



Installation Guide

**Adaptive Server[®] Enterprise
Cluster Edition 15.7 ESD #2**

IBM AIX

DOCUMENT ID: DC01086-01-1572-01

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	7
Managing Adaptive Server Enterprise with Sybase	
Control Center	7
Client Applications and Utilities	8
CHAPTER 4: System Requirements	11
Making Input/Output Completion Port API Available for	
Adaptive Server	13
Cluster Edition System Requirements for Using Private	
Interconnect Technology	14
System Requirements for Clients	14
CHAPTER 5: Planning Your Adaptive Server	
Installation	17
Adaptive Server Release Bulletin	17
Obtaining a License	17
Understanding License Generation	17
Determining Your License Deployment Model	18
Comparing License Deployment Models	19

- Fault Tolerance, License Grace, and Redundancy20
- Served License Deployment Model20
- Determining Host IDs21
 - Determining Host IDs for Machines with Multiple Network Adapters22
 - Using Alternate Host IDs on Windows22
- Knowing the Product License Type22
 - Optional Feature Licenses23
- Generating Licenses at SPDC23
 - Logging in to SPDC and Beginning License Generation24
 - Regenerating, Renewing, and Rehosting Licenses28
- Generating Licenses at SMP29
 - Generating License Keys30
- Installation Directory Contents and Layout30**
 - PC-Client Product Descriptions and Layout32
- Performing Administrative Tasks34**
 - Creating the Sybase User Account34
- Preparing to Install Adaptive Server35**
 - Adjusting the Operating System Shared Memory Parameters38
 - Managing Java in the Database During Installations and Upgrades38
 - Multipathing39
 - Setting Up Multipathing40
- CHAPTER 6: Installing Adaptive Server43**
 - Mounting the CD43**
 - Installing Adaptive Server with the GUI Wizard43**
 - Installing Adaptive Server in Console Mode47**
 - Installing Adaptive Server Using a Response File47**
 - Creating a Response File48

Installing in Silent Mode	48
Command Line Options	49
Uninstalling Adaptive Server	50
Removing an Existing Adaptive Server	50
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: Creating and Starting the Cluster	57
Differences Between Private and Shared Installations	57
Cluster Edition Preinstallation Checklist	58
Creating a Simulated Cluster on a Single Node	59
Before Creating a Cluster	60
Creating a Cluster with Storage Foundation for Sybase Cluster Edition	60
Worksheet for Creating a Cluster	60
Create a Cluster Using sybcluster	64
Configuring the Cluster Using sybcluster	64
Starting and Stopping a Cluster Using sybcluster	67
Verifying the Cluster Configuration	68
Starting the Cluster When the Operating System Starts	68
Cleaning Up After a Failed Cluster Installation	69
Auxiliary Servers	69
Configuring Backup Server Using sybcluster	70
Installing Job Scheduler	71

Configuring XP Server Using sybcluster 72

CHAPTER 10: Postinstallation Tasks75

Verifying That Servers are Running75

Verifying Your Connection to the Servers76

Testing the Installation and Network Connections76

 Creating a Test Environment77

 Configuring libtcl.cfg for LDAP77

 Adding a Server to the Directory Service78

Setting the System Administrator Password78

Installing Sample Databases79

 Default Devices for Sample Databases79

 Running the Database Scripts80

 Installing the interpubs Database81

 Installing the jpubs Database81

 Maintaining Sample Databases82

I/O Fencing82

 Preparing to Set Up I/O Fencing83

 Setting Up Raw Disks For Multi-Node Access
 and I/O Fencing83

 I/O Fencing via SCSI-3 PGR84

 Enabling I/O Fencing85

Creating Logical Volumes86

Manual Cluster Setup and Management86

 Setting the Environment87

 Raw Devices87

 The Cluster Input File87

 Manually Configuring the Cluster91

 Configuring the Interfaces File91

 Building the Master and Quorum Devices92

 System Stored Procedures93

 Creating Runserver Files93

 Setting Up Local System and Temporary
 Databases94

Automatic Takeover of a Cluster	95
Starting the Cluster	95
Starting the Cluster After a System Failure	96
Postconfiguration Tasks	96
Shutting Down the Cluster or the Instance	97
Shutting Down an Instance	97
Reconfiguring the Cluster	97
Enabling sybcluster and Sybase Control Center After a Manual Configuration	98
A sybcluster Sample Session Configuration Values	98
CHAPTER 11: Adaptive Server Upgrades	107
Upgrading Adaptive Server	108
Considerations for Component Integration Services ...	110
Preparing to Upgrade	110
Changes in Adaptive Server Directories	111
Checking System and Upgrade Requirements	112
Performing Preupgrade Tasks	113
Upgrade System Tables and Stored Procedures	115
The runserver File Location	115
Reserved Words	116
Running a Reserved Word Check	116
Addressing Reserved Words Conflicts	116
Quoted Identifiers	117
Upgrading to a Private Installation	117
The sybprocsdev device	121
Increasing the Size of the subsystemprocs Database	121
Increasing Device and Database Capacity for System Procedures	122
Upgrade to Adaptive Server 15.7 ESD #2	124
Upgrading from One Version of Adaptive Server Cluster Edition to Another	124

- Upgrade a Noncluster Version of Adaptive Server to
the Cluster Edition Using sybcluster128
- Verifying the Cluster Edition Server for Upgrade
.....128
- Upgrading the Cluster Edition Server Using an
Input File129
- Upgrading the Cluster Edition Server
Interactively130
- Responding to upgrade server Prompts130
- Installation of Version 15.7 ESD #2 over an Existing
15.7 or 15.7 ESD #1 Cluster Edition132
- Determining the Adaptive Server Version132
- Backing Up Adaptive Server132
- Installing Adaptive Server Using Binary Overlay
.....133
- Postupgrade Tasks134**
- Updating JAR and XML Files135
- Running the instmsgs.ebf Script136
- Restoring Functionality in Adaptive Server After
Upgrading136
- Reenabling Auditing137
- Updating Threshold Procedures for Audit
Segments137
- Reenabling Replication Server After the Dataserver
Upgrade138
- Restoring Replication on Replicate Databases
.....138
- Restoring Replication on Primary Databases139
- Migrate139**
- Migrating Data Using a Dump-and-Load Method140
- Migrating a Server Configured for High Availability140
- Using bcp to Migrate Data141
- Adaptive Server Components and Related Products ...141**
- Upgrading Job Scheduler141
- Upgrading Job Scheduler Templates142

Upgrading Java in the Database	146
Enabling Java in the Database in a High Availability System	146
Upgrading Backup Server	147
Upgrading Databases Using Dump and Load	147
How Compiled Objects Are Handled When Upgrading Adaptive Server	147
Finding Compiled Object Errors Before Production	148
CHAPTER 12: Adaptive Server Downgrades	151
Preparing to Downgrade Adaptive Server	151
Downgrading from Adaptive Server 15.7 ESD #2	152
Downgrading to an Earlier Version of Adaptive Server 15.7	155
Dumping a 15.7 ESD #2 Database to Load into a 15.7 or 15.7 ESD #1 Adaptive Server	156
Additional Considerations for New Features Used	157
Downgrading Job Scheduler	159
Post-Downgrade Tasks for Adaptive Server	159
CHAPTER 13: Troubleshooting SySAM Errors	161
Where to Find License Error Information	161
Problems and Solutions	161
Installing for the First Time	167
Possible Causes of a License Server Failure	167
Solution for Problem: Product Cannot Find License for Optional Features	168
Unserved License Deployment Models	168
Served License Deployment Models	170
Contacting SySAM Technical Support	171
CHAPTER 14: Troubleshoot the Server	173

Error Logs for Installation Utilities	174
Error Logs for Sybase Servers	174
Troubleshooting Common Installation Problems	175
Stopping Adaptive Server After a Failure	177
Recovering from a Failed Installation	177
If the Installation Quits While You Are Configuring Adaptive Server	177
If Adaptive Server Fails the Preupgrade Eligibility Test	178
If the Cluster Edition Configuration Fails	178
When an Upgrade Fails	178
If You Can Identify the Cause of the Upgrade Failure .	178
Restoring Databases After a Failed Upgrade	179
Re-running the Cluster Edition Upgrade	179
If You Cannot Identify the Cause of an Upgrade Failure	180
If Lack of Space Prevents an Upgrade	181
CHAPTER 15: Obtaining Help and Additional Information	183
Technical Support	183
Downloading Sybase EBFs and Maintenance Reports .	183
Sybase Product and Component Certifications	184
Creating a MySybase Profile	184
Accessibility Features	184
Index	187

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, configure, and use the Cluster Edition of Adaptive Server® Enterprise, use this installation guide, as well as the *Cluster Users 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 *Cluster Users Guide* provides descriptions of the features available in the Cluster Edition, as well as instructions for installing and configuring an Adaptive Server cluster system.

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. Configure the cluster server.
4. 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 50.

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.

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.

The Cluster Edition allows a single installation, a private-mode installation, and multiple installations of Adaptive Server to operate on multiple nodes as a shared-disk cluster environment with a single-system view. Each server runs on a separate node called a "server instance" or "instance." You can use a single configuration file to determine the configuration for all instances (a shared installation), or each instance can use a separate configuration file (a private installation).

CHAPTER 3: Adaptive Server Components

The Cluster Edition allows you to scale workloads using multiple physical and logical clusters. If an instance in a cluster fails, one or more of the still-running instances take over the workload of the failed instance. The cluster determines the instance to which each client connects. If one instance is overloaded, the cluster balances the load by shifting clients to other, more available instances.

The primary advantages of the Cluster Edition include:

- Improved availability – applications can continue to run as long as a single cluster member remains viable, even after several others have failed.
- Simple management – data is shared by all instances and does not require data repartitioning through cluster membership changes.

The Cluster Edition provides a distributed architecture. Internode communication takes place via a network interconnects, rather than via shared memory. Applications that minimize internode messaging yield the best performance in the Cluster Edition environment.

Single-System Presentation

The Cluster Edition supports a single-system presentation. That is, multiple instances that make up the cluster appear to clients as a single system. New client technologies allow clients to connect logically to a cluster while remaining physically connected to individual instances. This logical connection allows Adaptive Server to redirect clients to various instances in the cluster and to dynamically provide clients highly available failover data.

Workload Management

The Cluster Edition workload manager can customize workload management and failover for your business applications so that each performs most efficiently. The logical cluster provides individualized working environments.

Installation Options

The Cluster Edition offers these configurations:

- Shared installation – requires a shared file system created using a Network File System (NFS) or a cluster file system. A shared installation supports a single \$SYBASE installation directory, Adaptive Server home directory, and server configuration file.
- Private installation – does not use a shared file system. A private installation supports a separate \$SYBASE installation directory, Adaptive Server home directory, and server configuration file for each instance.

See the *Cluster Users Guide* for detailed information about the Cluster Edition.

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, DE, or CE 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.

CHAPTER 3: Adaptive Server Components

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.

Note: In a cluster environment, install the SCC remote command and control agent on each node of a cluster. You need only install the SCC management UI, however, in one place.

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

CHAPTER 3: Adaptive Server Components

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.

CHAPTER 4 **System Requirements**

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:

```
lslpp -L | grep devices.fcp.disk.rte
```

You see:

```
devices.fcp.disk.rte 6.1.0.3 COMMITTED FC SCSI CD-ROM, Disk,
```

Note: The Java Virtual Machine (JVM) and associated Adaptive Server support require at least 250MB of virtual memory space to start, and may need more depending on the requirements of the individual Java executions. For this reason, you may need to adjust the memory parameters when running Java to ensure that it has enough virtual memory space to start, and that both Adaptive Server and Java tasks can successfully coexist.

In particular, systems that have Adaptive Server total memory set to a value higher than 1.5GB (or 2.5GB on some enterprise servers) may experience problems, and you may have to reduce Adaptive Server total memory.

IBM AIX Requirements	Description
Platform	<ul style="list-style-type: none">• Processor 64-bit – Power6, Power5, Power4, and PowerPC 970• Kernel – 64-bit• Default user stack size – 148KB

IBM AIX Requirements	Description
Operating system requirements	Adaptive Server 15.7 Cluster Edition ESD #2 runs on AIX 6.1 and AIX 7. <ul style="list-style-type: none"> • Hardware – IBM Power, System p, System i, System p5, System i5, eServer p5, eServer i5, eServer pSeries, and BladeCenter blades • Operating system <ul style="list-style-type: none"> • AIX 6.1 runs based on POWER4, PPC970, POWER5 and POWER6 processors. • AIX OS IBM Power, System p, System i, System p5, System i5, eServer p5, eServer pSeries and eServer i5 servers, and BladeCenter blades. • Updates – Fileset <code>devices.fcp.disk.rte 6.1.0.3</code> or higher • Preferred RAM – at least 1G
Memory	<ul style="list-style-type: none"> • Minimum RAM required for Adaptive Server – 272MB • Minimum RAM per additional user – Approximately 364kB
Disk space	
Operating system patches	Adaptive Server 15.7 ESD #2 runs on AIX 6.1 and AIX 7.x. AIX 6.x and 7.x require a patch that fixes APAR IV10828, a known AIX issue. Contact IBM directly to obtain the patch.

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://www.ibm.com/java>.

See the *Clusters Users Guide* for hardware requirements for using Infiniband Interconnect on a production system. Sybase does not support file system devices when running on multiple nodes.

If you intend to run the cluster under Symantec Storage Foundation for Sybase Cluster Edition, see *Using the Cluster Edition with the Veritas Cluster Server*, in the *Clusters Users Guide*. Veritas Cluster Server for Sybase Cluster Edition is supported only on Solaris Sparc64 and Linux x86-64; it is not available for Solaris x86-64.

Database devices in the Cluster Edition must support SCSI-3 persistent group reservations (SCSI PGRs). Cluster Edition uses SCSI PGRs to guarantee data consistency during cluster membership changes. Sybase cannot guarantee data consistency on disk subsystems that do not support SCSI PGRs (such a configuration is supported for test and development environments that can tolerate the possibility of data corruption).

See also

- *Making Input/Output Completion Port API Available for Adaptive Server* on page 13

Making Input/Output Completion Port API Available for Adaptive Server

Configure the Input/Output completion port (IOCP) API so that Adaptive Server can start.

The IOCP API allows Adaptive Server to process multiple simultaneous asynchronous I/O operations in a scalable fashion. If you do not have IOCP version 6.1.5.0 or later configured as available on your system, Adaptive Server does not start when you run **dataserver**, and you see an error message such as:

```
% ./dataserver -v
exec(): 0509-036 Cannot load program dataserver because of the
following errors:
    0509-130 Symbol    resolution failed for dataserver because:
    0509-136 Symbol    CreateIoCompletionPort (number 579) is not
exported from dependent module /unix.
    0509-136 Symbol    GetMultipleCompletionStatus (number 580) is
not exported from dependent module /unix.
    0509-136 Symbol    PostQueuedCompletionStatus (number 581) is
not exported from dependent module /unix.
    0509-136 Symbol    ReadFile (number 582) is not exported from
dependent module /unix.
    0509-136 Symbol    WriteFile (number 583) is not exported from
dependent module /unix.
    0509-192 Examine  .loader section symbols with the 'dump -Tv'
command.
```

To check the status of IOCP on your machine, run the **lsdev** command at the UNIX prompt:

```
% lsdev -Cc iocp
```

If you installed IOCP and made it available, the system displays "Available" and provides additional information about the installation, such as:

```
% lsdev -Cc iocp
iocp0 Available  I/O Completion Ports
% lslpp -l bos.iocp.rte
Fileset                               Level  State           Description
Path: /usr/lib/objrepos
    bos.iocp.rte                       6.1.5.0  COMMITTED      I/O Completion
Ports API
%
```

If IOCP is not available, running the **lsdev** command returns "0 Defined," such as:

```
% lsdev -Cc iocp
iocp0 Defined  I/O Completion Ports
%
```

To make IOCP available for Adaptive Server:

CHAPTER 4: System Requirements

1. Log in to the IBM machine as root and issue:

```
# smitty iocp
```
2. Select **Change/Show Characteristics of I/O Completion Ports**.
3. Change the state of IOCP from Defined to Available.
4. Restart the machine.
5. Run **Isdev** again to confirm the updated status of IOCP.

Cluster Edition System Requirements for Using Private Interconnect Technology

The Cluster Edition supports only the UDP network protocol for private interconnects; do not use the TCP network protocol.

A private interconnect is a physical connection that allows internode communication, and is an essential component of a shared-disk cluster installation. A private interconnect can be a simple crossover cable with Ethernet, or it can be a complex solution. When you configure more than two nodes, a switch that enables high-speed communication between the nodes in the cluster is required.

To handle the amount of traffic created by contention, use scalable interconnect technology to connect nodes. The amount of traffic is directly proportional to the number of interinstance updates and transfers. Sybase recommends that you implement the highest bandwidth, lowest-latency interconnect available.

The Cluster Edition supports the current standards for interconnects. Sybase recommends that you research the available interconnects to find the one that works best for your site.

The Cluster Edition supports Infiniband in IP over IB (internet protocol over Infiniband) mode – the server uses a standard IP interface to communicate with the Infiniband interconnect. This is the simplest mode to configure.

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 `$$SYBASE / 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
The Server Installation Package	<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
<p>Shared Disk Cluster Administration Tools</p>	<ul style="list-style-type: none"> • sybcluster – interactive command line interface for configuring and managing the cluster installed in the \$SYBASE/\$SYBASE_ASE/bin directory. • Job Scheduler templates and utilities – predefined templates from which the database administrator can create and schedule useful, time-saving jobs. They are installed into the \$SYBASE/\$SYBASE_ASE/jobscheduler directory. • Interactive SQL – interactive command line interface for configuring and managing the cluster. It is installed into the SYBASE/DBISQL directory. • JRE – the Java Runtime Environment (JRE) is a runtime Java virtual machine used to execute Java-based programs such as Sybase Control Center. Installed into the \$SYBASE/shared/JRE-* directory. • The Cluster Edition Administration Utility is installed in \$SYBASE/SDCADMIN-15_0. • Sybase Central 6.x is a Java-based framework used by the system management tools. Installed into the \$SYBASE/shared/sybcentral600 directory. • Adaptive Server plug-in. Sybase Central plug-in that provides cluster configuration and complete administration functionality. Installed into the \$SYBASE/ASEP directory.
<p>Software Developer Kit (SDK)</p>	<p>Connectivity:</p> <ul style="list-style-type: none"> • Open Client™ (Client-Library, dlib) • Embedded SQL/COBOL 15.0 • XA interface for Adaptive Server <p>Installed into the DataAccess directory:</p> <ul style="list-style-type: none"> • 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.
<p>Sybase Control Center</p>	<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 \$SYBASE/SCC-3_2 directory.</p>
<p>Language modules</p>	<p>Installed into the \$SYBASE/locales and \$SYBASE_ASE/locales directories. 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.
Java Client Utility	Installed into the <code>jutils-3_0</code> directory and is a collection of Adaptive Server utility programs including ribo , a tool to trace TDS traffic between client program and Adaptive Server.
Java Database Connectivity (JDBC)	Installed into the <code>jConnect-7_0</code> directories. Provides a Java Database Connectivity (JDBC) driver for Adaptive Server.
Uninstaller	Installed into the <code>sybuninstall/ASESuite</code> directory.

Note: Sybase recommends that you:

- Install the ECDA DirectConnect option or MainframeConnect™ DirectConnect™ for z/OS, including DirectConnect Manager, into its own directory
- Do not install Sybase IQ 15.1 in a directory that includes the Adaptive Server Enterprise 15.7 ESD #2 Cluster Edition.

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)	Installed into the <code>OCS-15_0</code> directory: <ul style="list-style-type: none">• Open Client™ (Client Library, dlib)• ESQ/C• ESQ/COBOL• XA• Adaptive Server Enterprise extension module for Python• Adaptive Server Enterprise database driver for Perl• Adaptive Server Enterprise extension module for PHP

Product	Description
<p>DataAccess</p>	<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.
<p>Windows Cluster Server Admin utilities</p>	<p>Windows Cluster Administrator is a GUI tool for managing Microsoft Cluster Server (MSCS). Use it to create, modify and display group, resource, and cluster information. There is also an alternative command line tool called <code>Cluster.exe</code> for cluster administration.</p>
<p>PC-Client Administration Tools</p>	<ul style="list-style-type: none"> • sybcluster – interactive command line interface for configuring and managing the cluster installed in the <code>%SYBASE%\SDCADMIN-15_0\bin</code> directory. • Interactive SQL – interactive command line interface for configuring and managing the cluster. It is installed into the <code>%SYBASE%\DBISQL</code> directory. • Java Runtime Environment (JRE) – is a runtime Java virtual machine used to execute Java-based programs such as Sybase Control Center. Installed into the <code>%SYBASE%\Shared\JRE-6_0*</code> directory. • The Administration Utility is installed in <code>%SYBASE%\SDCADMIN-15_0</code>. • jutils-3_0 – a collection of Adaptive Server utility programs including ribo, a tool to trace TDS traffic between client program and Adaptive Server. • Sybase Central 6.x – is a Java-based framework used by the system management tools. Installed into the <code>Shared</code> directory. The <code>ASEPlugin.jar</code> file is located in <code>ASEP\lib</code>, while the scjview.exe executable to start Sybase Central with the ASE plug-in is located in <code>shared\Sybase Central 6.0.0\ [win32,win64]</code>. • Adaptive Server plug-in – Sybase Central plug-in that provides cluster configuration and administration functionality up to Adaptive Server 15.5. Installed into the <code>ASEP</code> directory. <p>Note: Sybase Central and the Adaptive Server plug-in do not support the new features and functionality introduced in Adaptive Server 15.7. For this reason, Sybase recommends that you instead use Sybase Control Center.</p>

Product	Description
Language Modules	Installed into the <code>locales</code> directory. Provides system messages and date/time formats.
Character Sets	Installed into the <code>charsets</code> directory. Provides character sets available for use with Adaptive Server.
jConnect	Installed into the <code>jConnect-7_0</code> directory. Provides a Java Database Connectivity (JDBC) driver for Adaptive Server.
Uninstaller	Installed into the <code>sybuninstall/PCClient</code> directory.

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. Ensure that there is at least 1GB of space available in the `temp` directory.
2. Manage permissions and privileges:
 - a) Ensure that the current shell has the appropriate inheritable privileges.

```
sudo lssecattr -p $$
    487528 eprivs= mprivs= iprivs=PV_KER_RAS lprivs=PV_ROOT
uprivs=
    > $SYBASE/ASE-15_0/bin/iofenceutil /dev/rhdisk2
    /dev/rhdisk2
```

If it is not fence capable, grant the current shell the appropriate inheritable privileges.

```
sudo setsecattr -p iprivs=+PV_KER_RAS $$
Then restart SCC agent
```

- b) Use the **ls -l** commands to verify paths and file permissions.
 - c) Use the **dd** utility to verify that the sybase account can read and write to the devices.
 - d) If you must change the permissions to install Adaptive Server, use **chmod** or **chown** to correct write permissions for `/dev/sg*` files. You can change the access permissions for these files to root only after restarting the machine.
 - e) Log in to the machine as the “sybase” user. Maintain consistent ownership and privileges for all files and directories. A single user—the Sybase system administrator with read, write, and execute permissions—should perform all installation, upgrade, and setup tasks.
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 account should own all devices and files, and must have permission to read and write all devices the cluster uses.
 - Ensure that all disk devices are accessible from all the nodes in the cluster.
 - Ensure that the account used to start the cluster has permissions to read and write to all of the disk devices.
 - Ensure that the cluster and the SCSI generic driver have write permission on the `/dev/sg*` files that correspond to configured database devices
 - The SCSI driver expects write access to `/dev/sg*` files for SCSI-3 PGR commands used in I/O fencing.

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

CHAPTER 5: Planning Your Adaptive Server Installation

5. Log in, as the "sybase" user, to the node on which you plan to install the first instance of the server.
6. If you do not have open administrative privileges, create a destination directory before running InstallAnywhere.
7. Decide where to install Adaptive Server:
 - Ensure there are no spaces in the path name of the directory.
 - If this is a shared installation, the \$SYBASE location must be on a shared file system, accessible from all cluster nodes using the same path.
 - If this is a private installation, install the Adaptive Server software on each node of the cluster.
 - Each instance in the cluster must have its own \$SYBASE directory.
 - The private installation mode does not use a Network File System (NFS) or cluster file system.
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.

The Cluster Edition requires your network to be configured for the nodes you plan to include in the cluster.

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

10. Create a \$HOME directory on the node on which you will run the installer.
11. Ensure that all nodes are running on the same operating system version.

The number of processors and the amount of memory can vary from node to node but the operating system version cannot.

12. Ensure that the quorum resides on its own device.
13. Create the local system temporary databases on a shared device using the Adaptive Server plug-in or **sybcluster**. Do this for each instance during the initial start-up of the cluster and later on whenever you add an instance to the cluster.

You can create or drop a local system temporary database from any instance, but you can access it only from the owning instance.

14. Ensure that all database devices, including quorum devices, are located on raw partitions. Do not use the Network File System (NFS).

Warning! Do not use file system devices for clusters. The Cluster Edition is not designed to run on a file system; mounting a nonclustered file system on multiple nodes immediately causes a corruption, leading to a total loss of the cluster and all of its databases. For this reason, Sybase does not support file system devices when running on multiple nodes.

15. Ensure that the raw partitions are accessible from each node using the same access path. Sybase recommends storage area network (SAN) connected devices.

Note: Local user temporary databases do not require shared storage and can use local file systems created as private devices—unlike local system temporary databases, which do require shared storage.

For test environments, use a single node or machine to run multiple instances of the Cluster Edition in a cluster configuration. When you do this, you must use the local file system (not NFS) or SAN storage for the database devices.

16. Ensure that all hardware nodes use Network Time Protocol (NTP) or a similar mechanism to synchronize clocks.
17. If you are using a shared installation, ensure that all Adaptive Server Enterprise software and configuration files (including the \$SYBASE directory, the interfaces file) are installed on a Network File System (NFS) or a clustered file system (CFS or GFS) that is accessible from each node in the cluster using the same access path.

If you are using a private installation, each node must have its own installation on a cluster file system.

18. Ensure that you have a high-speed network interconnect (for example, a gigabit Ethernet) providing a local network that connects all hardware nodes participating in the cluster.
19. Sybase recommends that each node in the cluster have two physically separate network interfaces; a primary network and a secondary network, both for cluster interconnect traffic.

The primary and secondary networks should be physically separated from each other, and are needed for security, fault-tolerance, and performance reasons. For fault-tolerance, the two network cards should be on different fabrics so that a cluster survives network failure.

20. Private interconnect fabrics should not contain links to any machines not participating in the cluster (that is, all cluster nodes should have their primary interconnect connected to the same switch, and that switch should not be connected to any other switches or routers).
21. 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. No special steps are necessary to configure for shared memory, however, you can modify the global ulimit information in `/etc/security/limits` with the following, in case the default hard limits for file sizes (*fsize*) and numbers of files (*nofiles*) may be too low:
 - **ulimit -d** – displays current process execution memory limits
 - **ulimit -d unlimited** – sets process execution memory limits to unlimited
2. You can also edit `/etc/security/limits` by changing the default data value to `"-1"`:

```
default:
    ...
    data = -1
    ...
```

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

Run the `installpcidb` script only on the first node and not on any other nodes in the cluster. In addition, create only the PCI device on the first node. A single instance of the PCI configuration is shared between the nodes of a cluster.

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.

CHAPTER 5: Planning Your Adaptive Server Installation

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
- lpfcdriver 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.

On AIX, Multiple Path I/O (MPIO) is installed and configured as part of BOS installation. No further configuration is required, but you can add, remove, reconfigure, enable, and disable devices (or device paths) using SMIT, Web-based System Manager, or the command line interface.

By default, MPIO is enabled on all disks and LUNs (logical unit numbers) that have this capability, which prevents third-party multipathing drivers (such as Veritas DMP or EMC

PowerPath) from managing the paths to such devices. To allow a third-party multipathing driver to manage multipathing instead of MPIO, install suitable Object Data Manager (ODM) definitions for the devices on the host.

To configure multiple paths for all LUNs, add all of the adapters and fibre-channel cables and run the `cfgmgr` command. You might have to run `cfgmgr` multiple times. See the IBM System Storage® Multipath Subsystem Device Driver User's Guide for details about using the `cfgmgr` command.

The `/dev/rhdiskN` devices are persistent. These are the device names that should be used to access the multipathed devices.

The following commands help manage MPIO paths:

- `lspath` – displays multipath device configuration
- `mkpath` – adds a path for a multipath device
- `chpath` – configures multipath device
- `rmpath` – removes a path for a multipath device
- `lsattr` – displays device configuration
- `lsdev` – displays system devices
- `mkdev` – creates and enables system devices
- `chdev` – configures system devices
- `rmdev` – remove and disables system devices
-

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.

Log in as "sybase" and issue:

```
mount -v 'cdrfs' -r device_name /mnt/cdrom
```

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

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,

CHAPTER 6: Installing Adaptive Server

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.

If you are installing the Cluster Edition in private installation mode, Sybase recommends that you install Adaptive Server on each node in the cluster in its own directory, and use the installer for each instance in the cluster. This ensures that environment variables, soft links, and so on are set correctly for each instance.

Sybase recommends installing products as a system administrator; however, root permissions are not required to run the installer. The installer creates the target directory (if necessary) and installs the selected components into that directory. At the end of the installation, you can verify the product installation. You may need to perform additional configuration procedures before you can use some products.

Warning! If you install the Cluster Edition into the same directory as earlier versions of the same components, the Cluster Edition overwrites the older versions. Installing other products on top of the Cluster Edition is likely to result in one or more of the products not working correctly.

InstallAnywhere supports Section 508 functionality, with these limitations in the wizard:

- **Choose Install Set** – the installer does not support the use of a keyboard shortcut to select the type of installation. You must use Tab to change the focus and the spacebar to choose.
- **Choose Product Features** – does not allow you to choose an option using a keystroke. You must use a mouse to choose features.

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.

Log in as "sybase" and issue:

```
mount -v 'cdrfs' -r device_name /mnt/cdrom
```

4. Launch the installer:

```
cd /device_name
./setup.bin
```

where:

- *device_name* 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 11, Adaptive Server Upgrades* on page 107.

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 installation type for the Adaptive Server suite. If you are updating a preexisting server, you do not see this step. The available types to be selected are different for different platforms.

- Licensed copy of Adaptive Server Enterprise Cluster Edition Suite – choose this if you have a licensed copy of Adaptive Server.
- Evaluation Edition of Adaptive Server Enterprise Cluster Edition Suite – choose this if you would like to evaluate Adaptive Server. If you choose the Evaluation Edition, the software works for 30 days from the original installation date.

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

11. If you are installing a licensed copy of Adaptive Server Enterprise Cluster Edition, select the license type:

CHAPTER 6: Installing Adaptive Server

- **CP** – CPU license
 - **SF** – standby CPU license
 - **DT** – development and test license
 - **AC** – OEM application deployment CPU license
 - **BC** – application deployment standby CPU license
 - **Unknown** – without a license
12. Configuring your server for e-mail notification enables designated users to receive information about license management events requiring attention. Provide the following information:
- SMTP server host name
 - SMTP server port number
 - E-mail return address
 - Recipient e-mail addresses
 - Message severity that triggers e-mail messages
13. 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.
14. Select **enable** or **disable** to remember the Adaptive Server password, then click **Next** to specify whether to have the ASE plug-in remember passwords.
15. 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.
16. Enable and order the security login module.
17. Enter the RMI port.
18. Choose whether to enable SCC shared-disk mode. If you do, enter the SCC instance name.
19. 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
When connecting to sybcluster, use this same password.
20. 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 ...
```
-

21. Source the environment variables needed for Adaptive Server products by running the `SYBASE . csh` script file after exiting the installer.

Next

You have installed Adaptive Server and its related products. See the *Cluster Users Guide* to set up your cluster, *Sybase Control Center Installation Guide* to install the SCC management UI on your server, and the *System Administration Guide* for more advanced topics.

See also

- *Managing Adaptive Server Enterprise with Sybase Control Center* on page 7

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.

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.

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:

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

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) agent for Adaptive Server Enterprise Cluster Edition enables the distributed management of the cluster. Set the \$SYBASE environment variables and start an SCC agent on each node to host instances in the cluster.

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
```


You can configure and start the cluster server after successfully installing Adaptive Server.

1. Read *Before Creating a Cluster* on page 60.
2. Set the \$SYBASE environment variable, and start the Sybase Control Center remote command and control agent on each node of the cluster.
3. Configure the cluster using Sybase Control Center, or the **sybcluster** utility. Sybase recommends that you use either option to create and manage your shared-disk cluster. However, you can also manually configure and manage the cluster.
If you have numerous clusters to set up, you can save your configuration parameters to a file, and use sybcluster to import them.
4. (Optional) Configure the auxiliary servers: XP Server, Backup Server, and Job Scheduler.

If the installation or start-up fails, see *Cleaning Up After a Failed Installation* on page 69.

See also

- *Manual Cluster Setup and Management* on page 86

Differences Between Private and Shared Installations

The installation process differs between private and shared installations.

Regardless of installation type, each instance of Adaptive Server in the cluster shares:

- All databases and database devices (for example, all instances share the same master database).
- A quorum device, which coordinates instances with the cluster.

Shared Installation	Private Installation
<p>Each instance of Adaptive Server in the cluster shares:</p> <ul style="list-style-type: none"> • A common \$SYBASE directory • The interfaces file, unless LDAP is used for server lookup • All configuration files, including the cluster input file • All server binaries (for example, <code>datasever</code>) and all scripts (for example, <code>installmaster</code>) 	<p>Each instance of Adaptive Server in the cluster maintains its own:</p> <ul style="list-style-type: none"> • \$SYBASE directory • The interfaces file, unless LDAP is used for server lookup • Server configuration file

See the *Cluster Users Guide*.

Cluster Edition Preinstallation Checklist

Before you begin the installation process, identify your settings and determine how to set up your environment.

Use this checklist to determine the configuration for an installation of Adaptive Server Cluster Edition.

Configuration	Value
Will the cluster installation support a shared (NFS), or a private installation?	
Will the cluster use LDAP (recommended for private installations), or interfaces files for server lookup information?	
Will the cluster support Java in the database via a third-party JVM?	
Will the cluster supports Veritas Cluster Server (VCS) for managing application servers in the clustered environment?	
Will the cluster use a single or multiple Backup Servers?	
Where will the \$SYBASE installation directory reside? (Private installations require a separate \$SYBASE installation directory for each instance.)	
How many instances will be in the cluster?	
What will the names of the clusters in the instance be?	

Configuration	Value
What is the Domain Name Service for the network on which the cluster will run? (The Sybase Control Center agent and the sybcluster utility do not function properly if the DNS entries are entered incorrectly for each node.)	
How many raw devices will each database device use? (For the Cluster Edition, you must create all devices, database devices, and quorum devices as raw devices on shared disks.)	
How many agents will there be in the cluster? (Sybase recommends that you use one node for each of the instances in the cluster, with two or more engines per instance.)	
What raw devices will each database device use? (You must create all devices, database devices, and quorum devices as raw devices on shared disks.)	
<p>What is the range of port numbers each instance will use to exchange messages with other instances via the private interconnect? (Choose unique port numbers that are not in use by other applications.)</p> <p>Local user temporary databases do not require shared storage and can use local file systems created as private devices. However, local system temporary databases you create during the cluster configuration can use only shared disks. See <i>Using Temporary Databases</i> in the <i>Cluster Users Guide</i>.</p> <hr/> <p>Note: The Adaptive Server plug-in, Sybase Control Center, and sybcluster supply default values if your system does not have an existing standard.</p> <hr/>	
What is the query or listening port number for each instance? (Choose unique port numbers that are not in use by other applications.)	
What are the IP addresses or network names of the private primary and secondary network cards on each node? (The only network protocol currently supported is UDP.)	

Creating a Simulated Cluster on a Single Node

Test the server configuration and installation.

Sybase recommends that you configure each instance on a different node, that is, one node for each instance. However, in a testing environment, you can create a simulated cluster on a single node with all instances running on that node.

1. For best performance, make sure the number of engines in all instances running on a single node do not exceed the number of CPUs on that node.

2. Verify the value for **runnable process search count**. Sybase recommends a value of 3 (the default). You can set this using the **sp_configure** stored procedure.

Before Creating a Cluster

After you install Sybase Control Center, you can create a cluster.

The **sybcluster** utility does not create a `run_server` file when it creates a cluster. You must start the cluster and each instance using **sybcluster** or Sybase Control Center. You cannot start the cluster from the command line using `run_server` files.

Creating a Cluster with Storage Foundation for Sybase Cluster Edition

There are steps you must perform for the Veritas Storage Foundation, and for your operating system.

Before you create a cluster with Storage Foundation for Sybase Cluster Edition, review *Using the Cluster Edition with the Veritas Cluster Server* in the *Cluster Users Guide*.

Worksheet for Creating a Cluster

Before creating a cluster, gather your cluster information.

Table 4. Configuration Values for Creating a Cluster

Parameter and Description (Default Values)	Your Information
Cluster name: <ul style="list-style-type: none"> • Number of instances (4) • Number of agents (4) 	
Is the cluster to be configured using a private \$SYBASE installation? (N)	
The page size in kilobytes (2KB)	
Full path and name for the master device: <ul style="list-style-type: none"> • Size of master device (30MB) • Size of master database (13MB) 	

Parameter and Description (Default Values)	Your Information
Sybase system procedure device – full path and name for <code>sybsystemprocs</code> device: <ul style="list-style-type: none"> • Size of sybsystemprocs device (152MB) • Size of sybsystemprocs database (152MB; minimum 140MB) 	
System database device – full path and name of system database device: <ul style="list-style-type: none"> • Size of system database device (6MB) • Size of system database (6MB) • Does this cluster have a secondary network? (Y) 	
Quorum device: <ul style="list-style-type: none"> • Full path and name for the quorum device • Trace flags 	
(Optional) PCI device: <ul style="list-style-type: none"> • Full path to the PCI database device • Size of the PCI database device (24MB) • Size of the PCI database (24MB) 	

Table 5. Verifying Locations for Shared Installations

Parameter (Default Values)	Value
The location of the home directory (<code>\$SYBASE</code>)	
Full path to environment variable script (<code>\$SYBASE/SYBASE.sh</code>)	
Path to <code>\$SYBASE_ASE (ASE-15_0)</code> .	
Path to interfaces file directory (<code>\$SYBASE</code>).	
Path to the dataserver configuration file (<code>\$SYBASE/cluster_name.cfg</code>)	

Table 6. Information for Shared Installations

Information	Instance 1	Instance 2	Instance 3	Instance 4
Node name				
Instance name				
Query port number for instance				
Primary protocol address for instance				
Secondary protocol address for instance				

Table 7. Information for Private Installations

Information	Instance 1	Instance 2	Instance 3	Instance 4
Node name				
Instance name				
Full path for \$SYBASE installation directory for instance				
Full path to environmental shell script for instance				
Full path to the server configuration file for instance				

Table 8. Local System Temporary Database Information

Parameter	Instance 1	Instance 2	Instance 3	Instance 4
Local system temporary database device name				
Path to local system temporary database device				
Local system temporary database device size (MB)				
Local system temporary database name				
Local system temporary database size (MB)				

Note: For the local system temporary database device name, enter the name of the Adaptive Server database device for the local system temporary database. The device for a local system temporary database must be a shared disk.

The Veritas Cluster Server support options appear only if VCS is supported on your system.

Table 9. Veritas Cluster Server Support

Information (Default Value)	Your Information
Integrate the Cluster Edition server with the VCS? (Y)	
The path to the interfaces file on the cluster; appears only if LDAP is supported	
Check if this device (master device, system procedure device, and system database device) is managed by the Veritas Cluster File system or Veritas Volume Manager (Y)	

Information (Default Value)	Your Information
Check each cluster device (master device, system procedure device, and system database device) for I/O fencing capability (Y)	

Create a Cluster Using sybcluster

Use **sybcluster** to create and configure a shared-disk cluster.

Before using **sybcluster**, read *Before You Begin the Installation* on page 58.

To view a complete **sybcluster** session that includes answers to all the prompts, see *Sample sybcluster Session* on page 102.

To create logical clusters, see *Managing the Workload* in the *Cluster Users Guide*.

Configuring the Cluster Using sybcluster

Configure the cluster using the information you gathered for your worksheet.

See the *Clusters Users Guide* for complete syntax and usage information for **sybcluster**.

1. Start **sybcluster** using **uafadmin**, the Unified Agent administrative login:

```
sybcluster -U uafadmin -P password -C clustername -F "hostname:
9999"
```

2. Run **create cluster**.

sybcluster prompts you for all necessary information one parameter at a time. If there is a default value, **sybcluster** displays it in the command prompt. To accept the default value, press **Enter**. Otherwise, enter the correct value, then press **Enter**.

3. Complete the fields in **sybcluster** using the information from *Worksheet for Creating a Cluster* on page 60.

Fields	Description
Cluster name	The name of the cluster to create if you did not set the default cluster on the command line.
Instances	The maximum number of instances to create for the cluster.
SCC agents	<ul style="list-style-type: none"> • The number of agents in the cluster – sybcluster displays a list of available host machines. In this list are all nodes specified by the -F and -d parameters on the sybcluster command line with a configured and running a Sybase Control Center remote command and control agent. Select agents from this list. • The number representing the cluster agent – sybcluster assigns numbers to the remaining agents in the cluster based on this number.
Configuration type	Is the cluster using private installation mode? The default value is no [N]; the cluster is to be configured for a shared installation.

Fields	Description
Quorum device	The full path for the quorum device (such as <code>/dev/raw/raw11</code>).
Trace flags	Any required trace flags.
Page size	The page size of the master database, in kilobytes.
Master device	<ul style="list-style-type: none"> The full path for the master device. For example, <code>/dev/raw/raw12</code>. The size of the master device. The size of the master database.
Sybase system procedure device	<ul style="list-style-type: none"> The full path to the system procedure database device, sysprocsdev. For example: <code>/dev/raw/raw13</code>. The size of the system procedure database device. The size of the system procedure database.
System database device	<ul style="list-style-type: none"> The full path to the Sybase system database device, systemdbdev. For example: <code>/dev/raw/raw14</code>. The size of the system database device. The size of the system database.
PCI device	<p>(Optional) Enable the pluggable component interface (PCI) to support Java in the database? If you enter Y, enter:</p> <ul style="list-style-type: none"> The path to the PCI database device path. For example: <code>/dev/raw/raw20</code>. The PCI database device size (24MB). The PCI database size (24MB).
Secondary network	<p>Does this cluster have a secondary network? (Y).</p> <ul style="list-style-type: none"> If you enter Y, sybcluster uses the default port number of 15100 as the starting port number for the range of port numbers required so that the instances can exchange messages, and reserves that number of ports, starting with the default value. <hr/> <p>Note: Adaptive Server uses several sockets for each interconnect. The number of ports required for each instance is 5, multiplied by the maximum number of instances.</p> <hr/> <ul style="list-style-type: none"> If you enter N, sybcluster asks for the beginning port number, and calculates the number of additional ports required and reserves that number of ports. The default value is 15100.

4. Skip this step if the cluster uses a private installation. For shared installations, provide:
- The `$SYBASE` home directory (such as `/remote/var/sybase`).
 - The full path to the “.sh” or “.csh” environment shell script (such as `/remote/var/sybase/SYBASE.sh` or `/remote/var/sybase/SYBASE.csh`).

CHAPTER 9: Creating and Starting the Cluster

- The Adaptive Server home directory (such as `/remote/var/sybase/ASE-15_0`).
- The directory containing the interfaces file (such as `/remote/var/sybase`). **sybcluster** adds the correct cluster and instance information during configuration.

Note: Ensure that the interfaces file does not already contain the cluster or instance information.

- The full path to the dataserver configuration file (such as `/remote/var/sybase/mycluster.cfg`). **sybcluster** must be able to find this file during configuration.
5. **sybcluster** displays the node name and prompts for values for each instance, one node at a time. If you are using:

- A shared configuration, provide:
 - Instance name
 - Query port number for the instance. Ensure that it is available and not in use by other applications.
 - The primary protocol address for the instance. For example, 10.0.1.1.
 - The secondary protocol address for the instance (if you answered Y to a secondary network). For example, 10.0.1.2.
- A private configuration, provide:
 - Instance name
 - The `$SYBASE` home directory
 - The environmental shell script path
 - The Adaptive Server home directory
 - The full path to the server configuration file

Note: The path for the server configuration file can be the same or different for each instance.

- Interfaces files for server lookup, provide:
 - Instance name.
 - The interface file query port number
 - The primary protocol address
 - The secondary protocol address
- A local system temporary database device, provide:
 - The name of the Adaptive Server database device for the local system temporary (LST) database.

Note: You must create local system temporary database on a shared disk.

- The LST device path. For example, `/dev/raw/raw15`.
- The LST device size.

Note: If you place all LST databases on the same device, the device size must be large enough for all LST databases.

- The LST database name.
- The LST database size.

Note: **sybcluster** asks if you want to add another instance. If you enter Y, **sybcluster** repeats this step for the next instance.

6. sybcluster asks if you want to save the cluster configuration.

If you enter Y, **sybcluster** stores the configuration in an XML file that you can edit and replay using the **sybcluster** command:

```
create cluster cluster_name file file_name
```

7. sybcluster asks if you want to create the cluster.

If you enter Y, the configuration process begins, and **sybcluster**:

- Checks to see if VCS is supported on your system. If so, **sybcluster** asks whether to integrate the cluster database with VCS [Y]. If yes, **sybcluster** prompts for:
 - The path to the interfaces file for each instance, unless LDAP is supported.
 - Whether to check if the master database, system procedure device, and system database devices are managed by the VCS Volume Manager or Veritas Cluster File System [Y]. If yes, **sybcluster** reports status for each device and asks if you want to continue.
- If you answer no, **sybcluster** prompts for the path to the interfaces file for each instance, unless LDAP is supported.
- sybcluster** asks if you want to check if each cluster device supports I/O fencing. If you answer yes, **sybcluster** checks all devices and reports the I/O fencing capability of each one.

Starting and Stopping a Cluster Using sybcluster

The **sybcluster** utility allows you to start and stop a cluster.

1. Start sybcluster if it is not already running:

```
sybcluster -U uafadmin -P password -C cluster_name -F  
"node_name[:port_num]  
[,node_name[:port_num]]..."
```

This statement identifies the default cluster and the Sybase Control Center agents on each node in the cluster. If you do not enter this information on the **sybcluster** command line, you can enter it in the following step. See the *Clusters Users Guide*.

This example starts “mycluster” on “blade1”, “blade2”, and “blade3”:

```
sybcluster -U uafadmin -P -C mycluster  
-F "blade1,blade2,blade3"
```

2. Connect to the cluster:

```
connect
```

3. Start the cluster:

```
start cluster
```

Next

To stop a cluster, enter:

```
shutdown cluster
```

Adaptive Server waits for all statements and transactions to complete before shutting down the cluster.

Verifying the Cluster Configuration

Once **sybcluster** is running, and you connect to and start the cluster, you can verify that both the cluster and cluster configuration.

1. Verify that the cluster is running:

```
show cluster status
```

```
lunch> show cluster status
INFO - Listening for the cluster heartbeat. This may take a
minute. Please wait
... (lunch::AseProbe:434)
```

Id	Name	Node	State	Heartbeat
1	burger	tigger.sybase.com	Up	Yes
2	fries	tigger.sybase.com	Up	Yes

```
lunch>
```

2. Verify the cluster configuration:

```
show cluster config
```

Starting the Cluster When the Operating System Starts

You can configure the host system to start the Adaptive Server cluster automatically; the operating system starts by running the Sybase Control Center agent on the host system and executing the **sybcluster** commands using a shell script.

1. Start Sybase Control Center, and verify that it started successfully.
2. Use **sybcluster -i** to pass a command file (with a name such as `asecel_startup`) that provides instructions to start the instances on your system, such as:

```
connect to asecel5
start instance asecel
quit
```

3. To start the cluster, use a **sybcluster** command such as:

```
sybcluster -U uafadmin -P -F host1:9999,host2:9999 -i
asecel_startup
```

Cleaning Up After a Failed Cluster Installation

If your cluster installation is unsuccessful, remove any residual files or operating system processes to prevent unexpected errors in subsequent installations.

1. Terminate **srvbuildres** or **dataserver** processes if either are left running.
2. Stop the SCC agents on all nodes.
3. Remove `$$SYBASE/SCC-3_2/instances/instance_name`.
4. Remove any entries for the cluster or cluster instances from the `interfaces` file.
5. Remove the `cluster_name.cfg` file if it exists after the last attempt.
6. Restart the Sybase Control Center agent.

Auxiliary Servers

You can configure the auxiliary servers such as Backup Server, XP Server, and Monitor Server with the **sybcluster** utility.

Cluster Edition version 15.5 and later allows a cluster to use multiple Backup Servers with one of these methods:

- Dedicated method – each instance is assigned a specific Backup Server.
- Round-robin method – at the time of the **dump** or **load** command, the Cluster Edition assigns the instance a Backup Server from a group according to availability.
- A single Backup Server named SYB_BACKUP.

See *Using Backup Server in a Clustered Environment* in the *Cluster Users Guide*.

Backup Servers in a shared-disk cluster environment are available across all nodes of the cluster as a single Backup Server, or to operate on only one instance in the cluster. You can configure the Backup Server on any node of the cluster, and you can provide additional hosts and port numbers on which Backup Server can start if its current node goes down.

Execute the **dump** and **load** commands from any node in the cluster. Once issued, the local instance services the **dump** and **load commands** and routes them to the cluster's Backup Server. The instances attempt to connect to Backup Server using the order specified in the `interfaces` file. If Backup Server is not running on any of the nodes in the cluster, the instance on which you issued the **dump** or **load** command starts Backup Server, if Backup Server is configured to run on this node in the `interfaces` file. If multiple Backup Servers are configured, the Backup Server is assigned based on whether the configuration is in dedicated or round-robin mode.

You can install and start Backup Server using **sybcluster**, which prompts you for required information.

The Backup Server binary (`$SYBASE/$SYBASE_ASE/bin/backupserver`) is installed when you install the Cluster Edition.

Configuring Backup Server Using sybcluster

You can configure Backup Server with **sybcluster**.

1. Start **sybcluster**.

For example, to start **sybcluster** and specify Sybase Control Center agents on the nodes "blade1," "blade2," and "blade3," enter:

```
sybcluster -U uafadmin -P password -F  
"blade1:1234,blade2:2345,blade3:3456"
```

Note: Sybase Control Center agents must be running on all nodes in the cluster to verify that the selected ports are available.

2. Connect to the cluster. For example, to connect to "mycluster," enter:

```
connect to mycluster
```

3. If the cluster is not running, start it. Enter:

```
start cluster
```

4. To create one or multiple Backup Servers for the cluster, enter:

```
create backupserver
```

5. **sybcluster** prompts you:

```
Do you want to create multiple Backup Servers?
```

a) If you answer Y, **sybcluster** prompts for:

- A routing policy for the multiple Backup Servers. The choices are:
 - 1 – dedicated
 - 2 – round robin
- Each Backup Server name. The default value is "*cluster_name*_BS." For example, "mycluster_BS."
- The path to each Backup Server log file.
- The listening port for each Backup Server.

b) If you answer N, **sybcluster** prompts you to configure Backup Server on the hosts, one by one.

- The Backup Server name. The default value is "*cluster_name*_BS." For example, "mycluster_BS."
- The Backup Server listening port for each node in the cluster.

Sybase recommends that you configure Backup Server on all nodes so that it can be started on any node. In some cases where Backup Server is not configured on all nodes, Adaptive Server cannot start Backup Server when Backup Server is not running. This happens when the **dump** command is initiated on a node that does not have Backup Server configured.

You can add or remove Backup Server from nodes at any time.

Installing Job Scheduler

All instances in the cluster share a single Job Scheduler. Set up Job Scheduler so that, in the event the instance on which it is running fails, Job Scheduler can fail over to another node.

1. Create a device called **sybmgmtdev** with a size of at least 90MB on a shared raw device that is accessible to all instances in the cluster.
2. Run the `installjsdb` script:

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

Note: You must have the directory with the location of the `isql` executable (`$SYBASE/$SYBASE_OCS/bin`) in your path.

The `installjsdb` script looks for the `sybmgmtdb` database. If it exists, the script creates Job Scheduler tables and stored procedures. If it does not exist, the script looks for a `sybmgmtdev` device on which to create the `sybmgmtdb` database, tables, and stored procedures.

Note: If the `installjsdb` script finds neither a **sybmgmtdev** device nor a `sybmgmtdb` database, it creates a `sybmgmtdb` database on the `master` device. Sybase strongly recommends that you remove the `sybmgmtdb` database from the `master` device to make recovery easier in the case of a disk failure.

3. Create a directory services entry for the JSAGENT in the interfaces file using **dsccp**, **dsedit**, or a text editor as appropriate. Sybase suggests that you name the entry “`clustername_JSAGENT`”.

To enable high availability failover, the JSAGENT entry must contain master and query rows for each node in the cluster. For example, to add a JSAGENT entry for the cluster “mycluster” with two nodes, the syntax might be:

```
mycluster_JSAGENT
  master tcp /dev/tcp node_name1 17780
  query tcp /dev/tcp node_name1 17780
  master tcp /dev/tcp node_name2 16780
  query tcp /dev/tcp node_name2 16780
```

The host name must match the name returned by the **uname -n** command executed at the UNIX prompt. For example, on host “myxml1,” **uname -n** returns the value “myxml1.sybase.com,” and on host “myxml2,” **uname -n** returns the value “myxml2.sybase.com.” The correct entry for JSAGENT is:

```
mycluster_JSAGENT
  master tcp /dev/tcp myxml1.sybase.com 17780
  query tcp /dev/tcp myxml1.sybase.com 17780
  master tcp /dev/tcp myxml2.sybase.com 16780
  query tcp /dev/tcp myxml2.sybase.com 16780
```

CHAPTER 9: Creating and Starting the Cluster

The host name for the JSAGENT entry must be identical to the host name of the instances. For example, instance 1 has an entry with "asekernel1.sybase.com" and instance 2 has "asekernel2":

```
INSTANCE_1
  master tcp /dev/tcp asekernel1.sybase.com 17700
  query tcp /dev/tcp asekernel1.sybase.com 17700
INSTANCE_2
  master tcp /dev/tcp asekernel2 16700
  query tcp /dev/tcp asekernel2 16700
```

The correct entry for JSAGENT should have:

```
mycluster_JSAGENT
  master tcp /dev/tcp asekernel1.sybase.com 17780
  query tcp /dev/tcp asekernel1.sybase.com 17780
  master tcp /dev/tcp asekernel2 16780
  query tcp /dev/tcp asekernel2 16780
```

Note: You must specify a port not currently in use.

See *Directory Services* in the *System Administration Guide*.

4. Using **sp_addserver**, create an entry in the `sys.servers` table for the cluster. For example:

```
sp_addserver SYB_JSAGENT, null, mycluster_JSAGENT
```

See the *Reference Manual: Commands* for more information on **sp_addserver**.

5. Enable Job Scheduler:

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

6. To start Job Scheduler, you can either restart the server, or execute:

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

7. To determine the instance on which Job Scheduler is running, query the global variable `@@jsinstanceid`:

```
select @@jsinstanceid
go
```

Configuring XP Server Using sybcluster

Configure an XP Server for each instance in the cluster using the **sybcluster create xpserver** command, which performs the configuration steps for you.

1. Start **sybcluster**.

For example, to start **sybcluster** and specify Sybase Control Center agents on the nodes "blade1," "blade2," and "blade3," enter:

```
sybcluster -U uafadmin -P -F
"blade1:1234,blade2:2345,blade3:3456"
```

2. Connect to the cluster. For example, to connect to "mycluster," enter:

```
connect to mycluster
```

3. Start the cluster. Enter:

```
start cluster
```

4. To configure the XP Servers, enter:

```
create xpserver
```

Adaptive Server prompts you for the XP Server port number for each instance. See the *Utility Guide* for more information about the **xp server** utility.

CHAPTER 10 **Postinstallation Tasks**

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 79

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 *instance_name* is the Adaptive Server name:

```
isql -Usa -P<password or leave it blank> -Sinstance_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 78

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.

Creating a Test Environment

To create a test environment, create a simulated cluster on a single node with all instances running on that node.

Sybase recommends that you configure each instance on a different node, that is, one node for each instance. However, for a test environment, you can create a simulated cluster on a single node with all instances running on that node. For best performance, make sure the number of engines in all instances running on a single node does not exceed the number of CPUs on that node.

The default setting for the **runnable process search count** parameter is 3. Sybase recommends that you use this default setting, which can be changed using the **sp_configure** stored procedure.

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 `-I` 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=libsybldap.dll
```

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

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

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

```
[ DIRECTORY ]
ldap=libsybldap.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> <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

Determine the type (raw partition, logical volume, operating system file, and so on) and location of the device where you plan to store the `pubs2` and `pubs3` databases.

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

Task

1. Start the server instance.
2. Go to the Adaptive Server scripts directory at `$SYBASE/$SYBASE_ASE/scripts`.
3. Use `isql` to log in to the instance 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`), `Roman8`, 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 (`ucjis`), `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 -Jeuclis \  
-i $SYBASE/$SYBASE_ASE/scripts/euclis/installjpubs
```

```
isql -Usa -Ppassword -Sservername -Jeuclis \  
-i %SYBASE%\%SYBASE_ASE%\scripts\euclis\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.

I/O Fencing

Sybase cannot guarantee data integrity unless you enable I/O fencing. If you do not use I/O fencing, data loss or other data corruption may occur in rare instances. Only test and development environments that can accept this risk should be deployed without I/O fencing.

The shared-disk cluster can detect the presence of a noncooperating instance and remove it from the cluster. In rare situations, however, the cluster cannot stop the noncooperating instance from writing to the shared disk, even though the instance is no longer part of the cluster. For example, if an instance has been removed from the cluster, but has not released

resources and shut down, it may still be capable of writing to the shared disk. Use I/O fencing to prevent the noncooperating instance from making data writes.

In the Cluster Edition, Sybase supports the SCSI-3 Persistent Group Reservation (PGR) feature of SCSI-3 devices to provide I/O fencing. PGR is the SCSI-3 standard for managing disk access in an environment where a single disk is shared by multiple hosts for read and write access.

The I/O fencing provided by the SCSI-3 PGR feature operates only on devices, not on partitions. For example, `/dev/sda1` and `/dev/sda2` are partitions of the device `/dev/sda`. A fencing operation targeted to a raw device bound to `/dev/sda1` affects all partitions of `/dev/sda`, so any file systems or other applications (including another Adaptive Server) using partitions on that device are also affected. For this reason, the device must be used exclusively by the cluster instance.

Preparing to Set Up I/O Fencing

Enable I/O fencing.

The user or process must have permissions to manage the raw devices, using `PV_ROOT`, `PV_SU_`, or `PV_KER_RAS` privileges.

1. To enable I/O fencing on IBM AIX ensure that the users who start Adaptive Server have permission to access the SCSI fixed disks (`/dev/rhdisk#`), since Adaptive Server can use the databases or the quorum.
2. Ensure that all nodes are running AIX.
3. All disks are PR-capable (DS3000, DS4000, or newer), hdisks, and permit read/write access, using the following “**chmod**” commands on both the hdisk and the corresponding character access device rhdisk:

```
chmod 666 rhdisk#
```

```
chmod 666 hdisk#
```

Setting Up Raw Disks For Multi-Node Access and I/O Fencing

You can set up raw disks for multi-node access and I/O fencing.

In Adaptive Server version 15.7 ESD #2, you need not to set the *PR_key_value* for any disks.

1. Set each node with the following:
 - a) Set *reserve_policy* to *no_reserve* on the quorum device and all raw database devices:


```
chdev -l hdisk# -a reserve_policy=no_reserve
```
 - b) Each node, with a unique node ID for each node:


```
chdev -l hdisk2 -a reserve_policy=no_reserve
```

If you are using an IBM System Storage device, you must install the latest versions of the Subsystem Device Driver Path Control Module (SDDPCM) and the Host Attachment for SDDPCM on AIX for the version of AIX and IBM System Storage device you are using. Downloaded these directly from the IBM Web site.

Note: IBM DS8000 storage series requires SDDPCM 2.4.0.4 or higher.

2. Run the following as SU or SUDO before starting dataserver, and make sure you have the appropriate shell privileges to run **datasecattr**:

```
setsecattr -p iprivs=+PV_KER_RAS $$
```

You can also run **setsecattr** with user “sybase” if it has the appropriate roles, such as a non-network information services (non-NIS) user.

3. Create a user with privileges to run programs using OpenX:

```
mkuser sybase
mkrole authorizations=aix.device.manage.change role_disk_access
chuser roles=role_disk_access sybase
```

4. If you are setting up raw disks to provide multi-node access and I/O fencing on IBM AIX, follow these additional steps:

- a) Create the user "sybase": **mkuser sybase**

- b) Create role and grant to sybase:

```
mkrole authorizations=aix.device.manage.change
role_disk_access
chuser roles=role_disk_access sybase
```

- c) Add the security 'aix.security.proc.set' for set security attribute:

```
smit -> Security & Users -> RBAC -> Roles -> Change ->
'role_disk_access' -> AUTHORIZATIONS
+'aix.security.proc.set'
setkst
```

- d) Change user to sybase and set the security attribute (set this attribute as default):

```
su - sybase
setsecattr -p iprivs=+PV_KER_RAS,PV_ROOT,PV_SU_ $$
```

I/O Fencing via SCSI-3 PGR

AIX supports I/O fencing with SCSI-3 Persistent Group Reservation (PGR) enabled.

There are two ways that AIX supports I/O Fencing via SCSI-3 PGR:

- The SCSI-3 PGR API
- The AIX DKPRES API

When the disk **reserve_policy** is set to:

- **no_reserve** – Adaptive Server uses raw SCSI-3 PGR API, where # is the hdisk number:

```
chdev -l hdisk# -a reserve_policy=no_reserve
```

In this configuration, remote applications cannot accidentally create and remove registrations on disk.

- **PR_shared** – Adaptive Server uses AIX DKPRES API, where # is the hdisk number.:

```
chdev -l hdisk# -a reserve_policy=PR_shared
```

This **PR_shared** setting may allow any remote application (such as database devices) to open the disk that automatically creates and removes registrations on the device. This can make the database devices susceptible to data corruption.

Note: **syb_cluster**/UAF are known to open database disks and cause this problem, and therefore should not be left running when using AIX DKPRES.

The preferred method is raw SCSI-3 PGR API, however, some IBM storage devices like DS5800 do not support raw SCSI-3 PGR. For this reason, set the `reserve_policy` to `no_reserve` first. If the Adaptive Server instance fails to start due to I/O fencing issues, then the device does not support raw SCSI-3 PGR API. In such cases, the devices must be configured to use DKPRES API, in which case, set the `reserve_policy` to **PR_shared**.

Enabling I/O Fencing

Configuring I/O fencing affects all devices, system-wide, that have I/O fencing capability.

1. Each instance in the cluster must run on a separate node.
2. All storage devices (disks) that hold database devices must support the SCSI-3 standard and cannot be partitioned.

Note: SCSI-3 PGR functionality is available only for a physical SCSI disk device or a virtual disk device exported by a storage area network (SAN). Partitioning such a device at the operating system level does not provide SCSI-3 PGR on each partition. Rather, the SCSI-3 PGR, and thus the fencing support, is shared among all partitions on the device. Adaptive Server cannot independently fence database devices.

- The quorum must reside on its own device. You cannot create database devices on the quorum device.
 - In addition, if device partitions are used outside the cluster, any fencing performed by the cluster denies access to those other partitions for the external applications or file systems placed on them.
-

3. I/O fencing is based on a device driver API that is targeted to a specific device driver. The device driver is typically called the fence device. Use operating system commands to create the fence device on each of the nodes running the cluster. To turn on the **enable i/o fencing** configuration parameter, enter:

```
sp_configure "enable i/o fencing", 1
```

4. The SCSI-3 PGR feature is platform-dependent, and all devices used by the Cluster Edition should have this functionality. See your operating system documentation for complete syntax and setup information. As part of the cluster creation process, both the Adaptive Server plug-in and **sybcluster** let you check that each device is enabled for I/O fencing. You can also run the **qrmutil** utility:

```
qrmutil -Qquorum path --fence-capable=device path
```

Creating Logical Volumes

Adaptive Server supports the fencing of logical volumes and partitions.

See your AIX documentation from IBM for details on these commands.

Note: You must create a volume group from a single logical unit number (LUN). However:

- Do not make database devices using LUNs (`/dev/rhdisk#`) that support existing volume groups and logical volumes.
 - Do not make volume groups or logical volumes using any LUNs that are currently used as database devices.
-

1. Create a volume group on one node using the AIX **mkvg** command:

```
mkvg -y volume_group_name hdisk#
```

For example: `mkvg -y myvolgroup hdisk5`

2. Create a logical volume on one node:

```
mklv -y logical_volume_name volume_group_name size
```

For example: `mklv -y mylogicalvol myvolgroup 100M`

3. Clear any disk reservations:

```
/opt/rsct/vsd/bin/prClear hdisk
```

For example: `/opt/rsct/vsd/bin/prClear hdisk5`

4. Import the volume groups on the other nodes to align the logical volumes on each node, then configure the logical volume to be created automatically when imported:

```
chdev -l hdisk -a pv=yes
importvg -y volume_group_name hdisk
```

For example:

```
chdev -l hdisk5 -a pv=yes
importvg -y myvolgroup hdisk5
```

5. Give read and write permission to the logical volume:

```
chmod 666 /dev/rlogical_volume_name
```

For example: `chmod 666 /dev/rmylogicalvol`

Manual Cluster Setup and Management

If you set up the cluster manually, you cannot use the Adaptive Server plug-in or **sybcluster** to manage your cluster until you create and deploy the Sybase Control Center agent.

See also

- *Enabling sybcluster and Sybase Control Center After a Manual Configuration* on page 98

Setting the Environment

Set your Sybase environment. From the Sybase release directory, source the SYBASE . sh or SYBASE . csh file.

From the \$SYBASE directory where the environment variable scripts can be located, source the environment variable. For example:

```
. SYBASE.sh
```

or,

```
source SYBASE.csh
```

Raw Devices

Each raw device must be accessible from each node using the same path. See your operating system storage administrator or documentation for help in setting up raw devices.

The local system temporary database device and the quorum disk device are specific to the Cluster Edition. The other required devices are the same as those required for all Adaptive Servers.

- Master database device
- **sybstemprocs** database device
- System database device
- Local system temporary database device (you can create multiple devices to set up one local system temporary database per device)

Note: Local system temporary database must use shared disks.

- Quorum disk device (at least 4MB)
- All other database devices

The Cluster Input File

Before configuring the cluster, create a cluster input file that specifies the name of the cluster, the number of instances in the cluster, the path to the directories containing the interfaces file, log files, quorum disk devices, and other required configuration information. Choose any name for the cluster input file (for example, mycluster . inp).

When you configure the cluster, Adaptive Server reads the information from the cluster input file and stores it securely in the quorum device. Adaptive Server subsequently retrieves the cluster configuration information from the quorum device.

See *Reconfiguring the Cluster* on page 97 for information about changing configuration information after the cluster has been initialized.

Note: You can configure one cluster with each cluster input file.

CHAPTER 10: Postinstallation Tasks

The cluster input file is distinct from the server configuration file, which stores Adaptive Server configuration values associated with **sp_configure**.

The syntax for the cluster input file is:

```
# all input files must begin with a comment
[cluster]
name = cluster_name
max instances = number
master device = path_to_the_master_device
configuration file = common_path_to_all_server_configuration_files
primary protocol = udp | tcp | other
secondary protocol = udp | tcp | other
installation mode = shared | private
configuration file = Adaptive_Server_configuration_file_name
interfaces path = interfaces_file_path
traceflags = trace_flag_number, trace_flag_number, . . .
additional run parameters = any_additional_run_parameters

[management nodes]
hostname = node_name
hostname = node_name
hostname = node_name
hostname = node_name

[instance]
id = instance_ID
name = instance_name
node = name_of_node_on_which_this_instance_runs
primary address = primary_interconnect_address
primary port start = port_number
secondary address = secondary_interconnect_address
secondary port start = port_number
errorlog = file_name
interfaces path = interfaces_file_path
config file = path_to_server_configuration_file_for_this_instance
traceflags = trace_flag_number, trace_flag_number, . . .
additional run parameters = any_additional_run_parameters

[instance]
id = instance_ID
name = instance_name
node = name_of_node_on_which_this_instance_runs
primary address = primary_interconnect_address
primary port start = port_number
secondary address = secondary_interconnect_address
secondary port start = port_number
errorlog = file_name
interfaces path = interfaces_file_path
configuration file =
path_to_server_configuration_file_for_this_instance
traceflags = trace_flag_number, trace_flag_number, . . .
additional run parameters = any_additional_run_parameters
```

where:

- **name** = *cluster_name* – is the name of the cluster.
- **max instances** = *number* – is the maximum number of instances in the cluster that your version of Adaptive Server supports. See the release bulletin for the most current information.
- **master device** = *path* – is the path to the master device.
- **configuration file** = *common_path* – is the common path to all server configuration files.
- **primary protocol** = **udp | tcp | other** – specifies the protocol used for the primary interconnect.
- **secondary protocol** = **udp | tcp | other** – specifies the protocol used for the secondary interconnect.
- **installation mode** = **shared | private** – specifies whether the installation mode is shared or private.
- **config file** = *filename* – is the path to the Adaptive Server configuration file. This file is used by all instances in the cluster unless an individual instance overrides this setting. In a private installation where all configuration files share the same path name, this is the common path.
- **interfaces path** = *file_path* – is the path to the interfaces file. Omit this parameter if you are using an LDAP environment. This interfaces file is used by all instances unless an individual instance overrides this setting.
- **traceflags** = *trace_flag_number, trace_flag_number, ...* – is the comma-separated list of trace flags used at instance start-up.
- **additional run parameters** = *parameters* – are the additional parameters passed to instances during start-up.
- **hostname** = *node_name* – is the name of the node. This name should be the same as returned by the host name command when run on this node. There is one host name field for each node that must be registered. Specify this node only once in the management node section.
- **ID** = *name* – is the ID of the instance.
- **name** = *instance_name* – is the name of the instance.
- **node** = *name* – is the name of the node on which this instance runs.
- **primary address** = *address* – is the address of this instance on the primary interconnect.
- **primary port start** = *number* – is the starting port number on the primary interconnect.
- **secondary address** = *address* – address of this instance on the secondary interconnect. Required if a secondary is defined in the interconnect section or if a secondary protocol is specified. Ignored if a secondary is not defined.
- **secondary port start** = *port_number* – is the starting port number on the secondary interconnect. Required if secondary address or a secondary protocol is specified.
- **error log** = *file_name* – is the full path to the error log for this instance.
- **interfaces path** = *path* – is the path to the server-side `interfaces` file. This file overrides the `interfaces file` field in the **cluster** section of the cluster input file. Do not

include the `interfaces` file name in the path. Omit this parameter if you are using LDAP.

- **config file** = *path* – is the path to an Adaptive Server configuration file. This file overrides the configuration file field specified in the **cluster** section of the cluster input file. In a private installation where path names to individual server configuration files are not the same, this is the path to the current server configuration file.
- **traceflags** = *trace_flag_number, trace_flag_number, ...* – is the comma-separated list of trace flags used at instance start-up. These are used in addition to, not in place of, the trace flags specified in the **cluster** section of the cluster input file.
- **additional run parameters** = *parameter_name* – are the additional parameters passed to instances during startup.

The formula for finding the socket port range is:

$$\text{start_port_number} + (\text{max_instances} * 5) - 1$$

Note: If you do not use an ADO.NET, ensure that the port numbers you choose are not in use by any other process.

In this example, the cluster input file defines a cluster named "mycluster" with two instances, "ase1" on node "blade1," and "ase2" on node "blade2." The addresses on the private interconnects are 192.169.0.1 and 192.169.0.2. The name of the server configuration file is `mycluster.cfg`. The maximum instances is 2. "ase1" has a starting port range of 15015, and "ase2" has a starting port range of 16015. This adds additional information to the mycluster cluster:

```
#input for a 2 node / 2 instance cluster
[cluster]
name = mycluster
max instances = 2
master device = /opt/sybase/rawdevices/mycluster.master
config file = /opt/sybase/ASE-15_0/mycluster.config
interfaces path = /opt/sybase
primary protocol = udp
secondary protocol = udp

[management nodes]
hostname = blade1.sybase.com
hostname = blade2.sybase.com

[instance]
id = 1
name = ase1
node = blade1.sybase.com
primary address = 192.169.0.1
primary port start = 15015
secondary address = 192.169.1.1
secondary port start = 15015
errorlog = /opt/sybase/ASE-15_0/install/ase1.log
additional run parameter = -M/opt/sybase/ASE-15_0

[instance]
```

```

id = 2
name = ase2
node = blade2.sybase.com
primary address = 192.169.0.2
primary port start = 16015
secondary address = 192.169.1.2
secondary port start = 16015
errorlog = /opt/sybase/ASE-15_0/install/ase2.log
additional run parameter = -M/opt/sybase/ASE-15_0

```

For an example of a cluster input file where all instances are located on a single node, see the *Clusters Users Guide*.

Manually Configuring the Cluster

Sybase strongly recommends that the total number of engines in all instances does not exceed the number of CPUs.

After setting the environment, and creating both raw devices and a cluster input file, you can begin configuring the cluster. Manually configuring a cluster involves several tasks:

1. Configure the `interfaces` file for the cluster server and all instances.
2. Create the quorum and master devices on raw devices on the shared-disk storage.
3. Use disk init to initialize and create the `sybssystemprocs` database.
4. Run `InstallAnywhere` to install the system stored procedures.
5. After creating the master and quorum devices, create a `runserver` file for each instance in the cluster.
6. Set up a local system and temporary databases.

See also

- *Configuring the Interfaces File* on page 91
- *Building the Master and Quorum Devices* on page 92
- *Initializing and Creating sybssystemprocs* on page 93
- *System Stored Procedures* on page 93
- *Creating Runserver Files* on page 93
- *Setting Up Local System and Temporary Databases* on page 94

Configuring the Interfaces File

If you use an `interfaces` file, it must include entries for the cluster server and all instances.

The syntax for the `interfaces` file is:

```

instance_name:
  master network_protocol machine_name port_number
  query network_protocol machine_name port_number
  .
  .
cluster_server_name
  query network_protocol ether machine_name port_number

```

where:

- **instance_name** – is the instance for which you are making this interfaces file entry.
- **network_protocol** – is the network protocol used for the instance.
- **machine_name** – is the name of the machine on which the instance runs.
- **port_number** – is the port number used to connect to this instance.
- **cluster_server_name** – is the name of the cluster server.

This example includes the instances "ase1," "ase2," and "ase3," on machines "blade1," "blade2," and "blade3," running in the cluster "mycluster":

```
ase1
  master tcp ether blade1 19786
  query tcp ether blade1 19786
ase2
  master tcp ether blade2 19786
  query tcp ether blade2 19786
ase3
  master tcp ether blade3      19786
  query tcp ether blade3 19786
mycluster
  query tcp ether blade1 19786
  query tcp ether blade2 19786
  query tcp ether blade3 19786
```

Building the Master and Quorum Devices

Create the quorum and master devices on raw devices on the shared-disk storage.

When creating the quorum device, use the same device name and major and minor numbers on all machines. Each raw device must be accessible from each node using the same path. This example uses raw11 for the quorum disk.

For example:

```
dataserver
...
--quorum-dev /dev/raw/raw11
```

The syntax for creating Cluster Edition master and quorum devices is:

```
dataserver
--cluster-input= cluster_input_filename_and_path
--quorum-dev= quorum_device_and_path
--master-device-size= master_device_size
--logical-page-size= page_size
--instance= instance_name
--buildquorum
```

where:

- **--master-device-size=<size spec>** – specifies the size of the master device.
- **--cluster-input=<cluster input file>** – loads the cluster configuration specified in the input file onto the quorum device.
- **--quorum-dev= path_to_quorum_device** – specifies the full path to the quorum device.
- **--instance=instance_name** – specifies the name of an instance.

- **--logical-page-size= page_size** – specifies the page size.
- **--buildquorum** – specifies that a new quorum device should be built.

Create the master and quorum devices using **dataserver**. This example creates an instance named “ase1” that is configured for 4K pages, a master device of 500MB, and a quorum device:

```
/opt/sybase/ASE-15_0/bin/ dataserver\  
--quorum-dev=/dev/raw/raw11\  
--instance=ase1  
--cluster-input=/opt/sybase/mycluster.inp  
--master-device-size=500M  
--logical-page-size=4k  
--buildquorum
```

See the *Cluster Users Guide* and *Utility Guide* for more information about the **dataserver** utility.

System Stored Procedures

Run **installmaster** from the `$$SYBASE/$SYBASE_ASE/scripts` directory to install the system stored procedures. You can run **installmaster** from any instance.

```
isql -U sa -P sa_password -S server_name -n  
-i $$SYBASE/$SYBASE_ASE/scripts/installmaster  
-o output_file
```

Initializing and Creating sybssystemprocs

Use **disk init** to initialize a device for sybssystemprocs, and then create the sybssystemprocs database.

1. Initialize a raw device of at least 150MB for sybssystemprocs:

```
disk init name = "sysprocsdev",  
physname = "/dev/raw/raw13",  
size = "150M"
```

2. Create the sybssystemprocs database:

```
create database sybssystemprocs on sysprocsdev = 150
```

Creating Runserver Files

After creating the master and quorum devices, create a runserver file for each instance in the cluster. Later, you will use these files to start the instances.

1. Create a runserver file.

This example creates `RUN_ase1` for instance `ase1`. Enter the runserver file on a single line.

```
$$SYBASE/ASE-15_0/bin/dataserver  
--quorum-dev=/dev/raw/raw11  
--instance=ase1
```

Note: Create all devices, database devices, and quorum devices as raw devices on shared disks.

2. Make a copy of the runserver file for each of the Adaptive Servers in the cluster. For example, the three instances for the cluster named “mycluster” have runserver files named RUN_ase1, RUN_ase2, and RUN_ase3. Make sure you include the same quorum device for all instances. **--instance** should change in each file to point to the appropriate instance name.

Setting Up Local System and Temporary Databases

In a shared-disk cluster, each instance must have a local system temporary database.

In the Cluster Edition, if an instance does not have a local system temporary database, it starts only if it is the first instance in the cluster to start.

1. Start Adaptive Server. For example:

```
startserver -f $SYBASE/$SYBASE_ASE/install/RUN_ase1_coord
```

2. Log in to Adaptive Server.
3. If you are not using the master device for the temporary databases, create the devices for the local system temporary databases. You can create local system temporary database on shared disks only. See *Using Temporary Databases* in the *Cluster Users Guide*.

Sybase recommends that you use separate devices for log and data for these databases. For example:

```
disk init name="tempdbdev1",
physname="/dev/raw/raw14",
size="400M"
```

And for the log device:

```
disk init name="temp_log_dev1",
physname="/dev/raw/raw15",
size="200M"
```

4. Create a local system temporary database for each instance in the cluster. For example, to create three local system temporary databases "ase1_tdb1," "ase2_tdb1," and "ase3_tdb1" for instances "ase1," "ase2," and "ase3," respectively, enter:

```
create system temporary database ase1_tdb1 for instance ase1 on
tempdbdev1 = 100
log on temp_log_dev1 = 50
create system temporary database ase2_tdb1 for instance ase2 on
tempdbdev1 = 100
log on temp_log_dev2 = 50
create system temporary database ase3_tdb1 for instance ase3 on
tempdbdev1 = 100
log on temp_log_dev3 = 50
```

5. Shut down the cluster using the **shutdown cluster** command.

Automatic Takeover of a Cluster

Configure the cluster to automatically recover from an abrupt total cluster failure.

Automatic cluster takeover is triggered if an instance attempts to connect to the cluster and:

- You have enabled **automatic cluster takeover**.
- The quorum device indicates the cluster is running, but the instance detects no heartbeat.

The **automatic cluster takeover** configuration parameter enables the instance attempting to connect to the cluster to form a new cluster, starting the cluster coordinator and recovering the databases. The syntax for **automatic cluster takeover** is:

```
sp_configure "automatic cluster takeover", [1 | 0]
```

Setting the value to 1 (the default) to enable the automatic cluster takeover; set to 0 to disable **automatic cluster takeover**.

In environments that have I/O fencing enabled, **automatic cluster takeover** is guaranteed to be a safe operation. For environments that do not have I/O fencing, a malfunction of the algorithm may introduce data corruption. This configuration parameter disables the algorithm if it malfunctions. There is always a risk of data corruption in environments without I/O fencing, and disabling automatic cluster takeover does not mitigate all of those risks.

See also

- *Starting the Cluster* on page 95

Starting the Cluster

Start the cluster by starting all the instances associated with the cluster on each node.

You can start a cluster:

- Normally, after a graceful shutdown.
- After a system failure.

Note: Sybase recommends that you do not alter the runserver files used for normal cluster startup.

1. Start each cluster instance from the node on which it runs:

For example, to start the "ase1" instance, enter:

```
startserver -f $SYBASE/$SYBASE_ASE/install/RUN_ase1
```

2. Log in to each node on which another instance will run, and execute **startserver**.

For example, to start the "ase2" instance on "blade2," enter:

a) On "blade2," move to the \$SYBASE directory.

b) Issue:

```
startserver -f $SYBASE/$SYBASE_ASE/install/RUN_ase2
```

See also

- *Starting the Cluster After a System Failure* on page 96

Starting the Cluster After a System Failure

The manner in which you start the server after a system failure depends on whether or not you enabled the **automatic cluster takeover** configuration parameter.

If the cluster fails because of a system failure and you:

- Enabled **automatic cluster takeover** – the instance attempting to connect to the cluster restarts the cluster, acting as the cluster coordinator and recovering the databases.
- Did not enable **automatic cluster takeover** – you must restart the cluster with the **dataserver . . . --cluster-takeover** parameter. For example, to add the **-cluster-takeover** parameter to the runserver file that originally started the cluster, enter:

```
$SYBASE/ASE-15_0/bin/dataserver \  
--quorum-dev=/dev/raw/raw11\  
--instance=ase1\  
--cluster-takeover
```

Once the cluster is running, restart all the instances normally.

See also

- *Automatic Takeover of a Cluster* on page 95

Postconfiguration Tasks

Perform additional tasks after you determine that the configured instances are running.

To determine if the instances you configured are running, log in to the cluster and enter **sp_cluster show**, and verify that you can connect to the instances:

1. Source `SYBASE.sh`.
2. Use **isql** to connect to servers. At the command prompt, enter:

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

where *server_name* is the instance name. You see the command prompt if the login is successful.

3. To display the Adaptive Server version number, enter:

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

If you encounter errors, see the *Error Messages and Troubleshooting Guide*.

Shutting Down the Cluster or the Instance

Shutting down the cluster shuts down all the instances associated with the cluster.

1. Log in to any instance. For example:

```
isql -Usa -P -Sase2 -I$SYBASE/interfaces
```

2. Issue:

```
shutdown cluster
```

Shutting Down an Instance

Use the **shutdown** command to shut down an individual instance in the cluster. You can also shut down an instance from another instance in the cluster.

1. Log in to any instance. For example:

```
isql -Usa -P -Smycluster -I$SYBASE/interfaces
```

2. Issue:

```
shutdown ase2
```

Reconfiguring the Cluster

You can reconfigure a cluster to add instances, change trace flags, and so on. Edit the cluster input file, then create a new run file for the instance you intend to start first and that includes the **dataserver** with the **cluster-input** option. This option tells Adaptive Server to write the new configuration information to the quorum device.

Log in to any instance, and shut down the cluster.

Note: If you use **sybcluster** or the Adaptive Server plug-in, you do not need to shut down the cluster before reconfiguring it.

1. Extract the current configuration to a file using the **-extract-config** command to the **qrmutil** utility. See the *Clusters Users Guide*.
2. Edit the cluster input file.
3. Copy the run file for the instance you intend to start and add the **--cluster-input** option to the **dataserver** statement. For example, change `RUN_ase1` in this way:

```
$SYBASE/$SYBASE_ASE/bin/dataserver\  
--cluster-input=/input_file>\  
--quorum-dev=/dev/raw/raw11\  
--instance=ase1\  

```

4. Start the cluster.

Enabling sybcluster and Sybase Control Center After a Manual Configuration

To use **sybcluster** or Sybase Control Center to manage the cluster after a manual configuration, start Sybase Control Center remote command and control agent on each node in the cluster, and then deploy the connection agent to each node.

1. Start Sybase Control Center on each node used by the cluster. From the \$SYBASE directory, enter:

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

2. Start **sybcluster**. For example, enter:

```
sybcluster -U uafadmin -P -C mycluster  
-F "blade1,blade2,blade3"
```

3. Execute **deploy plugin**. For example:

```
deploy plugin agent "blade1,blade2,blade3"
```

See *sybcluster Utility* in the *Clusters Users Guide* for complete syntax and usage information for **sybcluster** and the *Sybase Control Center for Adaptive Server Enterprise*.

See also

- *Chapter 8, Starting and Stopping Sybase Control Center Agent* on page 55

A sybcluster Sample Session Configuration Values

Configure a typical shared-disk cluster using **sybcluster**.

This example assumes a primary and secondary network. If you do not specify a secondary network, **sybcluster** asks for a beginning port number and calculates the number of port numbers needed for communication among the instances. Ensure that the number you provide, and the following required port numbers, are not in use by another application. For information on about how **sybcluster** calculates the number of required ports, see *Chapter 9, Creating and Starting the Cluster* on page 57.

Table 10. sybcluster Prompts and Configuration Values for Sample Session

Parameter	Value
Cluster name	mycluster
Number of instances	
Number of agents	
Cluster node ID	1
Configuration type	

Parameter	Value
Is the cluster to be configured using a private \$SYBASE installation?	N
Quorum Device	
Full path and name for the quorum device	/dev/rhdisk11
Page Size	
The page size in kilobytes	2kB
Master device	
Full path and name for the master device	/dev/rhdisk12
Size of master device (MB)	30MB
Size of master database (MB)	13MB
PCI device (optional)	
Full path to the PCI database device	/dev/raw/raw20
Size of the PCI database device (MB)	24MB
Size of the PCI database (MB)	24MB
Sybase system procedure device	
Full path and name for sybssystemprocs device	/dev/rhdisk13
Size of sybssystemprocs device (MB)	160MB
Size of sybssystemprocs database (MB)	152MB
System database device	
Full path and name of system database device	/dev/rhdisk14
Size of system database device (MB)	6MB
Size of system database (MB)	6MB

Does this cluster have a secondary network (Y)	Yes
Verify file locations	
The location of the SYBASE home directory	/remote/var/sybase/

Parameter	Value
Full path to environment variable script	/remote/var/sybase/ SYBASE.sh
Path to \$SYBASE_ASE	/remote/var/sybase/ ASE-15_0
Path to interfaces file directory	/remote/var/sybase
Path to the dataserver configuration file	/remote/var/sybase/ mycluster.cfg
Instance information	
Node name	blade1
sybcluster displays: Cluster: mycluster - Node: blade1 - Agent blade1:9999	
Instance name	ase1
Query port number for ase1	19786
Primary protocol address for ase1	000.000.001.001
Secondary protocol address for ase1	000.000.001.002
Local system temporary database	
Local system temporary database device name. Enter the name of the Adaptive Server database device for the local system temporary database. You must use a shared disk for local system temporary database devices.	mycluster1_tempdb
Path to local system temporary database device	/dev/rhdisk15
Local system temporary database device size (MB)	40MB
Local system temporary database name	mycluster_tdb_1
Local system temporary database size (MB)	40MB

Do you want to add another instance (Y or N)	Yes
Node name	blade 2

Parameter	Value
sybcluster displays: Cluster: mycluster - Node: blade2 - Agent blade2:9999	
Instance name	ase2
Query port number for ase2	19786
Primary protocol address for ase2	000.000.002.001
Secondary protocol address for ase2	000.000.002.002
Local system temporary database	
Local system temporary database device name. Enter the name of the Adaptive Server database device for the local system temporary database. You must use a shared disk for local system temporary database devices.	mycluster2_tempdb
Path to local system temporary database device	/dev/rhdisk16
Local system temporary database device size (MB)	40MB
Local system temporary database name	mycluster_tdb_2
Local system temporary database size (MB)	40MB

Do you want to add another instance (Y or N)	Yes
Node name	blade3
sybcluster displays: Cluster: mycluster - Node: blade3 - Agent blade3:9999	
Instance name	ase3
Query port number	19786
Primary protocol address for ase3	000.000.003.001
Secondary protocol address for ase3	000.000.003.002
Local system temporary database	
Local system temporary database device name. Enter the name of the Adaptive Server database device for the local system temporary database. You must use a shared disk for local system temporary database devices.	mycluster3_tempdb

Parameter	Value
Path to local system temporary database device	/dev/rhdisk17
Local system temporary database device size (MB)	40MB
Local system temporary database name	mycluster_tdb_3
Local system temporary database size (MB)	40MB

Do you want to add another instance (Y or N)	N
Save configuration information in a file (Y(Yes
Full path for configuration file	/remote/var/sybase/ mycluster.xml
Create the cluster now (Y)	Yes
Veritas Cluster Server support (optional – appears only if VCS is supported on your system)	
Integrate the Cluster Edition server with VCS? (Y)	N
Check whether each cluster device has I/O fencing capability? (Y)	

See also

- *Create a Cluster Using sybcluster* on page 64

Sample sybcluster Session

A **sybcluster** sample session that assumes a shared configuration, an interfaces file for server lookup, and support for Java in the database.

```

sybcluster -U uafadmin -P -F
hpcblade2:9009,hpcblade1:9009,hpcblade3:9009,hpcblade4:9009
> create cluster
Enter the name of the cluster: mycluster
Cluster mycluster - Enter the maximum number of instances: [ 4 ]
How many agents will participate in this cluster: [ 4 ] 4
Verifying the supplied agent specifications...
1) hpcblade1.sybase.com 9009 2.5.0 Linux
2) hpcblade2.sybase.com 9009 2.5.0 Linux
3) hpcblade3.sybase.com 9009 2.5.0 Linux
4) hpcblade4.sybase.com 9009 2.5.0 Linux
Enter the number representing the cluster node 1: [ 4 ] 1
2) hpcblade2.sybase.com 9009 2.5.0 Linux
3) hpcblade3.sybase.com 9009 2.5.0 Linux
4) hpcblade4.sybase.com 9009 2.5.0 Linux
Enter the number representing the cluster node 2: [ 4 ] 2

```



```

3) hpcblade3.sybase.com 9009 2.5.0 Linux
4) hpcblade4.sybase.com 9009 2.5.0 Linux
Enter the number representing the cluster node 3: [ 4 ] 3
4) hpcblade4.sybase.com 9009 2.5.0 Linux
Enter the number representing the cluster node 4: [ 4 ] 4
Will this cluster be configured using private SYBASE installations?
(Y/N) : [ N ]
----- Quorum Device -----
The quorum device is used to manage a cluster. It contains
information shared between instances and nodes.
Enter the full path to the quorum disk: /hpcblade_cfs/q/pd16218942/
d3.dbs
Enter any traceflags:
----- Page Size -----
Enter the page size in kilobytes: [ 2 ] 8
----- Master Database Device -----
The master database device controls the operation of the Adaptive
Server and stores information about all user databases and their
associated database devices.
Enter the full path to the master device: /hpcblade_cfs/q/pd16218942/
d4.dbs
Enter the size the Master Device (MB): [ 120 ] 500
Enter the size the Master Database (MB): [ 52 ] 100
----- Sybase System Procedure Device -----
Sybase system procedures (sybssystemprocs) are stored on a device.
Enter the System Procedure Device path: /hpcblade_cfs/q/pd16218942/
d5.dbs
Enter System Procedure Device size (MB): [ 152 ] 200
Enter the System Procedure Database size (MB): [ 152 ] 200
----- System Database Device -----
The system database (sybssystemdb) stores information about
distributed transactions.
Enter the System Database Device path: /hpcblade_cfs/q/pd16218942/
d6.dbs
Enter the System Database Device size (MB): [ 24 ] 100
Enter the System Database size (MB): [ 24 ] 100
----- PCI Device -----
Pluggable Component Interface (PCI) provides support for Java in
database by loading off-the-shelf JVMs from any vendor. If you want
to use JVM, create a device for it.
Enable PCI in Adaptive Server (Y/N): [ N ] y
Enter the full path to the PCI device: /hpcblade_cfs/q/pd16218942/
pci.dbs
Enter the size the PCI Device (MB): [ 96 ]
Enter the size the PCI Database (MB): [ 96 ]
-----
Does this cluster have a secondary network: [ Y ] n
Enter the port number from which this range will be applied:
[ 15100 ] 17005
-----
Enter the SYBASE home directory: [ /remote/quasr5/adong/aries/
release/lamce_s1 ]
Enter the environment shell script path: [ /remote/quasr5/adong/
aries/release/lamce_s1/SYBASE.sh ]
Enter the ASE home directory: [ /remote/quasr5/adong/aries/release/
lamce_s1/ASE-15_0 ]

```

CHAPTER 10: Postinstallation Tasks

```
Enter path to the dataserver configuration file: [ /remote/quasr5/
adong/aries/release/lamce_sl/mycluster.cfg ]
-----
You will now be asked for the instance information on a node by node
basis.
-- Cluster: mycluster - Node: hpcblade1.sybase.com - Agent:
hpcblade1.sybase.com:9009 --
Enter the name of the cluster instance: instance1
Enter the interface file query port number for instance instance1:
10665
Enter the primary protocol address for instance1:
[ hpcblade1.sybase.com ]
----- Local System Temporary Database -----
The Local System Temporary Database Device contains a database for
each instance in the cluster.
Enter the LST device name: LST
Enter the LST device path: /hpcblade_cfs/q/pd16218942/d7.dbs
Enter LST device size (MB): 200
Enter the LST database name: [ mycluster_tdb_1 ]
Enter the LST database size (MB): [ 200 ] 50
Do you want to add another instance to this node? (Y or N): [ N ]
-- Cluster: mycluster - Node: hpcblade2.sybase.com - Agent:
hpcblade2.sybase.com:9009 --
Enter the name of the cluster instance: instance2
Enter the interface file query port number for instance instance2:
15465
Enter the primary protocol address for instance2:
[ hpcblade2.sybase.com ]
----- Local System Temporary Database -----
The Local System Temporary Database Device contains a database for
each instance in the cluster.
Enter the LST device name: [ LST ]
Enter the LST database name: [ mycluster_tdb_2 ]
Enter the LST database size (MB): [ 150 ] 50
Do you want to add another instance to this node? (Y or N): [ N ]
-- Cluster: mycluster - Node: hpcblade3.sybase.com - Agent:
hpcblade3.sybase.com:9009 --
Enter the name of the cluster instance: instance3
Enter the interface file query port number for instance instance3:
16730
Enter the primary protocol address for instance3:
[ hpcblade3.sybase.com ]
----- Local System Temporary Database -----
The Local System Temporary Database Device contains a database for
each instance in the cluster.
Enter the LST device name: [ LST ]
Enter the LST database name: [ mycluster_tdb_3 ]
Enter the LST database size (MB): [ 100 ] 50
Do you want to add another instance to this node? (Y or N): [ N ]
-- Cluster: mycluster - Node: hpcblade4.sybase.com - Agent:
hpcblade4.sybase.com:9009 --
Enter the name of the cluster instance: instance4
Enter the interface file query port number for instance instance4:
15220
Enter the primary protocol address for instance4:
[ hpcblade4.sybase.com ]
```

```

----- Local System Temporary Database -----
The Local System Temporary Database Device contains a database for
each instance in the cluster.
Enter the LST device name: [ LST ]
Enter the LST database name: [ mycluster_tdb_4 ]
Enter the LST database size (MB): [ 50 ]
Would you like to save this configuration information in a file? [
Y ]
Enter the name of the file to save the cluster creation information:
[ /hpcblade_cfs/q/s16218942/mycluster.xml ]
-----
Create the cluster now? [ Y ]
-----

```

Configuring the Cluster Using an Input File

At the end of a **sybcluster** session, you can save the values from the current session to an external file. You can use this file to re-create the same cluster, or edit the values in the file to create a different cluster.

The syntax is:

```
create cluster cluster_name file xml_input_file
```

This example creates "mycluster2," using an input file named `mycluster.xml`:

```
create cluster mycluster2 file ./mycluster.xml
```


CHAPTER 11 **Adaptive Server Upgrades**

You can upgrade from version 15.5 and later of a nonclustered Adaptive Server to version 15.7 of the Cluster Edition.

You can upgrade to the Cluster Edition from these noncluster versions of Adaptive Server:

- 15.0 through 15.7
- 12.5 through 12.5.4 ESD #10

You can upgrade to the current Cluster Edition from these earlier Cluster Editions:

- 15.0.3 Cluster Edition
- 15.0.1 Cluster Edition through 15.0.1 Cluster Edition ESD #4

Once you have upgraded to the 15.5 or later Cluster Edition, you cannot downgrade to a nonclustered 15.x or earlier server. You can, however, downgrade to 15.0.1 Cluster Edition through 15.0.1 Cluster Edition ESD #4, Cluster Edition 15.0.3, and Cluster Edition 15.5.

Note: The **sybcluster** utility does not have a built-in tool to perform an upgrade or downgrade between two versions of Adaptive Server Cluster Edition.

Upgrade Adaptive Server using the Adaptive Server plug-in or **sybcluster**, or manually.

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

To upgrade from a shared-installation mode to a private-installation mode in Adaptive Server Cluster Edition, see *Upgrading to a Private Installation* on page 117.

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.

Note: When upgrading a 12.5.4 or earlier database to the Cluster Edition, make sure you execute all steps in the upgrade sequence from the same node. That is, you must load the database and execute online database on the same node.

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 11. 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.

Current Version	Upgrading to	Special Upgrade Information
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.
Adaptive Server Cluster Edition 15.5.x or later	Adaptive Server Cluster Edition 15.7.x	In the installer destination directory pane, enter the 15.5.X directory, then choose Update . After your installation, install the new system procedures (from \$SYBASE/ASE-15_0/scripts/installmaster). See <i>"Installation of Version 15.7 ESD #2 over Existing 15.x Adaptive Server"</i> on page 132 for details.

Table 12. 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.

Current Version	Upgrading to	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 138
- *Postupgrade Tasks* on page 134
- *Installation of Version 15.7 ESD #2 over an Existing 15.7 or 15.7 ESD #1 Cluster Edition* on page 132

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 13. 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

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
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	
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.

Checking System and Upgrade Requirements

Verify that your system meets upgrade requirements.

1. Verify that the computer on which you plan to upgrade the Sybase products meets system requirements.
2. Determine whether the version of your server can be upgraded to the Cluster Edition.
3. Verify that you unloaded Adaptive Server in a different directory than your previous Adaptive Server installation. However, if you overwrite the previous installation:

- a) Restore it from your most recent backup.
 - b) Reinstall the product files for Adaptive Server in a different directory.
 - c) Continue with the upgrade.
4. Verify that your operating system is compatible with, and prepared for the Cluster Edition.

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

- The upgrade process modifies the `sysaudits` tables in the `sybsecurity` database. Sybase recommends that you archive any auditing data and truncate these tables before you upgrade, reducing the chances of a failed upgrade due to lack of space in the `sybsecurity` database.
- Stored procedure text in the `syscomments` table is required for upgrade.
- If you are upgrading to a private installation, see *Upgrading to a Private Installation* on page 117.

Task

1. Install Adaptive Server 15.7 ESD #2 into its own installation directory.
2. Check system and upgrade requirements.
3. 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.

4. While the upgrade process for Adaptive Server requires that the previously installed version of the server is running, the upgrade process for Backup Server, Historical Server, and XP Server, requires those servers to be shut down.
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. Make sure reserved words use quoted identifiers.

7. Verify that users are logged off.
8. Use **dbcc** to check for database integrity.
9. Back up databases.
10. Dump the transaction log.
11. Verify that `master` is the default database for the "sa" user.
12. 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 137.
 - e) Disable disk mirroring.

Note: Cluster Edition version 15.7 ESD #2 does not support disk mirroring.

- f) Verify that your SYBASE environment variable points to the location of the new server software files you just installed.

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.

13. If you are upgrading from a nonclustered 12.5.4, 15.0.2, or 15.5 server, and archive database access has been applied in the lower version server, disable the related functionality before upgrading.
14. If you are upgrading from a 12.5.3 installation with cache bindings on `sybssystemdb`, drop the cache bindings for `sybssystemdb` that are bound to user-defined caches before running **preupgrade**.

If you do not do this, you see this error:

```
Current process... infected with 11
```

If you receive this error, re-run **preupgrade** after you have dropped the cache bindings.

15. 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.
16. 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`

17. If you have enabled the Java in the database feature, create the `sybpcidb` database and disable the feature during the installation.
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`.
19. When you upgrade a server to the Cluster Edition, index and table-level statistics become inaccurate, and must be updated. Execute **update index statistics** on these tables:
 - `sysobjects`
 - `sysindexes`
 - `syscolumns`
 - `systypes`
 - `syslogins`
 - `sysusers`

See also

- *Managing Java in the Database During Installations and Upgrades* on page 38

Upgrade System Tables and Stored Procedures

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

When you upgrade the Cluster Edition, system tables are also updated, which may affect your existing application. See *System Changes* in the *Clusters Users Guide* for a complete list of catalogs that are affected.

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.

The runserver File Location

Be sure the `runserver` file for your current server is located in `$SYBASE/$SYBASE_ASE/install/RUN_servername`.

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

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.

See the *Reference Manual: Building Blocks* for a complete list of reserved words.

Running a Reserved Word Check

Run the reserved word check on the old Adaptive Server.

1. Install the Cluster Edition version of `installupgrade` (where `$$SYBASE` and `$$SYBASE_ASE` are the values for the Cluster Edition):

```
isql -Usa -Ppassword -Sserver_name
-i$$SYBASE/$$SYBASE_ASE/scripts/installupgrade
```

2. Install the Cluster Edition version of `usage.sql`:

```
isql -Usa -Ppassword -Sserver_name
-i$$SYBASE/$$SYBASE_ASE/upgrade/usage.sql
```

3. Log in to the old Adaptive Server and execute **sp_checkreswords** on all the databases. For example:

```
use sybssystemprocs
go
sp_checkreswords
go
```

4. Correct any errors from the reserved word check.

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**.

Upgrading to a Private Installation

Upgrade from a shared installation to a private installation.

In Adaptive Server Cluster Edition version 15.5 and later, you can configure your cluster as a "shared" or "private" installation. See *An Overview of the Cluster Edition* in the *Clusters Users Guide*.

You must manually upgrade a symmetric multiprocessing (SMP) version of Adaptive Server to a private installation of the Cluster Edition. First, upgrade your Adaptive Server to a Cluster Edition shared installation, then switch to a private installation using the steps below. Since private installations were introduced in version 15.0.3, cluster instances created with earlier versions of Adaptive Server Cluster Edition automatically continue as a shared installation.

Note: When deciding on the installation location for Adaptive Server Cluster Edition 15.7 ESD #2, choose the location where you will install the private installation for this node. This location need not be accessible from other nodes participating in the cluster.

1. Make sure each participating node in the cluster has its own \$SYBASE environment variable. Typically, the private installation is performed on a local file system, as there is no longer a need for other nodes participating in the cluster to access this installation.
2. Install the Cluster Edition on each node participating in the cluster. You can set up one of the nodes to use the existing installation if it satisfies the needs, otherwise you can discard it at the end of this process. You may need to discard an existing installation if, for example, it is on an NFS file system being used by nodes and you want to install on a local file system. See the installation guide for your platform for instructions on how to install a Cluster Edition for every node.

3. On each node, shut down the cluster and the Sybase Control Center agent.
4. On one of the nodes in the cluster, set up your environment by sourcing `SYBASE.csh` or `SYBASE.sh`, depending on the shell you are using. If the SYBASE installation location differs from the shared installation and the private installation, set up the environment from the shared installation area.

5. Extract the current cluster quorum configuration from the quorum device. For example:

```
% $SYBASE/$SYBASE_ASE/bin/qrmutil
  --extract-config=mycluster_shared.cfg
  --quorum-dev=/dev/raw/raw50m41

Executing command 'extract cluster configuration', argument
'mycluster_shared.cfg'...

Extracted input file 'mycluster_shared.cfg'

Command 'extract cluster configuration', argument
'mycluster_shared.cfg' succeeded.

qrmutil execution completed.
```

6. Create a new cluster configuration file and update the required information:

- a) Make a copy of the extracted configuration file, then edit the new file to change required configurations, such as: `cp mycluster_shared.cfg mycluster_private.cfg`

- b) Edit the new configuration file; in the `[cluster]` section, change:

```
installation mode = shared
```

to:

```
installation mode = private
```

- c) In the `[instance]` section:

1. Move the configuration file and interfaces entry from the `[cluster]` to the `[instance]` section.
2. If the SYBASE installation location has changed from shared to private, adjust the paths in the error log, config file, and interfaces path locations.
3. If you have more than one instance in the configuration file, perform these actions for each instance. For example:

```
% cat mycluster_private.cfg

# All input files must begin with a comment
[cluster]
name = mycluster
max instances = 4
primary protocol = udp
secondary protocol = udp
master device = /dev/raw/raw1g2
traceflags =
additional run parameters =
installation mode = private
membership mode =
```



```

[management nodes]
hostname = nuno1
hostname = nuno2

[instance]
name = mycluster_instance1
id = 1
node = nuno1
primary address = nuno1
primary port start = 15100
secondary address = nuno1
secondary port start = 15181
errorlog = /mysybase1/mycluster_inst1.log
config file = /mysybase1/mycluster.cfg
interfaces path = /mysybase1
traceflags =
additional run parameters =

[instance]
name = mycluster_instance2
id = 2
node = nuno2
primary address = nuno2
primary port start = 15100
secondary address = nuno2
secondary port start = 15181
errorlog = /mysybase2/mycluster_inst2.log
config file = /mysybase2/mycluster.cfg
interfaces path = /mysybase2
traceflags =
additional run parameters =

```

7. Load the updated cluster configuration file into the cluster quorum device. For example:

```

% $SYBASE/$SYBASE_ASE/bin/qrmutil
--quorum-dev=/dev/raw/raw50m41
--cluster-input=mycluster_private.cfg
Loaded a new quorum configuration.
qrmutil execution completed.

```

8. If you have:

- More than one node in the cluster or have changed the SYBASE installation location – copy the Adaptive Server configuration file—typically named `servername.cfg`—and interfaces file from the original shared installation cluster into the corresponding interfaces path and configuration file locations for each instance in the private installation cluster. You can find the locations in the [instance] section of the updated cluster configuration file.
- Only one node/instance in the cluster and are not changing the SYBASE installation directory – update the Sybase Control Center agent configuration information. The agent plug-in XML file is located at `$SYBASE/SCC-3_2/instances/[machine_name]/plugins/[cluster_name]/agent-plugin.xml`.

In it, replace:

```
<set-property property="ase.installation.mode"
value="shared" />
```

with:

```
<set-property property="ase.installation.mode"
value="private" />
```

9. Restart the Sybase Control Center agent on each node in the cluster using the private installation directories. From the \$SYBASE directory, enter `SCC-3_2/bin/scc.sh`.
10. If you have more than one node in the cluster or have changed the SYBASE installation location, deploy a Sybase Control Center agent plug-in for each node:

a) Start **sybcluster**. For example, enter:

```
sybcluster -U uafadmin -P -C mycluster
-F "blade1, blade2,blade2"
```

b) Deploy the plug-in on each node individually. For example, enter:

```
deploy plugin agent "blade1"
deploy plugin agent "blade2"
deploy plugin agent "blade3"
```

See the *sybcluster Utility* in the *Cluster Users Guide* for complete syntax and usage information for **sybcluster** and the Adaptive Server plug-in.

11. You have now upgraded your shared installation to a private installation. You can start the cluster using **start cluster** command, or the individual nodes using **start instance <instance name>** command.

While issuing either command, you may see the following, including an error message that the cluster ID on the quorum device does not match the master device:

```
INFO - Starting the cluster mycluster instance
mycluster_instancel using the operating system command:
/mysybase1/ASE-15_0/bin/dataserver --quorum_dev= /dev/raw/
raw50m41 --instance_name= mycluster_instancel
INFO - 01:00:00000:00000:2009/06/07 23:09:35.46 kernel Quorum
UUID: 00000000-0000-0000-0000-000000000000
INFO - 01:00:00000:00000:2009/06/07 23:09:35.46 kernel Master
UUID: 91f058aa-bc57-408d-854d-4c240883a6c9
INFO - 01:00:00000:00000:2009/06/07 23:09:35.46 kernel Unique
cluster id on quorum device does not match master device. You may
be using the wrong master device. If this is the correct master,
pass 'create-cluster-id' on the command line to pair the devices.
```

When this occurs, reissue the same command, but add **create-cluster-id** as suggested in the message to pair the master device, and start the node manually. For example, issue:

```
/mysybase1/ASE-15_0/bin/dataserver --quorum_dev= /dev/raw/
raw50m41 --instance_name= mycluster_instancel--create-
cluster-id
```

The command should now run without an error message.

You have now upgraded your shared installation to a private installation.

To add new nodes to this cluster, you can now use either Sybase Control Center or the **sybcluster** tool to do so. See *Clusters Users Guide* and *Sybase Control Center for Adaptive Server Enterprise*.

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 sybssystemprocs database.

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

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

```
1> sp_helpdb sybssystemprocs
2> go
```

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

Increasing Device and Database Capacity for System Procedures

If you cannot fit the enlarged sybssystemprocs 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("sybssystemprocs")
and d.status & 2 = 2
and not exists (select vstart
                from sysusages u2
                where u2.dbid != u.dbid
                and u2.vstart between d.low and d.high)
```

where:

- *d.name* – is the list of devices to remove from sysdevices.
- *d.phyname* – is the list of files to remove from your computer.

The **not exists** clause in this query excludes devices that are used by sybssystemprocs 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 sybsystemprocs
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 `sybsystemprocs`, 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 sybserverprocs database of the appropriate size on that device, for example:

```
1> create database sybserverprocs 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.

To upgrade from a noncluster version of Adaptive Server to Adaptive Server Cluster Edition, use the **sybcluster** utility.

To upgrade from an earlier version of Adaptive Server Cluster Edition to Adaptive Server Cluster Edition version 15.7 ESD #2, upgrade manually.

Upgrading from One Version of Adaptive Server Cluster Edition to Another

Use this manual method if you are upgrading from an earlier version of Adaptive Server Cluster Edition to version 15.7 Cluster Edition. Start the Cluster Edition with only one instance until the upgrade is complete.

1. If you are upgrading to Adaptive Server Cluster Edition 15.7 ESD #2 from Adaptive Server versions 12.5.4 through 15.x, perform preupgrade tasks.

Preupgrade tasks are not required if you are updating from Adaptive Server version 15.x and higher.

2. Back up all old databases.
3. Verify that you have the old installation, then install the new server into its own installation directory:
 - a) Start the earlier version of Adaptive Server:

Move to the old \$SYBASE directory:

```
cd $SYBASE
```

b) Run the **source** command on the SYBASE script file:

- Bourne shell – `source SYBASE.sh`
- C shell – `source SYBASE.csh`

c) Execute the runserver file:

```
$SYBASE/$SYBASE_ASE/install/RUN_server_name
```

You can also use **sybcluster** to bring up an earlier version of an Adaptive Server cluster. For example:

1. Enter `$SYBASE_UA/bin/uafstartup.sh`
2. Start **sybcluster**:

```
sybcluster -U uafadmin -P password -C testcluster -F
"ibmpoc01-p3:8888"
> start cluster
```

d) In another window, change to the new \$SYBASE directory and **source** SYBASE.sh (Bourne shell) or SYBASE.csh (C shell).

4. Run the reserved word check on the old Adaptive Server:

a) Install the Cluster Edition version of **installupgrade**:

```
isql -Usa -Ppassword -Sserver_name
-i$SYBASE/$SYBASE_ASE/scripts/installupgrade
```

b) Install the Cluster Edition version of **usage.sql**:

```
isql -Usa -Ppassword -Sserver_name
-i$SYBASE/$SYBASE_ASE/upgrade/usage.sql
```

c) Log in to the old Adaptive Server and execute **sp_checkreswords** on all databases:

```
1> use sybssystemprocs
2> go
1> sp_checkreswords
2> go
```

d) Correct any errors the reserved word check reveals.

5. If your "sa" password is set to NULL, create a new password, as Adaptive Server 15.7 ESD #2 requires a password for the "sa" login.

6. Shut down the old Adaptive Server using **isql**.

7. Copy the old Adaptive Server `mycluster.cfg` configuration file from the old \$SYBASE directory to the new \$SYBASE directory.

8. (Required only if you are upgrading from a nonclustered server) Create the cluster input file. For example `mycluster.inp`:

```
#all input files must begin with a comment
```

```
[cluster]
name = mycluster
max instances = 2
master device = /dev/raw/raw101
config file = /sybase/server_name.cfg
interfaces path = /sybase/
traceflags =
primary protocol = udp
```

```

secondary protocol = udp

[management nodes]
hostname = blade1
hostname = blade2

[instance]
id = 1
name = server_name
node = blade1
primary address = blade1
primary port start = 38456
secondary address = blade1
secondary port start = 38466
errorlog = /sybase/install/server_name.log
config file = /sybase/server_name.cfg
interfaces path = /sybase/
traceflags =
additional run parameters =

[instance]
id = 2
name = server_name_ns2
node = blade2
primary address = blade2
primary port start = 38556
secondary address = blade2
secondary port start = 38566
errorlog = /sybase/install/server_name_ns2.log
config file = /sybase/server_name.cfg
interfaces path = /sybase/
traceflags =
additional run parameters =

```

For an example of what this input file must contain, see *The Cluster Input File* on page 87.

Note: The first instance's *server_name* should be the name of the old server from which you are upgrading.

9. (Required only if you are upgrading from a nonclustered server) Add an entry to the interfaces file for each of the instances in your cluster input file (described in the previous step). See *Configuring the Interfaces File* on page 91.
10. Create the quorum device and start the new instance with the old master device.

```

$SYBASE/$SYBASE_ASE/bin/dataserver\
--instance=server_name\
--cluster-input=mycluster.inp\
--quorum-dev=/dev/raw/raw102
--buildquorum
-M$SYBASE

```

Note: The *server_name* you indicate with the **--instance** parameter must be the name of the server from which you are upgrading, and the interfaces file must contain an entry for

this instance. Any additional options such as **-M** must be present in the `RUN_FILE` as the `dataserver` does not read them from the quorum. For complete `dataserver` documentation, see the *Clusters Users Guide*.

If you are upgrading from a 15.0.1 or a 15.0.3 Cluster Edition to a Cluster Edition server version 15.5 or later, use the original quorum device and cluster input file, and specify **--buildquorum=force** to rebuild the quorum and to override the existing one. Determine the raw device used for the quorum device. For the version of the Cluster Edition, use a raw device on shared disks. Do not use a file-system device.

11. (Skip this step if you are upgrading from a 15.0.1, 15.0.3, or 15.5 Cluster Edition to a 15.7 ESD #2 Cluster Edition server) Log in to the instance. Create the local system temporary database devices and local system temporary databases for each of the instances in your cluster. The syntax is:

```
create system temporary database database_name
    for instance instance_name on device_name = size
```

12. Shut down the instance. Log in to the instance with **isql** and issue:

```
shutdown instance_name
```

13. Restart the cluster.

```
$SYBASE/$SYBASE_ASE/bin/dataserver \
--instance=server_name\
--quorum-dev=/dev/raw/raw102\
-M$SYBASE
```

14. Log in to the Cluster Edition and execute **sp_checkreswords** on all of databases. For example, log in to the instance and execute:

```
1> use sybssystemprocs
2> go
1> sp_checkreswords
2> go
```

15. Correct any errors from the reserved word check.

16. Copy and modify the old `run_server` file to new directory. You must edit it to point to binaries in the correct `$SYBASE` directories:

a) Add this argument to the `run_server` file: `--quorum-dev=<path to the quorum device>`

b) Remove these options, as the information is now stored in the quorum device.

- **-c**
- **-i**
- **-e**

17. Start each instance in the cluster:

```
cd $SYBASE/$SYBASE_ASE/install
startserver -fRUN_server_name
```

18. Install the system procedures:

```
isql -Usa -Ppassword -Sserver_name  
-i$SYBASE/$SYBASE_ASE/scripts/installmaster
```

19. If Adaptive Server includes auditing, run `installsecurity`:

```
isql -Usa -P password -S server_name  
-i$SYBASE/$SYBASE_ASE/scripts/installsecurity
```

20. Run `installcommit`:

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

See also

- *Setting Up Local System and Temporary Databases* on page 94
- *Creating Runserver Files* on page 93

Upgrade a Noncluster Version of Adaptive Server to the Cluster Edition Using `sybcluster`

To upgrade from a noncluster Adaptive Server to the Cluster Edition, perform preupgrade tasks, then upgrade using the **sybcluster** utility.

The command line-based **sybcluster** utility allows you to create and manage a cluster. The utility uses the SCC Agent Framework to "plug in" to the Sybase Control Center remote command and control agent on each node in the cluster. The SCC agent processes the **sybcluster** commands that let you manage the cluster. See the *Users Guide to Clusters* for detailed information about **sybcluster**, and the *Sybase Control Center for Adaptive Server* for the SCC Agent Framework.

Before you start the upgrade:

- Back up old databases.
- Install the Cluster Edition. Make sure that the nonclustered Adaptive Server that is to be upgraded and the Adaptive Server Cluster Edition version 15.7 ESD #2 installation directories are located on the same machine. This is the upgrade machine.
- Review how to create and start a cluster, and prepare your worksheet for information you need for the upgrade.

Before you upgrade:

1. Make sure that your server is ready.
2. Upgrade the server to Adaptive Server Cluster Edition version 15.7 ESD #2.
- 3.

After upgrading, whether manually or using **sybcluster**, perform the postupgrade tasks.

Verifying the Cluster Edition Server for Upgrade

Test to make sure the server is ready for the upgrade process.

1. Start **sybcluster**. For example, enter:

```
sybcluster -U uafadmin -P password -F "mynode:8999"
```

where:

- **-U** – is the login name for the Sybase Control Center agent. "uafadmin" is the default value.
- **-P** – is the SCC agent password.
- **-F** – specifies the node on which the Adaptive Server is running and the listening port for the SCC agent. The default value is 9999.

In this example, the node name is "mynode" and the SCC Agent Framework listening port is 8999.

Note: You must specify the node on which the upgrade is to take place.

2. Start the verification. From the **sybcluster** command line, enter:

```
upgrade server server_name checkonly
```

For example:

```
upgrade server myserver checkonly
```

3. Respond to the **upgrade server checkonly** prompts.

If the verification process succeeds, you can upgrade the server.

If **sybcluster** displays error messages, resolve all error messages resulting from the check before upgrading the server.

See also

- *Responding to upgrade server Prompts* on page 130

Upgrading the Cluster Edition Server Using an Input File

After successfully completing the verification procedure, you can upgrade the server using an input file.

1. Start **sybcluster**. For example, enter:

```
sybcluster -U uafadmin -P -F "mynode:8999"
```

2. Start the upgrade. For example, enter:

```
upgrade server myserver file "/data/myserver_ce.xml"
```

In this example, `/data/myserver_ce.xml` is the configuration file created during the validation procedure.

Note: The configuration file is similar to that used by **sybcluster** to create a cluster.

3. **sybcluster** displays these prompts:

- a) Enter the name of an sa login on server `server_name` [sa] – press **Enter** to accept the default, or enter the name of another user with sa privileges on the target server.

b) Enter a password for the sa login.

sybcluster begins the upgrade procedure, displays informational messages as the upgrade proceeds, and confirms the success of the upgrade.

Upgrading the Cluster Edition Server Interactively

Upgrade the old server interactively.

1. Start **sybcluster**:

```
sybcluster -U uafadmin -P -F "node_name:port_number"
```

2. Start the upgrade:

```
upgrade server server_name
```

3. Respond to the **upgrade server** prompts.

If **sybcluster** displays error messages, resolve all error messages resulting from the check before upgrading the server.

See also

- *Responding to upgrade server Prompts* on page 130

Responding to upgrade server Prompts

Run the **upgrade server** command to upgrade Adaptive Server, as well as to test its readiness for upgrading.

Respond to the **upgrade server** prompts. Default values appear in square brackets. The prompts are:

1. Enter the name of the cluster. [*server_name*]: Enter the existing Sybase installation directory for server *server_name*.
2. Enter the existing Sybase installation directory for server *server_name* using the complete path to the Sybase installation directory, such as `/myserver/sybase15`.
3. Enter the name of the subdirectory containing the ASE installation for server *server_name*, such as `ASE-15_0`.
4. Enter the name of the subdirectory containing the OCS installation for server *server_name*. For example, enter `"OCS-15_0"`.
5. Enter the name of an sa login on server *server_name*, such as "sa."
6. Enter the password for the sa login.

Note: Although you do not need a password for the sa login, you must have a password to create a cluster.

7. Enter the maximum number of instances. [4]
8. Enter the number of nodes (hardware hosts) that will participate in this cluster. [1]
9. Enter the number representing the cluster node 1. [1]

10. Enter the full path to the quorum disk.
11. Enter any trace flags.
12. Enter the full path to the existing master device.
13. Indicate whether this cluster will have a secondary network. [Y]
14. Enter the port number from which this range will apply. [15100]
15. Enter the SYBASE home directory. [Your default location.]
16. Enter the environment shell script path. [Your default location.]
17. Enter the ASE home directory. [Your default location.]
18. Enter the interfaces directory. [Your default location.]
19. Enter path to the dataserver configuration file. [Your default location.]
20. Enter the name of the first instance in this cluster:

Note: The upgrade process creates a cluster with a single instance; add instances later on using the **add instance** command.

21. Enter the interface file query port number for *instance_name*, making sure that the instance query port number is the same as the query port number for the old nonclustered server.
22. Enter the local system temporary database device name.
23. Enter the local system temporary database device path.
24. Enter local system temporary database device size (MB).[100]
25. Enter the local system temporary database name, for example, [*cluster_name_tdb_1*]
26. Enter the local system temporary database size (MB).[100]
27. Specify whether you want save this configuration information in a file? [Y]

Note: If you are running **upgrade server** to verify that the server is ready for the upgrade process, you can use this file as input to upgrade the server. You can also use it as input to validate the server for numerous subsequent upgrades. For example:

```
upgrade server server_name file "/data/myserver_ce.xml" checkonly
```

28. Enter the name of the file to save the cluster creation information.
29. Create the cluster now? [Y]

If you enter Y and you are running **upgrade server** to:

- Test the server to verify that it can be upgraded, **upgrade server** begins the verification process and displays informational messages as the process proceeds, but does not perform an actual upgrade.
- Upgrade Adaptive Server, **sybcluster** begins upgrading the server, displaying informational messages as the process proceeds.

Installation of Version 15.7 ESD #2 over an Existing 15.7 or 15.7 ESD #1 Cluster Edition

Use a binary overlay to install Adaptive Server 15.7 ESD #2 Cluster Edition over an existing 15.7 or 15.7 ESD #1 Adaptive Server Cluster Edition.

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 38

Determining the Adaptive Server Version

Before using the binary overlay, verify that your current version of Adaptive Server is at least 15.7.

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 15.7 or 15.7 ESD #1:

- 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 15.7, you cannot use the binary overlay, but must use an upgrade method instead. See *Chapter 11, Adaptive Server Upgrades* on page 107.

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.

- 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 15.7.x.

- From your SYBASE directory, shut down Adaptive Server through the **shutdown cluster** command in **isql**. 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.
- Shut down all Sybase Control Center agents on all nodes:
`$SYBASE/SCC-3_2/bin-scc.sh -stop`
- Use the installer to load the new software from the CD or DVD.
 Navigate to the CD or DVD drive and enter `./setup.bin`.
- Install Adaptive Server 15.7 ESD #2 into the existing `$SYBASE` installation path.
- At the end of the installation process, select **Configure Sybase Control Center**, and set a new password for both the SCC admin and agent.

Note: If the current Cluster Edition installation is created on a shared disk, enable SCC shared disk mode by selecting **Yes for Do you want to enable SCC Shared Disk mode?**

- Restart Adaptive Server by using **sybcluster**. For example, to restart a cluster on a 2-node cluster:

Start SCC agent on each node:

```
nohup $SYBASE/SCC-3_2/bin/scc.sh -instance linstrs3 > & scc3.out &
nohup $SYBASE/SCC-3_2/bin/scc.sh -instance linstrs4 > & scc4.out &
sybcluster -U uafadmin -Psybase123 -C testcluster -F
"linstrs3:9999,linstrs4:9999" > connect testcluster> start
cluster
```

Note: Before starting **sybcluster**, perform the workaround steps in CR #713282 as described in the "Known Problems" section of the Adaptive Server release bulletin.

- Run **select @@version**. The server should now be at version 15.7 ESD #2.
- 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**.
- 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 10, Postinstallation Tasks* on page 75

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.

Warning! Running **installpcidb** drops and re-creates the `sybpcidb` database and all of the tables and data. The configuration returns to a default state. If you have made any changes to the pluggable component interface (PCI) configuration, you must make the same changes again after running **installpcidb**.

1. Run **dbcc upgrade_object()** for each database to explicit re-compile the objects
2. After upgrading from Adaptive Server Cluster Edition version 15.0.1 to 15.5 Cluster Edition, you must execute a one-time procedure for each of your user-created nodes to update the JAR and XML files.
3. 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)
```

4. 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 147 if you are using **dbcc upgrade_object** to upgrade compiled objects such as check constraints, rules, triggers, and views.

5. 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.
6. 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 147

Updating JAR and XML Files

After the upgrade process completes, Sybase recommends that you configure the Sybase Control Center agent by updating its associated JAR and XML files to manage your cluster. This allows you to use the **sybcluster** utility and the Adaptive Server plug-in to manage the cluster.

1. Start the SCC agent.
2. Run the **sybcluster** utility.
3. Use the **sybcluster deploy plugin** command to configure the SCC agent for your cluster.

If your XML files are from the \$LOWVERSION_SYBASE, then refer to the lower-version installation location. If they are from the higher version number, \$HIGHVERSION_SYBASE, refer to the higher-version installation location.

4. Compare the versions, then add the lower version `agent-plugin.xml` unique properties to the higher version `agent-plugin.xml` properties.

```
$LOWVERSION_SYBASE/SCC-3_2/nodes/Host Name/plugins/Cluster OR InstanceName/agent-plugin.xml with $HIGHVERSION_SYBASE/SCC-3_2/nodes/Host Name/plugins/Cluster OR InstanceName/agent-plugin.xml,
```

For example, copy this property from the lower-version XML file to the higher-version XML file.

```
<set-property property="ase.database.devices"
value="[1,sysprocsdev,/work/river/ase1501ce/data/
sybsystemprocs.dat,150,135;0,master,/work/river/ase1501ce/data/
master.dat,80,60;2,systemdbdev,/work/river/ase1501ce/data/
sybsystemdb.dat,12,12;3,1stdb1,/work/river/ase1501ce/data/
1stdb1.dat,50,50]" />
```

5. Repeat this procedure for all of the nodes in the cluster.
6. Shut down the SCC agent.

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 147.
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 a nonclustered version to the Cluster Edition of Adaptive Server.

To migrate, you can:

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

Migrating Data Using a Dump-and-Load Method

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

1. Run **dbcc** checks on all databases in your original server and make sure they run cleanly.
2. Create a Cluster Edition server in a new directory.
3. Create devices and databases to match those in the Cluster Edition server. Make sure the *sysusages* mapping is correct.

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

4. Dump the databases from the original server.
5. Load the databases to the Cluster Edition server.
6. If you have partitioned tables, update partition statistics.
7. Run **dbcc** checks on the Cluster Edition server and make sure they run cleanly.

For information on upgrading compiled objects, see *Finding Compiled Object Errors Before Production* on page 148.

Migrating a Server Configured for High Availability

If you use high availability, migrate from an Adaptive Server configured with Sybase active-active or active-passive high availability, before upgrading to the Cluster Edition.

See *Using Sybase Failover in a High Availability System* for more information about Adaptive Server for high availability.

1. Use the appropriate cluster subsystem command to stop monitoring resources associated with Adaptive Server on each node.
2. If Adaptive Server is configured for active-active high availability, drop the high-availability companionship:

a) For an asymmetric configuration, issue this command on the secondary companion:

```
sp_companion primary_companion_name, "drop"
```

b) For a symmetric configuration, issue this command on the primary and the secondary companion:

```
sp_companion companion_name, "drop"
```

3. (If you are running an active-active configuration) Issue this command on both companions to verify that both are in single-server mode:

```
sp_companion
```

Each companion issues this message if it is in single-server mode:

```
Server 'server_name' is currently in 'Symmetric normal' mode.
```

4. Remove the resources in the cluster subsystem created for high availability such as resource groups. These are not necessary in the Cluster Edition.
5. Disable the configuration option **enable HA** on the companion server. Log in to the companion and issue:

```
sp_configure "enable HA", 0
```

In an active-active configuration, you must do this on both companions.

6. Follow the steps described in *Using bcp to Migrate Data* on page 141 to upgrade Adaptive Server to the Cluster Edition.

Using bcp to Migrate Data

You can use the **bcp** utility to migrate to the Cluster Edition from any Adaptive Server version.

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. Run **dbcc** checks on databases to verify data integrity.
2. Use **bcp** to extract all the data from all the tables in the databases.
3. Create a new Cluster Edition 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 Cluster Edition server and make sure they run cleanly.

Adaptive Server Components and Related Products

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

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 14. Changed Job Scheduler Templates

Template	Modified Files	Datatype for <parameter> Changed from <code>varchar(5)</code> to <code>int</code>
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
rebuild table	SybRebuildTableTemplate.xml jst_reorg_rebuild_tables	@dump_flag

Template	Modified Files	Datatype for <parameter> Changed from varchar(5) to int
reclaim indexes	SybReclaimIndexTemplate.xml jst_reclaim_index_spac	@dump_flag
reclaim tables	SybReclaimTableTemplate.xml jst_reclaim_table_space	@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 15. 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.
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.

Template	Modified File	Change Description	Job Modifications
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 Java in the Database

The pluggable component interface (PCI) configuration values are used by all of the nodes in a cluster. To share files between nodes, using working directories, you need shared directory locations in the file system that are visible to all nodes using the same path.

Each node has its own separate and local temporary directory usually `/tmp` on UNIX.

1. Run the `installpci` script on only the first node, not on any other nodes in the cluster. A single instance of the PCI configuration is shared between the nodes of a cluster.
2. Create the `sybpcidb` database before running the `installpcidb` script. You can choose the location and size of this database. As long as the database is named `sybpcidb` the `installpcidb` script works correctly. For example:

```
1> disk init
2> name = "sybpcidb_dev",
3> physname = "/dev/raw/raw20",
4> size = '24MB'
5> go
1> create database sybpcidb on sybpcidb_dev = 24
2> go
```

3. Once the `sybpcidb` database has been created, run the `installpcidb` script to create and populate the tables that hold the PCI and pluggable component adapter (PCA) configuration properties. Stored procedures and tables are created in the `sybpcidb` database. Use `isql` to run the `installpcidb` script that is included with the release. Save the output to an operating system file. For example:

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

4. 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.

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 124

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 determine whether a compiled object has been upgraded, look at the `sysprocedures.version` column. If the object was upgraded, this column contains the number 12500.

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 12 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 155.

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 Cluster Edition to Adaptive Server Cluster Edition versions 15.5, 15.0.3, and 15.0.1.5. Downgrading to earlier versions of Adaptive Server is not supported.

Note: The **sybcluster** utility does not have a built-in tool to perform an upgrade or downgrade between two versions of Adaptive Server Cluster Edition.

1. If you have used pluggable component interface (PCI) functionality, and are downgrading to a version of the 15.0.1 Cluster Edition or a 15.0.1 Cluster Edition ESD, drop the `sybpcidb` database and restart the old server binary.
2. If you are downgrading to version 15.0.3 Cluster Edition, re-run **installpcidb** and restart the old server binary.
3. Start the 15.7 ESD #2 Cluster Edition server in single-user mode, with only one instance running, using **dataserver -m** option to guarantee that no other user can access Adaptive Server while you are performing the downgrade. See *Utility Guide* for more details on starting the server in single-user mode.
4. 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'
```

For downgrade purposes, Cluster Edition is differentiated from a nonclustered Adaptive Server by a "5" at the end of the version string regardless of the version to which you want to downgrade. If you are downgrading to Adaptive Server Cluster Edition version 15.5, for example, specify your target version as "15.5.0.5". You may omit the periods and use "15505", but the number must still end with a 5. Otherwise, the procedure assumes "0" as the final digit, indicating a nonclustered Adaptive Server. This produces an error, as you cannot downgrade from Cluster Edition to any nonclustered server.

5. Make sure Adaptive Server 15.7 ESD #2 Cluster Edition is prepared for downgrade by executing **sp_downgrade** from the master database:

```
sp_downgrade 'prepare', @toversion=<'version'>
, @override = 1
```

Write the value of "version" as "15.0.1.5", "15015", "15.0.3.5", or "15035". The version you enter is the version to which you are downgrading.

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.

6. 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 to shut down the cluster.

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**.

7. Copy the RUN_SERVER file to the release area of the version to which you are downgrading. Modify the RUN_SERVER file to use the dataserver binary from the release area of the version to which you are downgrading. Remove the **-m** option when you modify the RUN_SERVER file or ensure the server is not running in single-user mode.

8. Extract the quorum device using:

```
qrmutil --quorum-dev=<dev> --extract-config=quorum.out
```

If you are downgrading to Cluster Edition version 15.0.1 or a 15.0.1 ESD, comment out these two lines in quorum.out:

```
#installation mode=shared
#membership mode=
```

9. Rebuild the new quorum device with the old quorum device backup you created by adding this to the **dataserver** parameters:

```
--buildquorum=force --cluster-input=quorum.out
```

10. Restart the old server by first connecting to the cluster in sybcluster by running `connect` to `cluster_name`.

After you connect to the cluster and the cluster name appears at the prompt, issue the following to start Adaptive Server:

```
start cluster
```

Note: If you restart a downgraded server using the 15.7 configuration file, the new options cause the message `Unknown parameter`. This message is reported only the first time the server is started. When downgrading to a 15.0.1 Cluster Edition server, user passwords may be reset and printed on the console. Save the output of **sp_downgrade 'downgrade', <version>** into a file so that the passwords are not lost. The old passwords do not work. Otherwise, you must restart the downgraded server with **-psa** to regenerate the sa password.

11. Depending upon whether you are starting from a fresh installation or from a previously upgraded 15.x installation, run previously saved `restore_msgs.ebf` or `instmsgs.ebf` from 15.0.1 Cluster Edition or 15.0.3 Cluster Edition if none were saved, before running **installmaster** for the Cluster Edition.
12. Restart all instances of the cluster.
13. 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
```

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 159
- *Additional Considerations for New Features Used* on page 157

- *Downgrading Job Scheduler* on page 159

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 152
- *Dumping a 15.7 ESD #2 Database to Load into a 15.7 or 15.7 ESD #1 Adaptive Server* on page 156

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 155

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.

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>ddl_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>
Unicode noncharacters	<p>If you used <code>sp_configure</code> 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.</p> <p>To avoid such errors, search for and remove those characters from the system.</p>
Disabled varbinary truncation	<p>If you used <code>sp_configure</code> to turn on the <code>disable varbinary</code> truncation 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.</p>
Sharable inline default	<p>If you used <code>sp_configure</code> to enable the sharable inline default setting in Adaptive Server 15.7 and downgrade to a version earlier than 15.7, running the <code>ddlgen</code> 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.</p>

Downgrading Job Scheduler

If you are downgrading to Adaptive Server versions 15.0.1 Cluster Edition, 15.0.1 Cluster Edition ESD, or 15.0.3 Cluster Edition, 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 152

Post-Downgrade Tasks for Adaptive Server

After you perform the basic downgrade steps, finish the process by handling residual tasks.

Sybase certifies that each version of Adaptive Server can connect to earlier releases through Component Integration Services. Component Integration Services is tested and certified to connect to earlier versions of Adaptive 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 not available in the earlier release.

For example, if you have a local and a remote server, both running Adaptive Server version 15.7 ESD #2 Cluster Edition, upgrade the local server first, if you are performing an upgrade

on each of them. If one is to be upgraded but not the other, the local server should be upgraded first.

1. After you run `installmaster`, `installcommit`, `installsecurity`, and `installmsgsvss` on your downgraded server, any changed system stored procedures are returned to their earlier form. Any new stored procedures introduced to Adaptive Server are not removed. Attempts to execute such stored procedures against an older binary may generate 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**, **xmtable**, 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 157
- *Reenabling Replication Server After the Dataserver Upgrade* on page 138

CHAPTER 13 **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 13: 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

CHAPTER 14 **Troubleshoot the Server**

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 SySAM license and if the grace period has expired. Check the Adaptive Server error log for the cause of the license checkout failure and fix the issues.
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 177
- *Recovering from a Failed Installation* on page 177

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
<p>Cannot use X-Windows.</p>	<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 data-bbox="467 510 1176 760"> % 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>
<p>Cannot eject the CD or DVD from the drive.</p>	<p>If you cannot eject the CD from the drive, check:</p> <ul data-bbox="467 921 1166 1046" 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.
<p>DISPLAY environment variable not set correctly.</p>	<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.	<p>Enter a different port number on the srvbuild window. The command netstat -a produces a list of port numbers in use.</p>
Adaptive Server fails to start.	<p>The shared memory of the operating system may not be set high enough. Adjust the shared memory value. Restart the installation or upgrade process.</p>
The installer fails to start.	<p>Restart the installer in debug mode. Set the environment variable LAX_DEBUG to true before running the installer.</p>
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 data-bbox="467 447 817 473">% 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.

See also

- *Error Logs for Installation Utilities* on page 174

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 **sybcluster** or a manual upgrade to complete the upgrade session.

If the Cluster Edition Configuration Fails

Check the location of the system shared libraries.

The shared libraries `libXt.so` and `libX11.so` are normally stored in `/usr/openwin/lib`. The shared library `libsocket.so` is located by default in `/usr/lib`.

If the shared libraries on your operating system are located in other directories, set the environment variable `LD_LIBRARY_PATH` to indicate the new locations.

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 **sybcluster** 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.

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.

Re-running the Cluster Edition Upgrade

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

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

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

2. 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:

```
dbcc traceon(990)
```

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

3. To restart a failed upgrade::

```
online database <failed_db_name>
```

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

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

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

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

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

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

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

If 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 15, Obtaining Help and Additional Information* on page 183

If Lack of Space Prevents an Upgrade

During the upgrade process, the transaction log may fill due to system table changes. A full transaction log may prevent a successful Adaptive Server upgrade.

If you cannot continue with the upgrade process because the transaction log becomes full, log in to the new server and in **isql**, issue:

```
dump tran dbname with no_log
```

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

In some cases, preupgrade space estimations are insufficient for the data copy phase of the upgrade process. In this case, an error is generated letting you know there is insufficient space in the *system* segment for the upgrade. The upgrade process stops and waits for space to be provided. Log in to the new server using **isql**, and increase the size of the database:

```
alter database dbname on device_name = "2m"
```

Note: Specify the size to alter the database with the unit specifier "m", or "M".

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 15: 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 34
- Adaptive Server
 - installing in command line mode 48
 - installing in silent mode, unattended 48
- Adaptive Server components 141
- Adaptive Server editions
 - determining which edition you have 6
- Adaptive Server Enterprise Cluster Edition Cluster Users Guide 3
- Adaptive Server plug-in for Sybase Central 7
- additional steps for downgrading new features 157
- adjusting
 - shared memory parameters in operating system 38
- administration tasks 34
- applications
 - changing after renaming objects 116
- auxiliary servers
 - configuring 69

B

- Backing up the server 132
- Backup Server
 - configuring 70
 - configuring with sybcluster 69
- basic downgrade steps using sp_downgrade 152
- bcp
 - migrating with 141
- binary overlay 132

C

- check out errors 161
- checklist before installing Cluster Edition 58
- choosing
 - the correct license 161
- client applications 8
- cluster
 - preinstallation checklist 58
- cluster, starting 95
- clusters
 - reconfiguring 97

- starting after system failure 96
 - taking over automatically 95
 - worksheet for creating a cluster 60
- command line installation
 - Replication Server 48
- command line options 49
- commands
 - showserver 75
- Component Integration Services
 - upgrading local and remote servers 110
- configuration
 - manual 91
- configuration parameters 159
- configure
 - XP server using sybcluster 72
- configuring
 - Backup Server 70
 - master device 92
 - quorum device 92
- configuring Adaptive Server
 - troubleshooting 178
- configuring the cluster 64
- conflicting object names 116
- creating
 - sybase user accounts 34
- creating a cluster with sybcluster 64
- creating the cluster
 - starting the cluster 57

D

- data compression
 - described 7
- databases
 - image data for sample 80
 - increasing capacity 122
 - jpubs 81
 - optional 79
 - setting up local system and temporary databases 94
- determining
 - installation process 3
- determining the current edition 6
- determining whether select* should be changed in views 150

Index

- devices
 - increasing capacity 122
- directory changes 111
- directory locations 32
- directory services
 - described 7
- downgrading
 - Adaptive Server 151
 - handling new features 157
 - sp_downgrade_esd, using 155
 - sp_downgrade, using 152
 - steps to prepare beforehand 151
- downgrading job scheduler 159
- dsedit utility 138
- dump and load 15.7 ESD #2 database into 15.7 or 15.7 ESD #1 156

E

- enabling SCC 98
- Enabling the Java in the database in a high availability system 146
- encrypted columns
 - described 7
- error
 - information, finding 161
 - license checkout 161
- error log files
 - troubleshooting and 177

F

- failed installation
 - cleaning up 69
- failed upgrade
 - recovery 181
- fails preupgrade eligibility 178
- failures
 - Cluster Edition configuration 178
 - rerunning a Cluster Edition upgrade 179
- features, separately licensed optional 23
- file descriptors 35
- finding
 - license error information 161
- finding compiled object errors before production 148

G

- generating
 - served licenses 26

- unserved licenses 26
- grace period 20, 161
 - cannot check out 161
 - license continues to run after fixing 161
- GUI installation 43

H

- heartbeat 20
- host IDs
 - determining 21
 - lmutil utility 21
 - on Windows platforms 22
 - with multiple network adapters 22

I

- I/O fencing
 - SCSI-3 PGR 84
 - setting up 83
- I/O Fencing
 - data integrity 82
 - enabling 85
 - setting up 83
- identifiers
 - conflicting with reserved words 116
- in-memory database
 - described 7
- increasing
 - device and database capacity 122
- infiniband 14
- input file
 - create 87
- input/output completion port 13
- installation
 - Adaptive Server CD, mounting 43
 - determining process via workflow 3
 - overview 3
 - warning, cannot find 161
- installation steps for PC-Client 51
- installer suddenly quits 177
- installing
 - command line 47
 - with response file 47
- installing Adaptive Server
 - GUI mode 43
- installmaster, installing system stored procedures
 - with 93
- installpix script 80

installpubs2 script 80
 installpubs3 script 80
 interfaces file 138
 IOCP status 13

J

Java in the database
 preparing before installing Adaptive Server
 38
 upgrading 146
 Job Scheduler
 installing 71
 jpubs database
 installing 81

K

keywords 159

L

LDAP libraries
 environment variables 77
 location of 77
 license
 check out errors 161
 license deployment models
 comparison of 19
 recommendations 19
 license error information, finding 161
 license server
 cannot find licenses for optional features 161
 continues to run in grace period 161
 license checkout error 161
 starts in grace period 161
 license servers 20
 licenses 29
 recommendations for license models 19
 regenerating 28
 rehosting 28
 renewing 28
 licensing 17
 licensing optional features 161

M

manual configuration 91
 enabling sybcluster 98

methods of installation
 Adaptive Server 43
 migrating
 dump and load method 140
 high availability 140
 methods 139
 modifying individual licenses 29
 Monitor Server
 configuring with sybcluster 69
 mount commands by platform 43
 mounting the Replication Server installation media
 43
 multipathing
 described 39
 setting up before installation 40
 multiple simultaneous asynchronous I/O operations
 requirements for processing 13

N

network protocol 14
 new features
 reverting to earlier version 157

O

objects
 conflicting names of 116
 renaming 116
 optional databases 79
 optional features
 described 7
 errors 161
 license server cannot find 161
 options for server editions
 packages 7
 overview
 installation 3
 overview of the upgrade process 107

P

partitions
 described 7
 PC-client
 CD 8
 PC-Client
 directory layout 32
 product description 32

Index

- system requirements for 14
 - planning the installation 17
 - platform
 - mount commands 43
 - post downgrade 159
 - preinstallation tasks 35
 - preupgrade tasks
 - before upgrade 113
- ### Q
- quoted identifier errors 149
 - quoted identifiers 117
- ### R
- raw devices, accessing 87
 - raw disks
 - I/O fencing 83
 - multi-node access 83
 - recommendations
 - license models 19
 - reconfiguring the cluster 97
 - recovering from a failed installation 177
 - redundancy 20
 - reenable Replication Server 138
 - related products 141
 - release bulletin 17
 - removing
 - registry entries 50
 - removing existing server 50
 - renaming objects 116
 - replicated databases
 - restoring 139
 - Replication Server 138
 - draining logs for Adaptive Server 138
 - reserved words 116
 - conflicts 116
 - response file 48, 52
 - restoring from backup 179
 - running in a grace period 161
 - runserver file 115
- ### S
- sample databases 79
 - SAP Service Marketplace 17
 - SCC
 - starting 55
 - scripts
 - installpix 80
 - installpubs2 80
 - installpubs3 80
 - location 79
 - SCSI-3 PGR 84
 - security services
 - described 7
 - separately licensed optional features 23
 - served licenses 26
 - license servers 20
 - server
 - advantages to the cluster edition 5
 - component descriptions 30
 - directory layout 30
 - installation options 5
 - overview 5
 - single-system presentation 5
 - workload management 5
 - servers, license 20
 - setting up
 - cluster manually 86
 - shared disk cluster
 - private interconnect 14
 - shared installation 57
 - shared memory parameters
 - adjusting 38
 - showserver command 75
 - silent installation 52
 - simulated cluster on a single node 59
 - source environment variables 43
 - sp_checkreswords system procedure 116
 - sp_downgrade system procedure 152
 - sp_downgrade_esd 155
 - dump and load database 156
 - sp_lmconfig
 - determining current edition 6
 - SPDC
 - generating served licenses 26
 - generating unserved licenses 26
 - starting
 - cluster using sybcluster 67
 - cluster when operating system starts 68
 - Sybase Control Center 55
 - stopping
 - Adaptive Server after a failure 177
 - cluster using sybcluster 67
 - servers 177
 - Sybase Control Center 55

- Storage Foundation
 - creating a cluster 60
 - stored procedures
 - changing after renaming objects 116
 - style conventions
 - fonts 1
 - Sybase Central 7
 - Sybase Control Center 7
 - starting 55
 - stopping 55
 - SYBASE directory 43
 - Sybase Product Download Center
 - login page for Sybase reseller sales 17
 - sybase user
 - creating accounts 34
 - sybcluster
 - configure XP Server 72
 - configuring auxiliary servers 69
 - configuring the cluster 64
 - creating a cluster 64
 - starting a cluster 67
 - stopping a cluster 67
 - upgrading a noncluster to cluster Adaptive Server 128
 - sybsystemprocs database
 - increasing the size of 121
 - SySAM licensing 17, 23
 - license servers 20
 - problems when Host ID is using virtual network adaptor address 161
 - troubleshooting 161
 - sysmessages 159
 - system administrator password
 - setting 78
 - system requirements
 - checking for upgrading 112
 - IBM 11
 - PC-Client 14
 - private interconnect technology 14
 - system stored procedures 159
 - installing using installmaster 93
 - system tables 159
 - changes due to upgrade 115
- T**
- tasks
 - after configuring 96
 - tasks, administration 34
 - test environment
 - creating 77
 - Tivoli Storage Manager
 - described 7
 - troubleshooting 161
 - an unrecognisable license 161
 - by using error log files 177
 - calling technical support 171
 - installation problems 175
 - installing license server components 161
 - license check out errors 161
 - licensing optional features 161
 - picking the correct license 161
 - running in a grace period 161
 - troubleshooting Adaptive Server installations 173
- U**
- uninstalling
 - PC-Client 53
 - server 50
 - the old server 50
 - unknown cause of failure 178, 180
 - unrecognized tokens or objects 159
 - unserved licenses 26
 - updating JAR files 135
 - updating XML files 135
 - upgrade requirements 112
 - upgrading
 - Adaptive Server 108
 - entire installations 108
 - failure 178
 - input file 129
 - interactively 130
 - Java in the database 146
 - single databases 108
 - sybcluster, with 128
 - tasks after you upgrade 134
 - using bcp 141
 - using upgrade server command 130
 - upgrading a noncluster to cluster Adaptive Server sybcluster 128
 - upgrading Adaptive Server
 - Cluster Edition to Cluster Edition 124
 - upgrading Cluster Edition
 - shared to private installation 117
 - upgrading compatibility
 - versions 107
 - upgrading compiled objects
 - dbcc upgrade_object 147
 - upgrading Job Scheduler templates 142

Index

utilities 8
 dsedit 138

V

verifying connection to the server 76
verifying readiness for upgrade 128
virtual network adaptor license problems 161

W

warning, installation, cannot find 161

X

XP Server
 configure with sybcluster 72
 configuring with sybcluster 69