



Installation Guide

Adaptive Server[®] Enterprise

15.7 ESD #2

Sun Solaris

DOCUMENT ID: DC35889-01-1572-01

LAST REVISED: July 2012

Copyright © 2012 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.

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.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

Java and all Java-based marks are trademarks or registered trademarks of Oracle and/or its affiliates in the U.S. and other countries.

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

IBM and Tivoli are registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

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

CHAPTER 1: Conventions	1
CHAPTER 2: Installation Task Overview	3
Installation Workflows	3
CHAPTER 3: Adaptive Server Components	5
Adaptive Server Editions	5
Determining the Edition You Have	6
Adaptive Server Options	6
Managing Adaptive Server Enterprise with Sybase	
Control Center	6
Client Applications and Utilities	7
CHAPTER 4: System Requirements	9
Memory Requirements	10
System Requirements for Clients	10
CHAPTER 5: Planning Your Adaptive Server	
Installation	13
Adaptive Server Release Bulletin	13
Obtaining a License	13
Understanding License Generation	13
Determining Your License Deployment Model	14
Comparing License Deployment Models	15
Fault Tolerance, License Grace, and	
Redundancy	16
Served License Deployment Model	16

- Determining Host IDs17
 - Determining Host IDs for Machines with Multiple Network Adapters18
 - Using Alternate Host IDs on Windows18
- Knowing the Product License Type18
 - Optional Feature Licenses19
- Generating Licenses at SPDC19
 - Logging in to SPDC and Beginning License Generation20
 - Regenerating, Renewing, and Rehosting Licenses24
- Generating Licenses at SMP25
 - Generating License Keys26
- Installation Directory Contents and Layout26**
 - PC-Client Product Descriptions and Layout29
- Performing Administrative Tasks30**
 - Creating the Sybase User Account30
- Preparing to Install Adaptive Server31**
 - Adjusting the Operating System Shared Memory Parameters32
 - Managing Java in the Database During Installations and Upgrades34
 - Multipathing34
 - Setting Up Multipathing36
- CHAPTER 6: Installing Adaptive Server37**
 - Mounting the CD37**
 - Installing Adaptive Server with the GUI Wizard37**
 - Installing Adaptive Server in Console Mode41**
 - Minimally Configuring Adaptive Server41**
 - Installing Adaptive Server Using a Response File46**
 - Creating a Response File46
 - Installing in Silent Mode47
 - Command Line Options48

Uninstalling Adaptive Server	48
Removing an Existing Adaptive Server	49
CHAPTER 7: Installing Components from the PC-Client	51
Creating a Response File for the Client	52
Installing the Client in Unattended, Silent Mode	52
Uninstalling PC-Client	53
CHAPTER 8: Starting and Stopping Sybase Control Center Agent	55
CHAPTER 9: Postinstallation Tasks	57
Verifying That Servers are Running	57
Verifying Your Connection to the Servers	58
Testing the Installation and Network Connections	58
Configuring libtcl.cfg for LDAP	59
Adding a Server to the Directory Service	59
Setting the System Administrator Password	60
Installing Sample Databases	60
Default Devices for Sample Databases	61
Running the Database Scripts	62
Installing the interpubs Database	62
Installing the jpubs Database	63
Maintaining Sample Databases	64
CHAPTER 10: Adaptive Server Upgrades	65
Upgrading Adaptive Server	65
Considerations for Component Integration Services	68
Preparing to Upgrade	68
Changes in Adaptive Server Directories	69
Performing Preupgrade Tasks	70

- Upgrade System Tables and Stored Procedures72
- Reserved Words73
 - Running a Reserved Word Check73
 - Addressing Reserved Words Conflicts74
 - Quoted Identifiers74
- Preparing the Database and Devices74
- The sybprocsdev device76
 - Increasing the Size of the subsystemprocs Database76
 - Increasing Device and Database Capacity for System Procedures77
- Upgrade to Adaptive Server 15.7 ESD #279**
 - Upgrading Interactively Using sqlupgrade79
 - Upgrading Noninteractively Using sqlupgraderes80
 - Manually Upgrading to Adaptive Server81
 - Installation of Version 15.7 ESD #2 over an Existing 15.x Adaptive Server82
 - Determining the Adaptive Server Version82
 - Backing Up Adaptive Server82
 - Installing Adaptive Server Using Binary Overlay83
- Postupgrade Tasks84**
 - Running the instmsgs.ebf Script85
 - Restoring Functionality in Adaptive Server After
 - Upgrading85
 - Reenabling Auditing86
 - Updating Threshold Procedures for Audit Segments86
 - Reenabling Replication Server After the Dataserver
 - Upgrade87
 - Restoring Replication on Replicate Databases ... 87
 - Restoring Replication on Primary Databases88
- Migrate88**
 - Migrating Data Using a Dump-and-Load Method89
 - Using bcp to Migrate Data89

Migrating Data by Replacing the Binary	90
Adaptive Server Components and Related Products	90
Upgrading Job Scheduler	90
Upgrading Job Scheduler Templates	91
Upgrading High Availability and Cluster Support	95
Upgrading Java in the Database	96
Enabling Java in the Database in a High Availability System	96
Upgrading Backup Server	97
Upgrading Databases Using Dump and Load	97
How Compiled Objects Are Handled When Upgrading Adaptive Server	97
Finding Compiled Object Errors Before Production	98
CHAPTER 11: Adaptive Server Downgrades	101
Preparing to Downgrade Adaptive Server	101
Downgrading from Adaptive Server 15.7 ESD #2	102
Downgrading to an Earlier Version of Adaptive Server 15.7	104
Dumping a 15.7 ESD #2 Database to Load into a 15.7 or 15.7 ESD #1 Adaptive Server	106
Additional Considerations for New Features Used	107
Downgrading Job Scheduler	108
Downgrading an Adaptive Server That Uses Encryption	109
Downgrade Considerations for Adaptive Server with Replicated Databases	110
Post-Downgrade Tasks for Adaptive Server	110
CHAPTER 12: Troubleshooting SySAM Errors	113
Where to Find License Error Information	113
Problems and Solutions	113
Installing for the First Time	119

- Possible Causes of a License Server Failure 119
- Solution for Problem: Product Cannot Find License for
Optional Features120
- Unserved License Deployment Models120
- Served License Deployment Models122
- Contacting SySAM Technical Support123**

- CHAPTER 13: Troubleshoot the Server125**
- Error Logs for Installation Utilities126**
- Error Logs for Sybase Servers126**
- Troubleshooting Common Installation Problems127**
 - Stopping Adaptive Server After a Failure 129
 - Recovering from a Failed Installation129
 - If the Installation Quits While You Are
Configuring Adaptive Server129
- If Adaptive Server Fails the Preupgrade Eligibility Test
.....130**
- When an Upgrade Fails130**
 - If You Can Identify the Cause of the Upgrade Failure
.....130
 - Restoring Databases After a Failed Upgrade131
 - Recovering from a Failed Upgrade131
 - If You Cannot Identify the Cause of an Upgrade Failure
.....133

- CHAPTER 14: Obtaining Help and Additional
Information135**
- Technical Support135**
- Downloading Sybase EBFs and Maintenance Reports
.....135**
- Sybase Product and Component Certifications136**
- Creating a MySybase Profile136**
- Accessibility Features136**

Index139

Contents

CHAPTER 1 Conventions

These style and syntax conventions are used in Sybase® documentation.

- In a sample window display, commands to be entered exactly as shown appear 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, 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. SQL syntax conventions

Key	Definition
command	Command names, command option names, utility names, utility flags, and other keywords are in a bold 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.
()	() 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, which are typed as part of the command.

To successfully install and configure Adaptive Server[®] Enterprise, use the configuration guide along with this installation guide.

The Adaptive Server installation guide explains how to unload the Adaptive Server software from the distribution media to your hard disk, and perform the minimal set of additional tasks so that you can start Adaptive Server on your machine.

The configuration guide provides instructions for:

- Reconfiguring certain attributes of your existing Sybase Adaptive Server Enterprise, Backup Server[™], Adaptive Server Enterprise Monitor[™], and XP Server[™] to meet your specific needs.
- Creating network connections.
- Configuring optional functionality.
- Performing operating system administrative tasks.
- Working with system administration issues that are relevant to Adaptive Server running on UNIX platforms. This manual supplements the *System Administration Guide* and the *Performance and Tuning Series*.

Installation Workflows

Workflows define a complete path for planning, installing, and upgrading.

Choose the workflow that best describes your scenario.

Tip: Print this topic and use it as a checklist.

Whether you plan to install or upgrade Adaptive Server:

1. Identify the components and options to install or upgrade.
2. Obtain a license.

Installing Adaptive Server for the First Time

1. Plan your installation and review system requirements.
2. Install Adaptive Server.
3. Perform postinstallation tasks.

Upgrading to a New Version

1. Determine your upgrade eligibility, plan your installation, and review system requirements.
2. Run the **preupgrade** utility to prepare Adaptive Server for upgrading.
3. Upgrade Adaptive Server.
4. Perform postinstallation tasks.

Uninstalling Adaptive Server

Uninstall Adaptive Server. See *Uninstalling Adaptive Server* on page 48.

Adaptive Server® Enterprise is based on a client/server model, communicating with its clients over the network via the Tabular Data Stream™ (TDS) protocol. Each client process may execute on one machine and communicate with a database server on the same or a different machine.

Adaptive Server runs as an application on top of an operating system. The hardware that the operating system runs on is completely transparent to Adaptive Server, which sees only the operating system's user interfaces. To enhance performance on multiprocessor systems, configure multiple processes (engines).

Adaptive Server is divided into a DBMS component and a kernel component. The kernel component uses the operating system services for process creation and manipulation, device and file processing, and interprocess communication. The DBMS component manages SQL statement processing, accesses data in a database, and manages different types of server resources.

Adaptive Server Editions

Sybase® offers various editions of Adaptive Server Enterprise.

- Enterprise Edition – has no limit on scalability and can run all available options that can be purchased separately.
- Small Business Edition – has limited scalability, and runs a limited set of options that are purchased separately.
- Developer Edition – has limited scalability and includes many of the options that are included in the Enterprise Edition.

See the *Sybase Software Asset Management Users Guide* for information about unlocking editions and optional features.

There is a change to the license quantity used by per-CPU and per-chip license types. Adaptive Server version 15.7 ESD #2 and later checks out the same number of licenses as the number of cores on the machine (or chips, depending on the license type) regardless of any configuration settings. This is a correction from earlier versions of the server, in which if Adaptive Server was licensed per CPU or CPU chip, the requested license quantity was reduced if the **max online engines** configuration parameter was set to fewer than the number of CPUs on the machine.

The Adaptive Server installer automatically installs a SySAM license server if you choose the full installation option or enter a served license when prompted for the SySAM license key. You can also install the license server using the installer's custom installation option. See the *Sybase Software Asset Management Users Guide* to generate licenses.

Determining the Edition You Have

If you already use Adaptive Server, use the **sp_lmconfig** system procedure to identify the version.

Enter:

```
sp_lmconfig 'edition'
```

Adaptive Server returns a value of EE, SE, or DE based on the edition that is currently running. See the *Reference Manual* for more information about **sp_lmconfig**.

Adaptive Server Options

Sybase offers various optional features for Adaptive Server, such as data compression, partitions, and encrypted columns.

Option	Description
Data compression	Enables compression for regular and large object data, which uses less storage space for the same amount of data, reduces cache memory consumption, and improves performance due to lower I/O demands.
Security and directory services	Provides lightweight directory services and network-based authentication and encryption using SSL and Kerberos.
Partitions	Enables semantic partitioning for table row data.
Encrypted columns	Increases security parameters and allows for addition of datatypes.
Tivoli Storage Manager	Enables the database to back up and restore operations to IBM Tivoli Storage Manager.
In-memory database	Provides zero-disk-footprint in-memory database support that is fully integrated with Adaptive Server for high-performance transactional applications. Provides performance enhancements to disk-resident databases with relaxed durability properties.

Adaptive Server editions and optional features are unlocked by SySAM licenses. See the *Sybase Software Asset Management Users Guide*.

Managing Adaptive Server Enterprise with Sybase Control Center

Sybase Control Center provides a single, comprehensive, Web-administration console for real-time performance, status, and availability monitoring of large-scale Sybase enterprise servers. Sybase Control Center combines a modular architecture, a rich client administrative console, agents, common services, and tools for managing and controlling Sybase products. It

includes historical monitoring, threshold-based alerts and notifications, alert-based script execution, and intelligent tools for identifying performance and usage trends.

Although the PC-Client CD installation includes the Adaptive Server plug-in for Sybase Central (which contains software for managing Sybase servers), Sybase recommends that you use Sybase Control Center to monitor Adaptive Server, as the Sybase Central plug-in does not support any of the new features introduced in this version of Adaptive Server.

The Adaptive Server installer installs the Sybase Control Center (SCC) remote command and control agent with Adaptive Server, however, does not install the SCC management UI, which manages and monitors activities.

To install the SCC management UI, use the SCC installation CD or DVD, or download it from <http://downloads.sybase.com>. In a production environment, Sybase recommends that you install the SCC server on a machine other than the one you plan to run Adaptive Server on.

Client Applications and Utilities

The PC-Client installer contains client applications and utilities you can use to access the Adaptive Server to execute queries or manage the server. It also contains the Sybase Open Client/ Open Server™ Software Developers Kit, which you can use to develop applications that access the server, and ODBC, OLE DB, and ADO.NET clients.

The Sybase PC-Client CD includes the Software Developer Kit (SDK) for Windows platforms, consisting of:

- Embedded SQL™
 - Embedded SQL™/C (ESQL/C)
 - Embedded SQL/Cobol (ESQL/Cobol) – 32-bit only
- XA-Library™ – XA Interface Library for ASE Distributed Transaction Manager
- Adaptive Server Enterprise extension module for Python – 64-bit only
- Additional connectivity language modules
- Open Client™ (CT-Library, DB-Library™)
- Microsoft Cluster Server Resource Type for ASE – 64-bit only
- Adaptive Server Enterprise database driver for Perl – 64-bit only
- Adaptive Server Enterprise extension module for PHP – 64-bit only
- Interactive SQL
- Sybase Central™
- Adaptive Server plug-in for Sybase Central
- ASE ADO.NET Data Provider
- ASE OLE DB Provider by Sybase
- ASE ODBC Driver by Sybase
- ASE plug-in

CHAPTER 3: Adaptive Server Components

- QPTune
- jConnect™ for JDBC™ 7.0
- SySAM License Utilities
- SDC Administration Utility

In addition to the SDK, the PC-Client CD includes PowerDesigner Physical Architect, a tool for data modeling, including database design, generation, maintenance, reverse engineering, and documentation for database architects.

Before installing Adaptive Server, ensure that your system is updated with the latest patches and system requirements. Do not use a patch that is earlier than the version suggested for your operating system. Use the patch recommended by the operating system vendor, even if it supersedes the patch listed.

To list all currently installed patches and display the operating system version level, enter:

```
showrev -p
```

Sun Solaris Requirements

All Solaris systems support the TCP protocol. Platform system requirements are:

- Sun4u Architecture – Solaris 10
- AMD Architecture – Solaris 10

Solaris 64-bit platform system requirements are:

- Processor – 64-bit
- Minimum RAM for Adaptive Server – 106MB
- Default user stack size – 94KB
- Memory per user connection – Approximately 245KB

Sun Solaris AMD platform system requirements are:

- Processor – 64-bit
- Minimum RAM for Adaptive Server – 106MB
- Default user stack size – 98KB
- Memory per user connection – Approximately 247KB

Note: Adaptive Server supports the 128-bit Zettabyte File System (ZFS) on Solaris 10. To guarantee data integrity, set DSYNC or Direct IO on. You can turn DSYNC or Direct IO off only for temporary databases that do not require recovery.

Sun system patch requirements:

- For Solaris 10 SPARC
 - 120753-05
 - 120048-03
 - 144190-03
- For Solaris 10 on AMD
 - 120754-05
 - 120049-03

CHAPTER 4: System Requirements

- 144191-02

Solaris disc space requirements:

- Typical installation
 - 64-bit: 985MB
 - AMD: 890MB
- Default databases created during installation – 306MB
- **Total**
 - 64-bit: 1141MB
 - AMD: 1046MB

Ensure that you have the required operating system patches for the Java Runtime Environment (JRE) version 6.

Information about the required operating system patches is available from <http://java.sun.com>.

Memory Requirements

Each operating system has a default maximum shared-memory segment. Make sure your operating system is configured to allow the allocation of a shared-memory segment at least as large as the maximum memory value.

The amount of memory Adaptive Server can use is:

- On Sun Solaris 10 (SPARC) – 16 EB
- On Sun Solaris 10 on AMD (x86) – chip-dependent

System Requirements for Clients

Check the system requirements for the machine on which you plan to install the PC-Client.

Type	Requirement
Product	PC-Client
Hardware	P4 1.0GHz
Operating system	Windows Server 2008 R2, Windows Vista, Windows 7, Windows XP
Suggested minimum RAM	512MB

Note: If you are using ODBC, OLE DB, or ADO.NET drivers, make sure that Microsoft .NET Framework 2.0 Service Pack 1 is installed on your Windows machine. To confirm that it is

installed, select **Control Panel > Add or Remove Programs** and verify that .NET Framework appears in the list of currently installed programs.

Planning Your Adaptive Server Installation

Before you install or upgrade, prepare your environment.

- Identify the components and options to install or upgrade.
- Obtain a license.

Note: If you are using a served license, you must install the SySAM license server version 2.1 or later.

- Make sure your system meets all requirements for your installation scenario and planned use.

Adaptive Server Release Bulletin

Obtain updated information from the release bulletin.

The release bulletin contains last-minute information about installing and upgrading Adaptive Server software.

The latest release bulletins are available at the Product Documentation Web page at <http://www.sybase.com/support/manuals>.

Obtaining a License

Sybase® Software Asset Management (SySAM) is the Sybase product licensing and asset management system, which is built on FLEXnet technology from Flexera Software.

This section includes information about the SySAM licensing system that may be useful when you generate product licenses using the procedures in this document.

For complete information on SySAM, see the *SySAM Users Guide*.


Understanding License Generation

When you purchase SySAM 2–enabled Sybase products, you must generate, download, and deploy SySAM product licenses.

- If you purchased your product from Sybase or an authorized Sybase reseller, go to the secure Sybase Product Download Center (SPDC) at <https://sybase.subscribenet.com> and log in to generate license keys. The license generation process may vary slightly, depending on whether you ordered directly from Sybase or from a Sybase reseller.

- If you ordered your product under an SAP® contract and were directed to download from SAP Service Marketplace (SMP), you can use SMP at <http://service.sap.com/licensekeys> to generate license keys for Sybase products that use SySAM 2-based licenses.

Figure 1: SySAM 2 Licensing Process



When you purchase SySAM 2-licensed products from a Sybase reseller, your product package may include a Web key certificate that contains the location of the SPDC Web key login page at <https://sybase.subscribenet.com/webkey> and the activation string to use for the login name.

Determining Your License Deployment Model

Before you generate licenses for SySAM 2-enabled products, determine the license deployment model to use.

There are two license deployment models from which to choose: the served license model or the unserved license model.

When you use the unserved license deployment model, you generate and download a separate license for each machine where the product runs.

Sybase recommends that you use a license server if you run products on more than several machines, as this simplifies and centralizes license administration providing more control over software assets.

After you download the unserved licenses, you can install your Sybase products.

Comparing License Deployment Models

Learn about license deployment models.

Unserved Licenses	Served Licenses
Licenses can be used only on the machine for which the license was generated.	Licenses can be distributed from a network license server to products running on any network machine.
<p>Generate licenses at SPDC or SMP for each machine that will run the product:</p> <ol style="list-style-type: none"> 1. Specify the host ID of the machine where the product will run. 2. Generate a license for that machine. 3. Save the license to the specified machine. 4. Repeat steps 1 – 3 for each machine where the product will run. 	<p>Generate licenses at SPDC or SMP for products running on multiple machines:</p> <ol style="list-style-type: none"> 1. Specify the host ID of the license server. 2. Specify the number of required licenses. 3. Save the licenses to the license server host machine.
No license administration is required. However, when new licenses are required for product updates, you must update and deploy each license for each machine where the product update will run.	The license server requires administration. When new licenses are required for product updates, SPDC or SMP lets you update all licenses for a specific license server in a single step.
No license reporting or asset management capabilities are available.	Allows license monitoring and reporting of license use, capacity planning, and asset management using SAMreport.
Installed locally and always available.	Requires a functioning license server and network. If the license server or network fails, you must fix the problem or install an alternate license server before the product grace period expires.
If a machine where the product is running fails, you must regenerate all of its licenses and deploy those licenses to the replacement machines.	<p>If a machine where the product is running fails, you can move the product to a new machine, and it will acquire licenses from the running license server.</p> <p>If the license server host machine fails, use the Manage License Hosts functionality at SPDC or SMP to move its licenses to a new network license server host.</p>
License files are distributed across each machine running a product, and therefore they are difficult to manage and control.	License files are centrally located and managed.

Unserviced Licenses	Served Licenses
Unserviced Standalone Seat (SS) licenses do not allow product use via Remote Desktop Connection or other terminal services clients.	Products can be used via Remote Desktop Connection or other Terminal Services clients, irrespective of the type of license in use.

Fault Tolerance, License Grace, and Redundancy

Sybase products check out licenses on start-up, and perform periodic heartbeat checks to verify that licenses are still available. If a license is unavailable, the product determines whether it can be awarded on grace.

The grace period starts from the time a license was last in use and persists for 30 days for server products and 15 days for tool products. At the end of the grace period, the product performs an orderly shutdown, if running, or fails to start up, unless the license (or a replacement license) becomes available, at which point the cycle continues with the heartbeat, noting the last use of the license.

This tolerance for temporary licensing failure is usually sufficient. However, in certain circumstances, you may choose to use “three-license-server redundancy.” For example:

- A grace period is not provided for Sybase Floating License (FL) type licenses.
- Standby copy systems rarely qualify for grace as they are unlikely to have been used within the last 30 days.
- Company policy dictates the use of redundancy.

For three-server redundancy, use three machines that each:

- Run the same version of the SySAM license server.
- Have good intermachine communication.
- Use a separate copy of the same license files.

A product that uses per-processor licensing checks out a license quantity equal to the number of processors it can use, or awards runtime grace if an insufficient quantity is available. If the number of processors is dynamically increased while the product is running and the product is unable to check out additional licenses, runtime grace is also awarded. If the additional licenses are not made available within the runtime grace period the product shuts down. Decreasing the number of processors that a product can use while it is running does not reduce the required license quantity. You must restart the product on the correct number of processors.

Served License Deployment Model

When you choose the served license deployment model, licenses are deployed to one or more SySAM license servers.

After you install the required license servers and generate served licenses, you can install your Sybase products. You can configure the products to obtain licenses from one or more license servers.

License Servers

The license server is a lightweight application that serves licenses to products running on any operating system.

A license server can run on a low-end machine or a machine with spare cycles. For example, a license server running on a Solaris UltraSparc-60 serving 100 different licenses to 200 product instances, used 50MB of memory, 5 minutes of CPU time per week, and 100MB of disk space per year.

To download the license server software and installation instructions, go to the SySAM Standalone License Server Install page at <http://www.sybase.com/sysam/server> and select the download link.

Note: You cannot start the license server until there is at least one served license saved to the `licenses` directory on the license server host machine.

Determining Host IDs

When you generate licenses at SPDC or SMP, you must specify the host ID of the machine where the licenses will be deployed.

- For unserved licenses, specify the host ID of the machine where the product will run. If you are running a product with per-CPU or per-CHIP licensing that supports SySAM sub-capacity, and you want to run the product in a virtualized environment, see *SySAM Sub-capacity Licensing* in the *SySAM Users Guide* for information about determining the host ID for unserved licenses.
- For served licenses, specify the host ID of the machine where the license server will run.

SPDC or SMP remembers the host information so that you can select the same license server when generating additional licenses.

To determine the host ID of the machine, run the **lmutil** utility from a terminal window or the Windows command prompt. For example:

```
lmutil lmhostid
```

Note: You can download the **lmutil** utility from the Flexera Software Web site at http://www.globes.com/support/fnp_utilities_download.htm.

You can also determine the host ID using native operating system commands. See the Frequently Asked Questions topic "What is my Host ID?":

- SPDC: <https://sybase.subscribenet.com/control/sybs/faqs#30-4>
- SMP: <https://websmp208.sap-ag.de/~sapidb/011000358700001006652011E>

Determining Host IDs for Machines with Multiple Network Adapters

On some platforms, the host ID is derived from the network adapter address.

If the machine where your product is installed or where the license server is hosted has multiple network adapters, running **lmutil lmhostid** returns one host ID for each network adapter, and returns output similar to this:

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

It is important that you:

- Choose only one host ID to enter during license generation.
- Use the value associated with the primary wired Ethernet adapter.
- Do not use values associated with internal loopback or virtual adapters.

If you cannot determine the host ID using the **lmutil lmhostid** output, use the native operating system commands to get additional details.

Using Alternate Host IDs on Windows

If your Windows machine does not have a network adapter, SySAM allows you to use an alternate host ID based on the hard disk serial number.

1. At a Windows command prompt on the machine where you will deploy your license, enter:

```
lmutil lmhostid -vsn
```

This returns output similar to:

```
The FLEXlm host ID of this machine is
"DISK_SERIAL_NUM=70ba7a9d"
```

2. Use the complete output (DISK_SERIAL_NUM=70ba7a9d) for the host ID value requested during license generation.

Knowing the Product License Type

Sybase sells products under different license types, which entitle the customer to different rights of use. For example, whether a product is entitled to be used in production, or standby, or development and test environments.

The license type determines the quantity of licenses required; dictating, for example, whether a license is required for each machine; or for each CPU, or CPU chip; or per terabyte of storage.

For example, a product purchased under the CPU License (CP) type requires one license per CPU on the machine, partition, or resource set where it will run. The same product purchased under the Server License (SR) type requires a single license for the machine, partition, or resource set.

Your account can have the same product licensed under multiple license types. Make sure you select the correct license type when you generate your SySAM license file.

When you generate licenses, each license type is presented on the License Information screen in bold font. For example:

License Type: CPU License (CP)

Note: The End User License Agreement for your region, available from the Sybase Software Licenses Web page at <http://www.sybase.com/softwarelicenses>, contains a definition for each license type. It also describes rights of use, such as whether licenses can be used only on a specific machine, partition, or resource set, or can float and whether you need more than one license for use on a specific machine, partition, or resource set. In addition, check the Product Specific License Terms Web page at <http://www.sybase.com/pslt>.

Optional Feature Licenses

Some Sybase applications are offered as base products with optional features that require a separate license.

Customers can mix license types. For example, customers can order Adaptive Server® Enterprise under the Server License (SR) license type and order an optional feature (for example, High Availability or Enhanced Full Text Search) under the CPU License (CP) type.

Optional features are licensed only with the base product of the same product edition. For example, if you order Adaptive Server Enterprise, you cannot use a Small Business Edition optional feature license with the Enterprise Edition base product. In addition, the entitled use of the license types must be compatible; for example, both types should allow use in a production environment.

Generating Licenses at SPDC

Before you log in to SPDC and generate licenses, use this information as a reminder of the information you should have available and the tasks that you should have completed.

Table 2. Information Needed Before License Generation

Required Information or Action	License Model		Description
	Served	Unserv ed	
License deployment model	X	X	Decide whether to use a served or unserved license deployment model. Typically, this is a company-wide decision that is made only once. Therefore, this is one of the most important things to determine before license generation.

Required Information or Action	License Model		Description
	Served	Unserv ed	
Product machine host ID		X	Determine the host ID of the machine, or machine partition where the product will run.
License server – download and install	X		Download and install the SySAM license server before you generate the product’s licenses, and before you install the product.
License server host ID	X		Determine the host ID of the machine where the license server will run.
License server host name	X		Determine the host name of the machine where the license server will run.
License server TCP/IP port numbers	X		Determine the port numbers on which the license server will listen for license requests. Note: If you do not specify the license server port numbers during license generation, the license server uses the first available ports in the range 27000 to 27009. If a firewall exists between the server and the client machines, fix the license server port numbers to allow access to the ports. See <i>Access Through a Firewall or VPN</i> in the <i>SySAM Users Guide</i> .

Logging in to SPDC and Beginning License Generation

Once you log in to SPDC and begin license generation, complete your product’s license generation using the procedure that applies to the license deployment model you chose, either for generating served licenses or for generating unserved licenses.

If your product has been ordered from a Sybase reseller, perform the “Web Key step” instructions, where noted.

1. Go to the SPDC login page at <https://sybase.subscribenet.com>.
 - Go to the SPDC Web Key Registration page at <https://sybase.subscribenet.com/webkey>.
2. Enter the login ID and password, then click **Login**. If you forget your password, click **Password Finder**.

Your password is sent to you in an e-mail message.

- Enter the authorization string from the Web Key Certificate provided by the reseller when your Sybase product was purchased, then click **Submit Web Key**.

Note: If you do not know your SPDC account login ID and password, or Web Key Certificate authorization string, contact the person who ordered your product.

- At the Web Key Registration page, enter your account information, then click one of these options:
 - **Submit My Registration Information** – to register your product using your direct account information.
 - **Anonymous Activation** – to activate your product anonymously.
- 3. Select the product family that includes the product for which you want to generate a license; for example, Adaptive Server Enterprise.
- 4. Depending on the product family you choose, you may see an additional Product Information page.
 1. Product suites – if your product is included in one or more suites, select the suite that includes your product; for example, ASE Small Business Edition.
 2. Product version and platform – select the product version, name, and operating system that matches your product order.
- 5. If this is the first time you selected a specific product edition and platform, you must accept the Sybase License Agreement before you are allowed to generate a license for that product.
- 6. If your product software installation requires a license key (also referred to as the license file), click **License Keys** on the Product Download page.
- 7. On the License Information page:
 1. Select the option button to the left of the product for which to generate a license; for example, “CPU License (CP), ASE Enterprise Edition 15.7 for Sun Solaris SPARC 64-bit.”
 2. Scroll down and click **Select to Generate**.
- 8. In the Generate Licenses wizard, choose a license deployment model:
 - **Served license** – go to *Generating Served Licenses* in the *Sybase Software Asset Management (SySAM) 2 Users Guide* to complete the license generation and download process.
 - **Unserved license** – go to *Generating Unserved Licenses* in the *Sybase Software Asset Management (SySAM) 2 Users Guide* to complete the license generation and download process.

Note: Some Sybase products or license types do not offer a choice of license deployment models and do not display this page. If this is the case, continue through the Generate Licenses wizard to generate your license.

Click **Next**.

Generating Unserved Licenses

Generate and download an unserved license for your product.

1. Enter the number of machines (up to 10) for which to generate unserved licenses, and click **Next**.
2. Enter:
 - **Node Host ID** – enter the host ID of the machine where the product will be run. If you do not know the host ID, select **What Is My Host ID?** or see *Determining Host IDs* in the *Sybase Software Asset Management (SySAM) 2 Users Guide*.
 - **Host Name** – enter the machine host name.

For some license types, you must enter the number of licenses to generate. If you are unsure of the license quantity, select **How Many Licenses Should I Generate?**.

3. Click **Generate**.
4. When the license has generated, review the information on the View Licenses page, then, if the license information is correct, select one of:
 - If you generated only one license, click **Download License File**.
 - If you generated several licenses, click **Download All Licenses for Host**.

Note: Before you download and save generated licenses, you can select **Print Friendly** to print a paper copy of the license, or **License Overview** to return to the License Information page and generate additional licenses.

- To correct license information before downloading the license, click **License Overview**, select the license to be corrected and click **Check In** to reset the license to its original state. Repeat the license generation process.
5. If you chose to download licenses, when the File Download dialog box opens, click **Save**.
6. Save the generated licenses with a `.lic` file name extension. Although `$$YBASE / SYSAM-2_0/licenses` is typically the directory where unserved licenses are located, your product may have a different requirement. See your product installation guide and release bulletin for product-specific information.

Note: If you do not save the license files with a `.lic` extension, SySAM does not recognize the licenses.

Next, install your licensed product using the instructions in the product installation guide and release bulletin.

Generating Served Licenses

Generate and download a served license for your product.

1. Enter the number of licenses to generate and click **Next**.

If you are unsure of the license quantity, select **How Many Licenses Should I Generate?**.

2. Select an existing license server host, or enter the host ID, and an optional host name and port number, of a new license server host.

When completing this step, use these recommendations:

- If you do not know the license server host ID, select **What Is My host ID?** or see *Determining Host IDs* in the *Sybase Software Asset Management (SySAM) 2 Users Guide*.
- Although the host name is optional, Sybase recommends that you provide the host name to ease future license administration.
- The port number is optional unless your product configuration is going to use three-server redundancy (see the next bullet item). Any unused port number between 0 and 64000 is valid. On UNIX, choose a port number greater than 1024; most port numbers less than 1024 are privileged port numbers. If you do not specify a TCP/IP port number, a default port between 27000 and 27009 is used.
- To generate a license for a three-server redundant configuration, enter the required information: license server host ID, host name, and port number for all three machines. Enter a port number outside of the 27000 to 27009 range. You must enter the fully qualified domain name (FQDN) as the host name if any client machine requires an FQDN to contact a license server machine.

Note: You cannot upgrade SySAM 1.0 licenses and use them in a three-server redundant configuration.

3. Click **Generate**.
4. When the license has generated, review the information on the View Licenses page, then, if the license information is correct and you do not need to generate additional licenses, select one of:
 - If you generated only one license, click **Download License File**.
 - If you generated several licenses, click **Download All Licenses for Host**.

Note: Before you download and save generated licenses, you can select **Print Friendly** to print a paper copy of the license.

- a) To correct license information, click **License Overview**, select the license to be corrected and click **Check In** to reset the license to its original state. Repeat the license generation process, starting with step 1 of this procedure.
 - b) To generate additional licenses, click **License Overview** and repeat the generation process for the additional product licenses.
5. When the File Download dialog box opens, click **Save**.
 6. Save the license files with a `.lic` file name extension to the `SYSAM-2_0/licenses` directory of the license server installation.

Warning! If you do not save the license files with a `.lic` extension, SySAM does not recognize the licenses.

7. After you save the license files to the license server, enter on the machine where your license server is running:

```
sysam reread
```

Your new licenses are registered with the license server.

Regenerating, Renewing, and Rehosting Licenses

Learn when licenses must be regenerated, renewed, and rehosted.

Licenses must be upgraded to the latest versions when:

- Your support contract is renewed. Updated licenses enable you to use the newest versions of the products that become available during your support period.
- A new product version is added to your entitlement.

Licenses need to be rehosted when:

- You entered incorrect machine information when generating the license.
- The host ID of the machine has changed due to a hardware upgrade.
- You transfer products to a new machine.

You can either upgrade or rehost all licenses that were previously generated for a specific host (see *Managing License Hosts* in the *Sybase Software Asset Management (SySAM) 2 Users Guide*), or you can modify individual licenses (see *Modifying Individual Licenses* in the *Sybase Software Asset Management (SySAM) 2 Users Guide*).

Managing License Hosts

Learn how to manage licenses of license hosts.

1. Select **Licenses** on the left pane of the SPDC main page, and then select **Manage License Hosts**.
2. Select or search for the existing license host.
3. Select one of these buttons to perform the desired operation on all licenses previously generated for the specified host:
 - **Upgrade All** – upgrades all licenses on the host to their latest versions.
 - **Re-host All** – migrates all licenses to a new host. You provide details on the next screen.
 - **Return All** – returns all licenses to the available licenses pool for future redeployment.

If you select **Upgrade All** or **Re-host All**, new licenses are sent via e-mail message to your e-mail address, and to any additional e-mail addresses you specify.

Modifying Individual Licenses

Learn about modifying individual licenses.

1. Follow instructions in *Logging in to SPDC and Beginning License Generation* in the *Sybase Software Asset Management (SySAM) 2 Users Guide* through the step for selecting options that identify the license you want to change.
2. On the License Information page, select the desired license order.
3. Select one of:
 - **Check in** – returns the license to the available license pool. To rehost the license, follow instructions in *Generating Unserved Licenses* in the *Sybase Software Asset Management (SySAM) 2 Users Guide* or *Generating Served Licenses* in the *Sybase Software Asset Management (SySAM) 2 Users Guide*.
 - **Upgrade** – upgrades the license to its latest version. Click **Upgrade** to generate a new license file. Download the license file, save it, and deploy it. The **Upgrade** option appears only if newer versions of the license are available.

Note: Each product limits the number of times you can check in and rehost its licenses. If you reach your check-in limit, and the **Check-in** option does not appear, contact your company's designated person in Sybase Technical Support.

Generating Licenses at SMP

Before you log in to SMP and generate licenses, use this information as a reminder of the information you should have available and the tasks that you should have completed.

Table 3. Information Needed Before License Generation

Required Information or Action	License Model		Description
	Served	Unserved	
License deployment model	X	X	Decide whether to use a served or unserved license deployment model. Typically, this is a company-wide decision that is made only once. Therefore, this is one of the most important things to determine before license generation.
Product machine host ID		X	Determine the host ID of the machine, or machine partition where the product will run.
License server – download and install	X		Download and install the SySAM license server before you generate the product's licenses, and before you install the product.

Required Information or Action	License Model		Description
	Served	Unserv ed	
License server host ID	X		Determine the host ID of the machine where the license server will run.
License server host name	X		Determine the host name of the machine where the license server will run.
License server TCP/IP port numbers	X		Determine the two port numbers on which the license server will listen for license requests.

Generating License Keys

If you have purchased Sybase products that use SySAM 2-based licenses under SAP contract and are directed to download from SAP Service Marketplace (SMP), you can use SMP to generate license keys.

1. Go to the SAP Marketplace main page at <http://service.sap.com>.
2. Select **SAP Support Portal**.
3. Log in using your SMP credentials.
4. Select **Keys & Requests > License Keys**.
5. Follow the instructions in the "How to generate license keys for SAP Sybase products" presentation available under the "Documentation and Helpful Resources" quick access link.

Installation Directory Contents and Layout

Adaptive Server includes server components that are installed into specific directories.

Product	Description
Adaptive Server	<p>Installed into the ASE-15_0 directory.</p> <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. • 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 ASE-15_0/jobscheduler/.

Product	Description
Software Developer Kit (SDK)	<p>Installed into the <code>OCS-15_0</code> directory.</p> <ul style="list-style-type: none"> • Open Client (<i>Client Library, dblink</i>) • ESQL/C • ESQL/COBOL • XA • ODBC (Windows, Solaris SPARC, Solaris x64, Linux Power, HP-UX Itanium, IBM AIX) – a driver used to connect to Adaptive Server from ODBC-based applications. • OLEDB (Windows only) – a provider used to connect to Adaptive Server from OLEDB-based applications. • ADO.NET (Windows only) – a provider used to connect to Adaptive Server from .NET based applications.
shared directory	<p>The shared directory contains components and libraries that are shared by several other components. In earlier versions, this directory was named <code>shared-1_0</code>.</p> <ul style="list-style-type: none"> • Sybase Central 6 – is a Java-based framework used by the system management tools. • Java Runtime Environment (JRE) – is a runtime Java virtual machine used to execute Java-based programs, such as Sybase Central. Adaptive Server includes the JRE. In typical installations, by default, the JRE is fully installed. In custom installations, if you select a component for installation that requires the JRE, the JRE is also automatically selected for installation. It is located in the <code>shared/JRE-6_0 *</code> directory.
Sybase Control Center	<p>Sybase Control Center logs and utilities – files related to the Web-based tool for monitoring the status and availability of Adaptive Servers. They are installed into the <code>\$\$SYBASE/SCC-3_2</code> directory.</p>
Adaptive Server plug-in	<p>Installed into the <code>\$\$SYBASE/shared/sybasecentral6*</code> directory.</p> <hr/> <p>Note: Sybase recommends that you use the Sybase Control Center to monitor Adaptive Server instead, as the plug-in does not support any of the new features of Adaptive Server 15.7 ESD #2. See <i>Managing Adaptive Server Enterprise with Sybase Control Center</i> on page 6.</p> <hr/>
Language modules	<p>Installed into the <code>locales</code> directory. Provides system messages and date/time formats.</p>

CHAPTER 5: Planning Your Adaptive Server Installation

Product	Description
Character sets	Installed into the <code>charsets</code> directory. Provides character sets available for use with Adaptive Server.
Collation sequences	Installed into the <code>collate</code> directory. Provides the collation sequences available for use with Adaptive Server.
Sybase Software Asset Management (SySAM)	Installed into the <code>SYSAM-2_0</code> directory. Provides asset management of Adaptive Servers and optional features.
Web Services	Installed into the <code>WS-15_0</code> directory. An optional product that provides access to Adaptive Server using HTTP/SOAP and WSDL.
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 <code>DBISQL</code> 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 <code>UA</code> directory <code>UAF-2_0</code> directory.
juutils-3_0 utilities	A collection of Adaptive Server utility programs including ribo , a tool to trace TDS traffic between client program and Adaptive Server.
jConnect™ for JDBC™	Installed into the <code>jConnect-7_0</code> directory. Provides a Java Database Connectivity (JDBC) driver for Adaptive Server.

Note: Sybase strongly recommends you install the ECDA DirectConnect option or MainframeConnect™ DirectConnect™ for z/OS, including DirectConnect Manager, into its own directory.

PC-Client Product Descriptions and Layout

An Adaptive Server installation includes additional products that are installed into specific directories.

Product	Description
Software Developer Kit (SDK)	<p>Installed into the <code>OCS-15_0</code> directory:</p> <ul style="list-style-type: none"> • Open Client™ (Client Library, dblink) • ESQ/L/C • ESQ/L/COBOL • XA • Adaptive Server Enterprise extension module for Python • Adaptive Server Enterprise database driver for Perl • Adaptive Server Enterprise extension module for PHP
DataAccess	<p>Installed into the <code>DataAccess [64]</code> directories:</p> <ul style="list-style-type: none"> • (Windows and Linux only) ODBC – a driver used to connect to Adaptive Server from ODBC-based applications. • (Windows only) OLEDB – a provider used to connect to Adaptive Server from OLE DB-based applications. • (Windows only) ADO.NET – a provider used to connect to Adaptive Server from .NET-based applications.
Shared directory	<p>The <code>shared</code> directory contains components and libraries that are shared by several other components.</p> <ul style="list-style-type: none"> • Sybase Central 6 – is a Java-based framework used by the system management tools. • Java Runtime Environment (JRE) – is a runtime Java virtual machine used to execute Java-based programs, such as Sybase Central.
Adaptive Server plug-in	<p>Installed into <code>ASEP</code> directory. This plug-in contains the management software for managing Adaptive Server and its associated servers.</p>
Language Modules	<p>Installed into the <code>locales</code> directory. Provides system messages and date/time formats.</p>
Character Sets	<p>Installed into the <code>charsets</code> directory. Provides character sets available for use with Adaptive Server.</p>
jConnect	<p>Installed into the <code>jConnect-7_0</code> directory. Provides a Java Database Connectivity (JDBC) driver for Adaptive Server.</p>

Product	Description
Uninstaller	Installed into the <code>sybuninstall/PCClient</code> 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 <code>DBISQL</code> directory.
jutils-3_0 utilities	A collection of Adaptive Server utility programs including ribo , a tool to trace TDS traffic between client program and Adaptive Server.

Performing Administrative Tasks

Complete the administrative tasks before you begin the installation process.

1. Back up your current system.
2. Create the “sybase” user account and make sure it has read, write, and execute permissions.
3. Verify that the directory location for the Sybase installation has sufficient space.
4. 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.

Creating the Sybase User Account

Create a Sybase user account to ensure that Sybase product files and directories are created with consistent ownership and privileges.

One user—typically the Sybase system administrator, who has read, write, and execute privileges—should perform all installation, configuration, and upgrade tasks.

1. To create a Sybase system administrator account, choose an existing account, or create a new account and assign a user ID, group ID, and password for it.

This account is sometimes called the “sybase” user account. See your operating system documentation for instructions on creating a new user account.

If you have installed other Sybase software, the “sybase” user should already exist.

2. Verify that you can log in to the machine using this account.

Preparing to Install Adaptive Server

Prepare your system before starting the installation.

1. For Sun Solaris SPARC 64-bit, before the installation, set the file descriptors limit to a specific value. After the installation is complete, you can set the file descriptors to **unlimited**.
2. Ensure that the account used to start the server has read and write permissions over the directory in which you are installing the server. You may need to create nonadministrative logins to install the server into personal directories.
3. Review the SySAM licensing procedures, and plan your server configuration using the configuration guide for your platform.
4. Create a Sybase account with administrative privileges for consistency and security. This user account could be "sybase" or any other user name. Use this account to perform all installation and device-creation tasks.

This 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. 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.

If you are installing Adaptive Server on more than one computer, create the "sybase" user account on each machine.

5. Log in to the machine as "sybase."

Maintain consistent ownership and privileges for all files and directories. A single user—the Sybase system administrator with read, write, and execute permissions—should perform all installation, upgrade, and setup tasks.

6. If you have enabled the Java in the database feature, create the `sybpcidb` database and disable the feature during the installation. See *Managing Java in the Database During Installations and Upgrades* on page 34.
7. Decide where to install Adaptive Server:
 - Make sure there is sufficient disk space and temporary directory space available.
 - Ensure there are no spaces in the path name of the directory.
8. Decide whether you want licensing events to trigger e-mail alerts, and the severity of the events that generate e-mail messages.

If you choose to have e-mail notifications for license events, know your:

- SMTP server host name
- Port number for an SMTP server

Note: If Sybase assigns any port numbers to you for use with Adaptive Server, exclude these numbers from any port-scanning software you run. Adaptive Server attempts to service each scan as a login attempt, which may lead to poor performance.

- E-mail return address
- Recipients of the notifications
- Severity level of an event that triggers mail. Your choices are:
 - None
 - Informational
 - Warning
 - Error

9. Verify that your network software is configured.

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

If you are having connection problems, or to verify your network configuration, ping the host.

10. Adjust the shared memory for your operating system.

See also

- *Adaptive Server Editions* on page 5

Adjusting the Operating System Shared Memory Parameters

Adjust shared memory segments if Adaptive Server cannot obtain large memory as a single segment, or if Backup Server stripes fail due to insufficient segments.

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

If any reconfiguration through **sp_configure** requires additional memory, Adaptive Server allocates shared memory segments after start-up. To account for these additional segments, allocate the maximum memory available to Adaptive Server by using the **allocate max shared memory** configuration parameter. See the *System Administration Guide*.

1. Adjust the operating system shared memory parameter.

Solaris 10 has shifted to a new method of controlling system resources. For this reason **shminfo_shmmax** is technically obsolete. If you still have an entry for it in `/etc/system` it will work. However, the new method requires an additional entry to the file `/etc/project`. For example:

```
project-sybase:200:For use by Sybase:sybase:sybase:  
project.max-shm-memory=\(privileged,17179869184,deny)
```

The new parameter is **project.max-shm-memory=(privileged,17179869184,deny)** — replaces the old parameter **shminfo-shmmax**. "Privilege" is a threshold value on a

resource control that constitutes a point at which local actions are triggered or global actions such as logging into a machine can occur.

Privilege levels:

- Basic – can be modified by the owner of the calling process.
- Privilege – can be modified by privileged callers.
- System – is fixed for the duration of the operating system instance.

In specifying the threshold value of "privilege," you can use the abbreviation "priv" For example:

```
project.max-shm-memory=(priv, 17179869184,deny)
```

- "17179869184" – is the threshold value (16GB) on the resource control.
- deny – denies attempts to use more than 16GB.

The default value for *project.max-shm-memory* is 25% of the physical memory on the system. The maximum value is UIN64_MAX, which works out to 18446744073709551615 bytes, so is limited only by the size of the physical memory.

To set the *project.max-shm-memory* while the system is running, use the **prctl** command; it sets the configuration permanently.

2. If the installer cannot update for Solaris 10 patch ID 120012-14, edit your */etc/user_attr* file with:

```
sybase:::type=normal;project=project-sybase
-----
```

To confirm that the update to this patch, use either of these commands:

```
# projects -d sybase
project-sybase
```

```
# id -p sybase
uid=204409(sybase) gid=1(other) projid=200(project-sybase)
```

This update allows changes to shared-memory parameters within Solaris 10.

3. Adjust shared memory segments.

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

Adaptive Server may allocate shared memory segments after start-up if any reconfiguration through **sp_configure** requires additional memory. You may need to account for these additional segments. Allocate the maximum memory available to Adaptive Server, by using the **allocate max shared memory** configuration parameter. The default max memory size is 138MB. See the *System Administration Guide* for more information.

To adjust shared memory segments of the operating system, add the following line to the configuration file where *x* is the number of shared memory segments. `/etc/system`:

```
set shmsys:shminfo_shmseg=x
```

Managing Java in the Database During Installations and Upgrades

If you have enabled the Java in the database feature, create the `sybpcidb` database before installing or upgrading to Adaptive Server version 15.7 ESD #2.

1. Create the `sybpcidb` database. The `sybpcidb` database stores configuration information for the pluggable component interface (PCI) and all pluggable component adapter (PCA) components. This database is used by the `installpcidb` script. For example:

```
1> disk init
2> name = "sybpcidb_dev",
3> physname = "${SYBASE}/data/sybpcidb_dev.dat",
4> size = "24M"
5> go
1> create database sybpcidb on sybpcidb_dev = 24
2> go
```

The device and database sizes depend on the Adaptive Server page size:

- 2K page size – 24MB
- 4K page size – 48MB
- 8K page size – 96MB
- 16K page size – 192MB

2. Disable the Java in the database feature:

```
1> sp_configure 'enable java', 0
2> go
```

3. Reenable the feature after you have successfully installed or upgraded to Adaptive Server 15.7 ESD #2:

```
1> sp_configure 'enable java', 1
2> go
```

Multipathing

Multipathing is the ability of a server to communicate with the same physical or logical block storage device across multiple physical paths between the host bus adapters (HBA) in the server and the storage controllers for the device, typically in Fibre Channel (FC) or iSCSI SAN environments. You can also achieve multiple connections with direct attached storage when multiple channels are available.

Multipathing provides connection fault tolerance, failover, redundancy, high availability, load balancing, and increased bandwidth and throughput across the active connections.

Multipathing automatically isolates and identifies device connection failures, and reroutes I/O to alternate connections.

Typical connection problems involve faulty adapters, cables, or controllers. When you configure multipath for a device, the multipathing driver monitors the active connection between devices. Because multipathing is managed at the device level, when the multipath driver detects I/O errors for an active path, it fails over the traffic to the device's designated secondary path. When the preferred path becomes healthy again, control can be returned to the preferred path. Multipathing prevents single points of failure in any highly available system.

A common example of multipath connectivity is a SAN-connected storage device. Usually one or more Fibre Channel HBAs from the host is connected to the fabric switch and the storage controllers are connected to the same switch. A simple example of multipath connectivity is: two HBAs connected to a switch to which the storage controllers are connected. In this case the storage controller can be accessed from either of the HBAs providing multipath connectivity.

All OS platforms provide their own solutions to support multipathing. There are numerous other vendors that also provide multipathing applications for all available platforms, such as:

- AIX – Multiple Path I/O (MPIO)
- HP-UX 11.31 – Native MultiPathing (nMP)
- Linux – Device-Mapper Multipath (DM)
- Solaris – Multiplexed I/O (MPxIO)
- AntemetA Multipathing Software for HP EVA Disk Arrays
- Bull StoreWay Multipath
- NEC PathManager
- EMC PowerPath
- FalconStor IPStor DynaPath
- Fujitsu Siemens MultiPath
- Fujitsu ETERNUS Multipath Driver
- Hitachi HiCommand Dynamic Link Manager (HDLM)
- HP StorageWorks Secure Path
- NCR UNIX MP-RAS EMPATH for EMC Disk Arrays
- NCR UNIX MP-RAS RDAC for Engenio Disk Arrays
- ONStor SDM multipath
- IBM System Storage Multipath Subsystem Device Driver (SDD)
- Accusys PathGuard
- Infortrend EonPath
- OpenVMS
- FreeBSD - GEOM_MULTIPATH and GEOM_FOX modules
- Novell NetWare
- Sun StorEdge Traffic Manager Software
- ATTO Technology multipath driver Fibreutils package for QLogic HBAs
- RDAC package for LSI disk controllers

- Ipfcdriver package for Emulex HBAs
- Veritas Dynamic Multi Pathing (DMP)
- Pillar Data Systems
- Axiom Path
- iQstor MPA

Setting Up Multipathing

If you choose to use it, set up multipathing before you install Adaptive Server.

The Solaris Multiplexed I/O (MPxIO) feature is a multipathing solution for storage devices that is part of the Solaris operating environment. The **stmsboot** program is an administrative command to manage enumeration of multipath-capable devices with Solaris I/O multipathing. Solaris I/O multipathing-enabled devices are enumerated under `scsi_vhci(7D)`, providing multipathing capabilities. Solaris I/O multipathing-disabled devices are enumerated under the physical controller. In the `/dev` and `/devices` trees, Solaris I/O multipathing-enabled devices receive new names that indicate that they are under Solaris I/O multipathing control. This means a device will have a different name from its original name (after enabling) when it is under Solaris I/O multipathing control. The **stmsboot** command automatically updates `/etc/vfstab` and dump configuration to reflect the device names changes when enabling or disabling Solaris I/O multipathing. One reboot is required for changes to take effect.

1. Run **stmsboot -e** to enable multipathing. The **stmsboot -L** option lists all available paths.
2. (Optional) You can also use the **mpathadm** command to manage multipathing.

Install Adaptive Server using your chosen method.

Prerequisites

Complete installation planning tasks.

Task

1. Choose an installation method:
 - GUI wizard (recommended)
 - Console mode
 - Response file
2. Follow the instructions for your chosen method.
3. Perform postinstallation steps.

Mounting the CD

If you are installing using the CD, mount the CD.

The location of the **mount** command is site-specific and may differ from the instructions shown here. If you cannot mount the CD drive using the path shown, check your operating system documentation or contact your system administrator.

The operating system automatically mounts the CD. Log in as "sybase." If CD-reading errors occur, check your operating system kernel to make sure that the ISO 9660 option is turned on. If you have previously installed a Sybase CD on your system, the # sign interferes with the installation process. Before installing your current CD, either:

- Restart your system, or,
- Eject the CD. Delete the *Volume Label* file in `/vol/dsk`, and reinsert the CD.

Installing Adaptive Server with the GUI Wizard

Use GUI mode to install Adaptive Server and related products.

Prerequisites

Shut down all programs before running the installer.

Task

The installer creates the target directory (if it does not yet exist) and installs the selected components into that directory.

Note: Adaptive Server Enterprise 15.7 ESD #2 includes new major versions of Adaptive Server and many supporting components. Installing Adaptive Server version 15.7 ESD #2 into the same directory with existing products should not impact the existing products. However, installing other products on top of Adaptive Server version 15.7 ESD #2 may result in one or more products not working correctly.

Sybase strongly recommends that you install Adaptive Server version 15.7 ESD #2 into its own directory if possible. If you must install other products into the same directory, install Adaptive Server version 15.7 ESD #2 last.

Adaptive Server version 15.5 and later uses InstallAnywhere for its installation, while versions of Adaptive Server earlier than 15.5 and other Sybase products use InstallShield Multiplatform. Do not use both installers to install products into the same directory, as this creates files to be installed incorrectly and overwritten without warning.

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

1. Insert the Adaptive Server media in the appropriate drive, or download and extract the Adaptive Server installation image from the Sybase Product Download Center (SPDC) or the SAP Service Marketplace (SMP).
2. If you downloaded the product from SPDC or SAP Service Marketplace, go to the directory where you extracted the installation image and start the installer:

```
./setup.bin
```

3. If you are installing using the CD or DVD, mount the disk.

The operating system automatically mounts the CD or DVD. Log in as “sybase.” If CD- or DVD-reading errors occur, check your operating system kernel to make sure that the ISO 9660 option is turned on. If you have previously installed a Sybase CD or DVD on your system, the # sign interferes with the installation process. Before installing your current CD or DVD, either:

- Restart your system, or,
- Eject the CD or DVD. Delete the Volume Label file in /vol/dsk, and reinsert the CD or DVD

4. Launch the installer:

```
cd /cdrom/Volume Label  
./setup.bin
```

where:

- *cdrom* is the directory (mount-point) you specified when mounting the CD or DVD drive.

- `setup.bin` is the executable file name to install Adaptive Server.

If there is not enough disk space in the temporary disk space directory, set the environment variable `IATEMPDIR` to `tmp_dir` before running the installer again, where `tmp_dir` is where the installation program writes the temporary installation files. When specifying the `tmp_dir`, include its full path.

5. Select the language.
6. On the Introduction screen, click **Next**.
7. Accept the default directory or enter a new path, then click **Next**.

If you see a **Choose Update Installation** pane, this means you have an earlier version of the server, and you need to perform an upgrade rather than a new installation. See *Chapter 10, Adaptive Server Upgrades* on page 65.

8. Select the type of installation:

Option	Description
Typical	(Default) Installs the default components. This is recommended for most users.
Full	Installs every Adaptive Server component, including all the supported language modules.
Customize	Lets you select the components to install. Some components are automatically installed if they are required to run your selected components.

9. Select the software license type for the Adaptive Server suite:

- **Install licensed copy of Adaptive Server Suite**
- **Install Free Developer Edition of Adaptive Server Suite**
- **Evaluate Adaptive Server Suite**

An evaluation is valid for 30 days from the installation date.

10. Choose the most appropriate region, read the license terms, then click **I agree**. Click **Next**.

11. If you have a licensed copy of Adaptive Server Suite:

- a) Specify the license type:

- Enter a license file. If it is a served license key, you are prompted to set up the license server.
- Use an existing license server.
- Proceed without specifying license information for the duration of 30 days from the installation date.

- b) Specify the product edition:

- **Enterprise Edition**
- **Small Business Edition**
- **Developer Edition**
- **Unknown** – select this if you are unsure about the license type.

Note: Choosing **Unknown** prevents you from being able to select **Optimize Adaptive Server** later in the installation process.

Depending on the edition, you can also select the license type under which the product is licensed.

- c) Configure your server to send e-mail notifications about license management events that require attention:
- SMTP server host name
 - SMTP server port number
 - E-mail return address
 - Recipient e-mail addresses
 - Message severity that triggers e-mail messages

In the summary window, click **Next**.

- 12.** In the preinstallation summary, verify the installation type and that you have enough disk space for the installation. Click **Next**.

The installation status window displays the result of the installation process.

- 13.** Select **enable** or **disable** to remember the Adaptive Server password, then click **Next** to specify whether to have the ASE plug-in remember passwords.

- 14.** If you chose a typical installation or selected the Sybase Control Center remote command and control agent for Adaptive Server in a custom installation, the installer asks if you want to configure SCC. If you choose to configure SCC, you see a message about discovery services, and you can configure the UDP or JINI adapters.

If you select JINI adapter, enter the host name, port number, and heartbeat period.

- 15.** Enable and order the security login module.
- 16.** Enter the RMI port.
- 17.** Choose whether to enable SCC shared-disk mode. If you do, enter the SCC instance name.
- 18.** Set the SCC administrator and agent password that is at least six characters long. This password need not be the same as the Adaptive Server sa login password
- 19.** Review the summary of selections you made for SCC, then click **Next**. The installer configures the SCC agent.

Note: Two commands for starting an SCC agent are:

- To enable the SCC agents:

```
$SYBASE/SCC-3_2/bin/sccinstance -enable
```
 - To start the SCC agents in debug mode:

```
$SYBASE/scc-3_2/bin/scc.sh -m DEBUG ...
```
-

Next

You have installed Adaptive Server, and you can perform some basic settings. See *Minimally Configuring Adaptive Server* on page 41.

See also

- *Managing Adaptive Server Enterprise with Sybase Control Center* on page 6

Installing Adaptive Server in Console Mode

Choose command line installation if you prefer a nonwindowing interface, or for developing custom installation scripts.

Prerequisites

Launch the installer in console mode. If the installer launches automatically, click **Cancel** to cancel the GUI installation, then launch the **setup** program from a terminal or console.

Task

The steps for installing components in an interactive text mode are the same as those described in installing in GUI mode, except you use **setup -i console** to execute the installer from the command line, and you enter text to select installation options.

1. At the command line, enter:

```
setup.bin -i console
```

The installation program starts.

2. The flow of the installation is identical to a GUI installation, except output is written to a terminal window and responses are entered using the keyboard. Follow the remaining prompts to install Adaptive Server, then specify the basic settings for your installation.

Minimally Configuring Adaptive Server

Use GUI mode to minimally configure Adaptive Server and related products for use.

Prerequisites

The Configure screen shows a list of everything you can minimally configure. By default, all products are selected. Unselect any products you do not want to configure now. If you are configuring in console mode, enter the corresponding number:

- 1. **Configure new Adaptive Server**
- 2. **Configure new Backup Server** – see the *System Administration Guide*
- 3. **Configure a new XP Server** – see the *Transact-SQL[®] Users Guide*
- 4. **Configure Job Scheduler** – see the *Job Scheduler Users Guide*
- 5. **Enable Self Management**
- 6. **Configure Web Services** – see *Web Services Users Guide*
- 7. **Configure Unified Agent** – see *Unified Agent and Agent Management Console*

CHAPTER 6: Installing Adaptive Server

When finished, click **Next** in the GUI wizard, or enter **0** in console mode. You are now ready to begin configuring Adaptive Server.

Task

1. Set the **Configure new Adaptive Server** option:

Option	Description
Adaptive Server name	Default is the machine name.
Port number	Default is 5000.
Error log	Name and location of the error log file. Defaults to <code>servername.log</code> .
Application type	The options are numbered only in console mode: <ul style="list-style-type: none">• 1. (Default) MIXED – both OLTP and DSS.• 2. OLTP – online transaction processing generally consists of smaller, less complex transactions.• 3. DSS – decision-support systems generally have less update activity with large complex queries.
Page size	<ul style="list-style-type: none">• 2KB• (Default) 4KB• 8KB• 16KB <p>The default page size selection changed to 4KB in Adaptive Server version 15.5. The default in earlier versions was 2KB. If you plan to load database dumps from an Adaptive Server with a page size that is not 4KB, enter the page size that matches the database dump.</p>
Language	Default: us-english. In typical installations, only us-english is available.
Character set	Default is iso-1.
Sort order	Default is bin-iso-1.

Option	Description
Optimize Adaptive Server configuration. If you select this option, you must also provide this information	<ul style="list-style-type: none"> Available physical memory for Adaptive Server – the value is 80 percent of the sum of combined default physical memory and default OS shared memory. The default is no. Available CPU for Adaptive Server – the value is 80 percent of the physical CPU, with a default minimum of 1. The default is yes. <p>Note: If the value specified is larger than the available resource for allocation to the server, the optimize configuration may fail, causing the server to not start.</p>
Create sample databases	Create or install sample databases. Make sure you calculate the additional space for your master device.

After customizing your Adaptive Server configuration, click **Next** to record the input fields. You see **Enter the custom configuration values:**

Configuration	Value
Master device	Name of master device and the path where it should be located
Master device size	The default options are: <ul style="list-style-type: none"> 2KB page size – 30MB 4KB page size – 73MB 8KB page size – 120MB 16KB page size – 240MB
Master database size	The default options are: <ul style="list-style-type: none"> 2KB page size – 13MB (Default) 4KB page size – 26MB 8KB page size – 52MB 16KB page size – 104MB
System procedure device path	File path name.
System procedure device and database size	The default for both is 172MB.
System device path	File path name.

Configuration	Value
System device size	The default options are: <ul style="list-style-type: none"> • 2KB page size server – 3MB • (Default) 4KB page size server – 6MB • 8KB page size server – 12MB • 16KB page size server – 24MB
System database size	The default options are: <ul style="list-style-type: none"> • 2KB page size server – 3MB • (Default) 4KB page size server – 6MB • 8KB page size server – 12MB • 16KB page size server – 24MB
Temporary database device	Name of device.
Temporary database device size	Default is 100MB for all page sizes.
Temporary database size	Default is 100MB for all page sizes.
Enable PCI in Adaptive Server	If you select this option, provide this additional information: <ul style="list-style-type: none"> • PCI device – default is sybpcidbdev_data.dat in \$SYBASE/data. • PCI device size – default is 48. • PCI database size – PCI device and PCI database size are the same: <ul style="list-style-type: none"> • 2KB page size server – 24MB • (Default) 4KB page size server – 48MB • 8KB page size server – 96MB • 16KB page size server – 192MB

After customizing your Adaptive Server configuration, click **Next** to record the input fields.

Note: The following steps are not necessary to complete an installation in console mode.

2. Set the **Configure new Backup Server** option:

- Backup Server name
- Port number
- Error log

Click **Next**.

3. Set the **Configure new XP Server** option:

- Port number
- Error log

Click **Next**.

4. Set the **Configure Job Scheduler option:**

- Job Scheduler name – the server name of the Job Scheduler. Default is `js_server_name`.
- Agent name – Default is `<ASE Server Name>_JSAGENT`.
- Port number – default is 4900.
- Management device.
- Management device size – the default is 75MB.
- Management database size – the default is 75MB.

Click **Next**.

5. Set the **Configure self management option:**

- Self management user name – default is "sa".
- Self management password – prompted only if you change the default user name.

6. Set the **Configure Web Service option:**

- Producer HTTP Port Number – default is 8181
- Producer HTTPS Port Number – default is 8182
- Producer Host Name
- Certificate Password
- Confirm Certificate Password
- Keystore Password
- Confirm Keystore Password
- Producer Log File
- Consumer Port Number – default is 8184
- Consumer Log File

7. (Optional) Set the **Configure Unified Agent option:**

Unified Agent is a Self-Discovery Service Adapter. If you choose to configure Unified Agent, you see:

- Do you want to configure UDP Adapter?
- Do you want to configure JINI Adapter? If you select this, enter:
 - JINI host name
 - JINI port number – default is 4160
 - JINI heart beat period – default is 900

Note: If you choose No to both, you see a message similar to:

You must choose at least one adapter to configure.

Security Login Modules – choose the security login modules for the Unified Agent, then click **Enable**:

- **Simple Login Module** – default user name is Sybase. To change the user name and password, choose **Modify Simple Login Module Properties**.
 - **ASE Login Module** – enabled by default.
 - **Unix Proxy Login Module** – not enabled by default. Select to enable the module.
8. Verify that the values that appear in the Configure New Servers Summary window are correct, then click **Next**. You see the Configure Server Progress window.
- You may see the restart window, which reminds you to either log out or restart your machine to complete the installation process.
9. In the **Installation Completed** window, click **Done** in GUI mode. In console mode, press **Enter** to exit the installation.

Adaptive Server and related products have been successfully installed, and minimally configured for use. See *Installing Sample Databases* on page 60, to begin experimenting with your servers, or see the *System Administration Guide* for more advanced topics.

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

Installing Adaptive Server Using a Response File

Unattended or silent installations are typically used to update multiple systems across the enterprise.

To perform a silent, or "unattended" installation, run the installer and provide a response file that contains your preferred installation configuration.

Creating a Response File

By saving installation configuration information into a response file during the initial GUI installation, you can perform subsequent installations of Adaptive Server using a silent—or unattended—method.

Note: Adaptive Server version 15.7 ESD #2 uses a rearchitected installation program that is incompatible with response files generated for versions 15.0.3 and earlier. Do not use response files from these earlier versions; instead, create a new one from a 15.7 ESD #2 installation.

To create a response file when installing in GUI or console mode, specify the **-r** command line argument. The **-r** argument records your responses to the installation wizard prompts and creates a response file when the InstallAnywhere wizard exits. The response file is a text file that you can edit to change any responses prior to using it in any subsequent installations. You can find a sample response file at *installer image/sample_response.txt*.

1. Create a response file during GUI installation by specifying the **-r** command line argument, and, optionally, specifying the name of your response file:


```
setup.bin -r response_file_name
```

where:

- **-r** – specifies the command line argument.
- (Optional) *response_file_name* – is the absolute path of the file into which the installation information is stored, such as `/tmp/responsefile.txt`.

Note: The directory path you specify must already exist.

2. Verify that the response file includes the passwords for the Adaptive Server sa login, Sybase Control Center admin, and SCC agent admin by including these lines in the file:

```
SY_CFG_ASE_PASSWORD=<ASE sa password>
CONFIG_SCC_CSI_SCCADMIN_PWD=<SCC admin password>
CONFIG_SCC_CSI_UAFADMIN_PWD=<SCC agent admin password>
```

Each password must be at least six characters long. The sccadmin and uafadmin logins need not be the same as the sa password.

Installing in Silent Mode

To perform a silent, or unattended installation, run the installer and provide a response file that contains your preferred installation configuration.

Prerequisites

Use `setup.bin -r responseFileName` to generate a installation response file during a console or GUI installation.

Task

1. Run the following, where *responseFileName* is the absolute path of the file name containing the installation options you chose:

```
setup.bin -f responseFileName -i silent
-DAGREE_TO_SYBASE_LICENSE=true -DRUN_SILENT=true
```

Note: You must agree to the Sybase License Agreement when installing in silent mode.

You can either:

- Include the option `-DAGREE_TO_SYBASE_LICENSE=true` in the command line argument, or,
 - Edit the response file to include the property `AGREE_TO_SYBASE_LICENSE=true`.
-

Except for the absence of the GUI screens, all actions of InstallAnywhere 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.

2. The installer for Adaptive Server 15.7 ESD #2 requires non-null passwords for the "sa" login in Adaptive Server, and uafadmin and sccadmin logins in Sybase Control Center. For this reason, add these rows to the response file:

CHAPTER 6: Installing Adaptive Server

```
CY_CFG_ASE_PASSWORD=<ASE sa password>  
CONFIG_SCC_CSI_SCCADMIN_PWD=<SCC admin password>  
CONFIG_SCC_CSI+UAFADMIN_PWD=<SCC agent admin password>
```

The passwords for sccadmin and uafadmin logins need not be the same as the password for the Adaptive Server "sa" login.

Command Line Options

Options for installing or uninstalling Adaptive Server in console mode.

Option	Purpose
-i swing	Uses GUI mode.
-i console	Uses the console interface mode, where messages appear on the Java console and the wizard runs in console mode.
-i silent	Installs or uninstalls the product in silent mode, where the installation or uninstallation is performed without user interaction.
-D	Passes custom variables and properties. For example, to override the default installation directory when you run the installer, enter: <i>install_launcher_name</i> -DUSER_INSTALL_DIR=/sybase
-r	Generates a response file and a reference.
-f	References response files.
-l	Sets the installer's locales.
-h?	Displays the installer help.

Uninstalling Adaptive Server

Run the uninstaller to remove Adaptive Server.

Prerequisites

Shut down all servers before uninstalling Adaptive Server.

Task

Note: The uninstallation process reverses only those actions originally performed by the installer, leaving intact any files or registry entries you created after the installation. Remove these after you complete the uninstallation.

1. Execute the uninstallation program:

```
SYBASE/sybuninstall/ASESuite/uninstall
```

2. Click **Next** in the Uninstall Welcome window.
3. Choose one of the following, then click **Next**:
 - Complete Uninstall – uninstalls everything originally installed by the installer.
 - Uninstall Specific Features – displays a list of features. If any unselected products/ components are dependent on features selected for uninstallation, you can still proceed with the uninstallation but dependent features are not uninstalled.

You see a window indicating that the uninstallation process is being performed. There is no progress bar.

Note: You may need to confirm the deletion of some files that have been modified since installation.

4. When you see the final window, click **Done** to exit the uninstallation program.

Removing an Existing Adaptive Server

Remove an existing Adaptive Server.

1. From `$SYBASE`, enter:

```
rm servername.*
```
2. Change to `$SYBASE/$SYBASE_ASE/install` and run:

```
rm RUN_servername.*
rm servername.*
```
3. Edit `$SYBASE/interfaces`, to remove all references to the Adaptive Server.
4. Remove any operating system files for existing database devices.

Installing Components from the PC-Client

The PC-Client CD contains several components that are each packaged with their own installer. When you insert the PC-Client CD, a menu program launches automatically. The menu program lists the components that can be installed from the CD and allows you to install them one at a time. Before installing any products, review the `readme.txt` file, which contains a brief description of each of the products, any product or component dependencies, and last-minute instructions or changes.

The PC-Client CD contains 32-bit and 64-bit SDKs. It installs the 32-bit SDK on 32-bit OS; and the 64-bit SDK on 64-bit OS.

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. Go to **Start > Run** to start the installation program, and enter the following, where *X* is the directory into which you have copied the PC-Client installation files:

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

5. Select **Client Components 15.7**. You see the Welcome window.
6. Select a country, and agree to the licensing terms.
7. Enter the directory path. If you specify a destination that contains an earlier version of the server, you see **Choose Update Installation**, and you can select the features to update.
8. If you are installing a new server, choose the type of installation:

Option	Description
Typical Install	Installs the default components that most users need.
Full Install	Installs every component on the CD.
Custom Install	Allows you to select the components to install. Certain components are automatically installed if they are required to run other selected components. Specify the components to install.

9. The Summary window displays every component to be installed by the installer, the required disk space, and the available disk space.
10. If the target directory does not have enough free space, the available space appears in red. Click **Previous** to return to the previous window and change your selections, or **Cancel** to quit the installer.

Creating a Response File for the Client

By saving installation configuration information into a response file during the initial GUI installation, you can perform subsequent installations of Adaptive Server using a silent—or unattended—method.

Note: Adaptive Server version 15.7 ESD #2 uses a rearchitected installation program that is incompatible with response files generated for versions 15.0.3 and earlier. Do not use response files from these earlier versions; instead, create a new one from a 15.7 ESD #2 installation.

To create a response file when installing in GUI or console mode, specify the `-r` command line argument. The `-r` argument records your responses to the installation wizard prompts and creates a response file when the InstallAnywhere wizard exits. The response file is a text file that you can edit to change any responses prior to using it in any subsequent installations. You can find a sample response file at `installer_image/sample_response.txt`.

Create a response file during GUI installation by specifying the `-r` command line argument, and, optionally, specifying the name of your response file:

```
setup.bin -r response_file_name
```

where:

- `-r` – specifies the command line argument.
- (Optional) `response_file_name` – is the absolute path of the file into which the installation information is stored, such as `/tmp/responsefile.txt`.

Note: The directory path you specify must already exist.

Installing the Client in Unattended, Silent Mode

After running the installer in GUI mode and recording the responses into a response file, edit the file to customize the responses.

To install in unattended, silent mode, run:

```
setupConsole.exe -f responseFileName -i silent
-DAGREE_TO_SYBASE_LICENSE=true
```

where `responseFileName` is the absolute path of the file name containing the installation options you chose.

Agree to the Sybase License Agreement when installing in silent mode by either:

- Including this text in the command line argument –
`DAGREE_TO_SYBASE_LICENSE=true`, or,
- Editing the response file to include the property
`AGREE_TO_SYBASE_LICENSE=true`.

The only difference when you install the client in silent mode is the absence of the GUI screens. All actions and responses are the same as using InstallAnywhere.

Warning! Sybase recommends that, when you run a silent installation, you use the `setupConsole.exe` executable, which runs in the foreground. The normal `setup.exe` executable runs in the background, giving you the impression that the installation has terminated, and results in additional installation attempts using the silent installation. Multiple installations at the same time may corrupt the Windows Registry and lead to a failure to restart the operating system.

Uninstalling PC-Client

Choose one of two methods to remove PC-Client from a Windows machine.

Actions performed by the installer are removed during an uninstallation process. Files, or registry entries created after installation are not removed and must be manually removed when the installation process has finished.

To uninstall PC-Client from a Windows machine:

- Run the **uninstall** executable at:
`%SYBASE%\sybuninstall\PCClient\uninstall.exe`
- Use **Control Panel > Add or Remove Programs**.

Starting and Stopping Sybase Control Center Agent

Sybase Control Center (SCC) for Adaptive Server is a Web-based tool for monitoring the status and availability of Adaptive Servers. Set the \$SYBASE environment variables and start an SCC agent on your Adaptive Server.

There are several ways to run SCC:

- From the command line in the foreground
- From the command line in the background
- Set up a daemon to run it as a service

For instructions, see **Get Started > Launching Sybase Control Center > Starting and Stopping Sybase Control Center** in the SCC online help.

1. Source either the `SYBASE.csh` or `SYBASE.sh` environment script.
2. Start the SCC agent:

```
$SYBASE/SCC-3_2/bin/scc.sh
```

The SCC agent creates and sends output to one of these log files:

- (Default) If SCC shared-disk mode is not enabled: `$SYBASE/SCC-3_2/log/agent.log`
 - If SCC shared-disk mode is enabled: `$SYBASE/SCC-3_2/instances/<hostname>/log/agent.log`
3. Verify that the SCC agent is running. After you run the SCC script, the SCC console prompt appears. At this prompt, enter:

```
scc-console> status
```

You should see a status message similar to:

```
Agent Home: /remote/perf_archive/olwen/Install_Testing/157CE_C3/
SCC-3_2/instances/solstrs3
Connection URL: service:jmx:rmi:///jndi/rmi://solstrs3:9999/agent
Status: RUNNING
```

From the SCC console, you can also shut down the agent:

```
scc-console> shutdown
```


After you install the server, set it up.

The Adaptive Server installation includes directories containing sample client library applications. These sample programs are included for training purposes only and are not intended to be installed in a production environment.

If you are setting up a production environment, remove these directories:

- `$SYBASE/OCS-15_0/sample`
- `$SYBASE/DataAccess/ODBC/samples`
- `$SYBASE/jConnect-7_0/sample2`
- `$SYBASE/jConnect-7_0/classes/sample2`
- `$SYBASE/ASE-15_0/sample`
- `$SYBASE/WS-15_0/samples`

See also

- *Installing Sample Databases* on page 60

Verifying That Servers are Running

Verify the servers are running.

Prerequisites

Before starting the server, make sure you have stopped, then started the services related to the server.

Task

If you have already started the server, do not run the command a second time. Running it more than once can cause problems.

1. From a UNIX command line, set the Sybase environment variables.

- In C-shell:

```
source ASE_install_location/SYBASE.csh
```

- In Bourne shell:

```
ASE_install_location/SYBASE.sh
```

2. Show all Adaptive Server-related processes that are on the system:

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

Verifying Your Connection to the Servers

Verify your connection to the servers.

Use **isql** to perform a quick test:

- At the command prompt, enter the following, where *server_name* is the Adaptive Server name:

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

You see the **isql** prompt if the login is successful.

- At the **isql** prompt, enter:

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

The output should show Adaptive Server at version 15.7 ESD #2.

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

See also

- *Setting the System Administrator Password* on page 60

Testing the Installation and Network Connections

After installing Adaptive Server, Sybase Central, and the Java Runtime Environment, test the installation and network connections. 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. Sybase stores this information in the `interfaces` file on Windows or in a Lightweight Directory Access Protocol (LDAP) server.

1. Select **Start > Programs > Sybase > Sybase Central 6.0.0**.
2. Select **Tools > Connect** from the Sybase Central menu bar to activate an Adaptive Server login window.
3. Log in using the default user ID, "sa," with a password (the password cannot be null). If you have changed the password, use the new password.
4. Select the Adaptive Server to which you want to connect.
5. To disconnect from a server, select **Tools > Disconnect**, or right-click the connected server icon and select **Disconnect**.
6. Exit Sybase Central.

Configuring libtcl.cfg for LDAP

Use the `libtcl.cfg` file to specify an LDAP server name, port number, DIT base, user name, and password to connect to an LDAP server.

When you specify an LDAP server in the `libtcl.cfg` file, the server information becomes accessible only from the LDAP server; the `interfaces` file is ignored. Open Client and Open Server applications that use the `-l` option at start-up override the `libtcl.cfg` file and use the `interfaces` file. See the *Adaptive Server Configuration Guide*.

1. Configure the `libtcl.cfg` file to use directory services, using any standard ASCII text editor:
 - 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.

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

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

Warning! Ensure that the LDAP URL is on a single line.

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

For example (entry shown on multiple lines only for readability):

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

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

2. Verify that the appropriate environment variable points to the required third-party libraries. The Netscape LDAP SDK libraries are in `%SYBASE%\%SYBASE_OCS%\dll`.

The Windows `PATH` environment variable must include this directory.

Adding a Server to the Directory Service

Use `dsedit` utility to add a server to the directory service.

1. From Windows, select **Start > Programs > Sybase > Connectivity > Open Client Directory Service Editor**.
2. Select **LDAP** from the list of servers, and click **OK**.

3. Click **Add New Server Entry**, and enter:
 - Server name
 - Security mechanism – (Optional) there is a list of security mechanism OIDs in %SYBASE%\ini\objectid.dat.
4. Click **Add New Network Transport**, and:
 - Select the transport type.
 - Enter the host name.
 - Enter the port number.
5. Click **OK** twice to exit the **dsedit** utility.

Setting the System Administrator Password

When you install Sybase software, Sybase creates a system administrator account called "sa," which can use any database on Adaptive Server, including *master*, with full privileges.

Immediately after a new installation, Sybase assigns a default password to "sa." For security purposes, make sure you must assign a password to "sa" before using Adaptive Server in a production environment, as it is not possible to have the "sa" password be NULL.

The Sybase system administrator should log in to Adaptive Server as "sa" and set a password:

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

where:

- **default** – is the non-null password.
- *new_password* – is the password that you are assigning to the "sa" account.

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

Installing Sample Databases

The sample databases, which contain fictitious information, are designed to help you learn how to use Adaptive Server.

Note: The sample databases are intended for training purposes only. Do not install them in an Adaptive Server production environment.

Database	Description
<code>installpubs2</code>	<p>Installs the <code>pubs2</code> 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 <code>pubs2</code> database.</p> <hr/> <p>Note: The <code>master</code> device size should be at least 30MB to install the full <code>pubs2</code> database, including the <code>image</code> data.</p>
<code>installpubs3</code>	<p>Installs the <code>pubs3</code> sample database. This updated version of <code>pubs2</code> uses referential integrity. In addition, its tables are slightly different than the tables used in <code>pubs2</code>. Where noted, the Adaptive Server documentation uses the <code>pubs3</code> database in its examples.</p>
<code>installpix2</code>	<p>Installs the <code>image</code> data that is used with the <code>pubs2</code> database.</p> <p>Run the <code>installpix2</code> script after you run <code>installpubs2</code>.</p> <p>The <code>image</code> data requires 10MB—there are six pictures, two each in the PICT, TIFF, and Sun raster file formats. Run <code>installpix2</code> script only to use or test the <code>image</code> datatype. Sybase does not supply any tools for displaying <code>image</code> data; you must use appropriate window graphics tools to display the images after you have extracted them from the database.</p>

Default Devices for Sample Databases

The Adaptive Server installation includes scripts for installing the `us_english` sample database, foreign language sample databases, and the image data associated with the US English `pubs2` sample database on a default device.

These scripts are located in `$(SYBASE)/$(SYBASE_ASE)/scripts`.

By default, these scripts install the sample databases on the `master` device. Because the databases use valuable space on the `master` device that is best reserved 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—Sybase recommends that you change the default to a device other than the `master` device.

To change the default location where the scripts install these databases, use `sp_diskdefault`. See `sp_diskdefault` in *Reference Manual: Procedures*. Alternatively, use a text editor to revise the script directly.

Running the Database Scripts

After you determine your default device, run the scripts to install the sample databases.

Prerequisites

Back up the original `installpubs2` and `installpubs3` scripts, in case you have problems with the edited scripts.

Task

1. Start Adaptive Server.
2. Go to the Adaptive Server scripts directory at `$SYBASE/$SYBASE_ASE/scripts`.
3. Use `isql` to log in to Adaptive Server and run the script:

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

where:

- `server_name` – is 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
```

4. Install the image data associated with `pubs2`:

```
isql -Usa -Ppassword -Sservername  
-i$SYBASE/$SYBASE_ASE/scripts/installpix2
```

The `pubs3` database does not use image data.

For more information about running these scripts, see the *Adaptive Server Configuration Guide*.

Installing the Interpubs Database

The `interpubs` database is similar to `pubs2`, and contains French and German data.

Prerequisites

Back up the original `installintpubs` script, in case you experience problems with the edited script.

Task

1. Set your terminal to display 8-bit characters.
2. Verify that `iso_1`, `iso_15`, `Roman8`, `Roman9`, or `UTF-8` is installed as either the default character set or as an additional character set.

The `interpubs` database contains 8-bit characters and is available for use at Adaptive Server installations using the ISO 8859-1 (`iso_1`), ISO 8859-15 (`iso_15`), Roman 8, or Roman9 (for HP-UX) 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 `interpubs` database. You will need to provide this information later.
4. 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
```

Installing the jpubs Database

If you installed the Japanese language module with your server, you can run the `installjpubs` script to install `jpubs`, which 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.

Prerequisites

Copy the original `installjpubs` script, in case you experience problems with the edited script.

Task

1. Set your terminal to display 8-bit characters.
2. Verify that the EUC-JIS, Shift-JIS, or the UTF-8 character set is installed either 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. 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
```

```
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
```

```
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 Sample Databases

The sample databases contain a guest user option that allows authorized users to access the database. Guest users have a wide range of privileges, including permissions to **select**, **insert**, **update**, and **delete** user tables.

Sybase recommends that you remove the guest user option from user databases in production systems. For more information about the guest user and a list of guest permissions, see the *System Administration Guide*.

1. If possible, and if space allows, give each new user a clean copy of the sample databases so that each user is not confused by other users' changes.
2. If space is a problem, instruct users to issue the **begin transaction** command before updating a sample database.
3. After users have finished updating one of the sample databases, instruct them to issue the **rollback transaction** command to undo the changes.

CHAPTER 10 **Adaptive Server Upgrades**

You can upgrade to Adaptive Server 15.7 from Adaptive Server versions 15.0 through 15.0.3, with the exception of versions 15.0.1 and 15.0.3 of the Cluster Edition.

You can upgrade to this version of Adaptive Server from:

- Adaptive Server 15.0 through 15.0.3 (but not the 15.0.1 or 15.0.3 of the Cluster Edition) – to install version 15.7 on top of an earlier 15.x version, see *Installing Adaptive Server Using Binary Overlay* on page 83.
- Adaptive Server 12.5.x – to upgrade from Adaptive Server 12.5.4 or earlier, execute the upgrade sequence from the same node on which the database resides.

You can upgrade from a 32-bit version to a 64-bit version, but you can neither upgrade nor move from a 64-bit to a 32-bit:

You cannot upgrade to this version of Adaptive Server from:

- Adaptive Server version 12.5.3a
- Adaptive Server version 12.0.x or earlier – Sybase recommends that you first upgrade to version 12.5.4 before upgrading to version 15.7

If your server has replicated databases, consult the *Replication Server Configuration Guide* before starting the preupgrade tasks.

Only upgrades to and from the same page size are supported. Use **sybmigrate** to re-create the schema and load data from one page size to another. See the *Utility Guide*.

Adaptive Server version 15.7 ESD #2 includes both new and changed system catalogs that may affect your existing applications. See the *Adaptive Server Enterprise New Feature Guide* for a complete list. Sybase recommends that you do not use the new features until you are committed to staying with 15.7.

Upgrading Adaptive Server

A minor change was made to the format of a log record in both the cluster and nonclustered editions of Adaptive Server 15.5 and later versions.

This change introduces an unlikely possibility for Adaptive Server to misinterpret this changed log record in cases where the upgraded server contains a database that is the primary database for replication.

This change should not affect your upgrade process; however, it requires you to strictly follow some steps when you move from Adaptive Server 15.0.x or earlier to Adaptive Server 15.5.x

and later (nonclustered editions). See the following table for all possible combinations of upgrades.

An upgrade is done by:

- Upgrading your entire installation by switching binaries.
- Upgrading a single database using **online database**, having first loaded it with a database dump and transaction logs taken on a server with an earlier version.

Table 4. Upgrading Your Entire Installation

Current Version	Upgrading to	Special Upgrade Information
Adaptive Server 15.0.x or earlier	Adaptive Server 15.7.x	If you are using Replication Server to replicate one or more databases in the version to be upgraded, ensure that all transactions have been replicated before the polite shutdown is done, by draining the log. See <i>Upgrading Adaptive Server in a Replication System</i> and <i>Upgrading Replication Server</i> in the <i>Replication Server Configuration Guide</i> for your platform.
Adaptive Server 15.0.x	Adaptive Server Cluster Edition 15.7.x	If you are using Replication Server to replicate one or more databases in the installation to be upgraded, ensure that all transactions have been replicated before the polite shutdown is done, by draining the log. See the <i>Upgrading Adaptive Server in a Replication System</i> and <i>Upgrading Replication Server</i> in the <i>Replication Server Configuration Guide</i> for your platform.
Adaptive Server 15.5.x or later	Adaptive Server Cluster Edition 15.7.x	Not supported.
Adaptive Server Cluster Edition 15.5.x or later	Adaptive Server 15.7.x	Upgrading any version of Adaptive Server Cluster Edition to a nonclustered version is not supported.
Adaptive Server 15.5.x	Adaptive Server 15.7.x	No special upgrade instructions.

Current Version	Upgrading to	Special Upgrade Information
Adaptive Server Cluster Edition 15.5.x or later	Adaptive Server Cluster Edition 15.7.x	<p>In the installer destination directory pane, enter the 15.5.X directory, then choose Update.</p> <p>After your installation, install the new system procedures (from <code>\$SYBASE/ASE-15_0/scripts/installmaster</code>).</p> <p>See <i>"Installation of Version 15.7 ESD #2 over Existing 15.x Adaptive Server"</i> on page 82 for details.</p>

Table 5. Upgrading a Single Database

Current Version	Upgrading to	Special Upgrade Instructions
Adaptive Server 15.0.x and earlier	Adaptive Server 15.7.x	When you use online database to upgrade a single database in Adaptive Server 15.7.x (either the cluster or noncluster edition) after loading a database dump and transaction logs from Adaptive Server 15.0.x or earlier, if the database you are upgrading is also replicated, make sure the database's transaction log is truncated before you turn replication back on.
Adaptive Server 15.5.x or later	Adaptive Server Cluster Edition 15.7.x	Upgrading a single database from the nonclustered edition of 15.5 or 15.5 ESD #1 to cluster edition of 15.5 or 15.5 ESD #1 is supported, and no additional steps are required.
Adaptive Server Cluster Edition 15.5.x or later	Adaptive Server 15.7.x	Upgrading a database from any version of the Cluster Edition to a nonclustered version is not supported.
Adaptive Server 15.5.x	Adaptive Server 15.7.x	No special upgrade instructions.
Adaptive Server Cluster Edition 15.5.x	Adaptive Server Cluster Edition 15.7.x	No special upgrade instructions.

See also

- *Reenabling Replication Server After the Dataserver Upgrade* on page 87
- *Postupgrade Tasks* on page 84

- *Installation of Version 15.7 ESD #2 over an Existing 15.x Adaptive Server* on page 82

Considerations for Component Integration Services

If you have both a local and a remote server running on Adaptive Server version 15.x and are upgrading both of them to version 15.7 ESD #2, upgrade the local server first. If you plan to upgrade one server but not the other, upgrade the local server.

Sybase does not certify that Component Integration Services running on earlier versions of Adaptive Server can connect to later versions. You may encounter errors if an earlier version of Adaptive Server maps a proxy table to a later version and the remote table uses functionality that is unavailable in the earlier version.

Sybase certifies that each version of Adaptive Server connects to earlier versions through Component Integration Services. Component Integration Services is tested and certified to connect to earlier versions of Adaptive Server.

Preparing to Upgrade

Before you upgrade, run the **preupgrade** utility. You must have system administrator privileges to perform an upgrade.

If your server is at a 15.x level, do not use the **sqlupgrade** or **sqlupgraderes** upgrade utilities.

Each new server version contains features that introduce parameters, commands, reserved words, and so on. **preupgrade** prepares the old server by making sure all necessary directories and settings are correct for the upgrade. When running **preupgrade**, manually stop and start the server. You need not start the server before running the **sqlupgrade** upgrade utility; the utility starts the server when needed.

- If you are upgrading from:
 - Version 12.5.3 with cache bindings on **sybssystemdb** – drop the cache bindings for **sybssystemdb** that are bound to a user-defined cache before you run **preupgrade**.
 - Earlier than version 15.x but later than 12.5.3 – from the Adaptive Server 15.7 ESD #2 installation directory, perform the preupgrade checks on the old server using the **preupgrade** utility, located in `$(SYBASE)/ASE-15_0/upgrade`.
 - Manually drop objects before you execute procedures for the first time. The first time you run the server after an upgrade, the server internally rebuilds procedures from the text in `syscomments`. The procedure may not execute properly if it contains code that drops, then re-creates, an existing object.
1. Go to the directory where your current version of Adaptive Server is installed.
 2. Enter **cd in** to go to the directory `ASE-version`, where `version` is your current version of Adaptive Server.

3. Perform **cd in** to the upgrade directory.
4. Enter **preupgrade** to run the **preupgrade** utility.

Changes in Adaptive Server Directories

The directory structures for Adaptive Server installations differ depending on version.

Table 6. Directory Changes for UNIX Platforms

Component	12.5.4 Location	15.0.2 Location	15.0.3 Location	15.5, 15.7, and 15.7 ESD #1 and ESD #2 Locations
Adaptive Server	\$SYBASE/ ASE-12_5	\$SYBASE/ ASE-15_0	\$SYBASE/ ASE-15_0	\$SYBASE/ ASE-15_0
Shared directory	\$SYBASE/ shared	\$SYBASE/ shared	\$SYBASE/ shared	\$SYBASE/ shared
Sybase Central	\$SYBASE/ shared/ sybcentral4 3	\$SYBASE/ shared/ sybcentral4 3	\$SYBASE/ shared/ sybcentral6 00	\$SYBASE/ shared/ sybcentral6 00
JRE	\$SYBASE/ shared/ jre142	\$SYBASE/ shared/ jre142_*	\$SYBASE/ shared/ JRE-6_0*	\$SYBASE/ shared/ JRE-6_0*
Shared JAR file	\$SYBASE/ shared/lib	\$SYBASE/ shared/lib	\$SYBASE/ shared/lib	\$SYBASE/ shared/lib
Language settings (locales)	\$SYBASE/ locales	\$SYBASE/ locales	\$SYBASE/ locales and \$SYBASE/ ASE-15_0/ locales	\$SYBASE/ locales and \$SYBASE/ ASE-15_0/ locales
Connectivity	\$SYBASE/ OCS-12_5	\$SYBASE/ OCS-15_0	\$SYBASE/ OCS-15_0	\$SYBASE/ OCS-15_0
Web Service	\$SYBASE/ WS-12_5	\$SYBASE/ WS-15_0	\$SYBASE/ WS-15_0	\$SYBASE/ WS-15_0
Replicator	\$SYBASE/ RPL-12_5	\$SYBASE/ RPL-15_0	\$SYBASE/ RPL-15_0	

Component	12.5.4 Location	15.0.2 Location	15.0.3 Location	15.5, 15.7, and 15.7 ESD #1 and ESD #2 Locations
SySAM	\$SYBASE/ SYSAM-1_0	\$SYBASE/ SYSAM-2_0	\$SYBASE/ SYSAM-2_0	\$SYBASE/ SYSAM-2_0
Job Scheduler	\$SYBASE/ JS-12_5	\$SYBASE/ ASE-15_0/ jobschedule r	\$SYBASE/ ASE-15_0/ jobschedule r	\$SYBASE/ ASE-15_0/ jobschedule r
Unified Agent		\$SYBASE/ UAF-2_0	\$SYBASE/ UAF-2_0	\$SYBASE/ UAF-2_5 In Adaptive Server 15.7 ESD #1 and later, UAF-2_5 is placed with SCC-3_2 for the Sybase Control Center.

Performing Preupgrade Tasks

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

Prerequisites

- You need the master and system procedures files to upgrade the server. By default, the master (master.dat) and system procedures device (sybsystemprocs.dat) files are installed in the \$SYBASE/data directory.
- The previously installed version of the server must be running. If you are upgrading Backup Server, Historical Server, or XP Server, those servers must not be running.

Task

1. Install Adaptive Server 15.7 ESD #2 into its own installation directory.
2. Check system and upgrade requirements.
3. Ensure that the directory to which your TEMP environment variable points exists. The installation program uses this directory to write files temporarily during the installation, which it locates through the TEMP environment variable.

4. Verify the name and location of the `runserver` file, as well as that you renamed it to `RUN_servername`, where *servername* is the name of the old server, as it appears in the `interfaces` file.

The default `RUN_servername` file for a server SYBASE is called `RUN_SYBASE`. If the `RUN_servername` file for your current server is named differently, use the actual name of your server during the upgrade process.

5. Verify that the text of all stored procedures you upgrade is available in *syscomments* by either:
 - Reinstalling the procedures with text, or,
 - Removing and reinstalling the procedures after the upgrade.

This step can allow you to inspect the stored procedures for any hidden or unwanted text.

6. Set the '**allow password downgrade**' password policy option to 1 to ensure that passwords are encrypted with both the old and new algorithms, in case you need to downgrade to an earlier version of Adaptive Server.
7. Make sure reserved words use quoted identifiers.
8. Verify that users are logged off.
9. Use **dbcc** to check for database integrity.
10. Back up databases.
11. Dump the transaction log.
12. Verify that `master` is the default database for the "sa" user.
13. Prepare the database and devices for upgrade using the **preupgrade** utility:
 - a) Create a `sybssystemdb` database if it does not already exist.
 - b) Disable auditing by running **sp_configure 'auditing', 0**.
 - c) Disable Job Scheduler.
 - d) Obtain the current audit settings for the pre-15.7 Adaptive Server using **sp_displayaudit**; you use this saved information to reenabling auditing after you complete the installation. See *Reenabling Auditing* on page 86.
 - e) Disable disk mirroring.
 - f) Verify that your SYBASE environment variable points to the location of the new server software files you just installed.

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**.

Fix any problems that **preupgrade** utility reports.

You can safely ignore any warnings Adaptive Server issues about configuration parameters not being set to the default, as these are only for your information.

14. Ensure that the procedure cache size is at least 150 percent of the default procedure cache size, or between 53,248 and 2,147,483,647 2K pages.
15. Copy these files from earlier versions of the server to their corresponding Adaptive Server 15.x installation locations:
 - `$SYBASE/interfaces`
 - `$SYBASE/$SYBASE_ASE/servername.cfg` – where *servername* is your server name.
 - `$SYBASE/$SYBASE_OCS/config/libtcl.cfg`
 - `$SYBASE/SYSAM-2_0/licenses/license.lic`
16. If you have enabled the Java in the database feature, create the `sybpcidb` database and disable the feature during the installation.
17. Set the environment variables by sourcing the `SYBASE.[csh, sh, env]` file.

Warning! Do not run the environment variable script more than once.

18. Set the `OLDSYBASE_ASE` variable to the `SYBASE_ASE` appropriate to your older server. If you are upgrading from 12.5 for example, it should be `ASE-12_5`.

If upgrading from Adaptive Server 12.5, and you are using **sqlupgraderes** to perform the upgrade, enter:

```
setenv OLDSYBASE <old_$SYBASE_directory>
setenv OLDSYBASE_ASE ASE-12_5
```

Sybase recommends you also set the `OLDSYBASE`, and `OLDSYBASE_OCS` environment variables to point to the old installation of Adaptive Server, so the upgrade utility automatically populates input fields with the correct information as you upgrade.

See also

- *Managing Java in the Database During Installations and Upgrades* on page 34

Upgrade System Tables and Stored Procedures

When you upgrade Adaptive Server, drop and re-create `syscomments` to include the new and changed tables.

If you deleted text from the **syscomments** system table, drop and re-create the stored procedures to again add the text. Sybase recommends that you use the **sp_hidetext** stored procedure to hide text, rather than deleting text.

If you modified any system stored procedures without altering their names, back them up before you upgrade Adaptive Server. The upgrade process overwrites these modified procedures with default versions.

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, you see errors when you run queries, stored procedures, or applications that use these identifiers in your user databases.

Note: Before upgrading, use **sp_renamedb** to change the name of any user database that is a reserved word.

If you change an object name, also change applications and stored procedures that refer to that object. 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 preupgrade process, you can have **sqlupgrade**, **sqlupgraderes**, or **preupgrade** perform a reserved-word check for you. See the *Reference Manual* for a complete list of reserved words.

Running a Reserved Word Check

Run the reserved word check on the old Adaptive Server.

Even if you plan to use **sqlupgraderes** to upgrade Adaptive Server noninteractively, you can still first run the interactive **sqlupgrade** utility to perform reserved-word checks and other potential upgrade eligibility problems, and then continue with the upgrade if no problems are found.

Both **sqlupgrade** and **sqlupgraderes** automatically install new reserved words and the **sp_checkreswords** system procedure, which detects and displays identifiers in your existing databases that conflict with reserved words in the new database. You can use **sp_checkreswords** at any time while performing preupgrade tasks.

Note: Do not use an older version of **sp_checkreswords** to check for reserved words, as it does not contain the most current list of reserved words.

1. 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 to change.
2. The Server Config utility installs **sp_checkreswords** and checks the reserved words during the upgrade. Server Config does not display the names or locations of identifiers that conflict with reserved words; only the number of conflicts.

Addressing Reserved Words Conflicts

Change any database names that are reserved words.

1. Use **sp_dboption** to set the database to single-user mode, then run **sp_renamedb**, specifying the new name.
2. 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 )
```
3. Run **sp_checkreswords** in *master* and in each user database to display the names and locations of conflicting identifiers.

See *Reference Manual: Procedures* for information about **sp_dboption**, **sp_rename**, and **sp_checkreswords**.

Quoted Identifiers

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.

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

See *Reference Manual: Commands* for more information on **set quoted_identifier**.

Preparing the Database and Devices

Before starting the upgrade process, confirm that you have enough space for the upgrade to succeed.

1. Double the size of the largest catalog that you are going to upgrade, then add approximately 10 percent to calculate the free space required for the upgrade to succeed. For example, if you have a large number of stored procedures or compiled objects in any database, the **syscomments** and **sysprocedures** system tables require additional space. Sybase recommends that you use the **preupgrade** utility, since it scans the system catalogs and performs this calculation automatically.
2. If you are:
 - Not using two-phase commit or distributed transaction management (DTM) applications – create a **sybsystemdb** with a minimum of 4MB if it does not already exist.
 - Using two-phase commit or DTM, either:

- Create a minimum-sized `sybsystemdb` for upgrade and expand it later, or
 - Create a `sybsystemdb` 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.
3. If it does not already exist, create the `sybsystemdb` database. Servers that use a 2K logical page size require a `sybsystemdb` that is at least 4MB. Servers that use a larger logical page size require a `sybsystemdb` that is at least as large as the `model` database:

```
1> create database sybsystemdb on default = '4M'
2> go
1> use sybsystemdb
2> go
1> sp_create_syscoordinations
2> go
```

Adaptive Server uses this database for tracking transactions and during recovery. It is also used by applications using two-phase commit and DTM.

4. Verify that the `sybsystemdb` database is large enough. For an upgrade, the default size for **`sybsystemdb`** is 124MB, or enough free space to accommodate the existing `sybsystemdb` database, and the largest catalog that is to be upgraded, plus 10 percent for logging upgrade changes. You may need more space if you are adding user-defined stored procedures.
5. If you use Java in the database, the **`preupgrade`** utility detects that you set **`sp_configure "enable java"`** to 1, and requests that you enable PCI and configure **`sybpcidb`** before you upgrade the server.
6. Increase the default database size:

- a) Use **`alter database`** to increase the size of the master database.

For example:

```
1> alter database tempdb on master=x
2> go
1> alter database model on master=x
2> go
```

where *x* is the megabytes of space you are adding.

Repeat this process for each temporary and `model` databases making sure that `model` is at no point larger than `tempdb`.

- b) Use **`sp_helpdb`** to verify the size of each system database you increase in size.
- c) Use **`sp_configure`** to update the values as indicated by the **`preupgrade`** utility. This example updates the number of locks available for all users on Adaptive Server to 6,000:

```
sp_configure "number of locks", 6000
```

The sybprocsdev device

Sybase system procedures are stored in the `sybssystemprocs` database, which is stored in the `sysprocsdev` device. You may need to increase the size of `sysprocsdev` before upgrading Adaptive Server.

The minimum/default `sybssystemprocs` size is 172MB for all page sizes when configuring a new server. For upgrades, you need an additional 10 percent of this.

You may need more space if you are adding user-defined stored procedures.

If your `sybssystemprocs` 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 `sybssystemprocs` database:

```
1> sp_helpdb sybssystemprocs
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, increase the size of `sysprocdev`.

Increasing the Size of the sybssystemprocs Database

Create a new database with sufficient space if your current `sybssystemprocs` database does not have the minimum space required for an upgrade.

Prerequisites

If you do not have a current backup of your old database, create one now.

Task

Although you can drop the old database and device and create a new `sysprocsdev` device, Sybase recommends that you leave the old database and device alone and add a new device large enough to hold the additional memory, and alter the `sybssystemprocs` onto the new device.

1. In **isql**, use **alter database** to increase the size of the `sybssystemprocs` database. For example:

```
1> use master
2> go
1> alter database sybssystemprocs on sysprocsdev=40
2> go
```

In this example, "sysprocsdev" is the logical name of the existing system procedures device, and 40 is the number of megabytes of space to add. If the system procedures device

is too small, you may receive a message when you try to increase the size of the `sybsystemprocs` database.

If there is space available on another device, expand `sybsystemprocs` to a second device, or initialize another device that is large enough.

2. Verify that Adaptive Server has allocated more space to `sybsystemprocs`:

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

When the database is large enough to accommodate the increased size of `sybsystemprocs`, continue with the other preupgrade 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.

Prerequisites

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*.

Task

This procedure involves dropping the database. For more information on **drop database**, see the *Reference Manual*.

1. Determine which device you must remove:

```
select d.name, d.phyname
from sysdevices d, sysusages u
where u.vstart between d.low and d.high
and u.dbid = db_id("sybsystemprocs")
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 devices that are used by `sybsystemprocs` and other databases.

Make a note of the names of the devices to use in subsequent steps.

Warning! Do not remove any device that is in use by a database other than `sybsystemprocs`, or you will destroy that database.

2. Drop `sybsystemprocs`:

```

1> use master
2> go
1> drop database sybssystemprocs
2> go

```

Note: In versions of Adaptive Server Enterprise earlier than 15.x, use `sysdevices` to determine which device has a low through high virtual page range that includes the `vstart` from step 2.

In version 15.x, select the `vdevno` from `sysusages` matching the `dbid` retrieved in step 1.

3. Remove the device:

```

1> sp_configure "allow updates", 1
2> go
1> delete sysdevices
    where name in ("devname1", "devname2", ...)
2> go
1> sp_configure "allow updates", 0
2> 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 that were returned in the `d.phyname` list.

Note: File names cannot be complete path names. If you use 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 = 200M
5> go

```

Note: Server versions 12.0.x and later accept, but do not require "vdevno=number". 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: Commands*.

6. Create a `sybsystemprocs` database of the appropriate size on that device, for example:

```
1> create database sybsystemprocs on sysprocsdev = 180
2> go
```

7. Run the **installmaster** script in the old server installation directory. Enter:

```
isql -Usa -Ppassword -Sserver_name -i$SYBASE/ASE-15_0/scripts/
installmaster
```

Upgrade to Adaptive Server 15.7 ESD #2

After you successfully run the **preupgrade** utility, you are ready to upgrade Adaptive Server.

Upgrading Interactively Using `sqlupgrade`

Use the interactive **sqlupgrade** tool to upgrade Adaptive Server using an X-Windows or Motif GUI.

Prerequisites

Set the `OLDSYBASE_ASE` variable to the `SYBASE_ASE` appropriate to your older server. If you are upgrading from 12.5 for example, `OLDSYBASE_ASE` should be `ASE-12_5`.

Sybase recommends you also set the `OLDSYBASE`, and `OLDSYBASE_OCS` environment variables to point to the old installation of Adaptive Server, so the upgrade utility automatically populates input fields with the correct information as you upgrade.

Source `SYBASE.csh` before running **sqlupgrade** if you have not already so.

Make sure your old server is running. If the server you are upgrading is not running, **sqlupgrade** prompts you to start the server.

Task

1. Enter:

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

2. Verify that you have backed up databases, devices, and so on, and click **Next**.
3. Click **OK**.
4. Enter the name of the old Sybase directory and the old Adaptive Server directory, then click **OK**.
5. From the list of server names, select the server to upgrade, and click **OK**.
6. Enter the password for the "sa" login, then click **OK**.

7. Specify an upgrade option, then click **OK**. **sqlupgrade** performs the upgrade eligibility test to verify that your current server meets the requirements for upgrading to the new version.

If your current server fails the upgrade eligibility test, fix the problem based on the information from the results of the test.

8. Click **OK** to upgrade your server.

You see the Status Output window, which shows the completion status and informational messages about the upgrade process.

Warning! Do not interrupt the upgrade, connect to Adaptive Server, or run any stored procedures while the upgrade is in progress.

You can also check progress by viewing the **sqlupgrade** log in `SYBASE/ SYBASE_ASE/init/logs/sqlupgrademMDD.NNN`, where:

- *MM* – is the month.
- *DD* – is the date.
- *NNN* – is a three-digit number identifying the **srvbuild** server upgrade session.

After all preupgrade checks are successful, **sqlupgrade** shuts the old server down, then starts the new Adaptive Server **dataserver** binary on the old `master` device.

When the upgrade succeeds, you can:

- Click **OK** to specify another server to upgrade, or
- Exit **sqlupgrade**, and go to *Post-upgrade tasks* on page 84.

The upgrade process:

- Creates the `RUN_servername` file, which contains the information required to restart the server
- Runs the `installmaster` script to create system procedures

To verify that the upgrade was successful, log in to the server and run:

- **select @@version** – Adaptive Server should return 15.7.
- **sp_configure "upgrade version"** – Adaptive Server should return 15000.

Upgrading Noninteractively Using sqlupgraderes

You can upgrade Adaptive Server without the graphical user interface (GUI), in noninteractive mode by using values from a resource file that defines the attributes for the server to be upgraded.

1. Edit the sample resource file included in your Adaptive Server distribution in `SYBASE/ SYBASE_ASE/init/sample_resource_files`. Use a text editor to edit the resource file as described in *Alternative Installation Methods*.

2. Run the **sqlupgraderes** utility, which creates a new file containing the values you specified for the upgrade session, and writes it to the `sqlupgrademmdd.NNN-server_name.rs` file in `$SYBASE/$SYBASE_ASE/init/logs/`.

where:

- *server_name* – is the server you are upgrading.
- *MM* – is the month.
- *DD* – is the date.
- *NNN* – is a three-digit number identifying the **srvbuild** session.

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.

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

All attributes are required, and all values are case-sensitive.

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

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

Manually Upgrading to Adaptive Server

You can upgrade manually without using a wizard or utility.

To upgrade Adaptive Server manually, use the **upgrade** executable located in the `$SYBASE/ASE-15_7/upgrade/`.

1. Shut down the pre-15.x server.
2. Copy the server entry from the old `interfaces` file to the new `interfaces` file.
3. Copy the `server_name.cfg` file from old directory to the new `$SYBASE` directory.
4. Start the 15.7 server with the old `master` device from the old server.
5. Run **upgrade**. If problems occur, fix them, and re-run the upgrade process.

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

Installation of Version 15.7 ESD #2 over an Existing 15.x Adaptive Server

Use a binary overlay to install Adaptive Server 15.7 ESD #2 over an existing version 15.x installation directory.

If you have enabled the Java in the database feature, create the `sybpcidb` database and disable the feature during the installation.

See also

- *Managing Java in the Database During Installations and Upgrades* on page 34

Determining the Adaptive Server Version

Before using the binary overlay, verify that your current version of Adaptive Server is at least 15.x.

If your server is at a 15.7 version level, you can begin installing Adaptive Server 15.7 ESD #2.

1. Verify that you have version 15.x:

- If your server is running:

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

- If your server is not running:

```
$_SYBASE/$_SYBASE_ASE/bin/dataserver -v
```

2. Source environment variables before executing **dataserver -v**.

If the result shows that the Adaptive Server is earlier than version 15.x, you cannot use the binary overlay, but must use an upgrade method instead. See *Chapter 10, Adaptive Server Upgrades* on page 65.

Backing Up Adaptive Server

Installing Adaptive Server version 15.7 and later overwrites the current Adaptive Server software. Before installing, verify that your databases are error free, and back up the Sybase directory.

1. To ensure that your database is error free, run **dbcc checkdb**, **dbcc checkcatalog**, and **dbcc checkstorage** before 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.

2. 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.
3. Adaptive Server provides the **uninstmsgs.ebf** script for backing up `sysmessages` before upgrading to 15.7 ESD #2. Use it to back up `sysmessages` before running **instmsgs.ebf**.

Installing Adaptive Server Using Binary Overlay

Use the installer to install Adaptive Server version 15.7 ESD #2 on top of an Adaptive Server that is no earlier than version 15.x.

1. From your SYBASE directory, shut down the servers that are using the binaries. Use a regular "polite" shutdown as opposed to using the **shutdown with nowait** option, both as the initial and final steps. Doing so flushes free space accounting figures, object statistics, and runs **checkpoint** on the database to minimize recovery work during the upgrade process.
2. Use the installer to load the new software from the CD or DVD.
Navigate to the CD or DVD drive and enter `./setup.bin`.
3. Install Adaptive Server 15.7 ESD #2 into the existing \$SYBASE installation path.

Note: Once files are loaded into your \$SYBASE directory, select not to configure the new server, and click **Continue** to finish the installation.

4. Restart Adaptive Server by using :
`$$SYBASE/$SYBASE_ASE/install/startserver [-f RUN_server_name file]`
5. Run **select @@version**. The server should now be at version 15.7 ESD #2.
6. In Adaptive Server version 15.7, several changes were made to the system stored procedures and many new error messages have been added. After performing a binary overlay, you must perform the postinstallation tasks before running **installmaster** or `instmsgs.ebf`.
7. Once the server has been brought up on the new binary and you have completed the upgrade of system tables, shut down, then restart Adaptive Server. This prevents incorrect table statistics from being stored in memory, and which may lead to poorly optimized query plans.

See also

- *Chapter 9, Postinstallation Tasks* on page 57

Postupgrade Tasks

After you have upgraded, make sure your new Adaptive Server is up and running.

The upgrade process does not cause any changes in existing statistics, so you need not run **update statistics** on any tables after the upgrade. If you are upgrading from Adaptive Server version 15.x, however, restart the server to make the statistics available.

1. Run **dbcc upgrade_object()** for each database to explicit re-compile the objects
2. If you upgraded from Adaptive Server version 12.5.2 or earlier, run **dbcc checkcatalog** with the **fix** option to ensure there are no issues with OAM pages:

```
dbcc checkcatalog (database_name, fix)
```

3. Install new system stored procedures before starting any application activities:

```
isql -Usa -Psa_password -Sserver_name  
-i$SYBASE/$SYBASE_ASE/scripts/script_name -ooutput_file
```

The output of the script is saved in `output_file`. See *How Compiled Objects are Handled When Upgrading Adaptive Server* on page 97 if you are using **dbcc upgrade_object** to upgrade compiled objects such as check constraints, rules, triggers, and views.

4. Run the scripts for different options or licensed features:
 - **installcommit** – if you use two-phase commit or distributed transactions, re-run **installcommit** to restore:
 - **sp_start_xact**
 - **sp_abort_xact**
 - **sp_remove_xact**
 - **sp_stat_xact**
 - **sp_scan_xact**
 - **sp_probe_xact**
 - **installsecurity** – if you used auditing in your earlier installation.
 - **installhasvss** – if you enabled high availability, and are using this feature in this installation.
 - **installmsgsvss** – if you enabled real-time messaging, and you are using this feature in this installation.
 - **installpcidb** – if you enabled the Java in the database feature in your earlier installation.
 - **installjsdb** – if you enabled Job Scheduler in your earlier installation.
5. If you have any system stored procedures that you had saved from before the upgrade (because you had modified them without altering their names), reload them now.

See also

- *How Compiled Objects Are Handled When Upgrading Adaptive Server* on page 97

Running the instmsgs.ebf Script

Run message-related scripts after upgrading from an Adaptive Server version 15.0.x to version 15.7 ESD #2.

1. If you upgraded from Adaptive Server version 15.0.x, run **uninstmsgs.ebf**:

```
isql -Usa -Ppassword -w1000 -iuninstmsgs.ebf -orestoremsgs.ebf
```

This safeguards your modified messages in the master database before you install the default version 15.7 ESD #2 messages.

2. Regardless of the Adaptive Server version you upgraded from, run the **instmsgs.ebf**:

```
isql -Usa -Ppassword -iinstmsgs.ebf
```

Note: To undo **instmsgs.ebf** changes, run this script once you have downgraded to the version from which you upgraded:

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

3. If you use localized files, install any localized language using **langinstall**, **sqlloc**, or **syconfig**.

If you run **instmsgs.ebf** after installing the 15.7 ESD #2 localized message, this script may delete some new messages.

Restoring Functionality in Adaptive Server After Upgrading

Restore functionality in the server after an upgrade.

1. If you changed any configuration parameters before upgrading, use **sp_configure** to set them back to their earlier values.
2. Use **sp_dboption** to reset any database options you disabled before upgrading.
3. Before you use the upgraded server, verify that all scripts developed at your site point to Adaptive Server 15.7 ESD #2.
4. Verify procedure cache allocation. The size should be the same as before upgrading, unless the original size was smaller than the default value.
5. Check the procedure cache requirements. Stored procedures, triggers, and other compiled objects require much more memory to run in Adaptive Server 15.7 ESD #2.

Use **sp_configure** to increase the **procedure cache size** during runtime, and **sp_configure verify** to verify any changes you make to the configuration file without having to restart Adaptive Server:

```
sp_configure "configuration file", 0, "verify",  
"full_path_to_file"
```

See *Reference Manual: Procedures* and the *Performance and Tuning Guide* for details about **sp_configure** and **sp_sysmon**, and the *System Administration Guide* for information on configuring memory.

6. Verify data cache allocation.

The server ensures that all data cache sizes remain the same after upgrading. Adaptive Server treats this size as an absolute value of 8MB and sets it in the `config` file..

During the upgrade process, the server ensures that the default data cache size remains the same. Therefore, during the preupgrade process, the size of the default data cache is obtained and written to the configuration file as an absolute value, not as a default. This enables the server to have the same default data cache size as before the upgrade. If this size is less than the default size of 8MB, the server allocates a default data cache of 8MB.

7. If you unmirrored devices, remirror them using the **disk remirror** command.

8. If you used compiled objects, see *How Compiled Objects Are Handled When Upgrading Adaptive Server* on page 97.

9. If you used two-phase commit in the earlier version of Adaptive Server, run the script to install the two-phase commit tables:

```
isql -Usa -Psa_password -Sserver_name  
-i$SYBASE/$SYBASE_ASE/scripts/installcommit
```

Reenabling Auditing

If the server you upgraded from was configured for auditing, reenabling auditing in the upgraded server.

1. Enter:

```
sp_configure 'auditing', 1
```

2. Reenable auditing for all system stored procedures for which auditing was enabled before the upgrade.

a) Use the output of **sp_displayaudit** recorded during preupgrade to help determine the system stored procedures for which auditing was enabled.

b) 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:

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

Updating Threshold Procedures for Audit Segments

An update is required for threshold procedures used to archive audit segments.

If your previous installation used a threshold procedure, similar to this, to archive the *sysaudits* table:

```
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 15.7 ESD #2, then you must alter *MyPre15SysAuditHistoryTable* to add a *nodeid* column using:


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

See the *Reference Manual: Tables* for system tables for details about the `sysaudits` tables.

Reenabling Replication Server After the Dataserver Upgrade

If you disabled replication before the upgrade, you must reenabling replication.

1. Remove older formatted log records from your database.
2. 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 Agent™ and other log readers from accessing the preupgrade portion of the transaction logs.

```
1> use master
2> go
1> dump database sales to dumpdev
2> go
1> dump transaction sales with truncate_only
2> go
```

3. Reenable replication.

Restoring Replication on Replicate Databases

Restore replication if you have upgraded a destination-only dataserver in a replication system.

Perform these steps for each replicate database and Replication Server System Database (RSSD).

1. Start Adaptive Server if it is not already running.
2. Log in to the Adaptive Server.
3. If you have zeroed out the locator for the database, go to step 4. Otherwise, stop Replication Server, then run:

```
1> use RSSD
2> go
1> rs_zeroltm dataserver, database
2> go
```

4. Restart the Replication Servers.
5. 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 15.7 ESD #2, 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*.

See also

- *Chapter 7, Installing Components from the PC-Client* on page 51

Restoring Replication on Primary Databases

Restore replication if you have upgraded a source database, or primary database, in a Replication Server system, or if the destination database is also a source for other Replication Servers.

1. If you have zeroed out the locator for the database, go to the next step. Otherwise, stop Replication Server, then run:

```
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 an RSSD, resume the Replication Server connection to the RSSD by issuing the following command to the Replication Server specifying the same string that was specified during the **'hibernate_on'** command:

```
1> sysadmin hibernate_off, 'Replication Server'
2> go
```

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 the Log Transaction Manager, restart it.

Migrate

You can migrate from 32-bit to 64-bit versions on different computers or partitions.

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

To migrate, you can:

- Use **dump** and **load**.
- Use the **bcp** utility.
- Replace the binary.

Migrating Data Using a Dump-and-Load Method

To perform migrations, back up and restore databases using the **dump** and **load** commands.

1. In the 32-bit Adaptive Server, run **dbcc** checks (**checkdb**, **checkalloc**, **checkcatalog**, and **checkstorage**) on all databases in the 32-bit Adaptive Server to ensure they are error-free.
2. Create a 64-bit server in a new directory.
3. Create devices and databases to match those in the 32-bit server. Make sure the *sysusages* mapping is correct.

Note: Allow 10 percent additional space for the *sysystemprocs* database.

4. Dump the databases from 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.

For information on upgrading compiled objects, see *Finding Compiled Object Errors Before Production* on page 98.

Using bcp to Migrate Data

If you use DDL scripts to create devices, databases, tables, rules, stored procedures, triggers, and views, you can use **bcp** to migrate data from a 32-bit Adaptive Server to a 64-bit Adaptive Server.

If you:

- Do not have DDL scripts, use the **ddlgen** utility to re-create the schema for the Adaptive Server you are about to migrate. See the *Utility Guide*.
 - 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.
1. In the 32-bit Adaptive Server, run **dbcc** checks (**checkdb**, **checkalloc**, **checkcatalog**, and **checkstorage**) on all databases in the 32-bit Adaptive Server to ensure they are error-free.
 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, and stored procedures.

7. Run **dbcc** checks on the 64-bit Adaptive Server server and make sure they run cleanly.

Migrating Data by Replacing the Binary

Migrate data from a 32-bit server to a 64-bit server by replacing the binary.

1. In the 32-bit Adaptive Server, run **dbcc** checks (**checkdb**, **checkalloc**, **checkcatalog**, and **checkstorage**) on all databases in the 32-bit Adaptive Server to ensure they are error-free.
2. Copy the files for the 64-bit Adaptive Server into a new directory.
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 `$SYBASE` directory and source the `SYBASE.csh` script (C shell).
9. Change to the 64-bit `$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 tables in the `dbccdb`. This may cause you to lose data.
12. Drop and re-create the 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.

Adaptive Server Components and Related Products

After you finish upgrading Adaptive Server, upgrade its components and related products.

Refer to *Using Sybase Failover in a High Availability System* for instructions on how to upgrade an Adaptive Server configured with high availability.

Upgrading Job Scheduler

After upgrading to a new Adaptive Server, upgrade Job Scheduler.

Note: You must have `$SYBASE/$SYBASE_OCS/bin` directory in `$PATH` so the **isql** executable is accessible. Use **isql** to perform all the steps in this task.

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

2. Make sure the new server is running.

3. Stop the Job Scheduler:

```
1> sybmgmtdb..sp_sjobcontrol @name=NULL, @option="stop_js"
2> go
```

4. Ensure that at least 9000 locks are configured. If the number of locks on the server is fewer than 9000, increase it:

```
1> sp_configure "number of locks", 9000
2> go
```

5. Upgrade the internal Job Scheduler SQL code using:

```
1> use sybmgmtdb
2> go
1> dbcc upgrade_object
2> go
```

6. Restart Adaptive Server.

7. (Optional) Add more log space. Some 64-bit platforms require additional space for the sybmgmtdb log:

```
1> use master
2> go
1> alter database sybmgmtdb LOG on sybmgmtdev=20
2> go
```

8. To upgrade sybmgmtdb, run the installjsdb script that is included and save the output to a file:

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

Note: When upgrading from Adaptive Server version 12.5.x to 15.5 and later, increase the size of sybmgmtdb from 50MB to 90MB.

9. Enable Job Scheduler to start when Adaptive Server starts:

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

10. To start Job Scheduler from **isql**, enter:

```
sybmgmtdb..sp_sjobcontrol @name=NULL, @option="start_js"
go
```

Upgrading Job Scheduler Templates

After upgrading to the new server, upgrade the templates and jobs created by Job Scheduler.

Note: Several changes impact the Job Scheduler templates. These changes make some templates incompatible with earlier versions of the server. The current templates are version 3.0 in the XML files.

1. Disable Job Scheduler.

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. The new location is `$SYBASE/$SYBASE_ASE/jobscheduler`.

Directories below `jobscheduler` remain the same.

3. Copy files from the `jobscheduler` directory to the ASE-15_0 directory. If you are installing the new server directories over the old server directories, the installer automatically moves your files to the new `jobscheduler` directory.
4. If you modified Sybase-supplied templates, stored procedures, or XML documents, do not overlay the new templates with your modified versions. If you do, you 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: If you modify Sybase-supplied templates, 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 must make to the SQL code in the jobs created from the templates.

Table 7. Changed Job Scheduler Templates

Template	Modified Files	Datatype for <parameter> Changed from varchar(5) to int
dump database	SybBackupDbToDiskTemplate.xml jst_dump_databases	@use_srvr_name
dump database log	SybBackupLogToDiskTemplate.xml jst_dump_log	@truncate_flag and @use_srvr_name
update statistics	SybUpdateStatsTemplate.xml jst_update_statistics	@index_flag
rebuild indexes	SybRebuildIndexTemplate.xml jst_reorg_rebuild_indexes	@dump_flag

Template	Modified Files	Datatype for <parameter> Changed from varchar(5) to int
rebuild table	SybRebuildTableTemplate.xml jst_reorg_rebuild_tables	@dump_flag
reclaim indexes	SybReclaimIndexTemplate.xml jst_reclaim_index_spaces	@dump_flag
reclaim tables	SybReclaimTableTemplate.xml jst_reclaim_table_spaces	@resume_flag

6. Some Job Scheduler templates have been modified to support new server functionality. These changes are related to new parameters for specifying a partition name, or a data change value for the new server commands that added these options. If you have jobs created from any of the enhanced templates, modify the SQL of the jobs for the 15.7 ESD #2 servers.

If you have a job scheduled to run on a pre-15.7 ESD #2 server and you need it to also run on 15.7 ESD #2 servers, leave the existing job alone and create a new job for 15.7 ESD #2 servers, as the job commands are different.

You need not modify jobs you run on pre-15.7 ESD #2 servers. The following table lists the templates that changed in version 3.0 and the changes you must 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 are scheduled on pre-15.0.1 servers; you must use the 2.1 or 2.2 versions for pre-15.0.1 servers.

Table 8. Modified Job Scheduler Templates

Template	Modified File	Change Description	Job Modifications
delete statistics	SybDeleteStatsTemplate.xml jst_delete_statistics	Added @ptn_name as third parameter.	Optional.

Template	Modified File	Change Description	Job Modifications
update statistics	SybUpdateStatsTemplate.xml jst_update_statistics	Added @ptn_name as fifth parameter, @datachg_threshold as the 10th parameter, and added reference.	Required. Include values (or NULL) for new parameters.
rebuild indexes	SybRebuildIndexTemplate.xml jst_reorg_rebuild_indexes	Added @ndx_ptn_name as third parameter.	Required. Include value (or NULL) for new parameter.
reclaim indexes	SybReclaimIndexTemplate.xml jst_reclaim_index_space	Added @ptn_name as third parameter.	Required. Include value (or NULL) for new parameter.
reclaim tables	SybReclaimTableTemplate.xml jst_reclaim_table_space	Added @ptn_name as second parameter.	Required. Include value (or NULL) for new parameter.
multiple	jst_get_freespace, jst_get_usedspace	Replace reserved_pgs and data_pgs with reserved_pages and data_pages .	No impact on job SQL.

7. Install the template stored procedures and and move them to the Job Scheduler template stored procedure directory. For example:

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

- a) Run the stored procedure installation script for each server being upgraded:

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

Note: Upgrade template stored procedures on all Job Scheduler servers and target servers that are upgraded to Adaptive Server version 15.5. Do not install them on pre-15.5 servers.

8. Install the template XML documents. Move to the JS template XML directory. For example:

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

- a) Run the XML installation 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 omit the parameter completely, as "en" is the default.

Note: Upgrade template XML 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 and Cluster Support

Upgrade the cluster subsystem.

Adaptive Server 15.7 ESD #2 supports these cluster platforms for high availability:

- HP-UX – MCSG 11.17
- HPIA – MCSG 11.18
- IBM AIX – HACMP 5.4
- Sun Solaris – VCS4.0, SunCluster 3.2
- Linux-AMD – VCS4.1
- Win2003 – Cluster Manager 5.2

There are two ways to upgrade a cluster subsystem:

- Perform a major upgrade, which involves cluster downtime, then shut down and restart all servers. You must:
 - a) Run **sp_companion suspend**, as described in *Using Sybase Failover in a High Availability System*.
 - b) Take offline the resource groups of primary and secondary companions on both nodes. Ensure that the companion server and corresponding resource groups are not automatically brought online until the cluster system upgrade is complete.
 - c) Upgrade the cluster subsystem following the instructions from the cluster system vendor. You may find options to migrate the current resource groups to the new cluster version. If such an option is not available (or if the resource groups get deleted or corrupted), re-create the resource groups and configure them appropriately after you have upgraded the cluster system.
 - d) Bring the resource groups online. This should bring primary and secondary companions online on their respective nodes.
 - e) Run **sp_companion resume** as described in *Using Sybase Failover in a High Availability System*.
- Perform a minor upgrade, to avoid cluster downtime. Nodes are failed over to other nodes and upgraded one at a time. For example, assuming ASE1 is the primary companion on node N1, and ASE2 is secondary companion running on node N2:
 - a) Upgrade the primary companion:
 1. Relocate the primary resource group to N2 or shut down ASE1. This causes ASE1 to fail over from N1 to N2.

2. Upgrade the cluster subsystem on N1 following the upgrade instructions provided by the vendor.
 3. Fail back ASE1 from N2 to N1. See the appropriate cluster chapter in *Using Sybase Failover in a High Availability System* for more details on Adaptive Server failback.
- b) Upgrade the secondary companion. If you are using a symmetric configuration, follow the steps in "upgrade the primary companion," above, for ASE2.
If you are using an asymmetric configuration:
1. Take the secondary resource group offline and verify ASE2 is shut down. ASE2 is unavailable during this upgrade.
 2. Upgrade the cluster subsystem on N2 following the upgrade instructions from the vendor.
 3. Start ASE2 by bringing the secondary resource group online to N2.

Upgrading Java in the Database

If Java in the database functionality is currently enabled, you must run the `installpcidb` script, which creates tables and stored procedures in the `sybpcidb` database.

1. Using `isql`, run the `installpcidb` script that is included with the release. Save the output to an operating system file.

```
isql -Usa -P<sa_password> -S<server_name>  
-i$SYBASE/$SYBASE_ASE/scripts/installpcidb -o<output_file>
```

2. Enable the Java in the database functionality:

```
1> sp_configure 'enable pci', 1  
2> go  
1> sp_configure 'enable java', 1  
2> go
```

You may need to increase '**max memory**' to enable these parameters. Restart the server for changes to take effect. You can configure the maximum size of the PCI bridge memory pool through the '**pci memory size**' configuration parameter. For more information, see *Java in Adaptive Server Enterprise*.

Enabling Java in the Database in a High Availability System

You can use the Java in the database feature on a system that also uses high availability.

Drop high availability companionship before installing the `sybpcidb` database, and reestablish companionship afterwards.

Java in the database functionality must be either enabled or disabled on both nodes of the high availability system.

Upgrading Backup Server

You can upgrade Backup Server at any point after you upgrade Adaptive Server using similar steps. XP Server does not have a formal upgrade process.

1. If you are upgrading during the initial server installation, select:
 - **Upgrade Existing Servers** when the installer prompts Build
 - **Upgrade Adaptive Server and Backup Servers**

This launches the **sqlupgrade** utility. Click **OK**.

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

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

3. Perform the upgrade steps as prompted.

See also

- *Upgrade to Adaptive Server 15.7 ESD #2* on page 79

Upgrading Databases Using Dump and Load

When you upgrade Adaptive Server, you can also upgrade databases and transaction logs from versions of Adaptive Server as early as 12.5 using the **dump** and **load** commands.

Some issues of which you should be aware:

- The upgrade process requires space for copying data and logging changes to the system tables. If the source database in the dump was full, the upgrade process might fail. 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.

How Compiled Objects Are Handled When Upgrading Adaptive Server

Adaptive Server upgrades compiled objects based on their source text.

Compiled objects include:

- Check constraints
- Defaults
- Rules
- Stored procedures (including extended stored procedures)
- Triggers
- Views

The source text for each compiled object is stored in the `syscomments` table, unless it has been manually deleted. The upgrade process verifies the existence of the source text in `syscomments`. However, 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 its source text is verified when you upgrade. 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.

You do not receive any notification if the compiled objects in your database dump are missing 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.

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_checksourc` and `sp_hidetext`, see *Reference Manual: Procedures*.

Note: If you are upgrading from a 32-bit to 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 preupgrade process calculates the exact size; increase your upgraded database size accordingly.

To determine whether a compiled object has been upgraded, and you are upgrading to a 64-bit pointer size in the same version, look at the `sysprocedures.status` column. It contains a hexadecimal bit setting of 0x2 to indicate that the object uses 64-bit pointers. If this bit is not set, the object is a 32-bit object, which means the object has not been upgraded.

To ensure that compiled objects have been upgraded successfully before they are invoked, upgrade them manually using the `dbcc upgrade_object` command.

Finding Compiled Object Errors Before Production

Use `dbcc upgrade_object` to identify potential problem areas that may require manual changes to achieve the correct behavior.

After reviewing the errors and potential problem areas, and fixing those that need to be changed, use `dbcc upgrade_object` to upgrade compiled objects manually instead of waiting for the server to upgrade the objects automatically.

Problem	Description	Solution
Missing, truncated, or corrupted source text	If the source text in <code>syscomments</code> has been deleted, truncated, or otherwise corrupted, dbcc upgrade_object may report syntax errors.	If: <ul style="list-style-type: none"> The source text was not hidden – use sp_helptext to verify the completeness of the source text. Truncation or other corruption has occurred – drop and re-create the compiled object.
Temporary table references	If a compiled object, such as a stored procedure or trigger refers to a temporary table (<code>#temp table_name</code>) that was created outside the body of the object, the upgrade fails, and dbcc upgrade_object returns an error.	Create the temporary table exactly as expected by the compiled object, then execute dbcc upgrade_object again. Do not do this if the compiled object is upgraded automatically when it is invoked.
Reserved word errors	If you load a database dump from an earlier version of Adaptive Server into Adaptive Server 15.7 or later and the dump contains a stored procedure that uses a word that is now reserved, when you run dbcc upgrade_object on that stored procedure, the command returns an 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.

Quoted Identifier Errors

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.

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.

For compiled objects created in version 11.9.2 or later, the upgrade process automatically activates or deactivates quoted identifiers as appropriate.

1. Activate quoted identifiers before running **dbcc upgrade_object**.

When quoted identifiers are active, use single quotes instead of double quotes around quoted **dbcc upgrade_object** keywords.

2. If quoted identifier errors occur, use the **set** command to activate **quoted identifiers**, and then run **dbcc upgrade_object** to upgrade the object.

Determining Whether to Change **select *** in Views

Determine whether columns have been added to or deleted from the table since the view was created.

Perform these queries when **dbcc upgrade_object** reports the existence of **select *** in a view:

1. Compare the output of `syscolumns` for the original view to the output of the table.

In this example, you have the following statement:

```
create view all_emps as select * from employees
```

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`.

2. Before upgrading the `all_emps` view, use these 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")
```

3. Compare the output of the two queries by running **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.

If the table contains more columns than the view, retain the preupgrade results of the **select *** statement. Change the **select *** statement to a **select** statement with specific column names.

4. 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.

CHAPTER 11 Adaptive Server Downgrades

An Adaptive Server that has been upgraded to 15.7 or later requires specific tasks before it can be downgraded.

Even if you have not used any of the new features in Adaptive Server 15.7 or later, the upgrade process added columns to system tables. This means you must use **sp_downgrade** to perform the downgrade.

The **sp_downgrade** procedure requires `sybase_ts_role`, and you must have `sa_role` or `sso_role` permissions. See *sp_downgrade* in *Reference Manual: Procedures*.

There are additional steps to perform if you are using encryption or replicated databases.

Note: You cannot downgrade a single database through **dump** and **load** directly from Adaptive Server 15.7 ESD #2 to an earlier version.

Preparing to Downgrade Adaptive Server

Prepare your system before starting the downgrade.

Note: If you are downgrading to Adaptive Server 15.7 ESD #1, which supports features such as compression and in-row LOB, skip the steps in this section.

Before you begin downgrading Adaptive Server, perform these steps for the features or configurations you enabled in Adaptive Server 15.7 ESD #2.

- If the logical page size for your Adaptive Server is larger than 8192 bytes and you have configured any databases to allow wide data-only-locked (DOL) rows, turn the option off:

```
sp_dboption @dbname, 'allow wide dol rows', false
```

Before you downgrade Adaptive Server, verify that no wide DOL rows exist in any table in these databases. Adaptive Server cannot find them, and thus does not warn you about their existence. When you then downgrade to a version earlier than 15.7 ESD #2, Adaptive Server treats these as corrupt.

This does not occur if your logical page size is 8192 bytes or smaller.

- If you have configured any database to use either row or page compression, turn it off:

```
alter database @dbname set compression none
```
- If any database has nonzero, in-row large object (LOB) lengths, set it to zero:

```
alter database @dbname set inrow_lob_length = 0
```
- If any table is configured to use compression, turn it off:

```
alter table @tablename set compression = none
reorg rebuild @tablename
```

- If any table uses LOB compression or in-row LOBs:
 - a) Copy the data from that table to a new table.
 - b) Drop the original table.
- If any table was altered to include nonmaterialized columns, convert those columns to standard columns:

```
reorg rebuild @tablename
```
- If there are any holes in the log from previous **alter database log off** commands, remove the holes by using **alter database log on** to extend the log.

If you do not perform this before the downgrade, you see a message similar to:

```
Error: Database 'dbname' contains num hidden pages that have to be
filled.
Please, use ALTER DATABASE LOG ON command to extend the log num
pages.
```

You can extend the log on any device, as long as you specify the number of indicated pages as being at least as large as the *num* value.

Downgrading from Adaptive Server 15.7 ESD #2

Use **sp_downgrade** to downgrade Adaptive Server 15.7 ESD #2 to an earlier version.

Prerequisites

Note: Do not use these steps if you are downgrading from Adaptive Server version 15.7 ESD #2 to either version 15.7 or 15.7 ESD #1. Instead, see *Downgrading to an Earlier Version of Adaptive Server 15.7* on page 104.

1. Save the current auditing settings for the 15.7 ESD #2 Adaptive Server by saving the output of the **sp_displayaudit** system procedure in all system and user databases where any auditing options have been turned on:

```
1> sp_displayaudit
2> go
```

2. Back up all of your databases and the \$SYBASE release area.

Task

Downgrade from Adaptive Server 15.7 ESD #2 to an Adaptive Server 15.0, 15.0.1, 15.0.2, 15.0.3, or 15.5. Downgrading to versions of Adaptive Server earlier than 15.0 is not supported.

1. Start the 15.7 ESD #2 server in single-user mode using **dataserver -m** to guarantee that no other user can access Adaptive Server while you are going through the downgrade steps. See the *Utility Guide* for more details on starting the server in single-user mode.

2. Make sure Adaptive Server 15.7 ESD #2 is prepared for downgrade by executing this command from the master database:

```
sp_downgrade 'prepare', @toversion='version'
```

The format for the value of *version* does not need periods, so these are all valid: "15.5", "155", "15.0", "150", "15.0.1", "1501", "15.0.2", "1502", "15.0.3", "1503". The version you enter should be the Adaptive Server version to which you want to downgrade.

Note: If you have used the 15.0.2 features for encrypted columns and you are downgrading to:

- Adaptive Server 15.0 or 15.0 ESD #1 – specify the version as "15.0."
 - Adaptive Server 15.0, 15.0 ESD #2, 15.0.1, or a 15.0.1 ESD – specify the version as "15.0.1."
 - Adaptive Server 15.0.2 or a 15.0.2 ESD – specify the version as "15.0.2."
-

sp_downgrade 'prepare' validates the readiness of Adaptive Server 15.7 ESD #2 for downgrade. It may print messages regarding manual changes required before the downgrade can be completed. Repeat this step, correcting any errors reported, and make sure you understand the implications of all warnings before proceeding.

3. Execute the following, where *version* is the same number you used in the previous step:


```
sp_downgrade 'downgrade', @toversion='version', @override=1
```

Once this procedure completes successfully, there can be no more activity on the 15.7 ESD #2 server. Run **checkpoint** and immediately issue the **shutdown** command.

Note: Since Adaptive Server 15.7 ESD #2 transaction logs may contain data that older server versions cannot interpret correctly, all transactions in all databases must be completed to prevent the older server from recovering any transactions. To ensure that all transactions are completed, issue the normal **shutdown** command, and not the **shutdown with nowait** command, after running **sp_downgrade**.

4. Copy the RUN_SERVER file to the release area for the version you are downgrading to. Modify the file to use the **dataserver** binary of the version to which you are downgrading, using **-e**, **-c**, and **-M** option point to the old server.

For example, if you are downgrading to 15.0.2, change 15.7-related information to 15.0.2-related information. From the SYBASE directory, issue:

```
/work_dirs/sybase/ase1502/ASE-15_0/bin/dataserver\  
-s old_server_name\  
-d downgraded_master_device\  
-e old_server_log_file\  
-c old_server_config_file\  
-M /work_dirs/sybase/ase1502\  

```

5. Restart the old server using the modified RUN_SERVER file. Ensure that the environment variable points to the old server release, such as \$SYBASE, \$SYBASE_ASE, or \$SYBASE_OCS.

Note: When downgrading to a 15.0 or a 15.0.1 server, you can reset user passwords by setting **sp_passwordpolicy 'allow password downgrade'** to **0**. User passwords are reset

and print on the console. Save the output of `sp_downgrade 'downgrade', <version>` into a file, so the passwords are not lost. To regenerate the sa password, restart the downgraded server with `-psa`. See `sp_passwordpolicy` in *Reference Manual: Procedures*.

6. To restore the original messages for the earlier version of Adaptive Server, run this script after downgrading the server:

```
isql -Usa -Psa_password -irestoremsgs.ebf
```

7. Run any post-downgrade scripts that may be applicable to your site, including `installmaster` and `instmsgs.ebf`.

If you have not used any 15.7 ESD #2 functionality, this completes the primary downgrade process.

See also

- *Post-Downgrade Tasks for Adaptive Server* on page 110
- *Additional Considerations for New Features Used* on page 107
- *Downgrading Job Scheduler* on page 108

Downgrading to an Earlier Version of Adaptive Server 15.7

Use `sp_downgrade_esd` to downgrade your installation of Adaptive Server 15.7 ESD #2 to Adaptive Server 15.7 or 15.7 ESD #1.

Prerequisites

To use `sp_downgrade_esd`, you must have `sa_role`, and be in the `master` database.

Task

Use the `sp_downgrade_esd` system procedure only when downgrading from Adaptive Server version 15.7 ESD #2 to either 15.7 or 15.7 ESD #1; `sp_downgrade_esd` does not run on any other version of Adaptive Server. To downgrade to a version earlier than 15.7, use `sp_downgrade`.

1. Restart Adaptive Server in single-user mode by specifying the `-m` option.
2. In a loop, run `sp_downgrade_esd` on each database. You need not downgrade temporary databases. These include `tempdb` as well as user-created temporary databases, as they are re-created from a template database when Adaptive Server starts. To downgrade your entire installation, run it several times, once per database in your installation.

Note: If you are downgrading the entire installation, perform `sp_downgrade_esd` on all of your non-temporary databases, then downgrade the `master` database last.

The syntax for `sp_downgrade_esd` is:

```
sp_downgrade_esd @db_name [, @esd [, @verbose]]
```

where:

- `@db_name` – is the name of the database you are downgrading
- `@esd` – is the ESD number you are downgrading to. The valid options are:
 - **"ESD1"** – for Adaptive Server version 15.7 ESD #1.
 - **"GA"** – for Adaptive Server version 15.7.
- **@verbose** – when used, displays the output in the verbose mode.

For example:

```

1> sp_downgrade_esd tempdb, esd1
2> go
Reverting database 'tempdb' to 'ESD1'.
Database 'tempdb' is now suitable for use by ESD1.
(return status = 0)
1> sp_downgrade_esd sybssystemprocs, esd1
2> go
Reverting database 'sybssystemprocs' to 'ESD1'.
Database 'sybssystemprocs' is now suitable for use by ESD1.
(return status = 0)
1>
2> sp_downgrade_esd sybssystemdb, esd1
3> go
Reverting database 'sybssystemdb' to 'ESD1'.
Database 'sybssystemdb' is now suitable for use by ESD1.
(return status = 0)
1> sp_downgrade_esd model, esd1
2> go
Reverting database 'model' to 'ESD1'.
Database 'model' is now suitable for use by ESD1.
(return status = 0)
1> sp_downgrade_esd MYASE_tdb_1, esd1
2> go
Reverting database 'MYASE_tdb_1' to 'ESD1'.
Database 'LUMINOUS_tdb_1' is now suitable for use by ESD1.
(return status = 0)
1> sp_downgrade_esd master, esd1
2> go
Reverting database 'master' to 'ESD1'.
Database 'master' is now suitable for use by ESD1.
(return status = 0)
1> shutdown
2> go

```

3. Shut down Adaptive Server 15.7 ESD #2.

You may now start the Adaptive Server 15.7 or 15.7 ESD #1 dataserver with Adaptive Server 15.7 ESD #2 quorum devices.

See also

- *Downgrading from Adaptive Server 15.7 ESD #2* on page 102
- *Dumping a 15.7 ESD #2 Database to Load into a 15.7 or 15.7 ESD #1 Adaptive Server* on page 106

Dumping a 15.7 ESD #2 Database to Load into a 15.7 or 15.7 ESD #1 Adaptive Server

Using **sp_downgrade_esd** on an Adaptive Server 15.7 ESD #2 database allows you to load it into Adaptive Server 15.7 or 15.7 ESD #1.

Prerequisites

To use **sp_downgrade_esd**, you must have `sa_role`, and be in the **master** database.

Task

If you have a database in Adaptive Server version 15.7 ESD #2 that you want to open in Adaptive Server 15.7 or 15.7 ESD #1, you can downgrade it with **sp_downgrade_esd** temporarily to perform the dump and load, after which you revert the database so that it is usable in the current version of Adaptive Server again.

Perform these one database at a time, and only for the databases you want to open in an earlier version of Adaptive Server 15.7.

1. Put a database (*@db_name*) in single-user mode.
2. Run **sp_downgrade_esd @db_name** on the database.

The syntax for **sp_downgrade_esd** is:

```
sp_downgrade_esd @db_name [, @esd [, @verbose]]
```

where:

- *@db_name* – is the name of the database you are downgrading.
 - *@esd* – is the ESD number you are downgrading to. The valid options are:
 - "1" – for Adaptive Server version 15.7 ESD #1.
 - "GA" – for Adaptive Server version 15.7.
 - **@verbose** – when used, displays the output in the verbose mode.
3. Use **dump database @db_name** to dump the database.
 4. Run **online database @db_name**. This brings the affected database back to its proper revision level of 15.7 ESD #2 after the downgrade procedure.
 5. Take the *@db_name* database out of single-user mode.

See also

- *Downgrading to an Earlier Version of Adaptive Server 15.7* on page 104

Additional Considerations for New Features Used

Additional downgrade steps may be necessary due to Adaptive Server 15.7 or later features that are unavailable in the Adaptive Server version to which you are downgrading.

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

Returning to 15.0.2 ESD #3 is easier than downgrading to other versions, because this version can handle a number of error conditions that may arise as a result of downgrading.

If you return to a pre-15.0.2 ESD #3 version of Adaptive Server, no error handling is provided. Using the new features can result in inexplicable behavior, including misleading error messages and stack traces, unless you are uninstall all usages of the new features reported by `sp_downgrade`.

Feature	Considerations
Wide, variable-length data-only-locked (DOL) rows	<p>You cannot downgrade versions that contain wide, variable-length DOL rows to versions that do not support the feature.</p> <p>Use the <code>dol_downgrade_check</code> function to determine if a database includes tables with wide, variable-length DOL rows. See <i>Reference Manual: Building Blocks</i>.</p> <p>If you:</p> <ol style="list-style-type: none"> 1. Disable the allow wide dol rows database option before performing a database or transaction log dump, then 2. Load these database dumps onto an earlier version of Adaptive Server <p>Data returned from queries on the wide rows may be incorrect, or queries might fail because the earlier version of Adaptive Server perceives the data rows as corrupt.</p>
Java in Adaptive Server	<ul style="list-style-type: none"> • User classes compiled by a Java version later than 1.2 do not work in pre-15.0.3 Adaptive Server versions. • Earlier versions of Adaptive Server ignore the enable pci parameter; you can remove it. • The <code>sybpcidb</code> database is not used by Adaptive Server versions earlier than 15.0.3, and may be dropped.
sp_addserver with filter parameter	<p>If you have any entries in <code>sys.servers.srvnetname</code> column that are longer than 32 bytes, delete these entries before downgrading, and add them back after downgrading.</p>

Feature	Considerations
Unicode noncharacters	If you used sp_configure to enable functionality group or permissive Unicode in Adaptive Server 15.7 ESD #2, and also stored Unicode noncharacters (that is, <code>u+FFFF</code> or <code>u+FFFE</code>) in a database, you may see a conversion error after you downgrade to a version earlier 15.7 when Adaptive Server retrieves the data and triggers a character set conversion. To avoid such errors, search for and remove those characters from the system.
Disabled varbinary truncation	If you used sp_configure to turn on the <code>disable varbinary truncation</code> configuration in Adaptive Server 15.7, the <code>varbinary</code> data stored in the system may have trailing zeros. When you downgrade to a version earlier than 15.7 ESD #2, these trailing zeros are not removed and become part of the select output. This does not affect query results involving comparisons.
Sharable inline default	If you used sp_configure to enable the sharable inline default setting in Adaptive Server 15.7 and downgrade to a version earlier than 15.7, running the ddlgen utility to generate data definition language can cause Adaptive Server to generate extra create default statements. You can remove such statements before using the output to create new schema.

Downgrading Job Scheduler

If you are downgrading to Adaptive Server versions 15.0 or 15.0.1, run **installjsdb** script from the earlier version.

1. Disable Job Scheduler:

```
1> sp_configure "enable job scheduler", 0
2> go
1> sybmgmtdb..sp_sjobcontrol @name=NULL, @option="stop_js"
2> go
```

2. Follow the downgrade steps for your platform.

3. After downgrading, enter this command from the earlier version of Adaptive Server:

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

Note: The directory with the **isql** executable (`$SYBASE/$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:

```
1> use sybmgmtdb
2> go
1> sp_sjobcontrol @name=NULL, @option="start_js"
2> go
```

See also

- *Downgrading from Adaptive Server 15.7 ESD #2* on page 102

Downgrading an Adaptive Server That Uses Encryption

Downgrade an Adaptive Server versions 15.7 and later if you enabled 256-bit key encryption keys.

1. Set the **downgrade_kek_size** configuration option to true. This means you must change the key encryption key size to 128 bits.
2. Run **sp_downgrade** to find all column encryption keys (CEKs) and key copies that are protected by master keys or by dual control.
3. Drop dual control for each CEK:

```
alter encryption key keyname
    [with passwd password
    modify encryption
    with no dual_control
```

4. Use **drop encryption key** to drop user and login password-protected key copies.
5. Create a system encryption password for each database that has master keys:

```
sp_encryption system_encr_passwd, <password>
```

6. Reencrypt each CEK with the master key:

```
alter encryption key keyname
    with passwd master key
    modify encryption
    with passwd system_encr_passwd
```

Note: The key owner must perform the reencryption for each CEK, or else the key ownership changes.

7. Reencrypt all base keys that are protected by user passwords:

```
alter encryption key keyname
    with passwd old_password
    modify encryption
    with passwd new_password
```

Note: The key owner must perform the reencryption for each base key; otherwise, the key ownership changes.

Downgrade Considerations for Adaptive Server with Replicated Databases

Special considerations exist for downgrading Adaptive Server version 15.7 that uses replicated databases,

If the server you are downgrading from has replicated databases and SQL replication has been enabled, you must follow different procedures depending upon the version of Adaptive Server you are downgrading to.

If your replication system is using RepAgent and you are downgrading to a 15.0.2 ESD #3 or later, there are no additional steps required. Adaptive Server 15.0.2 ESD #3 can keep the data synchronized.

If you are not using RepAgent, consult the specific documentation for the product you are using to replicate data.

If you downgrade to a version earlier than 15.0.2 ESD #3, Sybase recommends that you use the standard documented procedure to downgrade an Adaptive Server with replicated databases.

The standard procedure to downgrade Adaptive Server with replicated databases consists of ensuring that all replicated data in the primary database transaction log has been successfully transferred to the standby or replicate database. Once RepAgent is started again after the downgrade process, you must ensure that it does not read logs written with version 15.7:

1. Suspend transaction processing and replication activities.
2. Drain the transaction logs for primary databases.
3. Disable the secondary truncation point.

Post-Downgrade Tasks for Adaptive Server

After you perform the basic downgrade steps, finish the process by handling residual tasks.

1. Installing system stored procedures over the old version drops auditing information about system stored procedures. Use the output of **sp_displayaudit** recorded during predowngrade to help determine the system stored procedures for which auditing was enabled. Reenter the audit options using **sp_audit**. For example, if **sp_addlogin** was audited in your server before upgrading, run this command to reenabling auditing on **sp_addlogin**:

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

After you run `installmaster`, `installmodel`, `installcommit`, `installsecurity`, `installhasvss`, `installjsdb`, and `installmsgsvss` on your downgraded server, return system stored procedures to their earlier form. New

stored procedures introduced in releases subsequent to the downgraded server are not removed.

Warning! Executing new stored procedures against an older binary may cause unpredictable results.

2. During the upgrade to Adaptive Server 15.7 ESD #2, you should have created a `restore_msgs.ebf` file. Run this script now to restore your messages back to the version from which you upgraded:

```
isql -Usa -P <sa_password> -S <server_name> -i
<restore_msgs.ebf>
```

Note: If you upgraded from a version 15.0 or 15.0.1 and are reverting back to 15.0 or 15.0.1, run `instmsgs.ebf` file only from the relevant release.

3. 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 version to which you are downgrading.
4. If during the upgrade process you ran **update all statistics** on `syslogins`, you must **delete statistics** for `syslogins` and re-create them.

spt_values are dropped and re-created when you run **installmaster** from the release area you downgraded to. Any new types are eliminated from this table.

5. Run **installmaster** from the release area of the server to which you downgraded to remove configuration parameters that belong to 15.7 ESD #2, by deleting `sysconfigures` rows for configuration parameters that do not exist in `syscurconfigs`. After running **installmaster**, the error messages no longer appear when you start the server.

If you start a downgraded server using the 15.7 ESD #2 configuration file, the new options generate an `Unknown parameter` message. The unknown options are reported the first time you restart the server. You can ignore these messages; the configuration file is rewritten without the unknown options.

Note: Keywords **decrypt_default**, **xmltable**, and **path** were added in Adaptive Server 15.5 Cluster Edition, making it impossible to create identifiers using these names. You must change applications if you used these names.

See also

- *Additional Considerations for New Features Used* on page 107
- *Reenabling Replication Server After the Dataserver Upgrade* on page 87

CHAPTER 12 **Troubleshooting SySAM Errors**

Troubleshoot the most common SySAM errors.

See your Sybase product documentation for product-specific information. See the **SySAM FAQ** at <http://www.sybase.com/sysam> for the most recent troubleshooting information.

If a license problem occurs, solve the problem as quickly as possible. Products that cannot acquire a valid license during installation may still install or run under grace. If you do not solve the problem or acquire a valid license before the grace period expires, the product stops functioning.

Where to Find License Error Information

Typically, server products list problems in an error log, and optionally allow you to configure e-mail notifications.

GUI tool products generally show messages in a status window. Additionally, some products may support a command or menu option, such as **Help > About** to display the current license status.

If you use served licenses, and license servers, all license server status and error messages are written to the `SYBASE.LOG` debug file, which is located, by default, in the `LOG` subdirectory.

Problems and Solutions

If the product does not install or does not function after installation, contact SySAM Technical Support.

Error	Possible Causes	Solution
Installation Warning: Cannot Find a Valid License	You may not have the required licenses installed. If there are installed licenses, the licenses may be incorrect for the product or feature you are trying to install.	See <i>Installing for the First Time</i> in the <i>SySAM Users Guide</i> .

Error	Possible Causes	Solution
Updating an Existing Installation	If you are updating an existing installation, make sure your license authorizes you to install the update. See <i>Installing Product Updates, EBFs, and Support Renewal</i> in the <i>SySAM Users Guide</i> . If your license does not allow you to perform the update, the product may not be usable.	If you are installing an update that is authorized by the license, see the solution for <i>Product Cannot Checkout a License and Starts in Grace Period</i> error below before you proceed with the update.
License Server Executables and Scripts not Installed	When you installed your product, the license server was not installed. Some product installations offer the option to install a license server; however when available, this option is not selected, by default, for installation. You may need to explicitly install a license server. See your product's installation guide and release bulletin to determine if the product installer offers this option.	Depending on the options provided by your Sybase product's installation wizard, use one of the following solutions: <ul style="list-style-type: none"> • If your product offers this option, use the instructions in your product's installation guide to install the license server. • If your product does not include the option to install a license server, go to http://www.sybase.com/sysam and click Download the SySAM Standalone License Server – Free!.
License Server Does not Start	See <i>Possible Causes of a License Server Failure</i> in the <i>SySAM Users Guide</i> .	Go to SPDC or SMP, generate a valid served license for your product, and copy it into the <code>licenses</code> directory on the machine where the license server is installed. See <i>Generating Licenses at SPDC</i> or <i>Generating Licenses at SMP</i> .

Error	Possible Causes	Solution
License Server Does not Recognize the License Files	<ul style="list-style-type: none"> • The license was generated for a different machine or generated using an incorrect host ID. • The license has been altered. You cannot change any information in a generated license. • If the host ID for your platform is based on the network adapter identification, the most common problem occurs when you have used an ID associated with the address of a virtual network adaptor. 	<ul style="list-style-type: none"> • Verify that the host ID recorded in the license file matches the host ID of the actual machine for which the license was issued. If the host IDs do not match, go to SPDC or SMP, check in the license, then regenerate the license with the correct host ID. • If your license is being created by typing from a paper copy, verify whether errors occurred when the license information was entered. You can also download a new copy of the activated license from SPDC or SMP. • If the host ID for your platform is based on a network adapter, verify that the ID you are using is associated with a valid NIC and that the ID is not associated with loopback or virtual adapters. If the ID you are using is associated with a removable network adapter, verify that the adapter is actually attached to the computer.
Linux virtual machine does not appear to be part of the Microsoft Hyper-V host.	/usr/sbin/dmidecode in Linux virtual machine cannot read from /dev/mem.	Log in to Linux virtual machine as root, then execute chmod 4555 /usr/sbin/dmidecode .

Error	Possible Causes	Solution
<p>Product Does not Start, License Checkout Error</p>	<ul style="list-style-type: none"> • You have not generated and deployed the valid licenses for the product requesting licenses. • The required license does not exist and the product does not award a license on grace. • The product is configured to use the wrong edition or license type. • You have the wrong host ID for an unserved license. • When multiple product editions include optional features, the features are offered as separately licensed items for each edition. Licensed optional features work only with the licensed base product of the same edition. For example, if you order Adaptive Server Enterprise, you cannot use a Small Business Edition optional feature license with an Enterprise Edition base product. • You are using an unserved Standalone Seat (SS) type license on a terminal server. • The license is for a different operating system. • The license is a Floating License (FL) but is currently in use elsewhere. 	<p>At a command prompt or in a terminal window, execute the following commands, where <i>feature_name</i> is the name of the feature for which SySAM failed to check out a license:</p> <pre>sysam diag feature_name</pre> <p>If the SySAM script is unavailable, enter:</p> <pre>lmutil lmdiag -c license_directory_location feature_name</pre> <p>Go to SPDC or SMP and generate your product's required licenses. If you are trying to use a separately licensable optional feature you must have a license for both the base product and the option, and if the product has more than one edition, the edition of the base product and option must be the same.</p> <p>If you generated an invalid license, check in the license at SPDC or SMP, and regenerate the license with the correct information.</p>

Error	Possible Causes	Solution
Product Cannot Check Out a License and Starts in Grace Period	<p>To help you determine the possible causes when your product cannot check out a license, execute the following command from the <code>SYSAM-2_0/bin</code> directory, at a Windows command prompt or in a UNIX system terminal window, where <i>feature_name</i> is the name of the feature license that could not be checked out.</p> <pre>sysam diag feature_name</pre> <p>If the command output says that no licenses are available for checkout, this may be due to one of the reasons below, which are separated into served and unserved license deployment model causes and solutions.</p>	See <i>Served License Deployment Models</i> or <i>Unserved License Deployment Models</i> , both in the <i>SySAM Users Guide</i> .
Product Continues to Run in Grace Period After Fixing a License Problem	The license status has not yet been updated. When the product periodically performs license checks, the license status is not immediately updated.	Wait up to 6 hours for server products, and up to 1.5 hours for tool products.
Product Cannot Find Licenses for Optional Features	You either have not installed the license for the optional feature, or the license exists, but cannot be checked out.	See <i>Solution for Problem: Product Cannot Find Licenses for Optional Features</i> in the <i>SySAM Users Guide</i> .

Error	Possible Causes	Solution
Product Obtains the Wrong License	<p>These locations are searched, in the order specified, until a suitable license is found. When a license directory is specified, license files in that directory are loaded in the directory sort order. When a product looks for a license, it looks in:</p> <ul style="list-style-type: none"> • The locations that represent the values set for the SYBASE_LICENSE_FILE and LM_LICENSE_FILE variables. Sybase does not advocate using environmental variables and recommends that all licenses be located centrally in the expected licenses directory. • All files with a .lic extension in the licenses directory. This location is product specific, though typically the directory is \$SYBASE/SYSAM-2_0/licenses. <p>The first license that matches the feature name, version, edition, and license type filter is used; however, this license may not be the intended license your product requested.</p>	<ul style="list-style-type: none"> • Configure your product to choose a license for a specific edition and license type. • If you are using served licenses, use the options file to ensure that the correct license is used.
License checkout problems with unserved license	See <i>Unserved License Deployment Models</i> in the <i>SySAM Users Guide</i> .	
License checkout problems with served license	See <i>Served License Deployment Models</i> in the <i>SySAM Users Guide</i> .	

Installing for the First Time

The installer or product cannot find a valid license and issues a warning.

1. If the installation fails, go to SPDC or SMP to generate and install the required valid licenses, then restart the installation.
2. Verify that you have generated and deployed the correct license. The error message should indicate what the problem is:
 - An incorrect host ID for unserved licenses
 - An incorrect license server reference file, or the license server is not running on the machine and listening on the specified port
 - A license generated for a specific product edition and license type, but the product is configured to use a different product edition and license type.
 - Using Standalone Seat (SS) unserved licenses on a terminal server
 - Check the `SYBASE.LOG` file to make sure the license server has started.
3. Ignore the warning. If the product continues to install under grace, complete the installation. Generate and install the required SySAM licenses before the grace period expires.

Possible Causes of a License Server Failure

The most common reason for a license server failing to start is that no served licenses are installed on the license server.

There must be at least one served license in the `licenses` directory for the license server to start. When you install a new license server, there are, by default, no served licenses in the `licenses` directory. Generate and deploy a served license to this directory.

Additional reasons for a license server failing to start include:

- Using an unserved license with a license server – the license activated from SPDC or SMP is an unserved license. Examine the license files. Served licenses always start with a `SERVER` header. If you do not see a line starting with `SERVER`, you are using an unserved license that does not work with the license server.
- The license server port number is already in use – if you are using a specific port number for the license, that port number may already be in use. Use `netstat -a` to verify that the port number is free. If not, reassign the port or use a different port number for the license server.
- Mismatch between the host name in the license file and the actual host name – the host name is recorded next to the `SERVER` keyword. If it does not match the actual host name, correct the host name in the license file or set the value next to the `SERVER` to `this_host` which is a keyword that works with any host name.
- Header mismatch – when you have multiple license files, each one must have the same header, host name, port, and so on.

- The licenses on one machine were generated for a different machine – examine the host ID recorded next to the value in the license file host name in the SERVER header. It must match the host ID for the machine on which the license server will run.

Solution for Problem: Product Cannot Find License for Optional Features

If a license for the optional feature exists, but cannot be checked out by the product, perform these diagnostic tasks to determine the specific problem.

To ensure that the license for the optional feature exists and can be checked out from the machine where the product is running, execute:

```
sysam diag feature_name
```

Compare the feature and the license to:

- Verify that the optional feature is for the same edition of the feature’s base product.
- (Products and features that support active and standby licensing only) Verify that the active or standby flags match for the base product license and the optional feature license.
- Compare the product and optional feature editions, and active and standby status, then verify that the available license matches.

See *Sybase License Attributes* in the *Sybase Software Asset Management (SySAM) 2 Users Guide*.

If any of the above items do not match, go to SPDC or SMP to generate and download the correct licenses or adjust product configuration.

Unserviced License Deployment Models

Learn about license checkout problems with unserved licenses.

Possible Causes	Solutions
<p>The license for the product or feature you are using is not in the local <code>licenses</code> directory. The local license directory is product-specific, and the license must be deployed to the correct location for the product. Typically this directory is <code>\$SYBASE/SYSAM-2_0/licenses</code>.</p>	<p>Generate the required unserved licenses from SPDC or SMP and deploy them.</p>
<p>Your product’s license may have been activated with the incorrect host ID for the machine on which you are running the product.</p>	<ol style="list-style-type: none"> 1. Check the license back in to SPDC or SMP. 2. Regenerate the license at SPDC or SMP with the correct host ID. 3. Install the license to the local machine where you are running your instance of the SySAM 2-enabled Sybase product.

Possible Causes	Solutions
<p>The available licenses are for a different operating system or architecture than the machine on which you are running your product feature.</p>	<p>Obtain a license for the correct platform from SPDC.</p> <ol style="list-style-type: none"> 1. Check the license back in to SPDC. 2. Generate a license for the correct platform, or install the product on the correct platform. 3. Deploy a license to the local machine where you are running your instance of the Sybase product. <hr/> <p>Note: Licenses generated from SMP are platform-independent.</p>
<p>You are running your SySAM 2-enabled program in a terminal services (TS) environment with an unserved Standalone Seat (SS) license.</p>	<p>Set up served licenses for a terminal server environment:</p> <ol style="list-style-type: none"> 1. Log in to SPDC or SMP and check in your old license. 2. Generate served licenses and deploy them to a license server. 3. Configure the product to obtain licenses from the license server.
<p>If the output from running <code>sysam diag feature_name</code> states that a license for your product or feature is available for checkout, you still may be unable to check out the license because your product has been configured as a specific edition or with a specific license type. For example, the product is configured to look for an Adaptive Server® Enterprise, Enterprise Edition license, but only an Adaptive Server Enterprise, Developer Edition license is available; or, the product is configured to use a Server License (SR) license type, but only a CPU License (CP) is available.</p> <p>Another incompatibility may occur if your product has several editions, and you are trying to run a separately licensed feature from one edition with the base product of a different edition. You can run an Enterprise Edition base product only with an Enterprise Edition separately-licensed feature; you cannot run a Developer Edition feature with an Enterprise Edition base product.</p>	<p>If the problem is incompatible editions or license types, reconfigure your product, or go to SPDC or SMP and check in the incorrect license and regenerate a license for the correct edition or license type.</p>

Served License Deployment Models

Learn about license checkout problems for served licenses.

Possible Cause	Solution
The license server may not be running.	<p>Verify that the license server is running by entering this command from the <code>bin</code> directory on the license server host:</p> <pre>sysam status</pre> <p>If the license server is not running, restart it.</p>
The license server is running, but does not contain the desired license.	<p>Enter the following command to determine if the license server is trying to grant a license for the licensed feature you are trying to use, where <i>feature_name</i> is the name of the separately licensed product feature for which the license server failed to check out a license:</p> <pre>sysam status -f feature_name</pre> <p>If the license is generated for one edition or license type, and the product configured to use another, go to SPDC or SMP and generate the correct licenses for your product or feature.</p>
All licenses may be in use; that is no licenses are available for checkout.	<p>If the output from <code>sysam status -f feature_name</code> indicates that there are no available licenses:</p> <ul style="list-style-type: none"> • Go to SPDC or SMP and generate additional licenses; or, • If the product is using Floating License (FL) type licenses, you can shut down other instances of the product or feature you are using, running on other machines, to free licenses. See <i>Controlling License Use with the SySAM Options File</i> in the <i>Sybase Software Asset Management (SySAM) 2 Users Guide</i>.

Possible Cause	Solution
Cannot connect to the license server	<p>Ensure that the client can connect to both TCP/IP ports used by the license server. If there is a firewall, fix the ports used by both lmgrd and the SYBASE vendor daemon (the license server's constituent processes), and set up the firewall or VPN policies to allow access to both ports. The license server's SYBASE.log file shows the port numbers to be fixed:</p> <pre>19:04:47 (lmgrd) lmgrd tcp-port 27010 19:04:47 (lmgrd) Starting vendor daemons ... 19:04:47 (lmgrd) Starting vendor daemon at port 27011</pre> <p>You can telnet from the client machine to check that the ports can be accessed:</p> <pre>telnet keyserver 27010 telnet keyserver 27011</pre> <p>See <i>Access Through A Firewall or VPN</i> in the <i>Sybase Software Asset Management (SySAM) 2 Users Guide</i>.</p>

Contacting SySAM Technical Support

If you cannot resolve a SySAM problem, contact Sybase Technical Support for licenses generated from SPDC, or visit SMP if your license was generated from SMP. Have available as much information about your system, and the conditions leading to the problem, as possible.

To report a SySAM problem at SMP, go to <https://service.sap.com/support>, select **Help + Support**, then **Report a Product Error**.

Typical information for an unserved license:

- The Sybase product name, version, edition (if any)
- Optional product features that are enabled
- The product error log or debug log output (if the product produces one), or a screenshot or copy of error message text if not
- The SySAM 2-enabled product or feature licenses saved in the `$SYBASE/SYSAM-2_0/licenses` directory, which is on your local machine or in the product-specific license location

Typical information for a served license:

CHAPTER 12: Troubleshooting SySAM Errors

- The output from:
`sysam diag feature_name`
- The license server software version
- The license server debug log file
- The licenses that are saved in the `licenses` subdirectory of the `SYSAM-2_0` directory on the license server host machine

Determine the cause of problems and apply the recommended solution.

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.

This table 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*.

Problem	Solution
The installation program cannot start Adaptive Server.	<ul style="list-style-type: none"> • 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. • 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. • Verify that you are logged in as system administrator. You must log in as an administrator to start Adaptive Server. • Shut down Monitor Server before restarting Adaptive Server. • Adaptive Server may not start if it cannot find a SysSAM 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.
After upgrading Adaptive Server, srvbuild does not run.	Exit and restart srvbuild .
The installation program cannot connect to the upgraded Adaptive Server.	Exit and restart srvbuild .

See also

- *Stopping Adaptive Server After a Failure* on page 129
- *Recovering from a Failed Installation* on page 129

Error Logs for Installation Utilities

The information in the error logs may help determine the reason and possible solutions for problems with installation-related utility programs.

Locations of Error Logs for Installation-Related Utilities

Utility	Default Location and File Name
InstallAnywhere	\$SYBASE/log/ASE_Suite.log
srvbuildres	\$SYBASE/\$SYBASE_ASE/init/logs/srvbuildMMDD.NNN 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.
sqlupgraderes	<ul style="list-style-type: none"> • \$SYBASE/\$SYBASE_ASE/init/logs/sqlupgrademmdd.NNN • \$SYBASE/\$SYBASE_ASE/upgrade/upgrade.NNN – is a temporary file created by the upgrade process.

Error Logs for Sybase Servers

The information in the error logs may help to determine the reason and possible solution for an error message.

Locations and File Names of Error Logs for Sybase Servers

Server	Default Location and File Name
Adaptive Server	\$SYBASE/\$SYBASE_ASE/install/servername.log
Backup Server	\$SYBASE/\$SYBASE_ASE/install/servername.log
Monitor Server	\$SYBASE/\$SYBASE_ASE/install/servername.log
XP Server	\$SYBASE/\$SYBASE_ASE/install/servername.log

Troubleshooting Common Installation Problems

Identify the cause of installation problems and apply the recommended solutions.

Problem	Resolution
Cannot use X-Windows.	<p>If the setup and configuration utilities do not show correctly, you may need to adjust the resolution on your monitor.</p> <p>To change to a smaller font size, issue the following UNIX commands:</p> <pre>% 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 %</pre> <p>The installation utilities will now use approximately 25 percent less window space.</p>
Cannot eject the CD or DVD from the drive.	<p>If you cannot eject the CD from the drive, check:</p> <ul style="list-style-type: none"> • To see whether the disk drive path is the current directory (pwd) in a UNIX terminal window. If it is, change (cd) to another directory. • For sybhelp processes. If these processes exist, kill them using the UNIX kill command.
DISPLAY environment variable not set correctly.	<p>The resolution to the DISPLAY variable problem should instruct you to open Exceed (under Hummingbird Connectivity) and set DISPLAY equal to <code>hostname :b</code>, where <i>b</i> is the number in the parentheses following Exceed in the Exceed window.</p>

Problem	Resolution
Client not authorized to connect to server.	<p>This error message means the remote machine does not have permission to display the user interface on the local machine where you start working:</p> <pre>Xlib: connection to "host_name" refused by server Xlib: Client is not authorized to connect to Server xhost: unable to open display "host_name"</pre> <p>To correct the problem:</p> <ol style="list-style-type: none"> 1. Enter the following command at the UNIX prompt of your <i>local</i> machine, where <i>remote_machine</i> is the machine on which you are running the installer: <pre>xhost +remote_machine</pre> 2. Restart the installer.
Address already in use.	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. Adjust the shared memory value. Restart the installation or upgrade process.
The installer fails to start.	Restart the installer in debug mode. Set the environment variable LAX_DEBUG to true before running the installer.
Cannot start XP Server.	<p>You may receive the following message from XP Server when it is invoked by <code>xp_cmdshell</code> or some other extended stored procedure:</p> <pre>Msg 11018, Level 16, State 1: Procedure "xp_cmdshell", Line 2: XP Server must be up for ESP to execute. (return status = -6)</pre> <p>Verify that there is an XP Server entry in the Adaptive Server <i>sys.servers</i> table. If you created XP Server in a different srvbuild session than Adaptive Server and you did not specify a related Adaptive Server, srvbuild cannot update the <i>sys.servers</i> table. Verify that the XP server exists in the interfaces file or LDAP server.</p> <p>Use sp_addserver to add an entry to the <i>sys.servers</i> table.</p>

Problem	Resolution
Troubleshooting resource file installations.	<p>If you encounter problems during the build, configuration, or upgrade process while using the <code>srvbuild[res]</code>, <code>sqlloc[res]</code>, or <code>sqlupgrade[res]</code>, it may be that these utilities did not allow enough time for Adaptive Server to shut down properly. Set the <code>SYBSHUTWAIT</code> 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 before proceeding with the next task:</p> <pre>% setenv SYBSHUTWAIT 120</pre>

Stopping Adaptive Server After a Failure

If the installation or upgrade session fails after you start Adaptive Server for any reason, use the **shutdown** command.

1. Log on as "sa".
2. Shut down Adaptive Server using the **shutdown with nowait** command. This stops the Adaptive Server immediately, without waiting for currently executing SQL statements to finish:

```
1> shutdown with nowait
2> go
```
3. If the installation or upgrade session fails after you start Adaptive Server, try using Sybase Control Center to shut down the server. If Sybase Control Center cannot shut down the server, use the **shutdown** command.

Recovering from a Failed Installation

Review error messages and your Adaptive Server error log to determine the cause of an installation failure.

If the Installation Quits While You Are Configuring Adaptive Server

Perform specific steps if the installation suddenly quits.

1. View the contents of the log file generated by Adaptive Server.
2. Take any suggested actions to correct the problem. If the installation fails after the installation program:
 - Created any operating system files, such as the master device or system procedures device files, delete those files.
 - Starts the Adaptive Server that you are attempting to install, shut down that server.
3. Use Server Config to restart the configuration.

See also

- *Error Logs for Installation Utilities* on page 126

If Adaptive Server Fails the Preupgrade Eligibility Test

Examine the log file to determine why Adaptive Server is ineligible for an upgrade.

If Adaptive Server fails the preupgrade test, Server Configuration 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 in `$$SYBASE/$$SYBASE_ASE/init/logs`.

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

When an Upgrade Fails

If the upgrade process fails, the installation program displays error messages.

Once you start Adaptive Server on the new version, you cannot start Adaptive Server on the earlier version. Attempts to do so result in having to restore from backups.

If You Can Identify the Cause of the Upgrade Failure

If the error logs or messages clearly indicate the cause of failure, and you are confident that your databases have not been damaged, you can attempt to fix the problem and immediately re-run the upgrade.

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 `$$SYBASE/$$SYBASE_ASE/install/<servername>log`.

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. If needed, 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.

Restoring Databases After a Failed Upgrade

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 **Server Config**, but do not attempt to restart the upgrade session until you have restored the databases from backup. After restoration is complete, retry the upgrade.

Recovering from a Failed Upgrade

Upgrade issues may be due to a failure to upgrade an individual database, or a failure to complete configuration changes after all databases have been upgraded.

Use the output from the upgrade utility to correct the problem that caused the failure.

- If the upgrade failed because the process ran out of some resource, such as data or log space, locks, or auxiliary scan descriptors, add space to the database using the **alter database** command.

You may also correct other resource failures by changing the server's configuration using the **sp_configure** stored procedure.

- If an upgrade failure leaves a database offline, and the failure can be corrected only 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)
```

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

Note: This trace flag grants access only to user "sa"; "sa_role" does not work. If you have disabled the "sa" login, reenable it to get access using this method.

To restart a server that has not successfully upgraded, use:

```
online database failed_db_name
```

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

- If the failure occurs after all databases have been upgraded, or if a failure somehow causes the upgrade utility to fail, manually re-run the utility. After you diagnose and correct the failure, 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, 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
```

- Reinstall system stored procedures using the scripts from the new version:

```
isql -Usa -Ppassword -i $SYBASE/$SYBASE_ASE/scripts/installmaster
```

Note: When a server fails during an upgrade process, it attempts to upgrade the databases when it restarts. When the installation is fully recovered, re-run the upgrade utility to ensure that all post-upgrade configuration work is done. Then reinstall the system stored procedures.

- During the upgrade process, the log may become too full due to recorded catalog changes. If so, log in to the new server using **isql**, and issue:

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

- a) If **truncate_only** indicates that it could not truncate the log, retry the command as **no_log** instead.
 - b) If this does not truncate the log successfully, use the **alter database** command to extend the database as explained below.
- In some cases, pre-upgrade space estimations are insufficient for the data copy phase of upgrade. You see a message that there is insufficient space in the system segment for the upgrade. When this happens, the upgrade process stops responding, waiting for space to be provided. To increase the size of the database, log in to the new server using **isql**, and use **alter database**:

```
isql> alter database dbname on device_name = "2m"
```

alter database allows you to specify the size to alter the database with the unit specifier "m", or "M".

Note: There are certain upgrade failures from which the server cannot recover. For example, attempts to upgrade system tables to version 15.x 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 You Cannot Identify the Cause of an Upgrade Failure

Your attempts to upgrade may continue to fail.

1. If attempts to upgrade continue to fail, check the error log file to identify where and where the upgrade failed.

By default, the log file is located in `$SYBASE/$SYBASE_ASE/install/<servername>.log`.

2. Contact Sybase Technical Support with the information.

See also

- *Chapter 14, Obtaining Help and Additional Information* on page 135

Obtaining Help and Additional Information

Use the Sybase Getting Started CD, Product Documentation site, and online help to learn more about this product release.

- The Getting Started CD (or download) – contains release bulletins and installation guides in PDF format, and may contain other documents or updated information.
- Product Documentation at <http://sybooks.sybase.com/> – is an online version of Sybase documentation that you can access using a standard Web browser. You can browse documents online, or download them as PDFs. In addition to product documentation, the Web site also has links to EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, Community Forums/Newsgroups, and other resources.
- Online help in the product, if available.

To read or print PDF documents, you need Adobe Acrobat Reader, which is available as a free download from the *Adobe* Web site.

Note: A more recent release bulletin, with critical product or document information added after the product release, may be available from the Product Documentation Web site.

Technical Support

Get support for Sybase products.

If your organization has purchased a support contract for this product, then one or more of your colleagues is designated as an authorized support contact. If you have any questions, or if you need assistance during the installation process, ask a designated person to contact Sybase Technical Support or the Sybase subsidiary in your area.

Downloading Sybase EBFs and Maintenance Reports

Get EBFs and maintenance reports from the Sybase Web site or the SAP® Service Marketplace (SMP). The location you use depends on how you purchased the product.

- If you purchased the product directly from Sybase or from an authorized Sybase reseller:
 - a) Point your Web browser to <http://www.sybase.com/support>.
 - b) Select **Support > EBFs/Maintenance**.

CHAPTER 14: Obtaining Help and Additional Information

- c) If prompted, enter your MySybase user name and password.
- d) (Optional) Select a filter, a time frame, or both, and click **Go**.
- e) Select a product.

Padlock icons indicate that you do not have download authorization for certain EBF/Maintenance releases because you are not registered as an authorized support contact. If you have not registered, but have valid information provided by your Sybase representative or through your support contract, click **My Account** to add the “Technical Support Contact” role to your MySybase profile.

- f) Click the **Info** icon to display the EBF/Maintenance report, or click the product description to download the software.
- If you ordered your Sybase product under an SAP contract:
 - a) Point your browser to <http://service.sap.com/swdc>.
 - b) Select **Search for Software Downloads** and enter the name of your product. Click **Search**.

Sybase Product and Component Certifications

Certification reports verify Sybase product performance on a particular platform.

To find the latest information about certifications:

- For partner product certifications, go to http://www.sybase.com/detail_list?id=9784
- For platform certifications, go to <http://certification.sybase.com/ucr/search.do>

Creating a MySybase Profile

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

1. Go to <http://www.sybase.com/mysybase>.
2. Click **Register Now**.

Accessibility Features

Accessibility ensures access to electronic information for all users, including those with disabilities.

Documentation for Sybase products is available in an HTML version that is designed for accessibility.

Vision impaired users can navigate through the online document with an adaptive technology such as a screen reader, or view it with a screen enlarger.

Sybase HTML documentation has been tested for compliance with accessibility requirements of Section 508 of the U.S Rehabilitation Act. 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.

Note: You may 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 the Sybase Accessibility site: <http://www.sybase.com/products/accessibility>. The site includes links to information about Section 508 and W3C standards.

You may find additional information about accessibility features in the product documentation.

Index

A

- accounts, creating sybase user 30
- Adaptive Server
 - installing in command line mode 47
 - installing in silent mode, unattended 47
- Adaptive Server components 90
- Adaptive Server Configuration Guide for UNIX 3
- Adaptive Server editions 5
 - determining which edition you have 6
- Adaptive Server plug-in for Sybase Central 6
- additional steps for downgrading new features 107
- adjusting
 - shared memory parameters in operating system 32
- administration tasks 30
- applications
 - changing after renaming objects 73

B

- Backing up the server 82
- basic configuration 41
- basic downgrade steps using sp_downgrade 102
- bcp
 - migrating with 89
- binary overlay 82

C

- check out errors 113
- choosing
 - the correct license 113
- client applications 7
- command line installation
 - Replication Server 47
- command line options 48
- commands
 - showserver 57
- Component Integration Services
 - upgrading local and remote servers 68
- configuration after installation
 - basic settings 41
- configuration parameters 110
- conflicting object names 73

- creating
 - sybase user accounts 30

D

- data compression
 - described 6
- databases
 - image data for sample 62
 - increasing capacity 77
 - jpubs 63
 - optional 60
- determining
 - installation process 3
- determining the current edition 6
- determining whether select* should be changed in
 - views 100
- devices
 - increasing capacity 77
- directory changes 69
- directory locations 29
- directory services
 - described 6
- downgrading
 - Adaptive Server 101
 - handling new features 107
 - replicated databases, with 110
 - sp_downgrade_esd, using 104
 - sp_downgrade, using 102
 - steps to prepare beforehand 101
- downgrading Adaptive Server
 - encrypted columns 109
- downgrading job scheduler 108
- dsedit utility 87
- dump and load 15.7 ESD #2 database into 15.7 or 15.7 ESD #1 106

E

- Enabling the Java in the database in a high
 - availability system 96
- encrypted columns
 - described 6
 - downgrading Adaptive Server 109

Index

error

- information, finding 113
- license checkout 113

error log files

- troubleshooting and 129

F

fails preupgrade eligibility 130

features, separately licensed optional 19

file descriptors 31

finding

- license error information 113

finding compiled object errors before production 98

G

generating

- served licenses 22
- unserved licenses 22

grace period 16, 113

- cannot check out 113
- license continues to run after fixing 113

GUI installation 37

H

heartbeat 16

host IDs

- determining 17
- lmutil utility 17
- on Windows platforms 18
- with multiple network adapters 18

I

identifiers

- conflicting with reserved words 73

in-memory database

- described 6

increasing

- device and database capacity 77

installation

- Adaptive Server CD, mounting 37
- determining process via workflow 3
- overview 3
- warning, cannot find 113

installation steps for PC-Client 51

installer suddenly quits 129

installing

- command line 41
- with response file 46

installing Adaptive Server

- GUI mode 37

installpix script 62

installpubs2 script 62

installpubs3 script 62

interfaces file 87

J

Java in the database

- preparing before installing Adaptive Server 34

- upgrading 96

jpubs database

- installing 63

K

keywords 110

- reserved by Sybase 79

L

LDAP libraries

- environment variables 59
- location of 59

license

- check out errors 113

license deployment models

- comparison of 15
- recommendations 15

license error information, finding 113

license server

- cannot find licenses for optional features 113
- continues to run in grace period 113
- license checkout error 113
- starts in grace period 113

license servers 16

licenses 25

- recommendations for license models 15
- regenerating 24
- rehosting 24
- renewing 24

licensing 13

licensing optional features 113

M

- memory requirements 10
- methods of installation
 - Adaptive Server 37
- migrating
 - dump and load method 89
 - methods 88
- minimal configuration
 - Adaptive Server 41
- modifying individual licenses 25
- mount commands by platform 37
- mounting the Replication Server installation media
 - 37
- multipathing
 - described 34
 - setting up before installation 36

N

- new features
 - reverting to earlier version 107

O

- objects
 - conflicting names of 73
 - renaming 73
- optional databases 60
- optional features
 - described 6
 - errors 113
 - license server cannot find 113
- options for server editions
 - packages 6
- overview
 - installation 3
- overview of the upgrade process 65

P

- partitions
 - described 6
- PC-client
 - CD 7
- PC-Client
 - directory layout 29
 - product description 29

- system requirements for 10
 - planning the installation 13
 - platform
 - mount commands 37
 - post downgrade 110
 - preinstallation tasks 31
 - preupgrade tasks
 - before upgrade 70

Q

- quoted identifier errors 99
- quoted identifiers 74

R

- re-running the upgrade 131
- recommendations
 - license models 15
- recovering from a failed installation 129
- redundancy 16
- reenable Replication Server 87
- related products 90
- release bulletin 13
- removing
 - registry entries 48
- removing existing server 49
- renaming objects 73
- replacing the binary 90
- replicated databases
 - restoring 88
- Replication Server 87
 - draining logs for Adaptive Server 87
- reserved words 73
 - conflicts 74
- resource files 80
- response file 46, 52
- restoring from backup 131
- running in a grace period 113

S

- sample databases 61
- SAP Service Marketplace 13
- SCC
 - starting 55
- scripts
 - installpix 62
 - installpubs2 62

Index

- installpubs3 62
 - location 61
- security services
 - described 6
- separately licensed optional features 19
- served licenses 22
 - license servers 16
- server
 - component descriptions 26
 - directory layout 26
 - overview 5
- servers, license 16
- shared memory parameters
 - adjusting 32
- shared-memory segments
 - allocating 10
- showserver command 57
- silent installation 52
- source environment variables 37
- sp_checkreswords system procedure 73
- sp_downgrade system procedure 102
- sp_downgrade_esd 104
 - dump and load database 106
- sp_lmconfig
 - determining current edition 6
- SPDC
 - generating served licenses 22
 - generating unserved licenses 22
- sqlupgrade utility 79
- sqlupgraderes utility 80
- starting
 - Sybase Control Center 55
- stopping
 - Adaptive Server after a failure 129
 - servers 129
 - Sybase Control Center 55
- stored procedures
 - changing after renaming objects 73
- style conventions
 - fonts 1
- Sybase Central 6
- Sybase Control Center 6
 - starting 55
 - stopping 55
- SYBASE directory 37
- Sybase Product Download Center
 - login page for Sybase reseller sales 13
- sybase user
 - creating accounts 30

- sybsytemprocs database
 - increasing the size of 76
- SySAM licensing 13, 19
 - license servers 16
 - problems when Host ID is using virtual network adaptor address 113
 - troubleshooting 113
- sysmessages 110
- system administrator password
 - setting 60
- system requirements
 - PC-Client 10
 - Solaris 9
- system stored procedures 110
- system tables 110
 - changes due to upgrade 72

T

- tasks, administration 30
- Tivoli Storage Manager
 - described 6
- troubleshooting 113
 - an unrecognisable license 113
 - by using error log files 129
 - calling technical support 123
 - installation problems 127
 - installing license server components 113
 - license check out errors 113
 - licensing optional features 113
 - picking the correct license 113
 - running in a grace period 113
- troubleshooting Adaptive Server installations 125

U

- uninstalling
 - PC-Client 53
 - server 48
 - the old server 49
- unknown cause of failure 130, 133
- unrecognized tokens or objects 110
- unserved licenses 22
- upgrading
 - Adaptive Server 65
 - entire installations 65
 - failure 130
 - Java in the database 96
 - single databases 65

- tasks after you upgrade 84
- using bcp 89
- upgrading compatibility
 - versions 65
- upgrading compiled objects
 - dbcc upgrade_object 97
- upgrading high availability and cluster support 95
- upgrading Job Scheduler templates 91
- utilities 7
 - dsedit 87
 - sqlupgrade 79

- sqlupgraderes 80

V

- verifying connection to the server 58
- virtual network adaptor license problems 113

W

- warning, installation, cannot find 113

