



Installation Guide

**Adaptive Server[®] Enterprise
Cluster Edition 15.7**

IBM AIX

DOCUMENT ID: DC01086-01-1570-02

LAST REVISED: June 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

Conventions	1
Installation Task Overview	3
Installation Workflows	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
Planning Your Adaptive Server Installation	11
Adaptive Server Release Bulletin	11
Obtaining a License	11
SySAM License Server	13
IPv6 Configuration	14
SySAM Licensing Checkout	14
System Requirements	14
Making Input/Output Completion Port API Available for Adaptive Server	16
Cluster Edition System Requirements for Using Private Interconnect Technology	17
System Requirements for Clients	18
Installation Directory Contents and Layout	18
PC-Client Product Descriptions and Layout	20
Performing Administrative Tasks	22
Creating the Sybase User Account	22
Preparing to Install Adaptive Server	23
Adjusting the Operating System Shared Memory Parameters	26
Managing Java in the Database During Installations and Upgrades	26

- Multipathing27
- Installing Adaptive Server31**
 - Mounting the CD31
 - Installing Adaptive Server with the GUI Wizard31
 - Installing Adaptive Server in Console Mode34
 - Installing Adaptive Server Using a Response File35
 - Creating a Response File35
 - Installing in Silent Mode36
 - Command Line Options36
 - Uninstalling Adaptive Server37
 - Removing an Existing Adaptive Server37
- Installing Components from the PC-Client39**
 - Creating a Response File40
 - Installing the Client in Unattended, Silent Mode40
 - Uninstalling PC-Client41
- Creating and Starting the Cluster43**
 - Differences Between Private and Shared Installations
.....43
 - Cluster Edition Preinstallation Checklist44
 - Creating a Simulated Cluster on a Single Node
.....45
 - Starting and Stopping Sybase Control Center Agent ... 45
 - Before Creating a Cluster46
 - Creating a Cluster With Storage Foundation for
Sybase Cluster Edition46
 - Worksheet for Creating a Cluster47
 - Create a Cluster Using sybcluster50
 - Starting the Cluster When the Operating System
Starts55
 - Cleaning Up After a Failed Cluster Installation55
 - Auxiliary Servers55
 - Configuring Backup Server Using sybcluster56
 - Installing Job Scheduler57
 - Configuring XP Server Using sybcluster59
- Postinstallation Tasks61**

Verifying That Servers are Running	61
Verifying Your Connection to the Servers	62
Testing the Installation and Network Connections	62
Creating a Testing Environment	63
Configuring libtcl.cfg for LDAP	63
Adding a Server to the Directory Service	64
Setting the System Administrator Password	64
Installing Sample Databases	65
Default Devices for Sample Databases	65
Running the Database Scripts	66
Installing the interpubs Database	66
Installing the jpubs Database	67
Maintaining Sample Databases	68
I/O Fencing	68
Preparing to Set Up I/O Fencing	69
Enabling I/O Fencing	70
Creating Logical Volumes	71
Manual Cluster Setup and Management	72
Setting the Environment	72
Raw Devices	72
The Cluster Input File	73
Manually Configuring the Cluster	76
Automatic Takeover of a Cluster	80
Starting the Cluster	80
Postconfiguration Tasks	81
Shutting Down the Cluster or the Instance	81
Reconfiguring the Cluster	82
Enabling sybcluster and Sybase Control Center After a Manual Configuration	82
Adaptive Server Upgrades	91
Upgrading Adaptive Server	92
Considerations for Component Integration Services	94
Preparing to Upgrade	94
Changes in Adaptive Server Directories	95
Checking System and Upgrade Requirements	96

Performing Preupgrade Tasks	96
Upgrade System Table and Stored Procedure ...	99
The runserver File Location	99
Reserved Words	99
Upgrading to a Private Installation	101
The sybprocsdev device	104
Upgrade to Adaptive Server 15.7	108
Upgrading from One Version of Adaptive Server Cluster Edition to Another	108
Upgrade a Noncluster Version of Adaptive Server to the Cluster Edition Using sybcluster	111
Installation of Version 15.7 over an Existing 15.x Adaptive Server	115
Postupgrade Tasks	116
Updating JAR and XML Files	117
Running the instmsgs.ebf Script	118
Restoring Functionality in Adaptive Server After Upgrading	119
Reenabling Auditing	120
Reenabling Replication Server After the Dataserver Upgrade	120
Migrate	122
Migrating Data Using a Dump-and-Load Method	122
Migrating a Server Configured for High Availability	123
Adaptive Server Components and Related Products .	123
Upgrading Job Scheduler	123
Upgrading High Availability and Cluster Support	128
Upgrading Java in the Database	129
Upgrading Backup Server	130
Upgrading Databases Using Dump and Load ...	130

How Compiled Objects are Handled When Upgrading Adaptive Server	131
Adaptive Server Downgrades	135
Preparing to Downgrade Adaptive Server	135
Downgrading from Adaptive Server 15.7	136
Additional Considerations for New Features Used	138
Downgrading Job Scheduler	140
Post-Downgrade Tasks for Adaptive Server	140
Troubleshoot SySAM	143
Troubleshooting SySAM	144
If the Installation Program Cannot Find a Suitable License	147
If Adaptive Server Cannot Check Out a License and Starts Using a License Under a Grace Period	148
E-Mail Notifications	149
Before You Contact Sybase Technical Support	149
Troubleshoot the Server	151
Error Logs for Installation Utilities	152
Error Logs for Sybase Servers	152
Troubleshooting Common Installation Problems	153
Stopping Adaptive Server After a Failure	155
Recovering from a Failed Installation	155
If Adaptive Server Fails the Preupgrade Eligibility Test	155
If the Cluster Edition Configuration Fails	156
When an Upgrade Fails	156
If You Can Identify the Cause of the Upgrade Failure	156
Restoring Databases After a Failed Upgrade	157
Re-running the Cluster Edition Upgrade	157
If You Cannot Identify the Cause of an Upgrade Failure	158
If Lack of Space Prevents an Upgrade	158
Obtaining Help and Additional Information	161
Technical Support	161

Contents

Downloading Sybase EBFs and Maintenance Reports	161
Sybase Product and Component Certifications	162
Creating a MySybase Profile	162
Accessibility Features	162
Index	165

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.

Installation Task Overview

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

These 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 37.

Adaptive Server Components

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 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).

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

Adaptive Server Components

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 one of 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 is a server application that uses a Web browser-based client to deliver an integrated solution for monitoring and managing Sybase products.

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.

Adaptive Server Components

The Adaptive Server installer installs the Sybase Control Center 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 disc or download it from <http://downloads.sybase.com>. In a production environment, Sybase recommends that you install this 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

In addition to the SDK, the PC-Client CD includes products that install separately:

- InfoMaker – a tool for personal data access, management, and reporting, used by both developers and end users. InfoMaker allows you to create InfoMaker presentation-quality reports and powerful queries. It complements client/server business applications and development tools as well as desktop productivity suites.
- PowerDesigner Physical Architect – a tool for data modeling, including database design, generation, maintenance, reverse engineering, and documentation for database architects.

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

Before installing your product, choose a SySAM license model, determine the license server information, and obtain license files.

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.

1. Choose a SySAM license model:

License Model	Description
Unserviced license model	Licenses are obtained directly from the license file. If you are using an unserved license, save the license file to the machine on which you install the product.
Served license model	A license server manages the allocation of licenses among multiple machines.

- For the served license model, decide whether to use an existing or new license server.

The license server and your product installation need not share the same machine, operating system, or architecture.

- If you choose the served license model:

- Decide whether to use an existing or new license server.

The license server and your product installation do not need to share the same machine, operating system, or architecture.

- To install on a machine that is running a SySAM 1.0 license server, follow migration instructions in the *SySAM Users Guide* to migrate to the newer SySAM version.

Note: There can be only one instance of a SySAM license server running on a given machine. To set up a SySAM 2 license server on a machine that is already running a SySAM 1.0 license server, migrate the old license server to SySAM 2. A migrated license server can serve licenses for both SySAM 1.0 and SySAM-2-enabled products.

- Get your host IDs – when you generate licenses at the Sybase Product Download Center (SPDC) or SAP Service Marketplace (SMP), you must specify the host ID of the machine where the licenses will be deployed. For:

- Unserviced licenses – 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 want to run the product in a virtualized environment, see *SySAM Sub-capacity Licensing* in the *SySAM Users Guide*.
- Served licenses – the host ID of the machine where the license server will run.

To use an existing license server, you must know the server host name and port number.

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.

- Before installing your product, get license files from SPDC or from SMP using the access information you receive from Sybase or your Sybase reseller.

Log in to SPDC or SMP by using the information in your welcome e-mail message.

Note: If you purchased your Sybase software from a Sybase reseller, you receive a Web key rather than an e-mail message.

If you plan to use sub-capacity licensing, see the *SySAM Users Guide* for information on how to configure using the **sysamcap** utility.

SySAM License Server

Make sure that you have installed the required SySAM license server version. Replication Server 15.5 and later include FLEXnet Publisher license server manager, which is platform-specific.

- License server version

You must install the SySAM license server version 2.1 or later. To determine the current license server version, use the **sysam version** command.

Note: License server version 2.0 and earlier do not include this command.

The SySAM license server installer is available in your *ASE_installer_image/sysam_setup* directory. You can also download the latest license server from the SySAM License Server and Utilities Install Web site at <http://www.sybase.com/sysam/server>.

- FLEXnet Publisher version

Table 2. FLEXnet Publisher Versions Supported on UNIX Platforms

Platforms	FLEXnet Publisher Version
HP-UX Itanium (64-bit)	11.6
IBM AIX (64-bit)	10.8.5
IBM p-Series (Linux on POWER) (64-bit)	11.5
Red Hat Enterprise Linux (RHEL) (32-bit and 64-bit)	11.5
SuSE Linux Enterprise Server (SLES) (32-bit and 64-bit)	11.5
Sun Solaris SPARC 8, 9, and 10 (64-bit)	11.5
Sun Solaris 10 x86-64 (64-bit)	11.5

If you use a SySAM license server, you must update the SySAM license server to at least version 2.1, which includes the FLEXnet Publisher version 11.6.1 license server components on all platforms, before installing Replication Server 15.7. To check the version of your license server, examine the license server log or execute:

```
cd $SYBASE/SYSAM-2_0/bin
./lmutil lmver lmgrd
```

IPv6 Configuration

Use the **sysam configure** command to choose the version of the license server based on your environment.

Adaptive Server includes IPv4 and IPv6 versions of the SySAM license server software. By default, the installer configures SySAM for IPv4 support.

The IPv4 version of the license server cannot be used properly if you have enabled IPv6 on the license server host. If this is the case, you can:

- Disable IPv6 TCP/IP protocol on the license server host.
- Use another Windows host that does not have IPv6 enabled.
- Use a UNIX host for the license server. You can use the IPv4 version of the license even if both IPv4 and IPv6 are enabled on UNIX.

If you are using the IPv6 version of the license server on a Windows host, you can access this license server only through the IPv6 protocol. Any host with an IPv4 only network stack cannot obtain licenses from this IPv6 license server. To resolve this issue:

- Use a UNIX host for the license server. License servers on UNIX can serve licenses to both IPv4 and IPv6 client hosts.
- Use two separate Windows machines. Set up two license servers on two different hosts in the network; one to serve the IPv4 network and the other one to serve the IPv6 network.

You can use the SySAM script to configure the appropriate version of the license server:

```
sysam configure [IPv6|IPv4]
```

For example, to configure use of the IPv4/IPv6 dual stack binaries, use:

```
sysam configure IPv6
```

SySAM Licensing Checkout

If Adaptive Server was licensed using a per-processor license type in versions earlier than 15.7, the license quantity was determined at start-up.

Adaptive Server periodically checks the number of processors it can use and attempts to check out additional licenses for any increase. Replication Server shuts down after a 30-day grace period expires if additional licenses are not made available within this period.

See the *SySAM Users Guide*.

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 Requirements for Cluster Edition

Adaptive Server hardware and memory specifications for IBM AIX:

- Processor 64-bit – Power6, Power5, Power4, and PowerPC 970
- Kernel – 64-bit
- Minimum RAM required for Adaptive Server – 272MB
- Default user stack size – 148KB
- Minimum RAM per additional user – Approximately 364kB

Operating system requirements for IBM AIX:

- Hardware – IBM Power, System p, System i, System p5, System i5, eServer p5, eServer i5, eServer pSeries, and BladeCenter blades
- Operating system
 - AIX 6.1 runs on 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 `devices.fcp.disk.rte 6.1.0.3` or higher
- Preferred RAM – at least 1G

AIX 6.x and 7.x require a patch that fixes APAR IV10828, a known AIX issue. Contact IBM directly to obtain the patch.

Before you can install a 64-bit Adaptive Server version 15.7 to a Red Hat Enterprise Linux 6.x x86_64, install the 32-bit GNU C Libraries package (`glibc-2.xx-x.xx.el6.i686.rpm`).

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

Planning Your Adaptive Server Installation

Information about the required operating system patches is available from the <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 Chapter 11, “Using the Cluster Edition with the Veritas Cluster Server,” in the *Clusters Users Guide*. Veritas Cluster Server for Sybase Cluster Edition is only supported 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).

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
% lspp -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:

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 **lsdev** 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 proprietary interconnect with a specialized proprietary communications protocol. 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 amount of inter-instance updates and inter-instance 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.

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/job-scheduler/.

Product	Description
Shared Disk Cluster Administration Tools	<ul style="list-style-type: none"> • Sybase Control Center logs and utilities – files related to the Web-based tool for monitoring the status and availability of Adaptive Servers. They are installed into the <code>\$SYBASE/SCC-3_2</code> directory. • sybcluster – interactive command line interface for configuring and managing the cluster installed in the <code>\$SYBASE/\$SYBASE_ASE/bin</code> 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 <code>\$SYBASE/\$SYBASE_ASE/jobscheduler</code> directory. • Interactive SQL– interactive command line interface for configuring and managing the cluster. It is installed into the <code>SYBASE/DBISQL</code> 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 <code>\$SYBASE/shared/JRE-*</code> directory. • The Cluster Edition Administration Utility is installed to <code>\$SYBASE/SDCADMIN-15_0</code>. • Sybase Central 6.x is a Java-based framework used by the system management tools. Installed into the <code>\$SYBASE/shared/sybcentral600</code> directory. • Adaptive Server plug-in. Sybase Central plug-in that provides cluster configuration and complete administration functionality. Installed into the <code>\$SYBASE/ASEP</code> directory.
Software Developer Kit (SDK)	<p>Connectivity:</p> <ul style="list-style-type: none"> • Open Client™ (Client-Library, dblib) • Embedded SQL/COBOL 15.0 • XA interface for Adaptive Sever <p>Installed into the <code>DataAccess</code> 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.
Language modules	<p>Installed into the <code>\$SYBASE/locales</code> and <code>\$SYBASE_ASE/locales</code> directories. Provides system messages and date/time formats.</p>
Character sets	<p>Installed into the <code>charsets</code> directory. Provides character sets available for use with Adaptive Server.</p>

Product	Description
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 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
DataAccess	<p>Installed into the DataAccess[64] 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.
Windows Cluster Server Admin utilities	<p>Windows Cluster Administrator is a GUI tool for managing Microsoft Cluster Server (MSCS). Use it to create, modify and displays group, resource, and cluster information. There is also an alternative command line tool called <code>Cluster.exe</code> for cluster administration.</p>
PC-Client Administration Tools	<ul style="list-style-type: none"> • sybcluster – interactive command line interface for configuring and managing the cluster installed in the %SYBASE%\SDCADMIN-15_0\bin directory. • Interactive SQL – interactive command line interface for configuring and managing the cluster. It is installed into the %SYBASE%\DBISQL 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 %SYBASE%\Shared\JRE-6_0* directory. • The Administration Utility is installed to %SYBASE%\SDCADMIN-15_0 • <code>jutils-3_0</code> – 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 Shared directory. <p>The <code>ASEPlugin.jar</code> file is located in <code>ASEP\lib</code>, while the <code>scjview.exe</code> executable to start Sybase Central with the ASE plug-in is located in <code>shared\Sybase Central 6.0.0\[win32,win64]</code>.</p> <ul style="list-style-type: none"> • Adaptive Server plug-in – Sybase Central plug-in that provides cluster configuration and administration functionality up to Adaptive Server 15.5. Installed into ASEP directory. <hr/> <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> <hr/>
Language Modules	<p>Installed into the <code>locales</code> directory. Provides system messages and date/time formats.</p>

Product	Description
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 before beginning the installation.
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 need to change the permissions, 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.

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

Planning Your Adaptive Server Installation

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

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

5. Log in, 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:
 - Make sure there is sufficient disk space and temporary directory space available.
 - Ensure that 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, make sure to 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 ensure that clocks are synchronized.
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. Supported versions of clustered file system are detailed in the next section.

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.

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.

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:

```
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. When multipathing is configured and running, it 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 multipath 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. So multipathing prevents single points of failure in any highly available system.

Planning Your Adaptive Server Installation

A common example of multipath connectivity is a SAN connected storage device. Usually one or more Fibre Channel HBAs from the host will be connected to the fabric switch and the storage controllers will be connected to the same switch. A simple example of multipath connectivity could be: 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 and hence we have multipath connectivity.

All OS platforms provide their own solutions to support multipathing. In addition 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 multipathing on your system, set it up 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
-

Installing Adaptive Server

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 for use.

Prerequisites

Shut down all programs before running the installer.

Task

Note: Adaptive Server Enterprise 15.7 includes new major versions of Adaptive Server and many supporting components. Installing Adaptive Server version 15.7 into the same directory with existing products should not impact the existing products. However, installing other

Installing Adaptive Server

products on top of Adaptive Server version 15.7 may result in one or more products not working correctly.

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

If you already have a working Adaptive Server on your system, do not install the Cluster Edition in the same \$SYBASE directory. If you must install other products into the same directory, install the Cluster Edition last.

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! 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 be overwritten without 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 will likely 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 hot key 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 installation 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. Create a /cdrom directory.
3. Launch the installer:
/cdrom/setup.bin

4. Select the language.
5. On the Introduction screen, click **Next**.
6. 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 *Adaptive Server Upgrades* on page 91.

7. 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.

8. 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 will work for 30 days from the original installation date.
9. Choose the most appropriate region, read the license terms, then click **I agree**. Click **Next**.
10. If you are installing a licensed copy of Adaptive Server Enterprise Cluster Edition, select the license type:
 - **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
11. 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

Installing Adaptive Server

- Recipient e-mail addresses
 - Message severity that triggers e-mail messages
12. In the pre-installation summary, verify the installation type and that you have enough disk space for the installation. Click **Next**.
The installation status window displays the result of the installation process.
 13. Select **enable** or **disable** to remember the Adaptive Server password, then click **Next** to specify whether to have the ASE plug-in remember passwords.
 14. If you chose a typical installation or selected the Sybase Control Center remote command and control agent for Adaptive Server in a custom installation, the installer asks if you want to configure SCC. If you choose to configure SCC, you see a message about discovery services, and you can configure the UDP or JINI adapters.
If you select JINI adapter, enter the host name, port number, and heart beat period.
 15. Enable and order the security login module.
 16. Enter the RMI port.
 17. Choose whether you want to enable SCC shared disk mode. If you do, enter the SCC instance name.
 18. Review the summary of selections you made for SCC, then click Next. The installer configures the SCC agent.
 19. 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.

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 in .

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*.

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 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

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.

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 during installation appear on the Java console and the wizard is run in console mode.
<code>-i silent</code>	Installs or uninstalls the product in silent mode, where the installation or uninstillation is performed with no user interaction.
<code>-D</code>	Passes custom variables and properties. For example, to override the default installation directory when you run the installer, enter: <pre>install_launcher_name -DUSER_INSTALL_DIR=/sybase</pre>
<code>-r</code>	Generates a response file and a reference.
<code>-f</code>	References response files.

Option	Purpose
-l	Sets the installer's locales.
-h?	Displays the installer help.

Uninstalling Adaptive Server

Run the uninstaller to remove Adaptive Server.

Prerequisites

Shut down all servers before uninstalling Adaptive Server.

Task

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

1. Execute the uninstall 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 uninstall, you can still proceed with the uninstall 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:

Installing Adaptive Server

```
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 SDK. It installs 32-bit SDK on 32-bit OS; and 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:

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

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**.

Installing Components from the PC-Client

Creating and Starting the Cluster

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

1. Read *Before Creating a Cluster* on page 46.
2. Set the \$SYBASE environment and start the Sybase Control Center remote command and control agent on each node of the cluster.
3. Configure the cluster using:
 - The **sybcluster** utility
 - Sybase Control Center
 - The Adaptive Server plug-in

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.
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 55.

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 <code>interfaces</code> file, unless LDAP is used for server lookup • All configuration files, including the cluster input file • All server binaries (for example, <code>data-server</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 • 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 a 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?	
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 not entered correctly 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.)	

Configuration	Value
<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 syb-cluster supply default values if your system does not have an existing standard.</p>	
<p>What is the query or listening port number for each instance? (Choose unique port numbers that are not in use by other applications.)</p>	
<p>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.)</p>	

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, for 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 that you use a value of 3 (the default). You can set this using the **sp_configure** stored procedure.

Starting and Stopping Sybase Control Center Agent

Sybase Control Center 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. You can:

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

Creating and Starting the Cluster

For instructions, see **Get Started > Launching Sybase Control Center > Starting and Stopping Sybase Control Center in UNIX** 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 the following 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 SCC script, the SCC console prompt appears. At this prompt, enter the following:

```
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
```

Before Creating a Cluster

After you install Sybase Control Center, perform these tasks so you can create a cluster.

The **sybcluster** utility does not create a `run_server` file when they create a cluster. After you create a cluster, you must start this cluster and each instance using **sybcluster** or Sybase Control Center. You cannot start this 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 3. 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) 	
Sybase system procedure device – full path and name for sybssystemprocs device: <ul style="list-style-type: none"> • Size of sybssystemprocs device (152MB) • Size of sybssystemprocs 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 	

Parameter and description (default values)	Your Information
(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 4. Verifying Locations for Shared Installations

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

Table 5. Information for Shared Installations

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

Table 6. 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 7. 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 8. 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)	
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 44.

To view a complete **sybcluster** session that includes answers to all the prompts, see *Sample sybcluster Session* on page 87.

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 -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 47.

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.
Quorum device	The full path for the quorum device (such as /dev/raw/raw11).
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, /dev/raw/raw12. 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: /dev/raw/raw13. 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: /dev/raw/raw14. 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: /dev/raw/raw20. The PCI database device size (24MB) The PCI database size (24MB)

Fields	Description
Secondary network	<p>Does this cluster have a secondary network? (Y).</p> <ul style="list-style-type: none"> If you enter Y to a secondary network, 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 to a secondary network, 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.

- 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” environment shell script (such as /remote/var/sybase/SYBASE.sh).
 - 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.
- 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 -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 asecel15
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

Creating and Starting the Cluster

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 Adaptive Server version 15.7 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 -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 the Backup Server 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 a later stage.

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*".

Creating and Starting the Cluster

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 node 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
```

```
mycluster_JSAGENT
  master tcp /dev/tcp hostname1 17780
  query tcp /dev/tcp hostname1 17780
  master tcp /dev/tcp hostname2 16780
  query tcp /dev/tcp hostname2 16780
```

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* for more information on directory services.

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.

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`

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.

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

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," without a password. 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 Testing Environment

To create a testing 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 testing 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, and can be set up 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 for your platform for more information.

1. Configure the `libtcl.cfg` file to use directory services, using any standard ASCII text editor:
 - Remove the semicolon (;) comment markers from the beginning of the LDAP URL lines in the `libtcl.cfg` file under the [DIRECTORY] entry.
 - Add the LDAP URL under the [DIRECTORY] entry. See the *Configuration Guide* for supported LDAP URL values.

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

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

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

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

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

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

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

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

The Windows PATH environment variable must include this directory.

Adding a Server to the Directory Service

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

1. From Windows, select **Start > Programs > Sybase > Connectivity > Open Client Directory Service Editor**.
2. Select **LDAP** from the list of servers, and click **OK**.
3. Click **Add New Server Entry**, and enter:
 - Server name
 - Security mechanism – (Optional) there is a list of security mechanism OIDs in `%SYBASE%\ini\objectid.dat`.
4. Click **Add New Network Transport**, and:
 - Select the transport type.
 - Enter the host name.
 - Enter the port number.
5. Click **OK** twice to exit the **dsedit** utility.

Setting the System Administrator Password

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

Immediately after a new installation; the initial default value for the "sa" password is NULL, allowing anyone to log in as "sa." For security purposes, make sure you assign a password to "sa" before using Adaptive Server in a production environment.

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 null, new_password
2> go
```

where:

- **null** – is the default 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
install-pubs2	<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>
install-pubs3	<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>
install-pix2	<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,

Postinstallation Tasks

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

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

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.

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`, `Roman 9`, or `UTF-8` is installed as either the default character set or as an additional character set.

The `interpubs` database contains 8-bit characters and is available for use at Adaptive Server installations using the ISO 8859-1 (`iso_1`), ISO 8859-15 (`iso_15`), Roman 8, or Roman9 (for HP-UX) character set.

3. Determine the type (raw partition, logical volume, operating system file, and so on) and location of the device where you will be storing the `interpubs` database. You will need to provide this information later.
4. Execute the script, using the `-J` flag to ensure that the database is installed with the correct character set:

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

Installing the jpubs Database

If you installed the Japanese Language Module with your server, you can run the `installjpubs` script to install `jpubs`, which is a database similar to `pubs2` that contains Japanese data. `installjpubs` uses the EUC-JIS (`eucjis`), UTF-8 (`utf8`), or the Shift-JIS (`sjis`) character sets.

Prerequisites

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

Task

1. Set your terminal to display 8-bit characters.
2. Verify that either the EUC-JIS, Shift-JIS, or the UTF-8 character set is installed as the Adaptive Server default character set or as an additional character set.
3. Determine the type (raw partition, logical volume, operating system file, and so on) and location of the device where you will be storing the `jpubs` database. You will need to provide this information later.
4. Execute the `installjpubs` script, using the `-J` flag to ensure that the database is installed with the correct character set:

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

```
isql -Usa -Ppassword -Sservername -Jeucjis \  
-i %SYBASE%\%SYBASE_ASE%\scripts\eucjis\installjpubs
```

or:

Postinstallation Tasks

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

```
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:

Postinstallation Tasks

```
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 with I/O fencing capability.

1. Each instance in the cluster must run on a separate node.

- 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. Adaptive Server does not allow you to 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.
-

- 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
```

- 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.
-

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

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

For example: `mkvg -y myvolgroup hdisk5`

- Create a logical volume on one node:

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

Postinstallation Tasks

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.

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 setup one local system temporary database per device).

Note: Local system temporary database must use shared disks.

- Quorum disk device (at least 4 MB)
- 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 82 for information about changing configuration information after the cluster has been initialized.

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

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
```

Postinstallation Tasks

```
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* – the maximum number of instances in the cluster that your version of Adaptive Server supports. See the release bulletin for the maximum number of instances your version of Adaptive Server supports.
- **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 startup. 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:

$$start_port_number + (max_instances * 5) - 1$$

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]
```

Postinstallation Tasks

```
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 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 sets of steps:

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.

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
```

Postinstallation Tasks

```
--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 **InstallAnywhere** 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. The “\” indicates that the line continues without a space.

```
$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 to boot 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
```

Postinstallation Tasks

```
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 **automatic cluster takeover** to 1 enables the automatic cluster takeover (the default). Setting it to 0 disables **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.

Starting the Cluster

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

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, start the "ase1" instance:

```
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":

1. On "blade2," move to the \$SYBASE directory.
2. Issue:

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

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, this example adds the **-cluster-takeover** parameter to the runserver file that originally started the cluster:

```
$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.

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. The command prompt displays 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
```

Postinstallation Tasks

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 boot 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 the chapter titled *sybcluster Utility* in the *Clusters Users Guide* for complete syntax and usage information for **sybcluster** and the Sybase Control Center for Adaptive Server for information on .

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 the Creating the Cluster Server Using **sybcluster** topic.

Table 9. sybcluster Prompts and Configuration Values Entered for Sample Session

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

Postinstallation Tasks

Parameter	Value
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/
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

Parameter	Value
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
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

Postinstallation Tasks

Parameter	Value
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
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

Parameter	Value
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)	

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

```

Postinstallation Tasks

```
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 ]
Enter path to the dataserer configuration file: [ /remote/quasr5/
adong/aries/release/lamce_s1/mycluster.cfg ]
-----
You will now be asked for the instance information on a node by node
basis.
-- Cluster: mycluster - Node: hpcblad1.sybase.com - Agent:
hpcblad1.sybase.com:9009 --
Enter the name of the cluster instance: instancel
Enter the interface file query port number for instance instancel:
10665
Enter the primary protocol address for instancel:
[ hpcblad1.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:

Postinstallation Tasks

```
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
```

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 15.5
- 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 preupgrade tasks.

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

To upgrade from a shared-installation mode to a private-installation mode in Adaptive Server 15.5 Cluster Edition, see to the Release Bulletin for special upgrade instructions.

Adaptive Server version 15.7 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 10. 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, then 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> in the <i>Replication Server Configuration Guide</i> for your platform. Also, see <i>Upgrading Servers with Replicated Databases</i> in the <i>Adaptive Server Installation 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, then 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> in the <i>Replication Server Configuration Guide</i> for your platform. Also, see <i>Upgrading Servers with Replicated Databases</i> in the <i>Adaptive Server Installation 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 non-clustered version is not supported.
Adaptive Server 15.5.x	Adaptive Server 15.7.x	No special upgrade instructions.
Adaptive Server Cluster Edition 15.5.x	Adaptive Server Cluster Edition 15.7.x	In the installer destination directory panel, enter the 15.5.X directory, then choose Update . After your installation, install the new system procedures (from \$SYBASE/ASE-15_0/scripts/installmaster) and follow instruction in <i>Postupgrade Tasks</i> on page 116.

Table 11. 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, and if the database you are upgrading is also replicated, make sure the database's transaction log is truncated before you turn replication back on. <i>See Reenabling Replication Server After the Dataserver Upgrade in the Adaptive Server Installation Guide for your platform.</i>
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 non-clustered 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.

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, 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 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.

Changes in Adaptive Server Directories

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

Table 12. Directory Changes for UNIX Platforms

Component	12.5.4 Location	15.0.2 Location	15.0.3 Location	15.5 and 15.7 Locations
Adaptive Server	<code>\$\$SYBASE/ASE-12_5</code>	<code>\$\$SYBASE/ASE-15_0</code>	<code>\$\$SYBASE/ASE-15_0</code>	<code>\$\$SYBASE/ASE-15_0</code>
Shared directory	<code>\$\$SYBASE/shared</code>	<code>\$\$SYBASE/shared</code>	<code>\$\$SYBASE/shared</code>	<code>\$\$SYBASE/shared</code>
Sybase Central	<code>\$\$SYBASE/shared/syb-central43</code>	<code>\$\$SYBASE/shared/syb-central43</code>	<code>\$\$SYBASE/shared/syb-central600</code>	<code>\$\$SYBASE/shared/syb-central600</code>
JRE	<code>\$\$SYBASE/shared/jre142</code>	<code>\$\$SYBASE/shared/jre142_*</code>	<code>\$\$SYBASE/shared/JRE-6_0*</code>	<code>\$\$SYBASE/shared/JRE-6_0*</code>
Shared JAR file	<code>\$\$SYBASE/shared/lib</code>	<code>\$\$SYBASE/shared/lib</code>	<code>\$\$SYBASE/shared/lib</code>	<code>\$\$SYBASE/shared/lib</code>
Language settings (locales)	<code>\$\$SYBASE/locales</code>	<code>\$\$SYBASE/locales</code>	<code>\$\$SYBASE/locales and \$\$SYBASE/ASE-15_0/locales</code>	<code>\$\$SYBASE/locales and \$\$SYBASE/ASE-15_0/locales</code>
Connectivity	<code>\$\$SYBASE/OCS-12_5</code>	<code>\$\$SYBASE/OCS-15_0</code>	<code>\$\$SYBASE/OCS-15_0</code>	<code>\$\$SYBASE/OCS-15_0</code>
Web Service	<code>\$\$SYBASE/WS-12_5</code>	<code>\$\$SYBASE/WS-15_0</code>	<code>\$\$SYBASE/WS-15_0</code>	<code>\$\$SYBASE/WS-15_0</code>
Replicator	<code>\$\$SYBASE/RPL-12_5</code>	<code>\$\$SYBASE/RPL-15_0</code>	<code>\$\$SYBASE/RPL-15_0</code>	
SySAM	<code>\$\$SYBASE/SY-SAM-1_0</code>	<code>\$\$SYBASE/SY-SAM-2_0</code>	<code>\$\$SYBASE/SY-SAM-2_0</code>	<code>\$\$SYBASE/SY-SAM-2_0</code>

Component	12.5.4 Location	15.0.2 Location	15.0.3 Location	15.5 and 15.7 Locations
Job Scheduler	\$SYBASE/ JS-12_5	\$SYBASE/ ASE-15_0/ jobschedu- ler	\$SYBASE/ ASE-15_0/ jobschedu- ler	\$SYBASE/ ASE-15_0/ jobschedu- ler
Unified Agent		\$SYBASE/ UAF-2_0	\$SYBASE/ UAF-2_0	\$SYBASE/ UAF-2_5 In Adaptive Server Cluster Edition 15.7, 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 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

If you are upgrading Adaptive Server:

- 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. This reduces 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.
- To a private installation, see *Upgrading to a Private Installation* on page 101.

Task

1. Shut down the clusters that use an earlier version of the Cluster Edition.
2. Install the Cluster Edition onto your system in a new location.
3. Check system and upgrade requirements.
4. Verify the name and location of the `runserver` file, as well as that you renamed it to `RUN_servername`, where `servername` is the name of the old server, as it appears in the `interfaces` file.

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

5. 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.
6. 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.

7. Make sure reserved words use quoted identifiers.
8. Verify that users are logged off.
9. Use `dbcc` to check for database integrity.
10. Back up databases.
11. Dump the transaction log.
12. Verify that `master` is the default database for the "sa" user.
13. Prepare the database and devices for upgrade using the **preupgrade** utility:
 - a) Create a `sybsemdb` 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 120.
 - e) Disable disk mirroring.

Note: Cluster Edition version 15.7 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.

- 14.** 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 upgrade.
- 15.** If you are upgrading from a 12.5.3 installation with cache bindings on `sybsemdb`, drop the cache bindings for `sybsemdb` 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 do receive this error, re-run **preupgrade** after you have dropped the cache bindings.

- 16.** 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.
- 17.** 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`
- 18.** If you have enabled the Java in the database feature, create the `sybpcidb` database and disable the feature during the installation.
- 19.** 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`.
- 20.** 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`

Upgrade System Table and Stored Procedure

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 sybsystemprocs  
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, and 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.

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, 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/rawlg2
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_instance1 using the operating system command:
/mysybase1/ASE-15_0/bin/dataserver --quorum_dev= /dev/raw/
raw50m41 --instance_name= mycluster_instance1
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_instance1--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*.

The sybprocsdev device

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

For an upgrade, the recommended minimum size is the larger of 140MB, or enough free space to accommodate the existing `sybsystemprocs` database, and the largest catalog that is to be upgraded, plus an additional 10 percent of the largest catalog's size. The additional 10 percent is for logging upgrade changes.

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

If your `sybsystemprocs` database does not meet these requirements and you have enough room on the device to expand the database to the required size, use the **alter database** command to increase the database size.

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

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

Use **sp_helpdevice** to determine the size of the `sysprocsdev` device:

```
1> sp_helpdevice sysprocdev
2> go
```

If the `db_size` setting is less than the required minimum, increase the size of `sysprocdev`.

Increasing the Size of the sybsystemprocs Database

Create a new database with sufficient space if your current `sybsystemprocs` 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 `sybsystemprocs` onto the new device.

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

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

In this example, "sysprocsdev" is the logical name of the existing system procedures device, and 40 is the number of megabytes of space to add. If the system procedures device is too small, you may receive a message when you try to increase the size of the `sybsystemprocs` database.

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

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

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

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

Increasing Device and Database Capacity for System Procedures

If you cannot fit the enlarged `sybsystemprocs` database on the system procedures device, increase the size of the device and create a new database.

Prerequisites

This procedure removes all stored procedures you have created at your site. Before you begin, save your local stored procedures using the `defncopy` utility. See the *Utility Guide*.

Task

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

1. Determine which device you must remove:

```
select d.name, d.phyname
from sysdevices d, sysusages u
where u.vstart between d.low and d.high
and u.dbid = db_id("sybsystemprocs")
and d.status & 2 = 2
and not exists (select vstart
                from sysusages u2
                where u2.dbid != u.dbid
                and u2.vstart between d.low and d.high)
```

where:

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

The **not exists** clause in this query excludes devices that are used by `sybsystemprocs` and other databases.

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

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

2. Drop `sybsystemprocs`:

```
1> use master
2> go
1> drop database 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 sybssystemprocs, where /sybase/work/ is the full, absolute path to your system procedures device:**

```

1> use master
2> go
1> disk init
2> name = "sysprocsdev",
3> physname = "/sybase/work/sysproc.dat",
4> size = 50M
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 sybssystemprocs database of the appropriate size on that device, for example:

```

1> create database sybssystemprocs on sysprocsdev = 112
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

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.

From an earlier version of Adaptive Server Cluster Edition, to Adaptive Server Cluster Edition version 15.7, 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. Perform preupgrade tasks.
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 -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. Shut down the old Adaptive Server. Alternatively, you can bring down the old Adaptive Server cluster using **sybcluster**, such as:

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

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

7. (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
```

Adaptive Server Upgrades

```
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 73.

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

8. 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.

9. Run the **upgrade** utility, where **instance_name** is the first instance in your cluster that has the same name as the server from which you are upgrading:

```
$SYBASE/$SYBASE_ASE/upgrade/upgrade
-S instance_name -Ppassword
```

10. (Skip this step if you are upgrading from a 15.0.1, 15.0.3, or 15.5 Cluster Edition to a 15.7 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
```

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

```
shutdown instance_name
```

12. Restart the cluster.

```

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

```

13. 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 sybserverprocs
2> go
1> sp_checkreswords
2> go

```

14. Correct any errors from the reserved word check.

15. 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`

16. Start each instance in the cluster:

```

cd $SYBASE/$SYBASE_ASE/install
startserver -fRUN_server_name

```

17. Install the system procedures:

```

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

```

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

```

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

```

19. Run `installcommit`:

```

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

```

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. If you already have a clustered Adaptive Server, upgrade manually.

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

Adaptive Server Upgrades

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 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 a list of the information you need for the upgrade.

Complete these tasks before you upgrade:

1. Make sure that your server is ready for upgrade.
2. Upgrade the server to Adaptive Server Cluster Edition version 15.7.
3. Perform the postupgrade tasks after the upgrade, whether you upgrade the server manually or use **sybcluster**.

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 -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. The default password for "uafadmin" is null or blank.
- **-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.

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**. Enter:

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

2. Start the upgrade. Enter:

```
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.

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 are displayed 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`.

Adaptive Server Upgrades

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.
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 over an Existing 15.x Adaptive Server

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

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

Determining the Adaptive Server Version

Verify that your current version of Adaptive Server is not earlier than 15.x before using the binary overlay.

If your server is at a 15.x version level, you can begin installing Adaptive Server 15.7.

1. Verify that your version of Adaptive Server you have is version 15.x:

- If your server is running:

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

- If your server is not running:

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

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

If the result shows that the Adaptive Server is lower than version 15.x, you cannot use the binary overlay, but use an upgrade method instead. See *Adaptive Server Upgrades* on page 91.

Backing Up Adaptive Server

Installing Adaptive Server version 15.7 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.

Adaptive Server Upgrades

2. After you have verified that your database is error free, back up your \$SYBASE directory in case you need to roll back to the original version of the software.
3. Adaptive Server provides the **uninstmsgs.ebf** script for backing up `sysmessages` before upgrading to 15.7. 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 on top of an Adaptive Server that is no earlier than version 15.x.

1. From your SYBASE directory, shut down the servers that are using the binaries.
2. Use the installer to load the new software from the CD or DVD.
Navigate to the CD or DVD drive and enter `./setup.bin`.
3. Install Adaptive Server 15.7 into the existing \$SYBASE installation path.

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

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

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. After upgrading from Adaptive Server Cluster Edition version 15.0.1 to 15.5 Cluster Edition, you must a one-time procedure for each of your user-created nodes to update the JAR and XML files.

2. If you upgraded from Adaptive Server version 12.5.2 or earlier, run **dbcc checkcatalog** with the **fix** option to ensure there are no issues with OAM pages:

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

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

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

The output of the script is saved in `output_file`. See *How Compiled Objects are Handled When Upgrading Adaptive Server* on page 131 if you are using **dbcc upgrade_object** to upgrade compiled objects such as check constraints, rules, triggers, and views.

4. Run the scripts for different options or licensed features:

- **installcommit** – if you use two-phase commit or distributed transactions, re-run **installcommit** to restore:
 - **sp_start_xact**
 - **sp_abort_xact**
 - **sp_remove_xact**
 - **sp_stat_xact**
 - **sp_scan_xact**
 - **sp_probe_xact**
- **installsecurity** – if you used auditing in your earlier installation.
- **installhasvss** – if you enabled high availability, and are using this feature in this installation.
- **installmsgsvss** – if you enabled real-time messaging, and you are using this feature in this installation.
- **installpcidb** – if you enabled the Java in the database feature in your earlier installation.
- **installjsdb** – if you enabled Job Scheduler in your earlier installation.

5. If you have any system stored procedures that you had saved from before the upgrade (because you had modified them without altering their names), reload them now.

6. Restore functionality in Adaptive Server.

7. Reenable auditing.

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.

Adaptive Server Upgrades

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,lstdb1,/work/river/ase1501ce/data/
lstdb1.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.

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 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 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.
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.

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 131.
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 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, 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_zerolrm 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, 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*.

Restoring Replication on Primary Databases

Restore replication if you have upgraded a source database, or primary database, in a Replication Server system, or 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_zerolrm dataserver, database
2> go
```

2. Log in to each replicated primary, and replicated RSSD and issue:

```
1> use database
2> go

1> dbcc settrunc ('lrm', 'valid')
2> go
```

Adaptive Server Upgrades

- Restart the Replication Server.
- 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
```

- 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.

- 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
```

- If you are using the Log Transaction Manager, restart it.

Migrate

You can migrate from a nonclustered to the Cluster Edition of Adaptive Server.

To migrate, you can:

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

Migrating Data Using a Dump-and-Load Method

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

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

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

- Dump the databases from the original server.
- Load the databases to the Cluster Edition server.
- If you have partitioned tables, update partition statistics.
- 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 132.

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 the following sections to upgrade Adaptive Server to the Cluster Edition.

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.

Adaptive Server Upgrades

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 the number of locks:

```
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.

- 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.

- Copy files from the `jobscheduler` directory to the new 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.
- 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.

- 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 13. Changed Job Scheduler Templates

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

Template	Modified Files	Datatype for <parameter> Changed from varchar(5) to int
rebuild table	SybRebuildTableTemplate.xml jst_reorg_rebuild_tables	@dump_flag
reclaim indexes	SybReclaimIndexTemplate.xml jst_reclaim_index_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 servers.

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

You need not modify jobs you run on pre-15.7 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 14. Modified Job Scheduler Templates

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

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

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

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

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

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

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

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

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

- a) Run the XML installation script on 15.0.1 servers with Job Scheduler installed:

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

Use "en" for the *language_code* or omit the parameter completely, as "en" is the default.

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

Upgrading High Availability and Cluster Support

Upgrade the cluster subsystem.

Adaptive Server 15.7 supports these cluster platforms for high availability:

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

There are two ways to upgrade a cluster subsystem:

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

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

Upgrading Java in the Database

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, tmp on Windows.

1. Run the **installpci** script only on 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
```

Adaptive Server Upgrades

```
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 for 15.7*.

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**, and reestablish companionship afterwards.

Java in the database feature 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** from the drop-down menu 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.

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 pre-upgrade 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 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 pre-upgrade 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.

Adaptive Server Downgrades

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

Even if you have not used any of the new features in Adaptive Server 15.7, the upgraded server adds 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 to an earlier version.

Preparing to Downgrade Adaptive Server

Prepare your system before starting the downgrade.

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

- 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, 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.

Adaptive Server Downgrades

- 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, **sp_downgrade** displays an error 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

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

Prerequisites

1. Save the current auditing settings for the 15.7 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 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** 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 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 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.

sp_downgrade 'prepare' validates the readiness of Adaptive Server 15.7 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.

5. Make sure Adaptive Server 15.7 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 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 server. Run **checkpoint** and immediately shut down the 15.7 server by issuing the **shutdown** command to shut down the cluster.

Note: Since Adaptive Server 15.7 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

Adaptive Server Downgrades

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

Note: 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 functionality, this completes the primary downgrade process.

Additional Considerations for New Features Used

Additional downgrade steps may be necessary due to Adaptive Server 15.7 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. • sybpcidb 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 non-characters	<p>If you used <code>sp_configure</code> to enable functionality group or permissive Unicode in Adaptive Server 15.7, 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 truncation</code> configuration in Adaptive Server 15.7, the <code>varbinary</code> data stored in the system may have trailing zeros. When you downgrade to a version earlier than 15.7, these trailing zeros are not removed and become part of the <code>select</code> 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_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
```

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 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**, **installhasvss**, and **installmsgsvss** on your downgraded server, any changed

system stored procedures are returned to their earlier form. Any new stored procedures introduced to Adaptive Server 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 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, 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 configuration file, the new options generate an `Unknown parameter` message. The unknown options are reported the first time you restart the server. You can ignore these messages; the configuration file is rewritten without the unknown options.

Note: Keywords **decrypt_default**, **xmltable**, and **path** were added in Adaptive Server 15.5 Cluster Edition, making it impossible to create identifiers using these names. You must change applications if you used these names.

Troubleshoot SySAM

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

There may be more recent information published in the *SySAM Asset Management Users Guide* online at the *SySAM FAQ*. If your Sybase product cannot find a required license, it continues to operate in grace mode, if a grace period is available. Determine the license status of your product using the output from **sp_lmconfig** command. The `Status` column shows graced for any license that is operating within a grace period.

Typically, you have 30 days during which to resolve any SySAM-related issues. When a license is awarded in grace mode, the Sybase product error log entry and optional e-mail message include the date and time when the grace period expires. In addition, you can execute **sp_lmconfig** to determine the expiration date.

All SySAM-related errors are printed in your Sybase product error log and are prefixed with "kernel SySAM. " For example:

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

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

You can always check the current state of SySAM licenses in your Sybase product using the **sp_lmconfig** command. The output shows the current SySAM configuration and the status of each license.

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

The Sybase product log includes information about licenses awarded under grace mode. If e-mail notifications are set up for a severity of "warning," e-mail messages are sent for this event. Warnings in the Adaptive Server error log and e-mail messages are repeated while the license is in grace mode.

Troubleshooting SySAM

Identify and troubleshoot SySAM issues.

Description	Action
<p>Cannot find license type during the Adaptive Server installation.</p>	<p>The Adaptive Server installation program displays only the most commonly used license types. If you are using a legacy license or a new license type, your license may not be listed.</p> <p>If you do not find the license type you need, select Unknown and complete the installation. After the installation is complete, use sp_lmconfig to change the license type. See the <i>Reference Manual:Procedures</i> for sp_lmconfig usage information.</p> <hr/> <p>Note: If you set the license type value to Unknown, Adaptive Server uses the first available license. If you have licenses for more than one type, Adaptive Server may pick the incorrect license. Set the correct license type.</p> <hr/> <p>See <i>If the Installation Program Cannot Find a Suitable License</i> on page 147 if you select a served license model that the program cannot find.</p>
<p>The license server executables and scripts are not installed.</p>	<p>The <code>sysam</code> wrapper script and the executables for the license server (Imgrd and SYBASE) are not installed or found in the <code>SYSAM-2_0 bin</code> directory.</p> <p>By default, the license server component is not selected for installation. To install the license server, choose the custom installation option and select the license server component.</p> <p>Alternatively, download a standalone SySAM license server installation program from http://www.sybase.com/products/allproductsa-z/sysam.</p>

Description	Action
<p>The license server does not start.</p>	<p>For a license server to start, there must be at least one served license in the <code>licenses</code> directory. When you install a new license server, there are no served licenses in the <code>licenses</code> directory and the license server does not start until you activate a license from the Sybase Product Download Center (SPDC) or SAP Service Marketplace (SMP) and copy it into the <code>licenses</code> directory.</p> <p>Other causes for license server not to start include:</p> <ul style="list-style-type: none"> • The license activated from SPDC or SMP was an unserved license. Examine the license files. Served licenses always start with a <code>SERVER</code> header. If you do not see this header, you are using an unserved license, which does not work with the license server. • If you are using a specific port number for the license, that port number may already be in use. Use <code>netstat -a</code> to verify that the port number is free. If not, free the port or use a different port number for the license server. • The host name recorded in the <code>SERVER</code> header does not match the actual host name. The host name is recorded next to the <code>SERVER</code> keyword. If it does not match, correct the host name in the license file or set it to <code>"this_host"</code>, which is a keyword that works with any host name. • The licenses copied may be activated for a different machine. The host ID next to the host name in the <code>SERVER</code> header must match the host ID for the machine where the license server is running.
<p>The license server does not recognize the license files.</p>	<p>If the license server refuses to serve the licenses, it may be because:</p> <ul style="list-style-type: none"> • The license was activated for a different machine or activated using an incorrect host ID. Compare the host ID recorded in the license with the host ID assigned to the machine. If they do not match, check in and regenerate the license with the correct host ID. • The license has been altered. If the license is being created by typing from a paper copy, verify if there are any errors entering the license. You can also download a new copy of the activated license from SPDC or SMP. <p>If the host ID for your platform is based on the network adapter identification, verify that you have used an ID associated with a valid NIC. IDs associated with loopback adapters are invalid. If the ID associated with a removable network adapter has been used, ensure that the adapter is attached to the computer.</p>
<p>Adaptive Server does not start, and generates a license checkout failure error.</p>	<p>If Adaptive Server cannot check out the required license, it determines whether the license can be issued in grace mode. If a grace period cannot be given, the license is not granted. If the base license for Adaptive Server (<code>ASE_CORE</code>) has been running in a grace period, and the issue is not fixed by the end of the grace period, Adaptive Server fails to start.</p> <p>See "Adaptive Server cannot check out a license and starts with graced license instead."</p>

Description	Action
<p>Adaptive Server shows graced licenses even after the issue has been fixed.</p>	<p>Adaptive Server periodically performs license checks, and the license status is not immediately updated. It is updated only after the next heartbeat cycle is completed, which may take a few hours.</p>
<p>Adaptive Server cannot find license for an optional feature, even though the license exists.</p>	<p>Execute <code>sysam diag feature_name</code> to ensure that the license for the optional feature exists and can be checked out from the machine where Adaptive Server is running. If the feature exists, but cannot be checked out from Adaptive Server, this could be because:</p> <ul style="list-style-type: none"> • The optional feature is for a different edition. • The status field in the license, which indicates whether a license is active or not, does not match between Adaptive Server and the optional feature. <p>Execute <code>sp_lmconfig</code> to determine the edition, and active and standby (that is, inactive) configuration values. The information is located in the Property Name/Property Value table. The edition value is in the PE row and active and standby value is in the AS row.</p> <p>The values of PE and AS for the Adaptive Server must match the values of the <code>VENDOR_STRING</code> and <code>ISSUER</code> fields for the optional feature license. Some features, such as high availability (<code>ASE_HA</code>), are not offered with standby pricing, so there is no "<code>AS=A S</code>" flag in the license for this feature. This field is not considered when searching for a license without active and standby capability.</p>
<p>Adaptive Server does not start with the expected edition or license type.</p>	<p>If the edition and license type configuration parameters are not set, Adaptive Server uses the first available <code>ASE_CORE</code> license. If you have multiple <code>ASE_CORE</code> licenses with different editions and license types, the license that is the first available depends on many factors, such as license file directory sort order, available licenses on the license server, and so on.</p> <p>Sybase recommends that you set the edition and license type configuration parameters. When these parameters are set, you are assured that Adaptive Server starts with that configuration. If a matching license is unavailable, Adaptive Server starts under grace (if available) to allow you to resolve the licensing issue.</p>

Description	Action
<p>Adaptive Server picks an incorrect license.</p>	<p>Adaptive Server looks for licenses in these locations:</p> <ul style="list-style-type: none"> • Any saved values in the <code>registry</code> on Windows platforms, or the <code>.flexlmrc</code> file on UNIX platforms. Use <code>lmpath -status</code> to view the saved values. • The values set for <code>SYBASE_LICENSE_FILE</code> and <code>LM_LICENSE_FILE</code> environment variables. • All license files (any file with a <code>.lic</code> extension) in the licenses directory under the <code>SYSAM-2_0</code> directory. <p>All of the above locations are searched in the order specified until a suitable license is found. When a directory is specified, license files in that directory are loaded in the directory sort order.</p> <p>The first license that matches the feature name, version, edition, and license type filter is used. This license may not be what you were expecting to use. Adjust the above locations to control this.</p>

If the Installation Program Cannot Find a Suitable License

If you select a served license model, the installer checks for a license with the selected edition and license type. You see a warning message if no suitable license is available.

If you are performing a new Adaptive Server installation or upgrading an existing Adaptive Server version 12.5.x or earlier, proceed with the installation. Adaptive Server provides a 30-day grace period to resolve licensing issues. Once the installation is complete, debug the license checkout failure. See "Adaptive Server cannot check out a license and starts with graced license instead" in *Troubleshooting SySAM* on page 144.

If you are upgrading an existing Adaptive Server version 15.0 or later, Adaptive Server may not start after the upgrade if you do not have a suitable license. Exit the installation program and see if the existing Adaptive Server can check out licenses. If it cannot, troubleshoot this issue and then proceed with the upgrade.

If the existing Adaptive Server can check out licenses, verify that the license you are using authorizes the update you are applying. This can be determined using the date-based version of your license file and the date this update was released. You must have a license with a date-based version that is later than or equal to the date the update was released. Adaptive Server provides a grace period for updating the licenses after support is renewed.

Determine the date-based version of the license from the *Version* column of the output from `sp_lmconfig`, by looking at the license checkout message in the Adaptive Server error log, or by examining the license itself. See *Anatomy of a License* in the *Sybase Software Asset Management Users Guide*.

Determine the Adaptive Server release date by looking at the version string, cover letter, or download link for the Sybase Product Download Center (SPDC) or SAP Service Marketplace (SMP). The release date is also posted in the error message.

You can download an updated license from SPDC or SMP, if your support has been paid for the date the update was released. If your support has not been paid, you are not authorized to install an update.

See *Working with Date-Based Versioning* and *Getting and Using your Licenses* in the *Sybase Software Asset Management Users Guide*.

If Adaptive Server Cannot Check Out a License and Starts Using a License Under a Grace Period

There are several reasons why Adaptive Server cannot check out a license.

Prerequisites

Execute the `sysam diag feature_name` command from the `SYSAM-2_0 bin` directory, where `feature_name` is the SySAM feature name for Adaptive Server or the optional feature that is in grace mode. The feature name is printed in the Adaptive Server error log and in the optional e-mail notifications.

- If you are using a **served** model and **diag** shows no license available to check out:
 - Verify that the license server is running and can be reached from the machine you are using. Use `"sysam status"` to verify that you can access the license server. If you cannot access the license server, verify that the license server is running.
 - If the license server is running, use `"sysam status -f feature_name"` to determine whether the license server is serving a license for the given feature. If not, obtain the correct licenses from the Sybase Product Download Center (SPDC) or SAP Service Marketplace (SMP).
 - All licenses on the license server may be in use. If `"sysam status -f feature_name"` indicates that no licenses are available, obtain additional licenses or shut down existing instances of Adaptive Server.
- If you are using an **unserved** model and **diag** shows no license available to check out:
 - The license for the given feature may not be copied in the local `licenses` directory. Obtain the correct licenses from SPDC or SMP, and copy them to the local `licenses` directory.
 - The license may have been activated for a different machine or activated with incorrect host ID. Check in and reactivate the license with the correct host ID.
 - The available licenses are for a different operating system or architecture. Obtain license for the correct platform from SPDC or SMP.

- If you are running in a terminal server environment, you cannot use unserved licenses. Set up served licenses for a terminal server environment.
- The available license may not authorize this version of Adaptive Server. For information about date-based versioning, see the *SySAM Users Guide*.
- If the **diag** command shows the license for the given feature is available for checkout, this may be because the product edition or license type does not match the requirement. If the **edition** and **license type** configuration parameters are set, Adaptive Server uses only the license that matches these settings:
 - a) Run **sp_lmconfig** to see the configured values.
 - b) Check the **VENDOR_STRING** attribute of the licenses to verify that a matching license exists.

If a matching license does not exist, obtain a suitable license from SPDC or SMP, or adjust the Adaptive Server configuration.

If Adaptive Server cannot start because the grace period expired, examine the `servername.properties` file located in the `ASE-15_0/sysam` directory under your SYBASE directory for values. These values are stored in two lines that start with `PE=` and `LT=`. These lines are present only if the configuration is set.

E-Mail Notifications

Use **sp_lmconfig** to configure e-mail notification. To view the current notification status and configuration, execute **sp_lmconfig** without any parameters.

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

To change the severity that triggers e-mail notification, set the **email severity** property to: "ERROR", "WARNING", "INFORMATIONAL", or "NONE".

To change the list of e-mail recipients, set the **email recipients** property to a comma-separated list of e-mail addresses. E-mail addresses can include anything that is capable of receiving SMTP messages.

Before You Contact Sybase Technical Support

Get information about your situation before you call Sybase Technical Support about SySAM-related issues.

- The Adaptive Server error log.
 - `$(SYBASE)/ASE-15_0/install/<servername>.log`
- If Adaptive Server does not start, the server's properties file:
 - `$(SYBASE)/ASE-15_0/sysam/<servername>.properties`

Troubleshoot SySAM

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

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

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 .

Error Logs for Installation Utilities

The information in the error logs may help determine the reason and possible solution for a problem with installation problems.

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="413 512 1180 763"> % 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="413 920 1189 1050" style="list-style-type: none"> • To see whether the CD drive path is the current directory (pwd) in a UNIX terminal window. If it is, change (cd) to another directory. • For sybhelp processes. If these processes exist, kill them using the UNIX kill command.
<p>DISPLAY environment variable not set correctly.</p>	<p>To correct the problem, enter the following command at the UNIX prompt of the <i>remote</i> machine, where <i>host_name</i> is the name of the machine on which you want the installer to appear (that is, on your <i>local</i> machine).</p> <p>For C shell:</p> <pre data-bbox="413 1232 1180 1258">setenv DISPLAY host_name:0.0</pre> <p>For Bourne shell:</p> <pre data-bbox="413 1319 1180 1345">DISPLAY=host_name:0.0; export DISPLAY</pre>

Problem	Resolution
<p>Client not authorized to connect to server.</p>	<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.
<p>Address already in use.</p>	<p>Enter a different port number on the srvbuild window. The command netstat -a produces a list of port numbers in use.</p>
<p>Adaptive Server fails to start</p>	<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>
<p>The Installer fails to start</p>	<p>Restart the installer in debug mode. Set the environment variable LAX_DEBUG to true before running the installer.</p>
<p>Cannot start XP Server.</p>	<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>
<p>Troubleshooting resource file installations.</p>	<p>If you encounter problems during the build, configuration, or upgrade process while using the srvbuild[res], sqlloc[res], or sqlupgrade[res], it may be that these utilities did not allow enough time for Adaptive Server to shut down properly. Set the SYBSHUTWAIT environment variable to force the utilities to wait for Adaptive Server to shut down. For example, this command forces the utility to wait for two minutes to allow Adaptive Server to shut down before proceeding with the next task:</p> <pre>% setenv SYBSHUTWAIT 120</pre>

Stopping Adaptive Server After a Failure

If the installation or upgrade session fails after you start Adaptive Server for any reason, use the **shutdown** command.

1. Log on as "sa".
2. Shut down Adaptive Server using the **shutdown with nowait** command. Using **with nowait** 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

If the installation fails, you see error messages. Review the error messages and your Adaptive Server error log to determine the cause of the 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.
3. 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.
4. If the installation fails after the installation program starts the Adaptive Server that you are attempting to install, shut down that server.

If Adaptive Server Fails the Preupgrade Eligibility Test

Examine the log file to determine why Adaptive Server was not eligible 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 normally located in `/usr/lib`.

If the shared libraries on your operating system are located in other directors, 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/errorlog`.

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

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

```
online database <failed_db_name>
```

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

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.

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**.
 - 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 22
- Adaptive Server
 - installing in command line mode 36
 - installing in silent mode, unattended 36
- Adaptive Server components 123
- 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 138
- adjusting
 - shared memory parameters in operating system 26
- administration tasks 22
- applications
 - changing after renaming objects 99
- auxiliary servers
 - configuring 55

B

- Backing up the server 115
- Backup Server
 - configuring 56
 - configuring with sybcluster 55
- basic downgrade steps using sp_downgrade 136

C

- checklist before installing Cluster Edition 44
- client applications 8
- cluster
 - preinstallation checklist 44
- command line installation
 - Replication Server 36
- command line options 36
- commands
 - showserver 61
- Component Integration Services
 - upgrading local and remote servers 94
- configuration
 - manual 76

- configuration parameters 140
- configure
 - XP server using sybcluster 59
- configuring
 - Backup Server 56
 - master device 77
 - quorum device 77
- configuring the cluster 50
- conflicting object names 100
- creating
 - sybase user accounts 22
- creating a cluster with sybcluster 50
- creating the cluster
 - starting the cluster 43

D

- data compression
 - described 7
- databases
 - image data for sample 66
 - optional 65
- determining installation process 3
- determining the current edition 6
- determining whether select* should be changed in
 - views 133
- directory changes 95
- directory locations 20
- directory services
 - described 7
- downgrading
 - Adaptive Server 135, 136
 - handling new features 138
 - steps to prepare beforehand 135
- downgrading job scheduler 140
- dsedit utility 121

E

- Enabling the Java in the database in a high
 - availability system 130
- encrypted columns
 - described 7
- error log files
 - troubleshooting and 155

Index

F

- failed installation
 - cleaning up 55
- failed upgrade
 - recovery 158
- fails preupgrade eligibility 155
- file descriptors 23
- finding compiled object errors before production 132

G

- GUI installation 31

I

- I/O fencing
 - SCSI-3 PGR 70
 - setting up 69
- I/O Fencing
 - data integrity 68
 - enabling 70
 - setting up 69
- identifiers
 - conflicting with reserved words 99
- in-memory database
 - described 7
- increasing
 - device and database capacity 106
- infiniband 17
- input file
 - create 73
- input/output completion port 16
- installation
 - Adaptive Server CD, mounting 31
 - determining process via workflow 3
 - overview 3
- installation steps for PC-Client 39
- installer suddenly quits 155
- installing
 - command line 34
 - with response file 35
- installing Adaptive Server
 - GUI mode 31
- installpix script 66
- installpubs2 script 66
- installpubs3 script 66

- interfaces file 121
- IOCP status 16

J

- Java in the database
 - preparing before installing Adaptive Server 26
 - upgrading 129
- Job Scheduler
 - installing 57

K

- keywords 140

L

- LDAP libraries
 - environment variables 63
 - location of 63
- license model 11
- licenses
 - check number of processors 14
- licensing software
 - obtaining 11

M

- manual configuration 76
 - enabling SCC 82
 - enabling sybcluster 82
- methods of installation
 - Adaptive Server 31
- migrating
 - dump and load method 122
 - high availability 123
 - methods 122
- Monitor Server
 - configuring with sybcluster 55
- mount commands by platform 31
- mounting the Replication Server installation media 31
- multipathing
 - described 27
 - setting up before installation 29
- multiple simultaneous asynchronous I/O operations
 - requirements for processing 16

N

- network protocol 17
- new features
 - reverting to earlier version 138

O

- objects
 - conflicting names of 100
 - renaming 100
- obtaining
 - SySAM license 11
- optional databases 65
- optional features
 - described 7
- options for server editions
 - packages 7
- overview
 - installation 3
- overview of the upgrade process 91

P

- partitions
 - described 7
- PC-client
 - CD 8
- PC-Client
 - directory layout 20
 - product description 20
 - system requirements for 18
- planning the installation 11
- platform
 - mount commands 31
- post downgrade 140
- preinstallation tasks 23
- preupgrade tasls
 - before upgrade 96

Q

- quoted identifier errors 132
- quoted identifiers 100

R

- raw disks
 - I/O fencing 69

- multi-node access 69
- reconfiguring the cluster 82
- recovering from a failed installation 155
- reenable Replication Server 120
- related products 123
- release bulletin 11
- renaming objects 100
- replicated databases
 - restoring 121
- Replication Server 120
 - draining logs for Adaptive Server 121
- reserved words 99
 - conflicts 100
- response file 35, 40
- restoring from backup 157
- runserver file 99

S

- sample databases 65
- SCC
 - starting 45
- scripts
 - installpix 66
 - installpubs2 66
 - installpubs3 66
 - location 65
- SCSI-3 PGR 70
- security services
 - described 7
- server
 - advantages to the cluster edition 5
 - component descriptions 18
 - directory layout 18
 - installation options 5
 - overview 5
 - single-system presentation 5
 - workload management 5
- setting up
 - cluster manually 72
- shared disk cluster
 - private interconnect 17
- shared installation 43
- shared memory parameters
 - adjusting 26
- showserver command 61
- silent installation 40
- simulated cluster on a single node 45
- source environment variables 31
- sp_checkreswords system procedure 100

Index

- sp_lmconfig
 - determining current edition 6
 - starting
 - cluster using sybcluster 54
 - cluster when operating system starts 55
 - Sybase Control Center 45
 - stopping
 - Adaptive Server after a failure 155
 - cluster using sybcluster 54
 - servers 155
 - Sybase Control Center 45
 - Storage Foundation
 - creating a cluster 46
 - stored procedures
 - changing after renaming objects 99
 - style conventions
 - fonts 1
 - Sybase Central 7
 - Sybase Control Center 7
 - starting 45
 - stopping 45
 - SYBASE directory 31
 - sybase user
 - creating accounts 22
 - sybcluster
 - configure XP Server 59
 - configuring auxiliary servers 55
 - configuring the cluster 50
 - creating a cluster 50
 - starting a cluster 54
 - stopping a cluster 54
 - upgrading a noncluster to cluster Adaptive Server 111
 - sybsystemprocs database
 - increasing the size of 105
 - SySAM
 - configure IPv6 14
 - FLEXnet Publisher 13
 - grace period 143
 - license checkout 14
 - license server version 13
 - obtaining a license 11
 - sub-capacity licensing 13
 - SySAM license server 13
 - sysmessages 140
 - system administrator password
 - setting 64
 - system requirements 96
 - IBM 14
 - PC-Client 18
 - private interconnect technology 17
 - system stored procedures 140
 - system tables 140
 - changes due to upgrade 99
- ## T
- tasks, administration 22
 - technical support
 - calling 149
 - Tivoli Storage Manager
 - described 7
 - troubleshooting
 - by using error log files 155
- ## U
- uninstalling
 - existing server 37
 - PC-Client 41
 - registry entries 37
 - server 37
 - the old server 37
 - unknown cause of failure 156, 158
 - unrecognized tokens or objects 140
 - updating JAR files 117
 - updating XML files 117
 - upgrade requirements 96
 - upgrading
 - Adaptive Server 92
 - entire installations 92
 - failure 156
 - input file 113
 - interactively 113
 - Java in the database 129
 - single databases 92
 - sybcluster, with 111
 - tasks after you upgrade 116
 - using upgrade server command 113
 - upgrading a noncluster to cluster Adaptive Server sybcluster 111
 - upgrading Adaptive Server
 - Cluster Edition to Cluster Edition 108
 - upgrading Cluster Edition
 - shared to private installation 101
 - upgrading compatibility
 - versions 91

- upgrading compiled objects
 - dbcc upgrade_object 131
- upgrading high availability and cluster support 128
- upgrading Job Scheduler templates 124
- utilities 8
 - dsedit 121

V

- verifying connection to the server 62

- verifying readiness for upgrade 112

X

XP Server

- configure with sybcluster 59
- configuring with sybcluster 55

