SYBASE:/*instance_name*[/*database_name*]

The **-N save_set_name** option or the database name option for the **nsrsybsv** program did not begin with the characters "SYBASE:". All Sybase server save sets must begin with this name.

PRECMD or *POSTCMD* did not return a result. It needs to return zero on success and non-zero on failure.

The *PRECMD* or *POSTCMD* did not return a status value.

process *process_number* running command *PRECMD* or *POSTCMD* completed with a result of *n*

The *PRECMD* or *POSTCMD* exited with a nonzero result code. Check the *PRECMD* or *POSTCMD* exit code for details. You should also verify that the settings in the *PRECMD* or *POSTCMD* are valid. For details, see [“Using Pre- and Post-Processing Commands” on page 30](#).

SQL Server *server_name* version is too old. It must be 11.0 or later, and it is *version_number*.

NetWorker Module for Sybase is supported on SQL Server 11.x or later and Adaptive Server 11.5 or later.

Sybase server version *version_number* does not support the “checkstorage” option. Versions 11.5 and later support it

The database consistency check checkstorage option only works with Adaptive Server 11.5 and later. SQL Server 11.x does not support this option.

The -T option is not valid for an incremental backup where the log is on a separate device. It will be ignored. To truncate the transaction log, use the same command with a level of full instead of incremental.

The -T option only works for a full backup when the transaction log is on a separate device.

Unable to close temporary file that has environment variables. Check for disk full or privilege errors in /nsr/tmpdir.

The temporary file used to pass environment variables between **nsrsybsv** and **libbms** could not be closed. The permissions might be incorrect, or the disk might have insufficient space to write the file.

unable to create directory /nsr/tmp

The access privileges for the user running **nsrsybsv** are insufficient to create the */nsr/tmp* directory. Re-create the directory manually or change the permissions so that this directory can be created by this user.

unable to create directory entries

The directory entries could not be created. Check the */nsr/applogs/xbsa.messages* file for the specific reason that the entries could not be created.

unable to determine whether database and log are on separate segments

The database to be backed up is not in a state in which it can be queried to determine whether incremental backups are allowed. The error message from the Sybase server that was displayed prior to this message indicates the reason the database cannot be queried.

unable to dump database *database_name* in instance *server_name*

The **dump** database command failed. The error message from the Sybase server that was displayed prior to this message indicates the reason the database was not dumped.

unable to dump the transaction log for database *database_name* in instance *server_name*

The command to dump the transaction log failed. The error message from the Sybase server that was displayed prior to this message indicates the reason the transaction log was not dumped.

unable to dump the transaction log without truncating it for database *database_name*

The command to dump the transaction log with the *no_truncate* option failed. The error message from the Sybase server that was displayed prior to this message indicates the reason the transaction log was not truncated.

unable to execute the command *PRECMD* or *POSTCMD*
contents

The PRECMD or POSTCMD could not be found. You should make sure the command exists in one of the directories specified in \$PATH.

unable to open temporary file to pass environment
variables

The temporary file used to pass environment variables between **nsrsybsv**, **nsrsybrc**, and **libbms** could not be opened. Check for file access or disk problems.

unable to print savegrp completion message

After the backup occurred, NetWorker could not find the save sets in the media database.

Unable to print summary. One or more parameters are not
set

The parameters that NetWorker expected to find for the function that prints the **savegrp** summary were not supplied.

unable to spawn process to issue the *PRECMD* or *POSTCMD*
command

The PRECMD or POSTCMD could not be run because a process needed to run them was not available.

unable to truncate the transaction log for database
database_name

The command to truncate the transaction log failed. The error message from the Sybase server that was displayed prior to this message indicates the reason the transaction log was not truncated.

unable to truncate the transaction log for database *database_name* with the *no_log* option

The command to truncate the transaction log failed. The error message from the Sybase server that was displayed prior to this message indicates the reason the transaction log was not truncated.

unable to write environment variables to the temporary file

The system could not write to the temporary file used to pass environment variables between **nsrsybsv** and **libbms**. Check for file access or disk problems.

User name and password were not found in the client *client_name* resource in group *group_name*. Edit the client resource for saveset *save_set_name* to include these.

The NetWorker client resource for this save set and group did not have the Sybase user ID and password set. This means that the scheduled backup could not log in to Sybase server.

user name is required and was not supplied

You must supply a user name for Sybase login. This user name can be queried from the client resource in the NetWorker server, entered on the command line, or obtained from the environment variable, \$USER.

savegrp

Aborted

This informational message only occurs when you stop a running **savegrp**. The session for this save set might not disappear immediately, especially if the program's attempt to kill the save session fails. The save set will be retried when you restart the **savegrp**.

Access violation - unknown host: *client*

This message is caused when the client's hostname and IP address are not correctly listed in one or more of */etc/hosts*, NIS, or DNS on the server. You need to perform one of the following actions:

- Change the appropriate host table (depending on which ones are in use on your server) to list the client's primary name as it is known to the NetWorker server.
- Add the name listed at the end of the error message to the aliases attribute of the client's Client resources.

asm: cannot open */path*: I/O error

This message generally means that there are bad blocks on the disks containing the specified file or directory. You should immediately run a filesystem check on the named client filesystem and check your client's system error log. If there are bad blocks, repair them or move the filesystem to a different disk.

asm: cannot stat */path*: Stale NFS file handle
asm: cannot stat */path*: Missing file or filesystem

These informational messages (or variants of them for other operating systems) mean that when **save** attempted to test the named directory to determine whether it was a different filesystem from the one currently being saved, the filesystem was, in fact, NFS-mounted, but the mount point was bad. Although this message does not affect the saved data, it does mean you have a network or NFS problem between the specified client and one or more of its file servers. You might need to remount filesystems on the client, or perhaps reboot it, to correct the problem.

asm: external ASM '*asm2*' exited with code 1

This message generally accompanies another message reporting a specific problem while saving a file or directory on the named save set. The backup attempts to continue and to save other data. Also, generally, the backup is not listed in the failed save sets section of the completion mail if any files on the save set are saved successfully, even if it only saves the top directory of the save set.

asm: getwd failed

This message means that while backing up the specified save set, an attempt to determine the current directory's name failed. This occurs on clients that are generally running older versions of the NetWorker ClientPak™, on which the **getwd(3)** library call is broken. You might want to contact Legato to find out whether there is a patch available for your client platform to work around this vendor-specific bug. Otherwise, contact your operating system vendor to see whether a more recent operating system version addresses this problem.

asm: missing hard links not found

This message is generated when a backed-up file had one or more hard links that were not found. The message is followed by a list of one or more filenames that were backed up minus some links. The message means that the files were either created (with multiple hard links) while the backup was occurring (so some of the links were missed due to the order of filesystem tree walking), or the file or some links were removed while the backup was occurring. Only those links that were found can be recovered; additional links have been lost. You can do an additional incremental backup of the affected filesystem if a consistent state for the affected file is essential.

asm: /path was not successfully saved

This message generally accompanies one or more specific messages for the save set. The specified path within the current save set was not saved successfully. The backup continues trying to back up other files and directories on the save set.

asm: xdr_op failed for /path

This error can be caused by several possible conditions (for example, insufficient memory, faulty networking software in the operating system, an external ASM unexpectedly exiting, or a lost network connection). If it was due to a lost network connection, the NetWorker server most likely exited (due to **nsr_shutdown**). After restarting the server, rerun the group. If it was due to an ASM exiting unexpectedly (in which case, an accompanying message describes which ASM exited unexpectedly), you might have a bad block on the disk or a bug. Check whether the client ran out of memory (there might be console messages), and verify that there are no bad blocks on the save set's disk. If there were network errors, there might also be messages logged by other programs on the system console (client or server), or to system log files.

connect to address *AA.BB.CC.DD: message*
Trying *AA.BB.CC.DD...*

These informational messages are displayed only when you use the **-v** option. They mean that the connection to the client failed on the address specified in the first line of the message. If the client has more than one IP address, **savegrp** has attempted the address listed in the second line. Looking at subsequent lines of the completion mail shows whether this second address succeeded. You might want to check and change your network routing tables to avoid getting these messages.

Connection refused

This means the client machine is up, but it is not accepting new network connections for **nsrexecd** (or **rshd**). This could mean the client was in the process of booting when the **savegrp** attempted to connect, or that the client exceeded some resource limit and was not accepting any new connections. You should attempt to log in to the client and verify that it is accepting remote connections.

Connection timed out

This usually means the client has crashed or is hung. Make sure that the client has rebooted and that **nsrexecd** is running on it (if you are using **nsrexecd**).

filesystem: No such file or directory

You named an explicit save set in the Client resource for the specified client, and that save set does not exist (or is not currently mounted) on the client. Make sure that you spelled the save set name correctly (and that it is capitalized correctly), and log in to the client and verify that the save set is mounted.

group groupname aborted, savegrp is already running

This message is only delivered by itself. It occurs when the named group has already been started or restarted (for example, after a reboot or when requested through the Group Control Window of **nwadmin**(8)), either automatically by **nsrd**(8) or manually from the command line. You can use **ps**(1) to find out the process ID of a running **savegrp**. You can determine the existence of a running group by looking for a file named */nsr/tmp/sg.groupname*, which, if existing and locked, means a **savegrp** is running.

*has been inactive for n minutes since time
client:saveset is being abandoned by savegrp*

A backup of the specified save set started, but after minutes of no activity, **savegrp** gave up on the save set. Generally, this means that the client is hung, waiting for an NFS partition. NetWorker (or any other program) has no way of reliably telling whether an NFS partition will hang until after it tries to access the partition. When the partition comes back on line, the **save** is completed, even though **savegrp** abandoned it. You should check the client, however, because you sometimes need to reboot the client to unhang NFS partitions. Non-UNIX clients also hang for other reasons that include, most notably, bugs in the implementation of their network protocols.

Host is unreachable

The NetWorker server cannot make TCP/IP connections to the client. This generally means the network itself is not configured correctly. Most commonly, one or more gateways or routers are down, or the network routes were not set up correctly. You should verify that the server can connect to the client. If the server cannot connect to the client, check your routers, gateways, or routing tables and reconfigure them, if needed.

no cycles found in media db; doing full save
This informational message is added by savegrp to any save set that is saved at level full instead of the level found in the client's schedule. Due to timing problems, you might occasionally see this message when the clocks on the client and server are out of sync, or when savegrp starts before midnight and ends after midnight.
No 'NSR client' resource for client <i>clienthostname</i> savefs: cannot retrieve client resources
This pair of messages occurs if the client's hostname has changed (in <i>/etc/hosts</i>, NIS or DNS). You might also have deleted the client's Client resource while savegrp was running. In the former case, you need to add the client's new name to the aliases attribute of the client (this is a hidden attribute) using nsradmin(8) (selecting the Hidden display option) or nwadmin(8) (selecting the Details View option for the Client window). In the latter case, no additional action is required if you deliberately deleted the resource (the next run of savegrp does not attempt to save the client). If you did not intend to delete the client, you should re-create the client resource and add the client back into the appropriate groups. The next time savegrp runs, it backs up the client, just as if the client had been down the previous day.
no output
The save set was completed, but it returned no status output. The most common reasons are that the client crashed or lost its network connection (that is, a router between the client and server crashed) while the client was being backed up. Another possible reason is that the disk on which the client status was being logged filled up (you can perform a df /nsr/tmp to see whether this was the case). To determine whether the save set was saved, you can use mminfo(8). For example, run mminfo -v -c clientname -t '1 day ago' and look at the flags column for the completion status. An 'a' flag means it aborted. Use a more distant time (the -t option) to look further back in time.

n retries attempted
1 retry attempted

One of these informational messages is included in a save set's output if **savegrp** could not back up the data on the first try and if the client's *retries* attribute for the group has a value greater than zero. In this case, the specified number of retries was performed before the backup of the save set succeeded or was finally marked as failed.

/path: This data set is in use and cannot be accessed at this time

This message is generated by save sets on PC clients running DOS or NetWare®. The NetWorker client software on these systems cannot back up files open for writing, due to the interface provided by the OS. This message comes from the Novell® TSA and is not changeable.

printer: unknown printer
/path/savegrp: printing bootstrap information failed (reproduced below)
/path/savegrp: printing bootstrap information failed

These messages, or similar messages, accompany the bootstrap information when **savegrp** cannot print the bootstrap on the printer. You need to either specify a different printer in the *printer* attribute for the group or configure your print server to recognize the printer (by default, your system's default printer is used). The bootstrap information is listed as part of the **savegrp** completion mail. You should print out this information immediately in case your server has a disaster and loses a disk, and fix the printer name used by **savegrp**.

reading log file failed

After the specified save set was completed, **savegrp** could not read the log file of the output status from the save set. This generally means that someone, or an automated non-NetWorker administrative program or script, removed the log file. This message can also occur if the filesystem on which the client logs are stored has run out of space (use **df /nsr/tmp** to determine whether this is the case). Verify that no scripts remove files from */nsr/tmp* (which is where **savegrp** stores the save set log files).

RPC exec on *client* is unavailable. Trying RSH.

This informational message is only displayed if you use the **-v** flag for verbose information. This message means that **nsrexecd** is not running on the client and that **savegrp** is attempting to use the **rshd** service instead, for backward compatibility with older versions of **savegrp**.

savegrp: client rcmd(3) problem for command '*command*'

This error message usually accompanies another, more specific, error message. It is generated when the attempt to run the specified command (usually **save** or **saveefs** with several command-line parameters) failed on the specified save set. The previous line of error output should include the more specific error message (look for that message elsewhere in this section). Generally, the problem is a bad host table configuration or various “permission denied” errors (server not specified when starting **nsrexecd** or missing permissions in *.rhosts* if you are not using **nsrexecd**). If not, log into the NetWorker server as *root* and run the **savegrp -p -v -c clientname groupname** command, giving the appropriate client for *clientname* and *groupname*. This output should include the necessary additional information needed for fixing the problem.

Saving server index because server is not in an active group

This informational message, generated by **savegrp**, means that **savegrp** has noticed that the NetWorker server is not listed in any automatically started, enabled group. Because all the indexes are stored on the server, **savegrp** is saving the server’s index and bootstrap information in case a disaster occurs. You should add the server to a group with Autostart enabled, or enable one of the groups of which the server is already a member.

socket: All ports in use

The NetWorker server has run out of socket descriptors. This means that you have exceeded the socket resource limit on your server. To avoid such future messages, you should determine what other network services are running while **savegrp** is running. Also, consider rescheduling either **savegrp** or the other service(s). You can also reduce the parallelism in the **nsr_service(5)** resource to reduce the resource utilization.

Sybase Backup Server and libbms Messages

When the Sybase Backup Server encounters an error or condition requiring a warning, it writes a message to the Sybase Backup Server error log. The default error log location is the \$SYBASE/install directory for Sybase Server 11.x or . \$SYBASE/\$SYBASE_ASE/install on Sybase Server 12.0.

If an error with the **libbms** shared library occurs, a **libbms** message is written to the /nsr/applogs/xbsa.messages file and reported to the Sybase Backup Server. The Sybase Backup Server logs the **libbms** error messages in the Sybase Backup Server error log.

This section lists **libbms** error messages that are logged in the Sybase Backup Server error log. For details on other Sybase Backup Server errors, 5efer to the *Sybase Administration Guide*.

libbms opened with an unknown mode: internal error
The libbms shared library was opened with a mode other than read or write.
there is insufficient memory to continue
There is not sufficient memory to complete the operation.
The time stamp dddddddd has non digits in it. Timestamps are composed of digits in the form YYYYMMDDhhmmsslll.
The timestamp supplied for the load command from the isql command line has a timestamp with an incorrect format. The timestamp must have the format YYYYMMDDhhmmsslll.
time stamps are not valid for dump command
The isql command line specified a timestamp for a dump command. Timestamps are not valid with the dump command.
unable to close and create save set
The BSA call to create and close the save set for a database or transaction dump failed. Check the /nsr/applogs/xbsa.messages file for specific details.

unable to close save set

The call to close the save set failed during a load of a database or a transaction log. Check the */nsr/applogs/xbsa.messages* file for specific details.

unable to create environment variables

The resources required to create the internal environment variable array were not available. This might be due to access problems in the */nsr/tmp* directory.

Unable to create save set. There is likely a configuration or enabler problem. Set the debug level to at least 2, retry the operation, and check the */nsr/applogs/xbsa.messages* file for the underlying reason.

The save set could not be created on the NetWorker server. If the debug level is at least 2 (the default), check the */nsr/applogs/xbsa.messages* file for the error text. If the debug level is not set at 2, change the setting to 2 and retry the operation. Check the */nsr/applogs/xbsa.messages* file for specific details.

unable to create the save set on the server

The call to create the save set on the NetWorker server failed. Check the */nsr/applogs/xbsa.messages* file for specific details.

unable to end the current read session

During a load database or load transaction log operation, the read session of the data from NetWorker could not be closed. Check the */nsr/applogs/xbsa.messages* file for specific details.

Unable to find backup of the (database or transaction log) SYBASE:/server_name/database_name. Check the command line for errors in the instance or database name or use nsrinfo to see which save sets are available.

The item to be loaded could not be found. Use the **nsrinfo** program to check that the object-owner for the backup is the same as the process that launched the Sybase Backup Server and that backups exist for this database.

Unable to find full backup of the database *database_name* for the time supplied.
Unable to find incremental backup of the database *database_name* for the time supplied.
Unable to find backup of the database *database_name* for the time supplied.

No backup could be found in the NetWorker server. If no time was supplied, the time used is the current time, which means that no backup exists. Use the **nsrinfo** program to check which backups are available and make sure that the object owner shown there is the same as the user ID that launched the Sybase Backup Server.

unable to parse stripe specifier

The **isql** command line had a poorly formatted stripe specifier. For details on the specific syntax for **isql** commands, see [“Syntax for Dump and Load Commands from the isql Command Line” on page 109](#).

unable to read the requested number of bytes from the save set

During a load database or load transaction log operation, the save set could not be read. Check the */nsr/applogs/xbsa.messages* file for specific details.

unable to send data to save set

During a database or transaction log dump, the data could not be written to the save set. Check the */nsr/applogs/xbsa.messages* file for specific details.

unknown backup type supplied
The backup type supplied from NetWorker server was neither a database or a transaction log.

NetWorker XBSA and libbms Error Messages

During a backup or recovery, NetWorker attempts to record error messages generated by the XBSA library in the file assigned to the NSR_DEBUG_FILE environment variable. If the assigned location is invalid or unreachable, NetWorker writes the message to one of the following locations:

- The */nsr/applogs/xbsa.messages* alternate messages directory created during installation
- The directory assigned to the NSR_TMPDIR environment variable
- The */tmp* directory, if NSR_TMPDIR is not set

For descriptions of the NetWorker XBSA environment variables and values you can assign to them in the **nsrsyb** script, see “Appendix A: XBSA Environment Variables” on page 53.

NetWorker XBSA error messages appear in the format

XBSA-1.0 bms-1. process_id day month date hh:mm:ss year
function_name: BSA_RC_message_code: message

The **libbms** error messages are also written to the */nsr/applogs/xbsa.messages* file as well as reported to the Sybase Backup Server. The Sybase Backup Server prints the messages in the *stdout* file and logs them in the Sybase Backup Server error log.

BSA_RC_ABORT_ACTIVE_NOT_FOUND No active object matched the name that was specified for a BSAMarkObjectInactive
No active object matching the given search parameters was found in the NetWorker server being used by the NetWorker XBSA session.

BSA_RC_ABORT_SYSTEM_ERROR System detected error due to <i>explanation</i> . Operation aborted
A general system error occurred within a NetWorker XBSA function call. This error is returned for all NetWorker errors that do not map cleanly to XBSA errors.

BSA_RC_APP_OBJECTOWNER_TOO_LONG The appObjectOwner field contained too many characters ($n \geq n$)

The appObjectOwner field of an ObjectOwner parameter contains too many characters and might be corrupt.

BSA_RC_AUTHENTICATION_ERROR There was an authentication failure for ObjectOwner *ownername*

The routine failed to authenticate a BSAObjectOwner with NetWorker server used by the NetWorker XBSA session. The code is returned by the routine BSASetEnvironment to allow for the possibility of changing NetWorker servers during a single session by changing the value of the NSR_SERVER environment option. For more details about available settings, see “[Appendix A: XBSA Environment Variables](#)”. NetWorker permits all users to back up data and recover their files without passwords, so this should not occur.

BSA_RC_BAD_CALL_SEQUENCE The sequence of API calls is incorrect. Must call *item1* before *item2*

An API call sequence was made that does not conform to the XBSA Data Movement API State Diagram document.

BSA_RC_BAD_HANDLE The handle used to associate this call with a previous BSAInit() call is invalid because *explanation*

The value passed into the function for bsaHandle contained a NULL pointer.

BSA_RC_BAD_PARAMETER received parameter *parm* with value *value*, which is invalid

An invalid parameter was received.

BSA_RC_BSA_OBJECTOWNER_TOO_LONG The bsaObjectOwner field contained too many characters ($n \geq n$)

The appObjectOwner field of an ObjectOwner parameter contains too many characters and might be corrupt.

BSA_RC_BUFFER_TOO_SMALL Buffer is too small to hold the object entry to be returned. *n* bytes required for the object entry

The buffer is too small to hold the object entry to be returned.

BSA_RC_COPYGPNAME_TOO_LONG The copyGpName field contained too many characters (*n* >= *n*)

The copyGpName field in one of the supplied structures contained more BSA_MAX_COPYGPNAME characters, and the structure could not be used for the requested operation.

BSA_RC_DESCRIPTION_TOO_LONG The description field contained too many characters (*n* >= *n*)

The Description field in one of the supplied structures contained more than BSA_MAX_DESC characters, and the structure could not be used for the requested operation.

BSA_RC_INVALID_COPYTYPE the copyType field contained an unrecognized value of *n*

The copyType field in one of the supplied structures has a value that is not in the NetWorker XBSA libraries implementation of this enumerated type.

BSA_RC_INVALID_DATABLOCK the dataBlock parameter contained inconsistent values: bufferLength: *n*, bufferPtr: *n*, numBytes: *n*

The fields of a supplied DataBlock parameter are not internally consistent. This can occur when the bufferLen field is less than the numBytes field when data is being sent, or when the bufferLen field is nonzero and the bufferPtr field is NULL.

BSA_RC_INVALID_KEYWORD an entry in the environment structure is invalid (*variable=value*)

One of the environment strings passed into the function did not have a valid structure. The value structure of an environment keyword is KEYWORD = VALUE, where KEYWORD is a white space delimited string and VALUE is a white space delimited string followed by a null terminator. This can indicate a number of possible errors:

- The KEYWORD was not in the reserved word list. This error is not returned by the NetWorker XBSA libraries because other environment variables might be passed into the library along with valid keywords.
- The KEYWORD and VALUE strings were not separated by a '=' character. This type of error is also used to detect environment vectors that are not properly terminated with a (char *)NULL entry, as well as invalid KEYWORD VALUE pair formats.
- The VALUE string was invalid.
- The VALUE string could not be validated, as in the case of a hostname string that could not be found by the **gethostbyname()** function.

BSA_RC_INVALID_OBJECTSTATUS the objectStatus field contained an unrecognized value of *n*

The objectStatus field in one of the supplied structures has a value that is not in the NetWorker XBSA libraries' implementation of this enumerated type.

BSA_RC_INVALID_OBJECTTYPE the objectType is invalid (*n*)

One of the object type parameters was either passed in directly or contained in one of the following structures: ObjectDescriptor QueryDescriptor was not in the range of BSAObjectType_ANY to BSAObjectType_DIRECTORY.

BSA_RC_INVALID_TIME a time field contained an unrecognized value of *n*

An invalid time value was received.

BSA_RC_INVALID_VERSION the version field contained an unrecognized value of *n*

The version for a parameter passed into the function is not supported by this version of NetWorker XBSA. For routines that receive multiple parameters containing a version field, it does not indicate which parameter is not supported.

BSA_RC_LGNAME_TOO_LONG The LGName field contained too many characters (*n* >= *n*)

An LGName, passed in to the function, contained more than BSA_MAX_LGNAME_SIZE characters and might be corrupt. For routines that require multiple LGName parameters, it does not indicate which token was invalid.

BSA_RC_MATCH_EXISTS object matching the specified predicate already exists

The object already exists in the NetWorker server being used by the NetWorker XBSA session and that the requested operation cannot be completed.

BSA_RC_MORE_DATA more data is available. Data can be obtained through BSAGetData() or BSAGetNextQueryObject()

This has two meanings in the XBSA Data Movement API:

- *Object Data Retrieval*—There is more data available for an object being read from the NetWorker server than is being used by the NetWorker XBSA session. Use BSAGetData to retrieve the next DataBlock from the NetWorker server (see also BSA_RC_BUFFER_TOO_SMALL and BSA_RC_NO_MORE_DATA). This message is not returned by the BSAGetObjectF function because all data for an object is written to a file descriptor by this function.
- *Query Result Retrieval*—There are more objects matching the requested query descriptor from the NetWorker server than is being used by the NetWorker XBSA session. Use BSAGetNextQueryObject to retrieve the next object descriptor from Backup Services (see also BSA_RC_NO_MORE_DATA).

BSA_RC_NO_MATCH The ResourceType predicate value of *D* does not match the reference value of *L*

The client index and media database are out of synch. To resynchronize the client index and media database, run the **nsrck -X** command. Alternatively, you can wait for NetWorker to run **nsrck** automatically.

BSA_RC_NO_MATCH The *variable* predicate value of *value* does not match the reference value of *variable*

No objects matching the specified QueryDescriptor were found in the NetWorker server being used by the NetWorker XBSA session.

BSA_RC_NO_MORE_DATA there is no more data for the current object

This has two meanings in the XBSA Data Movement API:

- *Object Data Retrieval*—This is used when all the data for an object being retrieved from a NetWorker server was placed into the given DataBlock parameter for a function call (see also BSA_RC_MORE_DATA).
- *Query Result Retrieval*—This is used when the last (or only) object matching a query is returned to the caller (see also BSA_RC_MORE_DATA).

BSA_RC_NULL_APIVERSION an ApiVersion pointer is required

A pointer to an ApiVersion structure, passed into the function, was NULL and is required as input.

BSA_RC_NULL_BUFFER an buffer pointer is required

This is not used by NetWorker XBSA. A NULL buffer when reading an object's data (BSAGetData, BSAGetObject) results in no bytes being read and a BSA_RC_MORE_DATA code being returned.

BSA_RC_NULL_DATABLOCK a data block pointer is required

The DataBlock pointer parameter for the called function was NULL. The caller is responsible for allocating and passing in the DataBlock structure to the NetWorker XBSA library (see also BSA_RC_NULL_BUFFER and BSA_RC_INVALID_DATABLOCK).

BSA_RC_NULL_ENVIRONMENT an environment pointer is required

This is not used by NetWorker XBSA. An environment vector parameter that is NULL is not processed.

BSA_RC_NULL_NEWTOKEN a value must be entered for the new token. The old token has expired

The SecurityToken parameter, newToken, was found to be NULL and is required as input. See also BSA_RC_NULL_SECURITYTOKEN.

BSA_RC_NULL_OBJECTDESCRIPTOR an ObjectDescriptor pointer is required

The SecurityToken parameter, newToken, was found to be NULL and is required as input. See also BSA_RC_NULL_SECURITYTOKEN.

BSA_RC_NULL_OBJECTNAME an object name is required

The ObjectName parameter passed into the called function was NULL.

BSA_RC_NULL_OBJECTOWNER an ObjectOwner pointer is required

A pointer to an object-owner structure was NULL and is required as input.

BSA_RC_NULL_POINTER a required pointer parameter is NULL

The NetWorker XBSA library does not return this code. Instead, specific codes indicating that a required parameter was NULL are returned:

BSA_RC_NULL_APIVERSION
BSA_RC_NULL_BUFFER
BSA_RC_NULL_COPYGPNAME
BSA_RC_NULL_COPYID
BSA_RC_NULL_DATABLOCK (BSA_RC_NULL_DATABLKPTR)
BSA_RC_NULL_ENVIRONMENT
BSA_RC_NULL_LGNAME
BSA_RC_NULL_NEWTOKEN
BSA_RC_NULL_OBJECTDESCRIPTOR
BSA_RC_NULL_OBJECTNAME
BSA_RC_NULL_OBJECTOWNER
BSA_RC_NULL_OLDTOKEN
BSA_RC_NULL_QUERYDESCRIPTOR
BSA_RC_NULL_RULEID
BSA_RC_NULL_SECURITYTOKEN
BSA_RC_NULL_STREAM

BSA_RC_NULL_SECURITYTOKEN an SecurityToken pointer is required

A pointer to a SecurityToken parameter is NULL and is required as input. The NetWorker XBSA library uses this internally and should not be seen in normal use. The more specific codes BSA_RC_NULL_NEWTOKEN and BSA_RC_NULL_OLDTOKEN are used, as appropriate.

BSA_RC_OBJECTINFO_TOO_LONG The objectInfo field contained too many characters ($n \geq n$)

The ObjectInfo parameter passed into the function, either directly or in one of the following data structures, was found to have more than BSA_MAX_OBJINFO characters: ObjectDescriptor

BSA_RC_OBJECTSPACENAME_TOO_LONG The objectSpaceName field contained too many characters ($n \geq n$)

The string objectSpaceName contains more than BSA_MAX_OBJECTSPACENAME characters in an ObjectName structure.

BSA_RC_PATHNAME_TOO_LONG The pathName field contained too many characters ($n \geq n$)

The string pathname contains more than BSA_MAX_PATHNAME characters in an ObjectName structure.

BSA_RC_RESOURCE_TYPE_TOO_LONG The resourceType field contained too many characters ($n \geq n$)

The string resourceType contains more than BSA_MAX_RESOURCE_TYPE characters and might be corrupt.

BSA_RC_SECURITY_TOKEN_TOO_LONG The securityToken field contained too many characters ($n \geq n$)

A SecurityToken, passed in to the function, contained more than BSA_MAX_SECURITY_TOKEN characters and might be corrupt. For routines that require multiple tokens, it does not indicate which token was invalid.

BSA_RC_SUCCESS the function was successful

The called function did not fail and is returned by all NetWorker XBSA function calls.

BSA_RC_TRANSACTION_ABORTED the transaction was aborted

The current transaction was aborted by the BSAEndTxn function call. A transaction can either be aborted by an internal error or by user request through the Vote parameter to this function.

Appendix C: Command Line Utilities

This appendix provides information about the NetWorker Module command line utilities—**nsrsybsv**, **nsrsybrc**, and **nsrsybrc**—and their syntax. Additionally, this appendix presents details about how the behavior of the backup command options is affected by the backup level, and about the correct syntax for using the **dump** and **load** commands from the **isql** command line.

Syntax for the NetWorker Module Command Line Utilities

This section includes the syntax and options that you can use with the NetWorker Module command line utilities.

Conventions

The conventions presented in the command line are as follows:

- The command option not residing in any brackets must always be present in the command.
- The command option in square “[]” brackets is optional.
- For the command options in curly “{}” brackets, one of the items must exist in the command.

Sybase Usernames and Passwords

The NetWorker Module command line utilities—**nsrsybsv**, **nsrsybrc**, and **nsrsybrc**—search for Sybase usernames and passwords in the following locations and in the following order:

1. The command line.
2. The client resource that matches the client, group, and save set names specified at the command line.
3. The \$USER and \$PASSWORD environment variables.

nsrsybsv

nsrsybsv [-CGqRTv] [-s *server*] [-N *name*] [-b *pool*] [-g *group*]
[-l *level*] [-U *user* [-P *password*]] [-D *level*]
SYBASE: /*instance*[/*dbname*]

- C** Use XBSA compression.
- G** Perform a dump with the NO_LOG option.
- q** Quiet mode.
- R** Perform a dump without truncating the transaction log. This performs a dump transaction with no_truncate if it is an incremental backup, and it performs a dump database without the preceding dump transaction for a full backup.
- T** Perform a dump with the truncate_only option. This is valid with a full or an incremental backup.
- v** Verbose mode.
- s** Specifies the NetWorker server to use. If gethostbyname for the server fails, an error occurs. This overrides NSR_SERVER.
- N** Overrides the constructed save set name.
- b** Pool used for both data volume pool and log volume pool.
- g** Backup group. This overrides NSR_GROUP.
- l** Backup(save) level. Valid options are full and incremental. The default is full.
- U** Username to use to log in to Sybase server.
- P** Password to use to log in to Sybase server.
- D** Specifies the level of debugging (valid values are 1–5, with 5 being the most detailed).

nsrsybcc

```
nsrsybcc [-qv] [-s server] [-c client] [-g group] [-D level] [-U user] [-P password] [-o ckal] [-o ckdb] [-o ckdbnoidx] [-o ckcat] [-o ckstor] SYBASE:/instance[/dbname]
```

- q** Quiet mode.
- v** Verbose mode.
- s** Specifies the NetWorker server to use. If gethostbyname for the server fails, an error occurs. This overrides NSR_SERVER.
- c** NetWorker client hostname. Used in imported recoveries.
- g** Backup group. This overrides NSR_GROUP.
- D** Specifies the level of debugging (valid values are 1–5, with 5 being the most detailed).
- U** Username to use to log in to Sybase server.
- P** Password to use to log in to Sybase server.
- o ckdb | -o ckdbnoidx**
dbcc checkdb (... , skip_ncindex)
- o ckal**
dbcc checkalloc
- o ckcat**
dbcc checkcatalog
- o ckstor**
dbcc checkstorage (System 11.5 and later—dbccdb must also be set up)

nsrsybrc

```
nsrsybrc [-qv] [-s server] [-t date] [-c client] [-U user]  
[-P password] [-d SYBASE:/instance[/dbname]]  
SYBASE:/instance[/dbname]
```

- | | |
|----------|--|
| q | Quiet mode. |
| v | Verbose mode. |
| s | Specifies the NetWorker server to use. If gethostbyname for the server fails, an error occurs. This overrides NSR_SERVER. |
| t | Specifies the time to recover data to. For Sybase server 11.5 or later, this option is used as the time for point-in-time recovery. |
| c | NetWorker client hostname. Used in imported recoveries. |
| U | Username to use to log in to Sybase server. |
| P | Password to use to log in to Sybase server. |
| d | Destination for recovery. If none is provided, the database is recovered to its original location, overwriting any existing data. If one is provided, the database (or instance) is recovered to another database (or instance). |

Backup Options

You can specify how NetWorker backs up the transaction logs with the following options:

- **-T** (truncate_only)
- **-G** (with no_log)
- **-R** (with no_truncate)

The behavior of these backup options is affected by the backup level (full or incremental) and by whether the database and transaction logs are on the same or separate devices. For information on how and when to use these backup options, refer to your Sybase administration guide.

Full Backups on Separate Devices

If the data and transaction log are on separate devices and you set one of these options, the following occurs:

- If the **-G** (with `no_log`) option is set, the NetWorker Module truncates the transaction log without logging the transaction. The NetWorker Module then performs a full backup.
- If the **-T** (`truncate_only`) option is set, the NetWorker Module truncates the transaction log. The NetWorker Module then performs a full backup.
- If the **-R** (`no_truncate`) option is set, the NetWorker Module backs up the transaction log but does not truncate it. The NetWorker Module then performs a full backup.
- If no option is set, the NetWorker Module just performs a full backup.

For more information on full backups, see [“What Is a Full NetWorker Backup?” on page 23](#).

Full Backups on the Same Device

If the data and transaction log are on the same device and you set one of these options, the following occurs:

- If the **-G** (with `no_log`) option is set, the NetWorker Module displays an error message saying that **-G** is not a valid option in this case and it is ignored. The NetWorker Module then performs a full backup, which in this case also truncates the transaction log.
- If the **-T** (`truncate_only`) options is set, the NetWorker Module truncates the transaction log. The NetWorker Module then performs a full backup, but does not truncate the transaction log.
- If the **-R** (`no_truncate`) options is set, the NetWorker Module performs a full backup, but does not truncate the transaction log.
- If no option is set, the NetWorker Module performs a full backup, which in this case also truncates the transaction log.

Incremental Backups on Separate Devices

If the data and transaction log are on separate devices and you set one of these options, the following occurs:

- If the **-G** (with `no_log`) option is set, the NetWorker Module displays an error saying that the **-G** option is not supported with incremental backups, and that a full backup must be performed instead. The NetWorker Module then performs an incremental backup.
- If the **-T** (`truncate_only`) option is set, the NetWorker Module displays an error saying that the **-T** option not supported with incremental backups, and that a full backup must be performed instead. The NetWorker Module then performs an incremental backup.
- If the **-R** (`no_truncate`) option is set, the NetWorker Module backs up the transaction log but does not truncate it.
- If no option is set, the NetWorker Module performs an incremental backup.

For more information on incremental backups, see [“What Is an Incremental NetWorker Backup?” on page 23](#).

Incremental Backups on the Same Device

The backup options are not valid for incremental backups of data and transaction log on the same device.

1. If the backup is of an entire database instance, the NetWorker Module displays a message saying that it is unable to perform an incremental backup, and skips this backup. However, this operation will perform an incremental backup of any databases which exist on different devices.
2. If the backup is of a single database, the NetWorker Module displays a message saying that it is unable to perform an incremental backup, and it performs a full backup instead. For more information, see [“Full Backups on the Same Device” on page 107](#).

Syntax for Dump and Load Commands from the isql Command Line

Use the NetWorker Module **nsrsybsv** and **nsrsybrc** programs rather than the Sybase **dump** and **load** commands to back up and recover NetWorker save sets. If you use the **dump** and **load** commands from the **isql** command line, you must follow a specific syntax.

Note: After setting XBSA variables in the Sybase shell, restart the Sybase backup server.

Dump Database Syntax

To dump a database from the **isql** command line, use the following syntax for each database to be dumped:

```
dump database database_name to "bms::"
```

If you want to specify the hostname and server name or specify that a notification should be sent to the operator console, use the following syntax:

```
dump database database_name to  
"bms::[[host_name][.[server_name][.[database_name]]]]"  
[with notify = {client| operator_console}]
```

Dump Transaction Log Syntax

To dump a transaction log from the **isql** command line, use the following syntax for each transaction log to be dumped:

```
dump transaction database_name to "bms::"
```

If you want to specify the hostname and server name or specify that a notification should be sent to the operator console, use the following syntax:

```
dump transaction database_name to  
"bms::[[host_name][.[server_name][.[database_name]]]]"  
[with {{truncate_only | no_log | no_truncate},}  
[notify = {client| operator_console}]]]
```

Finding the Timestamp for a Save Set

If you want to use a specific timestamp when you load a database or transaction, you can find the timestamp for a save set using either of the following methods:

- Run the following command to get a list of all the Sybase save sets for the NetWorker client.

```
nsrinfo -X all -n sybase client_name
```

- Use the Save Set Recover window in the NetWorker Administration program to select the save set you want to recover. The date and time are displayed in the Instances window. Then use the **load** command at the **isql** command line.

If you do not specify a timestamp, the most recent backup is recovered.

You must specify a timestamp in the following format:

YYYYMMDDhhmmss111

In this format, *Y* indicates the year, *M* indicates the month, *D* indicates the day, *h* indicates the hour, *m* indicates the minutes, *s* indicates the seconds, and *l* indicates the milliseconds. The *l* millisecond position is optional; alternatively, you can enter **000** for the milliseconds.



Important: You cannot use the Save Set Recover window to recover Sybase data. Use the **nsrsybrc** program to recover your databases and transaction logs. For details, see [“Recovering Data with nsrsybrc” on page 44](#).

Load Database Syntax

To load the most recent database backup from the **isql** command line, use the following syntax:

```
load database database_name from “bms:”
```

If you want to specify the hostname and server name, a timestamp, or that a notification should be sent to the operator console, use the following syntax:

```
load database database_name from  
“bms: : [[host_name] [. [server_name] [. [database_name]  
[. [timestamp]]]]]”  
[with {[headeronly,][notify = {client|  
operator_console}]}
```

For example, to load the master database from a backup performed on October 31, 1999 at 11:52:30 a.m., the syntax would be:

```
load database master from “bms:...1999103111523000”
```



Important: After you load the database, be sure to bring it back online.

Load Transaction Log Syntax

To load the most recent transaction log backup from the **isql** command line, use the following syntax:

```
load transaction database_name from "bms::"
```

If you want to specify the hostname and server name, a timestamp, or that a notification should be sent to the operator console, use the following syntax:

```
load transaction database_name from  
"bms::[[host_name][.server_name][.database_name]  
[.timestamp]]]"  
[with {[headeronly,][notify = {client|  
operator_console}]}
```


Glossary

This glossary contains terms and definitions found in this manual.

<i>Adaptive Server Enterprise</i>	Sybase-Adaptive Server Enterprise, considered by NetWorker as one or more instances of a Sybase database server.
<i>API</i>	An abbreviation for Application Programming Interface, an agreed-upon set of computer library routines to accomplish a task.
<i>autochanger</i>	A device that has the ability to move media among various components (including slots, media drives, media access ports, and transports) located in the device. Autochangers automate the media loading, labeling, and mounting during backups and recovers.
<i>backup group</i>	A NetWorker client or group of clients configured to start backing up files to the NetWorker server at a designated time of day.
<i>bootstrap</i>	At the end of a backup, NetWorker saves the server's media database, NetWorker configuration files, and part of the server index to a special save set named bootstrap. The bootstrap provides vital information needed for recovering from a disk crash.
<i>daemon</i>	A program that is not invoked explicitly, but lies dormant waiting for specified conditions to occur.

<i>fileserver</i>	A computer with disks that provide file storage to other computers on a network.
<i>filesystem</i>	1. A sub-tree of a UNIX file tree that is on a specific disk partition or other mount point. 2. A method for storing files.
<i>instance name</i>	The name of an individual SQL Server or Adaptive Server Enterprise running on a host. There can be several instances of a database server running on a single host.
<i>media database</i>	A database that contains indexed entries about the storage volume location and the life-cycle status of all data and volumes managed by the NetWorker server.
<i>media manager</i>	The NetWorker component that tracks save sets to backup volumes.
<i>mirroring</i>	A storage strategy that gives users uninterrupted access should the primary database fail.
<i>multiplex</i>	A method of transmitting multiple streams of data simultaneously through the same channel.
<i>NetWorker client</i>	A computer on a network running the NetWorker client software that can access storage management services from a NetWorker server. Clients can be workstations, PC desktops, file servers, or database servers with gigabytes of data.
<i>NetWorker server</i>	A computer on a network running the NetWorker server software, containing online client indexes, that provides storage management services to clients on a network.
<i>online client indexes</i>	Databases on a NetWorker server that contain information pertaining to client backups and backup volumes.

<i>resource</i>	A component of the NetWorker software that describes the NetWorker server and its clients. Devices, schedules, clients, groups, and pools are all examples of NetWorker resources. Each resource contains a list of attributes, defining the parameters to use for the configured NetWorker resource. Use the nwadmin or nsradmin program to configure NetWorker resources and their attributes.
<i>RPC</i>	An abbreviation for Remote Procedure Call, a protocol that allows a program running on one host to cause code to execute on another host without needing explicitly coded instructions.
<i>save set</i>	A set of files or a filesystem that NetWorker saves to media during a backup session. Save sets are assigned a save set ID, an internal number that identifies the backup session to NetWorker for subsequent recovery to primary disk.
<i>shell prompt</i>	1. UNIX command line prompt, either % or \$ (or # for superuser). 2. DOS command line prompt, for example, C:\.
<i>SQL Server</i>	Sybase-SQL Server, considered by NetWorker as one or more instances of a Sybase database server.
<i>storage manager</i>	An application that manages the storage devices and media used for backup and recovery requests. <i>NetWorker Module</i> for Sybase is a storage manager that connects NetWorker to the Sybase Backup Server through XBSA and a shared libbms library.
<i>XBSA</i>	An abbreviation for X/Open Backup Services Application Programming Interface, which connects NetWorker functionality to the Sybase Backup Server.

Index

A

Adaptive Server
 instance 20, 28
Adaptive Server Enterprise
 default settings for nsrsybcc 36
Admin Guide, typefaces and symbols
 explained 12

B

backup
 command options 106
 entire database instance 28
 group 21
 how nsrsyb works 17
 levels 23
 on-demand
 dump from isql command line 38
 NetWorker indexes 36
 nsrsybcc 36
 nsrsybsv 35
 options 34
 reports 29
 required Sybase roles 20
 scheduled
 defined 23
 selected databases 28
Backup Server. *See* Sybase Backup Server
BACKUP_OPT, environment variable 32
bootstrap
 description of 113
 recovering data 42
 saving after on-demand backup 37
browse policy 26

C

client resource
 creating for database server 27
command line utilities
 options 106
 Sybase username and password 103
 syntax 103
configuring. *See* customizing
customizing
 backup groups 21
 backup schedules 22
 client for database server 27
 label templates 25
 pools 24

D

daemon, description of 113
database consistency check program 36
database server, configuring as NetWorker
 client 27
Default backup group, description 21
disaster recovery. *See* recovering
dump command
 from isql 38
 syntax for isql command line 109
 See also nsrsybsv
dump database. *See* full backups
dump transaction logs. *See* incremental
 backups.

E

e-mail completion notice 29
environment variables, NetWorker XBSA
 NSR_BACKUP_LEVEL 55
 NSR_CLIENT 55

Index

F - L

NSR_COMPRESSION 56
 NSR_DATA_VOLUME_POOL 56
 NSR_DEBUG_FILE 57
 NSR_DEBUG_LEVEL 57
 NSR_GROUP 58
 NSR_LOG_VOLUME_POOL 58
 NSR_NO_BUSY_ERRORS 59
 NSR_SAVESET_NAME 59
 NSR_SERVER 60

error messages

NetWorker file 61
 NetWorker XBSA 57, 93
 NetWorker XBSA and libbms 93–101
 nsrck 62–63
 nsrexecd 63–64
 nsrindexd 64
 nsrmdbd 65
 nsrsyb 66
 nsrsybcc 66–68
 nsrsybrc 69–74
 nsrsybsv 75–82
 savegrp 84, 87, 89
 Sybase Backup Server and libbms 90–93
 while saving data 62–65, 82–89

F

full backups

defined 23
 dump from isql command line 38
 load from isql command line 47
 no_log (-G) option 106
 no_truncate (-R) option 106
 on-demand 35
 truncate_only (-T) option 106
See also incremental backups

G

Group Control feature, using 22
 groups, backup 21

H

how to

modify Sybase in nsrsyb 31
 perform disaster recovery 48
 recover expired data 41
 set up a volume pool 24, 26
 set up database server as NetWorker client 27
 use nsrinfo to generate report of index entries 30
 view results of scheduled backup 29

I

imported recoveries 46

incremental backups

dump command from isql 38
 load from isql command line 47
 threshold procedure 38
 when allowed 23

indexes

manually managing 26
 nsrck messages 62–63
 nsrindexd messages 64
 reducing disk space used by 26
 saving server after on-demand backup 37
 viewing entries in 42

instance backup

specifying in client resource 28

isql command line

dump command 38
 load command 47
 syntax 110–111

L

libbms

how NetWorker uses 16
 NetWorker XBSA error messages 93–101

Sybase Backup Server error messages
90–93

load command
from isql 47
syntax for isql command line 110–111
See also nsrsybrc

M

master database, recovering 49

messages. *See* error messages

modifying
BACKUP_OPT variable 32
nsrsyb shell script 30
POSTCMD variable 31
PRECMD variable 31
Sybase variable 31
USE_CONSISTENCY_CHECK
variable 31

N

NetWorker
client, description of 27, 114
error message file 61
featured highlights 16
resource, description of 115
server, description of 114

NetWorker XBSA
error message file 57, 93
error messages 93–101

no_log (-G) option 106

no_truncate (-R) option 106

notification, sending to client 29

NSR_BACKUP_LEVEL 55

NSR_CLIENT 55

NSR_COMPRESSION 56

NSR_DATA_VOLUME_POOL 56

NSR_DEBUG_FILE 57

NSR_DEBUG_LEVEL 57

NSR_GROUP 58

NSR_LOG_VOLUME_POOL 58

NSR_NO_BUSY_ERRORS 59

NSR_SAVESET_NAME 59

NSR_SERVER 60

nsrck error messages 62–63

nsrexecd error messages 63–64

nsrindexd error messages 64

nsrinfo 30

nsrmmdbd error messages 65

nsrsyb
connection to Sybase 17
default Sybase environment variable
settings 31
error messages 66
functionality overview 17

nsrsybcc
command line syntax 103
default settings 36
error messages 66–68
running from command line 36
Sybase roles 20
Sybase username and password 103
USE_CONSISTENCY_CHECK
variable 31

nsrsybrc
command line syntax 103
error messages 69–74
functionality overview 18
recovering data with 44
Sybase roles 20
Sybase username and password 103

nsrsybsv
command line syntax 103
error messages 75–82
NetWorker indexes 36
running from command line 35
Sybase roles 20
Sybase username and password 103

O

on-demand backup
NetWorker indexes 36
nsrsybcc 36
nsrsybsv 35

Index

P - T

saving NetWorker bootstrap 37
 using the dump command 38
 online indexes, description of 114

P

password, Sybase 103
 permissions, Sybase 20
 point-in-time recoveries 45
 pools
 customizing 24, 26
 description of 23
 using 24, 34
 valid pool types 23
 POSTCMD, description of 31
 PRECMD, description of 31
 preconfigured settings
 NetWorker Default group 21
 NetWorker XBSA environment
 variables
 NSR_BACKUP_LEVEL 55
 NSR_CLIENT 55
 NSR_COMPRESSION 56
 NSR_DATA_VOLUME_POOL 56
 NSR_DEBUG_FILE 57
 NSR_DEBUG_LEVEL 57
 NSR_GROUP 58
 NSR_LOG_VOLUME_POOL 58
 NSR_NO_BUSY_ERRORS 59
 NSR_SAVESET_NAME 59
 NSR_SERVER 60

R

recoveries
 imported 46
 options 43
 point-in-time 45
 recovering
 databases not on the master device 50
 expired data 41

from a disaster 48
 how nsrsybrc works 18
 master database 49
 with nsrsybrc 44
 retention policy 26

S

save sets
 finding timestamp 109
 specifying a database 28
 specifying entire database server 28
 specifying multiple instances 28
 savegrp
 error messages 84, 87, 89
 scheduled backup
 how nsrsyb works 17
 starting immediately 22
 schedules, customizing 22
 select database backup, specifying 28
 SQL Server
 default settings for nsrsybrc 36
 instance 20, 28
 Sybase
 username and password 103
 Sybase Backup Server
 error message file 90
 error messages 90–93
 Sybase environment variable, changing 31
 syntax
 command line utilities 103
 isql commands 109

T

threshold procedure
 last-chance 39
 sample for NetWorker Module 38
 timestamp 109
 transaction log dumps, *See* incremental backups.
 truncate_only (-T) option 106

types of recoveries
 imported 46
 point-in-time 45

NSR_NO_BUSY_ERRORS 59
NSR_SAVESET_NAME 59
NSR_SERVER 60

U

USE_CONSISTENCY_CHECK, description
 of 31
username, Sybase 103

V

variable
 BACKUP_OPT 32
 changing 31
 NSR_BACKUP_LEVEL 55
 NSR_CLIENT 55
 NSR_COMPRESSION 56
 NSR_DATA_VOLUME_POOL 56
 NSR_DEBUG_FILE 57
 NSR_DEBUG_LEVEL 57
 NSR_GROUP 58
 NSR_LOG_VOLUME_POOL 58
 NSR_NO_BUSY_ERRORS 59
 NSR_SAVESET_NAME 59
 NSR_SERVER 60
 POSTCMD 31
 PRECMD 31
 USE_CONSISTENCY_CHECK 31
volume label template 25
volume pools. *See* pools

X

XBSA environment variables
 NSR_BACKUP_LEVEL 55
 NSR_CLIENT 55
 NSR_COMPRESSION 56
 NSR_DATA_VOLUME_POOL 56
 NSR_DEBUG_FILE 57
 NSR_DEBUG_LEVEL 57
 NSR_GROUP 58
 NSR_LOG_VOLUME_POOL 58

