

Release Bulletin Adaptive Server[®] Enterprise Version 15.0 for HP-UX

Document ID: DC74500-01-1500-03

Last revised: October 27, 2005

Topic	Page
1. Accessing current release bulletin information	3
2. Product summary	3
2.1 Installation kit	3
2.2 Operating system requirements	4
3. Sybase Software Asset Management instructions	4
3.1 Monitoring Adaptive Server for SySAM issues	4
3.2 The SySAM network license server	6
4. Special installation instructions	7
5. Special upgrade instructions	7
5.1 Upgrading compiled objects	7
5.2 Upgrading high availability and Cluster support	7
6. Product and platform interoperability	12
7. Changes that may affect existing applications	13
7.1 Installation, upgrade and migration	13
7.2 Changes that will affect application behavior	13
7.3 Changes that affect database administration operations	17
7.4 Open Client and Open Server changes	21
8. Known problems	21
8.1 Highlighted known installation issues	21
8.2 Highlighted known Job Scheduler issues	22
8.3 Highlighted known ASE Replicator problems	23
8.4 Highlighted known XML problems	23

Copyright 1987-2005 by Sybase, Inc. All rights reserved. Sybase, the 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, Adaptive Warehouse, Afaria, Answers Anywhere, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-Translator, APT-Library, AvantGo Mobile Delivery, AvantGo Mobile Inspection, AvantGo Mobile Marketing Channel, AvantGo Mobile Pharma, AvantGo Mobile Sales, AvantGo Pylon, AvantGo Pylon Application Server, AvantGo Pylon Conduit, AvantGo Pylon PIM Server, AvantGo Pylon Pro, Backup Server, BizTracker, ClearConnect, Client Library, Client Services, Convoy DM, Copernicus, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DataWindow .NET, DB-Library, dbQueue, Developers Workbench, DirectConnect, DirectConnect Anywhere, Distribution Director, e-ADK, E-Anywhere, e-Biz Impact, e-Biz Integrator, E-Whatever, EC Gateway, ECMAP, ECRTP, eFulfillment Accelerator, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, eProcurement Accelerator, EWA, Financial Fusion, Financial Fusion Server, Gateway Manager, GlobalFIX, iAnywhere, iAnywhere Solutions, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, M2M Anywhere, Mach Desktop, Mail Anywhere Studio, Mainframe Connect, Maintenance Express, Manage Anywhere Studio, M-Business Channel, M-Business Network, M-Business Server, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, mFolio, Mirror Activator, MySupport, Net-Gateway, Net-Library, New Era of Networks, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, PB-Gen, PC APT Execute, PC DB-Net, PC Net Library, PocketBuilder, Pocket PowerBuilder, Power++, powerstop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, QAnywhere, Rapport, RemoteWare, RepConnector, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Report-Execute, Report Workbench, Resource Manager, RFID Anywhere, RW-DisplayLib, RW-Library, S-Designer, SDF, Search Anywhere, Secure SQL Server, Secure SQL Toolkit, Security Guardian, SKLS, smart.partners, smart.parts, smart.script, SOA Anywhere, SQL Advantage, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Ed-iTPTU, SQL Everywhere, SQL Modeler, SQL Remote, SQL Server, SQL Server Manager, SQL SMART, SQL Toolkit, SQL Server/CTT, SQL Server/DBM, SQL Server/SMP SubAgent, SQL Station, SQL STEP, SupportNow, S.W.I.F.T. Message Format Libraries, Sybase Central, Sybase Client/Server Interfaces, Sybase Financial Server, Sybase Gateways, Sybase IQ, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SybFlex, SyBooks, System 10, System 11, System XI (Iogo), SystemTools, Tabular Data Stream, TradeForce, Transact-SQL, Translation Toolkit, UltraLite, UltraLite.NET, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, 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, and XP Server are trademarks of Sybase, Inc. 06/05

Topic	Page
8.5 Highlighted known Web Services problems	24
8.6 Highlighted known Interactive SQL and ASE plug-in issues	24
8.7 Statistics in system tables	27
8.8 Running checkstorage on an upgraded master device may report faults with the syscharsets catalog	28
8.9 Monitor GUI does not start when the LDAP server is used	28
8.10 A stacktrace occurs when inserting into a table with identity column using select distinct	28
8.11 SySAM licenses are needed before upgrading Adaptive Server	28
8.12 Cross platform loads	29
8.13 Running dbcc checktable or dbcc checkdb before bring a database online	29
8.14 Recovering from an upgrade during an online database command	29
8.15 The configuration parameter allow backward scan does not work	29
8.16 Modifying dbccdb schema for large identifiers	30
8.17 Drop definition time table corruption	30
8.18 Manual upgrade with named cache fails	30
8.19 Running diskrefit to recover a corrupt master device may report errors	31
8.20 Using an equi-join clause between two different length columns	31
8.21 set statistics io does not display I/O generated by worker processes	32
8.22 alter table unpartition errors	32
8.23 Edit and Load windows	32
8.24 IPv6 support causes Adaptive Server to hang	32
8.25 Replication Server compatibility issues	33
8.26 Blank spaces in an ldap server entry	33
8.27 cis connect timeout and enable SNMP are not implemented	33
8.28 Using jConnect and Adaptive Server	33
8.29 Handling multibyte character sets during migration	33
9. Product compatibilities	34
9.1 Known compatibility issues	35
10. Technical support	36

Topic	Page
11. Other sources of information	36
11.1 Sybase certifications on the Web	36
11.2 Sybase EBFs and software maintenance	37

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 *Installation Guide*.
- 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:
 - *Installation Guide* for your platform

- *Configuration Guide* for your platform
- *Release Bulletin* for your platform

2.2 Operating system requirements

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

- HP-UX 11.11 – patch level B11.11.0412
HP-UX 11.23 – patch level B11.23.0409

3. Sybase Software Asset Management instructions

Starting with Adaptive Server 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 installation and configuration process. You should review the updated SySAM implementation and plan your SySAM deployment before installing Adaptive Server.

Warning! Adaptive Server works for an initial period of thirty days without proper SySAM configuration. If a valid license for your configuration is not found within this thirty day grace period, Adaptive Server will shut down.

See *What's New in Adaptive Server version 15.0?* for a brief description of SySAM changes. See the Chapter “Managing Sybase licenses for Adaptive Server: a SySAM overview” in the *Configuration Guide* for your platform for details on SySAM configuration and deployment options. See the Adaptive Server Installation Guide for your platform for information on pre-installation planning and SySAM installation procedures.

3.1 Monitoring Adaptive Server for SySAM issues

If Adaptive Server cannot obtain a license, then SySAM evaluates whether the license can be issued under grace. The grace periods evaluated are described in the SySAM Configuration section of *Configuration Guide*.

3.1.1 If the SySAM license cannot be acquired

If the license cannot be issued under grace, then the Adaptive Server will not start or the optional feature will not be enabled. Such errors are reported in the Adaptive Server error log as:

```
00:00000:00000:2005/07/05 16:09:12.96 kernel Sybase
Licensing: Using licenses from: /usr/u/sybase/
SYSAM-2_0/licenses
00:00000:00000:2005/07/05 16:09:13.06 kernel Sybase
Licensing: Failed to obtain 1 license(s) for ASE_CORE
feature from license file(s) or server(s).
00:00000:00000:2005/07/05 16:09:13.06 kernel Sybase
Licensing: Cannot find license file.
```

This error message indicates Adaptive Server's inability to obtain a valid license, and lists why license could not be obtained. If e-mail notification is enabled, an e-mail with this information is also generated. You must need fix the license failure before Adaptive Server can start, or the optional feature is enabled.

3.1.2 Acquiring SySAM licenses under grace

If the license can be issued under grace, then the issue is logged and Adaptive Server starts, or the optional feature is enabled. The Adaptive Server error log entry in such cases looks like:

```
00:00000:00000:2005/07/05 15:46:08.84 kernel Sybase
Licensing: Using licenses from: /opt/sybase/
SYSAM-2_0/licenses
00:00000:00000:2005/07/05 15:46:08.91 kernel Sybase
Licensing: Checked out graced license for 1 ASE_CORE
(2005.0703) will expire Tue Jul 5 15:47:02 2005.
00:00000:00000:2005/07/05 15:46:08.91 kernel Sybase
Licensing: Failed to obtain 1 license(s) for ASE_CORE
feature from license file(s) or server(s).
00:00000:00000:2005/07/05 15:46:08.91 kernel Sybase
Licensing: Cannot find license file
```

If e-mail notification is enabled, an e-mail with this information is generated. In such a case, Adaptive Server continues to operate normally until the issue causing license failure is fixed, or grace period ends. The error message and email notification indicate the date and time the grace period is scheduled to end. The error log entry and e-mail notifications are repeated periodically while Adaptive Server operates in grace period with increasing frequency. Adaptive Server continues to operate normally during the grace period, but you must fix the issue causing the license failure. When the issue is fixed, Adaptive Server automatically acquires the license and moves from grace period mode into normal mode. If Adaptive Server is unable to obtain the license before end of the grace period, then Adaptive Server shuts down, or optional features be disabled.

Note If Adaptive Server licenses are obtained from a network license server, Adaptive Server continuously performs a “handshake” with the license server. This is called a “heartbeat.” A successfully acquired license could be invalidated during the heartbeat. When this happens, Adaptive Server enters a grace period. This is logged in the Adaptive Server error log similar to what is logged at startup. An e-mail notification is sent, if you have configured your Adaptive Server for this service. You must fix the issue causing license failure. If the issue is not fixed before the grace period ends, Adaptive Server may shutdown or optional feature disabled.

Use the `sp_lmconfig` stored procedure without any arguments to display the current status of various licenses Adaptive Server is using. Licenses that issued under grace are marked “graced.”

3.2 The SySAM network license server

Note the following information about the SySAM network license server:

- 1 The SySAM Network license server cannot be started until there is at least one “Served” license copied into the *licenses* directory. Obtain a served license from SPDC before starting the license server.
- 2 There can be only one SySAM network license server running on a single machine. If you want to install the new license server on a machine that is already running a SySAM 1.0 license server, you must shut down the old license server and migrate the SySAM 1.0 licenses to SySAM 2.0. See the *Configuration Guide* for information about migrating the licenses.

4. Special installation instructions

Adaptive Server Enterprise 15.0 includes new major versions of Adaptive Server and many of the supporting components. All other current Sybase products (for example, Replication Server 12.6 or Enterprise Connect Data Access 12.6) contain earlier version of those same components. Installing Adaptive Server version 15.0 into the same directory with existing products should not impact the existing products. However, installing other products on top of Adaptive Server version 15.0 will likely result in one or more products not working correctly.

Sybase strongly recommends that Adaptive Server version 15.0 be installed into its own directory if possible. Where this is not practical and other products must be installed into the same directory, Adaptive Server version 15.0 should be installed last.

5. Special upgrade instructions

5.1 Upgrading compiled objects

Sybase recommends that as part of the post-upgrade tasks, you should see the “Upgrading compiled objects with dbcc upgrade_object” section in Chapter 6 of the *Adaptive Server version 15.0 Installation Guide* for your platform.

5.2 Upgrading high availability and Cluster support

Adaptive Server version 15.0 supports following cluster platforms for Sybase High Availability configuration:

- HP Unix – MCSG 11.15
- IBM AIX – HACMP 5.2
- Sun Solaris – VCS4.0, SunCluster3.0/3.1
- Linux – VCS4.0
- Win2000 – Cluster Manger 5.0
- Win2003 – Cluster Manger 5.2

Adaptive Server does not support SunCluster2.2 and VCS1.3/2.0 HA agents on Sun Solaris.

If you currently have these clusters configured, upgrade the respective cluster versions to configure Adaptive Server 15.0 for High Availability on Sun/Solaris.

5.2.1 Upgrading the cluster subsystem

The cluster subsystem can be upgraded in two ways:

1 Major upgrade which involves downtime need cluster shutdown restart. In this case you must:

- a Suspend the companionship by running `sp_companion suspend` command as described in *Using Sybase Failover in A High Availability System*.
- b Offline the resource groups of primary and secondary companions on both nodes. Make sure this shuts down both the companion server and their corresponding resource groups are not automatically brought online until the cluster system upgrade is complete.

Otherwise, when a node is booted after or during the cluster system upgrade, automatically all resource groups may get started on that node leading to unnecessary failovers. For example, in SunCluster3.0 you should “unmanage” the resource groups besides “offline” so that the resource groups are not brought online automatically during intermittent restarts of the cluster system upgrade.

- c Upgrade the cluster subsystem per instructions provided by the cluster system vendor.

You may find upgrade options to migrate the current resource groups to the new or desired cluster version. If such an option is not available (or the resource groups get deleted or corrupted), then you will need to recreate the resource groups and configure them appropriately after the cluster system upgrade.

See the Cluster documentation for more information on creating resource groups or migrating resource groups during upgrade.

- d Online the resource groups. This should bring primary and secondary companions on their respective nodes.
- e Resume the companionship by running `sp_companion resume` command as described in *Using Sybase Failover in A High Availability System*.

- 2 A minor upgrade does not involve a total cluster downtime. In this case, a node is failed over to another node and upgraded one at a time. Assuming ASE1 is the primary companion on node N1, and ASE2 is secondary companion running on node N2, upgrade steps are:
 - a Upgrade Primary Companion:
 - 1 Failover ASE1 from N1 to N2. This can be achieved by relocating the primary resource group to N2 or by shutting down ASE1.
 - 2 Upgrade cluster subsystem on N1 (based on the upgrade instructions provided by the vendor).
 - 3 Failback ASE1 from N2 to N1. Refer to corresponding cluster chapter in *Using Sybase Failover in A High Availability System* for more details on Adaptive Server failback.
 - b Upgrade Secondary Companion:
 - 1 If you are using asymmetric configuration:
 - Offline secondary resource group and make sure ASE2 is shutdown. Note that ASE2 is not available during this upgrade.
 - Upgrade cluster subsystem on N2 (based on the upgrade instructions provided by the vendor).
 - Start ASE2 by bringing online secondary resource group on N2.
 - 2 If you are using symmetric configuration, follow the steps in “Upgrade Primary Companion” mentioned above for ASE2:
 - 1 Failover ASE2 from N2 to N1.
 - 2 Upgrade cluster subsystem on N2 (based on the upgrade instructions provided by the vendor).
 - 3 Failback ASE2 from N1 to N2.

5.2.2 Upgrading Adaptive Server configured with high availability

Using the following steps to upgrade high availability-enabled Adaptive Server in active-active configuration:

- 1 Drop the high availability companionship.

In asymmetric configuration, on the secondary server, use `isql`:

```
sp_companion <primary-server-name>, "drop"
go
```

In symmetric onfiguration, run the above command on both servers.

Be sure both servers are in Single Server mode. Use `isql` to verify this on both servers:

```
sp_companion
go
```

- 2 Use the cluster system command to stop monitoring resources associated with Adaptive Server on each cluster node.
- 3 Follow the upgrade steps in the Upgrade Chapter of the *Installation Guide* for your platform to upgrade each of the Adaptive Servers separately.
- 4 Run the new Adaptive Server *installmaster* script against the newly upgraded Adaptive Servers.
- 5 Run the new Adaptive Server *installhasvss* script against the newly upgraded Adaptive Servers.
- 6 In the new Adaptive Server installation area, follow the instructions in *Using Sybase Failover in A High Availability System* to configure the permission and ownership for `$SYBASE/$SYBASE_ASE/bin/sybha` and `$SYBASE/$SYBASE_ASE/install/sybhauser`.
- 7 In the new Adaptive Server installation area, modify high availability-related files such as the *RUN_server_file*, and the *SYBASE.csh* and *SYBASE.sh* files, if those files are required on that cluster platform.
- 8 Reconfigure each resource associated with Adaptive Server depending on platform-specific requirements. For example, on Veritas Cluster, the *HAase* resource properties, the *RUN_server_file*, and *Sybase_home* need to be updated.
- 9 Manually restart Adaptive Server on each cluster node with trace flag 2209. Use the Adaptive Server command-line option `-T2209`.

Warning! Trace flag 2209 must not be used after the Adaptive Server companionship is re-established.

- 10 Use the cluster system command to restart monitoring resources associated with Adaptive Server on each cluster node.
- 11 Re-establish companionship by using `isql` to issue:

```
sp_companion <primary-server-name>, configure
```

```
go
```

If user databases exist on the secondary server, warning message may display on the console. There is no need to take any action. You can safely ignore them. The messages look similar to:

```
Msg 18739, Level 16, State 1:
Server 'svr2', Procedure 'sp_hacmpcfgvrfy', Line 102:
Database 'svr2_db1': a user database exists. Drop this
database and retry the configuration again.
```

- 12 Use the cluster command to offline and then online resource associated with Adaptive Server. Use `isql` to connect to each Adaptive Server and enter the command to verify the correct server companionship:

```
sp_companion
go
```

Console output on primary server is similar to the following:

```
Server 'svr1' is alive and cluster configured.
Server 'svr1' is configured for HA services.
Server 'svr1' is currently in 'Primary normal' mode.
(return status = 0)
```

Console output on secondary server is similar to the following:

```
Server 'svr2' is alive and cluster configured.
Server 'svr2' is configured for HA services.
Server 'svr2' is currently in 'Secondary normal'
mode.
(return status = 0)
```

Use the following steps to upgrade high availability-enabled Adaptive Servers in active-passive configuration:

- 1 Use the cluster system command to offline Adaptive Server.
- 2 Use the cluster system command to stop monitoring resources associated with Adaptive Server.
- 3 Follow the upgrade steps provided in the Upgrade Chapter of the *Installation Guide* for your platform to upgrade the Adaptive Server to the new version.
- 4 Run the new Adaptive Server *installmaster* script against the newly upgraded Adaptive Server.

- 5 In new Adaptive Server installation area, modify the high availability-related files such as the *RUN_server_file*, and the *SYBASE.csh* and *SYBASE.sh* files, if those files are required on that cluster platform.
- 6 Use the cluster system command to reconfigure resource properties associated with Adaptive Server to reflect the new Adaptive Server installation location.
- 7 Manually restart Adaptive Server.
- 8 Use the cluster system command to restart monitoring resource associated with Adaptive Server.
- 9 Use the cluster command to offline and then online resource associated with Adaptive Server on the same node. Use the cluster command to switch resource associated with Adaptive Server to another node and then switch back.

6. Product and platform interoperability

This section shows the interoperability of Adaptive Server, Replication Server, and Open Client/Server (OCS) products across versions. For specific platform or operating system level information, see the respective product certification reports at <http://www.sybase.com/>.

Table 1: Interoperability between Replication Server, Open Client/Server, and Adaptive Server

Adaptive Server 15.0	Adaptive Server/Open Client Server 15.0	Adaptive Server/ Open Client Server 12.5.x	Adaptive Server / Open Client Server 12.0	Replication Server 12.6	Replication Server 12.5
IBM AIX	X	X	X	X	X
Sun Solaris	X	X	X	X	X
Windows 32-bit	X	X	X	X	X
HP-UX	X	X	X	X	X
Linux	X	X	N/A	X	X

Legend

- X = compatibility
- N/A = incompatibility, or the product is not available for that version/platform combination

Note 1 Even though two or more products may 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, and HP-UX are big-endian platforms.

7. Changes that may affect existing applications

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

7.1 Installation, upgrade and migration

Sybase has made the following directory structure changes:

- Adaptive Server, Open Client and Open Services components are in *ASE-15_0* and *OCS-15_0* directory respectively.
- In Adaptive Server version 15.0, all product components use *JRE-1_4* and *JRE-1_3* is removed. *JRE-1_3* is replaced with *JRE-1_4* in the *\$SYBASE/shared* directory.
- The Job Scheduler component (*JS-12_5* in ASE 12.5 release) is now under *ASE-15_0/jobscheduler* directory.
- The licensing component (SYSAM) is now available in *SYSAM-2_0* directory.
- The Web Services offering directory has changed from *WS-12_5* to *WS-15_0*.

7.2 Changes that will affect application behavior

7.2.1 Long identifier changes

- Adaptive Server 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, and index name and so on. Due to the expanded limits, some system tables (catalogs) and built-in functions have been expanded. For variables, “@” count as 1 byte, and the allowed name for it is 254 bytes long.

- Identifier names should be changed with corresponding application changes for binding values. Please make sure that your application is not binding names of identifiers with only 30 bytes (previous limit). This may cause unexpected behavior or Access violation or bus error or General Protection Fault.

7.2.2 Query changes

- Due to changes in parser, some queries may return general syntax error (message 102) instead of Syntax error at line # (message 156).
- Ordering of result sets in Adaptive Server version 15.0 is different unless there is an order by clause in the query.
- Query compilation time is increased as query processing engine looks for more ways to optimize the query to run better.
- See *Query Processing in Adaptive Server version 15.0* for details about query processing in Adaptive Server version 15.0.

7.2.3 Component Integration Services changes

- Functional compensation will not be done as much as in previous releases. Adaptive Server version 15.0 will send the query to the remote server and if the remote server raises error the same is returned to the client.

For example, in Oracle, the function charlength includes blank spaces. Adaptive Server version 15.0 does not include charlength.

- NULL behavior is different in Oracle, ASA and IQ – you must override the ANSI NULL behavior to have it work like Adaptive Server version 15.0
- Adaptive Server version 15.0 no longer pads char null, varchar, binary null and varbinary datatypes for proxy tables.
- Cursors by default are READONLY. Cursor should be declared with FOR UPDATE to update through them.
- An index must be created if updateable cursors are to be used.
- Component Integration Services engines no longer take special measures to re-declare cursors under the covers for back ends that close cursors on END TRAN.
- An explicit BEGIN TRAN and END TRAN must be declared around cursor statements for DB2 servers.

- When creating a column constraint on java, text, image, and unitext types, message 11074 is raised.
- create table, create existing table, or create proxy table statements that contain a location clause will be restricted to be the only statement in a batch. This applies to select into statements that contain a location clause as well.

7.2.4 Error message changes

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

7.2.5 jConnect version 5.5 no longer shipped with Adaptive Server

Starting with Adaptive Server version 15.0, jConnect 5.5 is no longer shipped. Only jConnect 6.05 is shipped. If your application is dependent on jConnect 5.5, Sybase recommends you migrate the application to jConnect 6.05 or use an existing jConnect 5.5 release area.

7.2.6 SQL Remote is no longer shipped on the Adaptive Server PC-Client CD

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

7.2.7 Table changes

- The following MDA tables have new columns:
 - monEngine
 - monCachedObject
 - monProcessObject
- monCachedObject now tracks the 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 will have 0 after upgrade. The column base_partition is obsolete and will have 0 after upgrade.

- The syspartitions table that existed prior to upgrade is renamed to syslices during upgrade. This new table will be empty and unused.
- syscomments gains new columns partitionid
- systabstats gains the following new columns:
 - partitionid
 - plldegree
 - statmoddate
- sysstatistics gains one new column called partitionid.
- There is a new type of object called partition condition object. This object will have a row in sysobjects. 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:

- identityburnmax
- spacestate
- erlchgts
- sysstatistics stores the data change counters with formatid=108. sysstatistics space requirement increases due to the additional rows stored.
- Most system catalogs have been converted to data rows lock scheme. The DDLs however continue to use the table level locks. The row locking of the system catalogs may require an increase in the configuration parameter number of locks, depending on the DDLs in the application.
- The system catalogs converted to datarows lock scheme do not have a clustered index with indid equal to 1. The clustered indices now have index id greater than or equal to 2.

7.3 Changes that affect database administration operations

7.3.1 Usage of system built-ins

Built-in functions that provide space information like `data_pgs`, `reserved_pgs`, `used_pgs`, `ptn_data_pgs` and `rowcnt` have been deprecated. They have been replaced with `data_pages`, `reserved_pages`, `used_pages`, and `row_count`. See the Adaptive Server Reference Manual: Volume 1 for detailed information.

7.3.2 DDL and DML changes

- `alter table... unpartition` is not allowed on table with indexes (use `alter table with 1 partition` to remove the partitions.)
- `alter table on a partitioned table with max parallel degree less than number of partitions` may succeed without raising message 326.
- Creating clustered index on an empty partitioned table will return a new informational message 1936 in Adaptive Server version 15.0.
- During view creation null column names are not allowed.
- In Adaptive Server version 15.0, a `select` that contains an aggregate from a proxy table that is mapped to an RPC with a parameter fails with error message 201.

7.3.3 System stored procedure changes

- The output of `sp_help` *object* has been changed. Specifically, “Data_Located_on_segment” has been removed. “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 are implemented for `sp_helppartition`, `sp_helpsegment` outputs.
- When `sp_who` is invoked, it returns “NULL” instead of blank in the `hostname` column for all system tasks.
- When `sp_who` is invoked, it will result in SQL command (insert or select) in the output. In the prior releases of Adaptive Server version 15.0, `sp_who` always return select for the tasks that execute `sp_who`.
- `sp_objectsegment` now displays the segment information for all the partitions of the table.
- `sp_lock` output has a new column called `partitionid`. This column has a value of 0 now and is reserved for future.
- `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.

7.3.4 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`'s output has changed to print partition ID instead of 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 of "REORG RECLAIM_SPACE <tablename> with RESUME" has changed.
- reorg rebuild will fail 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 are changed to varchar(30) and made nullable. Select from sysprocesses will return NULL for these columns instead of space for system tasks.
- The "first" column in sysindexes is moved to "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 costing of the optimizer now includes CPU cost.
- Parallel costing is done only on base tables/indexes that are greater than 20 pages.
- For information about trace flags and diagnostics information, see the *Query Processor Guide*.

7.3.5 System resources changes

- Adaptive Server version 15.0 uses more procedure cache for several reasons. Most important reasons are:
 - Adaptive Server version 15.0 query processing engine looks for more ways to optimize the query to run better
 - Adaptive Server version 15.0 execution engine avoids materialization of work table and tends to evaluate aggregations in memory as much as possible.
 - The data change counters maintained for the function datachange() allocates 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.

- In Adaptive Server 15.0, a new configuration parameter `max_repartition_degree` has been added. This parameter controls the maximum degree to which an intermediate data stream can be re-partitioned. The default value of this parameter is set to 1, which means that this parameter is not set and is bound by the number of online engines configured for Adaptive Server. For a query with a large number of tables Adaptive Server 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 such resource usage. Also, configuring a larger value for the auxiliary scan descriptor pool may be required.

- Adaptive Server version 15.0 avoid worktable materialization and Adaptive Server version 15.0 may incur more resources in auxiliary scan descriptors.
- During upgrade, `max memory` is increased by Adaptive Server 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 upgrade, 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.

7.3.6 Unpartitioning of user tables during upgrade

During upgrade, Adaptive Server unpartitions any partitioned tables. You are required to repartition the table post-upgrade. The unpartitioning occurs due to the requirement in 15.0 for each partition to have a different `partitionid`. The expensive operation of changing the `partitionid` for each page during upgrade has been avoided and hence the unpartition during upgrade.

7.3.7 Space management changes

- The 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.
- Upgrade requires free space to convert system catalogs to datarows lock scheme. The typical requirement for free space is 125 percent of the largest system catalog. If the upgrade is from a 32 bit to 64 bit binary, an addition space of 55 percent of `sysprocedures` is required.

7.3.8 Changes to transaction dump content from pre-version 15.0 Adaptive Server releases

A transaction dump from 11.9.2, 12.0x, or 12.5x can be loaded atop a database dump or transaction dump which has preceded it in the load sequence.

However, if that transaction dump contains a create index transaction, then load transaction recovery will treat that transaction's sort record as the end of the log. No more log records in that transaction dump will be processed, and no more transaction dumps in the load sequence will be allowed to load. This behavior is much like the point-in-time recovery that is effected by load transaction's `until_time` option. The following message is printed when load transaction terminates due to the above mentioned reason:

```
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.
```

In contrast, any database dump from one of those releases should be successfully loaded by Adaptive Server, 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.

7.4 Open Client and Open Server changes

For information about changes that affect Open Client and Open Server see "Known issues" section of the *Open Client and Open Server Release Bulletin*.

8. Known problems

This section describes known problems and workarounds for Adaptive Server.

8.1 Highlighted known installation issues

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

8.1.1 Creating the uninstaller

[CR #398447] The uninstall program is not created on HP-UX 11.23 platforms.

Workaround: Run the following command to uninstall:

```
$JAVA_HOME/bin/java -classpath
$SYBASE/uninstall/ASESuite/uninstall.jar -Dtemp.dir=/var/tmp
-Ddis.jvm.home=$JAVA_HOME -Dis.jvm.temp=1
-Ddis.launcher.file=$SYBASE/uninstall/ASESuite/uninstall -Xms20m
-Xmx60m run -console
```

Where:

- *\$JAVA_HOME* is the JRE 1.4.X directory
- *\$SYBASE* is the location where you installed Adaptive Server version 15.0.

8.1.2 Uninstalling Adaptive Server

[CR #400959] If you install Adaptive Server 15.0 on top of Replication Server version 12.6, uninstalling Adaptive Server 15.0 does not remove files.

Workaround: There is no workaround for this issue.

8.2 Highlighted known Job Scheduler issues

This section contains information about known Job Scheduler issues and their workarounds.

8.2.1 Modifying Job Scheduler SQL

[CR #400948] SQL Editor and Load buttons on the Add Scheduled Job, Job Command page of the Remote Job Scheduling wizard (accessed from Schedule button) do not allow the user to modify the job SQL.

Workaround: Do not click on these buttons. You are not allowed to edit the automatically generated SQL of a remotely scheduled job. After the job is created, the user may connect to the Job Scheduler server and modify the SQL for the job from within the Scheduled Jobs folder.

8.3 Highlighted known ASE Replicator problems

8.3.1 ASE Replicator 15.0 supports only ASE 12.5 datatypes

ASE Replicator 15.0 does not support the new datatypes that are introduced in Adaptive Server version 15.0: bigint, unsigned bigint, unsigned int, unsigned smallint, unitext, computed columns, XML, and encrypted columns. It also does not support the new large identifiers.

8.4 Highlighted known XML problems

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

[CR #400269] If you try to 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.

8.4.2 Cross platform *bcp* of parsed XML images

[CR #400250, 332012] For parsed XML image values, 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* built-in 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)
```

Then you can use *bcp* or replicate the *xmltext* column from one platform to the other.

8.5 Highlighted known Web Services problems

8.5.1 SSL unavailable

[CR #402082] SSL is unavailable for the ASE Web Services Engine on HP-UX. If SSL is configured using the *configssl* script on HP-UX, the following error appears in the log indicating that SSL is not loaded:

```
ERROR [main] - Could not start SSL -- VM Not recognized
```

There is no impact to the availability of Web services using HTTP or mapping Web services to proxy tables in Adaptive Server Enterprise.

8.5.2 Run scripts fail to find LDAP configuration file

[CR #401115] When the ASE Web Services Engine is started, the *runws* script looks for LDAP configuration information in the *libtcl.cfg* file. If LDAP has not been configured, the following messages appear in the *producer.log* file:

```
INFO [main] - Unable to locate LDAP configuration file-  
libtcl.cfg.  
INFO [main] - java.io.FileNotFoundException: LDAP  
config File does not exist
```

Workaround: These messages are informational and can be ignored if your system is not configured to use LDAP directory services.

8.5.3 Alias names limited to 255 characters

[CR #400949] Alias names may be no longer than 255 characters. If an alias longer than 255 characters is added using the *addalias* option of *sp_webservices*, the alias name is truncated to 255 characters. No notification of this truncation is provided.

Workaround: Do not create alias names longer than 255 characters.

8.6 Highlighted known Interactive SQL and ASE plug-in issues

The following section discusses known Interactive SQL and ASE plug-in issues and their workarounds.

8.6.1 Opening XML files with Interactive SQL

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

Workaround: There is no workaround for this issue.

8.6.2 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 such as `sp_help` with multiple results and messages.

Workaround: Use command-line `isql` instead.

8.6.3 Running a script file with Interactive SQL

[CR #401391] Interactive SQL cannot run script files without a `.sql` extension such as `installpubs2`.

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

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

[CR #400436] Adaptive Server’s connect to `<server_name>` and disconnect commands conflict with Interactive SQL’s own connect and disconnect commands.

Workaround: There is no workaround for this issue.

8.6.5 Interactive SQL’s Make Permanent button

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

Workaround: There is no workaround for this issue.

8.6.6 ASE plug-in shows incorrect status for Job Scheduler task

[CR #401265] When you open the Job Scheduler Administration window, the Job Scheduler task status always displays as stopped, even when the task is running. Because of this incorrect status, the Stop button is disabled.

Workaround: Start and stop Job Scheduler task from `isql`, using the `sp_js_wakeup` procedure.

To start a Job Scheduler task, run: `sp_js_wakeup 'start_js', 1`

To stop a Job Scheduler task, please run: `sp_js_wakeup 'stop_js', 1`

8.6.7 Interactive SQL may not generate plans for SQL statements using go

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

Workaround: Eliminate go between SQL statements.

8.6.8 Connecting ASE plug-in to RepConnect

[CR #400709] The ASE plug-in cannot connect to Adaptive Server if RepConnect 2.5 and ASE plug-in 15.0 are installed in the same *\$\$SYBASE* location.

Workaround: Install RepConnect and ASE plug-in into different *\$\$SYBASE* directories.

8.6.9 Query support in Interactive SQL

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

Workaround: Use isql instead.

8.6.10 ASE plug-in drag and drop server capability

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

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

8.6.11 Finding help in the Options dialog

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

Workaround: Start Interactive SQL online help through Help | Interactive SQL | Help. Open Using Interactive SQL Help and find the relevant section in Interactive SQL dialog boxes.

8.6.12 Unified Agent Framework RMI ports

[CR #398968] The Unified Agent Framework RMI ports other than 9999 are not supported by the ASE plug-in. The ASE plug-in does not allow users to change the default RMI port.

Workaround: Use 9999 as the default port

8.6.13 ASE 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 ASE 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.

8.6.14 Starting Interactive SQL from a command line fails

[CR #402435] When trying to start Interactive SQL from the command line and 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: Edit the script file and change the `-path` command line argument to `-Dpath` to correct the problem. The script file is located at:

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

8.7 Statistics in system tables

[CR #399624] When creating, configuring, or upgrading a server to Adaptive Server version 15.0, the system catalogs do not have the right index statistics or table level statistics.

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

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

For example, run:

```
update index statistics <system catalog name>
go
```

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

[CR #401779] A master device originally created either in 11.9.x or 12.0.x, upon upgrade to Adaptive Server version 15.0 may report errors on the *syscharsets* catalog. These faults are reported when checkstorage is run on the upgraded master device and *sp_dbcc_faultreport* is used to report the faults.

Workaround: This problem is harmless and exists in all previous releases but is not reported by checkstorage.

8.9 Monitor GUI does not start when the LDAP server is used

[CR #400784] When the Monitor Client GUI is started in an environment where LDAP is used instead of the *interfaces* file, the Monitor Client GUI is not able to connect to any servers.

Workaround: Create an *interfaces* file with entries for the servers that are to be monitored and use this with the Monitor Client GUI.

8.10 A stacktrace occurs when inserting into a table with identity column using *select distinct*

[CR #401753] If you have enabled parallel plan, and you insert into a table with an identity column using *select distinct* and one of the *select distinct* list items is an expression or an *implicit convert()* expression, this query may cause a stacktrace. 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 above example query, you can disable parallel plan for it using:

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

8.11 SySAM licenses are needed before upgrading Adaptive Server

[CR #401682] When you are upgrading Adaptive Server from a release earlier than 15.0, Adaptive Server may not correctly calculate the grace period. When this happens, Adaptive Server does not start under “Installation Grace.” This results in a failed upgrade.

Workaround: Before performing an upgrade, ensure that you have obtained suitable licenses from Sybase Product Download Center (SPDC). Check your Adaptive Server 15.0 installation to make sure that Adaptive Server is able to obtain the required licenses. See the SySAM Configuration chapter in *Adaptive Server Configuration Guide* for more details.

8.12 Cross platform loads

[CR #371289] Adaptive Server does not support cross platform dump and load from one Adaptive Server version 15.0 to another. However, pre-15.0 to 15.0 platform dump and load is supported.

Workaround: There is no workaround for this issue.

8.13 Running *dbcc checktable* or *dbcc checkdb* before bring a database online

[CR #356308] Running *dbcc checktable* or *dbcc checkdb* after the load of a database, but before online database, can show errors due to a mismatch in the number of rows in *systabstats* data and index.

Workaround: Do not run *dbcc checktable* or *dbcc checkdb* after the load and before online.

8.14 Recovering from an upgrade during an *online database* command

[CR #401660] In Adaptive Server version 15.0, after load database has loaded a database dump that was created by an Adaptive Server executable of a version preceding 15.0, if there should be a system failure during the Online Database command, then the ensuing boot recovery may see errors while recovering the same database.

Workaround: Reload the database dump.

8.15 The configuration parameter *allow backward scan* does not work

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

Workaround: If you are encountering increased deadlocks due to backward scans, break problematic descending scans into two steps. First, select the required rows into a temporary table in ascending order. Then, select from the temporary table in descending order.

8.16 Modifying dbccdb schema for large identifiers

[CR #400680] The dbccdb database does not support workspace names longer than 30 characters. Further exclusions cannot be defined for tables with names longer than 30 characters.

Workaround: Alter the table dbcc_exclusions in dbccdb to modify the table_name column from varchar(30) to varchar(255) and modify the stored procedure sp_dbcc_run_evaluatedb to modify the @wsname local variable from varchar(30) to varchar(255).

8.17 Drop definition time table corruption

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

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

8.18 Manual upgrade with named cache fails

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

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

8.19 Running diskrefit to recover a corrupt master device may report errors

[CR #399678] Certain operations done one after another, and certain combinations of circumstances, can cause apparent page allocation errors.

If you starts Adaptive Server with "-w master", and when that finishes, restart immediately with "-w model", the master database reports that a page should belong to object 8, syslogs, but instead belongs to some other object, and tempdb may report 806 errors.

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

8.20 Using an equi-join clause between two different length columns

[CR #401336] if you have a equi-join clause between two different length char() and char()/varchar() type columns 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(), unichar() and univarchar() with char_length().

8.21 *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 only displays the parent thread scan, logical I/O, and physical I/O counts.

Workaround: There is no workaround for this issue.

8.22 *alter table unpartition* errors

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

Workaround: Run "*CHECKPOINT <database>*" immediately following an unpartition operation. In addition, a polite shutdown avoids this form of corruption.

8.23 Edit and Load windows

[CR #400948] You cannot edit SQL for a remotely scheduled job. In the ASE plug-in, SQL Editor and Load buttons incorrectly display on the Add Scheduled Job page in the Remote Job Scheduling wizard. Clicking SQL Editor or Load button displays SQL the Edit or Load windows, but they do not function.

Workaround: Cancel the Add Scheduled Job page and close the corresponding SQL Editor or Load windows. After the job is created, you can connect to the Job Scheduler server and modify the SQL for the job from within the Scheduled Jobs folder.

8.24 IPv6 support causes Adaptive Server to hang

[CR #392074] Machines configured for IPv6 can hang.

Workaround: You must add the following line to the */etc/nsswitch.conf* file:

```
ipnodes:      files nis dns
```


8.25 Replication Server compatibility issues

[CR # 382874] 12.6 and earlier version of rs_init are unable to create an RSSD database on Adaptive Server version 15.0.

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

8.26 Blank spaces in an *ldap* server entry

[CR #333307] If you have a blank space after your ldap server entry, dscp defaults to using the interfaces driver and does not connect to an ldap server.

Workaround: When you want to use dscp to make an entry into an ldap server, enable ldap by editing the `SYBASE/OCS-15_0/config/libtcl.cfg` file to add the ldap server you plan to use.

8.27 *cis connect timeout* and *enable SNMP* are not implemented

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

Workaround: Do not use these configuration parameters.

8.28 Using jConnect and Adaptive Server

[CR #354109] jConnect does not support roman8 character set. Because this is the default character set on HP-UX platforms, jConnect clients cannot connect to Adaptive Server.

Workaround: Install the iso_1 character set on Adaptive Server. From the jConnect clients, specify “iso_1” as the character set using the “CharSet” connection property. Please note roman8 and iso_1 character sets are not identical, so there could be conversion errors in some cases, jConnect clients will need to handle such conversion exceptions.

8.29 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.

9. 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 DirectConnect 12.6 for Informix, MSSS, UDB does not support LDAP.

The following components have been tested for compatibility with Adaptive Server 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 6.05
- jConnect for JDBC 5.5
- Open Server™ 12.5.1
- SDK 12.5.1
- Open ServerConnect™ (CICS, IMS/MVS) 4.0
- Open ClientConnect™ (CICS, IMS/MVS) 4.0
- ODBC Driver by Sybase 12.5.1
- OLE DB Provider by Sybase 12.5.1

- ODBC Driver by Sybase 15.0
- OLE DB Provider by Sybase 15.0
- InfoMaker 7.0.3
- Replication Server® 12.6
- Physical Architect 8.0
- XA-Library for CICS/ENCINA/Tuxedo 12.5.1

9.1 Known compatibility issues

9.1.1 Installing Enterprise Connect Data Access (ECDA) or MainframeConnect DirectConnect for z/OS with other Sybase software

Sybase strongly recommends you install ECDA Direct Connect option or Mainframe Connect DirectConnect for z/OS, including Direct Connect Manager, into its own product directory.

9.1.2 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.

9.1.3 Using Job Scheduler templates

Sybase recommends upgrading to the Job Scheduler templates provided with Adaptive Server 15.0. The templates contain an important changes support new functionality and used new built-ins `reserved_pages` and `data_pages`. For information on upgrading the templates, see the Upgrading Adaptive Server, Chapter 5, in the *Installation Guide* for your platform.

Note If you have existing jobs created from templates delivered with a pre-12.5.2 Adaptive Server, you will need to refer to the 12.5.2 Release Bulletin for your platform and follow the upgrade steps there for the Job Scheduler templates. Then you can upgrade Job Scheduler templates using the above information.

10. 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.

11. 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 may also contain 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 may 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/>.

11.1 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 Select Products from the navigation bar on the left.
- 3 Select a product name from the product list and click Go.
- 4 Select the Certification Report filter, specify a time frame, and click Go.
- 5 Click a Certification Report title to display the report.

❖ 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.

11.2 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.

