

SYBASE®

Installation Guide

Adaptive Server® Enterprise

15.0.2

[Linux]

DOCUMENT ID: DC30119-01-1502-01

LAST REVISED: June 2007

Copyright © 1987-2007 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 trademarks can be viewed at the Sybase trademarks page at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

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., One Sybase Drive, Dublin, CA 94568.

Contents

About This Book	ix	
CHAPTER 1	Overview	1
	Product licensing with SySAM 2.0.....	2
	Product components	3
	Adaptive Server Enterprise	3
	Backup Server.....	4
	Stored procedures	4
	The client.....	4
	Adaptive Server editions	5
	Enterprise Edition	6
	Small Business Edition	7
	Developer Edition	7
	Express Edition	7
	Adaptive Server specifications	8
	Product descriptions and directory layout	9
	Sybase support for OLE DB, and ODBC.....	11
	PC-Client product descriptions and directory layout	12
	Separately installable PC-Client products	14
CHAPTER 2	Installing Adaptive Server	15
	Overview	15
	System requirements	16
	Pre-installation tasks for SySAM	17
	Accessing SPDC	18
	Getting your host ID	19
	Generating licenses at SPDC.....	20
	Installing a new license server	23
	Pre-installation tasks for Adaptive Server	24
	Installing server components	28
	Installation methods	28
	Installing components with the Installer in GUI mode	28
	Installing version 15.0.2 over an existing 15.x Adaptive Server	36

	Determining Adaptive Server version.....	36
	Backing up Adaptive Server	37
	Installing the Adaptive Server version 15.0.2 binary overlay...	37
CHAPTER 3	Post-Installation Tasks.....	39
	Verifying that servers are running	39
	Verifying that you can connect to servers	40
	Connecting to Adaptive Server via Sybase Central	40
	Setting the System Administrator password	41
	The SAMreport.....	42
	Installing sample databases.....	42
	Default devices for sample databases	43
	interpubs database.....	44
	jpubs database.....	45
	Maintaining the sample databases.....	46
CHAPTER 4	Installing Sybase PC-Client Products.....	47
	Overview	47
	Before you install the PC-Client products	48
	Installing PC-Client products.....	48
	Installing PowerDesigner, PowerTransfer, and InfoMaker	50
	Configuring network connections for client products.....	51
	Configuring libtcl.cfg for LDAP	51
	Adding a server to the interfaces file	52
	Testing the Sybase Central installation	53
CHAPTER 5	Upgrading Adaptive Server	55
	Overview of the upgrade process	55
	System catalog changes during upgrade	58
	Pre-upgrade tasks.....	58
	Checking system and upgrade requirements.....	60
	Checking RUN_server file location.....	60
	Procedure text is required for upgrade.....	61
	Reserved words	61
	Preparing the database and devices for the upgrade	63
	Upgrading to Adaptive Server 15.0.2	65
	Using sqlupgrade	66
	Using sqlupgraderes	68
	Post-upgrade tasks	70
	Run the instmsgs.ebf script.....	70
	Restoring functionality in Adaptive Server.....	71
	Reenabling Replication Server.....	74

	Reenabling auditing.....	76
	Shutting down SySAM servers.....	77
	Upgrading Backup Server, Monitor Server, and XP Server.....	77
	Upgrading Job Scheduler.....	79
	Upgrading Job Scheduler templates.....	80
	Upgrading high availability.....	84
	Migrating from 32-bit to 64-bit versions on Different Computers....	86
	Recovering from a failed upgrade.....	88
CHAPTER 6	Downgrading from Adaptive Server 15.0.2.....	91
	Downgrading from 15.0.2 release.....	91
	Basic downgrade steps.....	92
	Post-downgrade tasks.....	94
	System stored procedures.....	94
	System messages.....	94
	Unrecognized tokens or objects.....	95
	System tables.....	95
	Configuration parameters.....	95
	Keywords.....	96
	Additional steps required if you used 15.0.2 features before downgrading.....	96
	Real-Time Messaging.....	96
	New sort orders.....	97
	Native XML.....	97
	Instead of trigger.....	97
	SQL user defined functions.....	97
	Automatic database expansion.....	98
	Replication issues with downgrade.....	98
	Downgrading Job Scheduler.....	99
	Returning to Adaptive Server version 15.0.2.....	99
CHAPTER 7	Troubleshooting SySAM Issues.....	101
	Calling Sybase technical support for SySAM-related issues.....	101
	Where to look for SySAM-related errors.....	102
	Troubleshooting SySAM.....	103
	SySAM grace mode.....	108
	SySAM grace period.....	109
	Enabling and changing e-mail notifications.....	109
CHAPTER 8	Troubleshooting.....	111
	Overview.....	111
	Error log locations.....	112

	Solutions to common installation problems.....	113
	Cannot use X-Windows.....	113
	Cannot eject the CD from the drive.....	114
	DISPLAY environment variable not set correctly.....	114
	Client not authorized to connect to server.....	114
	Address already in use.....	115
	Adaptive Server fails to start.....	115
	The Installer fails to start.....	115
	Cannot start XP Server.....	115
	Stopping Adaptive Server after a failure.....	116
	Recovering from a failed installation.....	116
	Troubleshooting resource file installations.....	117
	If Adaptive Server fails the pre-upgrade eligibility test.....	117
	Recovering from a failed upgrade.....	118
	Upgrading compiled objects with dbcc upgrade_object.....	121
	Finding compiled object errors before production.....	122
	Using dbcc upgrade_object.....	125
	Using database dumps in upgrades.....	127
	Determining whether a compiled object has been upgraded.....	128
CHAPTER 9	Removing Adaptive Server.....	129
	Uninstalling Adaptive Server 15.0.2.....	129
	Removing an existing Adaptive Server.....	130
APPENDIX A	Alternative Installation Methods.....	131
	Installation on consoles (non-GUI mode).....	131
	Silent installation.....	132
	Configuration using resource files.....	134
	Configuring from a resource file.....	134
	Using srvbuildres.....	137
APPENDIX B	Creating Raw Partitions.....	139
	Overview.....	139
	Creating the partitions.....	140
	Red Hat Raw device administration.....	142
	SuSE raw device administration.....	143
	Accessing raw devices from Adaptive Sever.....	144
APPENDIX C	Upgrading Servers with Replicated Databases.....	147
	Introduction.....	147
	Suspending transaction processing and replication activities.....	148
	Draining the transaction logs for primary databases.....	149

	Draining the RSSD transaction log	151
	Disabling the secondary truncation point	153
APPENDIX D	Using sybssystemprocs	155
	Introduction	155
	Increasing the size of the sybssystemprocs database.....	156
	Increasing device and database capacity for system procedures	157
	Index	161

About This Book

The *Installation Guide for Sybase Adaptive Server Enterprise* for 15.0.2 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
- Instructions on how to install Adaptive Server with SySAM 2.0.

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 an overview of Adaptive Server, and describes the Enterprise Edition, Small Business Edition, Developer Edition, and Express Edition (on Linux) of Adaptive Server. It also provides product descriptions and directory layout information.
- Chapter 2, "Installing Adaptive Server," describes pre-installation procedures, SySAM related tasks, and how to install server components.
- Chapter 3, "Post-Installation Tasks," describes how to verify that servers are running, connect to servers, set passwords, install SAMreport, and install sample databases.
- Chapter 4, "Installing Sybase PC-Client Products," describes how to install Adaptive Server client products like the Adaptive Server plug-in to Sybase Central.

-
- Chapter 5, “Upgrading Adaptive Server,” describes how to upgrade an existing Adaptive Server to the current version.
 - Chapter 6, “Downgrading from Adaptive Server 15.0.2,” describes how to downgrade from Adaptive Server 15.0.2 to the previous version.
 - Chapter 7, “Troubleshooting SySAM Issues,” provides information about and solutions to SySAM related errors.
 - Chapter 8, “Troubleshooting,” provides installation error messages and possible solutions to installation problems.
 - Chapter 9, “Removing Adaptive Server,” describes how to uninstall Adaptive Server.
 - Appendix A, “Alternative Installation Methods,” describes how to install Adaptive Server in noninteractive mode.
 - Appendix B, “Creating Raw Partitions,” describes how to create partitions, bind partitions for raw disk I/O, and access raw devices from Adaptive Server.
 - Appendix C, “Upgrading Servers with Replicated Databases,” describes how to upgrade servers with replicated databases.
 - Appendix D, “Using sybssystemprocs,” describes how to use the sybssystemprocs database.

Related documents

The Sybase Adaptive Server Enterprise documentation set consists of the following:

- 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 Product Manuals at <http://www.sybase.com/support/manuals/>.
- The installation guide for your platform – describes installation, upgrade, and configuration procedures for all Adaptive Server and related Sybase products.
- *Sybase Software Asset Management User’s Guide* – describes asset management configuration concepts and tasks.
- *The New Feature Guide* – describes the new features in Adaptive Server version 15.0, the system changes added to support those features, and changes that may affect your existing applications.

- *ASE Replicator User's Guide* – describes how to use the Adaptive Server Replicator feature of Adaptive Server to implement basic replication from a primary server to one or more remote Adaptive Servers.
- *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.
- The *Configuration Guide* for your platform – provides instructions for performing specific configuration tasks for Adaptive Server.
- *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.
- *Glossary* – defines technical terms used in the Adaptive Server documentation.
- *Historical Server User's Guide* – describes how to use Historical Server to obtain performance information for SQL Server[®] and Adaptive Server.
- *Java in Adaptive Server Enterprise* – describes how to install and use Java classes as data-types, functions, and stored procedures in the Adaptive Server database.
- *Job Scheduler User's Guide* – provides instructions on how to install and configure, and create and schedule jobs on a local or remote Adaptive Server using the command line or a graphical user interface (GUI).
- *Messaging Service User's Guide* – describes how to use Real Time Messaging Services to integrate TIBCO Java Message Service and IBM WebSphere MQ messaging services with all Adaptive Server database applications.
- *Monitor Client Library Programmer's Guide* – describes how to write Monitor Client Library applications that access Adaptive Server performance data.
- *Monitor Server User's Guide* – describes how to use Monitor Server to obtain performance statistics from SQL Server and Adaptive Server.
- *Performance and Tuning Guide* – is a series of four books that explains how to tune Adaptive Server for maximum performance:
 - *Basics* – the basics for understanding and investigating performance questions in Adaptive Server.
 - *Locking* – describes how the various locking schemas can be used for improving performance in Adaptive Server.

-
- *Optimizer and Abstract Plans* – describes how the optimizer processes queries and how abstract plans can be used to change some of the optimizer plans.
 - *Monitoring and Analyzing* – explains how statistics are obtained and used for monitoring and optimizing performance.
 - *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.
 - *Reference Manual* – is a series of four books that contains the following detailed Transact-SQL[®] information:
 - *Building Blocks* – datatypes, functions, global variables, expressions, identifiers and wildcards, and reserved words.
 - *Commands* – commands.
 - *Procedures* – system procedures, catalog stored procedures, system extended stored procedures, and dbcc stored procedures.
 - *Tables* – system tables and dbcc tables.
 - *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.
 - *System Tables Diagram* – illustrates system tables and their entity relationships in a poster format. Available only in print version.
 - *Transact-SQL User's Guide* – documents Transact-SQL, the Sybase 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.
 - *Using Adaptive Server Distributed Transaction Management Features* – explains how to configure, use, and troubleshoot Adaptive Server DTM features in distributed transaction processing environments.
 - *Using Sybase Failover in a High Availability System* – provides instructions for using Sybase Failover to configure an Adaptive Server as a companion server in a high availability system.

- *Unified Agent and Agent Management Console* – Describes the Unified Agent, which provides runtime services to manage, monitor and control distributed Sybase resources.
- *Utility Guide* – documents the Adaptive Server utility programs, such as `isql` and `bcpl`, which are executed at the operating system level.
- *Web Services User's Guide* – explains how to configure, use, and troubleshoot Web Services for Adaptive Server.
- *XA Interface Integration Guide for CICS, Encina, and TUXEDO* – provides instructions for using the Sybase DTM XA interface with X/Open XA transaction managers.
- *XML Services in Adaptive Server Enterprise* – describes the Sybase native XML processor and the Sybase Java-based XML support, introduces XML in the database, and documents the query and mapping functions that comprise XML Services.
- *FLEXnet Licensing End User Guide* – this Macrovision manual explains FLEXnet Licensing for administrators and end users and describes how to use the tools which are part of the standard FLEXnet Licensing distribution kit from Sybase.
- *SAMreport User's Guide* – this Macrovision manual explains how to use SAMreport, a report generator that helps you monitor the usage of applications that use FLEXnet licensing.

Other sources of information

Use the Sybase Getting Started CD, the SyBooks™ CD, and the Sybase Product Manuals Web site to learn more about your product:

- The Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the SyBooks CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader, which you can download at no charge from the Adobe Web site using a link provided on the CD.
- The SyBooks CD contains product manuals and is included with your software. The Eclipse-based SyBooks browser allows you to access the manuals in an easy-to-use, HTML-based format.

Some documentation may be provided in PDF format, which you can access through the PDF directory on the SyBooks CD. To read or print the PDF files, you need Adobe Acrobat Reader.

Refer to the *SyBooks Installation Guide* on the Getting Started CD, or the *README.txt* file on the SyBooks CD for instructions on installing and starting SyBooks.

- The Sybase Product Manuals Web site is an online version of the SyBooks CD that you can access using a standard Web browser. In addition to product manuals, you will find links to EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

To access the Sybase 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.

v Finding the latest information on product certifications

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click Certification Report.
- 3 In the Certification Report filter select a product, platform, and timeframe and then click Go.
- 4 Click a Certification Report title to display the report.

v Finding the latest information on component certifications

- 1 Point your Web browser to Availability and Certification Reports at <http://certification.sybase.com/>.
- 2 Either select the product family and product under Search by Base Product; or select the platform and product under Search by Platform.
- 3 Select Search to display the availability and certification report for the selection.

v Creating 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.

Sybase EBFs and software maintenance

- v **Finding the latest information on EBFs and software maintenance**
 - 1 Point your Web browser to the Sybase Support Page at <http://www.sybase.com/support>.
 - 2 Select EBFs/Maintenance. If prompted, enter your MySybase user name and password.
 - 3 Select a product.
 - 4 Specify a time frame and click Go. A list of EBF/Maintenance releases is displayed.

Padlock icons indicate that you do not have download authorization for certain EBF/Maintenance releases because you are not registered as a Technical Support Contact. If you have not registered, but have valid information provided by your Sybase representative or through your support contract, click Edit Roles to add the “Technical Support Contact” role to your MySybase profile.
 - 5 Click the Info icon to display the EBF/Maintenance report, or click the product description to download the software.

Conventions

The following style conventions are used in this manual:

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

`this font`

- In a sample window 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 san-serif font.
<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.

Accessibility features

This document is available in an HTML version that is specialized for accessibility. You can navigate the HTML with an adaptive technology such as a screen reader, or view it with a screen enlarger.

Adaptive Server version 15.0 and the HTML documentation have been tested for compliance with U.S. government Section 508 Accessibility requirements. Documents that comply with Section 508 generally also meet non-U.S. accessibility guidelines, such as the World Wide Web Consortium (W3C) guidelines for Web sites.

The online help for this product is also provided in HTML, which you can navigate using a screen reader.

Note You might need to configure your accessibility tool for optimal use. Some screen readers pronounce text based on its case; for example, they pronounce ALL UPPERCASE TEXT as initials, and MixedCase Text as words. You might find it helpful to configure your tool to announce syntax conventions. Consult the documentation for your tool.

For information about how Sybase supports accessibility, see Sybase Accessibility at <http://www.sybase.com/accessibility>. The Sybase Accessibility site includes links to information on Section 508 and W3C standards.

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 product and installation of Adaptive Server and Adaptive Server products.

Topic	Page
Product licensing with SySAM 2.0	2
Product components	3
Adaptive Server editions	5
Adaptive Server specifications	8
Product descriptions and directory layout	9
PC-Client product descriptions and directory layout	12

This installation guide contains information about how to install and configure Adaptive Server version 15.0.2. See the following chapters for detailed information about how to install, upgrade, and troubleshoot your Adaptive Server:

- For information about SySAM 2.0 requirements, Adaptive Server editions and specifications, and product descriptions and directory layouts, see Chapter 1, “Overview.”
- For information about installing Adaptive Server and its components, see Chapter 2, “Installing Adaptive Server” and Chapter 4, “Installing Sybase PC-Client Products.” For post-installation tasks see Chapter 3, “Post-Installation Tasks.”
- For information about upgrading Adaptive Server from previous versions, see Chapter 5, “Upgrading Adaptive Server.”
- For information about troubleshooting, see Chapter 8, “Troubleshooting,” and Chapter 7, “Troubleshooting SySAM Issues.”
- For information about uninstalling Adaptive Server, see Chapter 9, “Removing Adaptive Server.”

Product licensing with SySAM 2.0

Starting with Adaptive Server Enterprise version 15.0, the Sybase Software Asset Management System (SySAM) implementation has changed. With these changes, SySAM configuration is no longer optional, and requires some changes in the Adaptive Server installation and configuration process. Review the updated SySAM implementation and plan your SySAM deployment before you install Adaptive Server Enterprise.

For information about how to install and configure SySAM 2.0 enabled products, see the following places for information:

- For simple installations, see the *Installation Guide* (this document,) for instructions. No additional information should be necessary.
- For more complex installations, and for general information about SySAM 2.0, see the *Sybase Software Asset Management User's Guide*.

Warning! SySAM 2.0 enabled products work for an initial period without a license. If you do not install a valid license for your configuration in this grace period, the product ceases to function. For general information about SySAM 2.0, see the *Sybase Software Asset Management User's Guide*.

Before you install Adaptive Server with SySAM 2.0, you must decide whether you are using a served or an unserved license. For information on how to determine which is appropriate for your site, see Chapter 2, “Choosing a License Model” in the *Sybase Software Asset Management User's Guide*.

The Adaptive Server *Installation Guide* (this document) has the following information about SySAM-related tasks:

- Installation instructions are described in “Installing components with the Installer in GUI mode” on page 28.
- If you have problems, see Chapter 7, “Troubleshooting SySAM Issues.”
- For information about using SySAM for asset management, see “The SAMreport” on page 42.

Product components

Adaptive Server Enterprise is based on the client/server model. Each Adaptive Server user is associated with a client process that communicates with a server process via a network interface called the Tabular Data Stream™ (TDS) protocol. The client process may execute on one machine and communicate with a database server on the same or a different machine.

Adaptive Server Enterprise

Adaptive Server runs as an application on top of an operating system and depends solely on the services exported by the operating system to function. Adaptive Server uses operating system services for process creation and manipulation, device and file processing, and inter-process communication. The hardware that the operating system runs on is completely transparent to Adaptive Server which sees only the operating system's user interfaces.

Adaptive Server has a virtual server architecture and has one or more operating system processes (engines) that service client requests. You can configure multiple processes to enhance performance on multiprocessor systems. A process is divided into two components, a DBMS component and a kernel component. The kernel component manages multiple tasks, while the DBMS component performs the database functions on behalf of the current task. Each client is associated with a task, and there are several tasks that perform specific services; for example, writing buffers or audit data to disk, and communicating with the network. The DBMS component of Adaptive Server manages the processing of SQL statements, accesses data in a database, and manages different types of server resources.

All server processes share two distinct areas of shared memory: a kernel area and a DBMS area. The kernel area contains data structures that relate to task management and operating system services, while the DBMS component contains caches of database buffers, object descriptors, and other resources used by the DBMS component of the server process.

The kernel layer of the server abstracts the operating system-specific functions so that the rest of the server has a consistent view of those functions regardless of the underlying operating system. Essentially, the kernel provides low-level functions, such as task and engine management, network and disk I/O, and low-level memory management to the rest of the server. The part of the server that processes a TDS request (the TDS engine) is part of the kernel layer.

The heart of Adaptive Server is the SQL/DBMS engine. The SQL/DBMS engine processes data manipulation statements (DML), utility commands, which are generally data definition statements (DDL) and administrative functions, and performs stored procedure executions. The DBMS engine uses the services provided by the resource manager routines and the access methods routines.

A resource manager is a set of subroutines that manage a particular system resource, such as pages, buffers, or a particular data structure. The DBMS engine uses these subroutines to manipulate various system resources. There are also a number of tasks within the server that perform specialized tasks. The resource managers, as well as these specialized tasks, rely on the kernel layer for low-level functions.

Access methods have two primary functions: the management of data on disk and manipulation of database constructs, such as tables and rows. Actual disk I/O routines are provided to the access methods routines by the kernel disk I/O routines.

Backup Server

Like Adaptive Server, Backup Server is also an operating system process. Backup Server is dedicated to backing up and restoring Adaptive Server files. Clients do not communicate directly with Backup Server. Only Adaptive Server communicates with Backup Server.

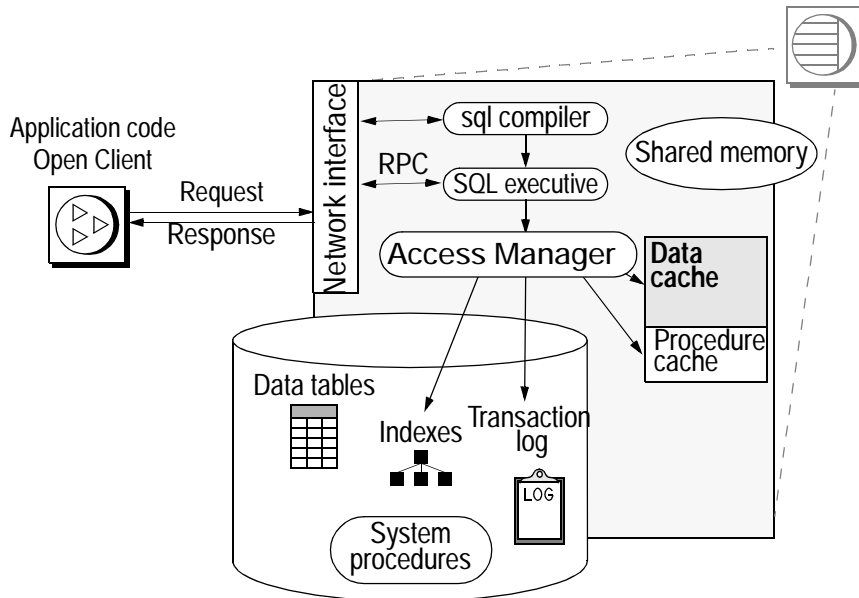
Stored procedures

Stored procedures encapsulate a series of SQL commands that the server can execute as a single command. System stored procedures (SSPs) are created when the server is installed.

The client

A client uses functions provided by the CT-Library (CT-Lib) and Common Library (CS-Lib) APIs to establish a connection to the server, submit requests to the server, and receive results back from the server.

Figure 1-1: Adaptive Server system model



Adaptive Server editions

Adaptive Server is available in the following editions:

- Enterprise Edition – has no limit on scalability and can run all available options that can be purchased separately.
- Small Business Edition – has some limit on scalability, and can run a limited set of options that can be purchased separately.
- Developer Edition – has a limit on scalability and includes many of the options.
- Express Edition – is a free release available only on Linux 32-bit.

Sybase also offers a series of options that you must purchase separately. Some of these options are intended for an enterprise-level site, and are not available in the Small Business Edition.

Table 1-1 describes the availability of each option. For a full description of the option contents, see your sales representative.

Table 1-1: Optional packages available for Adaptive Server

Available in the following editions:

Option	Enterprise Edition	Small Business Edition	Developer Edition	Express Edition	Description
Security and directory services	Optional	Not available	Included	Not available	Lightweight directory services and network-based authentication and encryption using DCE and Kerberos
High Availability	Optional	Not available	Included	Not available	Systems continuously available
SQL Expert	Optional	Optional	Not available	Not available	
Disaster Recovery option	Optional	Not available	Not available	Not available	Provides a disaster recovery package for Adaptive Server users
Partitions	Optional	Not available	Included	Not available	
Enhanced Full Text Search	Optional	Optional	Not available	Not available	

Adaptive Server editions and optional features are unlocked by providing SySAM licenses. For complete details, see the *Sybase Software Asset Management User's Guide*.

Use the `sp_lmconfig` command to determine which edition of Adaptive Server you are running. Enter:

```
sp_lmconfig 'edition'
go
```

Adaptive Server returns value of EE, SE, DE or XE based on the edition it is running. See the *Reference Manual* for more information about the `sp_lmconfig` command.

Enterprise Edition

The Enterprise Edition (EE) includes a full-featured Adaptive Server. In addition to the base server, you can purchase options, including High Availability. Your licensing agreement describes the features licensed for your site.

The options available for Adaptive Server Enterprise are listed in Table 1-1 on page 6.

Small Business Edition

The Small Business Edition (SE) includes features that meet the needs of small businesses. This edition of Adaptive Server includes most of the functionality of the Enterprise Edition, and allows for the following Adaptive Server configuration:

- Number of engines – maximum of 4.
- Number of connections – maximum of 256.

The options available for Adaptive Server Enterprise are listed in Table 1-1 on page 6.

Developer Edition

The Developer Edition (DE) is a competitively-priced version of Adaptive Server that developers can use to create applications in their development environment. The Developer Edition includes a full set of Adaptive Server Enterprise features and options (that is, DTM, High Availability, and so on), and allows for the following configuration:

- Number of engines – maximum of 1.
- Number of connections – maximum of 25.

The options available for Adaptive Server Enterprise are listed in Table 1-1 on page 6.

Express Edition

Sybase is providing a free release of Adaptive Server version 15.0.2 on Linux x86. This version, called Adaptive Server Enterprise Linux x86 Express Edition (XE), and is available for download at the Sybase Web site at <http://www.sybase.com> at no additional cost.

This Linux x86 Express Edition has the following configuration limitations:

- You can specify only one engine in the max online engines configuration parameter.
- You can configure up to 2GB of max memory.
- The total disk space is limited to 5GB.

This free edition does not include support. Support is available for purchase from Sybase eShop, which offers three plans:

- SupportNow Enterprise Developer Plan – provides fast response times for customers with critical development projects.
- SupportNow Developer Plan – provides unlimited access for customers with frequent support requirements.
- SupportNow Incident Plan – provides a case-based program that provides quality technical assistance so that you can easily scale your support needs and buy support as you need it.

Adaptive Server specifications

Table 1-2 provides Adaptive Server specifications for Linux.

Table 1-2: Adaptive Server specifications on Linux

Hardware		
Processor	32-bit Intel Pentium II	
Minimum RAM required for Adaptive Server	128MB	With default stack size, packet size, and user log cache size
Minimum RAM per additional user	Approx. 140KB	
Default user stack size	80KB	

Table 1-3: Adaptive Server specifications on Linux

Hardware		
Processor	AMD Opteron Processor or Intel Xeon with EM64T support	
Mimimum RAM required for Adaptive Server	128MB	
Mimimum RAM per addittional user	Approx. 140KB	
Default user stack size	80KB	

Note Adaptive Server 15.0.2 supports Pentium and higher chip sets.

Product descriptions and directory layout

The Adaptive Server installation installs a wide array of software products into the *SYBASE* directory. Table 1-4 provides a brief description of the server products installed, and the top-level directory structure.

Note Starting with ASE 15.0.2, the Enhanced Full-Text Search (EFTS) optional product will be not be shipped with ASE. It is now shipped separately but is installed into the same directory where Adaptive Server was installed. Please contact customer service for more information.

Table 1-4: Product description

Product	Description
Adaptive Server and Historical Server	Installed into the <i>ASE-15_0</i> directory. <ul style="list-style-type: none"> • Adaptive Server – the database server. • Backup Server – an Open Server-based application that manages all database backup (dump) and restore (load) operations. • Monitor Server – an Open Server-based application that obtains performance information on Adaptive Server and makes that information available to Monitor Server client applications. • XP Server – an Open Server application that manages and executes extended stored procedures (ESPs) from within Adaptive Server. • Job Scheduler – provides a job scheduler for Adaptive Server. Job Scheduler components are located in their own directory at <i>ASE-15_0/jobscheduler/</i>.
Software Developer Kit (SDK)	Installed into the <i>OCS-15_0</i> directory. <ul style="list-style-type: none"> • Open Client (Ctlib, dlib) • ESQ/C • Monitor Client Library • ESQ/COBOL • XA Installed into the <i>DataAccess</i> directory: <ul style="list-style-type: none"> • ODBC (Windows and Linux only) – a driver used to connect to Adaptive Server from ODBC based applications. • OLEDB (Windows only) – a provider used to connect to Adaptive Server from OLE DB-based applications. • ADO.NET (Windows only) – a provider used to connect to Adaptive Server from .NET based applications.
shared – on UNIX platforms Shared – on Windows platforms	The <i>shared</i> or <i>Shared</i> directory contains components and libraries that are shared by several other components. In earlier versions, this directory was named <i>shared-1_0</i> . <ul style="list-style-type: none"> • Sybase Central – Sybase Central is a Java-based framework used by the system management tools. In earlier versions, this component was placed into a top-level directory rather than in a shared directory. • JRE – the Java Runtime Environment (JRE) is a runtime Java virtual machine used to execute Java-based programs such as Sybase Central.
Adaptive Server Plug-in	Installed into <i>ASEP</i> directory. This plug-in contains the management software for managing Sybase servers and its associated servers.
Language Modules	Installed into the <i>\$\$SYBASE/\$SYBASE_ASE</i> directory. Provides system messages and date/time formats.
Character Sets	Installed into the <i>charsets</i> directory. Provides character sets available for use with Adaptive Server.

Product	Description
Collation Sequences	Installed into the <i>collate</i> directory. Provides the collation sequences available for use with Adaptive Server.
ASE Replicator	Installed into the <i>RPL-15_0</i> directory. Provides lightweight replication for those companies who do not need a full-featured replication product such as Sybase Replication Server.
Sybase Software Asset Management (SySAM)	Installed into the <i>SYSAM-2_0</i> directory. Provides asset management of Adaptive Servers and optional features.
Web Services	Installed into the <i>WS-15_0</i> directory. An optional product that provides access to Adaptive Server using HTTP/SOAP and WSDL.
Agent Management Console	This is a Sybase Central plug-in that manages the Unified Agent. It is installed into the <i>AMCP</i> directory.
Interactive SQL	Interactive SQL is a GUI tool that allows you to execute SQL statements, build scripts, and display data from Adaptive Server. It is installed into the <i>DBISQL</i> directory.
Unified Agent	Unified Agent is a management framework that provides runtime services to manage, monitor, and control distributed Sybase resources. Agents plug into the framework and provide the ability to manage specific components. It is installed into the new UA directory <i>\$SYBASE/UAF-2_0</i> .
jutils-2_0	A collection of Adaptive Server utility programs including: <ul style="list-style-type: none"> • <i>jisql</i> – a JDBC-based GUI <i>isql</i> tool. • <i>ribo</i> – a tool to trace TDS traffic between client program and Adaptive Server.
jConnect	Installed into the <i>jConnect-6_0</i> directory. Provides a Java Database Connectivity (JDBC) driver for Adaptive Server.

Sybase support for OLE DB, and ODBC

Adaptive Server version 15.0.2 contained new ODBC and OLE DB drivers developed by Sybase. The third-party rebranded ODBC and OLE DB Driver Kits included with previous versions are no longer shipped.

The retired ODBC Driver Kit was installed in *%SYBASE%\ODBC*, and registered with the ODBC Driver Manager as “Sybase ASE ODBC Driver”. The new ASE ODBC Driver by Sybase is installed in *%SYBASE%\DataAccess\ODBC*, and registered as “Adaptive Server Enterprise”. The version shipping with Adaptive Server and SDK 15.0 is version 15.0.0.50.

The retired OLE DB Driver Kit was installed in *%SYBASE%\OLEDB*, and used the provider short name of “Sybase.ASEOLEDBProvider” and the long name of “Sybase ASE OLE DB Provider”. The new ASE OLE DB Provider by Sybase is installed in *%SYBASE%\DataAccess\OLEDB*, and uses provider short name “ASEOLEDB”. The version shipping with SDK 15.0 is version 15.0.0.51.

Sybase recommends that you start your migration process the new ODBC and OLE DB drivers as soon as possible. New features introduced in Adaptive Server 15.0 are supported only by the new drivers.

See “New Features for OpenServer 15.0 and SDK 15.0 for Windows, Linux and UNIX”, for instructions on migrating to the new drivers.

Installing ODBC driver from ASE 15.x or SDK 15.x on top of an IQ 12.6 installation might produce a warning. You can ignore this warning. The ODBC driver will be functional at the end of the install.

PC-Client product descriptions and directory layout

The following products are installed as Adaptive Server PC-Client products. In addition to these products, there are several separately installed products listed in the following sections.

Table 1-5: Products and platforms

Product	Description
Software Developer Kit (SDK)	<p>Installed into the <i>OCS-15_0</i> directory.</p> <ul style="list-style-type: none"> • Open Client (Ctlib, dblib) • ESQ/C • Monitor Client Library • ESQ/COBOL • XA • jutils-2_0 • jConnect <p>Installed into the <i>DataAccess</i> directory:</p> <ul style="list-style-type: none"> • ODBC (Windows and Linux only) – a driver used to connect to Adaptive Server from ODBC based applications. • OLEDB (Windows only) – a provider used to connect to Adaptive Server from OLE DB-based applications. • ADO.NET (Windows only) – a provider used to connect to Adaptive Server from .NET based applications.
Windows Cluster Server Admin utilities	Windows Cluster Administrator is a GUI tool for managing Microsoft Cluster Server (MSCS). It can be used to create, modify and displays the information of groups, resources and the cluster itself. There is also an alternative command-line tool called <i>Cluster.exe</i> for cluster administration.
Shared	<p>The shared directory contains components and libraries that are shared by several other components. In earlier versions, this directory was named <i>shared-1_0</i>.</p> <ul style="list-style-type: none"> • Sybase Central – Sybase Central is a Java-based framework used by the system management tools. In earlier versions, this component was placed into a top-level directory rather than in a shared directory. • JRE – The Java Runtime Environment (JRE) is a runtime Java virtual machine used to execute Java-based programs such as Sybase Central.
Adaptive Server Plug-in	Installed into <i>ASEP</i> directory. This plug-in contains the management software for managing Adaptive Server and its associated servers.
Language Modules	Installed into the <i>locales</i> directory. Provides system messages and date/time formats.
Character Sets	Installed into the <i>charsets</i> directory. Provides character sets available for use with Adaptive Server.
jConnect	Installed into the <i>jConnect-6_0</i> directory. Provides a Java Database Connectivity (JDBC) driver for Adaptive Server.
Agent Management Console	This is a Sybase Central plug-in that manages the Unified Agent. It is installed into the <i>AMCP</i> directory.
Interactive SQL	Interactive SQL is a GUI tool that allows you to execute SQL statements, build scripts, and display data from Adaptive Server. It is installed into the <i>DBISQL</i> directory.

Product	Description
Unified Agent	Unified Agent is a management framework that provides runtime services to manage, monitor, and control distributed Sybase resources. Agents plug into the framework and provide the ability to manage specific components. It is installed into the new UA directory %SYBASE%\UAF-2_0.
juutils-2_0	A collection of Adaptive Server utility programs including: <ul style="list-style-type: none">• jisql – a JDBC-based GUI isql tool.• ribo – a tool to trace TDS traffic between client program and Adaptive Server.

Separately installable PC-Client products

The PC-Client CD also contains the following products, which are separately installable:

- InfoMaker® – a tool for personal data access, management, and reporting, used by both 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.
- PowerDesigner Physical Architect – a tool for data modeling, including database design, generation, maintenance, reverse engineering, and documentation for database architects.
- PowerTransfer – a plug-in for PowerDesigner™ that transfers table data from any supported database to Adaptive Server.
- SAMreport – a Macrovision report generator that helps you monitor the usage of applications that use FLEXnet licensing.
- Sun Development Kit - tool for creating software package and debugging

Installing Adaptive Server

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

Topic	Page
Overview	15
System requirements	16
Pre-installation tasks for SySAM	17
Pre-installation tasks for Adaptive Server	24
Installing server components	28
Installing version 15.0.2 over an existing 15.x Adaptive Server	36

Overview

The installation process copies all of the files from the installation media to the hard drive and then configures a base Adaptive Server. At the end of the installation, the Adaptive Server and associated software is ready to be customized for the user's environment – for example, creating devices, databases and tables, and loading data.

The basic process for installing Adaptive Server is:

- 1 Perform pre-installation tasks related to license administration as described in “Pre-installation tasks for SySAM” on page 17.

Warning! SySAM 2.0 enabled products work for an initial period without a license. If you do not install a valid license for your configuration in this grace period, the product ceases to function. For general information about SySAM 2.0, see the *Sybase Software Asset Management User's Guide*.

- 2 Perform the pre-installation tasks as described in “Pre-installation tasks for Adaptive Server” on page 24.

- 3 Perform the installation using any of the methods described in “Installation methods” on page 28.
- 4 Perform the post-installation tasks as described in Chapter 3, “Post-Installation Tasks.”

System requirements

Check to make sure your system has the required version-level, RAM, and network-protocol requirements for Adaptive Server. See the release bulletin for your platform the latest information about required system patches.

Table 2-1: System requirements for Linux 32-bit and 64-bit

Operating system	Hardware	Minimum RAM
Linux x86 RedHat Enterprise Linux 3.0 update: 8 <ul style="list-style-type: none"> • kernel 2.4.21-47.EL or higher • glibc-2.3.2-95.44 or higher • compat-libstdc++-7.3-2.96.128 or higher 	Intel x86 architecture Minimum: Pentium III, 300 MHz Recommend: at least 1 GHz	128 MB Recommend: at least 512 MB
Linux x 86 Red Hat Enterprise Linux 4.0 update: 0 <ul style="list-style-type: none"> • kernel -2.6.9-5.EL • glibc-2.3.4-2 	Intel x86 architecture Minimum: Pentium III, 300 MHz Recommend: at least 1 GHz	128 MB Recommend: at least 512 MB
Linux x86 SuSE Linux Enterprise Server SLES 9.0 Update: 2 <ul style="list-style-type: none"> • kernel 2.6.5-7.191 or higher • glibc 2.3.3-98.47 or higher • compat-2004.7.1-1.2 or higher 	Intel x86 architecture Minimum: Pentium III, 300 MHz Recommend: at least 1 GHz	128 MB Recommend: at least 512 MB
Linux x86 SuSE Linux Enterprise Server SLES 10 <ul style="list-style-type: none"> • kernel 2.6.13-15 or higher • glibc 2.3.5-40 or higher 	Intel x86 architecture Minimum: Pentium III, 300 MHz Recommend: at least 1 GHz	128 MB Recommend: at least 512 MB
Linux x64 RedHat Enterprise Linux 3.0 update 8 <ul style="list-style-type: none"> • kernel 2.4.21-47.EL or higher • glibc-2.3.2-95.44 or higher • compat-libstdc++-7.3-2.96.128 or higher 	AMD Opteron Processor or Intel Xeon with EM64T support	128 MB Recommend: at least 512 MB

Operating system	Hardware	Minimum RAM
Linux x64	AMD Opteron Processor or Intel	128 MB
RedHat Enterprise Linux 4.0 update: 0 <ul style="list-style-type: none"> kernel -2.6.9-5.EL glibc-2.3.4-2 	Xeon with EM64T support	Recommend: at least 512 MB
Linux PSeries	IBM pSeries Server with	256 MB
RedHat Enterprise Linux 3.0 update 8 <ul style="list-style-type: none"> kernel 2.4.21-47.EL or higher glibc-2.3.2-95.44 or higher compat-libstdc++-7.3-2.96.128 or higher vacpp.rte-7.0.0-0 or higher 	POWER5 CPU	Recommend: at least 512 MB
Linux PSeries	IBM pSeries Server with	256 MB
Red Hat Enterprise Linux 4.0 update: 0 <ul style="list-style-type: none"> kernel -2.6.9-5.EL glibc-2.3.4-2 compat-libstdc++-33-3.2.3-47.3 	POWER5 CPU	Recommended: at least 512 MB

Table 2-2: Disk space requirements for Linux 32-bit platforms

Product	Disk space requirements
Adaptive Server typical install	620MB
Default databases created during install	150MB
Total	1320MB

Pre-installation tasks for SySAM

Before you install Adaptive Server, you should:

- 1 Decide the SySAM license model you will use. The choices are:
 - The unserved license model – gets licenses directly from the license file. If you are using an unserved license, you should simply save the license to the machine on which you have installed Adaptive Server.
 - The served license model – uses a license server to manage the allocation of your licenses to multiple machines. See the following for more information.

For information about served and unserved licenses, and which model is best suited for your environment, see Chapter 2, “Choosing a License Model” in the *Sybase Software Asset Management User’s Guide*.

- 2 If you plan to use a served license, you need a license server. Decide whether it will be an existing license server or a new license server. The license server need not be on the same machine or running on the same operating system and architecture.
 - If you are going to use an existing license server, you must know the server host name and port number.
 - For information about creating a new license server before installing Adaptive Server, see “Installing a new license server” on page 23.

Note There can be only one instance of a SySAM license server running on a given machine. If you want to setup a SySAM 2.0 license server on a machine that is already running a SySAM 1.0 license server, the old license server must be migrated to 2.0. A migrated license server can serve licenses for both SySAM 1.0 and SySAM 2.0 enabled products.

Please see Appendix F of the *Sybase Software Asset Management User’s Guide* for instructions on migrating a license server.

For information about choosing a license server host, see the “Choosing a license server host” section in Chapter 2, “Choosing a License Model” of the *Sybase Software Asset Management User’s Guide*.

For information about license server administration, see Chapter 4, “License Server Administration” in the *Sybase Software Asset Management User’s Guide*.

- 3 Sybase recommends that you get your license files from the Sybase Product Download Center (SPDC) at <https://sybase.subscribenet.com> before you begin installing Adaptive Server. For information about accessing SPDC, see “Accessing SPDC” on page 18, and for information about generating your licenses at SPDC, see “Generating licenses at SPDC” on page 20.

Accessing SPDC

Once you have purchased a Sybase product, typically you are expected to download the product and generate its licenses from the online Web portal, the Sybase Product Download Center (SPDC.)

To access SPDC when you have purchased Sybase products from Sybase:

- 1 Once you have ordered a Sybase product you receive a welcome e-mail that contains a URL to the SPDC site, as well as a user name and password.

Note If you have purchased your Sybase software from a Sybase reseller, you will receive a web key rather than an e-mail.

- 2 Click on the URL, and when prompted, enter the designated user name and password. Your login is typically your e-mail address. If you have forgotten your password, use the password finder on the SPDC login page.

Getting your host ID

To get your host ID:

- 1 Change to `$$SYBASE/SYSAM-2_0/bin`. Where `$$SYBASE` is the Sybase installation directory.
- 2 Enter:

```
lmutil lmhostid
```

Note On some platforms, the host ID is derived from the network adapter address. If your machine has multiple network adapters, `lmutil lmhostid` returns one host ID for each network adapter. The output may look similar to:

```
The FLEXlm host ID of this machine
is "0013023c8251 0015c507ea90"
Only use ONE from the list of hostids.
```

Choose *one* of these host IDs. Sybase recommends using the value associated with the primary wired Ethernet adapter. Do not use values associated with internal loop back adapters. If you cannot determine which host ID to use from the `lmutil lmhostid` output, use the native operating system command to get additional details to help make the determination. See to the SPDC FAQ titled, “What’s my Host ID,” or “Appendix A” of the *FLEXnet Licensing End User Guide* for the exact commands for your platform.

- 3 Save the host ID so you can use it later at SPDC. Your host ID is platform-specific, but may be similar to:

00400516E525 on Linux platforms.

Generating licenses at SPDC

This section provides information about how to generate a license for a served and an unserved license model. To generate your license, regardless of license model:

- 1 At the Welcome screen at SPDC, you can view all of the product families to which you have access. Select the product family that contains the product for which you wish to generate a license.
- 2 A list of available Sybase products displays. Click the product with the appropriate edition and version. Make sure you select the product for the operating system you are using.
- 3 The license agreement displays. You must agree to the terms and conditions in the license in order to generate a license. To do so, click I Agree.

Note The license agreement only displays the first time you attempt to download a product. Once you have agreed to the license, you will not encounter the license agreement for all subsequent downloads of that product.

- 4 At the product download page, click the License Keys link.
- 5 The License Information page displays. Select the license you wish to generate. When you are selecting the license you wish to generate, consider:
 - The license *type* for the license you wish to generate. For more information about license types, see the “License types” in Chapter 3 “Getting and Using your License” of the *Sybase Software Asset Management User’s Guide*.
 - If you are unsure what you have licensed, your company’s purchase order should have the product name, edition, and license type.
- 6 Once you’ve determined the license you wish to generate, select it via the Generate radio button on the left side of the screen. Scroll to the bottom of the page, and click Select to Generate.

- 7 Step one of the SySAM license generation wizard asks you to choose between a served license and an unserved license. If you choose a served license, continue with the steps in “Generating a served license” on page 21, and if you choose an unserved license, continue with the steps in “Generating an unserved license” on page 22.

Note Some Sybase Products or specific license types do not give you a choice of license model to use, and therefore do not display this page. If this is the case, proceed with the SySAM license generation wizard to generate your license.

Generating a served license

To generate a served license:

- 1 Select Served License and click Next.
- 2 Indicate the quantity of licenses you wish to generate and click Next.
- 3 Enter the License Server Host ID and optionally the host name and port number.
 - a Sybase recommends that you provide the host name here to ease future license administration.
 - b A port number is not required unless you are using a 3-node redundant cluster. A valid number is any unused port number between 0 and 64000. On UNIX, choose a port greater than 1024, since those less than 1024 are privileged port numbers. If no TCP/IP port number is specified, one of the default ports in the range of 27000 and 27009 is used.
 - c If you want to generate a license for a 3-node redundant cluster, then you must specify the host ID, host name, and port number for the three nodes of the cluster. The port number is not optional for this configuration and should be outside of the 27000 to 27009 range.
 - d You should have gotten your server host ID before you began the process of generating your licenses. For information about how to get the host ID, see “Getting your host ID” on page 19. You have the option to specify redundant server information if you choose to configure your servers for three-server redundancy.
- 4 Click Generate.
- 5 Once you have generated the license, you can:

- a Download the license file by clicking the Download License File button.
 - b Print a paper copy of the license by clicking the Print Friendly button.
 - c Return to the license information page to generate additional licenses.
- 6 Once you have generated all the licenses you need, save them to the *licenses* directory in your license server installation.

Note You must save your license files with a *.lic* extension, or SySAM 2.0 will not recognize them.

Sybase recommends that you setup the license server and required licenses before you install your SySAM 2.0 enabled Sybase product.

Generating an unserved license

To generate an unserved license:

- 1 Select Un-served License and click Next.
- 2 Select the number of machines you want to license and click Next. You can generate licenses for a maximum of ten machines at a time.
- 3 Enter the host ID, and optionally the host name for each machine for which you want to generate a license. You should have determined the host ID before you began to generate your licenses. For information on getting your host ID and host name, see “Getting your host ID” on page 19.

For some license types, you are asked to enter the number of CPUs or licenses for the machine where this license will be used.

- 4 Click Generate.
- 5 Once you have generated the license, you can:
 - Download the license file by clicking the Download License File button.
 - Print a paper copy of the license by clicking the Print Friendly button.
 - Return to the license information page to generate additional licenses.

- 6 Once you have generated all the licenses you need, save them to the specific location required by your product.

Note You must save your license files with a *.lic* extension, or SySAM 2.0 does not recognize them.

Installing a new license server

Note If you have selected a served license model, you must install a license server.

For information about operating system requirements, see “System requirements” on page 16.

Perform the following steps to install a new license server.

- 1 Install the license server.
 - a Launch the Sybase installer program by entering:

```
./setup
```
 - b The Welcome Window displays. Click Next.
 - c Accept the license agreement.
 - d Enter or select the destination directory.
 - e Select the Custom installation.
 - f Select only “SySAM License Server” for installation.
 - g Deselect other components, to install the license server only.
 - h Click OK for the installer to install the license server.

Note The license server cannot be started until there is at least one valid served license installed in the *licenses* directory.

- 2 Get the host ID for the machine on which you will be running the license server. For information on how to generate the host ID, see “Getting your host ID” on page 19.
- 3 Go to the Sybase Product Download Center (SPDC) at <https://sybase.subscribenet.com>.

- 4 Generate the licenses for the products you want to install. For instructions on generating your licenses, see “Generating licenses at SPDC” on page 20.
- 5 From command prompt, navigate to installed directory to source the environment variables.
 - SYBASE.csh
- 6 Copy the license file in the `$SYBASE/SYSAM-2_0/licenses` directory on the network license server machine.
- 7 Refresh or restart the license server:
 - a Change to:
`$SYBASE/$SYBASE_SYSAM/bin`
 - b If the license server is not already started, start it, by entering:
`sysam start`

If the license server is already started, use the reread command to make the license server read the new license files:
`sysam reread`
- 8 Validate that the license daemon is running by entering:
`sysam status`

You can expect to see output for a running server that is similar to:

```
lmutil - Copyright (c) 1989-2006 Macrovision Europe
Ltd. and/or Macrovision Corporation. All Rights
Reserved. Flexible License Manager status on Wed
5/24/2006 15:04

License server status: 27000@mysysamserver
License file(s) on keyserver:
/opt/sybase/SYSAM_0/licenses/mysysamserver_60302031
5.lic
```

For more information, see Chapter 7, “Troubleshooting SySAM Issues.”

Pre-installation tasks for Adaptive Server

Before installing Adaptive Server:

- 1 Read the release bulletins for the latest information on the products (Adaptive Server, Monitor Server, and so on) that you are installing. See “Special Installation Instructions” in the release bulletin.
- 2 Install operating system patches, if required.
- 3 If you are having trouble starting the installer, make sure you have the required operating system patches for the Java Runtime Environment (JRE) version 1.4.2.

Information about the required operating system patches is available from the IBM Java Web site at <http://www.ibm.com/java>.

- 4 Review the SySAM procedures and plan your client/server configuration using the configuration guide for your platform.
- 5 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.

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

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.

- 7 Decide where the Adaptive Server software will be installed. Make sure there is sufficient available disk space. There cannot be any spaces in the path name of the directory.
- 8 Know what product edition you are installing. The options are:

- Enterprise Edition
- Small Business Edition
- Developers Edition
- Express Edition

If you are installing the Small Business or Enterprise Edition, you must know the license type under which Adaptive Server was licensed. For more information about license types, see the “License types” in Chapter 3 “Getting and Using your License” of the *Sybase Software Asset Management User’s Guide*.

9 Sybase recommends that you specify license types at the time of installation. However, you may choose to specify the product edition and license type later using `sp_lmconfig`. For information on `sp_lmconfig`, see *Reference Manual:Procedures*.

10 You must decide if you want licensing events to trigger e-mail alerts and the severity of the events that will generate e-mail messages.

If you choose to have e-mail notification of license events, you must know the:

- SMTP server host name
- Port number for an SMTP server
- E-mail return address
- Recipients of the notifications
- Severity level of an event that will trigger mail. Your choices are:
 - None
 - Informational
 - Warning
 - Error

11 Verify that the operating system meets the version-level, RAM, and network protocol requirements for your platform.

12 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 to verify your network configuration, ping the host.

13 The operating system shared memory default, for most Linux releases is 32MB. The minimum required by Adaptive Server is 64MB for default Server with 2K pages. A higher value is required if you plan to increase Adaptive Server's total memory.

Use the `sysctl(8)` method to check and adjust the operating system shared memory parameter.

To check the current shared memory size, enter:

```
# /sbin/sysctl kernel.shmmax
```

To adjust the shared memory size:

```
# /sbin/sysctl -w kernel.shmmax=nnn
```

where *nnn* is the new size in bytes (at least 64MB, which is 67108864 bytes).

To guarantee that this value is applied every time the system is started, add the above line to your */etc/rc.d/rc.local* file. On SuSE systems, the file is */etc/init.d/boot.local*.

- 14 On Linux RHEL Update 3.0 and greater, Adaptive Server version 15.0 running on multiple engines requires the security feature Exec-Shield to be disabled.

To disable Exec-Shield:

- a Add the following lines in */etc/sysctl.conf*:

```
kernel.exec-shield=0
kernel.exec-shield-randomize=0
```

- b As SU, enter:

```
/sbin/sysctl -P
```

for the action to take effect.

For additional information, see the Red Hat website at http://www.redhat.com/f/pdf/rhel/WHP0006US_Execshield.pdf.

If after adjusting the SHMMAX parameter the server fails to reboot you may also need to increase the value of another Kernel parameter, SHMALL, which is the maximum amount of shared memory that can be allocated. Its value is in the file */proc/sys/kernel/shmall*. This task requires root permission.

Installing server components

Follow the instructions in this section to install Adaptive Server, Backup Server, and Monitor Server.

Note Adaptive Server Enterprise 15.0.2 includes new major versions of Adaptive Server and many of the supporting components. Installing Adaptive Server version 15.0.2 into the same directory with existing products should not impact the existing products. However, installing other products on top of Adaptive Server version 15.0.2 will likely result in one or more products not working correctly.

Sybase strongly recommends that Adaptive Server version 15.0.2 be installed into its own directory if possible. Where this is not practical and other products must be installed into the same directory, Adaptive Server version 15.0.2 should be installed last.

Be sure that you are logged in as “sybase”, and that you have performed the “Pre-installation tasks for Adaptive Server” on page 24.

Installation methods

The following instructions are for using the Installer in GUI mode. You can also perform the installation in a non-GUI (text-based) mode or a completely silent installation with no interaction. Or, you can perform the installation, then configure Adaptive Server separately using either the GUI-based configuration tools or silently with the use of a resource file. For information on alternate methods of installation see Appendix A, “Alternative Installation Methods.”

Installing components with the Installer in GUI mode

Note Sybase recommends installing products as a sybase user. Root permissions are not required to run the installer.

The Installer creates the target directory (if necessary) and installs the selected components into that directory.

At the end of the installation, you can verify the product installation. You may need to perform additional configuration procedures before you can use some products.

As part of the installation, you must source the environment variables needed for Adaptive Server products by running the *SYBASE.csh* script file after exiting the Installer.

To install server components:

- 1 Stop and shut down all programs before running the installer.
- 2 Insert the CD into the CD drive, or download and extract the Adaptive Server install image from the Sybase Product Download Center (SPDC).

Launch the Adaptive Server Installer.

Start the installer by entering:

```
/cdrom/setup
```

- 3 The Welcome screen displays. Click Next.
- 4 When the license selection window displays, select the most appropriate location from the drop-down list, read the license terms, and then click “I agree” to proceed. Click Next.

Not every country is listed in the drop-down list. If the country you are located in is not listed, select the most appropriate area.

- 5 The Install directory window allows you to select a directory for the installation by clicking Browse. Click Next to accept the default of:

```
/opt/sybase
```

You can also enter a directory for installation.

- 6 You can choose from three types of installations in the Install Type window:

- Typical

Note Be aware that the Typical install does not install the license server. If you need to install the license server, select the Custom installation and select your components or select the Full installation.

- Full
- Custom

- 7 If you select Custom, the Products and Features Selection window displays. Select the products you want to install by checking the box next to each product name.

Note Some features are dependent on other features; therefore, you cannot unselect some items without first unselecting others.

After you have made your selection, click Next. Before proceeding to the next window, the Installer verifies the selections, and checks for dependencies and available disk space.

- 8 The Product Selection Summary window displays the selections that you have made.

Verify that you have selected the correct type of installation, and that you have enough disk space to complete the process. Click Next to proceed.

- 9 The Install Progress window shows the progress of the installation.
- 10 The Install Status window displays the result of the installation process. Click Next to configure the software.
- 11 The Sybase Software Asset Management License Server screen allows you to designate where Adaptive Server will find its licenses. Where you choose to store your license depends upon whether you have selected a served or unserved license model.

Note For information about served and unserved licenses, and about choosing the best license server for your environment, see Chapter 2, “Choosing a License Model” of the *Sybase Software Asset Management User’s Guide*.

You must specify whether the licenses will be obtained from a license server.

If you answer Yes, you must know the host name of the machine where the license manager is running, and the port number if you have used a non-default port number.

If you answer No, you will be using unserved licenses. You will be reminded to download and install the license file after you have finished the installation.

Note When you download the license file, copy it to `$$SYBASE/$SYBASE_SYSAM/licenses`.

- 12 The Sybase Software Asset Management Notification screen asks you to configure your server for e-mail notification. When configuration is enabled, designated users receive information about license management events requiring attention.

Provide the following information:

- SMTP server host name
 - SMTP server port number
 - E-mail Return Address
 - Recipient e-mail addresses
 - Message severity that triggers e-mail messages
- 13 The Product Licenses screen asks you to indicate the edition of Adaptive Server you are installing. You can choose from:
- Unknown
 - Enterprise Edition
 - Small Business Edition
 - Developers Edition
 - Express Edition

Note You must select only the Adaptive Server edition for which you have a license. For more information about troubleshooting license type and edition problems, see Chapter 7, “Troubleshooting SySAM Issues.”

- 14 If you are installing Enterprise Edition or Small Business Edition, you must also enter the License Type under which Adaptive Server is licensed.
- 15 The Configure New Servers window asks if you want to configure new servers, depending upon the products you installed. A Full or Custom installation allows you to choose to:
- Configure new Adaptive Server

- Configure new Backup Server
- Configure new Monitor Server
- Configure new XP Server
- Configure new Job Scheduler
- Enable Self Management
- Configure Web Services
- Configure Unified Agent

Select the servers you want to configure. If you do not want to configure the new servers at this point, unselect the item and click Next. The Custom Configure New Server Options window appears.

- 16 The Custom Configure New Server Options window allows you to select servers to custom configure. If you do not select any servers in this window, the Installer configures all the servers using default values.

If you choose the default values, the Installer chooses everything including the server name, port number, and master device location. It also installs Backup Server, Monitor Server, and the XP Server, choosing the name, port number, and error log.

If you accept the defaults, the server names are:

- Adaptive Server – *<host name>*
- Backup Server – *<host name>_BS*
- Monitor Server – *<host name>_MS*
- XP Server – *<host name>_XP*
- Job Scheduler Agent – *<host name>_JSAGENT*

If you configure the servers with default values, the Configure Server Attributes Summary window displays next.

If you select to configure the server with custom values, you have the following choices:

- Custom configure new Adaptive Server
- Custom configure new Backup Server
- Custom configure new Monitor Server
- Custom configure new XP Server

- Custom configure new Job Schedule
- Custom configure Self Management
- Custom configure Web Services
- Custom configure new Unified Agent

The Custom Configure Input window displays next.

17 The Custom Configure Input window prompts you for the custom Adaptive Server configuration information.

- Server name
- Port number
- Page size
- Error log – name of error log file and the path where it should be located.
- Master device – name of master device and the path where it should be located.
- Master device size (MB). The minimum master device sizes are:
 - 2K page size – 30MB
 - 4K page size – 60MB
 - 8K page size – 120MB
 - 16K page size – 240MB
- Master database size (MB). The minimum master database sizes are:
 - 2K page size – 13MB
 - 4K page size – 26MB
 - 8K page size – 52MB
 - 16K page size – 104MB
- System procedure device
- System procedure device and database minimum size are 132MB.
- System device
- System device size (in MB) The minimum sizes are:
 - 2K page size server – 3MB

- 4K page size server – 6MB
- 8K page size server – 12MB
- 16K page size server – 24MB
- System database size (in MB) The minimum sizes are:
 - 2K page size server – 3MB
 - 4K page size server – 6MB
 - 8K page size server – 12MB
 - 16K page size server – 24MB

Once you have customized the Adaptive Server configuration, select Next to record the input fields.

- 18 The Custom Configure Backup Server Input window prompts you for the Backup Server information.

Once you have entered the appropriate information, click Next to record the information.

- 19 The Custom Configure Monitor Server Input window prompts you to enter the Monitor Server configuration information.

Once you have entered the appropriate information, click Next to record the information.

- 20 The Custom Configure XP Server Input window prompts you to enter the XP Server configuration information.

Once you have entered the appropriate information, click Next to record the information.

- 21 The Custom Configure New Job Scheduler Input window prompts you to enter Job Scheduler information:

- Agent name
- Port number
- Management device
- Management device size
- Management database size

Once you have entered the appropriate information, click Next to record the information.

- 22 Custom configure Self Management by entering the following information:
- Self Management user name
 - Self Management password

Note If you choose the default user name “sa,” you cannot enter the password. The “sa” default password is null.

- 23 Security Login Modules – prompts you to choose the security login modules for the Unified Agent. You can choose:
- Simple Login Module

Note If you double-click Simple Login Module, a dialog asks you to change the user name and password.

- ASE Login Module
- Unix Proxy Login Module

Select the Enable box to enable a login module. To order the modules according to priority, select the module, and click Move up or Move down.

- 24 The Configure New Servers Summary window displays a summary of the values that are used to configure the servers. These values are either the default or custom values. After verifying the information, click Next, and the Installer proceeds with the server configuration.
- 25 The Configure Server Progress window displays the progress of the server configuration.
- 26 The Installation Complete window is the last window.
- If you encounter any errors, see the *Troubleshooting Guide*.

Adaptive Server and related products have been successfully installed, and minimally configured for use. See “Installing sample databases” on page 42, to begin experimenting with your servers, or see the *System Administration Guide* at Sybase Product Manuals at <http://www.sybase.com/support/manuals> for more advanced topics.

Adaptive Server

You have a running server, as well as various system databases and system tables. Some references that may help you get started using your new Adaptive Server include:

- “Initializing Database Devices” and “Creating User Databases” in the *System Administration Guide* – information about creating an Adaptive Server user database and its devices.
- *Transact-SQL User’s Guide* – learn how to write queries.

Backup Server	For information on developing a backup and recovery plan and backing up and restoring user databases, see the <i>System Administration Guide</i> .
Monitor Server	Monitor Server requires some additional configuration after installation. See the <i>Monitor Server User’s Guide</i> .
XP Server	For information on using extended stored procedures, see the <i>Transact-SQL User’s Guide</i> .
Job Scheduler	For information on using Job Scheduler, see the <i>Job Scheduler User’s Guide</i> .
Web Services	<i>Web Services User’s Guide</i> – explains how to configure, use, and troubleshoot Web Services for Adaptive Server.
Unified Agent	<i>Unified Agent and Agent Management Console</i> – Describes the Unified Agent, which provides runtime services to manage, monitor and control distributed Sybase resources.

Installing version 15.0.2 over an existing 15.x Adaptive Server

This section provides installation instructions for installing Adaptive Server 15.0.2 over and existing 15.x version.

Determining Adaptive Server version

Verify that your current Adaptive Server installation is at version 15.x by running the following query from isql:

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

If your server is not running, you can get a version string by running:

```
$$SYBASE/$SYBASE_ASE/bin/datasever -v
```

If this shows that the Adaptive Server is not at version 15.x, you must upgrade Adaptive Server. See Chapter 5, “Upgrading Adaptive Server.”

If your server is at a 15.x version level, you can begin installing Adaptive Server 15.0.2.

Backing up Adaptive Server

Adaptive Server 15.0.2 provides a script for backing up sysmessages before upgrading to 15.0.2. Use `uninstmsgs.ebf` to back up sysmessages before running `instmsgs.ebf`. Refer to the “Post-upgrade tasks,” section.

Installing Adaptive Server version 15.0.2 overwrites the current Adaptive Server software. Before installing, verify that your databases are error free and that your `$$SYBASE` directory is backed up.

To ensure that your database is error free, run `dbcc checkdb`, `dbcc checkcatalog` and `dbcc checkstorage` prior to loading any new Adaptive Server binaries, including the master database. If the `dbcc` commands reveal problems, check the *Error Messages and Troubleshooting Guide* for actions necessary to fix the problem. If the error is not listed in the manual, call Sybase Technical Support.

After you have verified that your database is error free, back up your `$$SYBASE` directory in case you need to roll back to the original version of the software.

Installing the Adaptive Server version 15.0.2 binary overlay

- 1 Read “Special Installation Instructions” in the release bulletin for the latest information that may affect specific 15.0.2 configurations.

This document may also contain last minute installation information for the 15.0.2 installation.

- 2 After backing up your databases, shut down Adaptive Server and back up your `SYBASE` directory.
- 3 Installing Adaptive Server 15.0.2 overwrites current Adaptive Server software. Prior to installing 15.0.2, make sure that your databases are error-free as described in “Backing up Adaptive Server” on page 37 and that your `SYBASE` directory is backed up.
- 4 Use InstallShield to load the new software from the CD.

For Windows installations, InstallShield starts automatically.

For UNIX installations, change to the CD drive and enter `./setup`.

- 5 Install Adaptive Server 15.0.2 over the `$$SYBASE` installation path.

Note Once files are loaded into your `$$SYBASE` directory, InstallShield asks if you want to configure the new installed server. Deselect configuring the new server, and click continue to finish the installation.

- 6 Restart Adaptive Server.

The shell command is:

```
$$SYBASE/$$SYBASE_ASE/install/startserver [ -f  
RUN_server_name file ].
```

- 7 Run `select @@version`. The server should now be at version 15.0.2.
- 8 In Adaptive Server version 15.0.2, several changes have been made to the system stored procedures and many new error messages have been added. You must perform the post-installation tasks to make these changes available.

Refer to the “Post-upgrade tasks” on page 70 before running *installmaster* or *instmsgs.ebf*.

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 configuration guide for your platform.

Topic	Page
Verifying that servers are running	39
Verifying that you can connect to servers	40
Connecting to Adaptive Server via Sybase Central	40
Setting the System Administrator password	41
The SAMreport	42
Installing sample databases	42

Verifying that servers are running

Note If you have already run the following command do not do it again. Running it more than once can cause problems.

To determine whether the servers you installed are running:

- 1 From a UNIX command line, set the Sybase environment variables by entering:


```
source $SYBASE/SYBASE.csh
```
- 2 Use `showserver` to display all Adaptive Server-related processes that are on the system


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

Verifying that you can connect to servers

Use `isql` to perform a quick test. Source `SYBASE.csh`. Then use `isql` to connect to servers:

- At the command prompt, enter:

```
isql -Usa -P<password or leave it blank>
      -Sserver_name
```

where `server_name` is the Adaptive Server name.

-
- 1 **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, change the System Administrator’s password. See “Setting the System Administrator password” on page 41.
-

The command prompt is displayed if the login is successful.

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

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

The Adaptive Server version number is displayed. The output should show Adaptive Server at version 15.0.2.

If you encounter errors, see the *Troubleshooting Guide*.

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 provides a hierarchical list of servers in the left pane and a list of details for the selected server in the right pane. To select a server, click on it in the left pane.

Note If you have already run the following command do not do it again. Running it more than once can cause problems.

- 1 From a UNIX command line, enter:

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

- 2 Then enter:


```
run $SYBASE/shared/sycen/sybcentral43
```
- 3 Select Tools | Connect, then select the server to which you want to connect from the drop-down list.
- 4 In 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.

The Adaptive Server plug-in to Sybase Central does not display all servers listed in the *interfaces* file. Instead, Sybase Central lists only those servers that you connected to earlier, or those servers that are started as Windows services.

To access a new server for the first time, select Tools | Connect to select a server listed in the *interfaces* file.

Setting the 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.

For greatest security, Sybase recommends that you set the password to have at least six characters, with a combination of letters and numbers.

The SAMreport

SAMreport is a SySAM 2.0 license reporting tool used in conjunction with the SySAM2.0 license server. It is a report generator that reports the license use of FLEXnet-licensed applications. SAMreport reports license use based on license activity recorded in the FLEXnet report logs, and on selected criteria that include product and feature name, user, date and time, and display or host.

The reports provided by SAMreport belong to three report families:

- Summary reports – summarize usage information about each feature used from the license server.
- Raw reports – show individual usage detail in textual format.
- Server reports – provide details about a given license server.

For more information on SAMreport, see Chapter 5, “Using SySAM for Asset Management” in the *Sybase Software Asset Management User’s Guide*.

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 43
- “Installing the interpubs database” on page 44
- “Installing the jpubs database” on page 45

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

Table 3-1 lists the scripts that you can use to install the sample databases.

Table 3-1: Sample database scripts

Script	Description
<i>installpubs2</i>	Installs the pubs2 sample database. 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.

Script	Description
<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. Run the <i>installpix2</i> script after you run <i>installpubs2</i>.</p> <hr/>

Default devices for sample databases

The *\$SYBASE/SYBASE_ASE/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. Sybase does not recommend this configuration because it uses valuable space that is best used for system tables. Each sample database requires 3MB on a 2K server, and multiples of 3MB on a 4K, 6K, 8K, and 16K server on your database device.

To avoid installing sample databases on the master device, do one of the following:

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

v 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 in to Adaptive Server and run the script:

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

Where:

- `server_name` – represents the destination server for the database
- `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 to use or test the image datatype. Sybase does not supply any tools for displaying image data. You must use appropriate window graphics tools to display the images after you have extracted them from the database.

For more information about running these scripts, see the configuration guide for your platform.

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`), ISO 8859-15 (`iso15`), `Roman8`, or `Roman9` (for HP-UX) character set. To display the French and German data correctly, you must set up your terminal to display 8-bit characters.

v **Installing the *interpubs* database**

- 1 Be sure `iso_1`, `iso15`, `Roman8`, `Roman 9`, or `UTF-8` 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 the EUC-JIS (eucjis), UTF-8 (utf8), or the Shift-JIS (sjis) character sets.

v **Installing the *jpubs* database**

- 1 Set your terminal to display 8-bit characters.
- 2 Verify that the EUC-JIS, Shift-JIS, or UTF-8 character set is installed as the Adaptive Server 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 43.
- 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 the *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*.

Note Sybase recommends that you remove the guest user from user databases in production systems.

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.

Installing Sybase PC-Client Products

Adaptive Server also includes PC-Client products, and this chapter describes how to install client products.

Topic	Page
Overview	47
Before you install the PC-Client products	48
Installing PC-Client products	48
Installing PowerDesigner, PowerTransfer, and InfoMaker	50
Configuring network connections for client products	51

Overview

Clients, such as Sybase Central and Open Client, are used to access Adaptive Server. See the *Open Client/Server Configuration Guide for Desktop Platforms*.

Sybase Central, ASE and SySAM plug-ins

Sybase Central allows you to connect to 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 51. 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:

- Software Developer’s Kit (SDK) for Windows platforms. SDK includes:
 - Open Client (CT-Library, DB-Library)
 - Embedded SQL/C (ESQL/C)
 - Embedded SQL/Cobol (ESQL/Cobol)

- Extended Architecture (XA)
- jConnect for JDBC
- ASE ODBC driver by Sybase
- ASE OLE DB provider by Sybase
- ASE ADO.NET Data provider by Sybase
- Language modules
- PowerDesigner
- PowerTransfer
- Infomaker

Monitor Client Library requires Open Client/C with the SDK option.

Before you install the PC-Client products

Before starting the installation process for PC-Client products, be sure that you have at least an extra 5MB of disk space, a `\temp` directory, and that you have set the 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.

Installing PC-Client products

The PC-Client CD contains several products that are each packaged with their own installer. When you insert the PC-Client CD, a menu program launches automatically. The menu program presents the list of products that can be installed from the CD and allows you to install them one at a time. The major products on the CD are described in Chapter 1, “Overview.”

These are the menu selections:

- View *readme.txt*
- Install PC-Client components 15.0.2

- Install PowerDesigner 12.0
- Install PowerTransfer 15.0
- Install InfoMaker 10.5
- View our (Sybase) Web site
- Exit

Before installing any products, review the *readme.txt* file, which contains a brief description of each of the products, lists dependencies, and contains any last-minute instructions or changes.

It is best to install the PC-Client components first because, other products use PC-client components.

Table 4-1: PC-Client system requirements

Operating system	Service pack level
Windows 2000 Professional, Server, Advanced Server, and Data Center Editions	Service Pack 4 or higher
Windows 2003	Service pack 2
Windows XP Professional	Service pack 2
Windows Server 2003 Edition x64	Service pack 2

v Installing PC-Client Components on Windows

- 1 Verify that your computer has sufficient disk space for each product.
- 2 If you are unloading components, log in using an account with administrator privileges.
- 3 Close any open applications or utilities to free memory and system resources.
- 4 Insert the PC-Client CD in the CD drive. The menu program should start automatically.

Alternatively, you can start the menu program from the Windows Start menu. Select Start | Run, and enter the following, where *X* is your CD drive:

```
X:\autorun.exe
```

- 5 Select Install PC Client Components 15.0.2 from the menu.
- 6 When the Installer starts, the Welcome window displays. Click next.
- 7 When the Sybase license agreement displays, select a country, click Agreement, then click Next.

- 8 Enter the directory path and click Next.
 - 9 Select the type of installation to be performed.
 - Typical Install – installs the default components that most users need.
 - Full Install – installs every component on the CD.
 - Customized Install – allows you to select the components to install. Certain components are automatically installed if they are required to run other selected components.
 - 10 If you select Customized Install, the next window is the Component Selection window, which allows you to specify which components to install.

Components that would be installed in a typical installation appear with a check in the check box to the left of the product name. You may select or unselect components from this list.
 - 11 Click Next.

The Summary window displays every component to be installed by the 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. You must either click Previous to return to the previous window and change your selections, or Cancel to quit the installer.
 - 12 Click Next.
 - 13 The Installer displays a progress indicator as it installs components.
- To configure client network connections to Adaptive Server, see the *Configuration Guide* for your platform.

Installing PowerDesigner, PowerTransfer, and InfoMaker

Install PowerDesigner by selecting Install PowerDesigner 12.0 from the menu.

Select Install InfoMaker 10.5 to install InfoMaker from the menu.

You must install PowerDesigner, PC-Client components, and the Microsoft .NET framework before installing PowerTransfer. To install PowerTransfer, select Install PowerTransfer 15.0 from the menu.

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 must know where the others reside on the network. This information is stored in the *interfaces* file on Windows or in a Lightweight Directory Access Protocol (LDAP) server.

Configuring *libtcl.cfg* for LDAP

In its simplest form on 32-bit LDAP drivers, the *libtcl.cfg* file is in this format:

For a 64-bit LDAP driver, the format is:

```
[DIRECTORY]
ldap=libsybdldap64.dll
```

When an LDAP server is specified in the *libtcl.cfg* file and the *libtcl64.cfg* file (when applicable to your platform), the server information is accessible only from the LDAP server. Adaptive Server ignores the *interfaces* file. Open Client and Open Server applications that use the *-i* option at start-up override the *libtcl.cfg* file and use the *interfaces* file.

To use a directory service, you must:

- 1 Configure the *libtcl.cfg* file, and the *libtcl64.cfg* file (when applicable,) 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 the *Configuration Guide* for supported LDAP URL values.

Warning! You must enter the LDAP URL on a single line.

Note On Windows x64, the *.dll* file is called *libsybdldap64.dll*.

- 2 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 | Connectivity | Open Client Directory Service Editor.
 - 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.
 - 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 twice to exit the *dsedit* utility.

For more information, see the *Configuration Guide* for your platform.

Adding a server to the *interfaces* file

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

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

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

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

Testing the Sybase Central installation

After you install Adaptive Server, Sybase Central, and the Java Runtime Environment, test the installation and network connections:

- 1 Select Start | Programs | Sybase | Sybase Central v4.3 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 window.
- 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, use the new password.
- 4 From the drop-down list, select the Adaptive Server to which you want to connect.

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, select Tools | Disconnect.
- 6 Exit Sybase Central.

If Adaptive Server fails any of these tests, see Chapter 8, "Troubleshooting." Before retrying the installation, follow the instructions in Chapter 3, "Post-Installation Tasks." For more details on Unified Agent installation and administration, see the *Unified Agent / Agent Management Console User Guide*.

Upgrading Adaptive Server

This chapter describes the various steps for upgrading Adaptive Server to version 15.0.2. If you are installing Adaptive Server 15.0.2 on top of an earlier version refer to section Installing version 15.0.2 over an existing 15.x Adaptive Server in Chapter 2, “Installing Adaptive Server”. If you need information on recovering from a failed upgrade refer to the section “Recovering from a failed upgrade” on page 118 in Chapter 8, “Troubleshooting.”

Topic	Page
Overview of the upgrade process	55
Pre-upgrade tasks	58
Upgrading to Adaptive Server 15.0.2	65
Post-upgrade tasks	70
Upgrading Backup Server, Monitor Server, and XP Server	77
Upgrading Job Scheduler	79
Upgrading high availability	84
Recovering from a failed upgrade	88

Overview of the upgrade process

You can upgrade to Adaptive Server 15.0.2 from any of these versions:

- 11.9.2.6 and later
- 12.0.x
- 12.5.x (except 12.5.3a)

Note Adaptive Server version 15.0.2 does not support upgrade from Adaptive Server version 12.5.3a.

For a server installation 11.5.x and older, Sybase recommends that you upgrade to version 12.5.4, then upgrade to version 15.0.2.

You can upgrade Adaptive Server from a 32-bit version to a 64-bit version, but you cannot move from a 64-bit version to a 32-bit version. Likewise, you can upgrade only from an earlier version of Adaptive Server to a more recent version.

Only upgrades to and from the same page size are supported. Changing the server schema from one page size to another is a database migration. See `sybmgmt` in the *Utility Guide* for information about migration.

Upgrading Adaptive Server consists of three processes:

- 1 Install Adaptive Server 15.0.2 into its own installation directory.

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

- 2 If your server is a pre-15.x version then, from the 15.0.2 server installation, perform the pre-upgrade checks on the old server using the `preupgrade` utility, which is located at `$$SYBASE/ASE-15_0/upgrade`. If necessary, fix any problems that pre-upgrade process reports. If your server is already at Adaptive Server 15.x or later most of this process is automatic when you install the new server.
- 3 `preupgrade` runs against the pre-15.0.x servers. It checks that all required databases exist and that they are large enough. It verifies that stored procedures have the text required for upgrade, and that configuration parameters are correct for the new server and for upgrading. It verifies that datatype definitions do not conflict with the new server's requirements. When `preupgrade` runs without reporting any errors, you can then:
 - a Shut down the pre-15.0.x server.
 - b Start the 15.0.2 server against the old installation.
 - c Copy files from old `$$SYBASE` directory to the new `$$SYBASE/ASE-15_0/upgrade` directory.
 - d Run the upgrade process. If problems occur, fix them and then run the upgrade process again.

- e Reinstall Sybase-supplied stored procedures to prevent any errors caused from changes to system catalogs.

Note Adaptive Server issues advisory warnings when configuration parameters are not set to the default. These do not require any action.

The preupgrade and upgrade utilities are internally called by `sqlupgrade` or `sqlupgraderes`. If your server is already at a 15.x level, you do not need to run `preupgrade` in order to move to version 15.0.2. All required changes happen automatically while ASE 15.0.2 is starting for the first time. However you must reinstall system stored procedures by running the install scripts noted in the steps above. If you are starting from a pre-15.x server you must run `preupgrade` before upgrading to the new 15.x server.

Each new version of Adaptive Server contains features that introduce new parameters, commands, reserved words, and so on. For this reason, `preupgrade` is needed to prepare the old server for upgrade. It is not necessary to start the server before running `sqlupgrade`. `sqlupgrade` will start the server when needed. However, if you are running `preupgrade`, then you will need to manually stop and start the server.

`preupgrade` runs various checks to determine how much space you must add to the old server to successfully upgrade.

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

During the pre-upgrade process, `preupgrade` returns informational messages as it checks the old server. You must fix all reported problems, and run `preupgrade` cleanly before beginning the upgrade process. Once the old server is eligible to be upgraded you need to manually shut down the old server, and start the new server against the existing databases. Then begin the upgrade process.

Following an upgrade to other platforms, Adaptive Server internally rebuilds procedures from the text in `syscomments` the first time Adaptive Server is executed. The procedure may not execute properly if it has code that drops an existing object and recreates it. This is a problem when the object already exists at the time the procedure is first executed. To avoid problems manually drop the object prior to executing the procedure for the first time.

System catalog changes during upgrade

Adaptive Server version 15.0.2 introduces some new system catalogs and some changes to existing catalogs.

The catalog upgrade in 15.0.2 may affect your existing applications. If you are upgrading, see the release bulletin for your platform for a complete list of catalogs that are affected.

Pre-upgrade tasks

Note Before you begin the upgrade process, install Adaptive Server 15.0.2 onto your system.

To ensure a successful upgrade, review the following pre-upgrade tasks and perform them as necessary. Depending on the old server configuration, you may not need to perform all pre-upgrade tasks.

- 1 Check system and upgrade requirements.
- 2 Check *RUN_server* file location.
- 3 Read and understand password encryption algorithm documentation for 15.0.2 found in the New Feature Guide. You should set the password policy option to 'allow password downgrade' to 1. Then passwords are encrypted with old and new algorithms aiding in transition should you ever downgrade from 15.0.2 version.
- 4 If you are upgrading Adaptive 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.
- 5 Stored procedure text in the syscomments table is required for upgrade.
- 6 Reserved words are resolved using quoted identifiers.
- 7 Verify that users are logged off.
- 8 Check database integrity.
- 9 Back up databases.
- 10 Dump transaction log.

- 11 Make sure that master is the default database for the “sa” user.
- 12 Prepare the database and devices for upgrade using the preupgrade utility.

- Create a subsystemdb database if it does not already exist.
- Disable auditing using the command:

```
sp_configure 'auditing', 0
```

Save the current audit settings for the pre-15.0.2 Adaptive Server using the command:

```
sp_displayaudit
```

- Disable disk mirroring.
- Verify that your SYBASE environment variable points to the location of the new Adaptive Server software files you just installed.

You also can set the OLDSYBASE, OLDSYBASE_ASE, and OLDSYBASE_OCS environment variables to the location of the server you are upgrading, to avoid having to type this path when you run sqlupgrade.

- Disable Job Scheduler.

If you do not disable Job Scheduler then the path for UNIX platforms is:

- *\$SYBASE/\$SYBASE_OCS/config/libtcl.cfg*
- *\$SYBASE/\$SYBASE_OCS/config/libtcl64.cfg*
- *\$SYBASE/\$SYBASE_SYSAM2/licenses/license.dat*

- 13 Ensure that the procedure cache size is at a minimum of 150 percent of the default procedure cache size, or between '53248' and '2147483647' 2k pages.

- 14 Copy the following files from earlier versions of Adaptive Server to their corresponding Adaptive Server 15.x installation.

For UNIX platforms:

- *\$SYBASE/interfaces*
- *\$SYBASE/\$SYBASE_ASE/<servername>.cfg*
- *\$SYBASE/\$SYBASE_OCS/config/libtcl.cfg*
- *\$SYBASE/SYSAM-2_0/licenses/license.dat*

For Windows platforms:

- `%SYBASE%\sql.ini`
- `%SYBASE%\SYBASE_ASE<servername>.cfg`
- `%SYBASE%\%SYBASE_OCS%\ini directory`

Checking system and upgrade requirements

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. Verify that your operating system is at the proper version level and has all operating system patches needed for Adaptive Server.
- 2 Determine whether your server supports an upgrade to Adaptive Server 15.0.2.
- 3 Set the environment variables by sourcing the `SYBASE.[csh, sh, env]` file. Running this script multiple times can cause problems. Only run it once.
- 4 Verify that you installed Adaptive Server 15.0.2 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.

Checking 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:

- `$SYBASE/install` directory for version 11.9.x
- `$SYBASE/$SYBASE_ASE/install` for version 12.0.x and higher

Verify that the file is renamed `RUN_servername`, where `servername` is the name of the old server. The `servername` must appear exactly as it appears in the `interfaces` file. The `RUN_servername` file for a server SYBASE is called `RUN_SYBASE`. If the `RUN_servername` file for your current Adaptive Server is named, you must change the name during the upgrade process.

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, the upgrade succeeds but you will encounter runtime problems. To upgrade stored procedures, the stored procedure text must be available in syscomments.

To upgrade successfully to Adaptive Server version 15.0.2:

- Reinstall the procedures with text, or
- Remove them and reinstall them after the upgrade.

You can also hide the text using the `sp_hidetext` stored procedure.

Reserved words

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. If you are upgrading Adaptive Server, and the identifiers in your user databases match new reserved words, errors result when you run queries, stored procedures, or applications that use these identifiers.

Note Before performing your upgrade, change the name of any user database that is a reserved word using the `sp_renamedb` stored procedure.

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. Rename all objects that use reserved words.

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

Note If you change an object name, 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`, which is 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 any time 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 have been 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.
- Quoted identifiers.
- Brackets around the identifier. For example:

```
create table [table] ( [int] int, [another int] int )
```

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 `set 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.

Preparing the database and devices for the upgrade

If you ran the pre-upgrade option in `preupgrade`, the utility scanned the system catalogs and calculated how much additional space you need 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.

v Increasing default database sizes

You can use `sp_helpdb` or Sybase Central to determine the current size of the master, model, `sybssystemprocs`, and `tempdb` databases.

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 The model database cannot be larger than the tempdb. If your changes would cause the model database to be larger than the tempdb, then you must first increase the size of tempdb, before increasing the size of the model database.

- 3 Verify the size of each database. Enter the following, where *database_name* is the name of the system database you are checking:

```
1> sp_helpdb database_name
2> go
```

Create a *sybssystemdb* database

In version 11.9.x, the *sybssystemdb* database was required only for servers using two-phase commit transactions. Beginning with 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).

If you have a *sybssystemdb* database

If you have a *sybssystemdb* and the database is at least 4MB, 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.

sybssystemprocs

sybprocsdev is the default name for the device-holding database sybssystemprocs. It is frequently referred to as the sybssystemprocs device in Adaptive Server.

Verify that the sybssystemprocs database is large enough. For an upgrade, the default size for sybssystemprocs is 124MB, or enough free space to accommodate the existing sybssystemprocs database, and the largest catalog that is to be upgraded, plus an additional 10 percent for logging upgrade changes. You may need more space if you are adding user-defined stored procedures. 124MB accommodates additional internal data structures, but does not account for the possibility of a larger number of user-defined system procedures.

Upgrading to Adaptive Server 15.0.2

To assist with the upgrade process, Sybase recommends you set the *OLDSYBASE*, *OLDSYBASE_ASE*, and *OLDSYBASE_OCS* environment variables to point to the old installation of Adaptive Server. If you set these environment variable to the old installation of the Adaptive Server, the upgrade utility automatically populates input fields with the correct information as you upgrade.

Set the *OLDSYBASE_ASE* variable to the *SYBASE_ASE* appropriate to your older server. If you are upgrading from 12.0 for example, it should be *ASE-12_0*. If your preupgrade version is earlier than 12.0, specify it as null:

```
setenv OLDSYBASE_ASE ""
```

This is required if you are using the *sqlupgraderes* utility to perform your 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 66.
- `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 68.

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 your old server is running.
- 2 Verify that you have performed the tasks mentioned on the window, before beginning the upgrade.
- 3 Source `SYBASE.csh` before running `sqlupgrade` if you have not already done that.
- 4 Execute `sqlupgrade`, by entering at the UNIX prompt:

```
    $SYBASE/$SYBASE_ASE/bin/sqlupgrade
```

A `sqlupgrade` window now displays.
- 5 The first screen warns you to back up the databases, devices, and so on. Click Next to proceed to the next screen.
- 6 Click OK.
- 7 The Next screen prompts you to enter the old Sybase directory, and the old Adaptive Server directory.
- 8 Click OK. The Adaptive Server Selection window displays.
- 9 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.
- 10 Click OK. The SA Password window is displayed.
- 11 Enter the password for the System Administrator (sa) login.
- 12 Click OK. The Specifying Upgrade Options window is displayed.

- 13 Specify an upgrade option.

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.

- 14 Adaptive Server 11.9.x did not automatically create a sybsystemprocs database and device. Unless you have created one, the eligibility test fails. The Back button launches the “Specify sybsystemprocs” window.

Enter the full path to the sybsystemprocs device and at least 124MB for the sybsystemprocs database. Click OK to continue with the eligibility test.

- 15 Click OK to upgrade your server.

sqlupgrade displays the Status Output window. View the Status Output window for the completion status and informational messages about the upgrade process.

Warning! 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.

You can also check the progress of the upgrade by viewing the sqlupgrade log in `$SYBASE/$SYBASE_ASE/init/logs/sqlupgradeMMDD.NNN MM` is the month, `DD` is the date, and `NNN` is a three-digit number identifying the srvbuild session.

After all pre-upgrade checks are successful, the old server is shut down and the new Adaptive Server dataserver binary is started on the old master device. sqlupgrade internally runs the upgrade binary to perform the upgrade.

If the upgrade is successful, a “Done” message is displayed in the Status Output window.

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, press Ctrl+C.

- 16 The upgrade process:
- Creates the *RUN_servername* file
 - Runs the *installmaster* script

- 17 Check that the upgrade has been successful by logging in to the server and checking `select @@version` to confirm that it is at 15.0.2. Also, `sp_configure "upgrade version"` should return 15000.
- 18 For information about compiled objects, see Chapter 8, “Troubleshooting.”

When the upgrade completes successfully, you can:

- Click OK to specify another server to upgrade, or
- Exit `sqlupgraderes`, and go to “Post-upgrade tasks” on page 70.

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:

- 1 Edit the sample resource file included in your Adaptive Server distribution in `$$SYBASE/$SYBASE_ASE/init/sample_resource_files` directory. Use an ASCII text editor to edit the resource file as described in Appendix A, “Alternative Installation Methods.”
- 2 Run the `sqlupgraderes` utility. This utility creates a new file containing the values you specified for that upgrade session, and writes it to the `sqlupgradeMMDD.NNN-server_name.rs` file in `$$SYBASE/$SYBASE_ASE/init/logs/`.

Where *server_name* is the name of the upgraded server, *MM* is the month, *DD* is the date, and *NNN* is a three-digit number identifying the `srvbuild` session.

For details on the attributes used in a resource file, see “Resource file attributes for upgrading Adaptive Server,” below.

Resource file attributes for upgrading Adaptive Server

Table 5-1 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/$SYBASE_ASE/init/sample_resource_files` 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 5-1: Resource file attributes for upgrading Adaptive Server

Attribute	Default value [other options]
<code>sybinit.release_directory</code>	<code>\$SYBASE [path = <i>_name_of_old_release</i>]</code>
<code>sybinit.product</code>	<code>sqlsrv</code>
<code>sqlsrv.server_name</code>	<code><i>server_name</i></code>
<code>sqlsrv.new_config</code>	<code>No</code>
<code>sqlsrv.sa_login</code>	<code><i>current_login</i></code>
<code>sqlsrv.sa_password</code>	<code><i>current_password</i></code>
<code>sqlsrv.do_upgrade</code>	<code>Yes</code>
<code>sqlsrv.do_reserved_word_check</code>	<code>Yes</code>

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.

Source the version 15.0.2 `SYBASE.csh` file before running `sqlupgraderes` and set `OLDSYBASE` and `OLDSYBASE_ASE`.

To execute `sqlupgraderes`, at the UNIX prompt, enter the following, where *resource_file* specifies the resource file containing the attributes that describe the server to upgrade:

```
$SYBASE/$SYBASE_ASE/bin/sqlupgraderes -r resource_file
```

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 you have upgraded, you do not need to run update statistics on any tables. The upgrade process does not cause any changes in existing statistics.

If you have modified any of the Sybase system stored procedures without changing the names, executing some of these scripts will overwrite your modifications. In this case, you must reload the modified stored procedures after you run these scripts.

- If the RTDS Adaptive Server Messaging functionality is enabled and the RTDS ASE MESSAGING feature is being used in this installation of Adaptive Server, then run the `installmsgsvss` command. Run the `installmsgsvss` command that is included with this release. Save the output of this step to an operating system file.

```
isql -Usa -P<sa password> -S<server name> -n- i
$SYBASE/$SYBASE_ASE/scripts/installmsgsvss -o<output file>
```

Run the *instmsgs.ebf* script

Follow these instructions if you are upgrading from Adaptive Server version 15.0 to Adaptive Version 15.0.2.

Adaptive Server 15.0.2 now has a way to restore error messages back to the 15.0.x release from which you are upgrading. Before installing the new messages for 15.0.2 you must run `uninstmsgs.ebf` to safeguard the modified messages in your master device.

Note If you are upgrading from a pre 15.x release then ignore this step to produce the *restore_msgs.ebf* file and proceed to run the *instmsgs.ebf* file.

To run `uninstmsgs.ebf` use the following command:

```
isql -S -Usa -P -w1000 -iuninstmsgs.ebf-orestore_msgs.ebf
```

Run the `instmsgs.ebf` script after you have completed running the `uninstmsgs.ebf` script.

If you are a localized customer:

- First run the `instmsgs.ebf` command, and then install any localized language using `langinstall`, `sqlloc`, or `syconfig`.
- If you run `instmsgs.ebf` after installing the 15.0.2 localized message, this script may delete some new messages.

If you need to undo the `instmsgs.ebf` changes, run the following script once you have downgraded to the version from which you upgraded.

```
isql -S -Usa -P <restore_msgs.ebf>
```

Restoring 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 upgrading, use `sp_configure` to set them back to their previous values.
- 2 Use `sp_dboption` to reset any database options you disabled before upgrading.
- 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 15.0.2.
- 4 Verify procedure cache allocation after upgrading.

This section is for information only. Adaptive Server ensures that procedure cache size after upgrade is the same as before upgrade, unless the original size was less than the default value.

If you are upgrading from a version of Adaptive Server that is earlier than version 12.5.x, 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 is set to the default value.

With the introduction of dynamically reconfigured memory configuration parameters in Adaptive Server 12.5.x, an increase in Adaptive Server's memory use does 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 much more memory to run Adaptive Server 15.0.2 than older versions.

You may increase the procedure cache size during runtime, using `sp_configure`. You can use the `verify` option in `sp_configure` to verify any changes you make to the configuration file without having to restart Adaptive Server. 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 Verify data cache allocation after upgrading. This section is for information only. Adaptive Server ensures that all the data cache sizes after upgrading are 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 was 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 allocated to the default data cache.

Adaptive Server 12.5.x and later 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.x and later.

In pre-12.5 Adaptive Server, it meant all the memory left after allocating memory for other configurations, procedures, and user-defined named caches. In Adaptive Server 12.5.x and later, it means a default value of 8MB.

During the upgrade process, Adaptive Server ensures that the default data cache size remains the same. Therefore, during the pre-upgrade process, the size of the default data cache is obtained and written to the configuration file as an absolute value, not as “DEFAULT”. This enables Adaptive Server 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 allocates a default data cache of size 8MB.

- 7 If you unmirrored devices, remirror them, using the disk remirror command.
- 8 For information about upgrading compiled objects, see “Upgrading compiled objects with dbcc upgrade_object” on page 121.
- 9 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 tran` command 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> dump database sales to dumpdev
2> go
1> dump transaction sales with truncate_only
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 in 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
1> rs_zeroltm dataserver, database
2> 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 15.0.x, 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 4, “Installing Sybase PC-Client Products.”

Restoring replication after upgrade

Restore the replication functionality:

- 1 Log in to 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 in to 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 in to 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 must log in to the replicate Replication Server.

- 6 If you are using Rep Agent, log in to 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.

Reenabling auditing

Reenable all auditing options for stored procedures using the audit options you recorded during pre-upgrade. Reenter the audit options using `sp_audit`.

For example, if you had enabled stored procedure auditing for `sp_addlogin` stored procedure in your server before upgrade, run the following `sp_audit` command to reenabling auditing on `sp_addlogin`:

```
sp_audit "exec_procedure", "all", "sp_addlogin", "on"
```

Adaptive Server 12.5.x and later includes the following global audit options for `sp_audit`:

- security
- dbcc

The functionality provided by the following `sp_auditoption` options is now covered by the `sp_audit security` global audit option:

- server boot
- roles

These options are turned off by default and are not turned on even if you used `sp_auditoption "all", "on"` to turn on all global audit options in a version earlier than 11.5.

To turn on these options, use:

```
1> sp_audit "security", "all", "all", "on"
2> go
1> sp_audit "dbcc", "all", "all", "on"
2> go
```

For more information, see the *Reference Manual*.

Special considerations for threshold procedures for audit segments

If your installation used a threshold procedure to archive the sysaudits table using T-SQL similar to the following example:

```
INSERT MyPre15SysAuditHistoryTable SELECT * FROM
sysaudits_0n
```

Where *n* corresponds to the sysaudits table number 1 – 8 and *MyPre15SysAuditHistoryTable* is a table defined prior to Adaptive Server version 15.0.2, then you must alter *MyPre15SysAuditHistoryTable* to add a *nodeid* column using the following command:

```
alter table MyPre15SysAuditHistoryTable
add nodeid tinyint NULL
```

See the *Reference Manual* for system tables sysaudits01 through sysaudits08 for further details about the sysaudits tables.

Shutting down SySAM servers

Shut down the old SySAM servers if they are no longer used by earlier versions of Adaptive Servers.

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.

v **Upgrading Backup and Monitor Servers**

The procedures for upgrading Backup Server and Monitor Server are similar.

- If you are upgrading during the initial installation of Adaptive Server 15.0.2, select Upgrade Existing Servers from the drop-down menu when the Installer prompts Build or Upgrade Adaptive Server, Backup, and Monitor Servers.

This launches the sqlupgrade utility. Click OK.

- If you are upgrading after the initial installation, start the sqlupgrade utility from the command line. Enter:

```
    $SYBASE/$SYBASE_ASE/bin/sqlupgrade
```

- 1 A warning message appears. Click OK to continue.
- 2 The next screen prompts you to enter the old Sybase directory and the old Adaptive Server directory.
- 3 From the Server selection window, select the server to upgrade. Click OK.
- 4 When you are prompted for the password, enter the password for the System Administration, and click OK.
- 5 The Status output window display the upgrade progress. When the upgrade is complete it displays, "Done." Click OK.
- 6 You return to the sqlupgrade window. You can either upgrade another server, or click Exit to leave the utility.

Repeat this procedure to upgrade Monitor Server.

Upgrading Job Scheduler

After upgrading to the new Adaptive Server, perform the following steps to upgrade Job Scheduler.

Note You must have `$SYBASE/$SYBASE_OCS/bin` directory in `$PATH` so the `isql` executable is accessible.

v Upgrading Job Scheduler

1 Copy the directory services entry for JSAGENT (or jsagent) from the old server to the new 15.0.2 server.

2 Make sure your 15.0.2 is running.

3 Stop the Job Scheduler. From `isql`, issue:

```
sybmgmtdb..sp_js_wakeup "stop_js", 1
go
```

4 Ensure that at least 9000 locks are configured. If the number of locks on the server is less than 9000, use `sp_configure` to increase the number of locks:

```
sp_configure "number of locks", 9000
```

5 Upgrade the internal Job Scheduler SQL using `isql`:

```
use sybmgmtdb
go
dbcc upgrade_object
go
```

There is no need to specify any parameters.

6 Restart Adaptive Server.

7 If the `sybmgmtdev` device is less than 100MB, change the device to be a minimum of 100MB. From `isql` enter the following, where *nn.nM* is the number of megabytes needed to get to 100 MB:

```
use master
go
disk resize name = 'sybmgmtdev' , size = 'nn.nM'
go
```

8 Optionally, add more log space. On some 64-bit platforms more space is need for the `sybmgmtdb` log. From `isql` enter:

```
use master
```

```
go
alter database sybmgmtdb LOG on sybmgmtdev=20
go
```

- 9 Upgrade the sybmgmtdb. Using isql, run the *installjsdb* script that is included with this release and save the output to a file:

```
isql -Usa -Psa_password -Sservername -n
-i$SYBASE/$SYBASE_ASE/scripts/installjsdb
-ooutput_file
```

- 10 Enable Job Scheduler to start when Adaptive Server starts. From isql enter the command:

```
sp_configure "enable job scheduler", 1
```

- 11 Start Job Scheduler. From isql enter the command:

```
sybmgmtdb..sp_js_wakeup "start_js", 1
go
```

Upgrading Job Scheduler templates

After upgrading to the new Adaptive Server, perform the following steps to upgrade templates and jobs created from Job Scheduler templates.

Note For the Adaptive Server version 15.0.2 release, there are several changes that impact Job Scheduler templates. These changes make some templates incompatible with earlier versions of Adaptive Server. These templates are version 3.0 in the XML files.

v Upgrading Job Scheduler templates

- 1 Disable Job Scheduler before upgrading Job Scheduler templates.
- 2 Update any environment variables, scripts, or applications that refer to the Job Scheduler directory path. The Job Scheduler directory has been renamed and moved beneath the *ASE-15_0* directory. New location is *\$SYBASE/\$SYBASE_ASE/jobscheduler*.

Directories below *jobscheduler* remain the same.

- 3 If you have any of your own files in the old Job Scheduler directory, copy them to the new 15.0.2 directory. If you are installing Adaptive Server over the old server directories, the installer will have moved your files to the new *jobscheduler* directory already.

- 4 If you have modified any of the Sybase-supplied templates stored procedures or XML documents, do not overlay the new templates with your versions. If you do, you will lose the enhancements provided by the revised templates. Carefully merge your template changes into the Sybase template files, or better yet, rename your modified templates.

Note Sybase recommends that if you modify Sybase-supplied templates that you save the changes under new names and in new files.

- 5 You may need to make minor modifications to jobs you created from 2.0 or earlier templates. In some templates, parameters changed from varchar(5) to int. The following table lists the templates that changed in version 2.1 and the changes you need to make to jobs created from them.

Table 5-2: Changed Job Scheduler templates

Template	Modified files	Change description	Modify jobs created from this template
dump database	<i>SybBackupDbToDiskTemplate.xml</i> <i>jst_dump_databases</i>	Changed the datatype for the @use_srvr_name parameter from varchar(5) to int	Required, modify job SQL to change @use_srvr_name to int datatype.
dump database log	<i>SybBackupLogToDiskTemplate.xml</i> <i>jst_dump_log</i>	Changed the datatype for the @truncate_flag and @use_srvr_name parameters from varchar(5) to int	Required, modify job SQL to change @truncate_flag and @use_srvr_name to int datatype.
update statistics	<i>SybUpdateStatsTemplate.xml</i> <i>jst_update_statistics</i>	Changed the datatype for the @index_flag parameter from varchar(5) to int	Required, modify job SQL to change @index_flag to int datatype.
rebuild indexes	<i>SybRebuildIndexTemplate.xml</i> <i>jst_reorg_rebuild_indexes</i>	Changed the datatype for the @dump_flag parameter from varchar(5) to int	Required, modify job SQL to change @dump_flag to int datatype.
rebuild table	<i>SybRebuildTableTemplate.xml</i> <i>jst_reorg_rebuild_tables</i>	Changed the datatype for the @dump_flag parameter from varchar(5) to int	Required, modify job SQL to change @dump_flag to int datatype.
reclaim indexes	<i>SybReclaimIndexTemplate.xml</i> <i>jst_reclaim_index_spac</i>	Required, modify job SQL to change @dump_flag to int datatype.	Required, modify job SQL to change @resume_flag to int datatype.
reclaim tables	<i>SybReclaimTableTemplate.xml</i> <i>jst_reclaim_table_space</i>	Changed the datatype for the @resume_flag parameter from varchar(5) to int	Required, modify job SQL to change @resume_flag to int datatype.

- 6 Some Job Scheduler templates were modified to support new 15.0.2 Adaptive Server functionality. The changes are related to new parameters for specifying a partition name or a datachange value for the Adaptive Server command that have added these options. If you have jobs created from any of the enhanced templates, modify the SQL of the jobs for 15.0.2 servers.

Note If you have a job scheduled to run on a pre-15.0.2 server and you need it to run on 15.0.2 servers too, leave the existing job alone and create a new job for 15.0.2 servers as the job commands will be different.

Jobs intended to run on pre-15.0.2 servers do not need to be modified. The following table lists the templates that changed in version 3.0 and the changes you need to make to jobs created from them.

Note All templates listed below, except Delete Statistics, are incompatible with pre-15.0.1 servers. Do not use them to create jobs that will be scheduled on pre-15.0.1 servers; you must use the 2.1 or 2.2 versions for pre-15.0.1 servers.

Table 5-3: Modified Job Scheduler templates

Template	Modified file	Change description	Modify jobs created from this template
delete statistics	<i>SybDeleteStatsTemplate.xml</i> <i>jst_delete_statistics</i>	Added @ptn_name as third parameter.	Optional.
update statistics	<i>SybUpdateStatsTemplate.xml</i> <i>jst_update_statistics</i>	Added @ptn_name as fifth parameter and @datachg_threshold as the 10th parameter and added reference	Required, modify job SQL to include values (or NULL) for new parameters.
rebuild indexes	<i>SybRebuildIndexTemplate.xml</i> <i>jst_reorg_rebuild_indexes</i>	Added @ndx_ptn_name as third parameter.	Required, modify job SQL to include value (or NULL) for new parameter.
reclaim indexes	<i>SybReclaimIndexTemplate.xml</i> <i>jst_reclaim_index_space</i>	Added @ptn_name as third parameter	Required, modify job SQL to include value (or NULL) for new parameter.
reclaim tables	<i>SybReclaimTableTemplate.xml</i> <i>jst_reclaim_table_space</i>	Added @ptn_name as second parameter	Required, modify job SQL to include value (or NULL) for new parameter.

Template	Modified file	Change description	Modify jobs created from this template
multiple	<i>jst_get_freespace,jst_get_usedspace</i>	Replace reserved_pgs and data_pgs with reserved_pages and data_pages	No impact on job SQL.

7 Install the Template stored procedures.

- Move to the JS template stored procedure directory. For example:

```
cd $SYBASE/$SYBASE_ASE/jobscheduler/Templates/sprocs
```

- Run the stored procedure install script for each server being upgraded:

```
installTemplateProcs <servername> <username> <password>
```

Note Template stored procedures should be upgraded on all Job Scheduler servers and target servers upgraded to Adaptive Server version 15.0.2. Do not install them on pre-15.0.2 servers.

8 Install the template XML documents.

- 1 Move to the JS template XML directory. For example:

```
cd $SYBASE/$SYBASE_ASE/jobscheduler/Templates/xml
```

- 2 Run the XML install script on 15.0.1 servers with Job Scheduler installed:

```
installTemplateXml <servername> <machinename> <serverport>  
<username> <password> [<language_code>]
```

Use “en” for the *language_code* or you may omit the parameter completely as “en” is the default.

Note Template XML should be upgraded on all Job Scheduler servers upgraded to Adaptive Server version 15.0.1. Do not install them on pre-15.0.1 servers or servers where Job Scheduler is not installed.

Upgrading high availability

Using the following steps to upgrade high availability-enabled Adaptive Server in active-active configuration:

- 1 If the High Availability (HA) functionality is enabled in your pre-15.0.2 server and the HA feature is being used in this installation of Adaptive Server, the *installhasvss* script must be run. Run the *installhasvss* script that is included with this release and save the output to an operating system file.

```
isql -Usa -P<sa password> -S<server name> -n
-i$SYBASE/$SYBASE_ASE/scripts/installhasvss -o<output file>
```

- 2 Drop the high availability companionship. On the secondary server, use isql:

```
sp_companion <primary_server_name>, "drop"
go
```

- 3 Use the cluster system command to stop monitoring resources associated with Adaptive Server on each cluster node.
- 4 Follow the upgrade steps in “Upgrading to Adaptive Server 15.0.2” on page 65 to upgrade each of the Adaptive Servers separately.
- 5 Run the new Adaptive Server *installmster* script against the newly upgraded Adaptive Servers.
- 6 Run the new Adaptive Server *installhasv* script against the newly upgraded Adaptive Servers.
- 7 In the new Adaptive Server installation area, follow the instructions in *Using Sybase Failover in A High Availability System* to configure the permission and ownership for *\$SYBASE/\$SYBASE_ASE/bin/sybha* and *\$SYBASE/\$SYBASE_ASE/install/sybhauser*.
- 8 In the new Adaptive Server installation area, modify high availability-related files such as the *RUN_server_file*, and the *SYBASE.csh* and *SYBASE.sh* files, if those files are required on that cluster platform.
- 9 Reconfigure each resource associated with Adaptive Server depending on platform-specific requirements. For example, on Veritas Cluster, the *HAase* resource properties, the *RUN_server_file*, and *Sybase_home* need to be updated.

- 10 Manually restart Adaptive Server on each cluster node with trace flag 2209. Use the Adaptive Server command-line option -T2209.

Warning! Do not use trace flag 2209 after the Adaptive Server companionship is re-established.

- 11 Use the cluster system command to restart monitoring resources associated with Adaptive Server on each cluster node.
- 12 Re-establish companionship by using isql to issue:

```
sp_companion <primary_server_name>, configure
go
```

If user databases exist on the secondary server, warning message may display on the console. There is no need to take any action. You can safely ignore them. The messages look similar to:

```
Msg 18739, Level 16, State 1:
Server 'svr2', Procedure 'sp_hacmpcfgvrfy', Line 102:
Database 'svr2_db1': a user database exists. Drop this
database and retry the configuration again.
```

- 13 Use the cluster command to offline and then online resource associated with Adaptive Server. Use isql to connect to each Adaptive Server and enter the command to verify the correct server companionship:

```
sp_companion
go
```

Console output on primary server is similar to the following:

```
Server 'svr1' is alive and cluster configured.
Server 'svr1' is configured for HA services.
Server 'svr1' is currently in 'Primary normal' mode.
(return status = 0)
```

Console output on secondary server is similar to the following:

```
Server 'svr2' is alive and cluster configured.
Server 'svr2' is configured for HA services.
Server 'svr2' is currently in 'Secondary normal'
mode.
(return status = 0)
```

Use the following steps to upgrade high availability-enabled Adaptive Servers in active-passive configuration:

- 1 Use the cluster system command to bring Adaptive Server offline.

- 2 Use the cluster system command to stop monitoring resources associated with Adaptive Server.
- 3 Follow the upgrade steps provided in “Upgrading to Adaptive Server 15.0.2” on page 65 to upgrade the Adaptive Server to the new version.
- 4 Run the new Adaptive Server *installmaster* script against the newly upgraded Adaptive Server.
- 5 In new Adaptive Server installation area, modify the high availability-related files such as the *RUN_server_file*, and the *SYBASE.csh* and *SYBASE.sh* files, if those files are required on that cluster platform.
- 6 Use the cluster system command to reconfigure resource properties associated with Adaptive Server to reflect the new Adaptive Server installation location.
- 7 Manually restart Adaptive Server.
- 8 Use the cluster system command to restart monitoring resource associated with Adaptive Server.
- 9 Use the cluster command to offline and then online resource associated with Adaptive Server on the same node. Use the cluster command to switch resource associated with Adaptive Server to another node and then switch back.

Migrating from 32-bit to 64-bit versions on Different Computers

sqlupgrade upgrades Adaptive Server only from one version to another. This section describes how to migrate a 32-bit version of Adaptive Server to a 64-bit version of Adaptive Server.

Following an upgrade and migration to other platforms, Adaptive Server internally rebuilds procedures from the text in *syscomments* the first time it is executed. The procedure may not execute properly if it has code that drops the object if it exists and recreates the object. This is a problem when the object already exists at the time the procedure is first executed. To avoid problems manually drop the object prior to executing the procedure for the first time.

To migrate Adaptive Server from the 32-bit to the 64-bit version, you must first install and configure the 64-bit operating system.

- Method #1: Dump and load
- 1 Run dbcc checks on all databases in the 32-bit Adaptive Server and make sure they run cleanly.
 - 2 Create a 64-bit Adaptive Server in a new directory.
 - 3 Create devices and databases to match those in the 32-bit server. Make sure that the sysusages mapping is correct. See document ID #1324 at <http://my.sybase.com/detail?id=1324>.

Note Allow 10 percent additional space for the sybsystemprocs database.

- 4 Dump the databases on the 32-bit server.
- 5 Load the databases to the 64-bit server.
- 6 If you have partitioned tables, update partition statistics.
- 7 Run dbcc checks on the 64-bit server and make sure they run cleanly.
- 8 For information on upgrading compiled objects, see “Upgrading compiled objects with dbcc upgrade_object” on page 121.

Method #2: BCP data out and in

If you have DDL scripts to create devices, databases, tables, rules, stored procedures, triggers, and views, you can use bcp to move data out of the old database and into the new.

If you do not have DDL scripts, use the ddlgen utility to re-create the schema for the Adaptive Server you are about to upgrade. See the *Utility Guide*.

- 1 Run dbcc checks on databases to verify data integrity.
- 2 Use bcp to extract all the data from all the tables in the databases.
- 3 Create a new 64-bit Adaptive Server in a new directory.
- 4 Create devices, databases, and tables.
- 5 Use bcp to bulk copy data into tables.
- 6 Re-create all views, triggers, stored procedures.
- 7 Run dbcc checks on the 64-bit server and make sure they run cleanly.

Method #3: Replacing the binary

The third method requires you to modify some key files, re-create stored procedures, then replace the binary.

- 1 Run dbcc checks (checkdb, checkalloc, checkcatalog and checkstorage) on all databases in the 32-bit Adaptive Server and make sure they run cleanly.
- 2 In a new directory, copy the files for the 64-bit Adaptive Server.
- 3 Shut down the 32-bit server.

- 4 Copy the *interfaces* file and the configuration file from the 32-bit `$$SYBASE` directory to the 64-bit `$$SYBASE` directory.
- 5 Copy the 32-bit `$$SYBASE/$$SYBASE_ASE/install/RUN_<server>` file to the equivalent 64-bit `$$SYBASE/$$SYBASE_ASE/install` directory.
- 6 Edit the `RUN_server` file to reflect the new location of the *interfaces*, configuration, and log files.
- 7 Remove all references to the 32-bit `$$SYBASE` directories from your `$PATH` definition.
- 8 Change to the 64-bit `$$SYBASE` directory and source the `SYBASE.csh` script (Bourne shell).
- 9 Change to the 64-bit `$$SYBASE/$$SYBASE_ASE/install` directory and execute:

```
startserver -f RUN_server
```
- 10 After the 64-bit server starts, run `installmaster`, `installmodel`, and `instmsgs.ebf`.
- 11 If you used `dbccdb` for `dbcc` checkstorage, run `installdbccdb`. This re-creates some tables in the `dbccdb`, and you may lose history data.
- 12 Drop and re-create all your compiled objects, such as stored procedures, triggers, views, and defaults.
- 13 If you have partitioned tables, update partition statistics.
- 14 Run `dbcc` again on all databases to verify that they run cleanly.

Recovering from a failed upgrade

The procedure for recovering from a failed upgrade is:

- 1 Fix the underlying problem
- 2 Rerun the upgrade utility, `$$SYBASE/$$SYBASE_ASE/upgrade/upgrade`
- 3 Reinstall system stored procedures using the scripts from the new version

```
isql -Usa -P<password> -i
```

```
$SYBASE/$SYBASE_ASE/scripts/installmaster
```

Note If Adaptive Server crashes (e.g., if there's a power failure) while upgrade is in progress, when it is restarted it will attempt to upgrade databases while it is booting. Messages for this upgrade will go to Adaptive Server's console (the terminal or window where it was started). After this finishes, and the installation is fully recovered, re-run the upgrade utility. This assures that all post-upgrade configuration work is done. Then reinstall system stored procedures as previously described

- During the upgrade process, the log may get full due to the catalog changes that are made. If so, log in to the 15.0.2 server using isql, and issue:

```
isql> dump tran dbname with truncate_only
```

If "truncate_only" indicates that it could not truncate the log, try again using option "no_log" instead. If this too does not truncate the log successfully, then you will need to use "alter database" to extend the database as explained below.

This frees the log space, and allow the upgrade process to continue.

- In some cases, the pre-upgrade space estimations might be insufficient for the data copy phase of upgrade. In this case, you may get an error that there is insufficient space in the system segment for the upgrade. The upgrade process hangs, 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 15.0.2 Adaptive Server, alter database allows you to specify the size to alter the database with the unit-specifier 'm', or 'M'.

Downgrading from Adaptive Server 15.0.2

This chapter provides instructions for downgrading from Adaptive Server 15.0.2.

Topic	Page
Downgrading from 15.0.2 release	91
Basic downgrade steps	92
Post-downgrade tasks	94
Additional steps required if you used 15.0.2 features before downgrading	96
Returning to Adaptive Server version 15.0.2	99

Downgrading from 15.0.2 release

This section describes how to downgrade from Adaptive Server 15.0.2 to an earlier Adaptive Server 15.0 or 15.0.1 release. Downgrading to a 12.5.x or earlier release is not supported. Instructions to downgrade to a 15.0 or 15.0.1 version include:

- Basic downgrade steps using new system procedure `sp_downgrade`.
- Post-downgrade tasks.
- Information regarding downgrading if you used 15.0.2 features.

- Returning to Adaptive Server version 15.0.2.

Note Adaptive Server 15.0.2 is not an upgrade release if you are starting from a 15.x version of Adaptive Server. For all pre-15.x versions Adaptive Server 15.0.2 is an upgrade release. Also, Sybase recommends that you do not use the new features offered in Adaptive Server version 15.0.2 until you are committed to staying with 15.0.2.

Note Downgrade is only available from Adaptive Server 15.0.2 to Adaptive Server versions 15.0 and 15.0.1. You cannot downgrade directly to version 12.5.x or earlier. If you must downgrade to a release earlier than 15.0, you need to restore databases from dumps made on that older server.

Basic downgrade steps

This section describes the basic steps for downgrading from Adaptive Server 15.0.2 to a 15.0 or 15.0.1 version. Additional actions are required if you used 15.0.2 features. If that is the case, follow the steps in the “Additional steps required if you used 15.0.2 features before downgrading,” section of this chapter.

Adaptive Server 15.0.2 provides a new procedure called `sp_downgrade` that is used to help with the downgrade process. This procedure requires 'sybase_ts_role' to be on. The 'sa_role' or 'sso_role' permissions are expected when running this procedure. For more detailed information on this new system procedure and its syntax see ‘The New System Procedures’ section in the ‘System Changes’ chapter of the *Adaptive Server New Feature Guide for 15.0.2*.

Note Returning to 15.0.1 ESD#3 is preferable, because this version can handle a number of error conditions that may arise as a result of downgrading for several reported errors.

If you return to a pre-15.0.1 ESD3 version of Adaptive Server, no error handling is provided. Using the new features can result in inexplicable behavior, misleading error messages and stack traces.

- 1 Back up all of your databases and the `$SYBASE` release area prior to the 15.0.2 downgrade.
- 2 Start the 15.0.2 server in single user mode using `dataserver -m` option to guarantee that no other user can access Adaptive Server while you are going through the downgrade steps. See *Utility Guide* for more details on starting the server in single user mode.
- 3 Make sure Adaptive Server 15.0.2 is prepared for downgrade by executing.

```
sp_downgrade 'prepare', <'version'>
```

The value of “version” may be written as “15.0” or “150” or “15.0.1” or “1501”. However, the version entered should be the Adaptive Server version you want to downgrade to.

Note If you have used the 15.0.2 features for encrypted columns and

- You are reverting back to Adaptive Server 15.0 GA or Adaptive Server 15.0 ESD #1 you must specify the version as '15.0'.
 - You are reverting back to Adaptive Server 15.0 EC, 15.0 ESD #2, 15.0.1 or later you must specify the version parameter as '15.0.1'.
-

`sp_downgrade 'prepare'` validates the readiness of Adaptive Server 15.0.2 for downgrade. This means that `sp_downgrade 'prepare'` may print messages regarding manual changes required before the downgrade can be completed. Errors reported must be corrected and implications of warnings reported must be understood before you can proceed with further downgrade steps. Refer to Table 6-1 to get more information on correcting the issues reported by `sp_downgrade 'prepare'`.

- 4 Repeat step 3 until `sp_downgrade 'prepare'` reports no errors. Your system is now ready to downgrade.
- 5 Begin the downgrade process by executing:

```
sp_downgrade 'downgrade', <'version'>
```

The version number should be the same as the version you entered in step 3. Once the above command completes successfully there can be no more activity on the 15.0.2 server. You must immediately shutdown the 15.0.2 server.

- 6 Copy the `RUN_SERVER` file to the 15.0 or 15.0.1 release area you are downgrading to. Modify the `RUN_SERVER` file to use the `dataserver` binary from the 15.0 or 15.0.1 release.

- 7 Restart the old server using the modified *RUN_SERVER* file.

Note If a downgraded server is booted using the 15.0.2 configuration file, the new options will cause the message ‘Unknown parameter ‘. The unknown option will be reported only the first time the server is booted.

- 8 Run any scripts that may be applicable to your site listed in the post-downgrade tasks.

These are the tasks needed for downgrade to Adaptive Server 15.0 or 15.0.1 if you have not used any 15.0.2 functionality. If you have used 15.0.2 functionality then follow the additional steps are described in the rest of the section.

Post-downgrade tasks

After you perform the basic downgrade steps, some changes from the 15.0.2 upgrade still persist.

System stored procedures

After you run *installmaster*, *installcommit*, *installsecurity*, *installhasvss*, and *installmsgsvss* on your downgraded server, any changed system stored procedures are returned to their earlier form. Any new stored procedures introduced to Adaptive Server in releases subsequent to 15.0 or 15.0.1 are not removed. Attempts to execute such stored procedures against an older binary will have unpredictable results.

System messages

During upgrade from a previous 15.0 or 15.0.1 version to Adaptive Server 15.0.2 you should have created an *restore_msgs.ebf* file. Now you must run this script to restore your messages back to the version from which you upgraded.


```
isql -S -Usa -P <restore_msgs.ebf>
```

Note If you upgraded from a pre-15.0 or 15.0.1 release and are reverting to a pre-15.x release then you must only run *instmsgs.ebf* file.

Unrecognized tokens or objects

The on-disk structures of stored procedures, triggers, and views may contain statement identity tokens, datatypes and object references that are not understood by the earlier Adaptive Server versions. You must drop all compiled objects that use features introduced into Adaptive Server subsequent to the release to which you are downgrading.

System tables

After returning to a pre-15.0.2 binary:

- If during the upgrade process you ran the 'update all statistics' on syslogins then you must delete statistics for syslogins and recreate them.
- `spt_values` is dropped and re-created when you run the 15.0 or 15.0.1 `installmaster`, so any new types are eliminated from this table.

Configuration parameters

New configuration parameters have been added in 15.0.2. If a downgraded server is booted using the 15.0.2 configuration file, the new options will cause the message 'Unknown parameter' the unknown option will be reported the first time the server is booted. These messages can be ignored because the configuration file will be rewritten without the unknown options.

Running the 15.0 or 15.0.1 version of `installmaster` removes configuration parameters that belong to 15.0.2, by deleting `sysconfigures` rows for configuration parameters that do not exist in `syscurconfigs`. After running `installmaster` the error messages will no longer appear when you start the server.

Keywords

Keywords `decrypt_default`, `xmltable`, `path`, were added in Adaptive Server 15.0.2, making it impossible to create identifiers using these names. You must change applications if you used these names.

See “Additional steps required if you used 15.0.2 features before downgrading” on page 96 for more information.

Additional steps required if you used 15.0.2 features before downgrading

If you are rolling back after having used any of the 15.0.2 features, additional steps can be necessary. You should take steps before you downgrade to 15.0 or 15.0.1, and others should be performed immediately after downgrading. You need to read the documentation on the features below that may require manual changes.

Table 6-1: List of topics and references for Unix

To correct errors related to	Refer to
Encrypted Columns	Chapter 2 “Encrypted Columns” in Adaptive Server 15.0.2 <i>New Feature Guide</i>
User Passwords	Chapter 10 “Supported Security Features” in Adaptive Server 15.0.2 <i>New Feature Guide</i>
Instead of Triggers	Chapter 22, “Instead of Triggers” in Adaptive Server 15.0.2 <i>New Feature Guide</i>
SQL user defined functions	Chapter 17 “User Defined SQL Functions” in Adaptive Server 15.0.2 <i>New Feature Guide</i>
Archived Databases	Chapter 4 “Archived Database Access” in Adaptive Server 15.0.2 <i>New Feature Guide</i>
Hash built-ins	Chapter 19 “Changes to System Procedures, Functions and Commands” in Adaptive Server 15.0.2 <i>New Feature Guide</i>

In general, no additional steps are required when you are returning to an Adaptive Server version in which the feature was already available.

Real-Time Messaging

Drop all stored procedures, views, and triggers that use the messaging built-ins for the Real-Time Messaging feature.

For more information about Real Time Messaging, see the *Real-time Data Services Messaging User's Guide*.

New sort orders

If a new nocase sortorder for Chinese or Japanese character sets is configured as Adaptive server's default sortorder. Before downgrading to a 15.0 or 15.0.1 release, switch to a sortorder that is compatible to that release. Switching a sortorder means all user indexes need to be reindexed. Please refer to *System Administration Guide Chapter 9 "Configuring Character Sets, Sort Orders, and Languages"* for details on how to change server's default sortorder. If `sp_downgrade` is called when the new nocase sortorder is in use, the error message "Cannot downgrade to '%1!' server, which does not support server's current default sortorder" will be raised and the downgrade process will be aborted.

Native XML

The XML Service feature of Adaptive Server includes the new `xmltable` function. If you create views or stored procedures using the `xmltable` function in 15.0.2, then return to 15.0 or 15.0.1, you will get an error.

Instead of trigger

Instead of Triggers are objects stored in the system catalogs. These objects should be removed before downgrading. When the 15.0 or 15.0.1 server is booted any Instead of Triggers that were not removed will remain in the system catalogs but will not execute.

SQL user defined functions

SQL user defined functions are also objects stored in the system catalogs. If you do not remove them before downgrading, they remain in the catalogs after downgrade. Attempts to drop or execute a SQL user defined function from a 15.0 or 15.0.1 version will result in misleading error messages.

Automatic database expansion

If you installed the 15.0.2 automatic database expansion procedures using `installdbextend`, and then applied the threshold procedure to one or more database segments, the thresholds might not work properly when applied to the log segment after a downgrade to 15.0 or 15.0.1.

To clear all auto-expansion thresholds that might exist on one or more segments before downgrading, use the command:

```
sp_dbextend 'clear', 'threshold'
```

Alternatively, before downgrading, you can disable the entire automatic expansion feature server-wide without changing any existing rules or clearing any thresholds. Execute the following commands using `sa_role`:

```
use master
go
sp_dbextend 'disable', 'database', 'server-wide'
go
```

This prevents threshold procedures from doing any work even if they were fired at runtime.

Sybase recommends that you leave all the policies and thresholds in place, and simply disable the entire feature server-wide before the downgrade. This simplifies re-enabling automatic expansion if you return to 15.0.2 later.

Replication issues with downgrade

When downgrading a server that has replication enabled on databases that contain encrypted data, you must do one of the following before you start the downgrade procedure:

- 1 Ensure that all replicated data in the primary database transaction log has been successfully transferred to the standby or replicate database. The process for doing this is application dependent.
- 2 Using the following commands truncate the transaction log in the primary database, and zero the RS locator for that database in the Replication Server. In the primary database run:

```
sp_stop_rep_agent primary_dbname
dbcc settrunc ('ltm', 'ignore')
dump tran primary_dbname with truncate_only
dbcc settrunc ('ltm', 'valid')
```

Shutdown Replication Server. In the RSSD for the Replication Server run:

```
rs_zeroltm primary_servername, primary_dbname
```

Downgrading Job Scheduler

If you are downgrading to Adaptive Server versions 15.0 or 15.0.1, you must run `installjsdb` script from the lower version anytime you downgrade Job Scheduler.

v Downgrading Job Scheduler

- 1 Disable Job Scheduler before you downgrade Adaptive Server:

```
sp_configure "enable job scheduler", 0
```

- 2 Downgrade Adaptive Server. See the installation guide for details.

- 3 Run the `installjsdb` script from the version you are downgrading to:

```
isql -Usa -Psa_password -Sservername -i  
$SYBASE_ASE/scripts/installjsdb
```

Note The directory with the location of the `isql` executable (`$SYBASE_OCS/bin`) must be in your path.

- 4 Enable Job Scheduler:

```
sp_configure "enable job scheduler", 1
```

- 5 To start Job Scheduler, either restart the server, or enter:

```
use sybmgmtdb  
go  
sp_js_wakeup "start_js",1  
go
```

Returning to Adaptive Server version 15.0.2

In the future, if you decide to return to Adaptive Server version 15.0.2, follow the instructions in Chapter 5, “Upgrading Adaptive Server.”

Troubleshooting SySAM Issues

Note If SySAM 2.0 related errors are not addressed in a timely fashion, Adaptive Server may cease to function.

For the most up to date information about troubleshooting SySAM, see the SySAM FAQ at <http://www.sybase.com/sysam>. For the most up-to-date troubleshooting information in this installation guide, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

This section includes information about troubleshooting the most common SySAM 2.0-related errors.

Topic	Page
Calling Sybase technical support for SySAM-related issues	101
Where to look for SySAM-related errors	102
Troubleshooting SySAM	103
SySAM grace mode	108
SySAM grace period	109
Enabling and changing e-mail notifications	109

Calling Sybase technical support for SySAM-related issues

If you call Sybase technical support for SySAM-related issues, be sure to have the following information available:

- The Adaptive Server errorlog located in the *install* directory.

On UNIX:

```
$SYBASE/ASE-15_0/install/<servername>.log
```

On Windows:

```
%SYBASE%\ASE-15_0\install\<servername>.log
```

- If Adaptive Server has started in grace mode – the output from the `sp_lmconfig` command.
- If Adaptive Server does not start – the server's properties file

On UNIX:

```
$(SYBASE/ASE-15_0/sysam/<servername>.properties
```

On Windows:

```
%SYBASE%\ASE-15_0\sysam\<servername>.properties
```

servername must correspond to the name of the Adaptive Server that does not start.

- The license files saved in the `SYSAM-2_0/licenses` directory on the machine running Adaptive Server.
- The value of `SYBASE_LICENSE_FILE` and `LM_LICENSE_FILE` environment variables, if they have been defined.
- Output from `lmutil lmpath -status` command. The `lmutil` binary is located in the `bin` directory under `SYSAM-2_0` folder.
- If you are using a served license model:
 - The licenses saved in the `$(SYBASE/SYSAM-2_0/licenses)` directory of the license server.
 - The license server log file in the `$(SYBASE/SYSAM-2_0/log)` directory.

Where to look for SySAM-related errors

All SySAM related errors are printed in the Adaptive Server errorlog. All SySAM-related messages are prefixed with "kernel SySAM:" in the Adaptive Server errorlog. For example:

```
<timestamp> kernel SySAM: Failed to obtain 1 license(s)
for ASE_CORE feature from license file(s) or server(s).
```

If e-mail alerts are configured, recipients receive an e-mail whenever a SySAM event occurs.

You can always check the current state of SySAM licenses in Adaptive Server using the `sp_lmconfig` command. The output shows the current SySAM configuration and the status of each license.

If you are using a served license model, the license server's debug log located in the *log* directory provides information about any license check-out issues.

Troubleshooting SySAM

Table 7-1 lists methods for troubleshooting SySAM.

Table 7-1: Troubleshooting SySAM

Description	What to do
Cannot find license type during the Adaptive Server installation	<p>The Adaptive Server installation program displays only the most commonly used license types. If you are using a legacy license or a new license type, your license may not be listed by the installation program.</p> <p>If you do not find the license type you need, select “Unknown” or “Other” from the drop-down list and complete the installation. After the installation is complete, use <code>sp_lmconfig</code> to change the license type. See the <i>Reference Manual:Procedures</i> for <code>sp_lmconfig</code> usage information.</p> <hr/> <p>Note If you set the license type value to “Unknown” or “Other,” Adaptive Server uses the first available license. If you have licenses for more than one type, Adaptive Server may pick the incorrect license. Set the correct license type when you have multiple license types on the license server.</p>

Description	What to do
<p>The installation program produces a warning about not finding a suitable license</p>	<p>If you select a served license model, the install program checks for a license with the selected edition and license type. A warning message is displayed if no suitable license is available.</p> <p>If you are performing a new Adaptive Server installation or upgrading an existing Adaptive Server version 12.5.x or earlier, proceed with the installation. Adaptive Server provides a 30 day grace period to resolve licensing issues. Once the installation is complete, debug the license check-out failure. For more information, see “Adaptive Server cannot check out a license and starts with graced license instead” in this table.</p> <p>If you are upgrading an existing Adaptive Server version 15.0 or later, Adaptive Server may not start after the upgrade if you do not have a suitable license. Exit the installation program and see if the existing Adaptive Server is able to check-out licenses. If it is not, troubleshoot this issue and then proceed with the upgrade.</p> <p>If the existing Adaptive Server is able to check-out licenses, verify the license you are using authorizes use of the update you are applying. This can be determined using the date based version of your license file and the date this update was released. You must have a license with a date based version that is greater than or equal to the date the update was released. Adaptive Server provides a grace period for updating the licenses after support is renewed.</p> <p>The date based version of the license can be determined from the Version column of the output from <code>sp_lmconfig</code>, by looking at the license checkout message in the Adaptive Server errorlog, or by examining the license itself. For more information, see Appendix A, “Anatomy of a License,” of the <i>Sybase Software Asset Management User’s Guide</i>.</p> <p>The release date of the Adaptive Server can be determined by the date displayed in the version string, cover letter, or SPDC download link. This data is also shown in the error message posted.</p> <p>You can download an updated license from SPDC, if your support was paid for the date the update was released. If you support was not paid, you are not authorized to install an update.</p> <p>See Appendix B of <i>Sybase Software Asset Management User’s Guide</i> for explanation of how SySAM uses date-based versioning. See “Updating your Sybase licenses” in Chapter 3, “Getting and Using your Licenses,” in the <i>Sybase Software Asset Management User’s Guide</i> for how to update your license.</p>
<p>The license server executables and scripts are not installed</p>	<p>The <i>sysam</i> wrapper script and the executables for the license server (<code>lmgrd</code> and <code>SYBASE</code>) are not installed or found in the <code>SYSAM-2_0 bin</code> directory.</p> <p>The license server component is not selected for installation by default. To install the license server, choose the Custom install option and select license server component.</p> <p>Alternatively, you can download a standalone SySAM license server install program from SPDC.</p>

Description	What to do
The license server does not start	<p>The most common cause for the license server failing to start is that there are no served licenses in the server. In order for license server to start, there needs to be at least one served license in the <i>licenses</i> directory. When you install a new license server, there are no served licenses in the licenses directory and the license server does not start until you activate a license from SPDC and copy it into the licenses directory.</p> <p>Other causes for license server not to start include:</p> <ul style="list-style-type: none"> • The license activated from SPDC was an unserved license. Examine the license files. Served licenses always start with a <code>SERVER</code> header. If you do not see a line starting with <code>SERVER</code> at the top of the license, you are using an unserved license which does not work with the license server. • If you are using a specific port number for the license, that port number may already be in use. Use <code>netstat -a</code> to verify that the port number is free. If not, free the port or use a different port number for the license server. • The host name recorded in the <code>SERVER</code> header does not match the actual host name. The host name is recorded next to the <code>SERVER</code> keyword. If it does not match, correct the host name in the license file or set it to <code>"this_host"</code>, which is a keyword that works with any host name. • The licenses copied may be activated for a different machine. Examine the Host ID recorded in the line. It is recorded next to the host name in the <code>SERVER</code> header. It must match the Host ID for the machine where license server will be running.
The license server does not recognize the license files	<p>If the license server refuses to serve the licenses, this may be because:</p> <ul style="list-style-type: none"> • The license was activated for a different machine or activated using incorrect host ID. Examine the host ID recorded in the license with the host ID for the machine. If they are mismatched, check-in and regenerate the license with the correct host ID. • The license has been altered. Nothing in the license except some designated information can be altered. The license is rejected if you change other information. If the license is being created by typing from a paper copy, verify if there are any errors entering the license. You can also download a new copy of the activated license from SPDC. <p>If the Host ID for your platform is based on the Network adapter identification, verify that you have used an ID associated with a valid NIC. IDs associated with loopback adapters are not valid. If the ID associated with a removable network adapter was used, ensure that the adapter is attached to the computer.</p>
Adaptive Server does not start, and generates a license check-out failure error	<p>If Adaptive Server cannot check out the required license, it determines whether the license can be issued in grace mode. If a grace period cannot be given, the license is not granted. If the base license for Adaptive Server (<code>ASE_CORE</code>) was in grace and the issue is not fixed by the end of the grace period, Adaptive Server fails to start.</p> <p>For more information, see “Adaptive Server cannot check out a license and starts with graced license instead” in this table. The troubleshooting tips for resolving licenses issued under grace apply to solve this issue as well.</p>

Description	What to do
<p>Adaptive Server cannot check out a license and starts with graced license instead</p>	<p>Execute the <code>sysam diag <i>feature_name</i></code> command from the <code>SYSAM-2_0 bin</code> directory, where <i>feature_name</i> is the SySAM feature name for Adaptive Server or the optional feature that is in grace mode. The feature name is printed in the Adaptive Server errorlog and in the optional e-mail notifications.</p> <p>If the <code>diag</code> command shows no licenses available for check-out, this may be due to one of the following reasons:</p> <ul style="list-style-type: none"> • If you are using a served model: <ul style="list-style-type: none"> • Verify that the license server is running and reachable from the machine you are using. Use <code>"sysam status"</code> to verify that you can access the license server. If you cannot access the license server, see if the license server is running. • If the license server is running, use <code>"sysam status -f <i>feature_name</i>"</code> to determine whether the license server is serving a license for the given feature. If not, obtain the correct licenses from SPDC. • All licenses on the license server may be in use. If <code>"sysam status -f <i>feature_name</i>"</code> indicates that no licenses are available, obtain additional licenses or shutdown existing instances of Adaptive Server. • If you are using an unserved model: <ul style="list-style-type: none"> • The license for the given feature may not be copied in the local <i>licenses</i> directory. Obtain the correct licenses from SPDC and copy them to the local <i>licenses</i> directory. • The license may have been activated for a different machine or activated with incorrect host ID. Check-in and reactivate the license with the correct host ID. • The available licenses are for a different operating system or architecture. Obtain license for the correct platform from SPDC. • If you are running in a terminal server environment, unserved licenses cannot be used. Set up served licenses for a terminal server environment. • The available license may not authorize use of this version of Adaptive Server. For information about date based versioning, see Appendix B, “Installing Product Updates, EBFs, and Support Renewal” in the <i>Sybase Software Asset Management User’s Guide</i>. For information about how to update your license, see the “Updating your Sybase licenses” section in Chapter 3, “Getting and Using your License,” in the <i>Sybase Software Asset Management User’s Guide</i>.

Description	What to do
Adaptive Server cannot checkout a license and starts with graced license instead (continued)	<p>If the <code>diag</code> command shows the license for the given feature is available for check-out, the reason that Adaptive Server cannot check-out the license could be:</p> <ul style="list-style-type: none"> • The product edition or license type does not match the requirement. If the edition and license type configuration parameters are set, Adaptive Server uses only the license that matches these settings. <ol style="list-style-type: none"> 1 Run <code>sp_lmconfig</code> to see the configured values. 2 Check the <code>VENDOR_STRING</code> attribute of the licenses to make sure a matching license exists. <p>If a matching license does not exist, obtain a suitable license from SPDC or adjust the Adaptive Server configuration.</p> <p>If Adaptive Server cannot start because the grace period expired, examine the <code>servername.properties</code> file located in the <code>ASE-15_0/sysam</code> directory under your <code>SYBASE</code> directory for values. These values are stored in two lines that start with <code>PE=</code> and <code>LT=</code>. These lines are present only if the configuration is set.</p> • Also see “Adaptive Server cannot find license for optional feature, even though the license exists” in this table for more information.
Adaptive Server shows graced licenses even after the issue has been fixed	<p>Adaptive Server performs license checks periodically, and the license status is not immediately updated. It is updated only after the next heartbeat cycle is completed. This may take a few hours.</p>
Adaptive Server cannot find license for an optional feature, even though the license exists	<p>Execute <code>"sysam diag feature_name"</code> to ensure that license for the optional feature exists and can be checked out from the machine where Adaptive Server is running. If the feature exists, but cannot be checked out from Adaptive Server, this could be because:</p> <ul style="list-style-type: none"> • The optional feature is for a different edition. • The Active or Standby flags do not match for base Adaptive Server and the optional feature. <p>Execute <code>sp_lmconfig</code> to determine the edition and active and standby configuration value. The edition and active and standby values are displayed in the Property Name / Property Value table. The edition value is displayed in the <code>PE</code> row and active and standby value is displayed in the <code>AS</code> row.</p> <p>Compare the values of <code>PE</code> and <code>AS</code> for the Adaptive Server with the same values recorded in the <code>VENDOR_STRING</code> and <code>ISSUER</code> fields of the optional feature license. They must match. Some features like <code>ASE_HA</code> are not offered with standby pricing, so there is no <code>"AS=A S"</code> flag in the license for this feature. This field is not considered when searching for a license without active and standby captability.</p>

Description	What to do
Adaptive Server does not start with the expected edition or license type	<p>If the edition and license type configuration parameters are not set, Adaptive Server uses the first available ASE_CORE license. If you have multiple ASE_CORE licenses with different editions and license types, which license is the first available depends on many factors such as license file directory sort order, available licenses on the license server, and so on.</p> <p>Because what edition and license type is granted cannot be guaranteed, Sybase recommends that you set the edition and license type configuration parameters. When these parameters are set, you are assured that Adaptive Server starts with that configuration. If a matching license were to be unavailable, Adaptive server starts under grace (if available) to allow you to resolve the licensing issue.</p>
Adaptive Server does not pick the correct license	<p>When Adaptive Server looks for a license, it does so in the following locations:</p> <ul style="list-style-type: none"> • Any saved values in the <i>registry</i> on Windows platforms, or the <i>.flexlmrc</i> file on UNIX platforms. Use "lmpath -status" to view the saved values. • The values set for SYBASE_LICENSE_FILE and LM_LICENSE_FILE environment variables. • All license files (any file with a <i>.lic</i> extension) in the licenses directory under the <i>SYSAM-2_0</i> directory. <p>All of the above locations are searched in the order specified until a suitable license is found. When a directory is specified, license files in that directory are loaded in the directory sort order.</p> <p>The first license that matches the feature name, version, edition, and license type filter is used. This license may not be what you were expecting to use. Adjust the above locations to control this.</p>

SySAM grace mode

If Adaptive Server cannot find a required license, it continues to operate in grace mode, if a grace period is available. Determine the license status of Adaptive Server using the output from `sp_lmconfig`. The `Status` column in the output shows `graced` for any license that is operating in grace.

The Adaptive Server errorlog includes information about licenses awarded under grace. If e-mail notifications are set up for warning severity, e-mail messages are sent for this event. See “Enabling and changing e-mail notifications” on page 109 for more information.

The warnings in the Adaptive Server errorlog and e-mail message are repeated while the license is under grace.

SySAM grace period

Typically, you get a 30-day grace period to resolve any SySAM-related issues. When a license is awarded under grace, the Adaptive Server errorlog entry and optional e-mail message includes the date and time when the grace period expires. In addition, you can execute `sp_lmconfig` to determine the expiration date. The Expiry Date column next to any row that shows `graced` status indicates the date and time the grace period expires. Adaptive Server disables the functionality corresponding to the license if the issue is not fixed within that time. Adaptive Server itself shuts down if the `graced` license is `ASE_CORE`.

Enabling and changing e-mail notifications

You can configure the e-mail notification setup using `sp_lmconfig`. Run `sp_lmconfig` command without any parameters to view the current status and configuration.

To enable e-mail notifications, set the `smtp host`, `smtp port`, `email sender`, `email recipients`, and `email severity` properties.

To change the severity that triggers e-mails, set the `email severity` property. Values are: "ERROR", "WARNING", "INFORMATIONAL" and "NONE".

To change the list of e-mail recipients, set the `email recipients` property to the comma separated list of e-mail addresses. E-mail addresses can include distribution lists, e-mail address for a pager device, and anything that is capable of receiving SMTP messages.

This chapter provides instructions for troubleshooting installation error messages.

If this chapter does not describe the error message you are encountering, see the *Error Messages and Troubleshooting Guide*.

Topic	Page
Overview	111
Error log locations	112
Solutions to common installation problems	113
Upgrading compiled objects with dbcc upgrade_object	121

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 112.

Table 8-1 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 `srvbuild` unexpectedly quits, or if you cannot correct the problem, see the *Error Messages and Troubleshooting Guide*.

Table 8-1: Troubleshooting guidelines for UNIX platforms

Problem	Possible cause and solution
The installation program cannot start Adaptive Server.	<p>Failure to start 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> <hr/> <p>Note Adaptive Server may not start if it cannot find a SySAM license and if the grace period has expired. Check the Adaptive Server error log for the cause of the license checkout failure and fix the issues.</p>
After upgrading Adaptive Server, you cannot use <code>srvbuild</code> .	After you begin upgrading a server, you may be unable to use the same <code>srvbuild</code> session for other tasks. Exit and restart <code>srvbuild</code> .
The installation program cannot connect to the upgraded Adaptive Server.	After you begin upgrading a server, you may be unable to use the same <code>srvbuild</code> session for other tasks. Exit and restart <code>srvbuild</code> .
The installation program detects reserved word conflicts.	See “Stopping Adaptive Server after a failure” on page 116.
The upgrade fails.	See “Recovering from a failed upgrade” on page 118.

Error log locations

The information in the error logs helps you determine the reason and possible solution for an error message.

Logs for installation utilities

Table 8-2 lists the default error log locations for the installation, upgrade, and setup utilities.

Table 8-2: Error log locations for installation utilities

Utility	Error log location
Installer	<code>\$SYBASE/log.txt</code>

Utility	Error log location
srvbuild	<code>\$\$SYBASE/\$SYBASE_ASE/init/logs/srvbuildMMDD.NNN</code>
srvbuildres	where: <ul style="list-style-type: none"> • <i>MM</i> is the month • <i>DD</i> is the date • <i>NNN</i> is a three-digit number identifying the srvbuild session
sqlloc	<code>\$\$SYBASE/\$SYBASE_ASE/init/logs/sqllocMMDD.NNN</code>
sqllocres	
sqlupgrade	• <code>\$\$SYBASE/\$SYBASE_ASE/init/logs/sqlupgradeMMDD.NNN</code>
sqlupgraderes	• <code>\$\$SYBASE/\$SYBASE_ASE/upgrade/upgrade.NNN</code> (a temporary file created by the upgrade process)

Logs for Sybase servers

Table 8-3 lists the default error log locations for each Sybase server.

Table 8-3: Error log locations for Sybase servers for UNIX

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	Default XP and Monitor server logs are <code>\$\$SYBASE/\$SYBASE_ASE/install/<servername>.log</code>
XP Server	<code>\$\$SYBASE/\$SYBASE_ASE/install/<servername>.log</code>

Solutions to common installation problems

If this section does not describe the problem you are experiencing, see the *Error Messages and Troubleshooting Guide*.

Cannot use X-Windows

If the setup and configuration 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/ASE-15_0
% 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 now use approximately 25 percent less window 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.

DISPLAY environment variable not set correctly

This error message means that the DISPLAY environment variable on the *remote* machine is not set correctly to display the Installer to your *local* machine:

```
The DISPLAY environment variable is not set correctly.
```

To correct the problem, enter the following command at the UNIX prompt of the *remote* machine, where *host_name* is the name of the machine on which you want the Installer to appear (that is, on your *local* machine):

For C shell:

```
setenv DISPLAY host_name:0.0
```

For Bourne shell:

```
DISPLAY=host_name:0.0; export DISPLAY
```

Client not authorized to connect to server

If you run the Installer and you get this error message, it means the remote machine does not have permission to display the user interface on the local machine where you start working:

```
Xlib: connection to "host_name" refused by server
Xlib: Client is not authorized to connect to Server
```

```
xhost: unable to open display "host_name"
```

To correct the problem:

- 1 Enter the following command at the UNIX prompt of your *local* machine, where *remote_machine* is the machine on which you are running the Installer:

```
xhost +remote_machine
```

- 2 Restart the Installer.

Address already in use

The following message from `srvbuild` means you entered a port number that is already in use:

```
kernel: ninit: bind, Address already in use
```

To correct the problem, enter a different port number on the `srvbuild` window. The command `netstat -a` produces a list of port numbers in use.

Adaptive Server fails to start

The shared memory of the operating system may not be set high enough.

To correct the problem, see Chapter 2, “Installing Adaptive Server” for instructions on adjusting the shared memory value. Restart the installation or upgrade process.

The Installer fails to start

If the Installer does not start, restart it, using:

```
"-is:javaconsole"
```

This displays any error messages to the console which user can then correct.

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 the Adaptive Server syssservers table. If you created XP Server in a different srvcbuild session than Adaptive Server and you did not specify a related Adaptive Server, srvcbuild cannot update the syssservers table. You also need to check verify that the XP server exists in the interfaces file or LDAP server.

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 8-3 on page 113.

If installation fails after files are created

If the installation program quits while you are configuring Adaptive Server:

- 1 View the contents of the log file generated by Adaptive Server. For default error log locations, see Table 8-3 on page 113.
- 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 starts the Adaptive Server that you are attempting to install, shut down that server.

Troubleshooting resource file installations

If you encounter problems during the build, configuration, or upgrade process 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 `SYBSHUTWAIT` environment variable to force the utilities to wait for Adaptive Server to shut down. For example, this command forces the utility to wait for two minutes to allow Adaptive Server to shut down before proceeding with the next task:

```
% setenv SYBSHUTWAIT 120
```

If Adaptive Server fails the pre-upgrade eligibility test

If Adaptive 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 failed the pre-upgrade eligibility test.

After you resolve any problems, shut down Adaptive Server and use `sqlupgrade` to complete the upgrade session.

Recovering from a failed upgrade

If the upgrade process fails, 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 8-2 on page 112.

Note Once you start ASE on the new version you cannot start ASE on the previous version. Attempts to do so will likely result in having to restore from backups.

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.

Re-running the upgrade

Failures while upgrading a server installation fall into one of two categories: failure to upgrade an individual database, or failure to complete configuration changes after all databases have been upgraded.

- 1 If an individual database fails to upgrade, you can retry that upgrade manually. Begin by correcting the problem that caused the failure. Output from the upgrade utility should identify the problem. The most common cause of failure is running out of some resource: space (either data or log), locks, auxiliary scan descriptors. You can add space to the database using the `alter database` command. Other resource failures may be correctable by changing the server's configuration via the `sp_configure` stored procedure.

If an upgrade failure leaves a database offline, and the failure can only be corrected by making data changes in the database, you can gain access to the failed database by using `isql` or a similar program to connect to the affected server as user "sa" and issuing this command:


```
dbcc traceon(990)
```

Note This trace flag grants access only to user “sa”. It is not sufficient to use an account having “sa_role”. If you have disabled the “sa” login, you must re-enable it to get access using this method.

Having set this trace flag, user “sa” can now use the offline database and make the necessary changes to correct the upgrade failure.

To restart a failed upgrade, use the online database command:

```
online database <failed_db_name>
```

The server restarts that database’s upgrade from the point of failure.

- 2 If the failure occurs after all databases have been upgraded, or if a failure somehow causes the upgrade utility to crash, you can re-run the utility manually. First diagnose and correct the failure, then run the upgrade utility:

```
$SYBASE/$SYBASE_ASE/upgrade/upgrade
```

When restarted in this way, the upgrade process says it is “verifying” the upgrade rather than “starting” it, but it makes all the same checks as for the original upgrade.

To verify that a database has upgraded successfully, you can check any database’s upgrade status using the online database command. If any upgrade is required for a database, this command performs it. You may also use a procedure such as this to check all databases at a given installation:

```
declare @dbname varchar(255)
select @dbname = min(name)
from master..sysdatabases
while @dbname is not null
begin
online database @dbname
select @dbname = min(name)
from master..sysdatabases
where name > @dbname
```

end

Note There are certain upgrade failures from which the server cannot recover. For example, attempts to upgrade system tables to version 15.0 format are quite sensitive to failures at certain points during the required changes. If you encounter such a failure, restore the failed database from backup. To prevent the upgrade from failing again, correct the problem that caused the original failure **before** issuing the online database command for that database. These catastrophic failures are nearly always caused by running out of resources, as described above, which then causes a failure to undo the aborted transaction.

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 re-run 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 You may need to shut down Adaptive Server.

Shutting down the server enables the installation program to start 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 process 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
`$$YBASE/$$YBASE_ASE/install/<servername>.log`.

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, if 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 15.0. 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 before being upgraded.

Compiled objects for which the source text was hidden using sp_hidetext are upgraded in the same manner as objects for which the source text is not hidden. For information on sp_hidetext, see the *Reference Manual:Procedures*.

Note If you are upgrading from 32-bit installations to use a 64-bit Adaptive Server, 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 process calculates the exact size. 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” on page 122.

Finding compiled object errors before production

Changes made in earlier versions of Adaptive Server may cause compiled objects to work differently in version 12.5.x 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 125.

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 that object is not upgraded. 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 re-create the compiled object.

For example, suppose you load a database dump from Adaptive Server 11.5 into Adaptive Server 15.0 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 version with quoted identifiers active (set quoted identifiers on).
- Quoted identifiers are not active (set quoted identifiers off) in the current session.

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 version 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*) that 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, then execute dbcc upgrade_object again. You need not do this 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, consider the following stored procedures:

```
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 upgrades the view and overwrites the information about the original column information in syscolumns.

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. You must fix these objects 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

Example 1 Upgrades all compiled objects in the active database:

```
dbcc upgrade_object
```

Example 2 Upgrades all stored procedures in the listdb database. Single quotes are used around procedure because set quoted identifiers is on:

```
dbcc upgrade_object(listdb, 'procedure')
```

Example 3 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, "rule")
```

Example 4 Upgrades all stored procedures named list_proc in the listdb database.

```
dbcc upgrade_object(listdb, list_proc)
```

Example 5 Upgrades the stored procedure list_proc, which is owned by the login "jkarrik".

```
dbcc upgrade_object(listdb,  
"listdb.jkarrik.list_proc")
```

Example 6 Returns an error because the value of *dbname* is master and the value of *database* is listdb. These values must match.

```
dbcc upgrade_object(master,  
"listdb.jkarrik.list_proc")
```

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. 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 window, a System Administrator can execute `dbcc traceon(3604)`. Sybase recommends that you use this command if you think the output of error messages might overflow the error log.

Using database dumps in upgrades

Upgrading using dump and load

You can load pre-12.5 database dumps and transaction logs and upgrade the databases.

Some issues of which you should be aware:

- Upgrading requires space for copying data and logging changes to the system tables during the upgrade process. If the source database in the dump was nearly full, the upgrade process might fail due to insufficient space. While this is expected to be uncommon, you can use `alter database` to extend the free space in the event of insufficient-space errors.
- After reloading an older dump, run `sp_checkreswords` from the new installation on the loaded database to check for reserved words.

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_checksourc` 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_checksourc`, see the *Reference Manual: Procedures*.

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 the object has not been upgraded.

Removing Adaptive Server

Only actions performed by the Installer are removed during an uninstallation process. This means that files or registry entries created after installation are not removed by the uninstaller and must be removed by the customer after the uninstall is completed. Where possible these actions are identified in this chapter.

Heading	Page
Uninstalling Adaptive Server 15.0.2	129
Removing an existing Adaptive Server	130

Uninstalling Adaptive Server 15.0.2

- 1 Verify that the server to be removed is not running. If it is running, shut it down. To check if it is running use:

```
"ps"
```

Then if it is running use isql commands: kill or shutdownnto shutdown the server.

- 2 Shut down all servers prior to running the uninstall program. See "Starting and Stopping Servers" in the *Configuration Guide*.
- 3 The uninstall program is located in the following product specific directories:

- *uninstall/ASESuite*
- *uninstall/PCClient*

Launch the uninstall program by executing the following steps:

- a Change to the *SYBASE* directory.
- b Execute the uninstall program by entering:

```
$SYBASE/uninstall/ASESuite/uninstall
```

- 4 Once the Uninstall Welcome window displays, click Next to continue.
- 5 The component selection window displays with everything selected. Deselect components by clicking the check box and erasing the check mark. Click Next to continue.

Note If there is SySAM Network License Server installed on this machine, it may be used by other Sybase products. Uncheck this component unless you know no one is using this License Server.

- 6 A summary window displays what is about to be uninstalled. Click Next to continue.
- 7 A window displays indicating that the uninstallation process is being performed. There is no progress bar.

Note One or more pop-up displays may occur indicating that a file was modified after it was installed and requesting confirmation that it should be deleted.

- 8 If there are any user created files left over from the uninstall, the uninstaller will prompt user if it can remove them.
- 9 When the uninstall process has completed a final window displays. Click Finish to exit the uninstall program.

Removing an existing Adaptive Server

- 1 To remove an existing Adaptive Server, from `$SYBASE`, enter:

```
rm servername.*
```
- 2 Change to the directory `$SYBASE/$SYBASE_ASE/install` and run the following commands:

```
rm RUN_servername.*  
rm servername.*
```
- 3 Edit `$SYBASE/interfaces`, to remove all references to the Adaptive Server.
- 4 If you used operating system files for database devices, remove those.

Alternative Installation Methods

Topic	Page
Installation on consoles (non-GUI mode)	131
Silent installation	132
Configuration using resource files	134

Installation on consoles (non-GUI mode)

In cases where there is no graphics display device available, or for some other reason you wish to run the installer without the graphical user interface (GUI), the Installer can be launched in console or non-GUI mode. When invoking the installer setup program add the `-console` command line argument. In cases where the Installer launches automatically, select Cancel to cancel the GUI install and then launch the setup program from a terminal or console. Then:

- 1 Change the default directory to the CD.
- 2 Run the setup program with the `-console` command line argument:

```
setup -is:javaconsole -console
```

The flow of the installation is identical to a regular GUI installation, except that the display is written to a terminal window and responses are entered using the keyboard.

Silent installation

A silent installation (sometimes referred to as an unattended install) is done by running the Installer and providing a response file that contains answers to all of the Installer's questions. There are two methods of generating a response file for the Installer. The first method is to run the Installer in GUI mode and record all of the responses into a response file. This response file then be edited with an editor to customize the responses if needed.

The second method is to run the Installer and have it write out a template file. The template file contains all of the questions, but no responses. It can then edited to provide responses for each of the questions. The Installer is then invoked with the response file and all questions are answered by reading from the response file.

To perform the silent installation using the resource file that you created by running the Installer through the installation process:

- 1 Change default directory to the CD.
- 2 Run the setup program with the -options -record *filename*:

```
./setup -silent -options '/tmp/neals_install.txt' -W 'is:javaconsole'  
SybaseLicense.agreeToLicense=true
```

- 3 Run through the installation saving all of the responses.
- 4 Edit the response file to customize it for your needs.
- 5 Run the setup program with the -options filename and -silent command line arguments. Additionally you must include the -W argument that indicates that you accept the license agreement.

```
./setup -options "/tmp/neals_install.txt" -W  
"is:javaconsole"  
SybaseLicense.agreeToLicense=true -silent
```

To generate a template response file, skip step 3 and alter step 2 above to do the following:

```
setup -is:javaconsole -options-template <response file in absolute path>
```

In either case it is possible to custom configure the servers as well as being able to select a typical install and use the default configuration values.

Except for the absence of the GUI screens all actions of the Installer are the same and the result of an installation in silent mode is exactly the same as one done in GUI mode with the same responses.

Note The value for `-options`, `-options-template`, and `-options-record` must specify an absolute path.

Installation in silent mode can fail when a response is required from the user. Such a response is required when files already exist on disk and appear to be newer than the files the installer is attempting to install, or when files have been modified by the user.

Set the following options at the command line level to specify actions to take when installing in silent mode:

- `replaceExistingResponse` – stores your response to whether you want to replace a file that currently exists on your system with the one being installed.
- `replaceNewerResponse` – stores your response to whether you want to replace a file that currently exists on your system with the one being installed if the existing file is newer than the file being installed.
- `removeExistingResponse` – stores your response to whether you want to remove a file that currently exists on your system.
- `removeModifiedResponse` – stores your response to whether you want to remove a file that has been modified since installation.

Any of these options can be specified on the command line or within the options file as follows:

```
-G <option>="<value>"
```

Valid values are:

- `yesToAll`
- `yes`
- `noToAll`
- `no`

For example:

```
./setup -is:javaconsole -silent -options <responseFileName>  
-W SybaseLicense.agreeToLicense=true  
-G replaceNewerResponse="yesToAll"
```

Configuration using resource files

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 `srvbuildres` utility. The `srvbuildres` 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 2, “Installing Adaptive Server.”

To install an Adaptive Server or Backup Server using resource files:

- 1 Edit a resource file as described in “Configuring from a resource file” on page 134.
- 2 Execute the `srvbuildres` utility, using the edited resource file as described in “Using `srvbuildres`” on page 137.

Configuring from a resource file

You can install Adaptive Server via a resource file. Use the sample resource file provided with the software.

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 that:

- 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.

- 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 minimum master device size must be:
 - 2K page size – 30MB
 - 4K page size – 60MB
 - 8K page size – 120MB
 - 16K page size – 240MB
- The system procedure device and sybssystemprocs database sizes must be at least 120MB.
- The auditing entries in the resource file are not supported and will not be processed by the utility. To enable auditing, see the configuration guide for your platform.

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 Server or Backup Server 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! Make sure to edit the resource file generated by *srvbuild* to change attributes, such as device names, that may differ 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*, and 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]
sybinit.release_directory	The value of \$SYBASE at your site
sybinit.product	sqlsrv
sqlsrv.server_name	<i>server_name</i>
sqlsrv.new_config	yes (required value)
sqlsrv.do_add_server (to <i>interfaces</i> file)	yes [no]
sqlsrv.network_protocol_list	tcp [spx]
sqlsrv.network_hostname_list	<i>hostname</i>
sqlsrv.addl_cmdline_parameterst	<i>null</i>
sqlsrv.network_port_list	<i>port_number</i>
sqlsrv.master_device_physical_name	<i>path_and_name_of_master_device</i>
sqlsrv.master_device_size	30 (MB)
sqlsrv.master_database_size	13 (MB)
sqlsrv.errorlog	<i>\$SYBASE/\$SYBASE_ASE/install/server_name.log</i>
sqlsrv.do_upgrade	no (required value)
sqlsrv.sybssystemprocs_device_physical_name	<i>path_and_name_of_sybssystemprocs_device</i>
sqlsrv.sybssystemprocs_device_size	120 (MB)
sqlsrv.force_buildmaster	No
sqlsrv.server_page_size	2K (required)
sqlsrv.sybssystemprocs_database_size	132 (MB)
sqlsrv.sybsystemdb_device_physical_name	<i>path_and_name_of_sybsystemdb_device</i>
sqlsrv.sybsystemdb_device_size	3 (MB)
sqlsrv.sybsystemdb_database_size	3 (MB)
sqlsrv.default_backup_server	<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*, and 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 Adaptive 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_BS</i>
bsrv.do_add_backup_server (to <i>interfaces</i> file)	yes [no]
bsrv.network_protocol_list	tcp [spx]
bsrv.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	roman8 [iso_1, cp850, cp437, deckanji, ascii_8, eucgb, eucjis, mac, roman8, sjis, utf8]
bsrv.tape_config_file	\$SYBASE/\$SYBASE_ASE/backup_tape.cfg
bsrv.errorlog	\$SYBASE/\$SYBASE_ASE/install/server_name_back.log

Using *srvbuildres*

Note When building a new sever with *srvbuild* or *srvbuildres*, configuration parameters can be set in the configuration file in \$SYBASE. This configuration file is used to start the new server.

Source *SYBASE.csh* before running *srvbuildres*.

To execute *srvbuildres* at the UNIX prompt enter the following, where *resource_file* specifies the resource file containing the attributes that describe the server to build:

```
$SYBASE/$SYBASE_ASE/bin/srvbuildres -r resource_file
```

Note Unlike *srvbuild*, the *srvbuildres* utility does not require any X libraries.

When you have completed resource file installation, go to Chapter 3, “Post-Installation Tasks.”

Creating Raw Partitions

Topic	Page
Overview	139
Creating the partitions	140
Red Hat Raw device administration	142
SuSE raw device administration	143
Accessing raw devices from Adaptive Sever	144

Overview

Adaptive Server 12.5 and later 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.

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 can be reused with the disk `resize` command.

- 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:

```
ls -l <device path>
```

v **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.
- 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.

Red Hat Raw device administration

Red Hat Enterprise Linux is fully equipped with the administration tools to set up raw devices and administrate them during system start. To set up a Red Hat Enterprise Linux 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
```

SuSE raw device administration

The raw disk partitions are administered in the */etc/raw* file. This is a plain text file containing comments and examples for possible configurations:

```
# /etc/raw
#
# sample configuration to bind raw devices
# to block devices
#
# The format of this file is:
# raw<N>:<blockdev>
#
# example:
# -----
# raw1:hdb1
#
# this means: bind /dev/raw/raw1 to /dev/hdb1
#
# ...
raw1:sda7
raw2:sda8
raw3:sda9
```

Once created, you need to bind the raw devices. You do this by starting them with the script `/etc/init.d/raw`:

```
# cd /etc/init.d
# sh raw start
bind /dev/raw/raw1 to /dev/sdb1... done
bind /dev/raw/raw2 to /dev/sdb2... done
bind /dev/raw/raw3 to /dev/sdb3... done
...
```

Use the `chkconfig(8)` utility to guarantee that the raw device binding occurs during any restart:

```
# /sbin/chkconfig raw on
```

Accessing raw devices from Adaptive Sever

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
```

- You should have root privileges to execute this command. In case you do not have root permission, you may see the following output for the `raw -qa` command:

```
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, `sybserverprocs`, `sybtempdb`, or any other device.

Upgrading Servers with Replicated Databases

Topic	Page
Introduction	147
Suspending transaction processing and replication activities	148
Draining the transaction logs for primary databases	149
Draining the RSSD transaction log	151
Disabling the secondary truncation point	153

Introduction

After you have upgraded, 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). This procedure will help to ensure that all replicated data from a replicated database has made it safely to the replicate database.

Warning! Be sure to have valid dumps of the database and its transaction log before executing the procedures in the following sections.

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:

- 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.

To determine whether your existing server contains replicated databases:

- 1 Use `isql` to connect to the Server you are upgrading.
- 2 If your server uses Replication Agent to replicate data, connect to the server and execute:

```
1> sp_config_rep_agent
2> go
```

The above command will display the list of databases enabled for replication agent. If you are using other applications to replicate data, LTM or Replicator run the following command in each database (including master databases):

```
1> dbcc gettrunc
2> go
```

- 3 If the command returns “1” for “secondary 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 the upgrade process while the subscriptions are being created so that Replication Server does not interfere by attempting to access 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_helpoute` in each Replication Server System Database (RSSD) being upgraded.

The status of all existing routes should be “Active.” If any routes are not active, resolve them 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, ensure that Replication Server completely processes the pre-upgrade log.

To be sure that the transaction log has been drained off:

- 1 Wait for all remaining transactions to be replicated.
- 2 Check that a manual update on a replicate table is forwarded to Replication Server:
 - a If the database is enabled for Replication Agent:

```
1>use <dbname>
2>go
1>sp_config_rep_agent <dbname>, 'traceon',
'9201'
2>go
```

- If you cannot access to the console where Adaptive Server was started, define additionally a trace log file so you can check the information Replication Agent will send to it:

```
1> sp_config_rep_agent <dbname>, 'trace log
file', 'path for the log file'
2>go
```

- b If the database is using a different method to replicate data, you will need to check if data is replicated by consulting Replication Server queues. In such case use the following procedure:

Run the following Replication Server command:

```
1>admin who, sqm
2>go
```

Find the entry that corresponds to the inbound queue for this database:

- 1 Look for the info field in the queue_number: queue_type entry. For an inbound queue, the queue type is 1. Note the Last Seg.Block entry for the queue.

- 2 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.

- 3 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 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.

- 4 In the primary Replication Server, execute the admin who, sqm command until the last segment:block entry for the inbound queue changes.
- 5 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.

- 6 Use Notepad or another text editor to examine the dump file to make sure it contains the transaction that corresponds to the update you performed in step 4.
 - 7 Repeat steps 5–7 until the transaction that corresponds to the update is in the dump file.
- 3 Stop the application or process reading the transaction log from the primary database:

If you are using Rep Agent, log into the Adaptive Server, and stop the Rep Agent:

```
1> sp_stop_rep_agent database  
2> go
```

- 4 If you are using LTM, shut down the LTM.

After draining the transaction logs, do not allow any other activity in the databases. If activity does occur, you will need to redrain the logs.

- 5 Instruct Replication Server to reject incoming connections from Replication Agent, LTM or Replicator:

Log in to the Replication Server and suspend the Log Transfer connection from that database:

```
1> suspend log transfer from server.database  
2> go
```

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 in to 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 in to 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 in to 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 in to the Adaptive 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, shut down 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 shut down (from the previous steps).

For each primary database and replicated RSSD, disable the secondary truncation point:

- 1 If this is a replicated RSSD, log in to 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.

- 3 Truncate the primary transaction log running the command:

```
1>dump tran primary_dbname with truncate_only  
2>go
```

- 4 Zero out the Replication truncation point as soon as you can stop Replication Server:

Using the RSSD for the Replication Server, run:

```
1>rs_zeroltm primary_servername, primary dbname  
2>go
```


Using sybsystemprocs

Topic	Page
Introduction	155
Increasing the size of the sybsystemprocs database	156
Increasing device and database capacity for system procedures	157

Introduction

Verify that the sybsystemprocs database is large enough. For an upgrade, the recommended minimum size for sybsystemprocs is the larger of 120MB, 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. 120MBs 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 requirements 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.

Use `sp_helpdb` to determine the size of the sybsystemprocs database:

```
1> sp_helpdb sybsystemprocs
2> go
```

Use `sp_helpdevice` to determine the size of the sysprocsdev device:

```
1> sp_helpdevice sysprocdev
2> go
```

If the `db_size` setting is less than the required minimum, you must increase the size of `sysprocdev`.

Increasing the size of the *sybsystemprocs* database

If your current `sybsystemprocs` database does not have the minimum space required, there are 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 or expand the current device.

v **Enlarging 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 sysprocdev=40
2> go
```

In this example, “`sysprocdev`” 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 157.

- 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. See the *Utility Guide* for more information.

v Expanding the device

Newer versions of ASE have a "disk resize" command that will extend a device if there is room enough. If your current Adaptive Server version does not have that command but you would like to take advantage of it, you can use these instructions to upgrade first, and enlarge your database later:

- 1 Use "dump database" to make a copy of your existing sybsystemprocs database.
- 2 Drop the sybsystemprocs database, then recreate it. This removes your existing stored procedures and ensures that sybsystemprocs has enough free space for upgrade.
- 3 Upgrade your installation.
- 4 Expand the device containing your sybsystemprocs database. This example adds 50 megabytes to a device whose logical name is "sysprocsdev":

```
1> disk resize name=sysprocsdev, size=[50M]
2>go
```

- 5 If you dumped sybssystemprocs earlier, use "load database" to restore it, then "online database" to make it ready for use. After doing this, you will need to recreate all your Sybase-supplied stored procedures by running the install scripts, such as "*installmaster*".

v **Creating a larger system procedures device (*sysprocsdev*)**

- 1 Determine which device or devices you need to remove.

Warning! Do **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
- *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 steps.

- 2 Drop sybssystemprocs:

```
use master
go
drop database sybssystemprocs
```

```
go
```

Note In versions of Adaptive Server Enterprise prior to 15.0, you use `sysdevices` to determine which device has a low through high virtual page range that includes the `vstart` from Step 2. The device fragment whose `vstart` you used is on that device

In the 15.0 version of Adaptive Server Enterprise select the `vdevno` from `sysusages` matching the `dbid` retrieved in Step 1.

- 3 Remove the device or devices using `sp_configure`:

```
sp_configure "allow updates", 1
go
delete sysdevices
    where name in ("devname1", "devname2", ...)
go
sp_configure "allow updates", 0
go
```

The `where` clause contains the list of device names returned by the query 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 cannot be complete path names. If they are relative paths, they are relative to the directory from which your server was started.

- 5 Find another existing device that meets the requirements for additional free space, or use a disk init command similar to the following to create an additional device for `sybssystemprocs` where `/sybase/work/` is the full, absolute path to your system procedures device:

```
1> use master
2> go

1> disk init
2> name = "sysprocsdev",
3> physname = "/sybase/work/sysproc.dat",
```

```
4> size = 51200
5> go
```

Note Server versions 12.0.x and later accept but do not require "vdevno=number". In server versions earlier than 12.0.x, 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 112MB (112 x 512 = 57344). For more information on disk init, see the *Reference Manual*.

- 6 Create a subsystemprocs database of the appropriate size on that device, for example:

```
1> create database subsystemprocs on sysprocsdev =
    112
2> go
```

- 7 Run the installmaster script in the *old* Sybase installation directory, as follows:

```
isql -Usa -Ppassword -Sserver_name -iold_dir/
    scripts/installmaster -oinstallmaster.out
```

Index

A

Adaptive Server
described 10
installing multiple 134
installing via resource file 134
language, changing 36
more information about 35
RAM required by 16
stopping 116
System Administrator 41
verifying environment for upgrade 60
additional steps for downgrading 15.0.2 features 96
Administrator
Sybase System 25
alter database command 155
applications
changing after renaming objects 61
testing current 60
auditing
feature 76
automatic database expansion 98

B

Backup Server
installing via resource file 134
more information about 36
basic downgrade steps using sp_downgrade 92
batch installation of servers 134

C

character devices
located in /dev directory 140
commands
alter database 155
showserver 39

configuration parameters 95
conflicting object names 62
connection problems
analyzing 26
conventions
used in this manual xv
current applications
testing 60

D

database devices 43
raw partitions for, preparing 139
databases
image data for sample 44
interpubs 44
jpubs 45
optional 42
sample 46
default
network protocols 52
directories
installation 48
specifying new for installation 60
disk init command 139
downgrade 91
downgrade from 15.0.2 91
downgrading job scheduler 99
dsedit utility 75

E

Enterprise Edition package 6
environment variables
PATH 47
TEMP 48
error log files
troubleshooting and 116

Index

extended stored procedures (ESPs) 36

G

guest user
in pubs2 and pubs3 databases 46

H

hardware requirements 8

I

identifiers
conflicting with reserved words 61
installation directory
for unloading products 25
installing
Open Client/C 11.9.2 for Windows 47
Sybase products 39, 43, 44
installintpubs script 45
installjpubs script 45
installpix script 44
installpubs2 script 43
installpubs3 script 43
interfaces file 51, 75
interpubs sample database 44
isql
checking for replicated databases 148
determining Adaptive Server release level 60

J

jpubs sample database 45

K

keywords 96
reserved by Sybase 67

L

LDAP
enabling 51
libtcl*.cfg file
format of 51
login
“sybase” 25
administrator, Windows NT 49

M

Monitor Server
more information about 36

N

network protocols
named pipes 52
SPX 16
TCP/IP 16
Winsock 52
network software configuration
verifying 26
new sort orders 97

O

objects
conflicting names of 62
renaming 62
Open Client/C
for Windows 47
operating system
kernel patches 25
option packages 5–8
descriptions 5
Enterprise Edition 6
Security and Directory Services package 6
Small Business Edition 7
optional databases 42
options package
Developer Edition 7
options packages ??–7

ownership
 of files and directories 25

P

partition map, avoiding damage to 140
 password
 setting System Administrator 41
 PATH environment variable 47
 post downgrade 94
 privileges 140

R

raw partitions
 availability of 140
 minimum size for database devices 140
 real-time messaging 96
 recovering from a failed installation 116
 reenable Replication Server 74
 renaming objects 62
 replicated databases
 checking for 148
 restoring 75
 replication issues with downgrade 98
 Replication Server 74
 draining logs for Adaptive Server 75, 149
 reserved words
 checking for conflicts 63
 using quoted identifiers 63
 resource files 68, 134
 installation using 134
 Returning to Adaptive Server Version 15.0.2 99

S

xd2 saxd3 user 41
 sample databases 42, 46
 scripts
 installintpubs 45
 installjpubs 45
 installpix 44
 installpubs2 43

installpubs3 43
 servers
 including in sql.ini files 52
 showserver command 39
 size
 recommended for database devices 155
 Small Business Edition package 7
 SPX network protocol 16
 sql.ini file 52
 sqlupgrade utility 67
 sqlupgraderes utility 68
 stopping servers 116
 stored procedures
 changing after renaming objects 61
 recommended database device size if adding user-
 defined 155
 testing 60
 Sybase installation directory 25
 Sybase products
 installing 39, 43, 44
 Open Client/C for Windows 47
 sybssystemprocs database
 preparing for upgrade 157
 syntax
 conventions for, in this manual xvi
 sysmessages 94
 sysprocsdev device
 preparing for upgrade 157
 System Administrator
 setting password 41
 system stored procedures 94
 System tables 95

T

TCP/IP network protocol 16
 TEMP environment variable 48
 troubleshooting
 by using error log files 116
 common installation problems and solutions 112

U

unrecognized tokens or objects 95

Index

utilities

dsedit 75

sqlupgrade 67

sqlupgraderes 68

W

Windows

client product installation on 48

X

XP Server

more information about 36

starts automatically with first ESP request 53