



Release Bulletin

Adaptive Server[®] Enterprise

15.5

IBM AIX

DOCUMENT ID: DC72410-01-1550-07

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
Adaptive Server Support for ASE Replicator	5
Password Compatibility with Replication Server	5
Replication Server Support for In-Memory Databases	5
Changes that Affect Existing Applications	10
Backup Server and directio	10
Open Client and Open Server	10
Changed Functionality	10
Change to job scheduler tasks Configuration Parameter	10
Change to directio and dsync Settings for Raw Devices	11
Support for Database Resynchronization in Replication Server	11
Support for db2 Server Class	11
Installation and Upgrade	13
Special Installation Instructions	13
Special Licensing Instructions	13
SySAM License Server May Fail to Find the Virtual Machine	13
Change to License Quantity Used by Per-CPU and Per-chip License Types	13
SySAM Licensing Checkout	14
SySAM Fails to Check Out License for Adaptive Server Small Business Edition	14
Sub-capacity licensing	15

Special Configuration Instructions	16
Number of Preallocated Extents	16
Special Upgrade and Downgrade Instructions	17
Upgrading Job Scheduler	17
Upgrading Adaptive Server	17
Migrating from Adaptive Server 12.5.1 to 15.5	19
Loading Database Dumps During Downgrade	19
Known Installation Issues for Adaptive Server	20
Known Issues for Adaptive Server	23
Documentation Changes	29
Adaptive Server Reference Manual: Building Blocks ...	29
Adaptive Server Reference Manual: Commands	29
Adaptive Server Reference Manual: Procedures	32
Adaptive Server Reference Manual: Tables	37
Adaptive Server System Administration Guide	38
Adaptive Server Utility Guide	38
Obtaining Help and Additional Information	41
Technical Support	41
Downloading Sybase EBFs and Maintenance Reports	41
Sybase Product and Component Certifications	42
Creating a MySybase Profile	42
Accessibility Features	42

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.

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, 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 Adaptive Server 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
Replication Server Manager plugin	15.2, 15.5	<ul style="list-style-type: none"> AIX 32-bit HP-UX IA 32-bit Linux x64 32-bit Linux on IBM PSeries 64-bit 	<ul style="list-style-type: none"> Solaris Sparc 32-bit Solaris 10 x86 Solaris 10 x64 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.

Adaptive Server Support for ASE Replicator

Beginning in version 15.5, Adaptive Server no longer supports ASE Replicator.

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.

Replication Server Support for In-Memory Databases

Learn about how Replication Server works with in-memory databases.

Replication Server version 15.2 supports as the replicate database only:

Product Summary

- In-memory databases
- Relaxed-durability databases set with durability at **no_recovery**

Note: Replicating to an in-memory database may not be faster than replicating to a relaxed-durability database. DML on in-memory databases depends on several factors; see the *In-Memory Databases Users Guide* for more information.

The primary database must be a traditional, full-durability, disk-resident database. For convenience, this document refers to relaxed-durability databases with durability set to **no_recovery** as "relaxed-durability databases."

Note: Replication performance with an in-memory or relaxed-durability database acting as the replicate database is commensurate with replication between traditional full-durability disk-resident databases acting as both the primary and replicate database.

You can initialize either an in-memory database, or relaxed-durability database as a new replicate database by obtaining data, object schema, and configuration information from one of:

- A template database that retains basic information.
- A database dump from another database. Load the dump to the target in-memory database or relaxed-durability database.

The dump source database can be another in-memory database, relaxed-durability database, or a traditional disk-resident database.

In-memory and relaxed-durability databases lose their object definition, data, and RepAgent configuration once the host data server shuts down or restarts. You must reinitialize the database from the template or database dump from a source database.

Using a template database to initialize in-memory and relaxed-durability databases

You can use the same disk-resident database as a template for multiple in-memory or relaxed-durability databases.

To set up replication using a template database:

1. Create the template database – the template database uses the name of the database with the outbound connection to Replication Server, which is usually the replicate database name.

```
create database tokyo_db on publicdev=10 log on publicdevlog=10
go
```

2. On Replication Server, use **rs_init** to configure the replicate databases.
3. Suspend the DSI thread to the template database by stopping RepAgent on the template database. For example:

```
suspend connection to TOKYO_DS.ttokyo_db
```

4. Rename the template database to **templatel**:

```
use master
go
sp_dboption tokyo_db, single, true
```

```
go
sp_renamedb tokyo_db, templatel
go
sp_dboption templatel, single, false
go
```

5. Create the in-memory or relaxed-durability database with durability set to **no_recovery** using the template created in step 1:

```
create inmemory database tokyo_db
use templatel as template
on imdb_cache_dev = '50' log on
imdb_cache_dev_log='50'
with DURABILITY=NO_RECOVERY
go
```

6. Connect to Replication Server and resume the connection to the replicate database:

```
resume connection to TOKYO_DS.tokyo_db
```

After a shutdown or termination and subsequent restart of Adaptive Server, Adaptive Server automatically re-creates the in-memory or relaxed-durability database from the template, wherever this was specified. When Replication Server resumes the connection to the replicate in-memory or relaxed-durability database that you created, Replication Server may reapply commands, as the information used by Replication Server to detect the last command applied is lost when you restart Adaptive Server.

Using a dump from another database to initialize in-memory and relaxed-durability databases

If you use a dump from another database to initialize the in-memory or relaxed-durability database as a replicate database, the replicate database inherits the configuration of the database from which you obtained the dump.

To set up replication using a database dump:

1. Create the in-memory or relaxed-durability database:

```
create inmemory database tokyo_db
on imdb_cache_dev2 = '50' log on
imdb_cache_dev_log2='50'
with DURABILITY=NO_RECOVERY
go
```

2. Create the objects such as tables and stored procedures, users, and permissions needed to receive replicated data; alternatively, you can load a database dump.
3. Use **rs_init** to create the Replication Server connection to the database.
4. Perform a dump to save the current state of the database:

1. Suspend the connection to the replicate in-memory or relaxed-durability database:

```
suspend connection to RDS.imdb1
go
```

2. Obtain a database dump of the database:

```
dump database imdb1 to '/databases/dump/tokyo_db.dump'
go
```

3. Resume the connection to the replicate database:

```
resume connection to RDS.imdb1
go
```

Re-creating relaxed-durability databases after an Adaptive Server shutdown

In-memory databases are re-created when you restart the server; you must therefore restore the data after each server restart:

1. Repopulate the re-created replicate in-memory or relaxed-durability database with a new dump from the dump source database or from any archived dumps taken on the replicate database.

Note: If the dump is not loaded from the dump of the source, there will be missing rows in the replicate tables.

For example, to load the `tokyo_db` database from the original `tokyo_db.source` dump when the host Adaptive Server restarts:

```
use master
go
sp_dboption tokyo_db, single, true
go
load database tokyo_db from
'/databases/dump/tokyo_db.dump'
go
online database tokyo_db
go
sp_dboption tokyo_db, single, false
go
```

2. Resume the connection to the re-created replicate in-memory or relaxed-durability database.

Enabling autocorrection

Enable autocorrection for the replication definitions used for subscriptions to the replicate in-memory or relaxed-durability database, to convert any update or insert into a delete and insert pair. You must set autocorrection before resuming a connection; this applies whether you used templates, or dump and load to create the replicate database. Enabling autocorrection allows Replication Server to continue replicating messages in Replication Server queues even if the Adaptive Server hosting a replicate in-memory or relaxed-durability database shuts down or terminates.

Minimal DML logging and replication

To optimize the log records that are flushed to the transaction log on disk, Adaptive Server can perform minimal to no logging when executing some data manipulation language (DML) commands—**insert**, **update**, **delete**, and slow **bcp**—on all types of low-durability databases, such as in-memory databases and relaxed-durability databases set with durability of **at_shutdown** or **no_recovery**. You can perform minimal logging for DMLs on a per-database, per-table, and session-specific basis. See Chapter 3, "Minimally-logged DML" in the *Adaptive Server Enterprise In-Memory Database Users Guide*.

Note: Minimal DML logging session-level settings take precedence over database-level settings and table-level settings.

Replication Server does not support a replication environment with in-memory databases, or where database durability set to **no_recovery** is acting as the primary database.

As replication uses full logging, replication and the minimal data manipulation language (DML) logging feature in Adaptive Server 15.5 are incompatible at the same level, such as the database level, or the table level. However, you can take advantage of the performance enhancements from minimal logging on some tables while allowing replication on others, as minimal DML logging and replication can coexist at different levels.

For example, if you set replication and minimal DML logging at the same level, such as table level, setting replication status fails and an error message appears, as described in these scenarios:

- If you create a database to use minimal DML logging:
 - And if you mark the database for replication using **sp_reptostandby**, the attempt fails:

```
Cannot set replication for database database_name
as it is minimally logged. Use ALTER DATABASE to
set full DML logging and try again.
```
 - You can mark tables and stored procedures for replication to replicate a subset of tables. When you mark a table in a database using minimal DML logging, you see:

```
Warning: database_name is using minimal logging.
Replicated objects will continue to use full DML logging.
```
- If a database is using full logging and you mark it for replication using **sp_reptostandby**, then try to alter the database to set minimal DML logging, the attempt fails:

```
Cannot alter database database_name to set minimal
logging because this database is marked for
replication. Remove replication and try again.
```
- If a database is using full logging and has objects marked for replication, you can set minimal DML logging at the database level, but you see:

```
Warning: Database database_name has objects marked for
replication. Replicated objects will continue to use full logging.
```
- If you create a table using minimal DML logging, and then mark the table for replication, the attempt fails:

```
Cannot set replication for table table_name because it is
using minimal logging. Use ALTER TABLE to set full logging and try
again.
```
- A database using full DML logging can contain tables defined to use minimal logging. If you mark the database for replication, you see:

```
Warning: Database database_name has tables that use
minimal DML logging. These tables will not be replicated.
```
- If you create a table using full logging and mark the table for replication, then set minimal DML logging for the table, the attempt fails:

```
Cannot alter the table table_name to set minimal DML
logging because this table is marked for
replication. Remove replication and try again.
```

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.

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.

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.

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.

Special Licensing Instructions

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

SySAM License Server May Fail to Find the Virtual Machine

When the **sysamcap** SySAM utility cannot determine the Virtual Machine, SySAM needs to be told to check for the Virtual Machine by setting these environment variables:

Workaround: Set the VMware environment variable to:
SYBASE_SAM_CPUINFO=VMCHECKVM.

Set the Xen environment to either variable:

- SYBASE_SAM_CPUINFO=Xen, or
- SYBASE_SAM_CPUINFO=XENCHECKXEN

Change to License Quantity Used by Per-CPU and Per-chip License Types

In version 15.5, Adaptive Server checks 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 Adaptive Server, in which if Adaptive Server was licensed per CPU or CPU chip, the license quantity requested was reduced if the **max online engines** configuration parameter was fewer than the number of CPUs on the machine.

SySAM Licensing Checkout

If Adaptive Server was licensed using a per-processor license type in previous releases, the license quantity was determined at startup.

In this release, Adaptive Server periodically checks the number of processors it can use and attempts to check out additional licenses for any increase. Adaptive Server shuts down after 30 days if additional licenses are not made available within this period. See the *SySAM Users Guide* for details regarding license grace.

SySAM Fails to Check Out License for Adaptive Server Small Business Edition

[CRs #628596, #628403] Under some circumstances with the Small Business Edition of Adaptive Server, you cannot install Adaptive Server, or it does not start in a virtual environment.

When this occurs, Adaptive Server reports an error similar to this:

```
Cannot determine the total number of processor
cores on the machine from within this virtual system.
The ASE_CORE ("CH") license type is restricted for use on
a machine with a maximum of 2 processor chips regardless
of whether the virtual system is limited to use fewer
processor chips. You need to run 'sysamcap' on the physical
machine (or virtual-machine/partition control domain) and
set the reported environment variable prior to running
Installer.
```

This error is raised when:

- Adaptive Server is running on a virtual system that does not have access to the processor counts of the physical machine (this applies to all virtual environments except AIX LPAR and HP-UX vPar), or
- Adaptive Server is using a license type that is restricted for use to a machine with no more than a specific number of processors; for example, you cannot use Adaptive Server Enterprise Small Business Edition on a machine with more than two processor chips.

Workaround:

1. Run the **sysamcap default** command on the physical machine or control domain for the virtualization technology (e.g., the VMware ESX console or XEN dom0).
2. Prior to running the product or product installer, set the environment variable reported.

See the SySAM Sub-capacity licensing section of the *SySAM Users Guide* for more information.

Sub-capacity licensing

Learn about sub-capacity licensing for this version of Adaptive Server.

Sybase now offers sub-capacity licensing options for Sybase IQ Enterprise Edition and Adaptive Server Enterprise Enterprise Edition. Sub-capacity licensing refers to licensing a Sybase product on a subset of the CPUs available on a physical machine.

Platform support

Table 2. SySAM virtualization sub-capacity compatibility

Vendor	Product	Platform support	Virtualization type
HP	nPar	HP IA 11.31	Physical partition
	vPar		Virtual partition
	Integrity Virtual Machines with Resource Manager		Virtual machine
	Secure Resource Partitions		OS containers
IBM	LPAR	AIX 6.1	Virtual partition
	dLPAR		Virtual partition
Sun	Dynamic System Domains	Solaris 10	Physical partition
	Solaris Containers/Zones with Solaris Resource Manager		OS partition
Intel, AMD	VMWare ESX Server* Guest OS: Windows	VMWare ESX 3.5, Guest OS: Windows 2003, RH 5.3, SuSe 10	Virtual machine
	VMWare ESX Server Guest OS: Linux	VMWare ESX 3.5, Guest OS: Windows 2003, RH 5.3, SuSe 10	Virtual machine
	Xen,** DomainU: Windows	Windows 2003	Virtual machine

Vendor	Product	Platform support	Virtualization type
	Xen, DomainU: Linux	RH 5.3, SuSe 10	Virtual machine

* VMWare ESX Server excludes Sun Solaris x64, VMWare Workstation, and VMWare Server.

** Xen excludes Sun Solaris x64.

Enabling Sybase sub-capacity licensing

You must have a sub-capacity licensing agreement with Sybase before you can enable sub-capacity licensing. As with other Sybase licenses, you must generate a license key. For specific steps, see the *SySAM Quick Start Guide*.

Note: Keep your license server up-to-date.

Although the installation media include a copy of the most recent SySAM license server, Sybase recommends that you periodically check for license server updates on the Maintenance Releases and Bug Fixes (EBFs) link at <http://www.sybase.com/support>. See *Obtaining Help and Additional Information* on page 41.

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 3. 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.

Current version	Upgrading to	Special Upgrade Information
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 4. Upgrading a single database

Current version	Upgrading to	Special Upgrade Instructions
Adaptive Server 15.0.x and earlier	Adaptive Server 15.5.x	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. See the section "Reenabling Replication Server After the Dataserver Upgrade" in the "Upgrade" chapter in the Adaptive Server installation guide for your platform.
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:xxxxx')"
```

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

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

CR #	Description
625837	<p>Cannot share machine-level license by different Linux PowerPC operating systems on IBM logical partitions.</p> <p>The machine ID for Linux PowerPCs on IBM logical partitions is the virtual machine's MAC address, which means that SySAM cannot identify a unique machine ID for Linux on IBM logical partitions.</p> <p>Workaround: None. You cannot share a machine license among different operating systems.</p>
625227	<p>SPDC-generated served partition-level licenses may not work.</p> <p>When you use a served partition-level license generated from the SPDC to start a SySAM server, you see a message similar to:</p> <pre>(SYBASE) Invalid license key (inconsistent authentication code) ... (SYBASE) License server system started on hpiabou (SYBASE) No features to serve, exiting (SYBASE) EXITING DUE TO SIGNAL 49 Exit reason 4 (lmgrd) SYBASE exited with status 49 (No features to serve) (lmgrd) SYBASE daemon found no features. Please correct (lmgrd) license file and re-start daemons. (lmgrd) (lmgrd) This may be due to the fact that you are using (lmgrd) a different license file from the one you expect. (lmgrd) Check to make sure that: (lmgrd) /remote/cat_fc/nli/iq152.hpia/SYSAM-2_0/licenses/ 18965_hpiabou_ (lmgrd) is the license file you want to use.</pre> <p>Workaround: Use the host IDs from the previous version for the license running on the license server.</p>

CR #	Description
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>
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>
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>

Known Issues for Adaptive Server

Learn about known issues and apply workarounds for Adaptive Server. 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
623788	<p>Adaptive Server cannot start Job Scheduler on IPV6-enabled AIX servers for IPV4 addresses.</p> <p>When you install Job Scheduler on an IPV6-enabled AIX system using an IPV4 address that starts with the first set of numbers smaller than 128, the installation may fail, and you may see a message similar to:</p> <pre>Job Scheduler Task failed to spawn an Agent.</pre> <p>When this occurs, the Job Scheduler log shows:</p> <pre>Socket bind failed. Errno:22</pre> <p>Workaround: Disable IPV6 on the system.</p>

CR #	Description
616224	<p>Setting the default sort order to nonbinary, single-file sort order may cause like pattern evaluation failures.</p> <p>If you configure Adaptive Server to use the nonbinary, single-byte sort order as its default sort order, you may encounter an issue where a query containing a like expression using a constant literal as the argument (with or without wildcards) does not correctly evaluate the value in the table is the same as the like argument prefix up to the wildcard character. For example:</p> <pre data-bbox="323 439 783 565">create table t1 (c1 varchar(20)) insert t1 values ('John') insert t1 values ('Johnson') insert t1 values ('Johnston')</pre> <p>Both of these queries fail to return the first record:</p> <pre data-bbox="323 626 868 673">select * from t1 where c1 like 'John' select * from t1 where c1 like 'John%'</pre> <p>This query, however, works correctly:</p> <pre data-bbox="323 734 854 756">select * from t1 where c1 like 'Joh%'</pre> <p>A query using variables or procedural parameters is not affected by this incorrect behavior, and the following returns an expected and correct result set:</p> <pre data-bbox="323 847 830 921">declare @var varchar(20) select @var = 'John' select * from t1 where c1 like @var</pre> <p>Workarounds:</p> <ol data-bbox="323 994 1190 1177" style="list-style-type: none"> 1. Enable the statement cache. See "Setting the statement cache" in Chapter 3, "Configuring Memory," of <i>System Administration Guide: Volume 2</i>. 2. Activate the autoperparameterization of literals: <pre data-bbox="360 1090 982 1112">-sp_configure 'enable literal autoparam', 1</pre> 3. Restart Adaptive Server using trace flag -T7739. This step parameterizes like constant literals. <hr/> <p>Note: Queries that are not placed in the statement cache, such as select into, may not benefit from this workaround.</p> <hr/> <p>For more information about CR #616224, contact your local Sybase Technical Support Center using the information from http://www.sybase.com/contactus/support#tech.</p>
607917	<p>Running transfer table may report SySAM errors.</p> <p>When you invoke the transfer table command, this error may be written to the error log:</p> <pre data-bbox="323 1437 989 1484">SySAM: Failed to obtain license(s) for ASE_DTU feature from license file(s) or servers(s).</pre> <p>Workaround: This is an erroneous report. As long as you have a valid license for the Adaptive Server in-memory database feature, transfer table executes successfully.</p>

CR #	Description
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> <pre>The database '<template_dbname>' is currently unavailable. It is being used to create a new database.</pre> <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>
607095	<p>If you restart Adaptive Server after using <code>sp_renamedb</code> to rename a database that is the template database for an in-memory database, Adaptive Server may create the in-memory database using the model database instead of the template database, or the in-memory database may become inaccessible.</p> <p>Workaround: After you run sp_renamedb, run sp_configure so that Adaptive Server can recognize the new name of the template database.</p> <p>The default command is:</p> <pre>sp_renamedb "tdb", "tdb_old"</pre> <p>where "tdb" is a template database, and "tdb_old" is the name you are renaming it to. This assumes that you have an in-memory database called "timb" that has "tdb" as its template database.</p> <p>After you run sp_renamedb, run:</p> <pre>sp_configure "allow updates", 1 go update master.dbo.sysdatabases set def_remote_loc='tdb_old' where def_remote_loc='tdb' go sp_configure "allow updates", 0 go</pre>

Known Issues for Adaptive Server

CR #	Description
606234	<p>If you enable ASE_JAVA for a commercial off-the-shelf Java Virtual Machine, then start Adaptive Server version 15.5 in a multiengine configuration, Adaptive Server may perform an incorrect setup of the JVM module responsible for physical file and network I/O protection.</p> <p>This incorrect setup allows Java methods executed from Adaptive Server to perform permission-sensitive physical file I/O operations, as well as launch network services while taking on the full identity of the user who restarted the Adaptive Server.</p> <p>Workaround:</p> <ul style="list-style-type: none"> • Do not enable the PCI subsystem. If you do, you cannot execute Java methods from Adaptive Server. • Restart a single-engine Adaptive Server first, then bring additional engines online before any requests are made to execute Java methods. • Configure the PCI bridge so that it defaults towards binding the commercial off-the-shelf JVM under engine 0.
595986	<p>Differing byte orders may generate an error during database dump.</p> <p>When you load a dump of a database that has a durability of no_recovery or at_shutdown, you may see:</p> <pre>Adaptive Server cannot load this database because the database that was dumped was not quiescent when the dump was performed.</pre> <p>This occurs when you load a dump of a database from one platform to another platform with a different byte order.</p> <p>Workaround: None.</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.
589269	<p>Adaptive Server may stop responding during start-up.</p> <p>If you are using the in-memory database or relaxed-durability database feature, Adaptive Server may run into a timeslice error during start-up if it cannot access the license server as it starts.</p> <p>Workaround: Make sure Adaptive Server can access the license server.</p>

CR #	Description
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>
567466	<p>When your statement cache is on, if an exception occurs while compiling a statement that includes temporary tables, error message 8203 may appear.</p> <p>Workaround: Use command-line traceflag 467, which prevents statements referencing temporary tables from being cached.</p>
550035	<p>Configuring Job Scheduler fails on IBM AIX</p> <p>srvbuild and srvbuildres fail to configure Job Scheduler.</p> <p>Workaround: Add <code>\$\$SYBASE_JRE6/bin/classic</code> to the LIBPATH environment variable before running srvbuild or srvbuildres.</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>

Known Issues for Adaptive Server

CR #	Description
439404	<p data-bbox="323 215 1137 241">Creating Job Scheduler schedules for Adaptive Server with Simplified Chinese.</p> <p data-bbox="323 258 1184 340">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 data-bbox="323 357 1184 409">Workaround: Localize the Adaptive Server for Simplified Chinese before configuring Job Scheduler.</p>
400784	<p data-bbox="323 440 1170 493">Starting Monitor Client GUI fails when the Lightweight Directory Access Protocol (LDAP) server is used.</p> <p data-bbox="323 510 1141 562">When you start the Monitor Client GUI in an environment using LDAP instead of an <i>interfaces</i> file, the Monitor Client GUI cannot connect to any servers.</p> <p data-bbox="323 579 1184 631">Workaround: Create an <i>interfaces</i> file with entries for the servers that are monitored and use this with the Monitor Client GUI.</p>
382874	<p data-bbox="323 666 1069 692">Creating RSSD database fails with earlier versions of Replication Server.</p> <p data-bbox="323 710 1170 762">rs_init in Replication Server version 12.6 and earlier cannot create an RSSD database in Adaptive Server Enterprise version 15.0 or later.</p> <p data-bbox="323 779 1184 861">Workaround: If you do not use an Embedded RSSD (ERSSD), manually create the RSSD database before using rs_init to create a new Replication Server. See the <i>Replication Server Administration Guide</i>.</p>

Documentation Changes

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

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_bindexclass 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_bindexclass** 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_addexclass** 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_dropexclass

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_unbindexclass**, then drop the execution class, using **sp_dropexclass**.

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_unbindexclass**, then drop the execution class, using **sp_dropexclass**.

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 ]]
    ]
```

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:

- 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 an 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**:

Documentation Changes

- 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