



Release Bulletin

**Adaptive Server[®] Enterprise
Cluster Edition 15.5**

Sun Solaris

DOCUMENT ID: DC00762-01-1550-04

LAST REVISED: July 2011

Copyright © 2011 by Sybase, Inc. All rights reserved.

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

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

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

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

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 Sun Microsystems, Inc. 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

EBF Downloads for Security Fixes	1
Product Summary	3
Adaptive Server Interoperability	3
Product Compatibility	5
Password Compatibility with Replication Server	5
Using Multipathing with I/O Fencing	5
Changes that Affect Existing Applications	5
Backup Server and directio	6
Open Client and Open Server	6
Changed Functionality	6
Change to job scheduler tasks Configuration Parameter	6
Change to directio and dsync Settings for Raw Devices	6
Support for Database Resynchronization in Replication Server	7
Support for db2 Server Class	7
Increase in lock hashtable size	7
Change to sp_configure	8
Change in Accessing Temporary Objects that are Shareable Across Sessions	8
Assigned Temporary Database Inside an Application	8
Run-time Change to Identity Gap	8
Unsupported Features and Utilities	9
Installation and Upgrade	11
Special Installation Instructions	11
Special Licensing Instructions	11
Sub-capacity Licensing not Supported in Cluster Edition	11

- Special Configuration Instructions11
 - Number of Preallocated Extents12
- Special Upgrade and Downgrade Instructions12
 - Upgrading Job Scheduler12
 - Upgrading Adaptive Server12
 - Migrating from Adaptive Server 12.5.1 to 15.514
 - Loading Database Dumps During Downgrade14
- Known Installation Issues for Adaptive Server Cluster Edition15
- Known Issues19**
 - Known Issues for Adaptive Server Cluster Edition19
 - Known Plug-In Issues for Adaptive Server Cluster Edition24
- Documentation Changes25**
 - Adaptive Server Clusters Users Guide25
 - Adaptive Server Reference Manual: Building Blocks ...26
 - Adaptive Server Reference Manual: Commands27
 - Adaptive Server Reference Manual: Procedures29
 - Adaptive Server Reference Manual: Tables35
 - Adaptive Server System Administration Guide35
 - Adaptive Server Utility Guide36
- Obtaining Help and Additional Information37**
 - Technical Support37
 - Downloading Sybase EBFs and Maintenance Reports37
 - Sybase Product and Component Certifications38
 - Creating a MySybase Profile38
 - Accessibility Features38

EBF Downloads for Security Fixes

Due to possible security vulnerabilities, Sybase® strongly recommends that you use the appropriate EBF to fix these issues.

For instructions, see the *Urgent Customer Notification* Web page.

Known issues described in this release bulletin may have been fixed in other EBFs released for this product. Get EBFs and maintenance reports describing the latest fixes from the Sybase Web site at <http://www.sybase.com/support>.

Product Summary

This release bulletin provides late-breaking information about Adaptive Server® Enterprise version 15.5 ESD #1. A more recent version may be available on the Web.

Adaptive Server Enterprise server and client components are distributed on separate CDs or DVDs.

See the Cluster Edition installation guide for operating system requirements, and <http://certification.sybase.com> for additional supported operating systems.

Installation kit

The installation kit includes:

- The Server media (CD or DVD)
- The PC-Client media
- The Getting Started media with the following documentation, specific to your platform:
 - Installation guide
 - Release bulletin (this document)

Adaptive Server Interoperability

Learn about interoperability of Adaptive Server against other Sybase® products, across different platforms, versions, and client products.

Interoperability between big-endian and little-endian platforms has been verified. Windows, Linux x86-32, Linux x86-64, Sun Solaris x86-32, and Sun Solaris x86-64 are little endian platforms. IBM AIX, Linux on Power, Sun Solaris Sparc, and HP-UX on PA-Risc and Itanium are big-endian platforms.

Note: To use new features of Adaptive Server Cluster Edition, make sure that your client supports them. See the client-specific documentation for information about the features your client supports. You may need to upgrade your client to use certain Cluster Edition features.

Table 1. Supported client platforms for Adaptive Server

Client name	Version	Supported platform	
Open Client™/ Open Server™	12.5.x	<ul style="list-style-type: none"> • AIX 32-bit • AIX 64-bit • HP-UX 32-bit • HP-UX 64-bit • HP-UX IA 32-bit • HP-UX IA 64-bit • Linux x86 32-bit • Linux Pseries 32-bit • Linux Pseries 64-bit • Linux x64 32-bit • Linux x64 64-bit • Linux IA 64-bit 	<ul style="list-style-type: none"> • Mac OSX 64-bit • SGI 32-bit • SGI 64-bit • Solaris 32-bit • Solaris 64-bit • Solaris x86 32-bit • Solaris x86 64-bit • TruUnix (Alpha) 64-bit • Windows x86 32-bit
	15.0.x , 15.5	<ul style="list-style-type: none"> • AIX 32-bit • AIX 64-bit • HP-UX 32-bit • HP-UX 64-bit • HP-UX IA 32-bit • HP-UX IA 64-bit • Linux x86 32-bit • Linux x64 32-bit • Linux x64 64-bit 	<ul style="list-style-type: none"> • Linux Pseries 32-bit • Linux Pseries 64-bit • Mac OS X Intel 32-bit • Solaris 32-bit • Solaris 64-bit • Solaris x86 32-bit • Solaris x86 64-bit • Windows x86 32-bit
jConnect™ for JDBC™	5.5, 6.0.x, 7.0.x	<ul style="list-style-type: none"> • All 	
ODBC by Sybase	12.5.x	<ul style="list-style-type: none"> • Linux x86 32-bit • Linux x64 32-bit 	<ul style="list-style-type: none"> • Mac OSX 32-bit • Windows x86 32-bit
	15.0.x , 15.5.x	<ul style="list-style-type: none"> • Linux x86 32-bit • Linux x64 32-bit • Linux x64 64-bit • Mac OS X Intel 32-bit 	<ul style="list-style-type: none"> • Windows x86 32-bit • Windows x64 32-bit • Windows x64 64-bit

Client name	Version	Supported platform	
OLE DB by Sybase	12.5.x	<ul style="list-style-type: none"> Windows x86 32-bit 	
	15.0.x , 15.5.x	<ul style="list-style-type: none"> Windows x86 32-bit Windows x64 64-bit 	<ul style="list-style-type: none"> Windows x64 32-bit
ADO.NET	1.x, 2.x	<ul style="list-style-type: none"> Windows x86 32-bit Windows x64 32-bit 	<ul style="list-style-type: none"> Windows x64 64-bit
Replication Server®	15.1, 15.2, 15.5	<ul style="list-style-type: none"> AIX 32-bit HP-UX IA 64-bit Linux x64 32-bit Linux on IBM PSeries 64-bit 	<ul style="list-style-type: none"> Solaris Sparc 32-bit Solaris x86 32-bit Windows x86 32-bit

Note: Adaptive Server is supported on Sybase Central™ 6.0 and Sybase Control Center 3.1.

Product Compatibility

Learn about Adaptive Server compatibility.

Password Compatibility with Replication Server

The master databases in Adaptive Server 15.5 and later can replicate only to Adaptive Server 15.0.2 and later installations.

Using Multipathing with I/O Fencing

Multipathing with I/O fencing enabled is not supported.

Using multipathing may result in I/O errors, which prevent the data server from running if you have I/O fencing enabled. To prevent I/O errors, disable multipathing on all affected database devices.

Changes that Affect Existing Applications

Learn about changes in Adaptive Server 15.5 ESD #1 that affect your existing applications.

Backup Server and directio

In version 15.5, both Adaptive Server and Backup Server can bypass the operating system buffer cache when you enable the **directio** parameter for the device using **disk init**, **disk reinit**, or **sp_deviceattr**. Adaptive Server passes the device options to Backup Server, which enables Backup Server to access the database device with the appropriate **directio** option.

Open Client and Open Server

For information about changes that affect Open Client and Open Server, see the most recent release bulletins for these products on the Sybase Web site.

Changed Functionality

Learn about late-breaking functionality changes in Adaptive Server version 15.5 ESD #1.

For complete information, see the *Clusters Users Guide*.

Change to job scheduler tasks Configuration Parameter

The default value for configuration parameter **job scheduler tasks** has been changed to 4 in Adaptive Server 15.5 ESD #1.

However, for compatibility with RAP - The Trading Edition R4, you must set **job scheduler tasks** to 32 using:

```
sp_configure "job scheduler tasks", 32
```

If you change the default value of the job scheduler tasks, you may also need to increase the **number of user connections** in the Adaptive Server by twice the value of **job scheduler tasks**. If the same Adaptive Server hosts the Job Scheduler and the jobs, you may need to increase **number of user connections** in the Adaptive Server by three times the value of **job scheduler tasks**. If you increase the **number of user connections**, you may need to increase **max memory** in the Adaptive Server.

Note: If you set the value of **job scheduler tasks** to "default" before you upgrade Adaptive Server, the server automatically sets the new default to 4.

If the initial value of **number of user connections** is 50, and you want to set **job scheduler tasks** to 10, increase **number of user connections** by $2*10$ and set the parameter to $50+20 = 70$. If the Job Scheduler is running on the same server, increase **number of user connections** by $3*10$ and set the parameter to $50+30 = 80$.

Change to directio and dsync Settings for Raw Devices

On raw devices, you cannot:

- Set the **directio** or **dsync** option in the **disk init** command to "**true**"
- Set the **directio** or **dsync** option via the **sp_deviceattr** stored procedure to "**true**"

Doing so will return a message such as:

```
You cannot set directio option for raw device '/dev/raw/
raw235' or You cannot set attribute dsync for raw device
'myrawdsk1'.
```

Support for Database Resynchronization in Replication Server

Adaptive Server 15.5 ESD #1 supports database resynchronization in Replication Server 15.5. The Adaptive Server Rep Agent can be configured to send a resync marker to the Replication Server when Adaptive Server is restarted. See *Replication Server New Features Guide* for more details.

Support for db2 Server Class

Adaptive Server does not support server class db2.

To use **db2**, migrate your **db2** server class to **direct_connect** class.

Increase in lock hashtable size

Increase **lock hash table size** to improve query response time with Adaptive Server 15.5 Cluster Edition.

You may experience delayed query response time with Adaptive Server 15.5 Cluster Edition if the size of the lock hash table is insufficient. To avoid this, set the value of **lock hashtable size** using:

$$\text{lock hashtable size} = ((\text{number of locks} * (4 + \text{cluster redundancy level})) + ((\text{total data cache size (in KB)} / \text{max database page size (in KB)}) * (2 + \text{cluster redundancy level}))) / 8$$

If the new value of **lock hashtable size** is more than the previously configured value, you must also increase the memory configuration of your Adaptive Server by 24 bytes for every byte of increase of **lock hash table size**.

Note: Adaptive Server may also warn you that allocation of regular Cluster Interprocess Communication (CIPC) messages fails due to lack of message buffers, and to increase the **CIPC regular message pool size** configuration parameter.

Example 1: Determining the lock hashtable size

If number of locks = 50000, cluster redundancy level = 1, total data cache size = 7000MB, and max database page size = 2KB, set **lock hashtable size** to: $((50,000 * (4 + 1)) + ((7,000,000/2) * (2+1)))/8 = 1343750$.

Example 2: Determining the amount of additional memory

If the previously configured value of **lock hashtable size** is 1,000 and the newly computed size is 11,000, then increase memory by: $24 * (11,000 - 1,000) = 240,000$ bytes.

Change to sp_configure

When using **sp_configure** with the Cluster Edition:

- Configure an instance-specific value that applies only to the instance to which you are currently connected.
- Configuration options that are “strictly cluster-wide”—that is, the small group of configuration options that are the same on all instances—cannot be made instance-specific.
- Reconfiguring a cluster-wide value does not overwrite an instance-specific configuration.

Change in Accessing Temporary Objects that are Shareable Across Sessions

Applications created since Adaptive Server version 12.5.0.3 may create shareable tables in user-created temporary databases. To enable existing applications to continue to work with the Cluster Edition, drop user-created temporary databases that include these shareable tables, and re-create them as global temporary databases with the same names.

Assigned Temporary Database Inside an Application

In a clustered environment, the default temporary database assigned to a session is a local, temporary, instance-specific database. It is not the system *tempdb* (with a *dbid* of 2). After upgrading from a nonclustered environment, you may need to modify applications that previously used the system *tempdb* (with a *dbid* of 2) as the default temporary database in their actions, so that these actions are now applied to the assigned local temporary database in the Cluster Edition.

For example, in a nonclustered environment, an application truncates the log of the default temporary database when you execute:

```
dump tran tempdb with truncate_only
```

In the Cluster Edition, modify the application to truncate the log of the default temporary database:

```
declare @tempdbname varchar(30)
select @tempdbname = db_name(@@tempdbid)
dump tran @tempdbname with truncate_only
```

Run-time Change to Identity Gap

In the Cluster Edition, you cannot use the **sp_chgattribute** stored procedure to change the value of **identity_gap** at runtime.

Unsupported Features and Utilities

Learn about features that are not supported in the Cluster Edition.

These utility programs are not supported, or have minimum support:

- **sqlupgrade**
- **sqlloc**
- **srvbuild**

These Adaptive Server features are not supported:

- In-memory databases introduced in Adaptive Server 15.5.
- Adaptive Server disk mirroring.
- High-availability. In the Cluster Edition, high-availability functionality is integrated into Adaptive Server.
- Replication using ASE Replicator.
- Maximum number of instances supported. The Adaptive Server Cluster Edition architecture provides for up to 32 instances in a cluster. However, 15.5 Cluster Edition is currently certified for, and supported on, 4 instances only.
- The **shutdown with wait = time** feature introduced in Adaptive Server 15.0.2. If you include with **wait = time** with your **shutdown** command, you see an error message similar to:

```
Command "WITH WAIT=nn:nn:nn" is unsupported
```

and the shutdown occurs without any wait.

Note: The ASE-15_0/bin directory includes a utility called **iofenceutil**, which duplicates the functionality of **qrmutil --fence_capable**, but does not require the quorum device and also returns an exit code. This is an internal utility used by **sybcluster** and is not designed for regular use.

Installation and Upgrade

Get last-minute information about installation and upgrading that was omitted from or incorrect in your installation guide, or that needs special emphasis.

Sybase strongly recommends that you install the Adaptive Server Enterprise 15.5 Cluster Edition into its own directory. If you cannot do this, install the Cluster Edition last.

Warning! Sybase recommends that you do not install an older version of a Sybase product in a \$SYBASE directory that includes a newer Sybase product. For example, installing Replication Server® 15.0 on top of the Cluster Edition may make either the Cluster Edition or Replication Server inoperable, depending on the choices you select during installation.

Special Installation Instructions

Learn about special installation instructions for this version of Adaptive Server.

Installing Enterprise Connect Data Access (ECDA) or MainframeConnect DirectConnect for z/OS with Other Sybase Software

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

Installing Adaptive Server Enterprise Cluster Edition with Sybase IQ 15.1

Sybase recommends that you not install Sybase IQ 15.1 in a directory that includes the Adaptive Server Enterprise 15.5 Cluster Edition.

Special Licensing Instructions

Learn about special licensing instructions for this version of Adaptive Server.

Sub-capacity Licensing not Supported in Cluster Edition

Adaptive Server 15.5 ESD #1 Cluster Edition does not support sub-capacity licensing.

Special Configuration Instructions

Learn about special configuration instructions for this version of Adaptive Server.

Number of Preallocated Extents

In earlier versions of Adaptive Server, the maximum value allowed for the number of preallocated extents was 31. In Adaptive Server version 15.5 and later, the maximum values of preallocated extents has been increased to 32.

Using a value of 32 for the **number of preallocated extents** has a special significance for configuration, and impacts the space allocations Adaptive Server performs internally. If you set the number of preallocated extents to 32, Adaptive Server attempts to reserve an entire allocation unit of extents for utility operations that use a large-scale allocation scheme of space reservation, such as **bcp in** and **select into**.

Using the maximum number of preallocated extents can greatly improve the performance of these utilities, particularly when you run them in parallel. Using a value of 32 greatly increases the likelihood that each engine running the utility can work independently on its own allocation unit without interference from other engines.

Special Upgrade and Downgrade Instructions

Learn about special upgrade and downgrade instructions for this version of Adaptive Server.

Upgrading Job Scheduler

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

Upgrade any jobs that are created from templates.

Upgrading Adaptive Server

Learn about how to upgrade Adaptive Server.

A minor change was made to the format of a log record in both the cluster and non-cluster editions of Adaptive Server 15.5 that 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 that some steps be strictly followed when you move from Adaptive Server 15.0.x or earlier to Adaptive Server 15.5.x (non-cluster editions). See the following table for all possible combinations of upgrades.

An upgrade is done by:

1. Upgrading your entire installation by switching binaries.
2. Upgrading a single database using **online database** having first loaded it with a database dump and transaction logs taken on a server on a lower version.

Table 2. Upgrading Your Entire Installation

Current version	Upgrading to	Special Upgrade Information
Adaptive Server 15.0.x or earlier	Adaptive Server 15.5.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 section, "Upgrading Adaptive Server in a replication system" in Chapter 3 of the configuration guide for your platform in the Replication Server documentation for more information on this. Also, see "Upgrading Servers With Replicated Databases" in the chapter on "Upgrade" in the Adaptive Server installation guide for your platform.
Adaptive Server 15.0.x	Adaptive Server Cluster Edition 15.5.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 section, "Upgrading Adaptive Server in a replication system" in Chapter 3 of the configuration guide for your platform in the Replication Server documentation for more information on this. Also, see "Upgrading Servers With Replicated Databases" in the chapter on "Upgrade" in the Adaptive Server installation guide for your platform.
Adaptive Server 15.5.x	Adaptive Server Cluster Edition 15.5.x	Not supported.
Adaptive Server Cluster Edition 15.5.x	Adaptive Server 15.5.x	Not supported.

Table 3. Upgrading a single database

Current version	Upgrading to	Special Upgrade Instructions
Adaptive Server 15.0.x and earlier	Adaptive Server 15.5.x	<p>When you use online database to upgrade a single database in Adaptive Server 15.5.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.</p> <p>See the section "Reenabling Replication Server After the Dataserver Upgrade" in the "Upgrade" chapter in the Adaptive Server installation guide for your platform.</p>

Current version	Upgrading to	Special Upgrade Instructions
Adaptive Server 15.5.x	Adaptive Server Cluster Edition 15.5.x	Upgrading a single database from the non-cluster edition of 15.5 or 15.5 ESD 1 to cluster edition of 15.5 or 15.5 ESD 1 is supported, and no particular steps must be emphasized in order to ensure a smooth upgrade.
Adaptive Server Cluster Edition 15.5.x	Adaptive Server 15.5.x	Upgrading a single database from Adaptive Server Cluster Edition version 15.5 or 15.5 ESD 1 to the non-cluster 15.5 or 15.5 ESD 1 versions is not supported and its prevention is enforced by Adaptive Server.

Migrating from Adaptive Server 12.5.1 to 15.5

Migrate from Adaptive Server version 12.5.1 to 15.5 using **sybmigrate** may generate an error message.

When you migrate from Adaptive Server version 12.5.1 to 15.5 using the **sybmigrate** utility, you may encounter an error message similar to:

```
"local server name not set for server 'source' on
host 'host1'(IP address + port number='xx.xx.xx.xxx:xxxx')"
```

Adaptive Server version 12.5.1 is no longer supported, and Adaptive Server does not automatically recognize the 12.5.1 server. Sybase recommends that you upgrade from 12.5.1 to 12.5.4, then from 12.5.4 to 15.5.

To make the 12.5.1 server visible to the 15.5 server, use:

```
1>sp_addserver 'servername', local
2>go
```

Restart Adaptive Server. It now recognizes the 12.5.1 server.

Loading Database Dumps During Downgrade

Adaptive Server does not automatically block attempts to downgrade to a minor version; however, do not attempt this, as the database dump can corrupt the server into which you are loading the dump.

When you downgrade from a version of Adaptive Server to a lower major version (such as version 15.x to 12.5x), Adaptive Server blocks the attempt with a message similar to:

```
ASE cannot load this database dump because the version in the
dump is newer than the version in the database. Database: log
version=7,
upgrade version=35. Dump: log version=7, upgrade version=40.
```

To load the dumps of 15.5 into a server of a minor version (such as 15.0.1):

1. Execute **sp_downgrade** in the 15.5 version Adaptive Server to downgrade the databases to be dumped.
2. After the downgrade succeeds, dump the downgraded databases in the 15.5 server.
3. Load the dumps of downgraded databases into the 15.0.1 server.

Known Installation Issues for Adaptive Server Cluster Edition

Learn about known installer issues and workarounds. Known issues are listed in descending order of Change Request (CR) numbers.

CR#	Description and Workaround
629400	<p>Installer incorrectly copies 64-bit JRE to Windows 32-bit.</p> <p>If the 32-bit Windows contains registry key HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node, the installer incorrectly installs the 64-bit JRE. This registry key should not exist in 32-bit Windows, and prevents Windows from running any Java applications.</p> <p>Workaround: Go to Control Panel > System Properties > Advanced and:</p> <ul style="list-style-type: none"> • Set the environment variables for 32-bit JRE: <ul style="list-style-type: none"> • SYBASE_JRE6=<Installed Directory>\Shared \JRE-6_0_6_32BIT • SYBASE_JRE6_32=<Installed Directory>\Shared \JRE-6_0_6_32BIT • Delete the 64-bit JRE environment variable:SYBASE_JRE6_64
624152	<p>SySAM utilities may generate error message in non-English environment.</p> <p>When running the SySAM utilities sysamcap or cpuinfo in a non-English environment, you may see an error message similar to this, in English, even though you have <code>sylapi.lcu</code> installed:</p> <pre>Failed to open the localised message file '../locales/unicode/sylapi/zh_cn/sylapi.lcu' (i18nuni_FileOpen returned -479).</pre> <p>Workaround: You can safely ignore this message.</p> <p>To prevent the message from displaying altogether, set this environment variable before you run sysamcap or cpuinfo: <code>LANG=en_US</code></p>

CR#	Description and Workaround
611780	<p>Installing Adaptive Server 15.5 Cluster Edition in the same directory as ECDA 15.0 ESD #2 and earlier versions results in errors.</p> <p>The Unified Agent does not start after this installation. Additionally, you cannot create a cluster from Sybase Central.</p> <p>Workaround: Sybase strongly recommends that you install the Adaptive Server Enterprise 15.5 Cluster Edition into its own directory. If you must install Adaptive Server 15.5 Cluster Edition in the same directory as ECDA 15.0, then:</p> <ul style="list-style-type: none"> • If you have installed ECDA 15.0 ESD #3 , you can install Adaptive Server 15.5 Cluster Edition on top of the ECDA installation. • If you have an ECDA installation that is version 15.0 ESD #2 or earlier, you must: <ul style="list-style-type: none"> • Uninstall the older ECDA version. • Delete directory \$SYBASE/ua. • Install Adaptive Server 15.5 Cluster Edition.
595614	<p>Installer cannot create installer directory on Microsoft Windows 2008.</p> <p>On Microsoft Windows 2008, the installer can create the Adaptive Server installation directory only when you are logged in as an administrator, even if your role has permission to create directories.</p> <p>Workaround: Manually create the installation directory before running the installer.</p>
595582	<p>Installer fails to launch if you specify a path to setup.bin that contains “..”</p> <p>Workaround: Ensure that the path to setup.bin does not contain “..”.</p>
595573	<p>Uninstallation program stops responding.</p> <p>The uninstallation program stops responding when you select Delete from the Delete User Files screen and then click Previous from the Uninstall Complete screen.</p> <p>Workaround: Do not return to the Delete User Files screen after you have selected Delete.</p>
593410	<p>Cannot run installer with nondefault "tar" tool.</p> <p>If the GNU tar utility is listed in the path before the operating system-supplied tar when you invoke the tar command, the installer fails with this error:</p> <p>The included VM could not be unarchived (TAR). Please try to download the installer again and make sure that you download using 'binary' mode. Please do not attempt to install this currently downloaded copy.</p> <p>Workaround: Define /usr/bin/tar in \$PATH.</p>

CR#	Description and Workaround
593389	<p>The 32-bit installer fails on Microsoft Windows Vista and Windows 2008 on x86-64 64-bit because the installer cannot set up the environment variables.</p> <p>Workaround:</p> <ol style="list-style-type: none"> 1. In Windows Explorer, right-click <code>setup.exe</code> or <code>setupConsole.exe</code>, and select Properties. On the Compatibility tab, select the Windows Vista compatibility mode. 2. Add the environment variables by typing the text, or using Ctrl+V to paste the content into the text area, as clicking on SySAM License Panel > Browse does not work.
593148	<p>Installation fails on Sun Solaris SPARC.</p> <p>When the file descriptor limit is set to "unlimited," the installer fails with this message:</p> <pre>awk: insufficient memory for string storage Context is: >>> <<<</pre> <p>Workaround: Set the file descriptor limit to a number.</p>
590282	<p>You cannot use the tab and arrow keys to navigate the Choose Product Features pane of the installation or uninstallation program.</p> <p>Workaround: Use the mouse to click in the Install or Uninstall window. This sets the focus on the window and allows you to navigate using your keyboard.</p>
588793	<p>Installing Adaptive Server 15.5 (which uses Install Anywhere) on older directories that have been installed with InstallShield Multiplatform can lead to incompatibility issues.</p> <p>These issues occur while using products installed with different installer technologies:</p> <ul style="list-style-type: none"> • If you install products using Install Anywhere or InstallShield Multiplatform on top of a product installed using the other installer, the same files included in both installers are silently overwritten by the later installation. • If you run either the Install Anywhere or InstallShield Multiplatform uninstaller, the same files installed by both installers are removed without any indicating messages. <p>Workaround: Do not install products that use Install Anywhere and products that use InstallShield Multiplatform in the same directory.</p>
583979	<p>Feature names are not validated when installing in silent mode.</p> <p>The installer does not validate feature names specified in the response file when you install in silent mode.</p> <p>Workaround: Ensure that the specified feature names are correct.</p>

CR#	Description and Workaround
579988	<p>"/setup.bin: !: not found" message appears.</p> <p>You see this error message when you install Adaptive Server:</p> <pre>Preparing to install... ./setup.bin: !: not found Extracting the installation resources from the installer archive... Configuring the installer for this system's environment...</pre> <p>Workaround: Ignore the message and continue with the installation.</p>

Known Issues

Learn about known issues and workarounds for this version of Adaptive Server and Adaptive Server plug-in.

Known Issues for Adaptive Server Cluster Edition

Learn about known issues and workarounds for the Adaptive Server Enterprise Cluster Edition. Known issues are listed in descending order of Change Request (CR) numbers.

Note: Sybase does not include system problem reports (SPRs) and closed problem reports (CPRs) with Adaptive Server Enterprise. You can search the Web site for solved cases. Click **Support > Services > Solved Cases**.

CR#	Description and Workaround
630077	<p>In rare cases, out-of-date log pages persist in cache after a node failure in cluster configurations if both these conditions occur:</p> <ul style="list-style-type: none"> • Two or more nodes survive the node failure • Failover recovery is needed for a multi-database transaction that has completed only the first phase of the two-phase commit protocol. <p>If the database is replicated, the stale log pages may also result in errors during replication.</p> <p>Workaround: None.</p>
618851	<p>Terminating a session that is executing a query using session temporary tables results in errors.</p> <p>When a session executing a query using session temporary tables is terminated by another session, and statement cache is enabled in Adaptive Server, the error log reports this error: <code>Current process infected with 11 in the module s_free</code>, and may include a stack trace with modules <code>clean_process</code> and <code>kill_proc</code>. Under some circumstances, Adaptive Server also reports error 6103: <code>Unable to do cleanup for the killed process</code>.</p> <p>Workaround: Use one of these:</p> <ul style="list-style-type: none"> • Turn off the statement cache. • Use <code>traceflag 467</code> in the command line. • If there is a successfully compiled plan in the cache, use <code>traceflag 299</code> in the command line.

CR#	Description and Workaround
615204	<p>If enable DTM is set, the distributed recovery feature is disabled and all databases are recovered only on the cluster coordinator, rather than on different instances of the cluster.</p> <p>Workaround: None.</p>
612683	<p>sybcluster and Adaptive Server Plug-in can validate Veritas Cluster Server (VCS) raw device and file systems with the actual path name, but cannot validate them for link paths.</p> <p>For example,</p> <ul style="list-style-type: none"> • Name of a VCS raw device is rawlg4. • The actual path to the device is dev/vx/rdisk/cfsdg/rawlg4 . • The link path to the device is /work1/rawlg4. <p>The link to the device is obtained by executing <code>ln -s /dev/vx/rdisk/cfsdg/rawlg4</code>.</p> <p>If you create a cluster using sybcluster or Adaptive Server plug-in and enter <code>dev/vx/rdisk/rawlg4</code> for the path to the master device or other database device, sybcluster and Adaptive Server plug-in correctly detect the device managed by VCS. If you create a cluster and enter <code>/work1/rawlg4</code> for the path to the master device, sybcluster or Adaptive Server plug-in prints:</p> <pre style="background-color: #f0f0f0; padding: 5px;">"This device is not managed by Veritas Cluster File System or Veritas Volume Manager. Do you want to continue (Y/N)? "</pre> <p>.</p> <p>Workaround: Ignore this message. sybcluster cannot validate a VCS raw device that uses a link path, but can use the link path to build the cluster.</p>
612447	<p>When you add any value into a bigtime field, and execute select dateadd using that field as a parameter, you see an overflow error.</p> <p>Workaround: Convert the bigtime value to bigdatetime before executing the Transact-SQL function dateadd.</p>
612286	<p>You cannot load database and transaction log dumps from Adaptive Server 15.5 Cluster Edition into a database on Adaptive Server 15.0.x or 15.5 (nonclustered) versions.</p> <p>Workaround: Use one of these methods to load the database or transaction dump from Adaptive Server 15.5 Cluster Edition:</p> <ul style="list-style-type: none"> • Load the database or transaction dumps into a database on Adaptive Server 15.5 Cluster Edition. • Use Replication Server to replicate the database or transaction dump. • Use sybmigrate, the Sybase migration tool. See the <i>Utility Guide</i>.

CR#	Description and Workaround
607464	<p>Recovering a template displays an error.</p> <p>In some cases, when a database that is used as a template for creating a database with a durability of no_recovery is being recovered, a 937 error may be reported to the error log:</p> <p>The database '<template_dbname>' is currently unavailable. It is being used to create a new database.</p> <p>Workaround: This error does not impact the correct recovery of any databases. Following the completion of recovery, however, the default data cache may have a memory pool configuration that differs from the configured values. Execute sp_poolconfig and reconfigure the default data cache to required values.</p>
595923	<p>There are two known issues with transfer table :</p> <ul style="list-style-type: none"> • When creating a unique index that is using the ignore_dup_key property, importing a row with a key that is already present in the table leads to an error, and the import is aborted. This differs from when a duplicate key is inserted with bcp or with a regular insert statement, since in those cases, the row is discarded, and the transaction continues. • When an insert trigger exists, the trigger is not fired when data is inserted through transfer table...from.
586499	<p>Adaptive Servers configured for Lightweight Directory Access Protocol (LDAP) user authentication may display error messages when authenticating several users simultaneously.</p> <p>The error is similar to:</p> <p>Unexpected AIO error(Invalid argument) returned from aio-wait().</p> <p>and occurs when LDAP user authentication has been enabled by setting configuration parameter enable ldap user auth to either 1 or 2.</p> <p>This may result in a non-terminating wait for some Adaptive Server client tasks.</p> <p>Workaround: Disable LDAP user authentication by setting enable ldap user auth to 0.</p>
582815	<p>Under certain conditions, you may need to re-create abstract plans associated with stored procedures that include the built-in function isnull.</p> <p>These are the conditions:</p> <ul style="list-style-type: none"> • One of the parameters to isnull is a char datatype and another is a literal. • enable literal parameterization is set to on. <p>Workaround: Re-create the abstract plans associated with the stored procedure.</p>

CR#	Description and Workaround
572710	<p>Taking an engine offline while Java clients are connected may cause an error.</p> <p>If Adaptive Server users execute Java-in-SQL statements while executing <code>sp_engine "offline"</code>, <code>engine_id</code>, Adaptive Server may report stack traces in its error log if the engine is hosting a JVM.</p> <p>Workaround: Do not take the engine offline until Java clients are disconnected. Instead, take the engine offline that is not hosting the JVM.</p>
545400	<p>Restarting Adaptive Server with Java in the database feature on Sun Solaris may fail.</p> <p>Adaptive Server configured with Java in the database on a Sun Solaris platform may fail to restart. Due to a Sun Solaris bug #6671882 JVM, may require an extremely large amount of memory.</p> <p>Workaround: A workaround fix is included Adaptive Server 15.5. The fix that can handle the incorrect memory request as long as the request is within a limit that can be satisfied by heap memory. A warning message prints in the error log when the bug #6671882 is encountered:</p> <pre data-bbox="323 704 1174 805">kernel PCI(T1): WARNING: pci_malloc; forwarded size 954437177 exceeding max request size 2097072 to the run-time for thread 13</pre>
487525	<p>Starting Unified Agent may report a permissions error.</p> <p>When you initially execute the <code>uafstartup.sh</code> script, you may receive an error stating that you do not have permissions to access some files. This problem is caused by latency issues in the NFS file system layer, and occurs only when you attempt to start multiple Unified Agents simultaneously in the current <code>\$SYBASE</code> directory.</p> <p>Workaround: Wait a few seconds, then retry.</p>
485070	<p>Restarting an unfinished Cluster Creation wizard session</p> <p>If you exit from the Cluster Creation wizard during a cluster configuration and then restart the wizard session using the same configuration parameters (cluster name, instance name, and so on), the wizard may have already created some configuration files and devices.</p> <p>Workaround: Before you restart a wizard session:</p> <ol data-bbox="323 1260 1174 1451" style="list-style-type: none"> 1. Stop the <code>srvbuildres</code> or <code>dataserver</code> utilities, if either are running. 2. Stop the Unified Agents on all nodes. 3. Remove the directory with the name of the cluster you tried to create from the <code>\$SYBASE_UA/nodes/node_name/plugins</code> directory. 4. Remove the interfaces file entries for the cluster you tried to create. 5. Restart the Unified Agents on all nodes.

CR#	Description and Workaround
483651	<p>Incorrect cluster failover may occur during diagnostic shared-memory dump.</p> <p>If you start an instance while the rest of the cluster is performing a diagnostic shared memory dump, the instance that is starting may incorrectly perform a cluster takeover. This only occurs if automatic cluster takeover is set to 1 or if the --cluster_takeover option is passed to the dataserver. In environments without i/o fencing enabled, this may lead to data corruption.</p> <p>Workaround: Do not start an instance during a diagnostic shared-memory dump. Configure automatic cluster takeover to 0.</p>
466370	<p>If Adaptive Server is configured for a large number of spinlocks, you may see a performance issue with sp_sysmon.</p> <p>This is because index creation on the worktable used by sp_sysmon fails due to duplicate key values. sp_sysmon reports error 1508:</p> <pre>Create index aborted on duplicate rows.</pre> <p>The large number of spinlocks on Adaptive Server account for a large number of rows in the table, and result in duplicate key values. sp_sysmon performance is slower due to the lack of an index on the worktable.</p> <p>Workaround: Run sp_sysmon with the clear option.</p> <hr/> <p>Note: The clear option will reset the monitor counter values on your server.</p>
462184	<p>Monitor Server and Historical Server do not work with an Adaptive Server that has set its net password encryption reqd options set to 1 or 2.</p> <p>Workaround: None.</p>
439406	<p>Configuring Job Scheduler Self Management for Simplified Chinese</p> <p>When you configure the Job Scheduler Self Management feature for an Adaptive Server localized for Simplified Chinese, you may receive an error saying that a <i>varchar</i> datatype cannot be converted to a date. The Job Scheduler configuration process stops at this point.</p> <p>Workaround: Do not configure the Self Management feature for Simplified Chinese.</p>
439404	<p>Creating Job Scheduler schedules for Adaptive Server with Simplified Chinese.</p> <p>When you attempt to create a Job Scheduler schedule using the Adaptive Server plug-in for an Adaptive Server that is localized for Simplified Chinese, you may receive an error indicating that the <i>varchar</i> datatype cannot be converted to a date.</p> <p>Workaround: Localize the Adaptive Server for Simplified Chinese before configuring Job Scheduler.</p>

Known Issues

CR#	Description and Workaround
400784	<p>Starting Monitor Client GUI fails when the Lightweight Directory Access Protocol (LDAP) server is used.</p> <p>When you start the Monitor Client GUI in an environment using LDAP instead of an <code>interfaces</code> file, the Monitor Client GUI cannot connect to any servers.</p> <p>Workaround: Create an <code>interfaces</code> file with entries for the servers that are monitored and use this with the Monitor Client GUI.</p>
382874	<p>Creating RSSD database fails with earlier versions of Replication Server.</p> <p><code>rs_init</code> in Replication Server version 12.6 and earlier cannot create an RSSD database in Adaptive Server Enterprise version 15.0 or later.</p> <p>Workaround: If you do not use an Embedded RSSD (ERSSD), manually create the RSSD database before using <code>rs_init</code> to create a new Replication Server. See the <i>Replication Server Administration Guide</i>.</p>

Known Plug-In Issues for Adaptive Server Cluster Edition

Learn about known issues and workarounds for Interactive SQL and the Adaptive Server Enterprise plug-in. Known issues are listed in descending order of Change Request (CR) numbers.

CR#	Description and Workaround
487524	<p>Raw device permission error from sybcluster while executing create cluster command.</p> <p>This happens when the name of the device is misspelled, or if no device by the specified name exists.</p> <p>Workaround : Reenter the name of an existing device.</p>
478875	<p>In the Adaptive Server plug-in, you cannot drop a connected instance; your connection is lost as soon as the instance is shut down.</p> <p>Workaround: Before shutting down the instance to drop, connect to another instance in the cluster.</p>

Documentation Changes

Read about updates, corrections, and clarifications to the documentation released with Adaptive Server Cluster Edition version 15.5.

Adaptive Server Clusters Users Guide

Read about updates, corrections, and clarifications for Adaptive Server Clusters Users Guide.

Using sp_showplan

Add this text to the "Additional Topics" chapter:

You cannot use **sp_showplan** across multiple nodes of a shared-disk cluster. It is specific to a single connection to a node.

Setting cluster redundancy level configuration parameter

Add this text to "Enabling multiple simultaneous failover" in the "Instance Recovery" chapter:

For the cluster to start, the value of the **cluster redundancy level** configuration parameter must be at least one less than the value of **maximum number of instances** as specified in `cluster.cfg` or the quorum file. Thus, the cluster cannot start if you set either of these:

- The value of **maximum number of instances** to a value that is equal to or less than the value of **cluster redundancy level**.
- The value of **cluster redundancy level** to a value equal to or greater than the value of **maximum number of instances**.

Upgrading from Adaptive Server 12.5.4 or earlier versions

Add this text to the "Upgrade" chapter:

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.

Changing the character set or sort order with the Veritas Cluster Server (VCS)

Add this text to the "Using the Cluster Edition with the Veritas Cluster Server" chapter:

To change the character set or sort order:

1. Use VCS commands to shut down Adaptive Server.
2. Change the membership mode to native:


```
>qrmutil -Q quorum_file --membership-mode="native"
```

Documentation Changes

3. Restart a single instance in the cluster either manually or using **sybcluster**.
4. To change the character set, run **charset**. For example:

```
>$SYBASE/ASE-15_0/bin/charset -Usa -P nocase.srt utf8
```
5. To change the sort order, reset the **default sortorder id** configuration parameter. For example:

```
>isql -Usa -P 1>sp_configure 'default sortorder id', 101, 'utf8'
```
6. From **isql**, shut down the cluster.
7. Restart an instance. Adaptive Server reconfigures the character set and sort order, and shuts down.
8. Change the membership mode back to VCS:

```
>qrmutil -Q quorum_file - -membership-mode='vcs'
```
9. Use VCs commands to restart the cluster.

Changing the number of user connections with the Veritas Cluster Server (VCS)

Add this text to the "Using the Cluster Edition with the Veritas Cluster Server" chapter:

When you increase the number of user connections, you must also increase the file descriptor limit. However, Veritas running over VCS does not propagate file descriptor limit changes from one user session to another.

For example, if you change the file descriptor limits in your current session using: `ulimit -n 8194`, subsequent UNIX sessions continue to use the default setting (1024) for the number of file descriptors. Consequently, while the Veritas VCS manages the cluster and brings instances online, it logs in to the Linux operating system as the user "sybase" (starting a new UNIX session), and uses the default value for the number of file descriptors.

To ensure that all sybase user sessions have the correct file descriptor limit, add these lines to `/etc/security/limits.conf`:

```
sybase hard nofile 8096
sybase soft nofile 8096
```

Adaptive Server Reference Manual: Building Blocks

Read about updates, corrections, and clarifications for Adaptive Server *Reference Manual: Building Blocks*.

@@curloid global variable

Chapter 3, "Global Variables," incorrectly defines *@@curloid* as:

```
Either no cursors are open, no rows qualify for the last opened
cursor, or the last open cursor
is closed or deallocated.
```

Replace with:

Returns the current session's lock owner ID.

time datatype

The **time** values in Adaptive Server are accurate to 1/300 second, however, some examples that include **date** and **time** values in the book show millisecond values.

The last digit of the fractional second is always 0, 3, or 6. Other digits are rounded to one of these three digits, so 0 and 1 round to 0; 2, 3, and 4 round to 3; 5, 6, 7, and 8 round to 6; and 9 rounds to 10.

Adaptive Server Reference Manual: Commands

Read about updates, corrections, and clarifications for Adaptive Server *Reference Manual: Commands*.

reorg

The reference page for **reorg** includes this incorrect note:

Note: You cannot run **reorg** on text that has an indid of 255 in sysindexes.

Replace with:

Note: **reorg** has no effect on space allocated to text or image columns.

alter table

The Auditing section of the **alter table** reference page includes this bullet item, which is incorrect:

- If the **with** parameter contains **set transfer table on**, Adaptive Server prints this to extrainfo: SET TRANSFER TABLE ON. If the **with** parameter contains **set transfer table off**, Adaptive Server prints this to extrainfo: SET TRANSFER TABLE OFF.

Replace with this information:

- If the **set** option for **set transfer table [on | off]** is:
 - **on** – Adaptive Server prints SET TRANSFER TABLE ON in the extra info in the audit record.
 - **off** – Adaptive Server prints SET TRANSFER TABLE OFF.

create table

Add this text to the "Information in extrainfo" column of the Auditing section of the **create table** reference page:

- If the **with** option for **with transfer table [on | off]** is:
 - **on** – Adaptive Server prints WITH TRANSFER TABLE ON in the extra info in the audit record.

- **off** – Adaptive Server prints `WITH TRANSFER TABLE OFF`.

dump database and dump transaction

There are two updates for **dump database** and **dump transaction**:

- The Parameters section of both **dump database** and **dump transaction** include the following:
 - **compress:: compression_level** is a number between 0 and 9, 100, or 101. For single-digit compression levels, 0 indicates no compression, and 9 provides the highest level of compression. Compression levels of 100 and 101 provide faster, more efficient compression, with 100 providing faster compression and 101 providing better compression. If you do not specify *compression_level*, Adaptive Server does not compress the dump.
The 100 and 101 compression levels do not work with **compress:: compression_level**. Specifying these values causes the command to return an error.
You may, however, continue to use 100 and 101 compression levels with the **with compression=** parameter.
- The maximum file path/name size for a physical device is 127 characters. This limitation is not explicitly stated in the Restrictions subsection of the Usage section of the reference pages for **dump database** and **dump transaction**.

execute

The Usage section includes this bullet item:

You cannot supply *string* and *char_variable* options to execute the following commands: **begin transaction**, **commit**, **connect to**, **declare cursor**, **rollback**, **dump transaction**, **dbcc**, **set**, **use**, or nested **execute** commands.

Adaptive Server currently uses a more relaxed list, which includes:

- **use**
- **exec(< string >)** – not the **execute** stored procedure
- **connect**
- **begin transaction**
- **rollback**
- **commit**
- **dbcc**

set

The definitions for the following **set** options should now read:

- **set literal_autoparam on | off** – is **on** by default. If the server-level setting for **literal_autoparam** is **on**, this option enables and disables use of that feature. If the server level setting is **off**, this setting has no effect.

- **set statement_cache on | off** – is **on** by default. If the server-level setting for **statement_cache** is on, this option enables and disables use of that feature. If the server level setting is **off**, this setting has no effect.

transfer table

Add the following information to the Usage section of the **transfer table** reference page:

When **transfer table** runs into an error (such as duplicate key), only the underlying error number is reported, but not the error message. This makes it difficult for a user to understand why the statement was aborted. For example:

```
Msg 2633, Level 20, State 1
Server 'SYB155', Line 1
TRANSFER TABLE failed to insert a row to table 'my_tab'. The
indicated error was 2601.
Msg 16025, Level 16, State 1
Server 'SYB155', Line 1
TRANSFER TABLE my_tab: command failed with status 2633.
```

To retrieve the error message, manually query **master..sysmessages**. For example, if 2601 is your error number, enter:

```
select * from master..sysmessages where error = 2601
```

See the *Troubleshooting Guide* for more information about error 2601.

Adaptive Server Reference Manual: Procedures

Read about updates, corrections, and clarifications for Adaptive Server *Reference Manual: Procedures*.

sp_addobjectdef

The Usage section of the **sp_addobjectdef** reference page includes this bullet item, which is incomplete:

- Use **sp_addobjectdef** before issuing any create table or create existing table commands.

Replace with:

- Use **sp_addobjectdef** before issuing any **create table** or **create existing table** commands. However, if a remote table exists, you need not use **sp_addobjectdef** before executing **create proxy_table**.

sp_bindexclass

The description of this stored procedure reads: "Associates an execution class with a client application, login, or stored procedure."

Update the description, with "Associates an execution class with a client application, login, stored procedure, or default execution class."

Syntax updates

The definitions of the parameters have been updated:

- **object_name** – the description should now read, "is the name of the client application, login, or stored procedure to be associated with the execution class, **classname**. If **object_type** is **df**, it should be null."
- **object_type** – the description should now read, "identifies the type of **object_name**. Use **ap** for application, **lg** for login, **pr** for stored procedure, or **df** for user-defined default execution class."
- **scope** – the description should now read, "is the name of a client application or login, or it can be null for **ap**, **lg**, or **df** objects. It is the name of the stored procedure owner (user name) for objects. When the object with **object_name** interacts with the application or login, **classname** attributes apply for the **scope** you set."

Example

Add this to the Example section:

This statement assigns 'CLASS1' attributes to all tasks that are running with default execution attributes:

```
sp_bindexeclass NULL, 'DF', NULL, 'CLASS1'
```

Usage updates

Add this new bullet item:

- When binding an execution class to a default execution class, all tasks running with default execution attributes run with attributes of the new class.

In addition, these two bullet items have been modified, and should now read:

- **sp_bindexeclass** associates an execution class with a client application, login, or stored procedure. It can also associate an execution class to the default execution class. Use **sp_addexeclass** to create execution classes.
- Due to precedence and scoping rules, the execution class being bound may or may not have been in effect for the **object_name**. The object automatically binds itself to another execution class, depending on other binding specifications, precedence, and scoping rules. If no other binding is applicable, the object binds to the user-defined default execution class. If a user-defined default execution class is not specified, then the object binds to the system-defined default execution class **EC2**.

sp_cacheconfig

The Data Cache Memory subsection in the Usage section of the reference page for **sp_cacheconfig** includes a bullet item that states, in part:

- A data cache requires a small percentage of overhead for structures that manage the cache. All cache overhead is taken from the default data cache.

Since you can configure the default data cache size separately from total memory, the cache overhead is taken from free memory, not the default data cache. For this reason, the bullet should read:

- A data cache requires a small percentage of overhead for structures that manage the cache. **All cache overhead is taken from free memory.**

sp_downgrade

The reference page for **sp_downgrade** should specify that you can execute this stored procedure only in the **master** database.

sp_dropexeclass

A bullet item in the Usage section reads as follows:

- **classname** must not be bound to any client application, login, or stored procedure. Unbind the execution class first, using **sp_unbindexeclass**, then drop the execution class, using **sp_dropexeclass**.

Replace with:

- **classname** must not be bound to any client application, login, stored procedure, or default execution class. Unbind the execution class first, using **sp_unbindexeclass**, then drop the execution class, using **sp_dropexeclass**.

sp_lmconfig

There is a correction to the **sp_lmconfig** stored procedure. The syntax in the book appears as:

```
sp_lmconfig
[ 'edition' [, edition_type ]]
[ , 'license type' [, license_type_name ]]
[ , 'smtp host' [, smtp_host_name ]]
[ , 'smtp port' [, smtp_port_number ]]
[ , 'email sender' [, sender_email_address ]]
[ , 'email recipients' [, email_recipients ]]]
[ , 'email severity' [, email_severity ]]
```

Replace with:

```
sp_lmconfig
[
  [ 'edition' [, edition_type ]]
  [ 'license type' [, license_type_name ]]
  [ 'smtp host' [, smtp_host_name ]]
  [ 'smtp port' [, smtp_port_number ]]
  [ 'email sender' [, sender_email_address ]]
  [ 'email recipients' [, email_recipients ]]
  [ 'email severity' [, email_severity ]]
]
```

Documentation Changes

The `sp_lmconfig` procedure, when executed without parameters, also displays the server name from where the license is checked out. Example #1 in the reference page should display the following, with the addition of the column "Server Name."

```
1> sp_lmconfig
2> go

Parameter Name      Config Value
-----
edition             EE
license type        CP
smtp host            null
email recipients     null
email severity       null
smtp port            null
email sender         null

License Name        Version   Quantity Status   Expiry Date
Server Name
-----
ASE_HA              2010.03314 2      expirable Apr 1 2010 12:00AM
cuprum
ASE_ASM             null       0      not used  null
null
ASE_EJB             null       0      not used  null
null
ASE_EFTS            null       0      not used  null
null
ASE_DIRS            null       0      not used  null
null
ASE_XRAY            null       0      not used  null
null
ASE_ENCRYPTION      null       0      not used  null
null
ASE_CORE            2010.03314 2      expirable Apr 1 2010 12:00AM
cuprum
ASE_PARTITIONS      null       0      not used  null
null
ASE_RLAC            null       0      not used  null
null
ASE_MESSAGING_TIBJMS null       0      not used  null
null
ASE_MESSAGING_IBMMQ null       0      not used  null
null
ASE_MESSAGING_EASJMS null       0      not used  null
null

Property Name      Property Value
-----
PE                 EE
LT                 CP
ME                 null
MC                 null
MS                 null
```

```
MM          null
CP          0
AS          A
(return status = 0)
```

sp_locklogin

Examples 2 and 4 use **all**, which **sp_locklogin** does not recognize. A note in the Usage section correctly states:

Note: A value of NULL for a login means all logins."

Example 2 locks all logins except those with the `sa_role`. The erroneous syntax is:

```
sp_locklogin "all", "lock", sa_role
```

Replace with:

```
sp_locklogin NULL, "lock", sa_role
```

Example 4 locks all login accounts that have not authenticated within the past 60 days. The erroneous syntax is:

```
sp_locklogin 'all', 'lock', NULL, 60
```

Replace with:

```
sp_locklogin NULL, 'lock', NULL, 60
```

sp_reportstats

In the reference page for **sp_reportstats**, disregard the following bullet, which no longer applies:

- **sp_reportstats** does not report statistics for any process with a system user ID (*suid*) of 0 or 1. This includes deadlock detection, checkpoint, housekeeper, network, auditing, mirror handlers, and all users with **sa_role**.

sp_showcontrolinfo

The definitions of the parameters have been updated:

- **object_name** – the description should now read, "is the name of the application, login, stored procedure, or engine group. Do not specify an **object_name** if you specify **ps** or **DF** as the **object_type**. If you do not specify an **object_name** (or specify an **object_name** of null), **sp_showcontrolinfo** displays information about all object names."
- **object_type** – the description should now read, "is **ap** for application, **lg** for login, **pr** for stored procedure, **eg** for engine group, **ps** for process, or **df** for user-defined default execution class. If you do not specify an **object_type** or specify an **object_type** of null, **sp_showcontrolinfo** displays information about all types."

Usage updates

Add these two bullets:

Documentation Changes

- If **object_type** is **df**, **sp_showcontrolinfo** shows information about the user-defined default execution class **any**.
- If **object_type** is **df**, **object_name** and **spid** should be null.

sp_unbindexclass

The description of this stored procedure reads: "Removes the execution class attribute previously associated with a client application, login, or stored procedure for the specified scope."

Update the description with: "Removes the execution class attribute previously associated with a client application, login, stored procedure, or default execution class for the specified scope."

Syntax updates

The definitions of the parameters have been updated:

- **object_name** – the description should now read, "is the name of the application, login, or stored procedure for which to remove the association to the execution class. If **object_type** is **df**, **object_name** should be null."
- **object_type** – the description should now read, "identifies the type of **object_name** as **ap**, **lg**, **pr**, or **df** for application, login, stored procedure, or default execution class."
- **scope** – the description should now read, "is the application name or the login name for which the unbinding applies for an application or login. For stored procedures, scope is the user name of the stored procedure owner. **scope** is null for object type **df**."

Usage updates

Add this bullet item:

- When unbinding from user-defined default execution class, all tasks running with user-defined default execution class attributes run with attributes of the system-defined default execution class **EC2**.

The following bullet item appears in the Usage section:

- "Due to precedence and scoping rules, the execution class being unbound may or may not have been in effect for the object called **object_name**. The object automatically binds itself to another execution class, depending on other binding specifications and precedence and scoping rules. If no other binding is applicable, the object binds to the default execution class, **EC2**."

The text has been modified to read:

- "Due to precedence and scoping rules, the execution class being unbound may or may not have been in effect for the **object_name**. The object automatically binds itself to another execution class, depending on other binding specifications, and precedence and scoping rules. If no other binding is applicable, the object binds to the user-defined default

execution class. If there is no user-defined default execution class, the object binds to class **EC2**."

Adaptive Server *Reference Manual: Tables*

Read about updates, corrections, and clarifications for Adaptive Server *Reference Manual: Tables*.

monIOQueue

The description of the IOType column in monIOQueue is missing the System category. The text should read:

"Category for grouping I/O. One of UserData, UserLog, TempdbData, TempdbLog, or System."

monLogicalClusterAction

Add new states **canceled** and **releasing** as possible values for the State column of the monLogicalClusterAction table.

sysobjects

These changes apply to the reference page for **sysobjects**:

- Remove the object type "L - log" under the type column. Remove this reference; it is not a valid object type.
- In Table 1-18 replace **O_LOG** with **O_PREDICATE**; **O_LOG** is not a valid object type.

sysprotects

Add action code 368 to the sysprotects table for permissions on the **transfer table** command. When a user issues **grant all** on a table, action code 368 is added to the sysprotects table.

sysprocesses

The table listing the column names for sysprocesses includes the column "clientport," which does not exist in Adaptive Server 15.5.

Adaptive Server *System Administration Guide*

Read about updates, corrections, and clarifications for Adaptive Server *System Administration Guide*.

Setting cluster redundancy level

Add this text to "Cluster redundancy level" section of Chapter 5, "Setting Configuration Parameters":

For the cluster to start, the value of the **cluster redundancy level** configuration parameter must be at least one less than the value of **maximum number of instances** as specified in `cluster.cfg` or the quorum file. Thus, the cluster cannot start if you set either:

- **maximum number of instances** to a value that is equal to or less than the value of **cluster redundancy level**.
- **cluster redundancy level** to a value equal to or greater than the value of **maximum number of instances**.

Adaptive Server *Utility Guide*

Read about updates, corrections, and clarifications for Adaptive Server *Utility Guide*.

bcp

Support for initialization strings

The **bcp** utility supports sending Transact-SQL commands, such as **set replication off**, to Adaptive Server before data is transferred.

Although you can use any Transact-SQL™ command as an initialization string for **bcp**, you must reset possible permanent changes to the server configuration after running **bcp**. You can, for example, reset changes in a separate **isql** session.

Syntax changes

The following **bcp** parameter supports Transact-SQL initialization commands:

```
-- initstring 'Transact-SQL command'
```

Result sets issued by the initialization string are silently ignored, unless an error occurs.

Example

In this example, replication is disabled when `titles.txt` data is transferred into the `pubs2..titles` table:

```
bcp pubs2..titles in titles.txt --initstring 'set replication off'
```

In this example, **set replication off** is limited to the current session in Adaptive Server; you need not explicitly reset the configuration option after **bcp** is finished.

Note: If Adaptive Server returns an error; **bcp** stops the data transfer and displays an error message.

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.

1. Point your Web browser to <http://www.sybase.com/support>.
2. From the menu bar or the slide-out menu, under **Support**, choose **EBFs/Maintenance**.
3. If prompted, enter your MySybase user name and password.
4. (Optional) Select a filter from the **Display** drop-down list, select a time frame, and click **Go**.
5. 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

Obtaining Help and Additional Information

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.

6. Click the **Info** icon to display the EBF/Maintenance report, or click the product description to download the software.

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.

Obtaining Help and Additional Information