



Installation Guide

Adaptive Server Enterprise

12.5

Linux/Intel

DOCUMENT ID: 30119-01-1250-03

LAST REVISED: July 2002

Copyright © 1989-2002 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase, the Sybase logo, AccelaTrade, ADA Workbench, Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise, Adaptive Server Enterprise Monitor, Adaptive Server Enterprise Replication, Adaptive Server Everywhere, Adaptive Server IQ, Adaptive Warehouse, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-FORMS, APT-Translator, APT-Library, Backup Server, BizTracker, ClearConnect, Client-Library, Client Services, Convoy/DM, Copernicus, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, Distribution Director, e-ADK, E-Anywhere, e-Biz Integrator, E-Whatever, EC-GATEWAY, ECMAP, ECRTP, eFulfillment Accelerator, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, eProcurement Accelerator, EWA, Financial Fusion, Financial Fusion Server, Gateway Manager, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, MainframeConnect, Maintenance Express, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, MySupport, Net-Gateway, Net-Library, New Era of Networks, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, PB-Gen, PC APT Execute, PC Net Library, Power++, power.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, PowerJ, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Rapport, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RW-DisplayLib, S-Designor, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILS, smart.partners, smart.parts, smart.script, SQL Advantage, SQL Anywhere, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Edit/TPU, SQL Everywhere, SQL Modeler, SQL Remote, SQL Server, SQL Server Manager, SQL SMART, SQL Toolset, SQL Server/CFT, SQL Server/DBM, SQL Server SNMP SubAgent, SQL Station, SQLJ, STEP, SupportNow, S.W.I.F.T. Message Format Libraries, Sybase Central, Sybase Client/Server Interfaces, Sybase Financial Server, Sybase Gateways, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, TradeForce, Transact-SQL, Translation Toolkit, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, Viewer, Visual Components, VisualSpeller, VisualWriter, VQL, WarehouseArchitect, Warehouse Control Center, Warehouse Studio, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web Deployment Kit, Web.PB, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, XA-Server and XP Server are trademarks of Sybase, Inc. 04/02

Unicode and the Unicode Logo are registered trademarks of Unicode, Inc.

All other company and product names used herein may be trademarks or registered trademarks of their respective companies.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Sybase, Inc., 5000 Hacienda Drive, Dublin, CA 94568.

Contents

| | |
|------------------------------|---|
| About This Book | 9 |
| CHAPTER 1 | Overview |
| | User roles |
| | Product descriptions |
| | Installation and setup utilities |
| | PC-Client product descriptions |
| | Sybase installation directory |
| CHAPTER 2 | Installation Requirements |
| | System requirements |
| | System requirements |
| | Product disk space requirements |
| | Language module sizes |
| | Adaptive Server specifications |
| | Adaptive Server devices and system databases |
| | master device |
| | sybsystemdb device and database |
| | sysprocsdev device |
| | Optional devices and databases |
| | Determining the location, type, and size of a database device |
| CHAPTER 3 | Sybase Software Asset Management (SySAM) |
| | SySAM overview |
| | Adaptive Server features licensed through SySAM |
| | How SySAM works |
| | Types of SySAM systems |
| | SySAM in the network environment |
| | SySAM administration |
| | Verify the software is running |
| | Starting the software manually |
| | Adding feature licenses |

| | | |
|----------------------|--|---------------|
| CHAPTER 4 | Installing Sybase Servers | 41 |
| | Overview | 41 |
| | Installation Definitions | 42 |
| | Pre-installation tasks | 42 |
| | Installing the Linux dummy-network interface | 46 |
| | Installing server components | 47 |
| | Installing using Red Hat Package Manager (RPM)..... | 47 |
| | Registering Adaptive Server features..... | 51 |
| | Configuring the servers | 53 |
| | Installing the documentation..... | 63 |
| CHAPTER 5 | Post-Installation Tasks..... | 67 |
| | Setting environment variables..... | 67 |
| | Setting environment variables | 69 |
| | Verifying that servers are running | 69 |
| | Verifying that you can connect to servers | 70 |
| | Connecting to Adaptive Server via Sybase Central | 70 |
| | Setting the System Administrator password | 71 |
| | Installing sample databases | 72 |
| | Default devices for sample databases | 72 |
| | interpubs database | 74 |
| | jpubs database..... | 75 |
| | Maintaining the sample databases..... | 75 |
| | jConnect 4.5, 5.5, and Java utilities..... | 76 |
| | Using open database connectivity..... | 77 |
| | Installing stored procedures | 77 |
| | Installing character sets | 77 |
| | Installing online help for Transact-SQL syntax..... | 80 |
| | Online syntax help: sp_syntax..... | 80 |
| | Default device for the sybsyntax database..... | 80 |
| | Installing sybsyntax | 81 |
| | Configuring XP Server after initial installation | 83 |
| CHAPTER 6 | Installing Sybase PC-Client Products..... | 85 |
| | Overview | 85 |
| | Installing PC-Client products..... | 86 |
| | Installing a Java Runtime Environment..... | 89 |
| | Installing the Java Development Kit | 89 |
| | Installing InfoMaker | 90 |
| | Installing PowerDynamo | 91 |
| | Configuring network connections for client products..... | 91 |
| | Adding a server entry to the libtcl.cfg file..... | 92 |
| | Adding a server to the sql.ini file | 93 |

| | | |
|------------------|--|------------|
| | Enabling TCP connections | 95 |
| | Testing the Sybase Central installation | 96 |
| CHAPTER 7 | Starting and Stopping Servers | 99 |
| | Overview | 99 |
| | Starting servers | 100 |
| | Using the Runserver file | 100 |
| | Using the startserver command | 100 |
| | Using the monserver command | 101 |
| | Starting servers when the operating system restarts | 101 |
| | Stopping servers | 102 |
| | Stopping Adaptive Server..... | 102 |
| | Stopping Backup Server..... | 103 |
| | Stopping Monitor Server..... | 104 |
| | Using the kill command | 104 |
| | Shutdown and shared memory files | 105 |
| CHAPTER 8 | Upgrading Sybase Servers | 107 |
| | Overview | 107 |
| | System catalog changes during upgrade | 109 |
| | Pre-upgrade tasks | 110 |
| | Check system and upgrade requirements..... | 110 |
| | Check RUN_server file location | 111 |
| | Test current applications and stored procedures | 111 |
| | Reserved words | 112 |
| | Verify that users are logged off | 114 |
| | Check database integrity | 114 |
| | Back up databases..... | 115 |
| | Dump transaction logs..... | 115 |
| | Preparing database and devices..... | 116 |
| | Turn off database options..... | 123 |
| | Preparing to upgrade servers with replicated databases | 124 |
| | Suspending transaction processing and replication activities | 125 |
| | Draining the transaction logs for primary databases | 126 |
| | Draining the RSSD Transaction Log | 128 |
| | Disabling the secondary truncation point | 129 |
| | Restoring replication after upgrade | 129 |
| | Upgrading to Adaptive Server 12.5 | 130 |
| | Using sqlupgrade | 131 |
| | Using sqlupgraderes | 134 |
| | Post-upgrade tasks | 135 |
| | Verifying that servers are running | 136 |
| | Restore functionality in Adaptive Server | 136 |

| | | |
|-------------------|---|------------|
| | Reenabling replication server | 139 |
| | Reenabling auditing..... | 140 |
| | Upgrading Backup Server, Monitor Server, and XP Server | 141 |
| | Upgrading compiled objects with dbcc upgrade_object | 142 |
| | Finding compiled object errors before production | 143 |
| | Using dbcc upgrade_object..... | 146 |
| | Upgrading compiled objects in database dumps..... | 149 |
| | Determining whether a compiled object has been upgraded | 149 |
| | Recovering from a failed upgrade | 149 |
| CHAPTER 9 | Troubleshooting | 151 |
| | Overview | 151 |
| | Error log locations | 152 |
| | For installation utilities | 152 |
| | For Sybase servers | 153 |
| | Solutions to common installation problems..... | 153 |
| | Unable to use X-Windows | 153 |
| | Cannot eject the CD from the drive | 154 |
| | DISPLAY environment variable not set correctly | 154 |
| | Client not authorized to connect to server..... | 155 |
| | Address already in use..... | 155 |
| | Adaptive Server failed to boot | 155 |
| | Cannot start XP Server | 155 |
| | Stopping Adaptive Server after a failure..... | 156 |
| | Recovering from a failed installation | 156 |
| | Troubleshooting resource file installations | 157 |
| | Recovering from a failed upgrade | 158 |
| | Troubleshooting SySAM | 160 |
| | Authorization code input error | 161 |
| APPENDIX A | Alternative Installation Methods | 163 |
| | Resource file installation | 163 |
| | Editing a resource file..... | 163 |
| | Using srvbuildres..... | 166 |
| | Troubleshooting resource file installations | 167 |
| APPENDIX B | Creating Raw Partitions | 169 |
| | Overview | 169 |
| | Creating the partitions | 171 |
| | Binding the partitions for raw disk I/O | 172 |
| | Accessing raw devices from Adaptive Sever | 175 |

| | |
|-------------------|------------|
| Index..... | 177 |
|-------------------|------------|

About This Book

This guide, *Installation Guide Sybase Adaptive Server Enterprise on Linux/Intel* provides:

- An overview of the Sybase® Adaptive Server™ Enterprise installation infrastructure
- Instructions for installing and upgrading Adaptive Server, and installing Backup Server™, Monitor Server, XP Server™, jConnect™ for JDBC, Java utilities, and client products, including the Adaptive Server plug-in for Sybase Central™
- Instructions for installing optional Adaptive Server functionality, such as auditing, and sample databases, and localization information

Audience

This guide is written for Sybase System Administrators and other qualified personnel who are familiar with their system's environment, networks, disk resources, and media devices.

How To Use This Book

This book contains the following chapters:

- Chapter 1, "Overview" provides product definitions and an overview of the installation process and the Adaptive Server installation infrastructure.
- Chapter 2, "Installation Requirements" provides descriptions and system-related information such as required RAM and disk space for all Adaptive Server server and client products for this release.
- Chapter 3, "Sybase Software Asset Management (SySAM)" provides installation instructions for Sybase Software Manager (SySAM).
- Chapter 4, "Installing Sybase Servers" describes pre-installation procedures, and how to install server components.
- Chapter 5, "Post-Installation Tasks" describes how to set environment variables, install sample databases, and initialize features.

- Chapter 6, “Installing Sybase PC-Client Products” describes how to install Adaptive Server client products like the Adaptive Server plug-in to Sybase Central.
- Chapter 7, “Starting and Stopping Servers” describes the different methods of starting and stopping Adaptive Server, Backup Server™, and other servers.
- Chapter 8, “Upgrading Sybase Servers” describes how to upgrade an existing Adaptive Server to the current release.
- Chapter 9, “Troubleshooting” provides installation error messages and possible solutions to installation problems.
- Appendix A, “Alternative Installation Methods” describes how to install Sybase servers in noninteractive mode.
- Appendix B, “Creating Raw Partitions” describes how to configure raw partitions for use with Adaptive Server.

Related documents

The following documents comprise the Sybase Adaptive Server Enterprise documentation:

- The release bulletin for your platform – contains last-minute information that was too late to be included in the books.

A more recent version of the release bulletin may be available on the World Wide Web. To check for critical product or document information that was added after the release of the product CD, use the Sybase Technical Library.
- The *Installation Guide* for your platform – describes installation, upgrade, and configuration procedures for all Adaptive Server and related Sybase products.
- *Configuring Adaptive Server Enterprise* for your platform – provides instructions for performing specific configuration tasks for Adaptive Server.
- *What's New in Adaptive Server Enterprise?* – describes the new features in Adaptive Server version 12.5, the system changes added to support those features, and the changes that may affect your existing applications.

- *Transact-SQL User's Guide* – documents Transact-SQL, Sybase's enhanced version of the relational database language. This manual serves as a textbook for beginning users of the database management system. This manual also contains descriptions of the pubs2 and pubs3 sample databases.
- *System Administration Guide* – provides in-depth information about administering servers and databases. This manual includes instructions and guidelines for managing physical resources, security, user and system databases, and specifying character conversion, international language, and sort order settings.
- *Reference Manual* – contains detailed information about all Transact-SQL commands, functions, procedures, and datatypes. This manual also contains a list of the Transact-SQL reserved words and definitions of system tables.
- *Performance and Tuning Guide* – explains how to tune Adaptive Server for maximum performance. This manual includes information about database design issues that affect performance, query optimization, how to tune Adaptive Server for very large databases, disk and cache issues, and the effects of locking and cursors on performance.
- *The Utility Guide* – documents the Adaptive Server utility programs, such as isql and bcp, which are executed at the operating system level.
- *The Quick Reference Guide* – provides a comprehensive listing of the names and syntax for commands, functions, system procedures, extended system procedures, datatypes, and utilities in a pocket-sized book. Available only in print version.
- *The System Tables Diagram* – illustrates system tables and their entity relationships in a poster format. Available only in print version.
- *Error Messages and Troubleshooting Guide* – explains how to resolve frequently occurring error messages and describes solutions to system problems frequently encountered by users.
- *Component Integration Services User's Guide* – explains how to use the Adaptive Server Component Integration Services feature to connect remote Sybase and non-Sybase databases.
- *Java in Adaptive Server Enterprise* – describes how to install and use Java classes as datatypes, functions, and stored procedures in the Adaptive Server database.

- *Using Sybase Failover in a High Availability System* – provides instructions for using Sybase's Failover to configure an Adaptive Server as a companion server in a high availability system.
- *Using Adaptive Server Distributed Transaction Management Features* – explains how to configure, use, and troubleshoot Adaptive Server DTM features in distributed transaction processing environments.
- *EJB Server User's Guide* – explains how to use EJB Server to deploy and execute Enterprise JavaBeans in Adaptive Server.
- *XA Interface Integration Guide for CICS, Encina, and TUXEDO* – provides instructions for using Sybase's DTM XA interface with X/Open XA transaction managers.
- *Glossary* – defines technical terms used in the Adaptive Server documentation.
- *Sybase jConnect for JDBC Programmer's Reference* – describes the jConnect for JDBC product and explains how to use it to access data stored in relational database management systems.
- *Full-Text Search Specialty Data Store User's Guide* – describes how to use the Full-Text Search feature with Verity to search Adaptive Server Enterprise data.
- *Historical Server User's Guide* – describes how to use Historical Server to obtain performance information for SQL Server and Adaptive Server.
- *Monitor Server User's Guide* – describes how to use Monitor Server to obtain performance statistics from SQL Server and Adaptive Server.
- *Monitor Client Library Programmer's Guide* – describes how to write Monitor Client Library applications that access Adaptive Server performance data.

Other sources of information

Use the Sybase Technical Library CD and the Technical Library Product Manuals Web site to learn more about your product:

- Technical Library CD contains product manuals and is included with your software. The DynaText browser (downloadable from Product Manuals at <http://www.sybase.com/detail/1,3693,1010661,00.html>) allows you to access technical information about your product in an easy-to-use format.

Refer to the *Technical Library Installation Guide* in your documentation package for instructions on installing and starting the Technical Library.

- Technical Library Product Manuals Web site is an HTML version of the Technical Library CD that you can access using a standard Web browser. In addition to product manuals, you will find links to the Technical Documents Web site (formerly known as Tech Info Library), the Solved Cases page, and Sybase/Powersoft newsgroups.

To access the Technical Library Product Manuals Web site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ For the latest information on product certifications

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Select Products from the navigation bar on the left.
- 3 Select a product name from the product list.
- 4 Select the Certification Report filter, specify a time frame, and click Go.
- 5 Click a Certification Report title to display the report.

❖ For the latest information on EBFs and Updates

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Select EBFs/Updates. Enter user name and password information, if prompted (for existing Web accounts) or create a new account (a free service).
- 3 Specify a time frame and click Go.
- 4 Select a product.
- 5 Click an EBF/Update title to display the report.

❖ To create a personalized view of the Sybase Web site (including support pages)

Set up a MySybase profile. MySybase is a free service that allows you to create a personalized view of Sybase Web pages.

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>
- 2 Click MySybase and create a MySybase profile.

Conventions

The following style conventions are used in this manual:

- In a sample screen display, commands you should enter exactly as shown are given in:

`this font`

- In a sample screen display, words that you should replace with the appropriate value for your installation are shown in:

this font

- In the regular text of this document, the names of files and directories appear in this font:

/usr/u/sybase

- The names of programs, utilities, procedures, and commands appear in this font:

`sqlupgrade`

- Commands for both the C shell and the Bourne shell are provided in this document, when they differ. The initialization file for the C shell is called *.cshrc*. The initialization file for the Bourne shell is called *.profile*. If you are using a different shell, such as the Korn shell, refer to your shell-specific documentation for the correct command syntax.

Table 1 shows the conventions for syntax statements in this manual.

Table 1: SQL syntax conventions

| Key | Definition |
|-----------------|--|
| command | Command names, command option names, utility names, utility flags, and other keywords are in bold. |
| <i>variable</i> | Variables, or words that stand for values that you fill in, are in <i>italic</i> . |
| { } | Curly braces indicate that you choose at least one of the enclosed options. Do not include braces in your option. |
| [] | Brackets mean choosing one or more of the enclosed options is optional. Do not include brackets in your option. |
| () | Parentheses are to be typed as part of the command. |
| | The vertical bar means you can select only one of the options shown. |
| , | The comma means you can choose as many of the options shown as you like, separating your choices with commas to be typed as part of the command. |

If you need help

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you cannot resolve a problem using the manuals or online help, please have the designated person contact Sybase Technical Support or the Sybase subsidiary in your area.

Overview

This chapter introduces important concepts about the installation of Adaptive Server and Adaptive Server products.

| Topic | Page |
|----------------------------------|------|
| User roles | 1 |
| Product descriptions | 2 |
| Installation and setup utilities | 4 |
| PC-Client product descriptions | 5 |
| Sybase installation directory | 8 |

User roles

The Adaptive Server installation and setup process defines various user roles. Different user roles have different responsibilities and privileges. These user roles clarify the way Adaptive Server is to be integrated into your system:

- *Operating System Administrator* – The individual who maintains the operating system. This individual has superuser or “root” privileges.
- *Sybase System Administrator* – The individual in charge of Adaptive Server system administration, creating user accounts, assigning permissions on databases, and creating new databases. At installation time, the System Administrator’s login name is “sa”. The “sa” login is not a UNIX login. The “sa” login is specific to Adaptive Server and is used to log in to Adaptive Server using the `isql` command.
- *Sybase Login* – The “sybase” login is a UNIX login that owns all the Sybase installation directories and files, sets permissions on those directories and files, and performs the installation and upgrading of Adaptive Server.

Product descriptions

Table 1-1 provides a description of the server products in your Adaptive Server package.

Table 1-1: Product descriptions

| Product | Description |
|---|---|
| <i>Adaptive Server</i> | <p>The relational database server. The default unloading process includes:</p> <ul style="list-style-type: none">• Adaptive Server• Backup Server• Monitor Server• XP Server• Adaptive Server utilities• Scripts and configuration files |
| <i>Adaptive Server Plug-In to Sybase Central Java Edition</i> | <p>Sybase Central is a common framework for managing servers. It allows you to manage Adaptive Server installations using the Sybase Central graphical management tool.</p> <p>For information on how to use Sybase Central to manage Adaptive Server, see the online help.</p> |
| <i>Sybase Monitor Server for Adaptive Server</i> | <p>An Open Server™ application that obtains performance statistics on Adaptive Server and makes those statistics available to Monitor Server client applications.</p> <p>Sybase Monitor Server includes:</p> <ul style="list-style-type: none">• Monitor Server for Adaptive Server Enterprise 12.5 – an Open Server application that obtains performance statistics on Adaptive Server and makes those statistics available to monitors in Monitor Historical Servers, and applications built with Monitor Client-Library.• Monitor Client-Library – a programming interface that provides access to Adaptive Server performance data.• Monitor Historical Server – an Open Server application that obtains performance statistics for many Adaptive Servers via Monitor Servers and records the data to specified file locations. |

| Product | Description |
|--|--|
| <i>Backup Server</i> | <p>Backup Server is an Open Server-based application that manages all database backup (dump) and restore (load) operations for Adaptive Server. Backup Server:</p> <ul style="list-style-type: none"> • Allows you to use up to 32 dump devices (this is called <i>dump striping</i>) in parallel to dump or load a single database or transaction log. • Allows one dump to span multiple tapes or allows multiple dumps to be made to a single tape. • Allows dumping and loading over the network to or from a device on another machine. • Provides automated determination of tape device characteristics, by using operating system commands, for a dump operation. • Supports dump and load command syntax specification for volume naming, dismount and load control, tape density, block size, tape capacity, days to retain, initialization, file naming for multi-dump volumes, and listing header or file information. <p>Install Backup Server if you plan to back up and restore databases in Adaptive Server. Backup Server is installed, by default, when you install Adaptive Server software.</p> |
| <i>Language Modules (Server)</i> | <p>Provides system messages and date/time formats to help you localize your applications. Default installation includes the us_english language module and the following character sets:</p> <ul style="list-style-type: none"> • cp437 – IBM CP437, U.S. code set • cp850 – IBM CP850, European code set • iso_1 – ISO 8859-1, Latin-1 • mac – Standard Macintosh coding • roman8 – HP Roman-8 <p>Other languages for Adaptive Server and Sybase client products are French, German, and Japanese. Language modules for Spanish, Korean, Brazilian Portuguese, and Simplified Chinese are also available for Adaptive Server only.</p> <p>For more information on server language modules and character sets, see the <i>Sybase Adaptive Server Enterprise Configuration Guide</i> for your platform.</p> |
| <i>Language Modules (Connectivity)</i> | <p>Provides messages and support files for running Open Client™ applications in various languages.</p> |
| <i>jConnect 4.5 and jConnect 5.5</i> | <p>Provides a Java database connectivity (JDBC) driver that works with both Sun and Microsoft virtual machines (VM).</p> <p>Provides support for Adaptive Server 12.5 extended limits by requesting support for wide tables. This request is ignored by servers other than Adaptive Server 12.5 and higher.</p> <p>For more information on jConnect™ for JDBC™, see the jConnect product page at http://www.sybase.com/support/manuals/.</p> |
| <i>jConnect Documentation</i> | <p>Contains the <i>Sybase jConnect for JDBC Programmer's Reference</i>.</p> |

| Product | Description |
|-----------------------|---|
| <i>Java utilities</i> | Includes: <ul style="list-style-type: none"> • The Cascade Gateway – A gateway that acts as a proxy to provide a path to the database server if it is running on a different host from the Web Server. • jisql – A graphical Transact-SQL editor written in Java that replaces SQL Advantage. • Ribo – A utility that captures, translates, and displays the Tabular Data Stream (TDS) protocol flowing between a TDS client and a TDS server. |
| <i>ODBC Driver</i> | Allows Windows NT client applications to access Adaptive Server data. |
| <i>Open Client</i> | Provides libraries and utilities for developing any Open Client application. |
| <i>XP Server</i> | An Open Server application that manages and executes extended stored procedures (ESPs) from within Adaptive Server. ESPs provide a method for calling procedural language functions from within Adaptive Server. XP Server is unloaded, by default, when you unload Adaptive Server software from the distribution media. Use the srvbuild utility to set up XP Server and to connect XP Server and Adaptive Server through the interfaces file. |

Installation and setup utilities

Table 1-2 lists the Adaptive Server installation and setup utilities you can use to unload, install, upgrade, or customize Adaptive Server and its product.

Table 1-2: Installation and setup utilities

| Utility | Use to | Environment |
|---------------|--|--------------------------------------|
| asecfg | Utility to access dsedit, srvbuild, sqlupgrade, sqlloc. | X-Windows/Motif GUI |
| dscp | Modify the interfaces file. | Interactive, text-based interface |
| dsedit | Modify the interfaces file. | X-Windows/Motif GUI |
| sqlloc | Modify localization settings, such as language and character set. For information on using sqlloc, see the <i>Adaptive Server Configuration Guide</i> . | X-Windows/Motif GUI |
| sqllocres | Modify localization settings, such as language and character set, using a resource file. Use of this utility is not described in this installation guide. For information on using sqllocres, see <i>Adaptive Server Enterprise Configuration Guide</i> . | Noninteractive, file-based interface |
| sqlupgrade | Upgrade Adaptive Server. | X-Windows/Motif GUI |
| sqlupgraderes | Upgrade Adaptive Server using a resource file. | Noninteractive, file-based interface |

| Utility | Use to | Environment |
|------------------|--|--------------------------------------|
| srvbuild | Build a functional server. | X-Windows/Motif GUI |
| srvbuildres | Build a functional server using a resource file. | Noninteractive, file-based interface |
| Studio Installer | Provide an interface from which to enter Sybase Server features. | Java 1.2.2 |

PC-Client product descriptions

Table 1-3 describes Adaptive Server client products, which can be installed on Windows 98, and Windows NT client computers.

Table 1-3: Products and platforms

| Product and Platforms | Description |
|---|--|
| <i>Open Client</i> version 12.5 | <p>Used to develop and deploy C language-based applications that access Adaptive Server data.</p> <p>Default installation includes:</p> <ul style="list-style-type: none">• Client-Library™• DB-Library• CS-Library• Bulk Library• bcp utility used to transfer data between files and Adaptive Server databases• isql utility used to connect to Adaptive Server for SQL queries.• defncopy utility used to copy stored procedure definitions from one database to another• Online help files for Open Client routines• Net-Library drivers for connecting 32-bit client applications to a server through most available network protocols. Default installation includes:<ul style="list-style-type: none">• dsedit for editing the <i>sql.ini</i> and <i>interfaces</i> files and testing client-server connections• Drivers for Named Pipes and Windows Sockets• Drivers for Named Pipes, Windows Sockets, Novell LAN Workplace, and Microsoft TCP/IP <p>Additional options include:</p> <ul style="list-style-type: none">• Programmer support (include files and libraries) for compiling and linking Client-Library and DB-Library applications• Sample programs for Client-Library and DB-Library• Net-Library driver protocols:<ul style="list-style-type: none">TCP/IP, SPX/IPX, Named Pipes, DECNet.TCP/IP and SPX/IPX require the WinSock API. DECNet requires DEC PathWorks.• Monitor Client Library – an application programming interface that provides access to Adaptive Server performance data. |
| <i>Language modules</i> version 12.5 | <p>Display system messages and datetime formats to help you localize your 32-bit applications. The Standard Installation option installs all of the character sets.</p> |

| Product and Platforms | Description |
|--|---|
| <i>Embedded SQL/C and Embedded SQL/COBOL</i> version 12.5 | A superset of Transact-SQL [®] that lets you place Transact-SQL statements in application programs written in languages such as C and COBOL. |
| <i>jConnect 4.5 and jConnect 5.5</i> | <p>Provides a Java database connectivity (JDBC) driver that works with both Sun and Microsoft virtual machines (VM).</p> <p>Provides support for Adaptive Server 12.5 extended limits by requesting support for wide tables. This request is ignored by servers other than Adaptive Server 12.5 and higher.</p> <p>For more information on jConnect[™] for JDBC[™], see the jConnect product page at http://www.sybase.com/support/manuals/.</p> |
| <i>jutils</i> (Cascade Gateway, jisql, Ribo) | <p>Includes:</p> <ul style="list-style-type: none"> • The Cascade Gateway – a gateway that acts as a proxy to provide a path to the database server if it is running on a different host from the Web Server. • jisql – a graphical Transact-SQL editor written in Java that replaces SQL Advantage[®]. • Ribo – a utility that captures, translates, and displays the Tabular Data Stream (TDS)[™] protocol flowing between a TDS client and a TDS server. |
| <i>jConnect Documentation</i> | Contains the <i>Sybase jConnect for JDBC Programmer's Reference</i> . |
| <i>Sybase Central ASE plug-in Java edition</i> | The Java-based, graphical administration utility for Adaptive Server. |
| <i>Sybase Central[™] ASE plug-in C++</i> version 12.0 | <p>SQL Remote and the Monitor Server GUI require the C++ version of this plug-in.</p> <hr/> <p>Warning! Do not use the Adaptive Server C++ plug-in for general administrative purposes.</p> |
| <i>Sybase ODBC drivers</i> version 3.5 | ODBC drivers have support for extended limits. |
| <i>PowerDynamo</i> version 3.5.2 | A suite of tools for building and managing database-hosted Web sites with dynamic content. |
| <i>InfoMaker</i> version 7.0.3 | A tool for personal data access, management, and reporting, for developers and end users. InfoMaker allows you to create presentation-quality reports and powerful queries without the complexities of programming. It complements client/server business applications and development tools as well as desktop productivity suites. |

| Product and Platforms | Description |
|--|--|
| <i>SQL Remote</i> version 7.0.2 | SQL Remote enables two-way replication between a database server and multiple portable computer databases through E-mail or dial-up connections. |
| <i>Physical Architect</i> version 8.0 | A tool for data modeling, including database design, generation, maintenance, reverse engineering, and documentation for database architects. |

Sybase installation directory

The Sybase installation directory structure is created by the installation process. Adaptive Server is placed in the directory you indicate when you run the the Red Hat Package Manager. Table 1-4 illustrates part of the Adaptive Server installation directory structure.

The Adaptive Server installation directory contains executable files and administrative tools, which are added as each product is installed.

Table 1-4: Installation directory for \$SYBASE

| Component directory | Subdirectory | Subdirectory |
|-----------------------|----------------------|-------------------|
| <i>/ASE-12_5</i> | | |
| | <i>/bin</i> | |
| | <i>/certificates</i> | |
| | <i>/debugger</i> | |
| | <i>/diag</i> | |
| | <i>/include</i> | |
| | <i>/init</i> | <i>/auditinit</i> |
| | | <i>/bsrv</i> |
| | | <i>/logs</i> |
| | | <i>/sqlsrv</i> |
| | <i>/install</i> | <i>/spr</i> |
| | <i>/lib</i> | |
| | <i>/sample</i> | <i>/esp</i> |
| | | <i>/JavaSql</i> |
| | | <i>/server</i> |
| | <i>/scripts</i> | |
| | <i>/sybhelp</i> | |
| | <i>/upgrade</i> | |
| | <i>/xappdefaults</i> | |
| <i>/charsets</i> | | |
| <i>/collate</i> | <i>/unicode</i> | |
| <i>/config</i> | | |
| <i>/EFTS-12_5</i> | | |
| <i>/installed</i> | | |
| <i>/Installer</i> | <i>/bin</i> | |
| | <i>/lib</i> | |
| <i>/jConnect_docs</i> | <i>/docs</i> | |

| Component directory | Subdirectory | Subdirectory |
|----------------------|---|--|
| <i>/jConnect-4_5</i> | <i>/classes</i> | <i>/com</i> <i>/gateway</i> <i>/sample</i> |
| | <i>/devclasses</i> | <i>/com</i> |
| | <i>/docs</i> | <i>/en</i> |
| | <i>/gateway</i> | |
| | <i>/sample</i> | |
| | <i>/sp</i> | |
| | <i>/tools</i> | |
| | <i>/devclasses</i> | |
| <i>/jConnect-5_5</i> | <i>/classes</i> | <i>/gateway2</i> <i>/sample2</i> |
| | <i>/devclasses</i> | |
| | <i>/docs</i> | <i>/en</i> |
| | <i>/gateway2</i> | |
| | <i>/sample2</i> | |
| | <i>/sp</i> | |
| | <i>/tools</i> | |
| | | |
| <i>/jutils-2_0</i> | <i>/cascade</i> | <i>/uk</i> |
| | <i>/jisql</i> | <i>/doc</i> <i>/HelpFiles</i> |
| | <i>/ribo</i> | <i>/doc</i> |
| <i>/locales</i> | <i>/us_english</i> | |
| | <i>/message</i> | |
| | <i>/unicode</i> | |
| <i>/OCS-12_5</i> | <i>/bin</i> | |
| | <i>/config</i> | |
| | <i>/devlib</i> | |
| | <i>/include</i> | |
| | <i>/lib</i> | |
| | <i>/lib3p</i> (<i>/lib3p64</i> for 64-bit platforms) | |
| | <i>/locales</i> | |
| | <i>/sample</i> | |
| | <i>/sybhelp</i> | |
| <i>/shared-1_0</i> | | |

| Component directory | Subdirectory | Subdirectory |
|---------------------|------------------|--------------|
| <i>/SQLRemote</i> | <i>/bin</i> | |
| | <i>/scripts</i> | |
| <i>/SYSAM-1_0</i> | <i>/bin</i> | |
| | <i>/licenses</i> | |
| | <i>/log</i> | |

Installation Requirements

This chapter covers important information regarding system requirements for installing Adaptive Server databases and devices:

| Topic | Page |
|--|------|
| System requirements | 13 |
| Product disk space requirements | 18 |
| Adaptive Server specifications | 19 |
| Adaptive Server devices and system databases | 23 |

For easier installation, make a copy of this chapter, so you have the requirements information available as you perform the pre-installation, installation, or upgrade procedures.

System requirements

Table 2-1 shows the required release-level, RAM, and network-protocol requirements for Adaptive Server. See the release bulletin for the latest information on required system patches.

System requirements

This section contains the hardware and software requirements for the server and client software included with Adaptive Server 12.5.

Adaptive Server 12.5.0.1 has been released on two versions of the Linux kernel. The release levels of the operating system components for both releases are indicated in the tables below.

Table 2-1: System requirements on kernel version 2.2.14

| Hardware | Operating system | Network protocols |
|---|--|-------------------------------|
| <i>Server</i> | | |
| Linux with Intel architecture, Pentium II, 230 MHz | <ul style="list-style-type: none"> Kernel – version 2.2.14-5.0 Run-time environment – glibc-2.1.3-15 RPM – rpm-3.0.4-0.48 X11R6.4 – XFree86-3.3.6-20 <p>If you plan to use raw disk I/O, you need kernel 2.2.18 or higher.</p> | TCP/IP sockets |
| <i>PC-client</i> | | |
| Windows NT | Windows NT 4.0, (Service Pack #6a, or later) | TCP IPX/SPX Named Pipes |
| Windows 98 | Windows 98 | TCP IPX/SPX Named Pipes |

Table 2-2: System requirements on kernel version 2.4.x

| Hardware | Operating system | Network protocols |
|---|---|-------------------|
| <i>Server</i> | | |
| Linux with Intel architecture, Pentium II, 230 MHz | <ul style="list-style-type: none"> Kernel – version 2.4.7-10 Run-time environment – glibc-2.2.4-13 <ul style="list-style-type: none"> Pthreads – libpthread-0.9 RT– librt-2.4.4 RPM – rpm-4.0.3-1.03 X11R6.4 – XFree86-4.1.0-3 | TCP/IP sockets |
| <i>PC-client</i> | | |

| Hardware | Operating system | Network protocols |
|-----------------|--|-------------------------------|
| Windows NT | Windows NT 4.0, (Service Pack #6a, or later) | TCP IPX/SPX Named Pipes |
| Windows 98 | Windows 98 | TCP IPX/SPX Named Pipes |

Table 2-3 shows the system requirements for installing jConnect for JDBC.

Table 2-3: System requirements for jConnect on UNIX

| | |
|----------------------------------|---|
| Hardware | Any UNIX platform that supports Java VM 1.2.x |
| Operating system | Linux/Intel, Sun Solaris, HP UX, IBM AIX, DEC Digital Unix, SGI IRIX, and so on. |
| Java Developer's Kit | Sun's JDK 1.2.x or later |
| Database | <ul style="list-style-type: none">• Adaptive Server Enterprise, from SQL Server 11.x through Adaptive Server Enterprise 12.x; or• Sybase SQL Anywhere 5.5.23 or later, running Open Server Gateway™ dbos50; or• Sybase SQL Anywhere Studio 6.0* <hr/> <p>Note Some features in jConnect version 4.5 and 5.5 require Adaptive Server Enterprise 12.x or later:</p> <ul style="list-style-type: none">• Support of wide tables• Support for outer join syntax• Support for Sybase Failover in high availability systems (“HA Failover”)• Implementation of a DynamicClassLoader (or Adaptive Server Anywhere 6.x and later)• Support for distributed transaction management —Java Transaction API (JTA), and Java Transaction Service (JTS) <p>For more information on these features, see the <i>Sybase jConnect for JDBC Programmer's Reference</i>.</p> <hr/> |
| Web browser | Any browser that supports JDK 1.1.x or later, such as Netscape 4.0.x, Internet Explorer 4.0.x, or HotJava 1.x |
| Other Sybase products (optional) | Adaptive Server IQ, OmniConnect, DirectConnect |
| Web server (optional) | For jConnect's TDS-tunnelling servlet, any Web server that supports the HTTPS protocol and javax.servlet interfaces |

Note A free evaluation version of SQL Anywhere Studio 6.0 (includes Adaptive Server Anywhere 6.0) is available from Sybase at <http://www.sybase.com/products/anywhere/sqlanyform.html>. The Open Server Gateway and the required DLLs are no longer needed with Sybase SQL Anywhere Studio version 6.0 and later.

Client products are supported on Windows NT, and Windows 98. Table 2-4 lists the RAM required for client products.

Table 2-4: RAM requirements for PC-client products

| Product | Memory requirements |
|--|---------------------------|
| Adaptive Server plug-in for Sybase Central, Java Edition | 16MB |
| Adaptive Server plug-in for Sybase Central, C++ | 16MB |
| Open Client/C Developer's kit | 16MB, 32MB recommended |
| Monitor Client Library | 16MB, 32MB recommended |
| Embedded SQL/COBOL and Embedded SQL/C | |
| jConnect 4.5 and 5.5 | 16MB each |
| SQL Remote 7.0.2 | 37MB |
| PowerDynamo 3.5.2 | |
| InfoMaker 7.0.3 | 32MB |
| Physical Architect 8.0 | 16MB |

Table 2-5 shows the installation requirements for the Adaptive Server plug-in for Sybase Central.

Table 2-5: Installation requirements

| Operating system | Product | Disk space |
|--|--|------------|
| Windows NT 4.0 Service Pack 6 or later, or Windows 98 | <i>Sybase Central ASE plug-in C++</i> | 12MB |
| <p>Warning! The C++ version of the plug-in is included only to provide the DDLGEN functionality required by the database migration utility. Do not use the C++ plug-in for general administrative purposes.</p> | | |
| Windows NT 4.0 Service Pack 6 or later, or Windows 98 | <i>Sybase Central ASE plug-in Java edition</i> | 16MB |

Product disk space requirements

Table 2-6 lists the disk space requirements for each product.

Table 2-6: Product disk space requirements

| Product and version level | Package size | Disk space required |
|--|--------------|---------------------|
| Adaptive Server 12.5, includes | 37MB | 115MB |
| <ul style="list-style-type: none"> • Monitor Server 12.5 • Historical Server 12.5 • Monitor Client Library 12.5 | | |
| Sybase Common package | 95MB | 100MB |
| Open Client/C 12.5 | 38MB | 120MB |
| Language Modules (Server 12.5) | | 2–7MB per language |
| Sybase-french | | |
| Sybase-german | | |
| Language Modules (Open Client 12.5) | | 3–6MB per language |
| Enhanced full-text search | 33MB | 65MB |
| Sybase-esql | 1.2MB | 4MB |
| Sybase Openserver | 13MB | 45MB |
| SQL Remote | 1.7MB | |

Language module sizes

Table 2-7 lists the sizes for the Adaptive Server and Open Client Language modules.

Table 2-7: Language module sizes

| Language | Size language module |
|-----------------------|----------------------|
| Brazilian Portuguese | 6.5MB |
| Chinese (simplified) | 2MB |
| Chinese (traditional) | 2MB |
| French | 6.5MB |
| German | 6.5MB |
| Japanese | 4.5MB |
| Korean | 2MB |
| Spanish | 6.5MB |

Adaptive Server specifications

Table 2-8: Adaptive Server specifications on Linux (approximate)

| | | |
|--|-------------------------|--|
| <i>Hardware</i> | | |
| Processor | 32-bit Intel Pentium II | |
| Minimum RAM required for Adaptive Server | 50MB | |
| Minimum RAM per additional user | 116KB | |
| Default user stack size | 64KB | |

Database specifications

| | | |
|---|--------------------------|---|
| Memory per user connection | 155K | With default stack size, packet size, and user log cache size |
| Databases per Adaptive Server | 32,767 | Practical limit is approximately 100 |
| Maximum database size | 2 ⁴³ (8TB) | Minus overhead for system databases |
| Minimum allowable sybsystemprocs database | 100MB | Required for an upgrade |

| | | |
|---|---------------------------|---|
| Maximum size of a database device (disk partition) | 2 ²⁴ (32GB) | <ul style="list-style-type: none"> • 2GB maximum device size if using operating system files • 32GB device size for raw partitions. |
| Maximum number of database devices per server | 256 | |
| Maximum number of devices or device pieces per database | Unlimited | |
| Maximum number of segments per database | 31 | |
| Maximum number of login IDs per server | 2147516416 | |
| Maximum number of users per database | 2146484223 | |
| Maximum number of groups per database | 1032193 | |
| pubs2 database | 2MB | Minimum required for a new installation |
| | 3MB | Minimum required for an upgrade |
| pubs3 database | 2MB | Minimum required for a new installation |
| | 3MB | Minimum required for an upgrade |
| interpubs database | 2MB | Minimum required for a new installation |
| | 3MB | Minimum required for an upgrade |
| jpubs database | 2MB | Minimum required for a new installation |
| | 3MB | Minimum required for an upgrade |

Table specifications

| | | |
|-----------------------------|------------------------------|-------------------------|
| User objects per database | 2 ³¹ – 100 | |
| Indexes per table | 250 (one clustered index) | |
| Rows per table | Limited by available storage | Maximum 2 ³² |
| Columns per composite index | 31 | |

| | | |
|-------------------------------------|---|---|
| Creation of clustered index | 1.2*(x + y) x = total data space in table, y = sum of space of all nonclustered indexes on table, and 20% overhead for logging | For sorted data, approximately 20% of the table size needed |
| Characters per database object name | 30 | |

Query specifications

| | | |
|---|-----|--|
| Maximum number of tables in a “union” query | 256 | |
| Maximum number of databases participating in one transaction | 16 | Includes database where transaction began, all databases changed during transaction, and <i>tempdb</i> , if it is used for results or worktables |
| Practical number of databases participating in one query | 16 | Includes each occurrence of each database queried and <i>tempdb</i> , if it is used for results or worktables |
| Maximum number of tables participating in a query | 64 | Maximum of 50 user tables, including result tables, tables referenced by views (the view itself is not counted) correlations and self-joins; maximum of 14 worktables. |
| Maximum number of tables with referential integrity constraints for a query | 192 | |

Procedure specifications

| | | |
|---|--------------|--|
| Number of buffers and procedure buffers | Configurable | Limited by amount of RAM and maximum size of shared memory segment |
| Minimum memory required per stored procedure | 2K | |
| Maximum number of parameters per stored procedure | 2048 | |

Adaptive Server 12.5 extended-limit capabilities vary by type of table and the database logical page size. Table 2-9 lists the column and row limits for allpages-locked (APL) tables.

Table 2-9: Allpages-locked (APL) tables

| Maximum APL table limits | Number of columns | Column size 2K page | Column size 4K page | Column size 8K page | Column size 16K page |
|--------------------------|-------------------|---------------------|---------------------|---------------------|----------------------|
| Fixed-length column | 1024 | 1960 bytes | 4008 bytes | 8104 bytes | 16296 bytes |
| Variable-length column | 254 | 1948 bytes | 3988 bytes | 8068 bytes | 16228 bytes |

Table 2-10 lists the column and row limits for data-only-locked (DOL) tables.

Table 2-10: Data-only-locked (DOL) tables

| Maximum DOL table limits | Number of columns | Column size 2K page | Column size 4K page | Column size 8K page | Column size 16K page |
|--------------------------|-------------------|---------------------|---------------------|---------------------|----------------------|
| Fixed-length column | 1024 | 1958 bytes | 4006 bytes | 8102 bytes | 16294 bytes |
| Variable-length column | 1024 | 1954 bytes | 4002 bytes | 8098 bytes | 16290 bytes |

Database space requirements depend upon the logical page size of the server. Table 2-11 lists the minimum size for each database.

Table 2-11: Database requirements for varying page sizes

| Databases | 2K page | 4K page | 8K page | 16K page |
|-----------------------|---------|---------|---------|----------|
| Default database size | 2MB | 4MB | 8MB | 16MB |
| master database | 6MB | 12MB | 24MB | 48MB |
| model database | 3MB | 4MB | 8MB | 16MB |
| tempdb database | 3MB | 4MB | 8MB | 16MB |

Larger logical page sizes can contain more data. Table 2-12 lists the maximum data for each logical page size.

Table 2-12: Data limits for tables according to page size

| Tables | 2K page | 4K page | 8K page | 16K page |
|---|---------|---------|---------|----------|
| Maximum number of data bytes per text or image file | 1800 | 3600 | 7650 | 16200 |
| Bytes per index key | 600 | 1250 | 2600 | 5300 |
| User-visible row length DOL table | 1958 | 4006 | 8102 | 16294 |
| User-visible row length APL table | 1960 | 4008 | 8104 | 16296 |

Adaptive Server devices and system databases

Devices are operating system files that are used to store databases and database objects. Adaptive Server uses the O_SYNC option with operating system files to ensure data is reliably written to the disk.

Adaptive Server requires the following devices:

- master device – to store system databases.
- sysprocsdev device – to store system procedures.

The master and sysprocsdev devices are created when you create a new Adaptive Server.

master device

The master device contains the following databases:

- master – controls the operation of Adaptive Server as a whole and stores information about all users, user databases, devices, objects, and system table entries. The master database is contained entirely on the master device and cannot be expanded onto any other device.
- model – provides a template for new user databases. The model database contains required system tables, which are copied into a new user database with the `create database` command.
- tempdb – the work area for Adaptive Server. Each time Adaptive Server is started the tempdb database is cleared and rebuilt from the model database.

- The sample databases are stored on the master device at installation, but should be moved to a user-defined device after installation. For more information, see “Sample databases” on page 25.

Note For recovery purposes, Sybase recommends that you do not create other system or user databases or user objects on the master device.

***sybssystemdb* device and database**

For new installations the master device also contains the sybssystemdb database. The sybssystemdb device stores the sybssystemdb database, which stores information about transactions in progress, and which is also used during recovery.

See Chapter 8, “Upgrading Sybase Servers” for instructions on how to create the sybssystemdb device and database for two-phase commit and Data Transfer Management.

Note For recovery purposes, it is recommended that you do not create other system or user databases or user objects on the master device.

The sample databases are stored on the master device at installation, but should be moved to a user-defined device after installation. For information on the sample databases, see “Optional devices and databases” on page 25.

***sysprocsdev* device**

The sysprocsdev device stores the sybssystemprocs database, which contains most of the Sybase-supplied system procedures. System procedures are a collection of SQL statements and flow-of-control statements that perform system tasks; for example, sp_configure.

The system procedures that are needed during recovery situations are stored in the master database.

Note sysprocsdev is the default system name for this device. However, it is frequently referred to as the sybssystemprocs device, since it stores the sybssystemprocs database.

Optional devices and databases

The following devices and databases are needed only if you configure Adaptive Server for optional functionality, like auditing or two-phase commit transactions.

***sybsecurity* device and database**

The sybsecurity device stores the sybsecurity database and the auditing system procedures with which you can configure auditing for your system.

The auditing system records system security information in an Adaptive Server audit trail. You can use this audit trail to monitor the use of Adaptive Server or system resources.

Install auditing using the auditinit utility. The sybsecurity device is created as part of the auditing installation process.

See the *Adaptive Server Configuration Guide* for instructions on how to configure Adaptive Server for auditing.

The auditing system is discussed in more detail in the *System Administration Guide*.

Sample databases

The pubs2 and pubs3 databases are sample databases provided as a learning tool for Adaptive Server. The pubs2 sample database is used for most of the examples in the Adaptive Server documentation; other examples use the pubs3 database. Both are available in U.S. English versions of Adaptive Server.

The interpubs database contain French and German data. jpubs contains Japanese data.

For information about installing the sample databases, see “Installing sample databases” on page 72.

For information on the contents of these sample databases, see the *Transact-SQL User’s Guide*.

***sybsyntax* database**

The syntax database, sybsyntax, contains syntax help for Transact-SQL commands, Sybase system procedures, Adaptive Server utilities, and Open Client routines. You can retrieve this information using the system procedure sp_syntax.

For example, to see the syntax of the Transact-SQL select command, enter:

```
sp_syntax "select"
```

Adaptive Server includes a script for creating the sybsyntax database.

For instructions on how to install sybsyntax, see “Installing online help for Transact-SQL syntax” on page 80.

For more information about sp_syntax, see the *Reference Manual*.

bccdb database

The database consistency checker (dbcc) provides commands for checking the logical and physical consistency of a database. The dbccdb database stores the results of dbcc when dbcc checkstorage or dbcc check verifying are used.

dbcc checkstorage records configuration information for the **target database**, operation activity, and the results of the operation in the dbccdb database. Stored in the database are dbcc stored procedures for creating and maintaining dbccdb and for generating reports on the results of dbcc checkstorage operations.

Information on installing dbccdb is located in the chapter “Checking Database Consistency,” in the *System Administration Guide*.

Determining the location, type, and size of a database device

Adaptive Server requires several database devices. Table 2-13 below shows baseline values for each of the devices. See the release bulletin for any last-minute changes to these values.

Table 2-13: Adaptive Server database devices

| Device | Purpose | Minimum size | Minimum recommended size |
|---|-------------------------------------|--|--|
| master | Stores system databases | 14MB for 2K pages 27MB for 4K pages 50MB for 8K pages 105MB for 16K pages | 30MB 60MB 120MB 240MB |
| sysprocsdev (also called the sybssystemprocs device) | Stores the sybssystemprocs database | 100MB | 100MB (plus any space for holding stored procedures that you have created) |

| Device | Purpose | Minimum size | Minimum recommended size |
|------------------------|------------------------|--------------|------------------------------------|
| sybsystemdb | Transaction processing | 4MB | 5 – 20MB |
| sybsecurity (optional) | Required for auditing | 0–5MB | 7MB; more for specialized auditing |

Note By default, Adaptive Server 12.5 opens database devices with the O_SYNC flag to ensure recoverability in the event of power failure, operating system crash or Adaptive Server crash. This significantly affects performance. You can turn off the O_SYNC flag by starting the server with -T1623 trace flag.

Sybase Adaptive Server supports the database devices on NFS- and CIFS-mounted devices with Network Appliance Filers for storing data. Network appliance filers provide the same performance and data integrity as raw devices. There are no changes needed to the operating system or to Sybase Adaptive Server to use network appliance filers.

NFS mounted devices have been tested on Solaris, HP/UX, IBM AIX, Windows NT, and Linux.

You may need to repartition the disk you choose. Contact your operating system administrator for assistance.

Note Sybase recommends that to configure raw devices, you use Adaptive Server version 12.5.0.1 released on kernel 2.4.7 and deploy it on a system with at least the same operating system release level.

For a description of how to configure Adaptive Server for raw devices, see Appendix B, “Creating Raw Partitions.”

Sybase Software Asset Management (SySAM)

This chapter describes licensing concepts that you need to know before you begin installing Adaptive Server.

Sybase recommends that you read through this chapter before beginning the installation of Adaptive Server in the enterprise environment.

| Topic | Page |
|-------------------------|------|
| SySAM overview | 29 |
| Types of SySAM systems | 33 |
| SySAM administration | 37 |
| Adding feature licenses | 38 |

This chapter describes important concepts regarding how to license optional Adaptive Server features, add new licenses to existing servers, set up the license manager in a network environment, and set up redundant servers for high-availability and failover using Sybase Software Asset Management (SySAM).

You will need to refer back to this chapter when you begin installing Adaptive Server in a network environment.

SySAM overview

SySAM (Sybase Software Asset Manager) is a licensing mechanism that:

- Provides System Administrators with a means to monitor their site's use of Sybase products and optional features.
- Enables select Adaptive Server 12.5 features.
- Enables failover to other registered installations of Adaptive Server in the event of failure.

Adaptive Server features licensed through SySAM

The Adaptive Server Enterprise license must be registered through SySAM before any of the optional features can be enabled.

Table 3-1 describes the Adaptive Server features that are licensed through SySAM.

Table 3-1: Licensed Adaptive Server features

| Feature name | License name | Description |
|------------------------------------|--------------|---|
| Adaptive Server | ASE_SERVER | The basic Adaptive Server Enterprise product, without optional features. |
| Java in Adaptive Server | ASE_JAVA | Java and XML classes in Adaptive Server databases. |
| Advanced Security Mechanisms | ASE_ASM | Network-based authentication and encryption using DCE and CyberSafe. |
| Distributed Transaction Management | ASE_DTM | Distributed transaction management features, including the Adaptive Server version 12.5 XA interface. |
| Enhanced Full-Text Search | ASE_EFTS | Enhanced full-text search engine |
| External file system | ASE_XFS | Management of non-relational Web content. |
| LDAP directory services | ASE_DIRS | Lightweight directory services. |

Note The Full-text search engine has been enhanced and now requires a license to install and configure.

If you plan to use Adaptive Server 12.5 with Replication Server, DirectConnect, or OpenSwitch, you must upgrade those products to the latest versions to ensure compatibility with new features from Adaptive Server 12.5.

How SySAM works

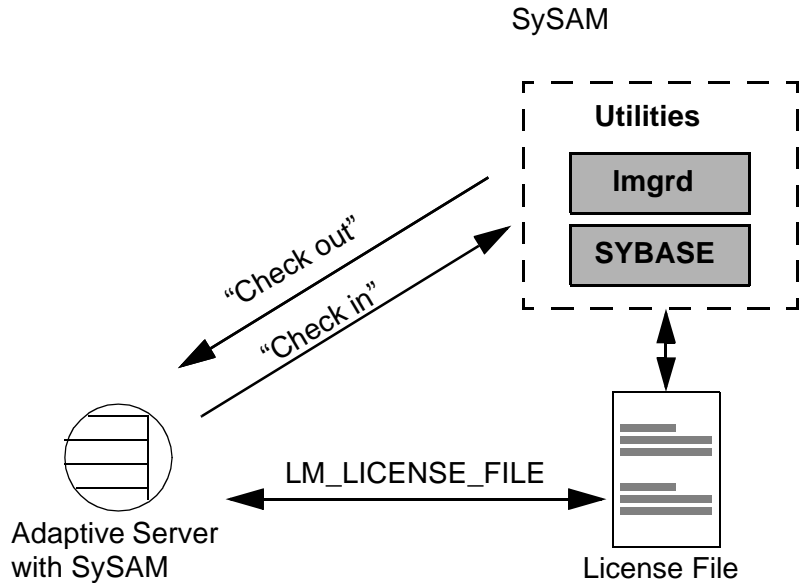
SySAM “checks out” licensed features to users, and checks them back in when they are no longer needed. The basic components of SySAM are:

- One or more Adaptive Server features
- One or more license files

- The SySAM software, which consists of a license management daemon and a Sybase daemon

Figure 3-1 shows the relationship between these components.

Figure 3-1: SySAM components



Adaptive Server check-out procedure

When you start Adaptive Server 12.5, it attempts to locate the license file *license.dat* located in `$SYBASE/$SYBASE_SYSAM/licenses`. If the file does not exist, Adaptive Server looks for the license file specified in the environment variable `LM_LICENSE_FILE`.

The license file contains a pointer to the SySAM software on the primary server known as the **license host**, or to a remote license host where SySAM is running. SySAM consists of a utility, `lmutil`, to manage licensing activities, and two daemons—the license management daemon, `lmgrd`, and the `SYBASE` daemon. The daemons handle requests to check in or check out licensed features, as shown in Figure 3-1.

Using information in the license file, Adaptive Server connects to SySAM and attempts to check out a license (ASE_SERVER) for the base Adaptive Server product. If the ASE_SERVER license is checked out successfully, Adaptive Server continues to boot and attempts to check out any optional features enabled in *server_name.cfg*.

If Adaptive Server is configured to use optional features, such as distributed transaction management (ASE_DTM) or Java in Adaptive Server (ASE_JAVA), it attempts to check out licenses for those features during the start-up process. If a license is not available for an optional feature, Adaptive Server still starts, but the feature cannot be used.

Starting Adaptive Server with optional features

Optional features can be turned on or off using *sp_configure* parameters. Table 3-2 lists the parameters used to configure optional features.

Table 3-2: Configuration parameters for optional features

| Feature name | License name | Configuration parameter |
|------------------------------------|--------------|--|
| Distributed Transaction Management | ASE_DTM | enable xact coordination enable DTM |
| Java in Adaptive Server | ASE_JAVA | enable java |
| Advanced Security Mechanisms | ASE_ASM | use security services |
| Enhanced full-text search | ASE_EFTS | enable full-text search |
| External file system | ASE_XFS | enable file access |
| LDAP directory services | ASE_DIRS | n/a |

Note Adaptive Server's support for distributed transaction management protocols: XA and Microsoft DTC protocols, requires configuration parameter *enable DTM* to be set and also a license to be available.

Distributed transaction management support through Adaptive Server Transaction Coordinator requires configuration parameter *enable xact coordination* to be set. However, this feature is available with no license requirements.

By default, the configuration parameters for optional features are set to 0 (off). To enable an optional feature, use `sp_configure` to set its configuration parameter to 1, and restart Adaptive Server.

Some features require additional preparation before you can use them. For more information, see Table 3-2 and the *Adaptive Server Configuration Guide*.

- For information about configuring Adaptive Server with Distributed Transaction, see *Using Adaptive Server Distributed Transaction Management Features*.
- For information about configuring Adaptive Server Full Text Search capabilities, see *Full Text Search Specialty Data Store User's Guide*.
- For information about Java in Adaptive Server, see *Java in Adaptive Server Enterprise*.
- For information about lightweight directory services, see the *System Administration Guide*.
- For information about advanced security mechanisms, see the *Component Integration Services User's Guide*.

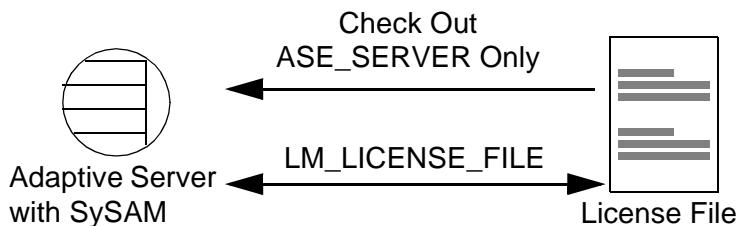
Types of SySAM systems

The basic components of SySAM can be arranged for a standalone server, for a collection of servers on a network accessing a primary license host, and with multiple, redundant servers across a LAN or a WAN, for back-up and load-balancing purposes.

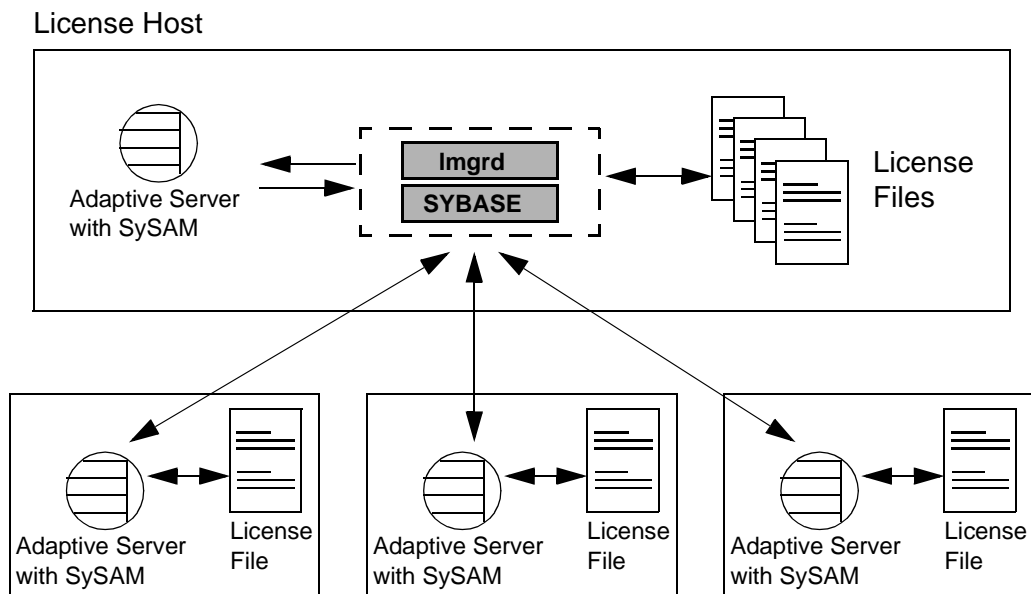
The procedures for registering licenses depend on which instance of Adaptive Server (primary or secondary license host) you are installing.

Standalone system

In a standalone system *all* of the SySAM components shown in Figure 3-2 reside on a single machine. The license file points directly to the asset management software, which runs on the same machine. This machine is called the **license host**.

Figure 3-2: Standalone system without optional features**Network system**

In a network system the license host runs SySAM, which handles check-in and check-out requests from all **secondary servers**, as shown in Figure 3-3.

Figure 3-3: Network system

The license files for secondary servers contain only pointers to the license host. Secondary servers check out licenses from the license host through the SySAM software.

Note If you are using networked software asset management system, you must install the software on the license host before installing it on secondary servers.

Redundant servers

As part of a failover plan or a load-balancing system, you can have multiple servers running the same Adaptive Server configuration, using a redundant licenses across all servers. The **redundant server system** can be connected by a local-area network (LAN) or by a wide-area network (WAN).

A redundant server system is an excellent load-balancing mechanism for physically distant servers. The redundant license allows you to specify local servers as the first in the queue, and make remote servers available as backup. The SySAM application attempts to check out a license from a license-file list, starting with the first server. If that server fails for any reason, the second server in the list is contacted, and so on.

The redundant server system also provides failover protection. If one server in the list goes down for any reason, the second server responds to the license request. If that server fails too, the next server in the list is contacted, and so on.

SySAM in the network environment

Before you begin installing additional Adaptive Servers in a network system, the license manager must be installed and running on the license host. To verify that the license manager is running, see “SySAM administration” on page 37.

To install Adaptive Server on secondary servers, you need to have your Sybase License Certificate available. Also, you need to know:

- The host name of the license host.
- The port number on the license host where the asset manager listens for connections.
- The asset management software installation directory on the license host.

When installing Adaptive Server on secondary servers, the license manager prompts you for licenses after the installer has installed the products on to the hard drive.

- 1 The license manager prompts: “Do you have Sybase Software Asset Management Certificates to register?”

Click No.

- 2 Next, it prompts: “Have you registered Sybase Software Asset Management Certificates at a central license host?”

Click Yes.

- 3 Enter the following information about the license host from which this machine will be checking out Adaptive Server licenses:
 - License Manager Host is the host name of the license host on which you installed the asset management software.
 - License Manager Port is an unused port number on the License host. The local Adaptive Server uses this port number to contact the primary network node when checking out licenses.
- 4 Click OK.

Continue with the installation procedures described in Chapter 4, “Installing Sybase Servers.”

See the Chapter 6, “Installing Sybase PC-Client Products” if you need help configuring client products.

Redundant servers

As part of a failover plan or a load-balancing system, you may have multiple servers running the same Adaptive Server configuration, using redundant licenses across all servers. The redundant servers can be connected by LAN or by WAN.

To configure redundant servers:

- 1 Follow the instructions in this guide to install Adaptive Server on each of the servers in the redundant system.
- 2 As you install the software, make a note of the machine names and the dedicated port number that SySAM uses on each machine.
- 3 For *each* machine in the redundant system, when the license manager prompts for license information, enter:
 - Order Number
 - Feature Name
 - Feature Count
 - Software Version
 - Authorization Code

The certificate information is provided in your Adaptive Server package on a printed Sybase Software Asset Management Certificate.

- 4 Click Done when you have entered all licenses.

Complete the installation of Adaptive Server. See Chapter 4, “Installing Sybase Servers”.

- 5 After the installation is completed, and for each server in the redundant system, set the LM_LICENSE_FILE environment variable to:

```
"port@machine:port@machine:port@machine".
```

For example, if you use port 29722 on each machine and the machine names are Huey, Dewey, and Louie, the environment variable that points to the license file would look like this:

```
LM_LICENSE_FILE="29722@huey:29722@dewey:29722@louie"
```

The first server in the license list is the first server queued by the licensing software.

SySAM administration

If you are installing servers in a network environment, SySAM must be installed and running on the license host before installing Adaptive Server on additional servers. This section describes how to start SySAM manually or as an automatic service, and other SySAM administration information.

Verify the software is running

To verify that the software is running:

When the lmgrd daemon starts, it automatically starts up the SYBASE daemon. To verify that the license management software is running on the system. Use:

```
$SYBASE/SYSAM-1_0/bin/lmutil lmstat -c
```

Make sure that both lmgrd and SYBASE are running, i.e. “up” before you continue the installation or start Adaptive Server.

If the SySAM software is not running, see Starting the software manually.

Starting the software manually

If the license manager is not running, you must start the license manager manually.

- Use the UNIX command:

```
$SYBASE/$SYBASE_SYSAM/bin/lmgrd -c  
$LM_LICENSE_FILE &
```

- Alternatively, run the *startd.sh* script in the bin directory:

```
$SYBASE/$SYBASE_SYSAM/bin/startd.sh $SYBASE/$SYBASE_SYSAM
```

Adding feature licenses

This section describes how to update the license file to enable Adaptive Server features.

To add features for which you have licenses:

- 1 Enable the new feature. See “Starting Adaptive Server with optional features” on page 32.
- 2 Shut down Adaptive Server. See Chapter 7, “Starting and Stopping Servers.”
- 3 Verify that the license manager software is running. See “SySAM administration” on page 37.
- 4 Log on to the machine where the license manager is installed (license host).
- 5 Launch the license manager:

```
$SYBASE/$SYBASE_SYSAM/bin/lmgr
```

- 6 Click Yes when prompted: “Do you have Sybase Software Asset Management Certificates to register?” The SySAM License Manager screen prompts you for:
 - Order Number
 - Feature Name
 - Feature Count
 - Software Version
 - Authorization Code
- 7 Click More until you have entered all licenses.
- 8 Click Done.

The new license is appended to the end of the *license.dat* file.

If you encounter problems with new licenses, check the *lmgrd.log* file in the `$SYBASE/$SYBASE_SYSAM/log` directory to see that there were properly appended to the license file.

Installing Sybase Servers

This chapter provides instructions for installing a new Adaptive Server, Backup Server and Monitor Server.

| Topic | Page |
|------------------------------|------|
| Overview | 41 |
| Installation Definitions | 42 |
| Pre-installation tasks | 42 |
| Installing server components | 47 |

Overview

The Red Hat Package Manager (RPM) installs the Adaptive Server and related files on to the hard drive. The installation process allows you to install one or more of the following Sybase servers:

- Adaptive Server
- Backup Server
- Monitor Server
- Enhanced Full-Text Search (FTS) engine

The basic process for installing Sybase servers is:

- 1 Perform the pre-installation tasks as described in “Pre-installation tasks” on page 42.
- 2 Perform the installation using methods described “Installing server components” on page 47.
- 3 Perform the post-installation tasks as described in Chapter 5, “Post-Installation Tasks.”

Installation Definitions

In this book, these terms are defined as follows:

- *Install* – to install the Adaptive Server files on to the hard drive.
- *Configure* – To configure attributes to make the Adaptive Server product fully functional.
- *Upgrade* – to configure an existing Adaptive Server to a more recent Adaptive Server version level.
- *Server* – provides a service in client/server computing. Examples include Adaptive Server, Backup Server™, Historical Server, Monitor Server, and XP Server™.
- *Client* – requests a service in client/server computing. Sybase Central, PowerDynamo, PowerDesigner®, SQL Modeler™, and end-user applications are clients.

Pre-installation tasks

By default, Adaptive Server installs to directory that it creates during the installation process. The new directory does not conflict with default installation directories of earlier versions of Adaptive Server.

Installation directory defaults for Adaptive Server are as follows:

| Adaptive Server version | Default installation directory |
|--------------------------------|--------------------------------|
| Adaptive Server 11.0.3.3 ESD#6 | /opt/sybase |
| Adaptive Server 11.9.2 | /opt/sybase-11.9.2 |
| Adaptive Server 12.5 | /opt/sybase-12.5 |

Warning! If you install Adaptive Server for an upgrade, you cannot install it into the same directory as the currently installed version. The new version overwrites the old one, causing the upgrade to fail. You must have both the old and the new versions of Adaptive Server to perform an upgrade.

Table 4-1 is a list of the environment variables and their default settings.

Table 4-1: Environment variables

| Variable | Set to... |
|---|---|
| \$SYBASE | The directory you specified for the installation. By default, <i>/opt/sybase-12_5</i> |
| \$SYBASE_ASE | ASE |
| \$SYBASE_FTS | FTS |
| \$SYBASE_OCS | OCS |
| \$SYBASE_SYSAM | SYSAM |
| \$DISPLAY | <i>your_display</i> :0.0 where <i>your_display</i> is the name of your machine. This variable must be set when performing a remote installation. |
| Note You must manually set the \$DISPLAY environment variable. | |
| \$USER | sybase |
| \$PATH | <i>:/sbin</i> This is an optional setting for use with the Linux dummy-network interface. |

Adaptive Server and its components on Linux are released as Redhat Package Manager (RPM) packages. You must have root permissions to install RPM packages. After you have installed the packages, you should log in as the user “sybase” to configure your servers.

When you log in as “sybase” for the first time after installation, verify that the SYBASE environment variable is correctly set and that *\$SYBASE/\$SYBASE_ASE/bin* and *\$SYBASE/\$SYBASE_OCS/bin* are in your search path.

Before installing Sybase servers:

- 1 Read the release bulletins for the latest information on the products (Adaptive Server, Monitor Server, and so on) that you are installing. See the “Special Installation Instructions” section.
- 2 Install operating system patches, if required. Required operating system patches are described in the release bulletin for the product.

Your Sybase product shipment includes printed release bulletins; they are also available at <http://www.sybase.com/support/manuals>.
- 3 Review the SySAM procedures and plan your client/server configuration using Chapter 3, “Sybase Software Asset Management (SySAM).”
- 4 Create a “sybase” account on your system to perform all installation tasks.

The “sybase” user must have permission privileges from the top (or root) of the disk partition or operating system directory down to the specific physical device or operating system file.

- 5 Log in to the machine as the “sybase” user.

It is important to maintain consistent ownership and privileges for all files and directories. A single user—the Sybase System Administrator with read, write, and execute permissions—should perform all installation, upgrade, and setup tasks.

- 6 Verify that you have a directory */opt* on your machine, and that it has sufficient space to accommodate the software. If an */opt* directory does not already exist, the Linux RPM creates it in the root directory.

If you are installing for an upgrade, see Chapter 8, “Upgrading Sybase Servers.”

- 7 Verify:

- That the directory location for the Sybase installation has sufficient space. Include an additional 2.4MB space for the installation program.
- That your operating system meets the release-level, RAM, and network-protocol requirements.

Space requirements for Adaptive Server Enterprise products are listed in Chapter 2, “Installation Requirements.”

- 8 If you are performing an installation on a remote machine, set the DISPLAY variable of the local X11 server.

- On the machine where you are monitoring the installation of Adaptive Server, enter:

```
% export DISPLAY=your_display:0.0
```

where *your_display* is the name of the local display.

- On the machine where you are installing Adaptive Server, enter:

```
$ xhost +
```

- 9 Verify that your network software is configured.

Sybase software uses network software even if Adaptive Server and Sybase client applications are installed on a machine that is not connected to a network.

If you are having connection problems, or if you want to verify your network configuration, complete these steps:

- a Log in to your host machine with this UNIX command:

```
telnet host
```

This command gives you a login prompt for the same machine.

- b Without logging out, log in to your host machine a second time, using the same command.

If you cannot log in to the machine a second time, your network software is not properly configured to run in loopback mode. See the network software documentation, or consult your vendor.

- c Log out twice.

- 10 Adjust shared memory segments.

Depending on the number and types of devices you use for backup (dump) and recovery (load), you may need to adjust the shared memory segment parameter in the operating system configuration file to accommodate concurrent Backup Server processes. The default number of shared memory segments available for process attachments is 6.

- 11 The operating system shared memory default, for most Linux releases is 32MB. The minimum required by Adaptive Server is 64MB. To tune the maximum shared memory segment size per request heap, go to the */proc* directory. Check the available operating system size.

```
# ipcs -lm | grep 'max seg size'
max seg size (kbytes) = 32768
```

Choose the new size for the operating system shared memory parameter. For example, to change the shared memory to 512MB, enter:

```
#echo '512 * 1024 * 1024' |bc >
proc/sys/kernel/shmmax
```

Verify that the change takes effect:

```
# ipcs -lm | grep 'max seg size'
max seg size (kbytes) = 524288
```

This change is a dynamic Linux kernel tuning. The operating system does not need to be restarted.

To avoid having to make this change after each restart, add this line to the */etc/rc.d/rc.local* file:

```
echo '512 * 1024 * 1024 | bc >
```

```
/proc/sys/kernel/shmmax
```

- 12 Determine the location, type, and size of each database device.

You must provide this information during the installation of Adaptive Server. Devices are operating system files, or raw partitions, used to store databases and database objects. For details, see “Determining the location, type, and size of a database device” on page 26.

Installing the Linux dummy-network interface

If your Linux system does not have network communication hardware installed, or if it uses the dynamic host configuration protocol (DHCP) for network IP address configuration, you can install Adaptive Server components using these procedures.

❖ Installing a dummy-network interface

- 1 Add */sbin* to your *\$PATH* environment setting:

```
# PATH=$PATH:/sbin
```

- 2 Add an IP address and host name to your */etc/hosts* table. Sybase recommends that you use an address in a free Internet address band.

For example, an Internet address of 10.0.0.1 would result in:

```
echo 10.0.0.1 $HOSTNAME >> /etc/hosts
```

Note Sybase recommends that you do not use “localhost” as an alternative.

- 3 Preload the dummy network driver on the machine.

```
# modprobe dummy
```

If this command is successful, you do not receive any messages from the server.

- 4 Bind the dummy established network interface to the *\$HOSTNAME* configured IP address.

```
# ifconfig dummy0 $HOSTNAME
```

- 5 For a DHCP environment, add a route other than “default” to the established virtual network setup.

```
# route add $HOSTNAME dummy0
```

6 Check the network binding:

```
# ifconfig dummy0
```

This command gives below listed output:

```
dummy0    Link encap:Ethernet  HWaddr 00:00:00:00:00:00
inet addr:10.0.0.1  Bcast:10.255.255.255  Mask:255.0.0.0
UP BROADCAST RUNNING NOARP  MTU:1500  Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
```

7 Check the route binding:

```
# route
```

This command shows the routing as it is set up.

```
Kernel IP routing table
Destination      Gateway         Genmask         Flags Metric Ref    Use Iface
HOSTNAME         * 255.255.255.255  UH              0      0          0 dummy0
10.0.0.0         * 255.0.0.0      U               0      0          0 dummy0
127.0.0.0        * 255.0.0.0      U               0      0          0 lo
```

Installing server components

This section describes how to install Adaptive Server version 12.5. Follow the instructions in this section to install Adaptive Server, Backup Server, and Monitor Server.

Note Installation of Adaptive Server 12.5 RPM packages require “root” privileges. Once the packages have been installed, log out as root, and log in as the “sybase” user to perform all configuration tasks.

Installing using Red Hat Package Manager (RPM)

On the Linux platform, you install Adaptive Server with the Red Hat Package Manager (RPM), which is included in the standard bundle by all Linux distributors.

To Install Adaptive Server products using RPM:

- 1 Verify that you are logged in as the “root” superuser and that your environment is set up as described in “Pre-installation tasks” on page 42.
- 2 Verify that you have a directory called */opt* on your server, and that it has sufficient space to accommodate unloading the software. If an */opt* directory does not already exist, the Linux RPM will create it in the root directory.
- 3 Place the Adaptive Server product CD in the CD-ROM drive.
- 4 If you need to mount the CD, enter:

```
# mount -t iso9660 /dev/cdrom /mnt/cdrom
```

To install Adaptive Server products using Red Hat Package Manager (RPM), you can either:

- Install packages using the graphical interface
- Install packages from the command line

For help on RPM, enter:

```
# /bin/rpm
```

❖ **Install packages using the graphical interface**

- 1 From the Linux task bar, select the Gnome or KDE icon that launches the Linux applications menu.
- 2 Select System | GnoRPM (or the KDE RPM).

This launches the RPM manager and displays a hierarchical view of the system applications.

- 3 Select Install.

This launches a new dialog box.

- 4 Select Add.

This displays a list of all RPM packages that have not yet been installed.

The RPMS packages come from the CDROM or from the location to which you download the packages. For CDROMs, the location is usually */mnt/cdrom/RPMS*.

If you have the X11 environment enabled when you install the CDROM, then a pop-up window displays, showing the contents of the CDROM.

To bring the desired packages to the GUI package installers like GNOME, gnorpm or KDE package, drag and drop the packages into the installer frontends.

- 5 Highlight each package and click Add after each selection. The new packages are placed in the appropriate directory under Packages | Applications | Databases | Sybase.

Note All products depend on the *sybase-common-12.5-*nn*.i386.rpm* package. Either install the “sybase-common” package before you install any other package, or install all packages at once.

- 6 As each package is added to the RPM interface, a check box appears to its right. To install all RPM package simultaneously, select All but Installed from the drop-down list.
- 7 Click Install.

RPM creates the */opt/sybase-12.5* target directory and installs all packages.
- 8 Once the installation is complete, select Close.
- 9 Unmount the CD from the drive at your # prompt by entering:


```
# umount /dev/cdrom
```


Remove the CD by entering:


```
# eject
```
- 10 RPM sets correct product permissions and ownership.

❖ **Install packages from the command line**

If you have any beta packages installed, follow steps one and two. If you do not have the beta packages installed, skip to step three.

Start command-line installation with the *sybase-common-12.5-*nn*.i368.rpm* and *sybase-ase-12.5-*nn*.i386.rpm* packages as they are a prerequisite to all other packages.

- 1 Check to see if you have any beta packages in the installation directory. Enter:


```
# rpm -qa | grep sybase
```
- 2 Remove the beta packages:


```
# rpm -e sybase-common-12.5-n \  
sybase-ase-12.5-n \  
sybase-monserver-12.5-n \  
sybase-openclient-12.5-n \  
sybase-sunjdk-1.1.8-n \  

```

```
sybase-doc-12.5-n \  
sybase-french-12.5-n \  
sybase-german-12.5-n
```

where *n* is the beta version numbers 1 or 2.

3 Mount the CDROM:

```
#mount /mnt/cdrom
```

4 Go to the mounted CDROM:

```
# cd /mnt/cdrom
```

5 To view the CDROM contents, (Cover letter, .PDF, and README) enter:

```
ls -l
```

6 From the */RPMs* directory, install the packages:

```
# rpm -Uhv sybase-common-12.5-nn.1DE-1.i386.rpm \  
sybase-ase-12.5-nn.i386.rpm \  
sybase-efts-12.5-nn.i386.rpm \  
sybase-esql-12.5-nn.i386.rpm \  
sybase-openclient-12.5.nn.i386.rpm \  
sybase-openserver-12.5.nn.i386.rpm \  
sybase-sqlremote-7.0.2.nn.i386.rpm \  
sybase-doc-12.5-nn.i386.rpm \  
sybase-french-12.5.nn.i386.rpm \  
sybase-german-12.5.nn.i386.rpm \  
sybase-chinese-12.5.nn.i386.rpm \  
sybase-japanese-12.5.nn.i386.rpm \  
sybase-korean-12.5.nn.i386.rpm \  
sybase-polish-12.5.nn.i386.rpm \  
sybase-portuguese-12.5.nn.i386.rpm \  
sybase-spanish-12.5.nn.i386.rpm \  
sybase-thai-12.5.nn.i386.rpm
```

where *nn* is the minor version number of the rpm package you are installing.

Warning! Do not use the upgrade (-U) option of RPM. This will adversely affect any existing installations of 11.9.2 and 11.0.3.3.

The packages can be relocated during installation using the rpm(1) --relocate option. Refer to the rpm(1) manual pages for correct command line instructions.

7 Unmount the CD from the drive at your # prompt by entering:

```
# umount /dev/cdrom
```

Remove the CD by entering:

```
# eject
```

- 8 RPM sets correct product permissions and ownership.

RPM sets the correct product permissions and ownership necessary to complete the installation and configuration.

Registering Adaptive Server features

If you install any components without the appropriate license information, only Adaptive Server 12.5, without licensed features, is enabled.

Warning! The following SySAM licence manager instructions are for the installation of Adaptive Server on the primary license host. If you are installing Adaptive Server on a secondary server, see the SySAM licensing instructions in “SySAM in the network environment” on page 35.

- 1 Before logging in as “sybase”, install a password on the sybase account that you created. To create a password, as “root” enter,

```
# password sybase
changing password for user sybase.
New password: *****
Retype new password: *****
```

Now log off, and login under the X11 environment as “sybase” with the newly created password.

- 2 Start the license manager:

```
$SYBASE/$SYBASE_SYSAM/bin/lmgr
```

- 3 The license manager prompts: “Please select a license type.”

Select ASE, and click Next.

- 4 The license manager prompts: “Do you have a Sybase Software Asset Management Certificate to register?”

Click Yes.

- 5 Enter information from the Sybase License Certificate for each Adaptive Server feature you have purchased:

Note The Full-text search engine has been enhanced and now requires a license to install and configure.

- Order Number: Enter your Sybase order number.
- Feature Name: Enter the name of the Adaptive Server feature. Valid Adaptive Server feature names are:
 - ASE_SERVER
 - ASE_JAVA
 - ASE_DTM
 - ASE_ASM
 - ASE_EFTS
 - ASE_XFS
 - ASE_DIRS
- Feature Count: Enter your license count number.
- Software Version: Enter the Adaptive Server software version.
- Authorization Code: Enter the license key for the purchased feature.

The installer records the information for the current feature in the license file and prompts you to enter information for an additional feature.

Note Entries are case sensitive. Enter the values exactly as they appear on the Sybase Software Asset Management Certificate. See Chapter 3, “Sybase Software Asset Management (SySAM)” for detailed information about using the license manager.

- 6 Click More. . . if you have purchased additional licensed features. The installer records the information for the current feature in the license file and prompts you to enter information for an additional feature.

- 7 Click Done after you have entered information from all of the Sybase License Certificates you have purchased.

Note After you click Done, a message displays stating that the `lmreread` command failed with this error:

```
license manager: can't initialize: No SERVER lines
in license file
```

This is expected behavior because the license daemon has not been started. The next step starts the daemon, and you can proceed with the configuration of the servers.

- 8 Start the SySAM license manager daemon:

```
$SYBASE/$SYBASE_SYSAM/bin/startd.sh $SYBASE/$SYBASE_SYSAM
```

The installer records all license information and prompts you to configure the components you have installed. If you encounter problems, check the installation log file to see a record of the installation process. The file is located in *\$SYBASE/Installer.log*.

Configuring the servers

This section describes how to configure a basic Adaptive, Backup, Monitor, XP Server, and EFTS. For information on advanced configuration topics, localization, and enabling additional functionality, see *Adaptive Server Enterprise Configuration Guide*.

Before configuring this installation of Sybase Adaptive Enterprise Server, Sybase recommends that you install a Sybase DBA account and correct product permissions. This post configuration can be performed at any time by running the *sybinstall.sh* shell script as root (`uid=0`).

```
${SYBASE}/install/sybinstall.sh'
```

❖ Configuring Adaptive Server

- 1 Log in as sybase.
- 2 Start the configuration utility:

```
$SYBASE/$SYBASE_ASE/bin/srvbuild
```

The Adaptive Server Setup & Configuration screen is displayed. Select Configure a new server.

- 3 The `srvbuild` screen displays.

Only the server types that were installed display on the `srvbuild` screen. For example, if you did not install Full-Text Search, it does not appear on the `srvbuild` screen.

- 4 Select, at a minimum, Adaptive Server. Click the check box to the left of the server name.

You can choose all server types at once or any combination of server types.

- 5 Provide a name for each server you install.

`srvbuild` provides a default server name using the name of the machine you are installing on. You can change the default name.

Server naming conventions are as follows:

- The name should not exceed 30 characters.
- The initial character must be a letter.
- The characters that follow must be letters, numbers, or underscores (`_`).
- Characters not allowed are: control characters, spaces, ampersands (`&`), asterisks (`*`), periods (`.`), and question marks (`?`).
- Use a name that is descriptive of your intended use and unique to your network.
- Use the extensions (`“_back”`, `“_mon”`, `“_XP”`) to relate other server types to Adaptive Server.
- The XP Server name must be all uppercase, followed by an underscore (`_`) and the two-letters `“XP”`. Although `srvbuild` does this for you automatically, you should be aware of this requirement in case you need to modify the XP Server name in the `interfaces` file.

All server names are stored in the *interfaces* file by `srvbuild` and are used by clients and other servers to establish communication.

- 6 Click OK.

Note `srvbuild` consecutively displays the attributes screen for each server you choose, in the order listed on the screen. For example, if you chose Adaptive Server, Backup Server, and XP Server, but did not select Monitor Server, the attribute screen for Adaptive Server appears first. When you finish choosing the attributes for Adaptive Server, and click OK, the Backup Server attribute screen appears. When you are finished with that screen, the XP Server attribute screen appears, and so on.

- 7 From the drop-down list at the top of the Adaptive Server attribute screen, select the logical page size for this installation of Adaptive Server.

All pages in this installation will have the logical page size you specify here. If this installation of Adaptive Server has a logical page size of 2K, and you decide that you want a logical page size of, for example, 8K, you must create a separate installation of Adaptive Server with the 8K page size.

The size of the master database is automatically allocated to accommodate the logical page size. You can increase the master database size if you want. The maximum size is the size of the file system, up to 2GB, the minimum is 25MB.

If you are adding additional server languages for localization, you need an additional 2–7MB of space, per language in the master database. See Chapter 2, “Installation Requirements” for language module sizes

- 8 Provide an operating system file for the master and sybsystemprocs device paths:

```
/op_system_dir/master.dat  
/op_system_dir/sybsystemprocs.dat
```

where *op_system_dir* is the name of the operating system directory you are going to use.

Note `sysprocsdev` is the default system name for this device. However, it is frequently referred to as the `sybsystemprocs` device, as in the Adaptive Server attribute screen, since it stores the `sybsystemprocs` database.

Sybase recommends that the `sybsystemprocs` device be located on a device separate from the master device. The `sybsystemprocs` database usually increases in size from release to release.

- 9 For the sybsystemprocs device and database sizes, enter 100.
- 10 Provide an error log path name.

The default is:

`$SYBASE/$SYBASE_ASE/install/errorlog`

- 11 For Directory Services, specify the transport type, host name, and port number.

These entities define how servers and clients find and communicate with each other on the network. `srvbuild` adds this entry to the `interfaces` file during the installation process.

- Transport type – the menu provides a list of network protocols supported by Sybase. TCP is the default. Some network protocols in the menu may not be valid for your platform. Check with your operating system administrator.
- Host name – do not change the host name. The host name is the machine name where you are installing Adaptive Server.
- Port number – if you did not change the default protocol, the port number displayed is the first available port address. The port number you specify cannot be used for any purpose other than installing Adaptive Server.

If you change the default port number, verify that the new port is not in use.

If you selected a protocol other than the default, you may not be able to use the first available port number supplied. Check with the operating system administrator to see what port numbers are available.

- 12 Click Edit Advanced Adaptive Server Attributes if you want to:
 - Specify a sybsystemdb device path for the sybsystemdb database. See *Sybase Adaptive Server Configuration Guide*.
 - Place shared memory files in a directory other than the default location under `$SYBASE`.
 - Change your default Backup Server.

- 13 To create the server, click Build Server. To continue configuring other Sybase servers, select OK.

`srvbuild` creates the *master.dat* and *systemprocs.dat* files for you.

srvbuild displays a warning regarding the use of an operating system file for the master and sybssystemprocs device paths. If this Adaptive Server is not intended for production use, you can ignore this warning.

❖ **Configuring Backup Server**

This procedure assumes that you have completed the “Pre-installation tasks” on page 42.

- 1 Provide the related Adaptive Server name, which is the Adaptive Server that will use this Backup Server.

If the related Adaptive Server is created in the same srvbuild session as Backup Server, this field does not appear. The Adaptive Server you are creating is used as the related Adaptive Server.

- 2 Provide the Sybase System Administrator (SA) user name for the related Adaptive Server, or accept the default value.

If the related Adaptive Server is created in the same srvbuild session as Backup Server, this field does not appear.

- 3 Provide the related Adaptive Server SA password.

If the related Adaptive Server is created in the same srvbuild session as Backup Server, this field does not appear.

- 4 Provide an error log path for Backup Server.

- 5 Tape backups are not supported. Leave this field blank.

- 6 Specify the language you want Backup Server to use for its error messages. The default language is U.S. English.

The language available depends on which language modules were installed on your system. By default, all character sets are copied from the CD when performing a standard or full install.

- 7 Specify the character set you want Backup Server to use.

The default character set is iso_1

Other options are:

- cp850: Code Page 850 (Multilingual) character set
- cp437: Code Page 437 (United States) character set
- mac: Macintosh default character set for Western European locales
- ISO 8859-1 (Latin-1) Western European 8-bit character set

- UTF-8 character set

8 Specify the maximum number of network connections.

This option specifies the total number of network connections available if you are backing up or restoring a database over the network to a remote Backup Server.

This option is not necessary if you are backing up to the default Backup Server, which is local. Set to 0.

9 Specify the maximum number of server connections. You can accept the default or change it.

This option specifies the number of server connections to Backup Server. Backup Server requires at least:

- Two connections for each dump operation
- One connection for each load operation
- One connection for volume change messages

Allow a maximum of three times the number of expected concurrent dump and load sessions.

10 For the Directory Services, specify the transport type, host name, and port number.

These entities comprise the interfaces file entry. This entry defines how Backup Server and Adaptive Server find and communicate with each other on the network. `srvbuild` adds this entry to the interfaces file during the installation process.

- Transport type – the menu provides a list of network protocols supported by Sybase. TCP is the default. Some network protocols on the menu may not be valid for your platform. Check with your operating system administrator.
- Host name – the host name is the machine name where you are installing Backup Server.
- Port number – if you did not change the default protocol, the port number displayed is the first available port address.

The port number you specify cannot be used for any purpose other than Backup Server.

If you change the default port number, verify that the new port is not already in use. Check with your operating system administrator.

If you selected a protocol other than the default, you may not be able to use the first available port number. Check with the operating system administrator to see what port numbers are available.

11 Click OK.

❖ **Configuring Monitor Server**

This procedure assumes that you have completed the “Pre-installation tasks” on page 42.

If you selected Monitor Server to be installed, the Monitor Server attribute screen displays.

- 1 Provide the related Adaptive Server name, which is the Adaptive Server that will be monitored by this Monitor Server.

This field does not appear if the related Adaptive Server is created in the same `srvbuild` session as Monitor Server.

- 2 Specify the Sybase System Administrator (SA) user name. This field does not appear if the related Adaptive Server is created in the same `srvbuild` session as Monitor Server.

This option is required unless the `USER` environment variable is set.

This name specifies the Monitor Server superuser. It must be a valid login account for the Adaptive Server to be monitored. For production systems, Sybase recommends that you use the “sa” or “sybase” user account as the Monitor Server superuser.

Monitor Server uses the superuser account to connect to Adaptive Server to run special stored procedures and to verify shutdown requests that can be performed only by the superuser.

- 3 Provide the Adaptive Server SA password. This is the password of the Monitor Server superuser. This field does not appear if the related Adaptive Server is created in the same `srvbuild` session as Monitor Server.

If you omit the password, Monitor Server prompts for a password during start-up.

- 4 Determine the maximum number of connections allowed to Monitor Server. Valid values are 1–20.

A connection to Monitor Server can be any of:

- A monitor in Sybase Central,

- An active Historical Server session or a scheduled Historical Server recording session, or
- A connection to Monitor Server from a Monitor Client-Library application.

For example, if a user starts a Historical Server recording session, and another user opens the Performance Summary Monitor in Sybase Central, two connections to Monitor Server are used. If Monitor Server is configured for five connections, there are only three more connections available.

5 Specify the error log path.

This is the path name of the Monitor Server log file, where informational and error messages are logged.

The default location is the directory from which Monitor Server is started and the error log name is *servername_mon.log*, where *servername* is the name of the server.

6 Specify the path of the configuration file.

This is the path name of the Monitor Server configuration file. A default configuration file is not provided. This file is optional, but if specified, the values in the configuration file override the default runtime parameters that are otherwise used by Monitor Server.

7 Specify the shared memory directory, where the *servername.krg* shared memory file is located and *servername* is the name of the Adaptive Server.

If you omit this parameter, Monitor Server looks for the shared memory in the *\$SYBASE* directory.

If you did not change the default location while configuring the Adaptive Server, the default value is correct. Otherwise, use the value you specified for the server.

8 For the Directory Services, specify the transport type, host name, and port number. These entities comprise the interfaces file entry. This entry enables Monitor Server and Adaptive Server to find and communicate with each other on the network. *srvbuild* adds this entry to the interfaces file during the installation process.

- Transport type – the menu provides a list of network protocols supported by Sybase. TCP is the default. Some network protocols in the menu may not be valid for your platform. Check with your operating system administrator.

- Host name – the host name is the machine name where you installed Monitor Server.
- Port number – if you did not change the default protocol, the port number displayed is the first available port address.

The port number you specify cannot be used for any purpose other than installing the Monitor Server.

If you change the default port number, verify that the new port is not already in use.

If you selected a protocol other than the default, you may not be able to use the first available port number. Check with the operating system administrator to see what port numbers are available.

9 Click OK.

❖ **Configuring XP Server**

If you installed XP Server, the XP Server Attribute Editor screen displays.

- 1 Provide the related Adaptive Server name, which is the name of the Adaptive Server that will be using XP Server.

If the related Adaptive Server is created in the same `srvbuild` session as XP Server, this field does not appear.

Note The XP server name must be in all capital letters, and must follow the format `ASE-NAME_XP`.

- 2 Provide the Sybase System Administrator (SA) user name for the related Adaptive Server, or accept the default value.

If the related Adaptive Server is created in the same `srvbuild` session as XP Server, this field does not appear.

- 3 Provide the related Adaptive Server SA password.

If the related Adaptive Server is created in the same `srvbuild` session as XP Server, this field does not appear.

- 4 For the Directory Services, accept the default values provided.

- 5 Click Build Servers.

`srvbuild` displays a Status Output screen that shows the status of the various installation tasks as they execute.

- 6 When the installation is complete, the Status Output screen displays “Done”. Click Exit to return to the main window. Click Exit, again.

The Adaptive Server and related products have been successfully installed, and minimally configured for use. See “Installing sample databases” on page 72, to begin experimenting with your servers, or see the *System Administration Guide* at <http://www.sybase.com/support/manuals> for more advanced topics.

❖ **Configuring the Full Text Search engine**

The following procedure assumes that you have completed the procedure in the section “Pre-installation tasks” on page 42. To continue installing FTS, complete the following procedure.

- 1 Errorlog: The default is `$SYBASE/$SYBASE_FTS/install/servername.log`
- 2 Collection directory: Accept the default, `$SYBASE/$SYBASE_FTS/collections`
- 3 Default Database: `text_db`
- 4 Language: `us_english`
- 5 Character set: `iso_1`
- 6 Minimum number of sessions: 10
- 7 Maximum number of sessions: 100
- 8 For Directory Services, specify the transport type, host name, and port number. These entities comprise the interfaces file entry. This entry enables FTS and Adaptive Server to find and communicate with each other on the network. `srvbuild` adds this entry to the interfaces file during the installation process.
 - Transport type – the menu provides a list of network protocols supported by Sybase. TCP is the default. Some network protocols in the menu may not be valid for your platform. Check with your operating system administrator.
 - Host name – the host name is the name of the machine where you installed Adaptive Server.
 - Port number – if you did not change the default protocol, the port number displayed is the first available port address.

The port number you specify cannot be used for any purpose other than installing the FTS.

If you choose to change the default port number, verify that the new port is not already in use.

If you selected a protocol other than the default, you may not be able to use the first available port number. Check with the operating system administrator to see what port numbers are available for use.

- 9 To create the server, click Build Server.

srvbuild displays a Status Output screen that shows the status of the various installation tasks as they execute.

See *Full-Text Search Specialty Data Store User's Guide* for information on the this feature.

❖ Installing Sybase Central Java Edition

This section describes how to install and register the Sybase Central Java plug-in.

- 1 Start Sybase Central. From the command line, enter:

```
$ $SYBASE/sybcent32/scjview
```

- 2 From the menu, select Tools | Plug-ins

This launches the Sybase Central Plug-ins dialog box.

- 3 Select Register

- 4 Enter the full path to the plug-in, including the plug-in name. For example:

```
opt/sybase-12.5/ASEP-1_0/ASEPlugin.jar
```

where *opt/sybase_12-5* is the Sybase installation directory.

- 5 Select OK.

- 6 Select Close. The Adaptive Server Enterprise icon should now appear in Sybase Central.

Installing the documentation

Adaptive Server version 12.5 for Linux comes with a separate documentation CD. This chapter describes how to install Sybase documents onto your system.

Online documentation

Documentation can be browsed using the xpdf utility available with standard Linux installations. If your MIME defaults are setup correctly you can open the PDF files launching this browser from the Gnome explorer gmc by double clicking on the document.

Log in as “sybase” under the X11 environment, and enter:

```
$ cd $SYBASE/doc/PDF
$ gmc
```

You can also browse these documents online by visiting the Sybase DynaWeb Web server at Product Manuals at <http://www.sybase.com/support/manuals>. For Linux-specific information, visit Adaptive Server Enterprise for Linux at <http://linux.sybase.com>.

Installing the documentation

While you can view the PDF files directly from the documentation CD, Sybase recommends that you install the files on your system using the Red Hat Package Manager (RPM). Using RPM provides better management of the documentation files, such as allowing you to uninstall all the files at once with a single command.

❖ Using RPM to install the documentation PDF files

- 1 Verify that you are logged in as “root”.
- 2 Verify that you have a directory called */opt* on your server, and that it has at least 40MB to accommodate the documentation PDF files. If an */opt* directory does not already exist, RPM creates it in the root directory.
- 3 Place the documentation CD in the CD-ROM drive.
- 4 At the # prompt, enter:

```
# mount -t iso9660 /dev/cdrom /mnt/cdrom
```

- 5 Unload all of your Sybase documentation files into the */opt/sybase-12.5/doc* directory by entering:

```
# rpm -Uhv sybase-doc-12.5.0.1ESD-1.i386.rpm
```

Note You can relocate the documentation files during installation using the `rpm(1) --relocate` option. For example, to install the documents under */opt/sybase/doc* directory, type:


```
# rpm -hiv --relocate "/opt/sybase-  
12.5"="/opt/sybase" \
```

```
/mnt/cdrom/sybase-doc-12.5-3.i386.rpm
```

- 6 When installation is complete, unmount the CD by entering this command at the # prompt:

```
# umount /mnt/cdrom
```

Remove the CD by entering:

```
# eject
```

You have now successfully installed documentation files onto your system.

Removing Sybase servers

There are several methods that you can use to remove Sybase servers.

- First, stop all Adaptive Servers.
- From a command line, enter:

```
# rpm -e sybase-ase-12.5-3.i386.rpm \  
sybase-doc-12.5-3.i386.rpm \  
sybase-efts-12.5-3.i386.rpm \  
sybase-sqlremote-7.0.2-1.i386.rpm \  
sybase-esql-12.5-3.i386.rpm \  
sybase-openserver-12.5-3.i386.rpm \  
sybase-openclient-12.5-3.i386.rpm \  
sybase-chinese-12.5-3.i386.rpm \  
sybase-french-12.5-3.i386.rpm \  
sybase-german-12.5-3.i386.rpm \  
sybase-japanese-12.5-3.i386.rpm \  
sybase-korean-12.5-3.i386.rpm \  
sybase-polish-12.5-3.i386.rpm \  
sybase-portuguese-12.5-3.i386.rpm \  
sybase-spanish-12.5-3.i386.rpm \  
sybase-thai-12.5-3.i386.rpm \  
sybase-common-12.5-3.i386.rpm
```

It is very important that you list the common package last, since it carries dependencies for all other packages.

- You can use `gnorpm(1)`, and select `uninstall`.

- Alternatively, to erase all Sybase packages in one command, invoke the `--nodeps` to force `rpm(1)` to ignore order dependencies.

Warning! Be sure to use `grep(1)` for “Sybase” to ensure that you erase only Sybase packages.

```
# rpm -e `rpm -qa | grep sybase` --nodeps
```

Post-Installation Tasks

After the installation process, you have a minimally configured Adaptive Server. This chapter describes post-installation tasks to administer Adaptive Server and install stored procedures and sample databases. For more information see the *Adaptive Server Configuration Guide* for your platform.

| Topic | Page |
|--|------|
| Setting environment variables | 67 |
| Verifying that servers are running | 69 |
| Verifying that you can connect to servers | 70 |
| Connecting to Adaptive Server via Sybase Central | 70 |
| Setting the System Administrator password | 71 |
| Installing sample databases | 72 |
| Installing stored procedures | 77 |
| Installing character sets | 77 |
| Installing online help for Transact-SQL syntax | 80 |
| Configuring XP Server after initial installation | 83 |

Setting environment variables

It is crucial to the operation of Sybase products that the system environment variables be set correctly. The following list describes the Sybase environment variables that are set by the Redhat Package Manager (RPM). RPM also generates a file that you can execute if you need to reset environment variables for any reason.

Environment variables are set in the user's environment either interactively or by including them in the user's *.login* and *.cshrc* files (for C shell) or *.profile* file (for Bourne shell).

You can set environment variables by running the *SYBASE.csh* or *SYBASE.sh* script file.

- 1 From a Bourne, enter:

```
. SYBASE.sh
```

2 From a C shell, enter:

```
source SYBASE.csh
```

The script file does not set jConnect or Java utility environment variables. See the product documentation at <http://www.sybase.com/support/manuals>.

The environment variables are:

- DSLISTEN – Defines the name Adaptive Server uses to listen for client connections if there is no name provided on the command line during start-up. If DSLISTEN is not set, and there is no name provided during start-up, the Adaptive Server name defaults to the server name given at installation.
- DSQUERY – Defines the Adaptive Server name that client programs try to connect to if an Adaptive Server is not specified by name during a connection attempt. If DSQUERY is not set, and you do not supply the Adaptive Server name, clients attempt to connect to the server name provided during installation.
- SYBASE – Defines the path of the Sybase installation directory. The installation program sets up the variable SYBASE to point to the release directory specified during installation.
- SYBASE_ASE – Defines the installation directory of Adaptive Server components.
- SYBASE_OCS – Defines the installation directory of Open Client.
- SYBASE_FTS – Defines the directory for the enhanced full-text search components. The enhanced full-text search engine must be licensed before you can use it, and it is only installed if you selected the “full” or “customized” installation option.
- SYBASE_SYSAM – Points to the license-management software directory.
- LM_LICENSE_FILE – Points to the *license.dat* file in the `$SYBASE/$SYBASE_SYSAM/licenses` (`%SYBASE%%SYBASE_SYSAM%\licenses` on Windows) directory.
- PATH – Specifies which directory path to search for executables. The Sybase executables are located in the installed component’s */bin* directory.

When using the source files: *SYBASE.csh* or *SYBASE.sh*, the following paths are added to the beginning of the PATH environment variable:

```
$SYBASE/$SYBASE_ASE/bin:$SYBASE/$SYBASE_OCS/bin:  
$SYBASE/$SYBASE_FTS/bin:$SYBASE/$SYBASE_SYSAM/bin:  
$SYBASE/shared/jre-1_2/bin
```

Note For instructions on setting jConnect environment variables, see the documentation for jConnect.

Setting environment variables

During the installation process, RPM generates C-shell and Bourne-shell files containing environment variables. The files are located in:

```
$SYBASE/SYBASE.sh, $SYBASE/SYBASE.csh
```

If you install multiple versions of Adaptive Server, or if you intend to access different Adaptive Servers on your network, you may want to change the environment variables PATH, DSLISTEN, DSQUERY, TEMP, CLASSPATH, SYBASE, SYBASE_ASE, SYBASE_OCS, and SYBASE_FTS.

Verifying that servers are running

To determine whether the servers you installed are running, do either of the following:

- 1 For each server, enter the following command at the UNIX prompt enter:

```
$SYBASE/$SYBASE_OCS/bin/isql -Usa -P -Sserver_name
```

where *server_name* is the name of the server you want to check. If the server is running, you will see the isql prompt:

```
1>
```

- 2 At the UNIX prompt, enter:

```
$SYBASE/$SYBASE_ASE/install/showserver
```

showserver shows all Adaptive Server-related processes are on the system.

Verifying that you can connect to servers

Use isql or dsedit to perform a quick test. Use isql to connect to servers:

- 1 At the command prompt enter:

```
isql -Usa -Ppassword -Sserver_name
```

where *server_name* is the Adaptive Server, Monitor Server, or Historical Server name.

The command prompt is displayed if the login is successful.

- 2 To display the Adaptive Server version number, enter:

```
1> select @@version  
2> go
```

Adaptive Servers's version number is displayed. The output should show Adaptive Server at version 12.5.

Connecting to Adaptive Server via Sybase Central

Sybase Central provides a graphical user interface where you can perform Adaptive Server administrative tasks. The Sybase Central interface is very similar to the Windows Explorer, with a hierarchical list of all servers in the left pane and a details list of the selected server in the right pane. To select a server, click on it with the mouse in the left pane.

To launch Sybase Central:

- 1 From a UNIX command line, enter:

```
source $SYBASE/SYBASE.csh
```

- 2 Then enter:

```
$SYBASE/sybcent32/scjview
```

- 3 In the left pane, click on the server you want to connect to.
- 4 At the Login window, enter, the System Administrator's user name and password.

You must log in as the System Administrator to perform administrative tasks.

- 5 Click OK.

Warning! The first time you log in to Adaptive Server, use the default “sa” user name and leave the password blank. After you log in for the first time, you should change the System Administrator’s password. See “Setting the System Administrator password” on page 71.

Sybase Central has online help files. To access the online help, right-click on the server icon, then select Online help from the drop-down list. To access online help for Sybase Central, select Help from the menu bar, then select Sybase Central Help.

Setting the System Administrator password

❖ Setting the Sybase system administrator password

A user account called “sa” is created for the Sybase System Administrator when you install the Sybase software. A user logged in as “sa” can use any database on Adaptive Server, including master, with full privileges.

Immediately after a new installation, there is no password on the “sa” account. The initial default value for the password is NULL. In a production environment, the Sybase System Administrator should always use a non-default password.

- The Sybase System Administrator should log in to the new Adaptive Server as “sa” and set a password using `sp_password`:

```
$SYBASE/$SYBASE_OCS/bin/isql -Usa -P -Sserver_name
1> sp_password null, new_password
2> go
```

where `null` is the default password and `new_password` is the password that you are assigning to the “sa” account.

Installing sample databases

This section describes how to install the U.S. English and international language sample databases. For installation instructions specific to each sample database, see the following sections:

- “Running the database scripts” on page 73
- “Installing the interpubs database” on page 74
- “Installing the jpubs database” on page 75

The sample databases contain information about a fictitious business. You can use this information to learn about the Sybase products, without affecting essential data while learning.

Table 5-1 lists the scripts that you can use to install the sample databases.

Table 5-1: Sample database scripts

| Script | Description |
|---------------------|--|
| <i>installpubs2</i> | <p>Installs the pubs2 sample database.</p> <p>This database contains data that represents a publishing operation. Use this database to test your server connections and to learn Transact-SQL. Most of the examples in the Adaptive Server documentation query the pubs2 database.</p> |
| <i>installpubs3</i> | <p>Installs the pubs3 sample database.</p> <p>This updated version of pubs2 uses referential integrity. In addition, its tables are slightly different than the tables used in pubs2. Where noted, the Adaptive Server documentation uses the pubs3 database in its examples.</p> |
| <i>installpix2</i> | <p>Installs the image data that is used with the pubs2 database.</p> <hr/> <p>Note The master device size should be at least 30MB to install the full pubs2 database, including the image data. Be sure to run the <i>installpix2</i> script after you run <i>installpubs2</i>.</p> |
| <i>psdemodb</i> | <p>Installs the Powersoft sample database for PC-clients.</p> <p>For information on installing psdemodb, see <i>Sybase Adaptive Server Configuration Guide</i>.</p> |

Default devices for sample databases

The `$$SYBASE/scripts` directory contains scripts for installing the us_english sample database, foreign language sample databases, and the image data associated with the U.S. English pubs2 sample database.

If you have not used `sp_diskdefault` to change the status of the master device or to specify another default device, the scripts install the sample databases on the master device. This configuration is not recommended because it uses valuable space that is best used for system tables. Each sample database requires 3MB on your database device.

To avoid installing sample databases on the master device, either:

- Use `sp_diskdefault` to specify a default device other than the master device. For information on `sp_diskdefault`, see the *Adaptive Server Reference Manual*.
- Modify each sample database installation script to specify a different device.

❖ Running the database scripts

- 1 Start Adaptive Server.
- 2 Determine the type (raw partition, logical volume, operating system file, and so on) and location of the device where you will be storing the pubs2 and pubs3 databases. You will need to provide this information later.
- 3 Make a copy of the original *installpubs2* and *installpubs3* scripts. Be sure you can access the copies, in case you have problems with the edited scripts.
- 4 Use a text editor to edit the script, if necessary, to specify a default device other than the master device, or use `sp_diskdefault`.

From the Adaptive Server scripts directory (`$SYBASE/$SYBASE_ASE/scripts`), use `isql` to log into Adaptive Server and run the script:

```
isql -Usa -P -Sserver_name -iscript_name
```

where *server_name* represents the destination server for the database and *script_name* is the full path to and file name of the script to run.

For example, to install pubs2 on a server named VIOLIN, enter:

```
isql -Usa -P -SVIOLIN  
-i $SYBASE/$SYBASE_ASE/scripts/installpubs2
```

- 5 To install the image data associated with pubs2 (pubs3 does not use image data) run:

```
isql -Usa -Ppassword -Sservername
```

```
-i $SYBASE/$SYBASE_ASE/scripts/installpix2
```

Note The image data requires a fair amount of space—there are six pictures, two each in the PICT, TIFF, and Sun raster file formats. Run *installpix2* script only if you want to use or test the image datatype. Sybase does not supply any tools for displaying image data. You must use appropriate screen graphics tools to display the images after you have extracted them from the database.

For more information about running these scripts, see *Configuring Adaptive Server Enterprise* at <http://www.sybase.com/support/manuals>.

***interpubs* database**

interpubs is a database similar to *pubs2* that contains French and German data. This data contains 8-bit characters and is available for use at Adaptive Server installations using the ISO 8859-1 (*iso_1*) or Roman8 (for HP-UX) character set. To display the French and German data correctly, you must set up your terminal to display 8-bit characters.

❖ Installing the *interpubs* database

- 1 Be sure *iso_1* or Roman8 is installed as the default character set or as an additional character set.
- 2 Determine the type (raw partition, logical volume, operating system file, and so on) and location of the device where you will be storing the *interpubs* database. You will need to provide this information later.
- 3 Make a copy of the original *installintpubs* script. Be sure you can access this copy, in case you experience problems with the edited script.
- 4 Use a text editor to edit the script, if necessary, to specify a default device other than the master device, or use *sp_diskdefault*.
- 5 Execute the script, using the *-J* flag to ensure that the database is installed with the correct character set:

```
isql -Usa -Ppassword -Sservername -Jiso_1  
-i $SYBASE/$SYBASE_ASE/scripts/iso_1/installintpubs
```

For more information on the *-J* option in *isql*, see the *Utility Guide*.

***jpubs* database**

If you installed the Japanese Language Module with your Adaptive Server, the scripts file contains the `installjpubs` script for installing the `jpubs` database. `jpubs` is a database similar to `pubs2` that contains Japanese data. `installjpubs` uses either the EUC-JIS (`eucjis`) or the Shift-JIS (`sjis`) character set.

❖ **Installing the *jpubs* database**

To install the `jpubs` database:

- 1 Set your terminal to display 8-bit characters.
- 2 Verify that either the EUC-JIS or Shift-JIS character set is installed as Adaptive Server's default character set or as an additional character set.
- 3 Determine the type (raw partition, logical volume, operating system file, and so on) and location of the device where you will be storing the `jpubs` database. You will need to provide this information later.
- 4 Make a copy of the original `installjpubs` script. Be sure that you can access this copy, in case you experience problems with the edited script.
- 5 Use a text editor to edit the script, if necessary, to specify a default device other than the master device, or use `sp_diskdefault`. See "Default devices for sample databases" on page 72.
- 6 Execute the `installjpubs` script, using the `-J` flag to ensure that the database is installed with the correct character set:

```
isql -Usa -Ppassword -Sservername -Jeucjis
-i $SYBASE/$SYBASE_ASE/scripts/eucjis/installjpubs
```

or:

```
isql -Usa -Ppassword -Sservername -Jsjis
-i $SYBASE/$SYBASE_ASE/scripts/sjis/installjpubs
```

For more information on the `-J` option in `isql`, see *Adaptive Server Enterprise Utility Guide*.

Maintaining the sample databases

The sample databases contain a "guest" user that allows access to the database by any authorized Adaptive Server user. The "guest" user has a wide range of privileges, including permissions to select, insert, update, and delete user tables. For more information about the "guest" user and a list of "guest" permissions, see the *System Administration Guide*.

If possible, and if space allows, give each new user a clean copy of the sample databases so that she or he is not confused by other users' changes.

If space is a problem, you can instruct the user to issue the begin transaction command before updating a sample database. After the user has finished updating one of the sample databases, he or she can issue the rollback transaction command to undo the changes.

jConnect 4.5, 5.5, and Java utilities

When you install jConnect 5.5, you receive classes that are compatible with JDBC 2.0, regardless of which JDK (1.1.x or 1.2) version you use. The following table lists the JDK and jConnect versions necessary to develop applications that are compliant with different versions of JDBC.

| If you are developing... | Use JDK version... | Use jConnect version |
|--------------------------------|-----------------------------|----------------------|
| JDBC 1.x applications/servlets | 1.1.x or compatible browser | 4.x |
| JDBC 2.x applications/servlets | 1.2.x or compatible browser | 5.x |

For information on developing applets with jConnect, see *jConnect for JDBC x.x Product Documentation* at <http://www.sybase.com/support/manuals>.

Note The Java VM security manager does not allow a version of jConnect installed in your local file system to create a connection to a server, even if the server is your local system. This can create a problem if you use Netscape to run an applet that uses jConnect.

When Netscape needs to load jConnect to create a connection to a server, it looks for a CLASSPATH defined in its environment to find the jConnect classes. If it finds a CLASSPATH that you have defined to point to your local jConnect installation, it loads jConnect directly from the local file system, rather than from the server, and tries to create a connection. This generates a security error and the connection fails.

Therefore, if you are using Netscape to run jConnect applets, do not set the CLASSPATH to the local jConnect directory.

Using open database connectivity

Some applications do not connect to Adaptive Server directly through the Open Client software but through the Open Database Connectivity (ODBC) driver, or the jConnect driver instead. For example, PowerDynamo connects through the ODBC driver.

ODBC and OLE DB require the Microsoft Data Access Component (MDAC). You can download it directly from Microsoft's Web site at <http://www.microsoft.com/data>

See *Configuring Adaptive Server Enterprise* for your platform for more information.

Installing stored procedures

By default, all Adaptive Server store procedures are installed during the installation process. If you need to re-install the jConnect stored procedures for any reason, see *jConnect for JDBC* documentation set at <http://www.sybase.com/support/manuals>.

Installing character sets

By default, when Adaptive Server and Backup Server are installed on Linux/Intel systems, the installation installs the character set files for iso_1 which supports the Western European languages.

In addition, the installation defaults include sort-order definition files for Western European character sets and U.S. English system message files.

If you want Adaptive Server error messages to display in a language other than U.S. English (us_english), you must install the appropriate language module.

Warning! Make all changes to the default character set and sort order for a new Adaptive Server before creating any user databases or making any changes to the Sybase-supplied databases. Changing the character set and sort order after data or data structures have been added to Adaptive Server can cause incorrect behavior. To change the character set or sort order after you have added data, see the *System Administration Guide*.

❖ Installing new language modules

- 1 Install the language modules into the same directory where you installed Adaptive Server.

Note Each language uses about 2MB of database space per module. If necessary, use the alter database command to increase the size of the master database before adding another language. For more information, see the *System Administration Guide*.

- 2 Use charset to load the default character set and sort order.

Note Use charset only at the time of installation. For alternative methods, see Localization in the *Adaptive Server Configuration Guide*.

- 3 To use charset, the server must be running and you must have System Administrator privileges. Use the *file name* of the sort order:

```
$SYBASE/SYBASE_ASE/bin/charset -Usa -Ppassword  
-Sserver_name sort_order_file character_set
```

- 4 Replace *sort_order_file* with the name of the sort order file. Replace *character_set* with the Sybase name for your character set.
- 5 Use charset to load any additional character sets. See Table 5-2 for charset syntax.

Table 5-2: charset syntax

| | |
|--------|--|
| Syntax | <code>charset [sort_order_file charset.loc] charset_directory</code> |
| Usage | <code>charset [-U username] [-P password] [-S server] [-I interfaces] [-v version] [sort_order_file charset.loc] [charset_directory]</code> |

❖ **Installing character-set definition files**

If you plan to use the Adaptive Server built-in character-set conversions, you must load the character-set definition files for all the character sets on your client platforms. If you are using the Unilib character-set conversions, you do not need to do this.

- 1 Using `isql`, log in to your server as “sa” and select the master database.

```
1> use master
2> go
```

- 2 Use the *ID* of the sort order to configure your server for the new character set and sort order.

```
1> sp_configure "default sort_order_id",
2> sort_order_id, "character_set"
3> go
```

- 3 Replace *sort_order_id* with the ID for your sort order. Replace *character_set* with the Sybase name for your character set.

- 4 Shut down the server to start the reconfiguration process.

Use the normal process on your UNIX system to restart the server, usually by invoking one of the *RUN_XXX* scripts from *\$SYBASE/\$SYBASE_ASE/install*.

- 5 The server starts, rebuilds all the system indexes, then shuts down. Restart a second time to bring the server up in a stable state.

To change the default character set and sort order of Adaptive Server after the initial installation, see the *System Administration Guide*.

Installing online help for Transact-SQL syntax

This section provides instructions for installing online help for Transact-SQL syntax.

Online syntax help: *sp_syntax*

In the *scripts* directory there are scripts for installing the syntax help database, *sybsyntax*. You can retrieve this data with the *sp_syntax* system procedure. The scripts are located in:

```
$$SYBASE/$$SYBASE_ASE/scripts/ins_syn_sql
```

For more information on *sp_syntax*, see the *Reference Manual*.

This script includes syntax information for Transact-SQL, the system procedures, and the Sybase utilities. When you execute this script, you install the SQL portion of the *sybsyntax* database.

When you first execute this script, it creates the *sybsyntax* database and the necessary tables and indexes. If you execute the script a second time, the previously installed rows of information are deleted from the table in the database and then reinstalled.

Default device for the *sybsyntax* database

By default, the *sybsyntax* installation scripts install the *sybsyntax* database on the device that is designated as the default database device.

Table 5-3 lists the minimum *sybsyntax* database requirements.

Table 5-3: *sybsyntax* requirements

| Logical page size | 2K | 4K | 8K | 16K |
|-----------------------|-----|-----|------------------------|------------------------|
| Minimum database size | 3MB | 5MB | Same size as the model | Same size as the model |

If you have not used *sp_diskdefault* to change the status of the master device (which is installed as the default disk) or to specify another default device, the scripts install *sybsyntax* on the master device. This configuration is not recommended because *sybsyntax* uses valuable space, which is best left available for future expansion of the master database.

To avoid installing *sybsyntax* on the master device, do one of the following:

- Use `sp_diskdefault` to specify a default device other than the master device. For information about `sp_diskdefault`, see the *Reference Manual*.
- Modify each `sybsyntax` installation script that you plan to execute to specify a different device, as explained in the following section.

Installing *sybsyntax*

For each `sybsyntax` installation script you want to execute:

- 1 Determine the type (raw partition, logical volume, operating system file, and so on) and location of the device where you plan to store the `sybsyntax` database. You will need to provide this information later.
 - 2 Make a copy of the original script. Be sure you can access this copy, in case you experience problems with the edited script.
 - 3 Use a text editor to edit the script, if necessary, to change the default device from the master device to the device created in step 1. For information on the default device, see “Default device for the `sybsyntax` database” on page 80.
- Comment out the following section, which specifies the default device:

```
/* create the database, if it does not exist */
if not exists (select name from sysdatabases
where name = "sybsyntax")
begin
    /* create the sybsyntax table if it doesn't exist */
    /* is the space left on the default database
    devices > size of model? */
    if (select sum (high-low +1) from sysdevices where status
    & 1 = 1) - (select sum(size) from sysusages, sysdevices
        where vstart >= sysdevices.low
        and vstart <= sysdevices.high
        and sysdevices.status &1 = 1) >
        (select sum(sysusages.size) from sysusages
        where dbid = 3)
    begin
        create database sybsyntax
    end
else
    begin
        print "There is not enough room on the default
        devices to create the sybsyntax database."
```

```
        return
    end
end
```

- After you have commented out this entire section, add a line like this to the script:

```
create database sybsyntax on device_name
```

where *device_name* is the name of the device where you want to install sybsyntax.

- 4 Execute the script with a command like the following:

```
isql -Usa -Ppassword -Sservername
$SYBASE/$SYBASE_ASE/scripts/ins_syn_sql
```

where *sa* is the user ID of the System Administrator, *password* is the System Administrator's password, and *servername* is the Adaptive Server where you plan to install the database.

If you have set the DSQUERY environment variable to the *servername*, you can replace the server name with DSQUERY. For example:

```
isql -Usa -Ppassword -S$DSQUERY
$SYBASE/$SYBASE_ASE/scripts/ins_syn_sql
```

- 5 To ensure that you have installed the sybsyntax database and that it is working correctly, use isql to log in to the server on which you installed the database, and execute sp_syntax. For example:

```
isql -Usa -Ppassword -Sservername

1> sp_syntax "select"
2> go
```

Adaptive Server displays a list of commands that contain the word or word fragment "select."

Configuring XP Server after initial installation

If you install Adaptive Server and XP Server in the same build session, `srvbuild` automatically adds information about the XP Server to the `syssservers` table of Adaptive Server. If you install XP Server in a different build session than the Adaptive Server installation, you are prompted during the XP Server installation process to supply the name of the related Adaptive Server and the System Administrator's name and password. This information is required by Adaptive Server to start XP Server.

If you do not supply this information during the installation process, XP Server cannot run. When you execute an extended stored procedure (ESP), you see an error message similar to the following:

```
Msg 11018, Level 16, State 1:  
Procedure 'xp_cmdshell', Line 2  
XP Server must be up for ESP to execute.
```

To add this required information to the `syssservers` table manually, enter:

```
sp_addserver servername_XP, NULL, servername_XP
```

where *servername* is the name of the Adaptive Server.

Installing Sybase PC-Client Products

Your UNIX-based Adaptive Server also includes PC-client products that you install on a PC and connect, via the interfaces file, with your Adaptive Server.

This chapter describes how to install client products on Windows 98 and Windows NT 4.0.

| Topic | Page |
|---|------|
| Overview | 85 |
| Installing PC-Client products | 86 |
| Installing a Java Runtime Environment | 89 |
| Installing the Java Development Kit | 89 |
| Installing InfoMaker | 90 |
| Installing PowerDynamo | 91 |
| Configuring network connections for client products | 91 |

Overview

Client plug-ins enable clients, such as Sybase Central and Open Client, to access Adaptive Server.

Sybase Central plug-In

The Sybase Central plug-in allows you to connect to of Adaptive Server via client connections. You can use Sybase Central to connect to remote servers, local servers (installed on the same machine), and multiple servers. See “Configuring network connections for client products” on page 91. Before you use the plug-in, restart your computer to implement changes made to the PATH and CLASSPATH environment variables during installation.

Open Client runtime

The Sybase PC-client CD includes Open Client software for Windows NT and Windows 98. Several products, such as the Adaptive Server plug-in and SQL Advantage require Open Client runtime. Monitor Client Library requires Open Client/C with the SDK option.

Open database connectivity

Some applications do not connect to Adaptive Server directly through the Open Client software but through the Open Database Connectivity (ODBC) driver, or the jConnect driver instead. For example, PowerDynamo connects through the ODBC driver.

See the *Adaptive Server Configuration Guide* for more information.

ODBC and OLE DB require the Microsoft Data Access Component (MDAC). You can download it directly from Microsoft's Web site at <http://www.microsoft.com/data>.

Installing PC-Client products

Before starting the installation process for PC-client products, be sure that you have at least an extra 2.4MB of disk space, a `\temp` directory and TEMP environment variable.

The installation program uses the extra space and directory to write files temporarily during the installation. The installation program frees this space after installation is complete.

Special instructions for Windows 98 platforms:

Environment Space is an area of memory set aside for each process in which environment variables can be defined. By default, there is a fixed amount of memory set aside to hold environment variables. This space contains the name of the variable (for example, TEMP) and its value (for example, `C:\Windows\temp`). Windows 98 by default provides for 256 bytes of environment space. Each environment variable defined takes up a specific amount of space depending upon the number of characters in the name and value. If you define an environment variable and there is not enough space to hold the variable and its definition, Windows will return an error indicating that you are "out of environment space."

If you are installing PC-client products on the Windows 98 platform, follow these steps to increase your environment space.

- Bring up a DOS window
- Right click on the MSDOS icon on the left end of the top window bar to bring up the properties window
- Select Memory
- Select Initial Environment and enter: 4096
- Click OK

To install client products on Windows 98 and Windows NT:

- 1 Verify that your computer has sufficient RAM for each product. See Chapter 2, “Installation Requirements”.
- 2 Log in using an account with Windows NT administrator privileges.
- 3 Close any open applications or utilities to free memory and system resources.
- 4 Insert the PC-client CD in to the drive. The Studio Installer should start automatically.

Alternatively, you can start the Studio Installer from the Windows Start menu, select Start | Run, and enter:

`X:\setup.exe`

where *X* is your CD-ROM drive.

- 5 Click OK.
- 6 When the Studio Installer starts, it prompts you for an installation directory. Click Next.
- 7 Select the type of installation to be performed.
 - *Standard Install* – installs the default components a user needs.
 - *Full Install* – installs every component on the CD.
 - *Customized Install* – allows you to select which components to install. Certain components will automatically be installed if they are required to run other selected components.

Note If you select Customized Install, the next window is the Component Selection screen, which allows you to specify which components to install.

Components that would be installed in a standard installation appear with a check in the check box to the left of the product name. You may select or deselect components from this list. Components with subcomponents have a More... button enabled. Clicking this button allows you to select or deselect subcomponents.

- 8 Click Next.

The Summary screen displays every component to be installed by the Studio Installer, the required disk space, and the available disk space.

If the target directory does not have enough free space, the available space appears in red. Clicking Next without sufficient hard disk space results in an error and stops the installation.

- 9 Click Next.

Note Optionally, you can select Save on the Summary screen to save all the installation information into a *cmdfile* to proceed with the installation in a noninteractive, silent install.

See Appendix A, “Alternative Installation Methods”.

- 10 If the target directory does not exist, Studio Installer prompts: “OK to create directory?” Click Yes.

The Studio Installer installs the components and displays a progress indicator.

Warning! If you are prompted to overwrite any DLLs, select Yes only if the version of the new DLLs is later (the version number is larger) than the one it is attempting to overwrite.

The Studio Installer generates an error message and stops the installation if you do not select any components for installation.

Click Back to select components, or Cancel to cancel the installation procedure.

- 11 If any components selected to be installed require a license, the Sybase Software Asset Management (SySAM) screen displays and prompts you for certificate information at this time.

- 12 When prompted to reboot the computer, leave the CD in the CD-ROM drive and select Yes. Restarting the PC updates the environment variables and registry keys.

Note If you are installing client products on the Windows 98 platform, you must manually restart the computer before proceeding with the configuration of client products. The Studio Installer does not automatically reboot the Windows 98 operating system.

To configure client network connections to Adaptive Server, see “Configuring network connections for client products” on page 91.

Note If you installed jConnect 4.5 or 5.5 from the PC-Client CD, see “Installing Stored Procedures for jConnect” in Chapter 5, “Post-Installation Tasks”.

Installing a Java Runtime Environment

Before you can use the Sybase Central monitoring features, you need to install the Java Runtime Environment (JRE):

To install the JRE:

- 1 Go to: %SYBASE%\ASEP_win32.
- 2 Double-click the file *jdk1_1_8-win.exe* and follow the installation instructions.
- 3 Restart your computer after you install the JRE.

Installing the Java Development Kit

Optionally, you can install the entire Java Development Kit (JDK) instead of the JRE. The JDK requires more disk space and includes files and features that support development of Java applets and applications. Sybase does not provide the JDK (see step 1 in the following instructions).

To install the JDK:

- 1 Download the JDK for your operating system from the Java Software Web site at <http://java.sun.com/products/jdk/1.1>.
- 2 Double-click the file you downloaded.
- 3 Follow the installation instructions.
- 4 Add the following path to your CLASSPATH environment variable:

```
x:\install_path\lib\CLASSES.ZIP
```

where *x:\install_path* is the location (drive and directory) where you installed the JDK; for example, *c:\javasoft*.
- 5 Add *x:\install_path\bin* to your PATH environment variable, where *x:\install_path* is the location (drive and directory) in which you installed the JDK; for example, *c:\javasoft*.

Installing InfoMaker

InfoMaker is distributed as part of the PC-Client package. If you want to use the InfoMaker Demo Database, you must download and install Adaptive Server Anywhere.

To install InfoMaker:

- 1 Use the Studio Installer to install the software from the CD as described in “Installing PC-Client products” on page 86.
- 2 Run the InfoMaker setup utility located in *%SYBASE%\INFO-7_0_2\install*.

From a DOS prompt, enter:

```
setup.exe "-EAS3"
```

or, from the Windows task bar, select Start | Run, and enter the full path to the executable, including the *-EAS3* argument.

```
X:\sybase\INFO-7_0_2\install\setup.exe "-EAS3"
```

where *X:\sybase* is the installation directory for the PC-client components.

- 3 When prompted for an installation directory, accept the default.
- 4 When prompted to reboot the computer, leave the CD in the CD-ROM drive and select Yes. Restarting the PC updates environment variables and registry keys.

- 5 See the Powersoft documentation for instructions on using InfoMaker.

Note If you cannot launch the Demo Database through your shortcut, please change the properties in it by setting the path to your SQL Anywhere win32 directory.

Installing PowerDynamo

PowerDynamo is distributed as part of the PC-Client package.

To install PowerDynamo:

- 1 Use the Studio Installer to install the software from the CD as described in “Installing PC-Client products” on page 86.
- 2 Run the PowerDynamo setup utility located in `X:\sybase\PowerDynamo\install\setup.exe` where `X:\sybase` is the installation directory for the PC-client components.
- 3 When prompted for an installation directory, accept the default.
- 4 When prompted to reboot the computer, leave the CD in the CD-ROM drive and select Yes. Restarting the PC updates environment variables and registry keys.

Configuring network connections for client products

Adaptive Server communicates with other Adaptive Servers, Open Server applications (such as Backup Server), and client software on your network. Clients can talk to one or more servers, and servers can communicate with other servers by remote procedure calls.

For Sybase products to interact with one another, each product needs to know where the others reside on the network. This information is stored in an interfaces file (`sql.ini` on Windows NT), or in the `libtcl.cfg` file.

Adding a server entry to the *libtcl.cfg* file

You use the *libtcl.cfg* files to specify an LDAP (Lightweight Directory Access Protocol) server name, port number, DIT base, user name, and password to connection to an LDAP server.

The default *libtcl.cfg* file is located in:

%SYBASE%\%SYBASE_OCS%\ini

In its simplest form, the *libtcl.cfg* file is in this format:

```
[DIRECTORY]
ldap=libldlap.dll ldapurl
```

where the *ldapurl* is defined as:

```
ldap://host:port/ditbase
```

When an LDAP server is specified in the *libtcl.cfg* file, the server information is accessible only from the LDAP server. Adaptive Server ignores the interfaces file. Open Client/Open Server applications that use the -l option at start-up override the *libtcl.cfg* file and use the interfaces file.

To use a directory service, you must:

- 1 Add the location of the LDAP libraries to the path environment variable for your platform.
- 2 Configure the *libtcl.cfg* file to use directory services.

Use any standard ASCII text editor to:

- Remove the semicolon (;) comment markers from the beginning of the LDAP URL lines in the *libtcl.cfg* file under the *[DIRECTORY]* entry.
- Add the LDAP URL under the *[DIRECTORY]* entry. See *Configuring Adaptive Server Enterprise* for supported LDAP URL values.

Warning! The LDAP URL must be on a single line.

```
ldap=libldlap.dll ldap://host:port/ditbase??scope??
bindname=username password
```

For example:

```
[DIRECTORY]
ldap=libldlap.dll ldap://huey:11389/dc=sybase,dc=com??
one??bindname=cn=Manager,dc=sybase,dc=com secret
```

- 3 Verify that the appropriate environment variable points to the required third-party libraries. The Netscape LDAP SDK libraries are located in *%SYBASE%\%SYBASE_OCS\lib3p*. The Windows PATH environment variable must point to this directory.
- 4 Once you have edited the *libtcl.cfg* file, use *dsedit* to add a server to the directory service:
 - a From the Windows task bar, select Start | Programs | Sybase | *dsedit*.
 - b Select LDAP from the list of servers, and click OK.
 - c Click Add New Server Entry.
 - d Enter:
 - The server name – this is required.
 - Security mechanism – optional. A list of security mechanism OIDs are located in *%SYBASE%\%SYBASE_OCS%\ini\objectid.dat*.
 - HA server name – optional. This is the name of the high-availability failover server, if you have one.
 - e Click Add New Network Transport.
 - Select the transport type from the drop-down list.
 - Enter the host name.
 - Enter the port number.
 - f Click OK two times to exit the *dsedit* utility.

For more information, *Configuring Adaptive Server Enterprise* for your platform.

Adding a server to the *sql.ini* file

During installation, the *srvbuild* process adds entries to the *interfaces* file for your new Adaptive Server, Backup Server, Monitor Server, and XP Server.

To access a Sybase server through the *interfaces* file, each server must be included in the *interfaces* file on the client computer.

To add a server entry to the *sql.ini* file on your PC-client computer:

- 1 Using the Microsoft Windows task bar, click Start | Programs | Sybase | *dsedit*.

- 2 Click OK on the first screen to open the InterfacesDriver screen.
- 3 Select ServerObject from the menu, and click Add.
- 4 In the Input Server Name box, enter the name of the server for which you are creating an entry.
Click OK.
- 5 In the Attributes column, double-click on the server address row you just added.
- 6 In the ProtocolNetwork Address, click Add.
- 7 From the drop-down list in the Protocol box, select TCP, NAMEPIPE, SPX, or VINESIP.
- 8 In the Network Address box, you can enter either the server name or the IP address, along with the server port number. For example:

```
machine_name, 4100
```

or,

```
machine_name 4100
```

where *machine_name* is the name of the computer, and *4100* is the port number the server is using to “listen” for clients.

Note The Windows platform now accepts TCP-style connection information.

```
[SERVER]
MASTER=NLWNSCK, heuy, 2222
QUERY=NLWNSCK, heuy, 2222
```

or, the preferred format:

```
[SERVER]
MASTER=TCP heuy 2222
QUERY=TCP heuy 2222
```

or

```
[SERVER]
MASTER=TCP, heuy, 2222
QUERY=TCP, heuy, 2222
```

The preferred format is to use TCP and a space between the *host_name* and *port_number*, because it is supported across all platforms. You can edit the *sql.ini* with any standard ASCII text editor.

- 9 To find the IP address for a machine, use `ypmatch`. For example:

```
ypmatch 'hostname' hosts
```

where *hostname* is the name of the machine.

To find a server's listener server port number:

- If the server is on another PC, check the `%SYBASE%\ini\sql.ini` file
- If the server is on a UNIX machine, check the `interfaces` file in `$SYBASE`

Note To connect to a UNIX server, the entries you add to `sql.ini` must match the entries in the `interfaces` file on the UNIX system.

For instructions on using `dsedit` or `dscp` to modify `interfaces` file entries or to create new `interfaces` file entries for existing servers, see the *Utility Guide*.

For additional information on using `dsedit` on the client computer, see the *Open Client/Server Configuration Guide for Desktop Platforms* and the *Open Client/Server Supplement*.

For specific information about entries for connections to Monitor Server and Historical Server, see:

- *Adaptive Server Enterprise Monitor Server User's Guide*
- *Adaptive Server Enterprise Monitor Historical Server User's Guide*

Enabling TCP connections

To connect through TCP instead of Named Pipes and Winsock network protocols, you must enable the TCP connections.

To enable TCP connections:

- 1 From the Windows task bar, click Start | Programs | Sybase | OC OS Config.
- 2 Click the NetLibrary tab.
- 3 Select Windows NT or 98 as the platform and TCP as the protocol.
- 4 Click OK.

This process maps NLWNSCK to TCP in your Sybase network configuration file (`libtcl.cfg`).

You have completed installing Adaptive Server client plug-ins. To find out more information about setting up network communications, or to add optional functionality to Adaptive Server, see *Configuring Adaptive Server Enterprise for Windows NT*.

After you have installed the Sybase products on your system, see their accompanying documentation for configuration and administration issues.

Testing the Sybase Central installation

After you install Adaptive Server, Sybase Central, and the Java Runtime Environment, it is a good idea to test the installation and network connections.

To test installation and network connections:

- 1 Select Start | Programs | Sybase Central Java Edition from the Windows task bar to start Sybase Central.
- 2 Select Tools | Connect from the Sybase Central menu bar to activate a Sybase Adaptive Server login screen.
- 3 Log in using the default user ID, “sa”, without a password. If you changed the password for Adaptive Server according to the post-installation instructions, you need to use the new password.
- 4 From the drop-down menu, select the Adaptive Server that you want to connect to.

If Sybase Central is installed on the same computer as the server, click the icon for the server.

If prompted to start Adaptive Server, click OK. When the traffic light icon representing the server is green, the server is running.

Warning! Do not start XP Server from Sybase Central. Adaptive Server starts XP Server automatically when Adaptive Server receives the first request for an extended stored procedure.

- 5 To disconnect from a server, choose Disconnect from the Tools menu.

Note If a server is running on the same computer as Sybase Central, disconnecting does not stop it.

- 6 Exit Sybase Central.

If Adaptive Server fails any of these tests, see Chapter 9, “Troubleshooting” Before retrying the installation, follow the instructions in Chapter 5, “Post-Installation Tasks.”

Starting and Stopping Servers

This chapter describes how to start and stop Adaptive Server, Backup Server, Monitor Server, and XP Server.

| Topic | Page |
|---|------|
| Overview | 99 |
| Starting servers | 100 |
| Starting servers when the operating system restarts | 101 |
| Stopping servers | 102 |

Overview

After you create the Adaptive Server or upgrade Adaptive Server using `srvbuild` or `sqlupgrade`, Adaptive Server, Backup Server, and Monitor Server are running. For information about running the Full-Text Search server, see the *Full-Text Search Specialty Data Store User's Guide*.

XP Server is not started by the installation process. XP Server is started only when any XP command is issued through `isql`.

The methods described here are used to start Adaptive Server, Backup Server, and Monitor Server after a shutdown for database maintenance, because of an operating system crash, or for some other reason.

You can start Adaptive Server, Backup Server, and Monitor Server on UNIX either by:

- Entering commands from the UNIX command line, or
- Configuring your operating system so that the servers start automatically when the operating system starts

Starting servers

You can start Adaptive Server, Backup Server, or Monitor Server from the UNIX command line by using runserver files and startserver commands. You can move the Adaptive Server shared memory files by adding a flag to the runserver file.

Using the Runserver file

Each time a new Adaptive Server, Backup Server, or Monitor Server is installed, the `srvbuild` program creates a runserver file that contains the information required to restart that server. Runserver files are created in the `$$SYBASE/$SYBASE_ASE/install` directory. For Adaptive Server, the runserver file is named `RUN_servername`, where *servername* is the name of the server.

For example, the runserver file name for an Adaptive Server named TEST is `RUN_TEST`. If the runserver file for your current Adaptive Server is named `RUNSERVER`, you must change the name to `RUN_TEST` during the upgrade process.

For Backup Server and Monitor Server, the runserver files are named `RUN_servername_back` and `RUN_servername_mon`, respectively, where *servername* is the name of the server.

Note Do not delete the runserver file that is created in `$$SYBASE/$SYBASE_ASE/install`. This file is used to restart servers when you customize your installation. If you need the runserver file in another location, make a copy of the original runserver file in the new location.

Using the *startserver* command

To start a server from the command line enter:

```
$$SYBASE/$SYBASE_ASE/install/startserver  
[ -f runserverfile ]
```

where `$$SYBASE/$SYBASE_ASE/install/startserver` is the full path name of the startserver utility, and *runserverfile* is the full path name of the runserver file (usually `$$SYBASE/$SYBASE_ASE/install/RUN_servername`).

If your Adaptive Server is named SYBASE, the runserver file specification is optional.

You must have read and write permissions on the master device for an Adaptive Server to start it with startserver. See *Adaptive Server Enterprise Utility Guide* for more information about startserver.

Although startserver runs Adaptive Server in the background, it does not return to the operating system prompt after restarting. After Adaptive Server has been restarted, press Return to return to the operating system prompt.

Using the *monserver* command

Use the monserver command to start Monitor Server from the UNIX command line. For details, see the *Adaptive Server Enterprise Monitor Server User's Guide*.

Starting servers when the operating system restarts

This section describes how to configure your operating system for automatic restart of Adaptive Server and Backup Server when the system restarts.

On production systems, Adaptive Server, Backup Server, and Monitor Server should restart automatically when the UNIX operating system starts. This is accomplished by making an entry for the server in the operating system start-up script.

Monitor Server must be listed *after* Adaptive Server in the start-up script. For more information on starting Monitor Server during the operating system restart, see the *Sybase Adaptive Server Enterprise Monitor Server User's Guide*.

The `$SYBASE/install/sybininstall.sh` script sets up a correct shell environment and installs a system run-level script in a Sybase user account. The script allows the automatic re-start and automatic shutdown of configured Sybase Server products such as Adaptive Server Enterprise, Backup Server, and Open Server.

To run *sybininstall.sh*:

- 1 Log in as root.

- 2 Run the sybinstall.sh script to perform the post-installation. Enter:

```
# $SYBASE/install/sybinstall.sh
```

- 3 Create the correct runtime links for shutdown and boot. Enter:

```
# /sbin/chkconfig sybase on
```

Stopping servers

Only the System Administrator has permission to issue a shutdown command. Using a shutdown command minimizes the amount of work that automatic recovery needs to do when the servers are restarted.

The preferred method of stopping Adaptive Server or Backup Server is to use the Transact-SQL shutdown command. For Monitor Server, use the sms_shutdown command.

Stopping Adaptive Server

To shut down Adaptive Server:

- 1 Use isql to log in to an Adaptive Server account with System Administrator privileges:

```
isql -Usa -Ppassword -Sserver_name
```

- 2 Enter the following command to shut down the server:

```
1> shutdown
2> go
```

The default for the shutdown command is with wait. The with wait option allows Adaptive Server to finish executing SQL statements or procedures, perform a checkpoint in each database, disable new logins, and perform other shutdown tasks.

Issuing the shutdown command prints a message like this to the *stderr* file:

```
Server SHUTDOWN by request.The SQL Server is terminating
this process.
CT-LIBRARY error:
```

This is normal behavior. If the message indicates that Adaptive Server is waiting for processes to complete, and you need to stop Adaptive Server immediately, you can use the shutdown with nowait command. shutdown with nowait does not wait for currently executing statements to finish and does not perform checkpoints in every database.

Stopping Backup Server

To shut down a Backup Server:

- 1 Use isql to log in to a server with System Administrator privileges:

```
isql -Usa -Ppassword -Sserver_name
```

- 2 Enter the following command to shut down the specified Backup Server:

```
1> shutdown SYB_BACKUP
2> go
```

After you shut down a Backup Server, you must wait at least 30 seconds before restarting it.

Issuing the shutdown command prints a message similar to the following to the *stderr* file:

```
Backup Server: 3.48.1.1: The Backup Server will go down
immediately.
Terminating sessions.
```

This is normal behavior. If a message indicates that Adaptive Server or Backup Server is waiting for processes to complete, and you need to stop Adaptive Server or Backup Server immediately, you can use the shutdown with nowait command. shutdown with nowait does not wait for currently executing statements to finish and does not perform checkpoints in every database.

Using shutdown with nowait for Backup Server can cause inconsistent or incomplete dumps and loads. Use this command only when necessary.

For more information on the shutdown command, see the *Reference Manual*.

Stopping Monitor Server

You must shut down Monitor Server before restarting Adaptive Server. If Adaptive Server stops, make sure that Monitor Server is shut down before you restart Adaptive Server. Monitor Server must be shut down to release resources. Otherwise, Adaptive Server may not be able to allocate enough resources to restart.

If the Monitor Server heartbeat feature is in effect, Monitor Server automatically detects the termination of Adaptive Server within the specified period and shuts itself down. Therefore, before attempting to restart Adaptive Server after a shutdown, either wait for the automatic shutdown of Monitor Server to occur, or explicitly stop Monitor Server.

For more information on stopping Monitor Server, see the *Adaptive Server Enterprise Monitor Server User's Guide*.

Using the *kill* command

Warning! Use the kill command in Adaptive Server and Backup Server only as a last resort.

When possible, use the Transact-SQL shutdown or shutdown with nowait command. Do not use kill with the -9 flag because it exits the server without running a checkpoint to ensure that all database changes are written to the database device. Adaptive Server may also exit without removing associated shared memory files and network handlers.

Because Adaptive Server and Backup Server are background processes, they can be killed from the operating system by their owner or by “root” with the UNIX kill command. The syntax is:

```
kill pid
```

where *pid* is the process identification of any *dataserver* or *backupserver* process, as determined by the *showserver* command. Killing one engine for a particular Adaptive Server kills all engines for that server.

If more than one Adaptive Server is running on the same system, you must be careful that the engine you kill is associated with the correct Adaptive Server. If your Adaptive Server is configured to use multiple engines (CPUs), each engine has an associated operating system process. The correct way to kill a multi-engine server is to specify the process ID for engine 0.

This `showserver` output shows the processes for a four-engine server:

```
showserver
```

| UID | PID | PPID | C | STIME | TTY | TIME | COMD |
|-------|------|------|----|----------|-------|--------|-----------------------|
| jorge | 3320 | 1 | 80 | 10:31:40 | pts/4 | 302:15 | dataserver -dteamster |
| jorge | 3321 | 3320 | 80 | 10:31:45 | pts/4 | 324:47 | dataserver -ONLINE:1 |
| jorge | 3322 | 3320 | 80 | 10:31:45 | pts/4 | 326:02 | dataserver -ONLINE:2 |
| jorge | 3323 | 3320 | 80 | 10:31:45 | pts/4 | 328:56 | dataserver -ONLINE:3 |

This example shows four running `dataserver` processes with operating system process identifications (PID) 3320, 3321, 3322, and 3323. (`dataserver` is the executable form of the Adaptive Server program.)

Child engine processes for the `dataserver` have the `-ONLINE:` argument.

Each child engine has a parent process identification (PPID) that is equal to the process identification (PID) of the parent. In the example above, the PID of the parent server is 3320. The other three engines spawned by the parent process have the same PPID.

If the PPIDs appear to be unrelated, and there is more than one `dataserver` process, then more than one Adaptive Server is running on the system.

Shutdown and shared memory files

When Adaptive Server is shut down in a normal manner, the shared memory files are automatically removed. If Adaptive Server crashes or is stopped with the `kill -9` command, these files are not deleted. You need read and write permissions on these files to restart Adaptive Server after a crash or a `kill -9` command, as Adaptive Server must be able to overwrite the previously created shared memory files.

Please consult the Linux `ipcs(1)` options for stale shared memory segment removal.

```
# ipcs -m
# ipcrm shm <shmid>
```


Upgrading Sybase Servers

This chapter describes how to upgrade approved versions of Adaptive Server, SQL Server, Backup Server, and Monitor Server to version 12.5.

| Topic | Page |
|--|------|
| Overview | 107 |
| Pre-upgrade tasks | 110 |
| Preparing to upgrade servers with replicated databases | 124 |
| Upgrading to Adaptive Server 12.5 | 130 |
| Post-upgrade tasks | 135 |
| Upgrading Backup Server, Monitor Server, and XP Server | 141 |
| Upgrading compiled objects with dbcc upgrade_object | 142 |
| Recovering from a failed upgrade | 149 |

Note If you are migrating applications from one version of jConnect to another, see Chapter 5, “Migrating jConnect Applications,” in the *jConnect for JDBC Programmer’s Reference* for instructions.

Overview

You can upgrade to Adaptive Server 12.5 from 11.9.2. If you have an earlier version of Adaptive Server, Sybase recommends that you upgrade to 11.9.2 before upgrading to 12.5.

For a server installation older than version 11.5.x, Sybase recommends that you upgrade to version 11.9.2, then upgrade to version 12.5.

Only upgrades from 2K pages to 2K pages are supported. Changing the server schema from a 2K page to *n*K page size is a database migration, not an upgrade.

Upgrading Adaptive Server consists of four processes:

- Installing the new server into its own installation directory.

Note You must have both the old server and the new server to perform an upgrade.

- Performing the pre-upgrade checks on the old server using the `preupgrade` utility, from the new server installation.
- If necessary, fixing any problems that pre-upgrade process reports.
- Running the upgrade utility from the new server against the databases to update the underlying schema so that their structures are correct for the new server.

The `preupgrade` and `upgrade` utilities are internally called by the `sqlupgrade` utility.

Each new version of Adaptive Server contains different features that introduce new parameters, commands, reserved words, and so on. For this reason, the new Adaptive Server to which you are upgrading is responsible for preparing the old server for the upgrade.

The new server provides a utility, `sqlupgrade`, that runs various checks, such as reserved word checks, to determine how much space you must add to the old server to successfully upgrade the old server to the new.

As part of the pre-upgrade tasks, `sqlupgrade` scans all databases and catalogs and determines how much free space is required for each to upgrade successfully. Essentially, `sqlupgrade` searches for the largest catalog, then calculates the required free space by doubling the size of the largest catalog, and adding approximately 10 percent for logging the upgrade changes for each catalog.

During the pre-upgrade process, `sqlupgrade` returns informational messages as it checks the old server. You must fix all reported problems, and run `sqlupgrade` cleanly before beginning the upgrade process. Once the old server is eligible for upgrade, `sqlupgrade` shuts down the old server, starts the new server against the existing databases, and begins the upgrade process.

System catalog changes during upgrade

Adaptive Server 12.5 introduces support for wider columns, more columns, a larger number of user logins, and multiple logical page sizes. To support these extended limits, there have been several changes to existing system catalogs. See *What's New in Sybase Adaptive Server 12.5* for a complete list of catalogs that are affected by the relaxed server limits.

Warning! As system catalogs are copied during the upgrade process, the size of these catalogs might cause the upgrade step to take a long time. Do not attempt to abort the process using Ctrl-C, as this can cause unknown corruptions in the catalogs.

Catalog changes that might affect existing applications

Any user-written stored procedures or applications that query system catalogs to obtain information must be changed to use the new datatypes for various columns. For instance, existing stored procedures and tools to regenerate table schema that look up information in syscolumns, will continue to work immediately following the upgrade. However, if new tables are created using the new limits, existing stored procedures are unable to retrieve the information from various catalogs.

For example, if a new table is created with 300 columns, and an existing stored procedure declares a tinyint datatype to retrieve the column ID from syscolumns, this procedure cannot return the right information for this table because the table has more than 255 columns.

Similarly, other tools and procedures that access catalogs such as sysusers, syslogins, sysprotects, sysconstraints, and so on, must be updated to reflect the new column definitions. In general, all local variables in user applications and procedures need to be re-defined to match the datatype of the columns in system catalogs that are modified during the upgrade.

All Sybase stored procedures have been upgraded to reflect this change.

Note Before you begin the upgrade process, install Adaptive Server 12.5 onto your system. See Chapter 4, “Installing Sybase Servers.”

Pre-upgrade tasks

Before beginning any upgrade tasks, read the “Special Upgrade Instructions” section in the release bulletin for Adaptive Server.

To ensure a successful upgrade, review the following pre-upgrade tasks and perform them as necessary. Depending on your current Adaptive Server configuration, it may not be necessary to perform all pre-upgrade tasks.

Pre-upgrade tasks include:

- “Check system and upgrade requirements” on page 110.
- “Check RUN_server file location” on page 111.
- “Test current applications and stored procedures” on page 111.
- “Reserved words” on page 112.
- “Verify that users are logged off” on page 114.
- “Check database integrity” on page 114.
- “Back up databases” on page 115.
- “Dump transaction logs” on page 115.
- “Preparing database and devices” on page 116.
- “Turn off database options” on page 123.

Check system and upgrade requirements

Note The max memory parameter must be at least the size indicated in “Pre-installation tasks” on page 42.

To verify that your system environment is configured correctly:

- 1 Verify that the computer on which you plan to upgrade the Sybase products meets the requirements described in Chapter 2, “Installation Requirements”.
- 2 Determine whether your server supports an upgrade to Adaptive Server 12.5.

To determine the release level of your current Adaptive Server (or SQL Server), enter the following command after login in to your server using isql:

```
1> select @@version
2> go
```

- 3 Set the environment variables to the values described in Chapter 4, “Installing Sybase Servers”.
- 4 Verify that you unloaded Adaptive Server 12.5 in a *different* directory than your previous Adaptive Server installation.

If you overwrote your previous server installation:

- a Restore it from your most recent backup.
 - b Reinstall the product files for Adaptive Server in a different directory.
 - c Continue with the upgrade.
- 5 Verify that your operating system is at the proper release level and has all operating system patches needed for Adaptive Server. See the *Adaptive Server Enterprise Release Bulletin* for your platform.

Check RUN_server file location

Verify the name and location of the *runserver* file. Be sure the *runserver* file for your current server is located in the old `$SYBASE/install` directory and is still named *RUN_server_name*, where *server_name* is the name of the current server. The *server_name* must appear exactly as it appears in the interfaces file. The *runserver* file for a server named SYBASE is named *RUN_SYBASE*. If the *runserver* file for your current Adaptive Server is named, you must change the name during the upgrade process.

Test current applications and stored procedures

If you have critical applications and stored procedures, perform an upgrade in a test environment, so that you can determine any adjustments you may need to make in your production environment before upgrading it.

Determine which Sybase products should be running.

If you are upgrading Adaptive Server (or SQL Server), the previously installed version of the server *must* be running.

If you are upgrading Backup Server, Historical Server, Monitor Server, or XP Server, those servers must *not* be running.

Procedure text is required for upgrade

If you created any compiled objects or stored procedures in a pre-11.9.3 server, then removed the syscomments text from sybsystemprocs, the upgrade will succeed but you will encounter runtime problems. To upgrade system procedures, the system procedure text must be available in syscomments. See “Upgrading compiled objects with dbcc upgrade_object” on page 142.

Reserved words

Check for reserved words that may conflict with identifiers currently in use by existing databases. Reserved words are pieces of SQL syntax that have special meaning when used as part of a command.

Transact-SQL does not allow words that are part of command syntax to be used as identifiers, unless they are enclosed in quotation marks, as explained in “Using quoted identifiers” on page 114. If you are upgrading Adaptive Server, and the identifiers in your user databases match new reserved words, errors can result when you run queries, stored procedures, or applications that use these identifiers.

Warning! If a user database name is a reserved word, upgrade fails when it tries to upgrade the database. You must use `sp_renamedb` to change the name of any user database that is a reserved word before performing the upgrade.

Conflicts between object names do not prevent the upgrade process from completing. However, applications that refer to conflicting object names may not work after the upgrade. Be sure to rename all objects that use reserved words.

As part of the pre-upgrade process, `sqlupgrade` can perform the reserved word check for you. See the *Adaptive Server Reference Manual* for a complete list of reserved words.

Note If you change an object name, be sure to change applications and stored procedures that refer to that object.

Running a reserved word check

If you plan to use `sqlupgrade` to perform the upgrade process, you do not have to check for reserved words as a separate step. In `sqlupgrade`, you can check for reserved words at the start of the upgrade process, and then continue with the upgrade if no problems are found.

If you plan to use `sqlupgraderes` to perform the upgrade process using a resource file, you can run `sqlupgrade` first to identify reserved word conflicts and other potential upgrade eligibility problems. You can specify that you do not want the upgrade to proceed even if no problems are found.

If you use either `sqlupgrade` or `sqlupgraderes`, they automatically install new reserved words and `sp_checkreswords`, a stored procedure used to detect and display any identifiers in your existing databases that conflict with reserved words in the new database. You can then run the new version of `sp_checkreswords` at anytime while performing pre-upgrade tasks.

Note You must use the `sp_checkreswords` stored procedure from the new version of Adaptive Server to verify that the old installation of Adaptive Server does not use any reserved words that were introduced with the new server.

Checking for reserved words generates a list of identifiers that conflict with reserved words, and the owners of those identifiers, in the file `$SYBASE/$SYBASE_ASE/init/logs/sqlupgradeMMDD.nnn`. Review this file to determine which identifiers must be changed.

Addressing reserved words conflicts

If any database names are reserved words, you must use `sp_renamedb` to change the database names before you can upgrade. Use `sp_dboption` to set the database to single-user mode, and then run `sp_renamedb`, specifying the new name. See the *Reference Manual* for more information on these procedures.

If other identifiers are reserved words, you can:

- Use `sp_rename` to change the name of the object, before or after the upgrade.
- Use quoted identifiers. See “Using quoted identifiers” on page 114 for more information.

Run `sp_checkreswords` in master and in each user database to display the names and locations of conflicting identifiers.

For more information about `sp_rename` and `sp_checkreswords` and methods for avoiding reserved word conflicts, see the *Reference Manual*.

Using quoted identifiers

You can enclose the identifiers that are reserved words in double quotation marks and invoke the `quoted_identifier` option of the `set` command in procedures and queries that include the reserved words. The `quoted_identifier` option tells Adaptive Server to treat any character string enclosed in double quotation marks as an identifier.

To avoid reserved word conflicts, all users on the server must invoke the `quoted_identifier` option in all stored procedures and queries that include the reserved words.

Verify that users are logged off

Use the following procedure to verify that users are logged off of the system:

- 1 Log in to your current server as the Sybase System Administrator:

```
isql -Usa -Ppassword -Sserver_name
```

- 2 Enter:

```
1> sp_who
2> go
```

This command returns a list of all users and processes currently on the server.

- 3 Log off all users on the server.

The upgrade will not work if you boot the server in single-user mode.

Check database integrity

Check logical and physical database consistency by running the following commands in each database. Use `isql` to log in to the server, and run:

- `dbcc checkalloc`
`dbcc checkalloc [(database_name [, fix | nofix])]`
- `dbcc checkdb`
`dbcc checkdb [{database_name [, skip_ncindex]}]`

- `dbcc checkcatalog`
`dbcc checkcatalog [(database_name)]`

For a description of the dbcc functions, see the *System Administration Guide*.

If you encounter any errors while running these commands, resolve them before continuing the upgrade. For help with resolving errors, see the *Troubleshooting and Error Messages Guide*.

If you are upgrading from Adaptive Server 11.5.1 or later, and you have set up your system to run dbcc checkstorage, you can use that command instead of the other dbcc commands to check database consistency. For details on using dbcc checkstorage, see the *System Administration Guide*.

If a database is marked “suspect,” drop that database with the dbcc dbrepair (*database_name*, dropdb) command.

Back up databases

Use the dump database command to back up your databases. If you encounter any problems during the upgrade, you may need these dumps to recover the previous installation. Be sure to dump the master database.

Warning! *The upgrade process permanently modifies your existing Sybase databases. Back up all your databases, including master, sybssystemprocs, tempdb, and model before beginning the upgrade so you can restore them if necessary.*

For information on backing up and recovering databases, see the *System Administration Guide*.

Dump transaction logs

The upgrade process may consume substantial space in the transaction log. Before beginning the upgrade process, run the dump transaction command in each database to free up log space. See “Backing Up and Restoring User Databases” in the *System Administration Guide*.

Preparing database and devices

Make sure that master is the default database for the “sa” user. The upgrade cannot complete successfully with a different database as your default. For information about specifying the default database, see the *System Administration Guide*.

If you ran the pre-upgrade option in sqlupgrade, the utility scanned the system catalogs and calculated how much additional space you needed to perform the upgrade. If you did not run the pre-upgrade check, you must manually calculate how much free space you must add to your system catalogs and databases.

As a general rule, you must double the size of the largest catalog that you are going to upgrade, and add approximately 10 percent more to calculate the free space required for the upgrade to succeed.

For example, if you have a large number of stored procedures or compiled object in any database, the syscomments and sysprocedures catalogs will require additional space.

Note Sybase recommends that you use preupgrade to determine how much free space is required for the upgrade.

An installation of Adaptive Server 12.0, without any additional user-defined data creates the following:

- tempdb with 2MB
- model with 2MB
- master database with 6MB

If you are upgrading an Adaptive Server installation that meets these minimum parameters, you must increase:

- tempdb to 4MB
- model to 4MB
- master database to 8MB

Use sp_help to determine the size of your existing databases. If any database does not have the space requirements as reported by preupgrade, use alter database to increase the size.

❖ Increasing default database sizes

Before you can successfully upgrade Adaptive Server, the master, model, and tempdb must, at a minimum, meet the size requirements listed in Table 2-11 on page 22.

To determine the current size of each of these databases, log in to Adaptive Server using isql:

```
isql -Usa -Ppassword -Sserver_name
```

Enter:

```
1> sp_helpdb database_name
2> go
```

where *database_name* is the name of the system database you are checking.

To enlarge the master, tempdb, and model databases:

- 1 In isql, use alter database to increase the size of the master database. For example:

```
1> use master
2> go
1> alter database master on master=x
2> go
```

This example increases the size of the master database; *x* is the number of megabytes of space added to the existing database size.

- 2 Repeat this step to increase the size of each of the tempdb and the model databases. For example:

```
1> alter database tempdb on master=x
2> go

1> alter database model on master=x
2> go
```

Note You must increase the size of tempdb before increasing the size of model.

- 3 Verify the size of each database. Enter:

```
1> sp_helpdb database_name
2> go
```

where *database_name* is the name of the system database you are checking.

Create a *sybssystemdb* database

In version 11.9.2, the *sybssystemdb* database was required only for servers using two-phase commit transactions. Starting in version 12.0, all servers must have a *sybssystemdb* database. Adaptive Server uses this database for tracking transactions and during recovery. In addition, it is used for applications using two-phase commit and distributed transaction management (DTM).

sybssystemdb does not need to be located on the master device.

If you have a *sybssystemdb* database

If you have a *sybssystemdb* and the database is at least 4 MB, you do not need to make any changes. If the database is smaller than 4MB, increase the size with `alter database`.

If you do not have a *sybssystemdb* database

If you do not have a *sybssystemdb*, and will not be using two-phase commit or DTM applications, create a *sybssystemdb* with a minimum of 4MB.

If you will be using two-phase commit or DTM, you can either:

- Create a minimum-sized *sybssystemdb* for upgrade and expand it later, or
- Create a *sybssystemdb* of 5 to 20MB, depending on your expected usage. Approximately 25 percent of the database should be data storage, and 75 percent should be log storage.

The following example initializes data and log devices, and creates a 20MB *sybssystemdb*:

```
1> disk init name = "sybssystemdb_data",
2> physname = "/sybase/data/sybssystemdb_data",
3> vdevno = 8,
4> size = 2560

1> disk init name = "sybssystemdb_log",
2> physname = "/sybase/data/sybssystemdb_log",
3> vdevno = 9,
4> size = 7680

1> create database sybssystemdb on sybssystemdb_data=5
2> log on sybssystemdb_log = 15
```

sybsystemprocs

Note `sysprocsdev` is the default system name for this device. However, it is frequently referred to as the `sybsystemprocs` device, as in the Adaptive Server attribute screen, since it stores the `sybsystemprocs` database.

Verify that the `sybsystemprocs` database is large enough. For an upgrade, the recommended minimum size for `sybsystemprocs` is the larger of 105MB, or enough free space to accommodate the existing `sybsystemprocs` database, and the largest catalog that is to be upgraded, plus an additional 10 percent of the largest catalog's size. The additional 10 percent is for logging upgrade changes.

You may need more space if you are adding user-defined stored procedures. 105MBs accommodate additional internal data structures, but does not account for the possibility of a large number of user-defined system procedures.

If your `sybsystemprocs` database does not meet these requirement and you have enough room on the device to expand the database to the required size, use the `alter database` command to increase the database size.

To determine whether the `sybsystemprocs` database is large enough, use `sp_helpdb` to determine the size of the `sybsystemprocs` database:

```
1> sp_helpdb sybsystemprocs
2> go
```

To determine whether the system procedures device is large enough, use `sp_helpdevice` to determine the size of the `sysprocdev` device:

```
1> sp_helpdevice sysprocdev
2> go
```

If the `db_size` setting is less than 105MB, you must increase the size of `sybsystemprocs`.

Increasing the size of the *sybsystemprocs* database

If your current `sybsystemprocs` database does not have the minimum space required, you have two options for creating a new database with sufficient space for the upgrade:

- You can drop the old database and the device and create a new `sysprocsdev` device, or

- You can leave the database and old device alone and add a new device large enough to hold the additional megabytes, and alter the sybsystemprocs onto the new device. Sybase recommends expanding the current device.

To enlarge the sybsystemprocs database:

- 1 If you do not have a current backup, create one.
- 2 In isql, use alter database to increase the size of the sybsystemprocs database. For example:

```
1> use master
2> go
1> alter database sybsystemprocs on sysprocsdev=40
2> go
```

In this example, “sysprocsdev” is the logical name of the existing system procedures device and “40” is the number of megabytes of space to add.

If the system procedures device is too small, you may receive a message similar to the following when you try to increase the size of the sybsystemprocs database:

```
Could not find enough space on disks to extend
database sybsystemprocs
```

If there is space available on another device, you can expand sybsystemprocs to a second device, or initialize another device that is large enough. For instructions on creating a larger sybsystemprocs device, see “Increasing device and database capacity for system procedures” on page 120.

- 3 To verify that Adaptive Server has allocated more space to sybsystemprocs, issue:

```
1> sp_helpdb sybsystemprocs
2> go
```

When the system procedures database is large enough to accommodate the increased size of the sybsystemprocs database, continue with the other pre-upgrade tasks.

Increasing device and database capacity for system procedures

If you cannot fit the enlarged sybsystemprocs database on the system procedures device, increase the size of the device and create a new database.

This procedure involves dropping the database. For more information on drop database, see the *Reference Manual*.

Warning! This procedure removes all stored procedures you have created at your site. Before you begin, save your local stored procedures using the defncopy utility.

To create a larger system procedures device (sysprocsdev):

- 1 Determine which device or devices you need to remove.

Warning! You must *not* remove any device that is in use by database other than sybssystemprocs, or you will destroy that database.

```
select d.name, d.phyname
from sysdevices d, sysusages u
where u.vstart between d.low and d.high
and u.dbid = db_id("sybssystemprocs")
and d.status & 2 = 2
and not exists (select vstart
                from sysusages u2
                where u2.dbid != u.dbid
                and u2.vstart between d.low and d.high)
```

where *d.name* is the list of devices to remove from sysdevices and *d.phyname* is the list of files to remove from your computer.

The “not exists” clause in this query excludes any devices that are used both by sybssystemprocs and other databases.

Note the names of the devices; you will need them in the coming steps.

- 2 Drop sybssystemprocs:

```
use master
drop database sybssystemprocs
```

- 3 Remove the device or devices using sp_configure:

```
sp_configure "allow updates", 1

delete sysdevices
where name in ("d.name")

sp_configure "allow updates", 0
```

where *d.name* is the name of each devices returned in step 1.

Note Each device name must have quotes. For example, “*devname1*”, “*devname2*”, and so on.

If any of the named devices are OS files rather than raw partitions, use the appropriate OS commands to remove those files.

- 4 Remove all files for the list of *d.phyname* returned in step 1.

Note File names may not be complete path names. If they are relative paths, they will be relative to the directory from which your server was started.

- 5 Find another existing device that has 105MB additional free space, or use a disk init command similar to the following to create an additional device for *sybsystemprocs*:

```
1> use master
2> go

1> disk init
2> name = "sysprocsdev",
3> physname = "$SYBASE/work/sysproc.dat",
4> vdevno = 9,
5> size = 51200
6> go
```

where *\$SYBASE/work* is the path to your system procedures device. You cannot use the environment variable *\$SYBASE* in the command.

Note The number for *vdevno* must be available. For information about determining whether *vdevno* is available, see the *System Administration Guide*.

The size you provide should be the number of megabytes of space needed for the device multiplied by 512. *disk init* requires the size to be specified in 2K pages. In this example, the size is 105MB (105 x 512 = 53760). For more information on *disk init*, see the *Reference Manual*.

- 6 Create a *sybsystemprocs* database of the appropriate size on that device, for example:

```
1> create database sybsystemprocs on sysprocsdev =
105
```

```
2> go
```

- 7 Run the installmaster script in the *old* Sybase installation directory, as follows:

```
isql -Usa -Ppassword -Sserver_name -i/old_dir/scripts/installmaster  
-oinstallmaster.out
```

Turn off database options

To turn off database options:

- 1 Run the `sp_helpdb` stored procedure to show the options that are set for each database.
- 2 Use the `sp_dboption` stored procedure and the results from `sp_helpdb` to turn off all options *except* select into/bulk copy on tempdb.

For example, to turn off the trunc log on chkpt database option, enter:

```
sp_dboption sybssystemprocs, 'trunc log on chkpt', false
```

Be sure to turn off the trunc log on chkpt database option to avoid upgrade failure.

For more information about using these system procedures, see the *Reference Manual*.

Note The database options for master cannot be changed and are not disabled.

Disable auditing

To upgrade a server that contains a database in which auditing is activated make sure that auditing is disabled before starting the upgrade:

- 1 Turn off auditing before upgrading.
 - a Verify that all Adaptive Server users are logged off.
 - b Disable auditing with the following Transact-SQL command:

```
sp_auditoption 'enable auditing', 'off'
```

- 2 Record system procedure audit options before upgrading.

When upgrading, all previous versions of system procedures are overwritten when sybssystemprocs is populated. Be sure to record the audit options for all system procedures. After the upgrade is complete, shut down and restart the server. Reenable auditing using:

```
sp_configure auditing, 1
```

You can reenter the audit options, using sp_audit. For more information, see the *System Administration Guide*.

Disable disk mirroring

Use the disk unmirror command to disable disk mirroring for your current Adaptive Server or SQL Server. For more information about this command, see the *Reference Manual*.

Preparing to upgrade servers with replicated databases

After upgrade, you will no longer be able to scan any part of the transaction log that existed before the upgrade, so you must follow the following process if your server contains replicated primary databases (this includes replicated RSSDs). The following procedure will help to ensure that all replicated data from a replicated database has made it safely to the replicate database.

Warning! It is not sufficient to just get the replicated data into the Replication inbound queue, because the inbound queue cannot be rebuilt after the upgrade.

The procedures described here do not upgrade Replication Server itself. For information on upgrading Replication Server, see your Replication Server documentation.

The database upgrade procedure consists of the following activities:

- Suspending transaction processing and replication activities.
- Draining transaction logs for primary databases.
- Draining the Replication Server System Database (RSSD) log.
- Disabling the log truncation point.

After upgrading to version 12.5, complete the post-upgrade tasks to reenable database replications functions.

For more information, see the *Replication Server Reference Manual* and the *Replication Server System Administration Guide*.

Warning! As a safeguard, perform a dump database and a dump transaction before executing the procedures in the following sections.

To determine whether your existing server contains replicated databases:

- 1 Connect to the Server you are upgrading via isql.
- 2 Run the following command in each database (including system databases):

```
1> dbcc gettrunc
2> go
```
- 3 If the command returns “1” for “lrm_trunc_state” in any database, replication is enabled in that database.

Suspending transaction processing and replication activities

To suspend replication of and transaction activity in the databases:

- 1 Verify that subscriptions that are being created by the create subscription command, with primary data in the databases being upgraded, have reached the “valid” state at the primary Replication Server before you begin the upgrade. Use the check subscription command to find the “valid” state.

Defer upgrade while the subscriptions are being created so that Replication Server does not interfere with the upgrade process by accessing the database being upgraded.

Make sure that no users create subscriptions for data in the database you are upgrading until the upgrade procedure is finished.

- 2 Execute rs_helproute in each Replication Server System Database (RSSD) being upgraded.

The status of all existing routes should be “Active”. If any route is not active, resolve the route before continuing. See the Replication Server documentation for help in diagnosing and fixing the problem with the route, and then go to step 3.

- 3 Shut down all applications that use the databases you are upgrading.
- 4 Use the `admin who` command in Replication Server to find the existing Data Server Interface (DSI) connections to the data server being upgraded.
- 5 Suspend all DSI connections to the non-RSSD databases you are upgrading by entering the following command in Replication Server for each database:

```
1> suspend connection to dataserver.database
2> go
```
- 6 Leave the DSI connections to the RSSD databases running.

Draining the transaction logs for primary databases

For each primary database you are upgrading, you need to ensure that Replication Server completely processes the pre-upgrade log.

To drain the transaction logs:

- 1 Wait for all remaining transactions to be replicated.
- 2 Run the following Replication Server command:

```
1> admin who, sqm
2> go
```
- 3 Find the entry that corresponds to the inbound queue for this database by looking for the Info field for the `queue_number:queue_type` entry. For an inbound queue, the queue type is 1. Note the Last Seg.Block entry for the queue.
- 4 Open a queue dump file by executing the following Replication Server command:

```
1> sysadmin dump_file, "file_name"
2> go
```

where *file_name* is the name of the file to which you will dump the queue.

- 5 Use `isql` to update one row in a single replicated table in the primary database:

```
1> update table set column = column
2> where key = unique_value
3> go
```

The update command helps to track whether all modifications to the replicated database have been sent to the Replication Server.

Note In Replication Server 10.1 or later, choose a table that does not use the replicate minimal columns clause, or use the alter replication definition command...replicate all columns command to change the replication definition before updating the row. If you alter the replication definition, be sure to change it back after you complete this upgrade procedure.

- 6 In the primary Replication Server, execute the admin who, sqm command until the last segment: block entry for the inbound queue changes.
- 7 Execute the following Replication Server command to dump the last block of the inbound queue to the dump file you created in step 3:

```
1> sysadmin dump_queue, queue_number,  
2> queue_type, last_seg, block, 1  
3> go
```

Use the *queue_number*, *queue_type*, *last_seg*, and *block* values found in the output of the last admin who, sqm command.

- 8 Examine the dump file to make sure it contains the transaction that corresponds to the update you performed in step 4. (You can use Notepad to examine the file.)
- 9 Repeat steps 5–7 until the transaction that corresponds to the update is in the dump file.
- 10 Log into the Replication Server and suspend the Log Transfer connection from that database:

```
1> suspend log transfer from server.database  
2> go
```

- 11 If you are using Rep Agent, log into the ASE server, and stop the Rep Agent:

```
1> use database  
2> go  
  
1> sp_stop_rep_agent database  
2> go
```

- 12 If you are using LTM, shutdown the LTM.

After draining the transaction logs, do not allow any other activity in the databases. If activity does occur, you will need to drain the logs.

Draining the RSSD Transaction Log

If the Replication Server has routes to other Replication Servers, you must ensure that Replication Server processes all transactions in the RSSD transaction log before you upgrade the databases.

To see whether the transaction log has been processed completely, create a replication definition in the primary Replication Server and then watch for it to appear in the replicate Replication Server's RSSD. When the replication definition is in the replicate RSSD, you can assume that the log is processed fully.

To ensure that the RSSD log is processed:

- 1 Log into the primary Replication Server and create a temporary replication definition:

```
1> create replication definition rep_def_name
2> with primary at dataserver.database
3> (column_a int)
4> primary key (column_a)
5> go
```

The data server and database names must be valid, but the replication definition does not have to reference an actual table.

- 2 Log into the replicate RSSD (not the primary RSSD) and execute the following query to find out if the replication definition has arrived from the primary RSSD:

```
1> select * from rs_objects
2> where objname = "rep_def_name"
3> go
```

If this select statement returns rows, the last replication definition created in step 1 has been sent successfully to the replicate RSSD. This means that the transaction log has been drained.

- 3 Log into the replicate Replication Server and suspend the Log Transfer connection from the primary RSSD:

```
1> suspend log transfer from server.database
2> go
```

- 4 If you are using Rep Agent, log into the ASE server, and stop the Rep Agent:

```
1> use database
2> go
```



```
1> sp_stop_rep_agent database
2> go
```

- 5 If you are using LTM, shutdown the LTM.

Disabling the secondary truncation point

When you upgrade a primary database, the Rep Agent or the Log Transfer Manager (LTM) must not be running, and the Secondary truncation point should be turned off for the duration of the upgrade. The Rep Agent or Log Transfer Manager should already be shutdown (from the previous steps).

For each primary database and replicated RSSD, disable the secondary truncation point:

- 1 If this is a replicated RSSD, log into the Replication Server of the RSSD, and issue:

```
1> sysadmin hibernate_on, 'Replication Server'
2> go
```

- 2 Disable the secondary truncation point in the Adaptive Server database that is being upgraded by issuing:

```
1> use database
2> go

1> dbcc settrunc('ltm', 'ignore')
2> go
```

If the dbcc settrunc command fails, make sure that the Rep Agent or LTM are not running. When the Rep Agent and LTM are disabled, repeat this step.

Restoring replication after upgrade

Restore the replication functionality:

- 1 Log into the RSSD for each replicated primary, and for each replicated RSSD, and issue:

```
1> use RSSD_name
2> go

1> rs_zeroltm dataserver, database
2> go
```

- 2 Log into each replicated primary, and replicated RSSD and issue:

```
1> use database
2> go

1> dbcc settrunc ('ltm', 'valid')
2> go
```

- 3 Restart the Replication Server.

- 4 If the database is also used as a RSSD, resume the Replication Server connection to the RSSD by issuing the following command to the Replication Server:

```
1> sysadmin hibernate_off, 'Replication Server'
2> go
```

Make sure to specify the same string that was specified during the 'hibernate_on' command.

- 5 Log into the Replication Server, and resume the Log Transfer connection for each replicated primary and each replicated RSSD:

```
1> resume log transfer from server.database
2> go
```

If this is a replicated RSSD, you will need to log into the replicate Replication Server.

- 6 If you are using Rep Agent, log into the Adaptive Server and restart the Rep Agent:

```
1> use database
2> go

1> sp_start_rep_agent database
2> go
```

- 7 If you are using LTM, restart the LTM.

Upgrading to Adaptive Server 12.5

To assist with the upgrade process, Sybase recommends that you set the OLDSYBASE environment variable to point to the old installation of Adaptive Server. If you set the OLDSYBASE environment variable to the old installation of the Adaptive Server, the upgrade utility will automatically populate input fields with the correct information as you perform the upgrade.

Back up all existing Sybase installations. Installing Adaptive Server files into an existing Sybase installation directory overwrites the previous installation.

Warning! If you install Adaptive Server files into the same directory as the currently installed version, the new version overwrites the old one. If this occurs, the upgrade cannot complete, because you must have both the old and the new versions of Adaptive Server to perform an upgrade.

Use either of the following methods to upgrade Adaptive Server:

- `sqlupgrade` – provides an X-Windows or Motif GUI for the upgrade process. See “Using `sqlupgrade`” on page 131.
- `sqlupgraderes` – provides a noninteractive, file-based interface for the upgrade process. `sqlupgraderes` allows you to upgrade Adaptive Server using a resource file. This is useful for customer sites that want to upgrade many similar servers. See “Using `sqlupgraderes`” on page 134.

Note Backup Server, Monitor Server, and XP Server cannot be upgraded using `sqlupgraderes`.

Using *sqlupgrade*

`sqlupgrade` verifies that the server is eligible for the upgrade. Perform all pre-upgrade tasks before running `sqlupgrade`.

To upgrade using `sqlupgrade`:

- 1 Make sure the server you are upgrading is running.
- 2 Verify that your SYBASE environment variable points to the location of the new Adaptive Server software files you just unloaded.

You also can set the OLDSYBASE environment variable to the location of the server you are upgrading, to avoid typing this path when you run `sqlupgrade`.

- 3 Verify that you have performed the tasks mentioned on the screen, before beginning the upgrade.

Note sqlupgrade automatically copies the *interfaces* file entry of the server you are upgrading from the old *interfaces* file to the new *interfaces* file.

- 4 Execute sqlupgrade in either of the following ways:

- At the UNIX prompt, enter:

```
$SYBASE/$SYBASE_ASE/bin/sqlupgrade
```

A sqlupgrade screen is displayed.

- Alternatively, you may access sqlupgrade through the asecfg utility. Then click on the Upgrade an Existing Server button.

- 5 Click OK.

- 6 The Specify Sybase Directories screen of sqlupgrade is displayed. Provide the installation directory location for both the new server release and your current server.

If the SYBASE environment variable is set, sqlupgrade displays the value of the SYBASE environment variable as the directory location for the new server release.

If the OLDSYBASE environment variable is set, sqlupgrade displays that value as the directory location for your current server.

- 7 Enter the name of the component directory that contains the new server version. For example, *ASE-12_5*.

- 8 Enter the full path of the old Sybase release directory. For example, */opt/sybase-11.9.2*.

If the OLDSYBASE environment variable is set, sqlupgrade displays that value as the directory location for the old server.

- 9 Leave the field, “Enter the name of the component directory which contains the old version” blank.

- 10 Click Continue. The Adaptive Server Selection screen displays.

- 11 From the list of server names provided, select the server you want to upgrade.

If the server you are upgrading is not running, sqlupgrade prompts you to start the server. The server you are upgrading *must* be running.

- 12 Click OK. The SA Password screen is displayed.
- 13 Enter the password for the System Administrator (sa) login.
- 14 Click OK. The Specifying Upgrade Options screen is displayed.
- 15 Specify an upgrade option.

You can run the upgrade eligibility tests before performing the actual upgrade, or you can upgrade immediately after the upgrade eligibility tests run successfully.

The upgrade eligibility tests verify that your current server meets the requirements for upgrading to the new version. (The online help provides a partial list of what the upgrade eligibility tests check.) If your current server fails the upgrade eligibility test, you must fix the problem.

Note Adaptive Server 11.9.x did not automatically create a sybsystemproc database and device. Unless you created one, the eligibility test will fail. The Back button launches the “Specify sybsystemprocs” screen.

Enter the full path to the sybsystemproc device and 105MB for the sybsystemproc database. Click OK to continue with the eligibility test.

- 16 Click OK to upgrade your server. sqlupgrade displays the Status Output screen.

View the Status Output screen for the completion status and informational messages about the upgrade process.

Do not interrupt the upgrade, and do not try to connect to Adaptive Server or run any stored procedures while the upgrade is in progress.

Note On some UNIX platforms, the sqlupgrade utility may not return you to the command prompt even though the upgrade has completed successfully. To exit the utility, force an exit by pressing Ctrl + C.

You can also check the progress of the upgrade by viewing the sqlupgrade log in `$SYBASE/$SYBASE_ASE/init/logs/sqlupgradeMMDD.NNN`.

If the upgrade is successful, a “Done” message is displayed in the Status Output screen.

- 17 When the upgrade completes successfully, you can:
 - Click OK to specify another server to upgrade, or
 - Exit sqlupgrade, and go to “Post-upgrade tasks” on page 135.

Using *sqlupgraderes*

You can upgrade Adaptive Server in noninteractive mode by using values from a resource file that defines the attributes for the server to be upgraded.

To upgrade an Adaptive Server using a resource file, either:

- Edit the sample resource file included in your Adaptive Server distribution in `$SYBASE/$SYBASE_ASE/init/sample_resource_files/upgrade.rs`. Use an ASCII text editor to edit the resource file as described in Appendix A, “Alternative Installation Methods.”.
- Run the `sqlupgrade` utility. This utility creates a new file containing the values you specified for that upgrade session, and writes it to:

`$SYBASE/$SYBASE_ASE/init/logs/sqlupgradeMMDD.NNN-server_name.rs`

where *server_name* is the name of the upgraded server.

Warning! Be sure to change device names in the resource file, if necessary, before executing the resource file using `sqlupgraderes`.

For details on the attributes used in a resource file, see “Resource file attributes for upgrading Adaptive Server” on page 134.

Resource file attributes for upgrading Adaptive Server

Table 8-1 on page 135 shows the resource file attribute names, their default values, and other options, which you use in the resource file for upgrading Adaptive Server. The sample resource file in the `$SYBASE_ASE/init/sample/upgrade` file included in your Adaptive Server distribution shows examples of these attributes.

All attributes are *required*. And all values are case sensitive.

If the resource file you are modifying was created by `sqlupgrade`, the prefix of the attribute name may vary. `sqlupgraderes` ignores this prefix when processing the resource file.

Table 8-1: Resource file attributes for upgrading Adaptive Server

| Attribute | Default value [other options] |
|---------------------------|--|
| sybinit.release_directory | \$SYBASE [<i>path_name_of_old_srv</i>] |
| sybinit.product | sqlsrv |
| sqlsrv.server_name | <i>server_name</i> |
| sqlsrv.new_config | no (required value) |
| sqlsrv.sa_login | <i>current_login</i> |
| sqlserv.sa_password | <i>current_password</i> |
| sqlsrv.do_upgrade | yes (required value) |

Upgrading using *sqlupgraderes*

When you have finished editing your resource file, you can run the file using *sqlupgraderes*.

Note *sqlupgraderes* is a command-line tool; therefore, unlike *sqlupgrade*, it does not require any X libraries.

To execute *sqlupgraderes*, at the UNIX prompt, enter:

```
$SYBASE/$SYBASE_ASE/bin/sqlupgraderes -r resource_file
```

where *resource_file* specifies the resource file containing the attributes that describe the server to upgrade.

Post-upgrade tasks

After you have upgraded to the new version of Adaptive Server, perform the following tasks to make sure your new Adaptive Server is up and running.

Note After upgrade, there is no need to run update statistics on any tables. Upgrade process does not cause any changes in existing statistics.

Verifying that servers are running

- 1 To find out whether Adaptive Server is currently running, try to log in to the server, using isql:

```
isql -Usa -P -Sserver_name
```

If the server is running, you see the isql prompt:

```
1>
```

You can also use the showserver command to determine whether any Adaptive Server processes are running. For example:

```
$SYBASE/$SYBASE_ASE/install/showserver
```

- 2 To verify that you are at the new version level, connect to Adaptive Server via isql and run the following commands:

```
1> select @@version
2> go
```

Look for “12.5” in the version string.

```
1> sp_configure "upgrade version"
2> go
```

sp_configure should return the Run Value “12500”.

Restore functionality in Adaptive Server

To reset or enable options or functionality you had to change before the upgrade:

- 1 Reset the configuration parameters.

If you changed any configuration parameters before the upgrade to ensure enough memory, use the sp_configure command to set them back to their previous values.
- 2 Use sp_dboption to reset any database options you disabled before the upgrade.
- 3 After you complete the upgrade, but before you use the upgraded Adaptive Server, verify that all scripts developed at your site point to Adaptive Server 12.5.
- 4 Procedure cache allocation after upgrade.

This section is for information only. Adaptive Server ensures that procedure cache size after upgrade is the same as before upgrade, unless the before upgrade size was less than the default value.

In Adaptive Server 12.5, procedure cache percent has been changed to procedure cache size. Adaptive Server now allocates memory dynamically and determines the size of the procedure and data caches in terms of megabytes. During the upgrade process, Adaptive Server converts the procedure cache percent to an absolute value and sets procedure cache size based on this value.

For example, if procedure cache percent is configured as 20 percent of 200MB of memory, then the absolute value of procedure cache percent is 40MB. During the upgrade process, Adaptive Server converts procedure cache percent at 20 percent to procedure cache size at 40MB. If the value for procedure cache size is less than the default value, Adaptive Server sets procedure cache size to the default value. For example, if procedure cache percent is set to 10 percent in the example above, and the absolute value of 10 percent is less than the default value of procedure cache size, then procedure cache size will be set to the default value.

With the introduction of dynamically reconfigured memory configuration parameters in Adaptive Server 12.5, an increase in Adaptive Server's memory use will not decrease the size of the procedure cache or the data cache. That is, if your default data cache is set to 40MB, and you increase your procedure cache, the default data cache is still 40MB.

For more information about `sp_configure`, see the *Reference Manual*.

5 Check the procedure cache requirements.

Stored procedures, triggers, and other compiled objects require more memory to run than older versions. The memory required to run a stored procedure increased by 20 percent between versions 10.x and 11.5. Adaptive Server 12.5 may need more procedure cache for the server to maintain the same performance.

You may increase the procedure cache size during run-time, using `sp_configure`. You may verify any changes you make to the configuration file without having to reboot Adaptive Server using the `verify` option in `sp_configure`. The syntax is:

```
sp_configure "configuration file", 0, "verify",  
"full_path_to_file"
```

For example:

```
sp_configure "configuration file", 0, "verify",
```

```
"/work2/Sybase/ASE125.cfg"
```

For more information about `sp_configure` and `sp_sysmon`, see the *Reference Manual* and the *Performance and Tuning Guide*. For more information about configuring memory, see the *System Administration Guide*.

- 6 Data cache allocation after upgrade. This section is for information only. Adaptive Server ensures that all the data cache sizes after upgrade will be the same as before upgrade.

In pre-12.5 versions of Adaptive Server, the default data cache available depended on the amount of memory available to Adaptive Server. If the default data cache was set to 2MB, Adaptive Server allocated 2MB to the default data cache. If, after memory was allocated to all user-defined data caches and to the procedure cache, there was more memory available, that memory was allocated to the default data cache. However, if there was not enough memory available for Adaptive Server to set up the user-defined caches and the procedure cache, then memory would be taken from the default data cache regardless of the designated configuration parameter. Therefore, it was possible to set default data cache to the default value and have all remaining available memory be allocated to the default data cache.

Adaptive Server 12.5 treats the default data cache size as an absolute value and sets it in the *config* file. The string `DEFAULT` in the *config* file has a different meaning for default data cache size in Adaptive Server 12.5.

In pre-12.5 Adaptive Server, it meant all the memory left-over after allocating memory for other configurations, procedure and user defined named caches. In Adaptive Server 12.5, it means a default value of 8MB.

During upgrade Adaptive Server ensures that the default data cache size before upgrade is the same as after upgrade. Therefore, during the preupgrade process, the size of the default data cache is obtained and written to the configuration file as an absolute value, not as “`DEFAULT`.” This step is done to enable Adaptive Server 12.5 to have the same default data cache size as before the upgrade. If this size is less than the default size of default data cache of 8MB, then Adaptive Server 12.5 allocates a default data cache of size 8MB.

- 7 If you unmirrored devices, remirror them, using the `disk remirror` command.
- 8 If you used two-phase commit in 11.9.x, run the script to install the two-phase commit tables:

```
isql -Usa -Psa_password -Sserver_name  
-i$SYBASE/$SYBASE_ASE/scripts/installcommit
```

You do not need to run this script if you do not use two-phase commit.

Reenabling replication server

If you disabled replication before the upgrade, you must reenabling replication. To do so, first remove any older format log records from your database, and then reenabling replication.

Removing old log records

Use the dump database command after the upgrade to dump the database and transaction logs to remove the older format log records from your database. This prevents Replication Server from accessing the pre-upgrade portion of the transaction logs.

For example, to dump the sales database:

```
1> use master  
2> go  
1> dump database sales to "//./backup"  
2> go
```

Reenabling replication

Follow the steps in this section to reenabling replication after upgrading the primary databases and primary RSSDs.

For each primary database and RSSD:

- 1 Start Adaptive Server if it is not already running.
- 2 Log on to the server.
- 3 Clear the locator for the database by executing the following command in the RSSD for this database:

```
1> use RSSD  
2> go  
3> rs_zeroltm dataserver, database  
4> go
```

- 4 Enable the truncation point for the database.

```
1> use database
```

```
2> go
1> dbcc settrunc("ltm", "valid")
2> go
```

5 Restart the Replication Servers and LTMs.

6 Resume the DSI connections that were suspended before the upgrade by executing the following Replication Server command for each suspended database:

```
1> resume connection to dataserver.database
2> go
```

The replication system is now ready for Adaptive Server 12.5, and applications can resume.

If you installed any Sybase client products, such as Open Client, use the dsedit utility to edit your interfaces file and to specify the servers to which you want to connect.

For more information about establishing client/server connections, see the *Open Client Configuration Guide* or Chapter 8, “Upgrading Sybase Servers.”

Reenabling auditing

Reenable all auditing options for stored procedures using the audit options you recorded during pre-installation. See “Disable auditing” on page 123). Reenter the audit options using `sp_audit`.

- Adaptive Server 12.5 includes the following global audit options for `sp_audit`:
 - security
 - dbcc
- These options are turned off by default and will not be turned on even if you used `sp_auditoption all` to turn on all global audit options in a version earlier than 11.5.

To turn on these options, use:

```
1> sp_audit security "on"
2> go
1> sp_audit "dbcc" "on"
2> go
```

For more information, see the *Reference Manual*.

- The functionality provided by these obsolete `sp_auditooption` options is now covered by the `sp_audit` security global audit option:
 - server boot
 - role toggle
- The security option is turned off by default.

If any of these options were turned off before the upgrade, reset the security option to on to achieve the same auditing actions.

Upgrading Backup Server, Monitor Server, and XP Server

Before you can upgrade Backup or Monitor Server, you must upgrade Adaptive Server.

XP Server does not have a formal upgrade process.

Note Back up all databases, including master, after successfully upgrading Adaptive Server and Backup Server.

❖ Upgrading Backup and Monitor Servers

The procedure for upgrading Backup Server and Monitor Server are similar.

- If you are upgrading during the initial installation of Adaptive Server 12.5, select Upgrade Existing Servers from the drop-down menu when the Studio Installer prompts Build or Upgrade ASE, Backup, and Monitor Servers.

This launches the `sqlupgrade` utility. Select OK.

- If you are upgrading after the initial installation, start the `sqlupgrade` utility from the command line. Enter:

```
$SYBASE/$SYBASE_ASE/bin/sqlupgrade
```

- 4 From the Server selection screen, select the server to upgrade. Click OK.
- 5 Provide the release path:

- Enter the path where the Sybase Server is installed.
 - Enter the path to the Backup Server you are upgrading. Leave this field blank if the Backup Server is older than 12.0.
 - Enter the path to the Open Client/Server Library that the old Backup Server uses. Leave this field blank if it is older than version 12.0.
 - Click OK.
- 6 The Status output screen display the upgrade progress. When the upgrade is complete it displays, “Done.” Select OK.
 - 7 You return to the sqlupgrade screen. You can either upgrade another server, or click Exit to leave the utility.

Repeat this procedure to upgrade Monitor Server.

Upgrading compiled objects with dbcc upgrade_object

Adaptive Server version 11.9.3 introduced the process of upgrading compiled objects based on their source text. Compiled objects are:

- Check constraints
- Defaults
- Rules
- Stored procedures (including extended stored procedures)
- Triggers
- Views

The source text of each compiled object is stored in the *syscomments* table, unless it has been manually deleted. When you upgrade the server, the existence of the source text in *syscomments* is verified during that process. However, the compiled objects are not actually upgraded until they are invoked.

For example, say you have a user-defined stored procedure named *list_proc*. The presence of source text for *list_proc* is verified when you upgrade to Adaptive Server 12.5. Then, the first time *list_proc* is invoked after the upgrade, Adaptive Server detects that the *list_proc* compiled object has not been upgraded. Adaptive Server recompiles *list_proc*, based on the source text in *syscomments*. The newly compiled object is then executed.

Upgraded objects retain the same object ID and permissions that they used prior to being upgraded.

Compiled objects for which the source text was hidden using `sp_hidetext` are upgraded like objects for which the source text is not hidden. For information on `sp_hidetext`, see the *Reference Manual*.

Note The size of each 64-bit compiled object in the sysprocedures table in each database increases by approximately 55 percent when the object is upgraded. The pre-upgrade calculates the exact size. Be sure to increase your upgraded database size accordingly.

To ensure that compiled objects have been upgraded successfully *before* they are invoked, you can upgrade them manually using the `dbcc upgrade_object` command. For details, see “Finding compiled object errors before production.”

Finding compiled object errors before production

Changes made in previous releases of Adaptive Server may cause compiled objects to work differently in version 12.5 and later. You can use `dbcc upgrade_object` to find the following errors and potential problem areas that may require manual changes to achieve the correct behavior:

- Reserved word errors
- Missing, truncated, or corrupted source text
- Quoted identifier errors
- Temporary table references
- `select *` potential problem areas

After reviewing the errors and potential problem areas, and fixing those that need to be changed, you can use `dbcc upgrade_object` to upgrade compiled objects manually instead of waiting for the server to upgrade the objects automatically. For details, see “Using `dbcc upgrade_object`” on page 146.

Reserved word errors

If dbcc upgrade_object finds a reserved word used as an object name in a compiled object, it returns an error, and the upgrade of that object fails. To fix the error, either manually change the object name or use quotes around the object name and issue the command set quoted identifiers on. Then, drop and recreate the compiled object.

For example, suppose you load a database dump from Adaptive Server 11.5 into Adaptive Server 12.5 and the dump contains a stored procedure that uses the word “lock.” When you run dbcc upgrade_object on that stored procedure, the command returns an error because, although “lock” was not reserved in version 11.5, it became a reserved word in version 11.9.2. With this advance notice, you can change the stored procedure and any related tables before they are used in a production environment.

Missing, truncated, or corrupted source text

If the source text in syscomments was deleted, truncated, or otherwise corrupted, dbcc upgrade_object may report syntax errors. If the source text was not hidden, you can use sp_helptext to verify the completeness of the source text. If truncation or other corruption has occurred, drop and re-create the compiled object.

Quoted identifier errors

dbcc upgrade_object returns a quoted identifier error if:

- The compiled object was created in a pre-11.9.2 release with quoted identifiers active (set quoted identifiers on).
- Quoted identifiers are not active (set quoted identifiers off) in the current database.

To avoid this error, activate quoted identifiers before running dbcc upgrade_object. When quoted identifiers are active, you must use single quotes instead of double quotes around quoted dbcc upgrade_object keywords.

If quoted identifier errors occur, use the set command to activate quoted identifiers, and then run dbcc upgrade_object to upgrade the object.

For compiled objects created in 11.9.2 or later, the upgrade process automatically activates or deactivates quoted identifiers as appropriate.

Note Quoted identifiers are not the same as literals enclosed in double quotes. The latter do not require you to perform any special action before the upgrade.

Temporary table references

If a compiled object such as a stored procedure or trigger refers to a temporary table (`#temp table_name`), which was created outside the body of the object, the upgrade fails, and `dbcc upgrade_object` returns an error. To correct this error, create the temporary table exactly as expected by the compiled object, and then execute `dbcc upgrade_object` again. This is not necessary if the compiled object is upgraded automatically when it is invoked.

select * potential problem areas

In Adaptive Server version 11.9.3 and later, the results of a `select *` clause in a stored procedure, trigger, or view that was created in an earlier version of Adaptive Server may be different from what you expect.

For more information about the changes, see the *Reference Manual*.

If `dbcc upgrade_object` finds a `select *` clause in the outermost query block of a stored procedure, it returns an error, and does not upgrade the object.

For example, if a stored procedure has the following statements:

```
create procedure myproc as
    select * from employees
go
create procedure yourproc as
    if exists (select * from employees)
        print "Found one!"
go
```

`dbcc upgrade_object` returns an error on `myproc` because `myproc` includes a statement with a `select *` clause in the outermost query block. This procedure is not upgraded.

`dbcc upgrade_object` does not return an error on `yourproc` because the `select *` clause occurs in a subquery. This procedure is upgraded.

Determining whether *select ** should be changed in views

If dbcc upgrade_object reports the existence of *select ** in a view, compare the output of syscolumns for the original view to the output of the table, to determine whether columns have been added to or deleted from the table since the view was created.

For example, suppose you have the following statement:

```
create view all_emps as select * from employees
```

Before upgrading the all_emps view, use the following queries to determine the number of columns in the original view and the number of columns in the updated table:

```
select name from syscolumns
  where id = object_id("all_emps")
select name from syscolumns
  where id = object_id("employees")
```

Compare the output of the two queries. If the table contains more columns than the view, and retaining the pre-upgrade results of the *select ** statement is important, change the *select ** statement to a *select* statement with specific column names. If the view was created from multiple tables, check the columns in all tables that comprise the view and rewrite the *select* statement if necessary.

Warning! Do not execute a *select ** statement from the view. Doing so will upgrade the view and the information about the original column information in syscolumns will be overwritten.

Another way to determine the difference between the columns in the view and in the new tables is to run sp_help on both the view and the tables that comprise the view.

This comparison works only for views, not for other compiled objects. To determine whether *select ** statements in other compiled objects need to be revised, review the source text of each compiled object.

Using dbcc upgrade_object

Syntax

```
dbcc upgrade_object [ ( dbid | dbname
[, ['database.[owner].]compiled_object_name' |
'check' | 'default' | 'procedure' | 'rule' |
'trigger' | 'view'
[, 'force' ] ] ) ]
```

where:

- *dbid* specifies the database ID. If you do not specify *dbid*, all compiled objects in the current database are upgraded.
- *dbname* specifies the database name. If you do not specify *dbname*, all compiled objects in the current database are upgraded.
- *compiled_object_name* is the name of a specific compiled object you want to upgrade. If you use the fully qualified name, *dbname* and *database* must match, and you must enclose the fully qualified name in quotes. If the database contains more than one compiled object of the same name, use the fully qualified name. Otherwise, all objects with the same name are parsed, and if no errors are found, upgraded.
- *check* upgrades all check constraints and rules. Referential constraints are not compiled objects and do not require upgrading.
- *default* upgrades all declarative defaults and the defaults created with the *create default* command.
- *procedure* upgrades all stored procedures.
- *rule* upgrades all rules and check constraints.
- *trigger* upgrades all triggers.
- *view* upgrades all views.

The keywords *check*, *default*, *procedure*, *rule*, *trigger*, and *view* specify the classes of compiled objects to be upgraded. When you specify a class, all objects in that class, in the specified database, are upgraded, provided that *dbcc upgrade_object* finds no errors or potential problem areas.

- *force* specifies that you want to upgrade the specified object even if it contains a *select ** clause. Do not use *force* unless you have confirmed that the *select ** statement will not return unexpected results. The *force* option does not upgrade objects that contain reserved words, contain truncated or missing source text, refer to nonexistent temporary tables, or do not match the quoted identifier setting. These objects must be fixed before they can be upgraded.

Note If *set quoted identifiers* is on, use single quotes around the keywords. If *set quoted identifiers* is off, you can use either double quotes or single quotes.

Examples

```
dbcc upgrade_object
```

Upgrades all compiled objects in the active database.

```
dbcc upgrade_object(listdb, 'procedure')
```

Upgrades all stored procedures in the listdb database. Single quotes are used around procedure because set quoted identifiers is on.

```
dbcc upgrade_object(listdb, "rule")
```

Upgrades all rules and check constraints in the listdb database. Double quotes are used around rule because set quoted identifiers is off.

```
dbcc upgrade_object(listdb, list_proc)
```

Upgrades all stored procedures named list_proc in the listdb database.

```
dbcc upgrade_object(listdb,  
    "listdb.jkarrik.list_proc")
```

Upgrades the stored procedure list_proc, which is owned by the login "jkarrik".

```
dbcc upgrade_object(master,  
    "listdb.jkarrik.list_proc")
```

Returns an error because the value of *dbname* is master and the value of *database* is listdb. These values must match.

Permissions

Only the Database Owner or a System Administrator can execute dbcc upgrade_object. The Database Owner can upgrade his or her own objects in the database.

Upgraded objects retain the same owner that they had prior to being upgraded.

Increasing the log segment size

You can specify that all compiled objects of a particular class should be upgraded in one execution of dbcc upgrade_object; for example, you can upgrade all triggers by using the trigger keyword. However, even though you use only one dbcc command, the upgrade of each object is recorded in a separate transaction; the old row is deleted from sysprocedures and a new row is written. Therefore, if you run dbcc upgrade_object on a large number of compiled objects, your system may run out of log space. Be sure to increase the size of the log segment in the databases in which you plan to run this command, to allow sufficient room to log all the upgrades.

Error reporting

To send all the output from dbcc upgrade_object to the screen, a System Administrator can start the server with dbcc traceon(3604). Using this command is recommended if you think the output of error messages might overflow the error log.

Upgrading compiled objects in database dumps

When you load a database dump that was created in an earlier version than the current Adaptive Server, you are not required to perform the pre-upgrade tasks before loading the dump. Therefore, you will not receive any notification if the compiled objects in your database dump are missing their source text. After loading a database dump, run `sp_checksource` to verify the existence of the source text for all compiled objects in the database. Then, you can allow the compiled objects to be upgraded as they are executed, or you can run `dbcc upgrade_object` to find potential problems and upgrade objects manually.

For information on using `sp_checksource`, see the *Reference Manual*.

Determining whether a compiled object has been upgraded

To determine whether a compiled object has been upgraded, do one of the following:

- Look at the `sysprocedures.version` column. If the object was upgraded, this column will contain the number 12500.
- If you are upgrading to a 64-bit pointer size in the same version, look at the `sysprocedures.status` column. It will contain a hexadecimal bit setting of 0x2 to indicate that the object uses 64-bit pointers. If the bit is not set, the object is a 32-bit object, which means it was not upgraded.

Recovering from a failed upgrade

- A common failure you might encounter, even after going through the pre-upgrade tasks, is that the log may get full due to the catalog changes that are being done during upgrade. If so, log in via `isql` to the new server, and issue:

```
isql> dump tran dbname with no_log
```

This will free up the log space, and allow the upgrade process to continue.

- In some cases, the space estimations done by preupgrade might be insufficient for the data copy phase of upgrade. In that case, you may get an error that there is insufficient space in the system segment for the upgrade step. The upgrade process will hang, waiting for space to be provided. You can log in to the new server using isql, and use `alter database` to increase the size of the database.

```
isql> alter database dbname on device_name = "2m"
```

Note In 12.5 Adaptive Server, `alter database` allows you to specify the size to alter the database with the unit-specifier 'm', or 'M'.

Troubleshooting

This chapter provides instructions for troubleshooting installation error messages.

If this chapter does not describe the error message you are encountering, see the *Troubleshooting and Error Messages Guide*.

| Topic | Page |
|---|------|
| Overview | 151 |
| Error log locations | 152 |
| Solutions to common installation problems | 153 |
| Troubleshooting SySAM | 160 |

Overview

For Sybase server products, there are two categories of errors:

- Errors generated by the installation, upgrade, and setup utilities
- Errors generated by the server (Adaptive Server, Backup Server, and so on)

To determine the cause of an error, first look in the log file of the utility being used, to identify the task the utility was performing when it failed. Then check the server error log. See “Error log locations” on page 152 for the location of the error log files for the installation utilities and the servers.

Table 9-1 lists possible causes and solutions for common problems that you might encounter during a first-time installation or upgrade. If you continue to have problems, retry the installation or upgrade.

If the installation program or `svrbuild` unexpectedly quits, or if you are unable to correct the problem, see the *Troubleshooting and Error Messages Guide*.

Table 9-1: Troubleshooting guidelines

| Problem | Possible cause and solution |
|--|--|
| The installation program cannot start Adaptive Server. | <p>Failure to boot Adaptive Server is generally caused by a lack of available RAM or disk space.</p> <p>Make sure you meet RAM requirements. If you have the required RAM, remove then reinstall all applications to the hard drive and restart the installation process.</p> <p>After Adaptive Server is installed, there should be 25MB of free disk space left in the disk drive. Adaptive Server needs approximately 18MB for creating shared memory files.</p> <p>Verify that you are logged in as system administrator. You must log in as an administrator to start Adaptive Server.</p> <p>Shut down Monitor Server before restarting Adaptive Server.</p> |
| After upgrading Adaptive Server, you cannot use svrbuild. | After you begin upgrading a server, you may be unable to use the same svrbuild session for other tasks. Exit and restart svrbuild. |
| The installation program cannot connect to the upgraded Adaptive Server. | After you begin upgrading a server, you may be unable to use the same svrbuild session for other tasks. Exit and restart svrbuild. |
| The installation program detects reserved word conflicts. | See “Starting and Stopping Servers” on page 99. |
| The upgrade fails. | See “Recovering from a failed upgrade” on page 158. |

Error log locations

The information in the error logs can help you determine the reason and possible solution for an error message.

For installation utilities

Table 9-2 lists the default error log locations for the installation, upgrade, and setup utilities.

Table 9-2: Error log locations for installation utilities

| Utility | Error log location |
|-----------------------------|---|
| License manager | <code>\$SYBASE/installer.log</code> |
| srvbuild srvbuildres | <code>\$SYBASE/\$SYBASE_ASE/init/logs/srvbuildMMDD.NNN</code> where <i>MM</i> is the month, <i>DD</i> is the date, and <i>NNN</i> is a three-digit number identifying the srvbuild session |
| sqlloc sqllocres | <code>\$SYBASE/\$SYBASE_ASE/init/logs/sqllocMMDD.NNN</code> |
| sqlupgrade sqlupgraderes | <ul style="list-style-type: none"> <code>\$SYBASE/\$SYBASE_ASE/init/logs/sqlupgradeMMDD.NNN</code> <code>\$SYBASE/\$SYBASE_ASE/upgrade/upgrade.NNN</code> (a temporary file created by the upgrade process) |

For Sybase servers

Table 9-3 lists the default error log locations for each Sybase server.

Table 9-3: Error log locations for Sybase servers

| Server | Default error log path and file name |
|-----------------|--|
| Adaptive Server | <code>\$SYBASE/\$SYBASE_ASE/install/servername.log</code> |
| Backup Server | <code>\$SYBASE/\$SYBASE_ASE/install/servername_back.log</code> |
| Monitor Server | The directory from which Monitor Server is started. The error log file name is <i>ms.log</i> . |
| XP Server | Writes to the Adaptive Server error log. |

Solutions to common installation problems

If this section does not describe the problem you are experiencing, see the *Troubleshooting and Error Messages Guide*.

Unable to use X-Windows

If the setup and installation utilities do not display correctly, you may have to adjust the resolution on your monitor.

To change to a smaller font size issue the following UNIX commands:

```
% cd $SYBASE
% chmod +w xappdefaults
% cd xappdefaults
% chmod +w *
% foreach i(*)
? cat $i | sed -e "s/140/100/g" | sed -e "s/^#D/D/g" | sed -e "s/^#S/S/g" > p
? mv p $i
? end
%
```

The installation utilities will use approximately 25-percent less screen space.

Cannot eject the CD from the drive

If you cannot eject the CD from the drive:

- Check to see whether the CD drive path is the current directory (pwd) in a UNIX terminal window. If it is, change (cd) to another directory.
- Check for sybhelp processes. If these processes exist, kill them using the UNIX kill command.

DISPLAY environment variable not set correctly

If you run license manager and you get this error message:

The DISPLAY environment variable is not set correctly.

it means that the DISPLAY environment variable on the *remote* machine is not set correctly to display the user interface to your *local* machine.

To correct the problem, enter the following command at the UNIX prompt of the *remote* machine:

For C shell:

```
setenv DISPLAY host_name:0.0
```

For Bourne shell:

```
DISPLAY=host_name:0.0; export DISPLAY
```

where *host_name* is the name of the machine on which you want the Studio Installer interface to appear (that is, on your *local* machine).

Client not authorized to connect to server

If you run the license manager and you get this error message:

```
Xlib: connection to "host_name" refused by server
Xlib: Client is not authorized to connect to Server
xhost: unable to open display "host_name"
```

it means the *remote* machine does not have permissions to display the user interface on the *local* machine where you are working.

To correct the problem, enter the following command at the UNIX prompt of your *local* machine:

```
xhost +remote_machine
```

where *remote_machine* is the machine on which you are running the license manager.

Restart the license manager.

Address already in use

If you receive the following message from `srvbuild`:

```
kernel: ninit: bind, Address already in use
```

it means you entered a port number that is already in use.

To correct the problem, enter a different port number on the `srvbuild` screen. The command `netstat -a` produces a list of port numbers in use.

Adaptive Server failed to boot

The shared memory of the operating system may not be set high enough.

To correct the problem, see Chapter 4, “Installing Sybase Servers” for instructions on adjusting the shared memory value and start the installation or upgrade process again.

Cannot start XP Server

You may receive the following message from XP Server, when it is invoked by `xp_cmdshell` or some other extended stored procedure:

```
Msg 11018, Level 16, State 1:  
Procedure "xp_cmdshell", Line 2:  
XP Server must be up for ESP to execute.  
(return status = -6)
```

Verify that there is an XP Server entry in Adaptive Server's `syssservers` table. If you created XP Server separately (in a different `srvbuild` session) from Adaptive Server and you did not specify a related Adaptive Server, `srvbuild` cannot update the `syssservers` table.

Use `sp_addserver` to add an entry to the `syssservers` table.

Stopping Adaptive Server after a failure

If the installation or upgrade session fails after you start Adaptive Server, use the shutdown command:

- 1 Log on as "sa".
- 2 Shut down Adaptive Server using the shutdown with nowait command. Using the with nowait option stops the Adaptive Server immediately, without waiting for currently executing SQL statements to finish:

```
1> shutdown with nowait  
2> go
```

Recovering from a failed installation

If the installation does not succeed, the installation program displays error messages. Review the error messages and your Adaptive Server error log to determine the cause of the installation failure. For default error log locations, see Table 9-3.

If installation fails after the installer creates files

If the installation program quits while Adaptive Server is being configured, perform the following steps:

- 1 View the contents of the log file generated by Adaptive Server. For default error log locations, see Table 9-3.
- 2 Take any suggested actions to correct the problem.

- 3 If the installation fails *after* the installation program has created any operating system files, such as the master device or system procedures device files, delete those files.
- 4 If the installation fails *after* the installation program boots the Adaptive Server that you are attempting to install, shut down that server. Follow the procedure under “Stopping Adaptive Server” on page 102 before performing step 5.
- 5 Use `asecfg` to restart the configuration.

If Adaptive Server or SQL Server fails the pre-upgrade test, `sqlupgrade` displays:

```
Server SERVER_NAME failed preupgrade eligibility test. See log for more
information.
```

- 1 From the Upgrade window, select Exit.
- 2 Examine the log file created in the `$SYBASE/$SYBASE_ASE/init/logs` directory to find out why Adaptive Server or SQL Server failed the pre-upgrade eligibility test.

If the log contains messages about insufficient space in `sybsystemprocs`, follow the instructions in “Increasing the size of the `sybsystemprocs` database” on page 119 to correct the problem.

After you resolve any problems, shut down Adaptive Server or SQL Server and use `sqlupgrade` to complete the upgrade session.

Troubleshooting resource file installations

If you encounter problems during the build, configure, or upgrade of Adaptive Server using the `srvbuild[res]`, `sqlloc[res]`, or `sqlupgrade[res]`, it may be that these utilities did not allow enough time for Adaptive Server to shut down properly.

You can set the environment variable `SYBSHUTWAIT` to force the utilities to wait for Adaptive Server to shutdown. For example:

```
% setenv SYBSHUTWAIT 120
```

This forces the utility to wait for 2 minutes to allow Adaptive Server to shut down before proceed with the next task.

Recovering from a failed upgrade

If the upgrade process does not succeed, the installation program displays error messages. Review the error messages and the Adaptive Server error log to determine the cause of the upgrade failure. For default error log locations, see Table 9-3.

Restoring from backup

You may need to restore your databases due to a failed upgrade.

If you think the upgrade failure or its cause may have damaged your databases, restore the databases from backups. For information on restoring databases, see the *System Administration Guide*.

If you are concerned about the possible corruption of your databases, exit `sqlupgrade`, but do not attempt to restart the upgrade session until you have restored the databases from backup. After restoration is complete, retry the upgrade.

Rerunning the upgrade

Whether you can safely re-run the upgrade depends on when the failure occurred in the upgrade process. If the failure occurs while the message: “Starting to upgrade Adaptive Server” is displayed, it is safe to re-run the upgrade program.

- *First*, try to fix the problem that caused the upgrade to fail.
- *Then*, run the upgrade again.

If the upgrade fails:

- *Before* returning the message “Setting upgrade version to 12.5”, you may need to restore your latest database backup, and restart the upgrade.
- *After* returning the message “Setting upgrade version to 12.5”, it is not necessary to restart the upgrade. The installation utility considers the upgrade to be complete.

Also, it is not necessary to restore a database from a backup unless that database failed during the upgrade.

Recording the upgrade manually

If Adaptive Server did not finish recording the upgrade in the sysattributes table before the failure occurred.

- 1 Fix the problem that caused the failure.

The first error message indicates the cause of the failure. If you can, solve the problem and proceed to step 2. For example, an 1105 error can usually be corrected with a dump transaction command. However, more complex problems may have to be referred to Sybase Technical Support.

- 2 Execute the following SQL statements to allow Adaptive Server to complete recording the upgrade:

```
1> declare @dbname varchar(30)
2> select @dbname = min(name)
3> from sysdatabases
4> while @dbname is not null
5> begin
6> online database @dbname
7> select @dbname = min(name)
8> from sysdatabases
9> where name > @dbname
10> end
```

If the cause of the failure is known

If the error logs or messages clearly indicate the cause of failure, and you do not believe your databases were damaged, you can attempt to fix the problem and rerun the upgrade immediately.

- 1 Exit the sqlupgrade program.
- 2 Perform the necessary actions to fix the problem.

For example, if the error log indicates that the upgrade failed because your existing databases do not contain enough space, use the alter database command to increase the available space.

- 3 It may be necessary to shut down Adaptive Server. Follow the instructions for “Stopping Adaptive Server” on page 102.

Shutting down the server enables the installation program to boot the server and re-run the upgrade session.

- 4 Start sqlupgrade again.
- 5 Select Upgrade Adaptive Server, and proceed with the upgrade.

If the cause of the failure is unknown

If the upgrade fails again, and you cannot determine the cause of failure, check the error log file to find out when and where the upgrade failed, and contact Sybase Technical Support.

By default, the log file is located in
`$$SYBASE/$SYBASE_ASE/install/servername.log`.

Troubleshooting SySAM

When you boot Adaptive Server with SySAM support, problems acquiring licenses or contacting the asset management software appear in the Adaptive Server error log file, *lmgrd.log* (`$$SYBASE/$SYBASE_SYSAM/log`).

Table 9-4 describes the SySAM error messages and suggests ways to correct common problems.

Table 9-4: SySAM error messages

| Message | Description | Action |
|--|--|--|
| Use license file <i>filename</i> . | This message identifies the license file that Adaptive Server uses. | This is a normal start-up message; no action is required. |
| Feature <i>feature_name</i> is not licensed. | There is no license available for the specified optional feature. | You must purchase and install a valid license for the feature before it can be used. See Chapter 3, “Sybase Software Asset Management (SySAM)”. |
| There is no valid license for ASE server product. Server is booting with all the optional features disabled. | Adaptive Server cannot find a valid license file, or it cannot contact the asset management software to check out an ASE_SERVER license. | When Adaptive Server cannot check out an ASE_SERVER license, the server boots but does not enable any optional features. If you have purchased an ASE_SERVER license, make sure it has been installed. Also, make sure the license file exists and has been defined in LM_LICENSE_FILE. |
| License manager initialization fails. | The asset management software did not initialize. | Use “Starting the software manually” on page 37. |
| Feature <i>feature_name</i> is over drafted <i>n</i> times. | All licenses for the specified feature have already been checked out by other servers. | Adaptive Server enables the optional feature, even though all available licenses have been used by other servers. Contact your authorized Sybase reseller to purchase an additional license for the feature. |

| Message | Description | Action |
|---|---|--|
| Warning: there is no valid license for ASE_ <i>feature</i> and therefore <i>feature</i> is not initialized. | There is no license available for the configured feature. | Adaptive Server does not enable the optional feature until it can obtain a valid license. Purchase and install a license for the feature. Contact your authorized Sybase reseller. |

Authorization code input error

When you purchase licenses for Sybase Adaptive Server products, you are issued a Sybase Software Asset Management Certificate. The certificate has the following information for each product:

- Order Number
- Feature Name
- Feature Count
- Software Version
- Authorization Code
- Product Description

This information is used by SySAM to build the license file, with new licensed features appended to the end of the file. Here is a sample license file, *license.dat*:

```
SERVER server1 ANY 4100
VENDOR SYBASE $SYBASE/$SYBASE_SYSAM/bin/SYBASE
USE_SERVER

INCREMENT ASE_SERVER SYBASE 12.5 PERMANENT 1000
123456789123 SN=10001 OVERDRAFT=10000 ck=0

INCREMENT ASE_JAVA SYBASE 12.5 PERMANENT 1000
123456789123 SN=10001 OVERDRAFT=10000 ck=0

INCREMENT ASE_DTM SYBASE 12.5 PERMANENT 1000
123456789123 SN=10001 OVERDRAFT=10000 ck=0

INCREMENT ASE_HA SYBASE 12.5 PERMANENT 1000
123456789123 SN=10001 OVERDRAFT=10000 ck=0
```

- ASE_<FEATURE> is the feature name, such as ASE_SERVER, ASE_JAVA, etc.
- 12.5 is the version number.

- *Feature Count* immediately follows the license type, *PERMANENT*.
- *SN=10001* is the Order Number.
- *OVERDRAFT= ###* is the maximum licenses that can be checked out.
- *123456789123* is a 12-digit number representing the authorization code.

The authorization code is case sensitive. If you make a mistake while entering the authorization code, correct it by accessing the license file with a text editor, making the necessary changes, and saving the file.

The file is located in `$SYBASE/$SYBASE_SYSAM/licenses/license.dat`.

Warning! Tampering with any portion of the licenses file other than the authorization code invalidates the license.

Alternative Installation Methods

Resource file installation

You can create an Adaptive Server or Backup Server, using values specified in a resource file that defines the attributes for the server.

Resource files are ASCII format template files that contain configuration variables. To use the template files, edit the resource file, replace variables with desired values, and execute the `svrbuidres` utility. The `svrbuidres` utility uses the variables to create servers.

Resource files allow you to create servers in batch mode. You can create servers on multiple systems from the same resource file by editing a few values in that file. Within a single resource file, you can create only one server at a time.

The server files must already exist on the computer before the servers can be created. Use procedures in Chapter 4, “Installing Sybase Servers”.

Note You cannot install Monitor Server, XP Server or FTS using resource files.

To install an Adaptive Server or Backup Server using resource files:

- 1 Edit a resource file as described in “Editing a resource file” on page 163.
- 2 Execute the `svrbuidres` utility, using the edited resource file as described in “Using `svrbuidres`” on page 166.

Editing a resource file

You can edit a resource file by:

- Editing a sample resource file

- Editing the resource file created by the `srvbuild` utility

Editing a sample resource file

Sample resource files for creating Adaptive Server and Backup Server are included in your Adaptive Server distribution in:

`$$SYBASE/$$SYBASE_ASE/init/sample_resource_files.`

Edit these files to specify the attributes for the servers you want to create.

Note the following:

- If you use `USE-DEFAULT` as the Adaptive Server name, the utility substitutes the name of the computer on which it is running for the server name. If you use `USE-DEFAULT`, do not create Backup Server entries in the resource file. They are created automatically using the default name.
- To create a server with a different name, change the Adaptive Server name and use the new name with the “_BK” extensions for Backup Server. For example, the Backup Server for `PIANO` should be `PIANO_BK`.
- The master device size must be at least 30MB.
- The system procedure device and sybsystemprocs database sizes must be at least 100MB.
- The auditing entries in the resource file are not supported and will not be processed by the utility. To enable auditing, see the *Adaptive Server Configuration Guide*.

Editing a resource file created by `srvbuild`

You can create a resource file by running the `srvbuild` utility. A new file containing the values you specified for the Adaptive or Backup servers is written to:

`$$SYBASE/$$SYBASE_ASE/init/logs/srvbuildMMDD.VVV-servername.rs.`

Edit these files to specify the attributes for the servers you want to create.

Warning! Be sure to edit the resource file generated by `srvbuild` to change attributes, such as device names, that may be different on another system.

Resource file attributes for Adaptive Server

Table A-1 shows the Adaptive Server resource file attributes, their default values, and other options.

The attributes in bold type are *required*.

All alpha values are case sensitive.

The prefix of the attribute name varies, depending on whether the resource file was created by `srvbuild` or by `sybinit` (from a pre-11.5 SQL Server). The prefix is ignored by `srvbuildres` when processing the resource file.

The attribute names in Table A-1 are the attribute names provided in the sample resource file included in your Adaptive Server distribution.

Table A-1: Resource file attributes for Adaptive Server

| Attribute | Default Value [Other Options] |
|---|--|
| <code>sybinit.release_directory</code> | The value of <code>\$\$SYBASE</code> at your site |
| <code>sybinit.product</code> | <code>sqlsrv</code> |
| <code>sqlsrv.server_name</code> | <i>server_name</i> |
| <code>sqlsrv.new_config</code> | yes (required value) |
| <code>sqlsrv.do_add_server</code> (to interfaces file) | yes [no] |
| <code>sqlsrv.network_protocol_list</code> | <code>tcp [spx]</code> |
| <code>sqlsrv.network_hostname_list</code> | <i>hostname</i> |
| <code>sqlsrv.network_port_list</code> | <i>port_number</i> |
| <code>sqlsrv.master_device_physical_name</code> | <i>path_and_name_of_master_device</i> |
| <code>sqlsrv.master_device_size</code> | 30 (MB) |
| <code>sqlsrv.master_database_size</code> | 5 (MB) |
| <code>sqlsrv.errorlog</code> | <code>\$\$SYBASE/\$\$SYBASE_ASE/install/server_name.log</code> |
| <code>sqlsrv.do_upgrade</code> | no (required value) |
| <code>sqlsrv.sybsystemprocs_device_physical_name</code> | <i>path_and_name_of_sybsystemprocs_device</i> |
| <code>sqlsrv.sybsystemprocs_device_size</code> | 90 (MB) |
| <code>sqlsrv.sybsystemprocs_database_size</code> | 90 (MB) |
| <code>sqlsrv.sybsystemdb_device_physical_name</code> | <i>path_and_name_of_sybsystemdb_device</i> |
| <code>sqlsrv.sybsystemdb_device_size</code> | 5 (MB) |
| <code>sqlsrv.sybsystemdb_database_size</code> | 5 (MB) |
| <code>sqlsrv.default_backup_server</code> | <i>server_name_back</i> |

Resource file attributes for Backup Server

Table A-2 shows the Backup Server resource file attributes, their default values, and other options.

The attributes in bold type are *required*.

All values are case sensitive.

The prefix of the attribute name varies, depending on whether the resource file was created by `srvbuild` or by `sybinit` (from a pre-11.5 SQL Server). The prefix is ignored by `srvbuildres` when processing the resource file.

The attribute names in Table A-2 are the attribute names provided in the sample resource file included in your Adaptive Server distribution.

Table A-2: Resource file attributes for Backup Server

| Attribute | Default value [other options] |
|---|--|
| sybinit.release_directory | \$SYBASE |
| sybinit.product | bsrv |
| bsrv.server_name | <i>server_name_back</i> |
| bsrv.do_add_backup_server (to interfaces file) | yes [no] |
| bsrv.network_protocol_list | tcp [spq] |
| bserv.network_hostname_list | <i>hostname</i> |
| bsrv.network_port_list | <i>port_number</i> |
| bsrv.language | us_english [chinese, french, german, japanese, spanish] |
| bsrv.character_set | iso_1[iso_1, cp850, cp437, dekanji, ascii_8, eucgb, eucjis, mac, roman8, sjis, utf8] |
| bsrv.tape_config_file | \$SYBASE/\$SYBASE_ASE/ <i>backup_tape.cfg</i> |
| bsrv.errorlog | \$SYBASE/\$SYBASE_ASE/ <i>server_name_back.log</i> |

Using `srvbuildres`

To execute `srvbuildres` at the UNIX prompt enter:

```
$SYBASE/$SYBASE_ASE/bin/srvbuildres -r resource_file
```

where *resource_file* specifies the resource file containing the attributes that describe the server to build.

Note Unlike `srvbuild`, the `srvbuildres` utility does not require any X libraries.

When you have completed resource file installation, go to Chapter 5, “Post-Installation Tasks”.

Troubleshooting resource file installations

If you encounter problems during the build, configure, or upgrade of Adaptive Server using the `srvbuild[res]`, `sqlloc[res]`, or `sqlupgrade[res]`, it may be that these utilities did not allow enough time for Adaptive Server to shut down properly.

You can set the environment variable `SYBSHUTWAIT` to force the utilities to wait for Adaptive Server to shutdown. For example:

```
% setenv SYBSHUTWAIT 120
```

This forces the utility to wait for 2 minutes to allow Adaptive Server to shut down before proceed with the next task.

Creating Raw Partitions

Adaptive Server 12.5 is certified and released on kernel 2.2.14. If you choose to configure your system to use raw partitions on a multi-CPU system, you may get better results with Linux kernel release level 2.2.18, since that is the kernel release that is most stable with raw disk partitions.

Adaptive Server version 12.5.0.1 has been released on kernel 2.4.7-10 in addition to kernel 2.2.14. Sybase recommends that you upgrade to this new release since 2.4.7-10 provides:

- More stable support for raw devices.
- Support for large file systems.
- Support for asynchronous I/O.
- Support for LDAP.
- Backup directly to tape devices.

Overview

Adaptive Server 12.5 allows you to create and mount its database devices on raw bound devices for raw disk I/O. Raw disk I/O has performance advantages since it enables direct memory access from user address space to the physical sectors on the disk, omitting needless memory copy operations from the user address space to the kernel buffers.

Raw disk I/O also assumes that logical and physical I/O are simultaneous, and writes are guaranteed to flush to the disk when the system `write()` call returns.

For Red Hat 6.2 distributions, you must use a distribution with kernel 2.2.18, at minimum. Correct kernel errata packages can be downloaded from the Red Hat Web site at <http://www.redhat.com>.

When preparing a raw partition device, follow these guidelines:

- Do not initialize a database device on the partition that contains your Sybase installation software. Doing so destroys all existing files on that partition.
- A raw partition designated for use by Sybase cannot be mounted for use by the operating system for any other purpose, such as for file systems or swap space.
- After a Sybase configuration utility or the disk init command has initialized a portion of a partition as a database device, the entire partition cannot be used for any other purpose. Any space left on the partition beyond the size specified for the device becomes inaccessible, unless you drop and re-create the device.
- For best performance, place the Sybase software and all disk devices, including the master device, on the same machine.
- To avoid any possibility of using a partition that contains the partition map, do not use cylinder 0.
- Place the database device on a character device, because the Adaptive Server recovery system needs unbuffered system I/O.

To determine whether a device is a block device or a character device, run this command on the */dev/raw* directory:

```
ls -l
```

“b” (block) or “c” (character) appears in the left-most column.

❖ **Choosing a raw partition**

- 1 Determine which raw partitions are available.
- 2 Determine the sizes of the raw partitions.
- 3 From the list of available raw partitions, select a raw partition for each device, based on the size recommendations in Table 2-13 on page 26.
- 4 Verify with the operating system administrator that the partition you have chosen is available.
- 5 Make sure the “sybase” user has read and write privileges to the raw partition.

Note For more information on choosing a raw partition, see your operating system documentation.

Creating the partitions

Specific system administration is required before raw devices can be enabled and used. The available tools to configure devices depend on the distribution.

Physical disk space needs to be allocated in partitions on the disks where you want to set up raw devices. The physical I/O subsystem can be on either SCSI or EIDE devices.

Note Creating the partitions can be done with the Linux default `fdisk(8)` utility. You must have “root” privileges to use the command `fdisk`. Refer to the `fdisk(8)` man pages for a complete description of the command.

This example shows how to set up partitions as raw devices, on four SCSI disks in the system—`sda`, `sdb`, `sdc`, and `sdd`.

- Start `fdisk` on `/dev/sdd`:
`# fdisk /dev/sdd`

The system returns:

```
The number of cylinders for this disk is set to 8683
....
Command (m for help):
```

- Enter “p” to print the current partition lay-out. The output is:

```
Disk /dev/sdd: 64 heads, 32 sectors, 8683 cylinders
Units = cylinders of 2048 * 512 bytes
Device Boot  Start    End  Blocks   Id  System
/dev/sdd1     1     7499   7678960   83  Linux
/dev/sdd2   7500    8012   525312    82  Linux swap
/dev/sdd4   8013    8683   687104    5  Extended
```

This sample shows the extended partition (`sdd4`) has 687104 free blocks, starting from 8013 and ending at 8683. The remaining partitions can be assigned later.

The next example assigns an additional partition for raw bound disk I/O:

- 1 Use the `n` command to create a new partition, and enter “l” at this prompt for logical.

```
Command (m for help):n
Command action
l   logical (5 or over)
p   primary partition (1-4)
```

- 2 Accept the default by pressing Enter on the keyboard when you are prompted:

```
First cylinder (8013-8683, default 8013):
```

- 3 Accept the default by pressing Enter on the keyboard when you are prompted:

```
Last cylinder or +size or +sizeM or +sizeK (8013-8683, default 8683): 8269
```

- 4 Use the `t` command, enter “5” at this prompt:

```
Partition number (1-8):5
```

- 5 Enter “60” at this prompt:

```
Hex code (type L to list codes): 60
```

- 6 The output is:

```
Changed system type of partition 5 to 60 (Unknown)
```

- 7 Repeat the above steps to create four partitions for raw device I/O.

Verify the setup using `p` to print the full partition table before writing it out. Make sure that there are no overlapping partitions and the type for the unassigned partitions is Unknown type 60.

The partition table can now be written to disk and you can quit the `fdisk(8)` utility.

Before continuing, restart your system to verify that the kernel running is in sync with the new disk partition layout. This is reported during system start when probing for disk spindles.

```
# /sbin/reboot
```

Binding the partitions for raw disk I/O

Red Hat 6.2 is equipped with the administration tools to setup raw devices but does not setup the raw device bindings during system start. For those environments a script named `sybdevices.sh` is available from `$SYBASE/install/sybdevices.sh`. This script contains four arguments:

- `show` – all dedicated partitions
- `create` – all dedicated raw device entries

- bind – all dedicated raw device entries to their partitions
- list – all dedicated raw device bindings

Follow below steps to administrate your system allowing Adaptive Server to initialize and gain direct access to the device mappings.

- 1 Execute `$SYBASE/install/sybdevices.sh`, and enter:

```
show
```

The output should be:

```
/dev/sdd5      88326 blocks
/dev/sdd6      88326 blocks
/dev/sdd7      48163 blocks
/dev/sdd8      40131 blocks
```

- 2 From the same directory, enter:

```
create
```

The output should be:

```
/dev/raw/raw1 c 162 1
/dev/raw/raw2 c 162 2
/dev/raw/raw3 c 162 3
/dev/raw/raw4 c 162 4
```

- 3 From the same directory, enter:

```
bind
```

The output should be:

```
/dev/raw/raw1:      bound to major 3, minor 8
/dev/raw/raw2:      bound to major 3, minor 9
/dev/raw/raw3:      bound to major 3, minor 10
/dev/raw/raw4:      bound to major 3, minor 11
```

- 4 From the same directory, enter:

```
list
```

The output should be:

```
/dev/raw/raw1:      bound to major 3, minor 8
/dev/raw/raw2:      bound to major 3, minor 9
/dev/raw/raw3:      bound to major 3, minor 10
/dev/raw/raw4:      bound to major 3, minor 11
```

- 5 To guarantee that the raw device binding occurs on any system restart, open an editor on the following file:

```
# vi /etc/rc.d/rc.local
```

- 6 Add the following commands to the end of this file:

```
export SYBASE=`rpm -q --queryformat '%{installprefix}' sybase-common-12.5`  
  
$SYBASE/install/sybdevices.sh bind  
$SYBASE/install/sybdevices.sh list
```

Red Hat 7.2 is fully equipped with the administration tools to set up raw devices and administrate them during system start. To set up a Red Hat 7.2 system the partitions are administrated in the file */etc/sysconfig/rawdevices*.

This is a plain text file contains comments and examples for possible configurations, as follows:

```
# raw device bindings  
  
# format: rawdev major minor  
#          rawdev blockdev  
# example: /dev/raw/raw1 /dev/sda1  
#          /dev/raw/raw2 8 5  
/dev/raw/raw1 /dev/sdd1  
/dev/raw/raw2 /dev/sdd2  
/dev/raw/raw3 /dev/sdd3  
/dev/raw/raw4 /dev/sdd4
```

- 1 Once created, you need to bind the raw devices. This can be done by starting them from */etc/rc.d/init.d/rawdevices*.

```
[root@legolas init.d]# cd /etc/rc.d/init.d  
[root@legolas init.d]# sh rawdevices start
```

Assigning devices:

```
/dev/raw/raw1 --> /dev/sdd5  
/dev/raw/raw1: bound to major 3, minor 5  
/dev/raw/raw2 --> /dev/sdd6  
/dev/raw/raw2: bound to major 3, minor 6  
/dev/raw/raw3 --> /dev/sdd7  
/dev/raw/raw3: bound to major 3, minor 7  
/dev/raw/raw4 --> /dev/sdd8  
/dev/raw/raw4: bound to major 3, minor 8
```

done

- 2 To guarantee that the raw device binding occurs during any restart, use the `chkconfig(8)` utility.

```
# /sbin/chkconfig rawdevices on
```

Accessing raw devices from Adaptive Server

Once the partitions are created and the devices are bound for raw disk I/O, Adaptive Server is able to use them. If Adaptive Server is running as user `sybase`, then apply read, write, and owner permissions to the `/dev/raw/raw#` device entries and the raw bound control device `/dev/rawctl` for user `sybase`.

See the `chown(1)`, `chgrp(1)`, and `chmod(1)` commands to apply correct permissions.

- To verify your settings, query the device binding using the `raw` command, enter:

```
$ raw -qa
```

The output should be:

```
/dev/raw/raw1: bound to major 3, minor 5
/dev/raw/raw2: bound to major 3, minor 6
/dev/raw/raw3: bound to major 3, minor 7
/dev/raw/raw4: bound to major 3, minor 8
```

- To verify your settings, query the device binding using the permissions command, enter:

```
<username@servername.sybase.com:4> raw -qa
```

The output should be:

```
Cannot open master raw device '/dev/rawctl'
(Permission denied)
```

Using the raw devices, Adaptive Server and the installation and configuration utility, `srvbuild`, detects and presents their size automatically. Just enter the absolute path to the raw device when creating the master, `sybtempprocs`, `sybtempdb`, or any other device.

Index

A

- Adaptive Server
 - base license 30
 - clients 2
 - communicating with Backup Server 58
 - database device requirements 26
 - described 2
 - disk space requirements 18
 - installation directory structure 8
 - installing multiple 163
 - installing via resource file 163
 - performance statistics for 2
 - preparing to upgrade 114
 - RAM required by 13
 - shutting down 102
 - starting from UNIX command line 100
 - starting with operating system 101
 - System Administrator 71
 - verifying environment for upgrade 110
- additional licenses
 - SySAM Certificate 38
- administrator
 - operating system 1
 - Sybase system 1, 44
- Advanced Security Mechanism
 - authentication and encryption 30
 - enabling 32
 - encryption using DCE and CyberSafe 30
- alter database command 119
- applications
 - changing after renaming objects 112
 - testing current 111
- auditing
 - feature 140
 - recommended database device size 27
- auditinit utility 25
- authentication
 - advanced security mechanism 30
 - encryption 30

- authorization code
 - troubleshooting 161

B

- Backup Server
 - character set for 57
 - communicating with Adaptive Server 58
 - default language for 57
 - described 3
 - when to install 3
 - installing via resource file 163
 - network connections for 57
 - server connections 58
 - specifying message language for 57
 - starting from UNIX command line 100
 - starting with operating system 101
 - using tape with 3
- base license, Adaptive Server 30
- batch installation of servers 163
- bcp utility 6
- buffer specifications 21

C

- character devices
 - located in /dev directory 170
- character sets
 - changing 78
 - for Backup Server 57
- client products
 - description of 2
- configuration file, network 95
- conflicting object names 113
- connection problems, analyzing 45
- connectivity 6
- conventions, used in this manual 14
- create database command, system tables created by 23

Index

current applications, testing 111

CyberSafe

advanced security mechanisms 30

encryption and authentication 30

D

database devices

See also individual device names 46

72, 80

master 23

raw partitions for, preparing 169

sysprocsdev 23, 24

databases

dbccdb 26

image data for sample 73

interpubs 20, 74

jpubs 20, 75

master 23, 24

model 23

optional 72

pubs2 20

pubs3 20

sample 20, 25, 75

sizes of 19, 20

specifications 19

suspect 115

sybsecurity 25

sybsyntax 25

sybsystemprocs 19, 24

tempdb 23

dataserver utility 105

dbcc checkalloc command 114

dbcc checkcatalog command 115

dbcc checkdb command 114

dbcc checkstorage, database for 26

dbcc checkverify command 115

dbccdb database 26

DCE

advanced security mechanisms 30

encryption and authentication 30

default

character set for Backup Server 57

language for Backup Server 57

directories

installation 8, 86

specifying new, for installation 110

disk init command 170

Distributed Transaction Management

enabling 32

recovery 30

documentation, online 64

drivers, ODBC 7, 77, 86

dscp utility 4, 93

dsedit utility 4, 93, 140

DSLISEN environment variable 68, 69

DSQUERY environment variable 68, 69

dump database command, backing up with 115

dump requests 3

dump tran command 115

Dynatex 64

DynaWeb 64

E

enabling features 32

advanced security mechanisms 32

distributed transaction management 32

Java 32

encryption

advanced security mechanisms 30

authentication 30

environment variables

DSLISEN 68, 69

DSQUERY 68, 69

PATH 68, 69, 85

SYBASE 8, 68, 69

TEMP 86

error log files, troubleshooting and 156

errors

resolving 115

extended stored procedures (ESPs) 4

F

feature licenses

adding additional licenses 38

troubleshooting 161

features

- licensed through SySAM 30
- files
 - runserver 100
 - shared memory 105
- french sample database 25

G

- german sample database 25
- guest user
 - in pubs2 and pubs3 databases 75

H

- hardware requirements 19
- Historical Server
 - description of 2
- host name 58
 - License Manager Host 36
 - primary network node 35

I

- identifiers, conflicting with reserved words 112
- InfoMaker
 - described 7
- installation directory
 - for unloading products 44
 - structure of 8
- installation utilities 4
- installing
 - multiple Adaptive Server versions 69
 - Sybase products 67, 73
- installing 11.9.2
 - Open Client/C for Windows NT/95 86
- installing documentation files using RPM 64
- installing Java 89
- installing Sybase products 74
- installintpubs script 74
- installjpubs script 75
- installpix script 73
- installpubs2 script 73
- installpubs3 script 73

- interfaces file 91, 140
- interfaces file entry attribute for custom installation
 - 58, 60
- interpubs sample database 25, 74
- isql
 - checking for replicated databases 125
 - determining Adaptive Server release level 110
 - verifying which servers are running 69

J

- japanese sample database 25
- Java
 - enabling 32
 - in Adaptive Server databases 30
 - XML classes and 30
- Java Development Kit 89
- Java Runtime Environment 89
- jpubs sample database 25, 75

K

- keywords, reserved by Sybase 133
- kill command 104, 105

L

- language functions, procedural 4
- language modules 3, 6, 78
 - disk space required for 18
 - for Backup Server 57
 - memory requirements for 78
 - sizes of 19
- LDAP
 - enabling 92
 - specifying in libtcl.cfg 92
- LDAP libraries
 - environment variables 93
 - location of 93
- ldapurl
 - defined 92
 - example 92
- libtcl*.cfg file 92

Index

- format of 92
- location of 92
- license files 30
 - sample license file 161
- license management daemons
 - lmgrd 31
 - lmutil 31
- License Manager Host, see host name 36
- License Manager Port, see port number 36
- licensed features 30
- licenses management daemons
 - SYBASE 31
- listen port number, finding 95
- lmgrd daemon
 - starting the 37
 - SYBASE daemon 37
- load requests 3
- localization
 - disk space needed for 57
- login
 - “root” 1
 - “sa” 1
 - “sybase” 1, 44
 - administrator, Windows NT 87
 - superuser 1
- loopback mode 45

M

- mapping NLWNSCK to TCP 95
- master database 23
- master device 23, 26
- model database 23
- Monitor Server
 - described 2
 - description of 2
 - disk space required for 18
 - monserver command 101
 - shutting down 104
 - starting from UNIX command line 100
 - starting with operating system 101
- monserver command 101
- multiple Adaptive Server products
 - SySAM support for 30
- multiple servers

- installing in network asset management 35

N

- Named Pipe network protocol 6
- Net-Library 6
- network asset management
 - installing multiple servers 35
- network configuration file 95
- network connections
 - Backup Server 57
 - Monitor Server 59
- network protocols
 - Named Pipes 95
 - SPX 13
 - TCP 95
 - TCP/IP 13
 - Winsock 95
- network software configuration, verifying 44, 45

O

- objectid.dat
 - location of 93
- objects
 - conflicting names of 113
 - renaming 113
- ODBC 7, 77, 86
- online documentation 64
- online syntax help 80
- Open Client
 - disk space requirements 18
- Open Client applications 4
- Open Client/C
 - for Windows NT/95 6, 86
- Open Database Connectivity. See ODBC 77, 86
- Open Server applications 3
- operating system
 - administrator 1
 - kernel patches 43
 - starting servers 101
- optional databases 72
- optional features
 - authorization code input error 161

- enabling all 33
- sample license file 161
- turning on and off 32
- ownership, of files and directories 44

P

- parent process identification (PPID) 105
- partition map, avoiding damage to 170
- password, setting System Administrator 71
- PATH environment variable 68, 69, 85
- PID. See process identification 105
- plug-ins
 - installation requirements for 17
- port number 35, 36
- port number, specifying 58
- PowerDesigner, described 8
- PowerDynamo
 - described 7
- PPID. See parent process identification 105
- preparing
 - Adaptive Server for upgrading 114
- primary network node 35
 - host name 35
 - port number 35
- privileges 170
- procedural language functions 4
- Procedure specifications 21
- process identification (PID) 105
- Product Manuals Web site 64
- pubs2 sample database 25
- pubs3 sample database 25

Q

- query specifications 21

R

- raw partitions 46
 - availability of 170
 - minimum size for database devices 170
- recovering from a failed installation 156

- recovery
 - distributed transaction management 30
- reenable Replication Server 139
- referential integrity constraint 21
- renaming objects 113
- replicated databases
 - checking for 125
 - restoring 129
- Replication Server 139
 - draining logs for Adaptive Server 126, 139, 140
- reserved word conflicts
 - checking for 114
- Reserved words
 - checking for conflicts 114
- reserved words
 - using quoted identifiers 114
- Resource files 134, 163
 - installation using 163
- RPM
 - documentation files, installing 64
- runserver file 100

S

- xd2 saxd3 user 71
- sample databases 72, 75
- scripts
 - installintpubs 74
 - installjpubs 75
 - installpix 73
 - installpubs2 73
 - installpubs3 73
- server connections for Backup Server 58
- setup utilities 4
- shared memory files 105
- showserver command 69, 136
- shutdown command 103
- size
 - interpubs database, minimum 20
 - jpubs database, minimum 20
 - language module 19
 - pubs2 database, minimum 20
 - pubs3 database, minimum 20
 - recommended for database devices 119
 - sybsystemprocs database, minimum required for

Index

- upgrade 19
 - sort orders
 - changing 78
 - sp_configure
 - enable all features at boot up. 33
 - SPX network protocol 13
 - SQL Modeler. See PowerDesigner 8
 - SQL Remote
 - described 8
 - sql.ini file 93
 - sqldedit utility 6
 - sqlloc utility 4
 - sqllocres utility 4
 - sqlupgrade utility 4, 133
 - sqlupgraderes utility 4, 134
 - srvbuild utility 5, 93
 - srvbuildres utility 5
 - Starting servers
 - from UNIX command line 100
 - with operating system 101
 - startserver utility 100
 - Stored procedures
 - testing 111
 - stored procedures
 - changing after renaming objects 112
 - recommended database device size if adding user-defined 119
 - suspect databases 115
 - Sybase Central, Adaptive Server plug-in to 2
 - Sybase DynaWeb Web server 64
 - SYBASE environment variable 8, 68, 69
 - Sybase installation directory 8, 44, 131
 - Sybase products
 - InfoMaker 7
 - installing 67, 73, 74
 - language modules 6
 - Net-Library 6
 - Open Client/C for Windows NT/95 6, 86
 - PowerDesigner 8
 - PowerDynamo 7
 - running during installation 111
 - sybsecurity
 - database 25
 - device 25, 27
 - sybsetup utility 5
 - sybsyntax database 25, 80
 - sybssystemdb
 - device for 27
 - sybssystemprocs
 - database 24
 - sybssystemprocs database 24
 - preparing for upgrade 120
 - Syntax
 - database of 25
 - online help for 25
 - syntax
 - conventions for, in this manual 14
 - SySAM 30
 - configuration 33
 - defined 29
 - features licensed through 30
 - licenses files 30
 - multiple license files 30
 - networked servers 33
 - standalone 33
 - starting manually 37
 - support for multiple Adaptive Server products 30
 - troubleshooting 160
 - SySAM Certidicate
 - sample license file 161
 - SySAM Certificates
 - adding additional licenses 38
 - sysprocsdev device
 - minimum size for 26
 - purpose of 23, 24
 - sysprocsdev device, preparing for upgrade 120
 - System Administrator
 - setting password 71
 - System databases
 - where stored 26
 - system procedures
 - storage location of 24
- ## T
- table specifications 20
 - tape, using with Backup Server 3
 - TCP/IP network protocol 13, 95
 - TEMP environment variable 86
 - tempdb database 23
 - transaction log, space in 115

- transport type menu 58
- troubleshooting
 - by using error log files 156
 - common installation problems and solutions 152
 - feature licenses 161
 - SySAM 160
 - SySAM authorization codes 161

U

- UNIX
 - system requirements 16
- utilities
 - bcp 6
 - dsedit 6, 140

W

- Windows 95 and NT
 - client product installation on 86
 - plug-in installation on 18
- Windows Socket network protocol 6

X

- XML classes
 - and Java classes in 30
- XP Server
 - described 4
 - starting 99
 - starts automatically with first ESP request 96
- xp_cmdshell command 99

