

Release Bulletin Adaptive Server[®] Enterprise Version 15.0 for Linux Opteron

Document ID: DC00498-01-1500-01

Last revised: June 1, 2006

Topic	Page
1. Accessing current release bulletin information	3
2. Product summary	4
2.1 Installation kit	4
2.2 Operating system requirements	4
3. Special installation instructions	5
3.1 Using SySAM 2.0	6
3.2 Runtime libraries required for Linux on Power installation	6
3.3 Monitoring and Diagnostic (MDA) tables	7
4. Special upgrade instructions	8
4.1 Upgrading compiled objects	8
5. Changes that affect existing applications	8
5.1 Directory changes	8
5.2 Changes to ODBC and OLE DB Driver support	9
5.3 Changes that affect application behavior	10
5.4 Changes that affect database administration operations	13
5.5 Open Client and Open Server changes	17
6. Known problems	18
6.1 Highlighted known installation issues	18
6.2 Highlighted known XML problems	19
6.3 Known alter table problems	20

Copyright 1987-2006 by Sybase, Inc. All rights reserved. Sybase, SYBASE (logo), ADA Workbench, Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise, Adaptive Server Enterprise Monitor, Adaptive Server Enterprise Replication, Adaptive Server Everywhere, Advantage Database Server, Afaria, Answers Anywhere, Applied Meta, Applied Metacomputing, App-Modeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-Translator, ASEP, Avaki, Avaki (Arrow Design), Avaki Data Grid, AvantiGo, Backup Server, Bay Cam, Beyond Connected, Bit-Wise, BizTracker, Certified PowerBuilder Developer, Certified SYBASE Professional, Certified SYBASE Professional Logo, ClearConnect, Client-Library, Client Services, CodeBank, Column Design, ComponentPack, Connection Manager, Convo/DM, Copernicus, CSP, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DataWindow .NET, DB-Library, dbQueue, DeJima, DeJima Direct, Developers Workbench, DirectConnect Anywhere, DirectConnect, Distribution Director, Dynamic Mobility Model, e-ADK, E-Anywhere, e-Biz Integrator, E-Whatever, EC Gateway, ECMAP, ECRTP, eFulfillment Accelerator, EII Plus, Electronic Case Management, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise Portal (logo), Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeller, eProcurement Accelerator, eremote, Everything Works Better When Everything Works Together, EWA, ExtendAssist, Extended Systems, ExtendView, Financial Fusion, Financial Fusion (and design), Financial Fusion Server, Formula One, Fusion Powered e-Finance, Fusion Powered Financial Destinations, Fusion Powered STP Gateway Manager, GeoPoint, GlobalFX, iAnywhere, iAnywhere Solutions, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InstaHelp, Intelligent Self-Care, InternetBuilder, iremote, iRLite, iScript, Jaguar CTS, jConnect for JDBC, KnowledgeBase, Legion, Logical Memory Manager, M2M Anywhere, Mach Desktop, Mail Anywhere Studio, Mainframe Connect, Maintenance Express, Manage Anywhere Studio, MAP, M-Business Anywhere, M-Business Channel, M-Business Network, M-Business Suite, MDI Access Server, MDI Database Gateway, media.splash, Message Anywhere Server, MetaWorks, MethodSet, mFolio, Mirror Activator, ML Query, MobiCATS, MobileQ, MySupport, Net-Gateway, Net-Library, New Era of Networks, Next Generation Learning, Next Generation Learning Studio, O DEVICE, OASIS, OASIS logo, ObjectConnect, ObjectCycle, OmniConnect, OmniQ, OmniSQL Access Module, OmniSQL Toolkit, OneBridge, Open Biz, Open Business Interchange, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Server, Open ServerConnect, Open Solutions, Optima++, Partnerships that Work, PB-Gen, PC APT Execute, PC DB-Net, PC Net Library, Pharma Anywhere, PhysicalArchitect, Pocket PowerBuilder, PocketBuilder, Power++, Power Through Knowledge, power.stop, SybaseAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, Powering the New Economy, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Pylon, Pylon Anywhere, Pylon Application Server, Pylon Conduit, Pylon PIM Server, Pylon Pro, QAnywhere, Rapport, Relational Beans, RemoteWare, RepConnector, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RFID Anywhere, RW-Display-Lib, RW-Library, SAFE, SAFE-PRO, Sales Anywhere, Search Anywhere, SDF, Search Anywhere, Secure SQL Server, Secure SQL Toolkit, Security Guardian, ShareSpool, ShareLinks, SKILLS, smart.partners, smart.parts, smart.script, SOA Anywhere, Trademark SQL Advantage, SQL Anywhere, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Edit/TPU, SQL Everywhere, SQL Modeller, SQL Remote, SQL Server, SQL Server Manager, SQL SMART, SQL Toolkit, SQL Server/CFT, SQL Server/DBM, SQL Server SNMP SubAgent, SQL Station, SQLI, Stage III Engineering, Startup.Com, STEP, SupportNow, S.W.I.F.T. Message Format Libraries, Sybase Central, Sybase Client/Server Interfaces, Sybase Development Framework, Sybase Financial Server, Sybase Gateways, Sybase IQ, Sybase Learning Connection, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Group, Sybase Synergy Program, Sybase Virtual Server Architecture, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SybFlex, SybMD, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, The Enterprise Client/Server Company, The Extensible Software Platform, The Future Is Wide Open, The Learning Connection, The Model For Client/Server Solutions, The Online Information Center, The Power of One, TotalFix, TradeForce, Transact-SQL, Translation Toolkit, Turning Imagination Into Reality, UltraLite, UltraLite.NET, UNIBOM, Unilink, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, Viafone, Viewer, VisualWriter, VQL, WarehouseArchitect, Warehouse Control Center, Warehouse Studio, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web Deployment Kit, Web.PB, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, XA-Server, XcelleNet, XP Server, XTNDAccess and XTNDConnect are trademarks of Sybase, Inc. or its subsidiaries. 01/06

Topic	Page
6.4 Running on RH3.0 64-bit distribution occasionally stops responding	20
6.5 Interactive SQL and Adaptive Server Enterprise plug-in issues	20
6.6 Statistics in system tables	23
6.7 Running checkstorage on an upgraded master device may report faults with the syscharsets catalog	24
6.8 Inserting into a table with identity column using select distinct	24
6.9 SySAM licenses needed to upgrade Adaptive Server Enterprise	24
6.10 Cross-platform dumps and loads	25
6.11 Running dbcc checktable or dbcc checkdb before bringing a database online	25
6.12 allow backward scan does not work	25
6.13 Modifying dbccdb schema for large identifiers	25
6.14 Drop definition timetable corruption	26
6.15 Manual upgrade with named cache fails	26
6.16 Running diskrefit to recover a corrupt master device can report errors	26
6.17 Using an equijoin clause between two different length columns	26
6.18 set statistics io does not display I/O generated by worker processes	27
6.19 IPv6 support causes Adaptive Server Enterprise to stop responding	27
6.20 Replication Server compatibility issues	28
6.21 Dumping or loading databases with asynchronous I/O	28
6.22 cis connect timeout and enable SNMP are not implemented	28
6.23 Handling multibyte character sets during migration	28
6.24 Network package reallocation during reconfiguration	28
6.25 bcp in or bcp out	29
6.26 Silent installation	29
6.27 Sybase Central	29
6.28 File names	29
6.29 Command line flag -k	30
6.30 table scan and index scan	30
6.31 Determining devices for a database	30

Topic	Page
7. Product compatibilities	30
7.1 Adaptive Server interoperability	32
7.2 Adaptive Server Replicator 15.0 support of Adaptive Server Enterprise 12.5 datatypes	35
7.3 Known compatibility issues	35
8. Documentation updates and clarifications	36
8.1 Reference Manual: Building Blocks	36
8.2 Reference Manual: Commands	38
8.3 Reference Manual: Procedures	39
8.4 Installation guide	45
9. Technical support	45
10. Other sources of information	45
10.1 Sybase certifications on the Web	46
10.2 Sybase EBFs and software maintenance	47
11. Accessibility features	47

1. Accessing current release bulletin information

A more recent version of this release bulletin may be available on the Web. To check for critical product or document information added after the release of the product CD, use the Sybase™ Technical Library Product Manuals Web site.

❖ Accessing release bulletins at the Technical Library Product Manuals Web site

- 1 Go to Product Manuals at <http://www.sybase.com/support/manuals/>.
- 2 Follow the links to the appropriate Sybase product.
- 3 Select the Release Bulletins link.
- 4 Select the Sybase product version from the Release Bulletins list.
- 5 From the list of individual documents, select the link to the release bulletin for your platform. You can either download the PDF version or browse the document online.

2. Product summary

Enclosed is Sybase Adaptive Server® Enterprise version 15.0. Server and client components are distributed on separate CDs.

For details on system requirements, including disk space and RAM, see the installation guide for your platform.

2.1 Installation kit

The installation kit includes:

- The Server CD – for contents list, see the installation guide for your platform.
- The PC-Client CD – contains software client components to be installed on Windows 2000, Windows XP Pro, and Windows 2003 computers.
- The Getting Started CD with the following documentation, specific to your platform:
 - Installation guide
 - Configuration guide
 - Release bulletin (this document)

2.2 Operating system requirements

Adaptive Server Enterprise version 15.0 has the following minimum operating system requirements:

For Linux Opteron:

- RHEL release 3 update 3:
 - kernel 2.4.2.1-20.EL
 - kernel-hugemem 2.4.21-20.EL (for Large Memory Support, LMS)
 - glibc-2.3.2.95-27
 - compat-libstdc++-7.3-2.96.128
- SLES 9 – patch level 2

- For details on different Linux distributions supported and operation system levels required, see the Red Hat Web site at <http://certification.sybase.com/ucr/search.do>.

Note The latest RH3 update that includes the required patch is RH3.0 Update 6 after the RH kernel 2.4.21-27.6.EL. The patch must be used, at the minimum, on a 64-bit kernel level. See “Running on RH3.0 64-bit distribution occasionally stops responding” on page 20 for more detail.

3. Special installation instructions

Adaptive Server Enterprise 15.0 includes new versions of many supporting components. The most current version of other Sybase products (such as Enterprise Connect™ Data Access 12.6) contain earlier versions of those same components. Installing Adaptive Server version 15.0 into the same directory with existing products should not impact the existing products.

Warning! Installing other products on top of Adaptive Server version 15.0 may result in one or more products not working correctly.

Sybase strongly recommends that you install Adaptive Server version 15.0 into its own directory. Where this is not practical, install Adaptive Server version 15.0 **last**.

See “Highlighted known installation issues” on page 18 for known installation problems.

Warning! Sybase recommends that you not install an older version of a Sybase product in a `$SYBASE` directory that already includes a newer Sybase product as this may not work. For example, attempting to install Replication Server 12.6 on top of Adaptive Server 15.0 may make either Adaptive Server Enterprise or Replication Server inoperable depending on the choices taken during install.

If the installation stops responding, see “ClearCase causes InstallShield to stop responding” on page 18 for more information.

3.1 Using SySAM 2.0

Starting with Adaptive Server Enterprise version 15.0, the Sybase Software Asset Management System (SySAM) implementation has changed. With these changes, SySAM configuration is no longer optional, and requires some changes in the Adaptive Server Enterprise installation and configuration process. Review the updated SySAM implementation and plan your SySAM deployment before you install Adaptive Server Enterprise.

See the *Sybase Software Asset Management User's Guide* for complete details.

Warning! Adaptive Server Enterprise works for an initial period of 30 days without proper SySAM configuration. If your configuration does not identify a valid license within this 30-day grace period, Adaptive Server Enterprise shuts down.

3.2 Runtime libraries required for Linux on Power installation

Warning! You cannot install Linux on Power without the following procedure.

❖ Installing Visual Age xlc compiler runtime libraries

Adaptive Server is compiled and linked with the Visual Age compiler, xlc version 7.0. Before installing Adaptive Server, you must install the runtime libraries for the Visual Age compiler.

- 1 Confirm that the runtime libraries are installed by entering:

```
rpm -q vacpp.rte-7.0.0-0
```

If this command returns “vacpp.rte-7.0.0-0”, the Visual Age runtime packages are installed and you can skip the following steps.

If the command returns the message that vacpp.rte-7.0.0-0 is not installed, you must install them.

- 2 Go to the IBM Web site at <http://www-1.ibm.com/support/docview.wss?uid=swg24007906>.

Follow the instructions to download and install the runtime libraries on your operating system. Select the package that is appropriate for your Linux distribution. For details on different Linux distributions supported and operating system level required, go to the Sybase certification Web site at <http://certification.sybase.com/ucr/search.do>.

- 3 After installing the xlc runtime libraries, verify that LD_LIBRARY_PATH is set to include the xlc runtime library path.

3.3 Monitoring and Diagnostic (MDA) tables

The *installmontables* script for installing Monitoring and Diagnostic tables (MDA) has been modified to now use the materialized parameter when defining the MDA proxy tables. This change results in improved performance for many queries using the MDA tables, making it no longer necessary to define a “loopback” remote server in your *syservers* table in order to use the MDA tables.

Sybase recommends you reinstall the MDA tables using the *installmontables* script contained in this release.

Note Do not make any changes to the *installmontables* script if you do not plan to use the MDA tables to monitor remote servers.

❖ Monitoring remote servers using MDA tables

Make these changes to use the MDA tables to monitor remote servers:

- 1 Update the *installmontables* script by changing all occurrences of “materialized” to “external procedure.”
- 2 Modify the *at* clause in the proxy table definitions in the *installmontables* script to add the name of the remote server followed by the string “...” before the MDA RPC name.
- 3 You can identify MDA RPC names by the “\$” (dollar sign) preceding the name. The remote server name must be the same as the remote server name as it appears in the *interfaces* file on the server where the MDA tables are installed.

For example, for the *monState* table definition, if your remote server name appearing in the *interfaces* file is SALESPROD, change:

```
at "$monState"
to:
at "SALESPROD...$monState"
```

If the remote server has not already been registered with the server on which you are installing the MDA tables, use the *sp_addserver* stored procedure to do so.

When installing the MDA tables in order to reference a remote server as described above, be sure to use the version of the *installmontables* script that came with the version of the Adaptive Server Enterprise installed at the remote server.

Use of the materialized directive has changed the way some error messages from the MDA tables are sent to the client application. If all configuration parameters required by an MDA table used in a query are not enabled, then the client application receives an error message.

4. Special upgrade instructions

This section describes special upgrade instructions for Adaptive Server.

4.1 Upgrading compiled objects

Sybase recommends that you read “Upgrading compiled objects with dbcc upgrade_object” in Chapter 6, “Troubleshooting,” in the *Adaptive Server Installation Guide for Linux*.

5. Changes that affect existing applications

Following are changes in Adaptive Server Enterprise that affect your existing applications.

5.1 Directory changes

Sybase has made the following directory structure changes:

- Adaptive Server Enterprise components are in *ASE-15_0*, and Open Client and Open Services components are in *OCS-15_0*.
- In Adaptive Server Enterprise version 15.0, all product components use JRE-1_4, which is *jre142* in the *\$SYBASE/shared* directory.
- The licensing component (SySAM) is now available in the *SYSAM-2_0* directory.

5.2 Changes to ODBC and OLE DB Driver support

Adaptive Server version 15.0 contains new ODBC and OLE DB drivers developed by Sybase. The third-party rebranded ODBC and OLE DB Driver Kits included with previous versions are no longer shipped.

The retired ODBC Driver Kit was installed in `%SYBASE%\ODBC`, and registered with the ODBC Driver Manager as “Sybase ASE ODBC Driver”. The new ASE ODBC Driver by Sybase is installed in `%SYBASE%\DataAccess\ODBC`, and registered as “Adaptive Server Enterprise”. The version shipping with Adaptive Server and SDK 15.0 is version 15.0.0.50.

The retired OLE DB Driver Kit was installed in `%SYBASE%\OLEDB`, and used the provider short name of “Sybase.ASEOLEDBProvider” and the long name of “Sybase ASE OLE DB Provider”. The new ASE OLE DB Provider by Sybase is installed in `%SYBASE%\DataAccess\OLEDB`, and uses provider short name “ASEOLEDB”. The version shipping with SDK 15.0 is version 15.0.0.51.

Sybase recommends that you update your ODBC Driver and OLE DB Provider to the new Sybase ODBC Driver and OLE DB Provider as soon as possible. Specifically, you should migrate:

- When upgrading to Adaptive Server Enterprise version 15 or later versions of Adaptive Server Enterprise version 12.5.x
- For older versions of Adaptive Server Enterprise, before July 30, 2007

For your convenience during the migration process, the third-party driver and provider support connectivity to the 15.x version of Adaptive Server Enterprise until November 29, 2007, but will not support new features added after ASE 12.5.3. Patches released after that date will perform a version check, and will not allow connection to ASE 15.x. This version-check behavior will only be incorporated in patches released after November 30, 2007. If patches released after November 30, 2007 are not applied, the third-party driver and provider can continue to connect to any versions of Adaptive Server available in the market.

See *New Features for OpenServer 15.0 and SDK 15.0 for Windows, Linux and UNIX*, for instructions on migrating to the new drivers.

For more information on migrating to the new OLE DB Provider, see *Open Server 15.0, Open Client 15.0 and SDK 15.0* and *New Features Open Server™ 15.0 and SDK 15.0 for Microsoft Windows, Linux, and UNIX* for more information.

5.3 Changes that affect application behavior

This section discusses changes that affect application behavior.

5.3.1 Long-identifier changes

- Adaptive Server Enterprise now supports long identifiers. There are new limits for the length of object names or identifiers: 255 bytes for regular identifiers, and 253 bytes for delimited identifiers. The new limit applies to most user-defined identifiers including table name, column name, index name and so on. Due to the expanded limits, some system tables (catalogs) and built-in functions have been expanded. For variables, “@” counts as 1 byte, and the allowed name is 254 bytes long.
- Change identifier names with corresponding application changes for binding values. Make sure your application is not binding names of identifiers with only 30 bytes, which was the previous limit. This may cause a variety of unexpected behaviors or error messages.

5.3.2 Query changes

- Some queries may now return general syntax error (message 102) instead of syntax error at line # (message 156).
- The order of result sets in Adaptive Server Enterprise version 15.0 differs unless there is an order by clause in the query.
- Query compilation time may increase as the query processing engine looks for more ways to optimize the query.
- See *Query Processor* for details about query processing in Adaptive Server Enterprise version 15.0.

5.3.3 Component Integration Services changes

- NULL behavior differs in Oracle, ASA, and IQ—you must override ANSI NULL behavior for it to work as in Adaptive Server Enterprise version 15.0.
- Adaptive Server Enterprise version 15.0 no longer pads char null, varchar, binary null, and varbinary datatypes for proxy tables.
- By default, cursors are READONLY. Declare cursors with for update to update through them.
- You must have an index if you are using updateable cursors.

- Component Integration Services engines no longer take special measures to re-declare cursors for back ends that close cursors on end tran.
- You must declare an explicit begin tran and end tran around cursor statements for DB2 servers.
- If you attempt to create a column constraint on java, text, image, and unitext datatypes, message 11074 is raised.
- create table, create existing table, or create proxy table statements that contain a location clause are restricted to be the only statement in a batch. This also applies to select into statements that contain a location clause.

5.3.4 Error message changes

- Many messages have been changed to specify “ASE” in the error message.
- When creating a temporary table that already exists, Adaptive Server Enterprise raises message 12822, instead of 2714.
- The identity column overflow message is now raised with message 587 instead of 4916.
- When you create a Java function that does not exist in the catalogs, message 14216 is raised instead of syntax error message 195.
- When a non-owner executes sp_procxmode to change the transaction mode associated with a stored procedure, error message 10354 is raised.
- Arithmetic overflow errors are now raised with message 3606 with severity 16.
- Message 2579 has been replaced with message 12907 in dbcc checktable output.

5.3.5 jConnect version 6.05 ships with Adaptive Server Enterprise

Adaptive Server Enterprise 15.0 ships only with jConnect™ for JDBC™ version 6.05. If you have applications that are dependent on jConnect 5.5, Sybase recommends you either migrate those applications to jConnect 6.05 or use an existing jConnect 5.5 release area.

5.3.6 SQL Remote no longer ships on the PC-Client CD

SQL Remote has been replaced with a more flexible and powerful technology called MobiLink that provides bidirectional synchronization between ASA/UltraLite clients and various back-end databases, including Adaptive Server Enterprise. To download the developer edition of MobiLink, go to the iAnywhere Web site at <http://www.iAnywhere.com/developeredition>.

5.3.7 Table changes

- The following Monitoring and Diagnostic Access (MDA) tables have new columns:
 - monEngine
 - monCachedObject
 - monProcessObject
- monCachedObject now tracks cached pages based on indid and partitionid.
- sysindexes has the following new columns:
 - partitiontype
 - conditionid
 - status3
- The following sysindexes columns are now maintained in syspartitions:
 - doampg
 - ioampg
 - first
 - root

These columns display 0 after you upgrade. The base_partition column is now obsolete and displays 0 after you upgrade.
- The syspartitions table is renamed to syslices during the upgrade process, then the new table is empty and unused.
- syscomments gains a new column called partitionid.
- systabstats gains the following new columns:
 - partitionid
 - plldegree

- statmoddate
- sysstatistics gains a new column called partitionid.
- There is a new type of object called a partition condition object, which has a row in sysobjects. A partition condition object is the representation of a tree for the partition table boundary conditions. The tree is stored in sysprocedures.
- sysobjects has the following new columns:
 - identburnmax
 - spacestate
 - erlchgts
- sysstatistics stores the data change counters with formatid=108. The space required by sysstatistics increases due to the additional rows stored.
- Most system catalogs have been converted to the datarows-locking scheme. However, DDLs continue to use the table-level locks. The row locking of the system catalogs can require an increase in the configuration parameter number of locks, depending on the DDLs in the application.
- The system catalogs converted to the datarows-locking scheme do not have a clustered index with indid equal to 1. The clustered indexes now have index ID greater than or equal to 2.

5.4 Changes that affect database administration operations

This section discusses changes that affect database administration operations.

5.4.1 Usage of system built-in functions

Built-in functions that provide space information, such as data_pgs, reserved_pgs, used_pgs, ptn_data_pgs, and rowcnt have been replaced with data_pages, reserved_pages, used_pages, and row_count, respectively. See *Reference Manual: Building Blocks* for detailed information.

5.4.2 DDL and DML changes

- alter table... unpartition is not allowed on tables with indexes (use alter table with 1 partition to remove the partitions).
- alter table on a partitioned table with max parallel degree less than the number of partitions may succeed without raising message 326.

- Creating a clustered index on an empty partitioned table returns a new informational message, 1936.
- null column names are not allowed during view creation.
- A select that contains an aggregate from a proxy table that is mapped to an RPC with a parameter fails with error message 201.

5.4.3 System stored procedure changes

- The output of `sp_help` *object* has been changed. Specifically, `Data_Located_on_segment` has been removed, while `Computed_Column_Object` and information related to partitions have been added.
- The order of `index_keys` and `index_description` in the output of `sp_helpindex` has been changed.
- Major changes have been made to the output of `sp_helppartition` and `sp_helpsegment`.
- `sp_who` returns “NULL” instead of blank in the `hostname` column for all system tasks.
- `sp_who` results in SQL command (insert or select) in the output. In earlier versions, `sp_who` returned select for the tasks that executed `sp_who`.
- `sp_objectsegment` now displays segment information for all partitions of the table.
- `sp_lock` output has a new column called `partitionid`, that currently has a value of 0, and is reserved for future use.
- `sp_monitorconfig` now accepts number of open partitions. When this procedure is executed with the parameter `all`, additional output for configuration parameter `number of open partitions` is displayed.
- The Metadata Cache Management section of `sp_sysmon` displays additional information pertaining to open partitions.

5.4.4 Common diagnostics changes including trace flag usage

- `dbcc listoam` output has changed substantially.
- The space state message printed in the old `dbcc listoam` output is now available as a new column `spacestate` in `sysobjects`.
- `dbcc tablealloc` output has changed.

- dbcc page output now displays the print partition ID instead of the object ID.
- dbcc checktable output on a partitioned table has been changed to provide partition-level information.
- dbcc checktable/tablealloc/indexalloc has new syntax for partition support.
- Output for this command has changed:

```
REORG RECLAIM_SPACE <tablename> with RESUME
```
- reorg rebuild fails with message 11051 when table is in use by other tasks.
- Output of sp_dbcc_faultreport has been enhanced to show partition ID information.
- The hostname, program_name, hostprocess, and cmd columns in sysprocesses have been changed to varchar(30) and made nullable. select from sysprocesses returns NULL for these columns instead of space for system tasks.
- The “first” column in sysindexes has been moved to the “firstpage” column in syspartitions.
- The name column in syspartitions has been changed from index_name+_+tableid to index_name+_+ptn_id for partition table with indexes.
- The basic cost of the optimizer now includes CPU cost.
- Parallel costing is now performed only on base tables and indexes that are larger than 20 pages.
- For information about trace flags and diagnostics, see the *Query Processor* document.

5.4.5 System resource changes

- Adaptive Server Enterprise version 15.0 uses more procedure cache for several reasons, including:
 - The query processing engine now looks for additional ways to optimize the query.
 - The execution engine avoids materialization of worktables and evaluates aggregations in memory as much as possible.

- The data change counters maintained for the `datachange()` function allocate memory from procedure cache. The partition condition tree is cached in the procedure cache. Partition boundary values are allocated in the procedure cache, resulting in a required increase in procedure cache resources.
- A configuration parameter, `max repartition degree`, has been added. This parameter controls the maximum degree to which an intermediate data stream can be repartitioned. The default value of this parameter is 1, which indicates that repartitioning is not set and is bound by the number of online engines configured for Adaptive Server Enterprise. For a query with a large number of tables, Adaptive Server Enterprise version 15.0 can put an increased demand on auxiliary scan descriptors.

Set the value of `max repartition degree` to a value lower than the number of engines to decrease resource usage. You may also need to configure a larger value for the auxiliary scan descriptor pool.

- Adaptive Server Enterprise version 15.0 avoids worktable materialization and incurs more resources in auxiliary scan descriptors.
- During the upgrade process, `max memory` is increased by Adaptive Server Enterprise if the new total logical memory is greater than `max memory`. The new value of `max memory` is set to the new value of total logical memory.
- During the upgrade process, number of open partitions is set to the same value as number of open indexes, resulting in increased memory usage. An open partition requires approximately 950 bytes.

5.4.6 Unpartitioning of user tables during upgrade

In version 15.0, each table partition must have a different partitionid. To avoid the expense of changing the partitionid for each page during the upgrade process, Adaptive Server Enterprise unpartitions any partitioned tables. You must repartition these tables when the upgrade process completes.

5.4.7 Space management changes

- System databases have increased in size by 2 allocation units. The master database has increased in size by 14 allocation units. The default user database size has increased by 2 allocation units.

- The upgrade process requires free space to convert system catalogs to the datarows-locking scheme. The typical requirement for free space is 125 percent of the largest system catalog. If you are upgrading from a 32-bit to 64-bit binary, additional space of 55 percent of sysprocedures is required.

5.4.8 Changes to transaction dump content

You can load a transaction dump from Adaptive Server Enterprise version 11.9.2, 12.0x, or 12.5x over a database dump or transaction dump that has preceded it in the load sequence.

However, if that transaction dump contains a create index transaction, then load transaction recovery treats that transaction's sort record as the end of the log. Adaptive Server Enterprise does not process any more log records in that transaction, and does not allow load transaction dumps in the load sequence. This message prints:

```
The pre-15.0 log for database '%.*s' includes a create
index incompatible with this server version and
operating environment. Only transactions completed
before create index will be recovered.
```

This behavior is much like the point-in-time recovery that is affected by the load transaction until_time option.

In contrast, Adaptive Server Enterprise should successfully load any database dump from one of the earlier versions, regardless of its content.

There is no content limitation in load database or load transaction when using a dump from a version 15.0 Adaptive Server Enterprise.

5.5 Open Client and Open Server changes

For information about changes that affect Open Client and Open Server, see most recent release bulletins for these products at the Sybase Web site at <http://sybooks.sybase.com>.

6. Known problems

This section describes known problems and workarounds for Adaptive Server Enterprise.

Note Sybase no longer includes System Problem Reports (SPR) and Closed Problem Reports (CPR) with your copy of Adaptive Server Enterprise. Known problems, which used to be in the SPR, are now described in this section of the release bulletin. Closed issues that were documented in the CPR are now searchable from Solved Cases at the Sybase Web site at <http://www.sybase.com>. To access the Solved Cases search form, click Support and Services | Solved Cases.

6.1 Highlighted known installation issues

This section contains information about known installation problems and their workarounds.

Warning! If you have multiple Sybase products installed in the same `$$SYBASE` directory, uninstalling one of them *may* make the other products inoperable. If this happens, you may need to reinstall those remaining products.

6.1.1 ClearCase causes InstallShield to stop responding

InstallShield may stop responding, and display this message:

```
Please wait while the installer checks your system
information.
```

There are two known causes:

[CR #404068] The JRE used by InstallShield is in conflict with ClearCase.

Workaround: Turn off ClearCase before installing Adaptive Server Enterprise. See your ClearCase administrator for details. You can turn ClearCase back on after installation.

[CR #402530] NFS-mounted device is inaccessible because the remote server is not responding. Verify this by executing `df -P` from a command prompt. You may instead see a message similar to: `NFS3 server not responding still trying`.

Workaround: Unmount the inaccessible device and retry the installation. You may need to restart the server.

Note You cannot use these workarounds on a 64-bit hardware system. There are no workarounds for the 64-bit systems.

6.2 Highlighted known XML problems

This section describes known problems with XML.

6.2.1 Storing a parsed XML document in a *varbinary* column

[CR #400269] If you store a parsed XML document in a varbinary column, the document is truncated by one byte at the end. As a result, you cannot query that document.

Workaround: Do not store a parsed XML document in a varbinary column.

6.2.2 Cross-platform bulk copying of parsed XML images

[CR #400250, 332012] You cannot use bcp or replicate parsed XML image data between a big-endian platform and a little-endian platform, for example between a Solaris platform (big-endian) and a Linux or Windows platform (little-endian).

The term “parsed XML image data” refers to data of datatype image that is generated by the `xmlparse` function.

Workaround: When you transfer XML data between platforms, transfer the character form of the XML documents rather than the parsed XML form. If you have not stored the character form of the XML documents, you can regenerate it from the parsed form. For example, if column `xmlindexed` of the `xmldocuments` table is an image column containing parsed XML image data, the following two commands append a new text column to the table and populate it with a character form of the documents contained in the `xmlindexed` column:

```
alter table xmldocuments add xmltext text
update xmldocuments set xmltext = xmlextract('/', xmlindexed)
```

You can then use bcp or replicate the `xmltext` column from one platform to the other.

6.3 Known alter table problems

This section contains information about known alter table problems and their workarounds.

6.3.1 *alter table unpartition* errors

[CR #400988] If a server crashes following an alter table unpartition operation on very large tables, the server recovery after restart may leave the table in an inconsistent or corrupted state. dbcc check table and other commands affecting a table's pages may result in incorrect page errors. In some cases, the recovery process encounters timestamp mismatches on pages allocated to the table being unpartitioned. This form of corruption may occur when you issue a shutdown no wait command, and can affect very large tables with multiple OAM pages in the OAM chain for each partition.

Workaround: Run checkpoint <database> immediately following an unpartition operation. You can also issue a shutdown command without the no wait clause to avoid this form of corruption.

6.4 Running on RH3.0 64-bit distribution occasionally stops responding

Due to a Red Hat 3.0 update 3 anomaly, Adaptive Server Enterprise occasionally stops responding when kernel asynchronous I/O is enabled. When this happens, Adaptive Server Enterprise generally displays a high CPU load—approximately 99%—without I/O. Adaptive Server Enterprise can still accept connections, but cannot recover; you must restart the system to clear the process.

This problem has been reported to Red Hat and addressed in kernel 2.4.21-27.6.EL.

For more information, see the Red Hat Web site at https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=138905.

6.5 Interactive SQL and Adaptive Server Enterprise plug-in issues

This section discusses Interactive SQL and Adaptive Server Enterprise plug-in issues and their workarounds.

6.5.1 Opening XML files with Interactive SQL

[CR #400825] Interactive SQL displays a stack trace trying to open an XML file.

Workaround: There is no workaround for this issue.

6.5.2 Displaying messages next to result sets with Interactive SQL

[CR #386931] Interactive SQL does not display messages next to result sets, making it difficult to read the output of certain stored procedures that have multiple results and messages, such as `sp_help`.

Workaround: From command line use `isql`.

6.5.3 Running a script file with Interactive SQL

[CR #401391] Interactive SQL cannot run script files without a `.sql` extension.

Workaround: Instead of using “Run Script,” use File | Open, and then run the script manually.

6.5.4 Connecting or disconnecting from pass-through servers using Interactive SQL

[CR #400436] The Adaptive Server Enterprise connect to `<server_name>` and disconnect commands conflict with the Interactive SQL connect and disconnect commands.

Workaround: There is no workaround for this issue.

6.5.5 Interactive SQL Make Permanent button

[CR #400053] Interactive SQL Make Permanent button does not save options permanently.

Workaround: There is no workaround for this issue.

6.5.6 Interactive SQL does not generate plans for SQL statements using go

[CR #400362] Interactive SQL cannot generate plans for SQL statements separated with `go`.

Workaround: Eliminate `go` between SQL statements.

6.5.7 Connecting Adaptive Server Enterprise plug-in to RepConnector

[CR #400709] The Adaptive Server Enterprise plug-in cannot connect to Adaptive Server Enterprise if RepConnector™ 2.5 and the Adaptive Server Enterprise plug-in 15.0 are installed in the same *SYBASE* location.

Workaround: Install RepConnector and the Adaptive Server Enterprise plug-in into different *SYBASE* directories.

6.5.8 Query support in Interactive SQL

[CR #398435] Interactive SQL does not support queries with a compute clause.

Workaround: Use `isql` from the command line instead.

6.5.9 Adaptive Server Enterprise plug-in drag-and-drop server capability

[CR #400226] The Adaptive Server Enterprise plug-in drag-and-drop server capability is sometimes not available.

Workaround: Use copy, cut, or paste instead, by right-clicking on the selected server and selecting Copy, Cut, or Paste.

6.5.10 Finding help in the Options dialog

[CR #399507] Interactive SQL cannot find the help topic `HELP_SCEDITOR_TAB_CUSTOMIZATION` in the Options dialog.

Workaround: Start Interactive SQL online help by selecting Help | Interactive SQL | Help. Open Using Interactive SQL Help and go through the list of topics for the information you need.

6.5.11 Unified Agent Framework RMI ports

[CR #398968] The Adaptive Server Enterprise plug-in does not allow users to change the default Unified Agent Framework RMI port to anything other than 9999.

Workaround: Use 9999 as the default port.

6.5.12 Adaptive Server Enterprise plug-in data panel for a table with a binary null column

[CR #382837] When a binary column has a length of less than 256, the Adaptive Server Enterprise plug-in data panel for a table with a binary null column displays a value with trailing nulls.

Workaround: There is no workaround for this issue.

6.5.13 Starting Interactive SQL from a command line fails

[CR #402435] When trying to start Interactive SQL from the command line, an error is returned starting with:

```
"unrecognized option -path=....."
```

On Windows, starting Interactive SQL from the Start menu results in a brief flash from a DOS window starting up and immediately closing.

Workaround: Change the `-path` command line argument to `-Dpath` in the script file. The script file is located at:

- `$SYBASE/DBISQL/bin/dbisql` on UNIX
- `%SYBASE\DBISQL\bin\dbisql.bat` on Windows

6.6 Statistics in system tables

[CR #399624] When you create or configure a server in, or upgrade a server to, Adaptive Server Enterprise 15.0, index and table-level statistics are incorrect.

Workaround: Execute update index statistics on the following system tables:

- `sysobjects`
- `sysindexes`
- `syscolumns`
- `systypes`
- `syslogins`
- `sysusers`

6.7 Running *checkstorage* on an upgraded master device may report faults with the *syscharsets* catalog

[CR #401779] During the upgrade to Adaptive Server Enterprise 15.0, a master device originally created either in Adaptive Server Enterprise 11.9.x or 12.0.x may report errors on the *syscharsets* catalog. Adaptive Server Enterprise reports these faults when you execute *checkstorage* on the upgraded master device. *sp_dbcc_faultreport* reports these faults.

Workaround: You can safely ignore these faults.

6.8 Inserting into a table with *identity* column using *select distinct*

[CR #401753] If you have enabled parallel plan, and use *select distinct* into a table with an *identity* column and one of the *select distinct* list items is an expression or an *implicit convert()* expression, this query may cause a stack trace. Here is an example query:

```
create table tt(id numeric(5) identity not null, c1 int,
               c2 int)
go
insert into tt select distinct c1, c2+1 from t
go
```

Workaround: Disable parallel plan for the query. For the example query, use:

```
insert into tt select distinct c1, c2+1 from t (parallel 1)
```

6.9 SySAM licenses needed to upgrade Adaptive Server Enterprise

[CR #401682] When you are upgrading from a version earlier than 15.0, Adaptive Server Enterprise does not correctly calculate the grace period, which results in a failed upgrade.

Workaround: Before performing an upgrade, make sure you have obtained suitable licenses from the Sybase Product Download Center (SPDC). Check your Adaptive Server Enterprise 15.0 installation to make sure that Adaptive Server Enterprise can obtain the required licenses. See Chapter 11, “SySAM Configuration,” in the configuration guide for your platform for more details.

6.10 Cross-platform dumps and loads

[CR #371289] Adaptive Server Enterprise does not support cross-platform dumps and loads from one Adaptive Server Enterprise version 15.0 to another. However, you can successfully use the dump and load commands from versions earlier than 15.0 to version 15.0.

Workaround: There is no workaround for this issue.

6.11 Running *dbcc checktable* or *dbcc checkdb* before bringing a database online

[CR #356308] If you run *dbcc checktable* or *dbcc checkdb* after loading a database, but before bringing it online, you may see errors due to a mismatch in the number of rows in *systabstats* data and index.

Workaround: Do not run *dbcc checktable* or *dbcc checkdb* until you have loaded and brought a database online.

6.12 *allow backward scan* does not work

[CR #401543] The configuration parameter *allow backward scan* does not work in Adaptive Server Enterprise version 15.0.

Workaround: If you are encountering increased deadlocks while executing backward scans, separate problematic scans:

- 1 Select the required rows into a temporary table in ascending order.
- 2 Select from the temporary table in descending order.

6.13 Modifying *dbccdb* schema for large identifiers

[CR #400680] The *dbccdb* database does not support workspace names longer than 30 characters.

Workaround:

- 1 In the *dbcc_exclusions* table in *dbccdb*, modify the *table_name* column from *varchar(30)* to *varchar(255)*.
- 2 In *sp_dbcc_run_evaluatedb*, modify the *@wsname* local variable from *varchar(30)* to *varchar(255)*.

6.14 Drop definition timetable corruption

[CR #397910] If an attempt to create a new table using select into fails, a subsequent dbcc checkcatalog in the target database may generate errors on missing rows between syscomments and sysprocedures. These errors occur for the definition timetable that was created as part of the previously failed select into statement. One common instance of such failures is an insufficiently configured number of locks parameter, which can result in a 1204 error message.

Workaround: Drop the table referenced by the object ID in the error message. Re-execute the failed select into command after increasing the number of locks configuration parameter.

6.15 Manual upgrade with named cache fails

[CR #401279] If a named cache exists in the installation being upgraded and if you use a configuration file that does not contain the named cache configuration during upgrade, the upgrade fails.

Workaround: Use the sqlupgrade utility to perform the upgrade. sqlupgrade uses the configuration file from the installation being upgraded.

6.16 Running *diskrefit* to recover a corrupt master device can report errors

[CR #399678] If you start Adaptive Server Enterprise using the -w master parameter, then immediately start it again using the -w model parameter, the master database reports that a page should belong to object 8, syslogs, but instead belongs to some other object. Additionally, tempdb may report 806 errors.

Workaround: Between these two operations, restart Adaptive Server Enterprise with -m and then shut down.

6.17 Using an *equijoin* clause between two different length columns

[CR #401336] If you have an equijoin clause between two char() and char()/varchar() type columns of different lengths, and also use char_length() on one of the columns, the result of the char_length() is indeterminate. For example:

```
create table t1 (sid char(8), styp char(5))
```

```
go
create table t2 (sid char(16))
go
select char_length(a.sid)
from t1 a, t2 b
where a.sid=b.sid and styp = 'R1'
go
```

Workaround: Use `convert()` in `char_length()` to state explicitly which type you need for the `char_length`.

For example, in the above query, use:

```
select char_length(convert(char(8), a.sid))
from t1 a, t2 b
where a.sid=b.sid and styp = 'R1'
```

The same is true for `binary()` and `varbinary()` with `data_length()`, and `unichar()` and `univarchar()` with `char_length()`.

6.18 *set statistics io* does not display I/O generated by worker processes

[CR #358654] After executing a parallel query, `set statistics io` does not display I/O counts generated by worker processes. It displays only the parent thread scan, logical I/O, and physical I/O counts.

Workaround: There is no workaround for this issue.

[CR# 401560] After executing a parallel query, `set statistics io` displays only shows the alpha thread scan, and logical and physical I/O counts. `set statistics io` does not display I/O counts generated by worker processes.

Workaround: Use `set statistics plancost on` to display the I/O counts from alpha thread scan and any child threads in a graphical format.

6.19 IPv6 support causes Adaptive Server Enterprise to stop responding

[CR #392074] Machines configured for IPv6 may stop responding.

Workaround: Add the following line to the `/etc/nsswitch.conf` file:

```
ipnodes: files nis dns
```

6.20 Replication Server compatibility issues

[CR # 382874] `rs_init` in Replication Server version 12.6 and earlier cannot create an RSSD database in Adaptive Server Enterprise version 15.0.

Workaround: Create the RSSD database manually before using `rs_init` to create a new Replication Server if you do not use an Embedded RSSD. See the *Replication Server Administration Guide* for additional information.

6.21 Dumping or loading databases with asynchronous I/O

[CR #335852] On an IA32 running Red Hat, running a dump or load database command can cause Backup Server to stop responding when using asynchronous I/O. Backup Server uses asynchronous I/O by default.

Workaround: Start Backup Server using trace flag `-D32` to force a synchronous I/O.

6.22 `cis connect timeout` and `enable SNMP` are not implemented

[CR #323177] The `cis connect timeout` and `enable SNMP` configuration parameters are not currently implemented.

Workaround: Do not use these configuration parameters.

6.23 Handling multibyte character sets during migration

[CR #353079] If you configure multibyte character sets after migrating data from system catalogs `sysattributes` and `sysxtypes`, the text columns in these catalogs are inconsistent with the multibyte character sets.

Workaround: Manually run `dbcc fix_text` on `sysattributes` and `sysxtypes` to make the text columns consistent with the multibyte character sets.

6.24 Network package reallocation during reconfiguration

[CR #409471] **On Linux on Power** – if the dynamic reconfiguration on demand configuration parameter is set to 0, increasing the number of user connections may result in network packet allocation failure. However, network packet allocation succeeds when you make an actual user connection to Adaptive Server Enterprise.

Workaround: Retry the same increase in number of user connections.

6.25 *bcp in or bcp out*

[CR #426810] Adaptive Server Plug-in using Sybase Central with the 150ESD#2 image gets a stack trace while doing bcp in or out.

Workaround: Once you complete the installation of 150ESD#2:

- 1 Remove `.scRepository` from `$SYBASE/shared/sybcentral43`.
- 2 Reregister Adaptive Server Plug-in with `$SYBASE/shared/sybcentral43/registerASEP`.

6.26 Silent installation

[CR #428343] When you perform a silent installation the versioning does not change to ESD#2. No errors are flagged after the install.

Workaround: In a silent installation, use:

```
./setup -is:javaconsole \  
-silent \  
-W SybaseLicense.agreeToLicense=true \  
-G overwriteJVM=yes \  
-P installLocation="filename"
```

Do not specify an options file.

6.27 Sybase Central

[CR #428292] Using the Adaptive Server Plug-in with an older Sybase Central release results in “SybPasswordField not found” stack trace.

Workaround: Update the existing Sybase Central version, or install Sybase Central in a new location.

6.28 File names

[CR #416569] You can update a file name multiple times if updates are not restricted.

Workaround: When updating the filename column, qualify the update to specific files by including the *filename* in the where clause.

6.29 Command line flag *-k*

[CR #425696] Using the *-k* command line option, to set the Kerberos Principal Name, produces usage errors.

Workaround: Set the environment variable `SYBASE_PRINCIPAL` to the principal name, rather than using the *-k* command line option.

6.30 *table scan* and *index scan*

[CR # 404698] Index Usage counters are not being updated correctly.

Workaround: None.

6.31 Determining devices for a database

[CR# 407828] The Installation Guide gives incorrect information on how to determine which device a database is using. The `infocenter.sybase.com` search function may show many hits but it does not show the documents where the hits occur.

Workaround:

- In versions of Adaptive Server Enterprise prior to 15.0, you use `sysdevices` to determine which device has a low through high virtual page range that includes the `vstart`. The device fragment whose `vstart` you used is on that device.
- Select the `vdevno` from `sysusages` matching the `dbid` retrieved in step1.

7. Product compatibilities

This section lists the Sybase components that are compatible with Adaptive Server Enterprise version 15.0. For information about operating system requirements, see individual component documentation.

Note Component Integration Services and SSL are not supported on Adaptive Server Enterprise version 15.0 for Linux.

The following components have been tested for compatibility with Adaptive Server Enterprise version 15.0:

- Sybase Character Sets 3.0

- Sybase Central viewer 4.3
- ECDA option for ODBC 12.6
- ECDA option for Informix 12.6
- MainframeConnect DirectConnect for z/OS 12.6
- ECDA option for Oracle 12.6
- ECDA option for Microsoft SQL Server 12.6
- ECDA option for DB2 Universal Database 12.6
- OpenSwitch 12.5, OpenSwitch 12.5.1, ESD #2

Note OpenSwitch 12.5 does not support SSL.

- jConnect for JDBC 5.5
- jConnect for JDBC 6.05
- Open Server™ 12.5.1
- Open ServerConnect™ (CICS, IMS/MVS) 4.0
- Open ClientConnect™ (CICS, IMS/MVS) 4.0
- InfoMaker™ 7.0.3
- Replication Server® 12.6
- PhysicalArchitect™ 8.0
- XA-Library™ for CICS/ENCINA/Tuxedo 12.5.1
- ADO.NET 1.15.x Windows x86 32-bit
- ADO.NET 1.1x Windows x86 32-bit
- ASE ODBC Driver by Sybase 12.5.x Window x86 32bit
- ASE ODBC Driver by Sybase 15.0.x Window x86 32bit and x64 64bit
- ASE OLE DB Provider by Sybase 12.5.x Windows x86 32-bit
- ASE OLE DB Provider by Sybase 15.x Windows x86 32-bit
- ASE ODBC Driver by Sybase Linux x86 32 bit

7.1 Adaptive Server interoperability

The following tables shows the interoperability of Adaptive Server version 15.0 against other Sybase products across different platforms and versions, as well as client products. For specific platform or operating system level information, perform a search at the Sybase Certification Web page at <http://certification.sybase.com/ucr/search.do> for the appropriate platform certification.

Note 1 Even though two or more products might be interoperable, features introduced in a newer version of a product are not likely to be supported with older versions of the same or other products.

Note 2 Interoperability between big-endian and little-endian platforms has also been verified. Windows and Linux-32 platforms are little-endian. IBM AIX, Sun Solaris, Linux on Power, and HP-UX are big-endian platforms.

Table 1: Adaptive Server version 15.0.x interoperability with other products

Platforms		OC/OS 12.0x	OC/OS 12.5X	OC/OS 15.0X	jConnect 5.5	jConnect 6.0.x	ASE ODBC by Sybase 12.5x	ASE ODBC by Sybase 15.0x	ASE OLEDB by Sybase 12.5x	ASE OLEDB by Sybase 15.0x	ADO.NET 1.1x	OEM ODBC Driver Kit 5.00.0096	OEM OLEDB Driver Kit 2.70.0063	Replication Server 12.6	Replication Server 15.0.x
AIX	32-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
HP-UX	32-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	IA 64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
Linux	x86 32-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	Pseries 64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	x64 64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	IA 64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
Solaris	32-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
	64-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X
Windows	x86 32-bit	X	X	X	X	X	X	X	X	X	X	+	+	X	X

+ Available until November 2007. See “Changes to ODBC and OLE DB Driver support” on page 9 for details.

Table 2 shows the supported platforms of client products listed in Table 1.

Table 2: Supported client platforms

Client name	Version	Supported platform	
Open Client/Open Server	12.0.x	<ul style="list-style-type: none"> • AIX 32-bit • HP-UX 32-bit • Solaris 32-bit 	<ul style="list-style-type: none"> • TruUNIX (Alpha) 64-bit • Windows x86 32-bit
	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 Opteron 64bit • TruUnix (Alpha) 64-bit • Windows x86 32-bit
	15.0x	<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 • Linux Pseries 32-bit • Linux Pseries 64-bit 	<ul style="list-style-type: none"> • Solaris 32-bit • Solaris 64-bit • Solaris Opteron 32-bit • Solaris Opteron 64-bit • Windows x86 32-bit • Windows x64 64-bit
jConnect	5.5	<ul style="list-style-type: none"> • All 	
	6.0x	<ul style="list-style-type: none"> • All 	
ASE ODBC by Sybase	12.5.x	<ul style="list-style-type: none"> • Linux x86 32-bit • Linux x64 32-bit • Mac OSX 32-bit 	<ul style="list-style-type: none"> • Windows x86 32-bit • Windows x64 64-bit
	15.0.x	<ul style="list-style-type: none"> • Linux x86 32-bit • Linux x64 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
ASE OLEDB by Sybase	12.5.x	<ul style="list-style-type: none"> • Windows x86 32-bit • Windows x64 32-bit
	15.x	<ul style="list-style-type: none"> • Windows x86 32-bit • Windows x64 32-bit
ADO.NET	1.1.x	<ul style="list-style-type: none"> • Windows x86 32-bit • Windows x64 32-bit
OEM ODBC Driver Kit	5.00.0096	<ul style="list-style-type: none"> • Windows x86 32-bit
OEM OLEDB Driver Kit	2.70.0063	<ul style="list-style-type: none"> • Windows x86 32-bit
Replication Server	12.6	<ul style="list-style-type: none"> • AIX 32-bit • HPUX 32-bit • HPUX IA 32-bit • Linux x86 32-bit • Linux x64 32-bit • Mac OSX 32-bit • SGI 32-bit • Solaris 32-bit • TruUnix (Alpha) 64-bit • Windows x86 32-bit
	15.0.x	<ul style="list-style-type: none"> • AIX 32-bit • HPUX 32-bit • HPUX IA 32-bit • Linux x86 32-bit • Linux x64 32-bit • Solaris 32-bit • Windows x86 32-bit

7.2 Adaptive Server Replicator 15.0 support of Adaptive Server Enterprise 12.5 datatypes

Adaptive Server Replicator 15.0 does not support large identifiers, or these datatypes that are introduced in Adaptive Server Enterprise version 15.0: bigint, unsigned bigint, unsigned int, unsigned smallint, unitext, computed columns, and XML.

7.3 Known compatibility issues

This section discusses known compatibility issues.

7.3.1 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 Mainframe Connect DirectConnect for z/OS, including DirectConnect Manager, into its own product directory.

7.3.2 jConnect 6.0 and *sqldbgr*

Due to missing components in jConnect 6.0, *sqldbgr* does not run correctly if you select a “Typical” installation. If you plan on using *sqldbgr*, select the “Full” installation option to install jConnect6.0.

7.3.3 Installing Replication Server with other Sybase software

Because there are incompatibilities among some Sybase components, Sybase strongly recommends that you install Replication Server, including Replication Server Manager (RSM) Server, in its own product directory, separate from other Sybase products.

8. Documentation updates and clarifications

This section describes updates and clarifications for Adaptive Server Enterprise documentation.

8.1 *Reference Manual: Building Blocks*

This section discusses updates and clarifications to *Reference Manual: Building Blocks*.

8.1.1 Correction in *derived_stat*

The information for the *derived_stat* built-in function in *Reference Manual: Building Blocks* is incomplete.

Syntax

The syntax as it currently appears in the reference page is:

```
derived_stat(object_name | object_id,
            index_name | index_id,
            [partition_name | partition_id,]
            "statistic")
```

Replace with the following:

```
derived_stat("object_name" | object_id,
            index_name | index_id,
            ["partition_name" | partition_id,]
            "statistic")
```

Parameters

The descriptions for *partition_name* and *partition_id* have been updated with the following:

- *partition_name* – is the name of the partition, belonging to the specific partition that you are interested in. *partition_name* is optional. When you use *partition_name* or *partition_id*, Adaptive Server returns statistics for the target partition, instead of for the entire object.
- *partition_id* – is an alternative to *partition_name*, and is the partition ID of the specified object that you are interested in. *partition_id* is optional.

Examples

The Examples section has these two new examples:

Example 4 Selects of the titles table. Selects the data page cluster ratio for index ID 2, partition ID 960003420, of the titles table. Note that you can use either "dpcr" or "data page cluster ratio":

```
select derived_stat("titles", 2, 960003420, "dpcr")
```

Example 5 Selects of the titles table. Selects the data page cluster ratio for index ID 2, partition name titles_ptn1, of the titles table. Note that you can use either "dpcr" or "data page cluster ratio":

```
select derived_stat("titles", 2, "titles_ptn1", "dpcr")
```

8.1.2 Correction in *xa_bqual* example

Example 1 in the reference pages for the *xa_bqual* built-in function shows an incorrect binary translation. Replace the example with the following.

Example 1 Returns “0x227f06ca80”, the binary translation of the branch qualifier for the Adaptive Server Enterprise transaction ID “0000000A_IphIT596iC7bF2#AUfkzaM_8DY6OE0”. The Adaptive Server Enterprise transaction ID is first obtained using *sp_transactions*:

```
1> sp_transactions
```

```
xactkey          type      coordinator starttime      state
  connection      dbid  spid  loid  failover  srvname  namelen
xactname
-----
-----
-----
0x531600000600000017e4885b0700 External XA          Dec 9 2005 5:15PM
      In Command Attached      7      20      877      Resident Tx      NULL      39
0000000A_IphIT596iC7bF2#AUfkzaM_8DY6OE0
```

```
1> select xa_bqual("0000000A_IphIT596iC7bF2#AUfkzaM_8DY6OE0", 0)
2> go
```

```
...
```

```
-----
0x227f06ca80
```

8.2 Reference Manual: Commands

This section discusses updates and clarifications to *Reference Manual: Commands*.

8.2.1 Correction in *create table*

The Usage section of the *create table* reference pages includes the following information, which erroneously lists the limit of user-defined columns per table as 250:

- There can be as many as 2,000,000,000 tables per database and **250** user-defined columns per table. The number of rows per table is limited only by available storage.

The correct number of user-defined columns per table is **1024**.

8.2.2 Clarification in *dbcc*

Example 13 for the *dbcc* command is described as being an abbreviated output for the *dbcc traceon(3604)* command. You can obtain the same abbreviated version using *dbcc serverlimits*; running *dbcc traceon(3604)* results in a complete listing of limits in the server.

8.2.3 Correction in *select*

In the reference section for the *select* command, a bulleted item describes how to add a new *IDENTITY* column and use precision to specify the numeric datatype:

- A specification to add a new *IDENTITY* column to the result table:

column_name = identity(precision)

Beginning with Adaptive Server Enterprise version 12.5.4, you can specify the *IDENTITY* column with *int*, *smallint*, and *tinyint* as well as precision.

- A specification to add a new *IDENTITY* column to the result table:

column_name = identity(int | smallint | tinyint | precision)

If you specify *int*, *smallint*, or *tinyint*, the resulting column is an integer. If you specify precision, the result is a numeric datatype.

8.2.4 Correction in set proc_output_params off

The example provided for set proc_output_params off in the set command reference pages includes two errors in the output:

In the first portion, the value of @x returned from the command is erroneously listed as “2.” The value is 1, and appears in bold in the following corrected version:

In the second portion, the return parameters output is erroneously published as “3.” The value is 2, and appears in bold in the following corrected version:

Example 14 Suppresses the output of parameter information:

```
1> create procedure sp_pout (@x int output) as select @x = @x + 1
2> go
1> set proc_output_params off
2> go
1> declare @x int
2> select @x = 1
3> exec sp_pout @x output
4> print "Value of @x returned from sproc is: %1!", @x
5> go
(1 row affected)
(return status = 0)
```

Value of @x returned from sproc is: **1**

If you do not perform set proc_output_params off, the output after (return status = 0) includes the following:

```
Return parameters:
-----
                2
```

8.3 Reference Manual: Procedures

This section describes updates and clarifications in *Reference Manual: Procedures*.

8.3.1 sp_sproc_columns column_type description

In the reference section for the sp_sproc_columns stored procedure, the table listing the result set includes the following description of the column_type column:

This field always returns a value:

- 0 = SQL_PARAM_TYPE_UNKNOWN
- 1 = SQL_PARAM_TYPE_INPUT
- 2 = SQL_PARAM_TYPE_OUTPUT
- 3 = SQL_RESULT_COL
- 4 = SQL_PARAM_OUTPUT
- 5 = SQL_RETURN_VALUE

This information is erroneous and should be removed from the description.

8.3.2 sp_poolconfig

The reference pages for sp_poolconfig have been revised. Replace the pages for sp_poolconfig with the following.

Description	Creates, drops, resizes, and provides information about memory pools within data caches.
Syntax	<p>To create a memory pool in an existing cache, or to change pool size:</p> <pre>sp_poolconfig cache_name [, "mem_size [P K M G]", "config_poolK" [, "affected_poolK"]]</pre> <p>To change a pool's wash size:</p> <pre>sp_poolconfig cache_name, "affected_poolK", "wash=size[P K M G]"</pre> <p>To change a pool's asynchronous prefetch percentage:</p> <pre>sp_poolconfig cache_name, "affected_poolK", "local async prefetch limit=percent "</pre>
Parameters	<ul style="list-style-type: none"> • <i>cache_name</i> – is the name of an existing data cache. • <i>mem_size</i> – is the size of the memory pool to be created or the new total size for an existing pool with the specified I/O size. The minimum size of a pool is 256 logical server pages. For a 2K logical page size server, the minimum size is 256K. Specify size units with P for pages, K for kilobytes, M for megabytes, or G for gigabytes. The default is kilobytes. • <i>config_pool</i> – is the I/O size performed in the memory pool where the memory is to be allocated or removed. <p>Valid I/O sizes are multiples of the logical page size, up to four times the amount.</p> <ul style="list-style-type: none"> • <i>affected_pool</i> – is the size of I/O performed in the memory pool where the memory is to be deallocated, or the pools attributes such as 'wash size' and 'prefetch limit' are to be modified. If <i>affected_pool</i> is not specified, the memory is taken from the lowest logical page size memory pool.

- `wash=size` – changes the wash size (the point in the cache at which Adaptive Server writes dirty pages to disk) for a memory pool.
- `local async prefetch limit=percent` – sets the percentage of buffers in the pool that can be used to hold buffers that have been read into cache by asynchronous prefetch, but that have not yet been used.

Examples

Example 1 Creates a 16K pool in the data cache `pub_cache` with 10MB of space. All space is taken from the default 2K memory pool:

```
sp_poolconfig pub_cache, "10M", "16K"
```

Example 2 Creates 16MB of space to the 32K pool from the 64K pool of `pub_cache`:

```
sp_poolconfig pub_cache, "16M", "32K", "64K"
```

Example 3 Reports the current configuration of `pub_cache`:

```
sp_poolconfig "pub_cache"
```

Example 4 Removes the 16K memory pool from `pub_cache`, placing all of the memory assigned to it in the 2K pool:

```
sp_poolconfig pub_cache, "0K", "16K"
```

Example 5 Changes the wash size of the 2K pool in `pubs_cache` to 508K:

```
sp_poolconfig pub_cache, "2K", "wash=508K"
```

Example 6 Changes the asynchronous prefetch limit for the 2K pool to 15 percent:


```
sp_poolconfig pub_cache, "2K", "local async prefetch limit=15"
```

Usage

- When you create a data cache with `sp_cacheconfig`, all space is allocated to the logical page size memory pool. `sp_poolconfig` divides the data cache into additional pools with larger I/O sizes.
- If no large I/O memory pools exist in a cache, Adaptive Server performs I/O in logical page size units, the size of a data page, for all of the objects bound to the cache. You can often enhance performance by configuring pools that perform large I/O. A 16K memory pool reads and writes eight data pages in a single I/O for a 2K logical page size server.
- The combination of cache name and I/O size must be unique. In other words, you can specify only one pool of a given I/O size in a particular data cache in `sp_poolconfig` commands.
- Only one `sp_poolconfig` command can be active on a single cache at one time. If a second `sp_poolconfig` command is issued before the first one completes, it sleeps until the first command completes.

- Figure 1 shows a data cache on a server that uses 2K logical pages with:
 - The default data cache with a 2K pool and a 16K pool
 - A user cache with a 2K pool and a 16K pool
 - A log cache with a 2K pool and a 4K pool

Figure 1: Data cache with default and user-defined caches



- You can create pools with I/O sizes up to 16K in the default data cache for a 2K page size server.
- The minimum size of a memory pool is 256 logical pages (for example, a 2K logical page size server, the minimum size is 512K). You cannot reduce the size of any memory pool in any cache to less than 256 pages by transferring memory to another pool.
- Two circumstances can create pool less than 512K:
 - If you attempt to delete a pool by setting its size to zero, and some of the pages are in use, `sp_poolconfig` reduces the pool size as much as possible, and prints a warning message. The status for the pool is set to “Unavailable/deleted”.

- If you attempt to move buffers to create a new pool, and enough buffers cannot be moved to the new pool, `sp_poolconfig` moves as many buffers as it can, and the cache status is set to “Unavailable/too small.”

In both of these cases, you can retry the command at a later time. The pool will also be deleted or be changed to the desired size when the server is restarted.

- You can create memory pools while Adaptive Server is active; no restart is needed for them to take effect. However, Adaptive Server can move only “free” buffers (buffers that are not in use or that do not contain changes that have not been written to disk). When you configure a pool or change its size, Adaptive Server moves as much memory as possible to the pool and prints an informational message showing the requested size and the actual size of the pool. After a restart of Adaptive Server, all pools are created at the configured size.
- Some `dbcc` commands and drop table perform only logical page size I/O. `dbcc checkstorage` can perform large I/O, and `dbcc checkdb` performs large I/O on tables and logical page size I/O on indexes.
- Most Adaptive Servers perform best with I/O configured for transactions logs that is twice the logical page size. Adaptive Server uses the default I/O size of twice the logical page size if the default cache or a cache with a transaction log bound to it is configured with a memory pool twice the logical page size. Otherwise, it uses the logical page size memory pool.
- You can increase the default log I/O size for a database using the `sp_logiosize` system procedure. However, the I/O size you specify must have memory pools of the same size in the cache bound to the transaction log. If not, Adaptive Server uses the logical page size memory pools.

Wash percentage

- The default value for the wash size is computed as follows:
 - If the pool size is less than 300MB, the default wash size is set to 20 percent of the buffers in the pool
 - If the pool size is greater than 300MB, the default wash size is 20 percent of the number of buffers in 300MB
- The minimum setting for the wash size is 10 buffers, and the maximum setting is 80 percent of the size of the pool.

- Each memory pool contains a wash area at the least recently used (LRU) end of the chain of buffers in that pool. Once dirty pages (pages that have been changed while in cache) move into the wash area, Adaptive Server initiates asynchronous writes on these pages. The wash area must be large enough so that pages can be written to disk before they reach the LRU end of the pool. Performance suffers when Adaptive Server needs to wait for clean buffers.

The default percentage, placing 20 percent of the buffers in the wash area, is sufficient for most applications. If you are using an extremely large memory pool, and your applications have a very high data modification rate, you may want to increase the size to 1 or 2 percent of the pool. Run `sp_sysmon` to look for recommendations, or contact Sybase Technical Support for more information about choosing an effective wash size.

Local asynchronous prefetch percentage

- The default value for a pool's asynchronous prefetch percentage is set by the configuration parameter `global async prefetch limit`. The pool limit always overrides the global limit.
- To disable prefetch in a pool (if the global limit is a nonzero number), set the pool's limit to 0.
- See the *Performance and Tuning Guide* for information on the performance impact of changes to the asynchronous prefetch limit.

Permissions

Only a System Administrator can execute `sp_poolconfig` to reconfigure memory pools within data caches. Any user can use `sp_poolconfig` to get information about memory pools.

Auditing

Values in event and extrainfo columns from the `sysaudits` table are:

Event	Audit option	Command or access audited	Information in extrainfo
38	<code>exec_procedure</code>	Execution of a procedure	<ul style="list-style-type: none"> • <i>Roles</i> – Current active roles • <i>Keywords or options</i> – NULL • <i>Previous value</i> – NULL • <i>Current value</i> – NULL • <i>Other information</i> – All input parameters • <i>Proxy information</i> – Original login name, if set proxy in effect

See also

System procedures `sp_cacheconfig`, `sp_cacheconfig`, `sp_helpcache`, `sp_helpcache`, `sp_logiosize`, `sp_unbindcache`, `sp_unbindcache_all`

8.4 Installation guide

This section describes updates and clarifications in the installation guide for your platform.

8.4.1 Determining devices for a database

[CR# 407828] The installation guide gives incorrect information on how to determine which device a database is using. The infocenter.sybase.com search function may show many hits but it does not show the documents where the hits occur. See Appendix 2 of the installation guide for complete details.

Workaround:

- In versions of Adaptive Server Enterprise prior to 15.0, you use `sysdevices` to determine which device has a low through high virtual page range that includes the `vstart`. The device fragment whose `vstart` you used is on that device.
- Select the `vdevno` from `sysusages` matching the `dbid` retrieved in step 1.

9. Technical support

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you have any questions about this installation or if you need assistance during the installation process, ask the designated person to contact Sybase Technical Support or the Sybase subsidiary in your area.

10. Other sources of information

Use the Sybase Getting Started CD, the SyBooks CD, and the Sybase Product Manuals Web site to learn more about your product:

- The Getting Started CD contains release bulletins and installation guides in PDF format, and also contains other documents or updated information not included on the SyBooks CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader, which you can download at no charge from the Adobe Web site using a link provided on the CD.
- The SyBooks CD contains product manuals and is included with your software. The Eclipse-based SyBooks browser allows you to access the manuals in an easy-to-use, HTML-based format.

Some documentation can be provided in PDF format, which you can access through the PDF directory on the SyBooks CD. To read or print the PDF files, you need Adobe Acrobat Reader.

Refer to the *SyBooks Installation Guide* on the Getting Started CD, or the *README.txt* file on the SyBooks CD for instructions on installing and starting SyBooks.

- The Sybase Product Manuals Web site is an online version of the SyBooks CD that you can access using a standard Web browser. In addition to product manuals, you will find links to EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

To access the Sybase Product Manuals Web site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

10.1 Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ Finding the latest information on product certifications

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click Certification Report.
- 3 In the Certification Report filter select a product, platform, and time frame and then click Go.
- 4 Click a Certification Report title to display the report.

❖ Finding the latest information on component certifications

- 1 Point your Web browser to Availability and Certification Reports at <http://certification.sybase.com/>.
- 2 Either select the product family and product under Search by Base Product; or select the platform and product under Search by Platform.
- 3 Select Search to display the availability and certification report for the selection.

❖ Creating a personalized view of the Sybase Web site (including support pages)

Set up a MySybase profile. MySybase is a free service that allows you to create a personalized view of Sybase Web pages.

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click MySybase and create a MySybase profile.

10.2 Sybase EBFs and software maintenance

❖ Finding the latest information on EBFs and software maintenance

- 1 Point your Web browser to the Sybase Support Page at <http://www.sybase.com/support>.
- 2 Select EBFs/Maintenance. If prompted, enter your MySybase user name and password.
- 3 Select a product.
- 4 Specify a time frame and click Go. A list of EBF/Maintenance releases is displayed.

Padlock icons indicate that you do not have download authorization for certain EBF/Maintenance releases because you are not registered as a Technical Support Contact. If you have not registered, but have valid information provided by your Sybase representative or through your support contract, click Edit Roles to add the “Technical Support Contact” role to your MySybase profile.

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

11. Accessibility features

This document is available in an HTML version that is specialized for accessibility. You can navigate the HTML with an adaptive technology such as a screen reader, or view it with a screen enlarger.

Adaptive Server Enterprise HTML documentation have been tested for compliance with U.S. government Section 508 Accessibility requirements. 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 might 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 Sybase Accessibility at <http://www.sybase.com/accessibility>. The Sybase Accessibility site includes links to information on Section 508 and W3C standards.