

Release Bulletin Adaptive Server® Enterprise Version 12.5.2 for HP Tru64 UNIX

Document ID: DC72157-01-1252-02

Last revised: November 10, 2004

Topic	Page
1. Accessing current release bulletin information	3
2. Product summary	3
2.1 Installation kit	4
2.2 Server components	4
2.3 PC-client components	5
2.4 Operating system updates	6
3. Special installation instructions	7
3.1 Running dbcc checkcatalog	7
3.2 Installing Adaptive Server version 12.5.2 with previous releases	7
3.3 General configuration steps for Kerberos	8
3.4 Known installation issues	8
4. Special upgrade instructions	9
4.1 Upgrading Job Scheduler	9
4.2 Using standby_access mode during migration	10
4.3 Upgrade support for Adaptive Server version 12.5.2	10
4.4 Adaptive Server 12.5.2 backward compatibility	11
5. Product and platform interoperability	11
6. New functionality in this version	12
6.1 Non-clustered indexes and the UNIQUE attribute set	12
6.2 PowerTransfer	12
6.3 HP Tru64 UNIX support for asynchronous I/O	12

Copyright 1989-2004 by Sybase, Inc. All rights reserved. Sybase, the Sybase logo, AccelaTrade, 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 Server IQ, Adaptive Warehouse, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-Translator, APT-Library, Backup Server, BizTracker, ClearConnect, Client-Library, Client Services, Convo/DM, Copernicus, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DataWindow.NET, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, 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 Modeller, eProcurement Accelerator, EWA, Financial Fusion, Financial Fusion Server, Gateway Manager, GlobalFIX, iAnywhere, iAnywhere Application Alerts, iAnywhere Mobile Delivery, iAnywhere Mobile Document Viewer, iAnywhere Mobile Inspection, iAnywhere Mobile Marketing Channel, iAnywhere Mobile Pharma, iAnywhere Mobile Sales, iAnywhere Pylon, iAnywhere Pylon Application Server, iAnywhere Pylon Conduit, iAnywhere Pylon PIM Server, iAnywhere Pylon Pro, iAnywhere Solutions, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, Mail Anywhere Studio, MainframeConnect, Maintenance Express, Manage Anywhere Studio, M-Business Channel, M-Business Network, M-Business Server, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, My iAnywhere, My iAnywhere Media Channel, My iAnywhere Mobile Marketing, MySupport, Net-Gateway, Net-Library, New Era of Networks, ObjectConnect, ObjectCycle, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, Open Client, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, Orchestration Studio, PB-Gen, PC APT Execute, PC DB-Net, PC Net Library, PocketBuilder, Pocket PowerBuilder, Power++, power.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, PowerJ, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Rapport, RepConnector, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Report-Execute, Report Workbench, Resource Manager, RW-DisplayLib, RW-Library, S-Designer, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKI/S, smart.partners, smart.parts, smart.script, SQL Advantage, SQL Anywhere, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Edit/TPU, SQL Everywhere, SQL Modeler, SQL Remote, SQL Server, SQL Server Manager, SQL SMART, SQL Toolset, SQL Server/CTF, SQL Server/DBM, SQL Server SNMP SubAgent, SQL Station, SQLJ, STEP, SupportNow, S.W.I.F.T. Message Format Libraries, Sybase Central, Sybase Client/Server Interfaces, Sybase Financial Server, Sybase Gateways, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, TotalFix, TradeForce, Transact-SQL, Translation Toolkit, UltraLite, UltraLite.NET, UNIBOM, Unilib, Unimull, 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 and XP Server are trademarks of Sybase, Inc. 05/04

Topic	Page
7. Changes that may affect existing applications	13
7.1 File location changes	13
7.2 Changes to commands, stored procedures, and global variables	13
7.3 Changes to system behavior	15
7.4 Changes that affect performance	16
7.5 Diagnostic changes	17
7.6 Changes to Job Scheduler	17
7.7 XML and Web Services changes	18
8. SDK compatibility issues	20
8.1 Moving Open Client/Server applications to version 12.5.1	20
9. Known problems	20
9.1 Some database devices do not display from ASE Plug-in	20
9.2 Setting the application name	20
9.3 Running a pre-12.5.1 XML processor in standalone mode	21
9.4 Blank spaces in an ldap server entry	21
9.5 alter table and triggers with if update() clauses	21
9.6 Truncation error using select with power() function and numeric argument	21
9.7 Some messages returned incorrectly on UTF-8 servers	22
9.8 Cross platform bcp of parsed XML images	22
9.9 Starting the Sybase License Manager Utility	22
9.10 Memory allocation failure	23
9.11 Migration of stored procedures	23
9.12 cis connect timeout and enable SNMP are not implemented	23
9.13 XP Server returns incorrect results	23
9.14 Network license daemon support	24
9.15 Configuring connectivity for clients accessing wide data	24
9.16 Upgrading check constraints	24
9.17 Handling multibyte character sets during migration	25
10. Product compatibilities	25
10.1 Known compatibility issues	27
11. Documentation updates and clarifications	28
11.1 Web Services User's Guide	28
11.2 Using Sybase Failover in a High Availability System	29
11.3 Utility Guide	30
11.4 Transact-SQL User's Guide	36

Topic	Page
11.5 Reference Manual: Building Blocks	36
11.6 Reference Manual: Commands	37
11.7 Reference Manual: Procedures	40
11.8 System Administration Guide	45
11.9 Performance and Tuning Guide	46
11.10 XA Interface Integration Guide for CICS, Encina, and TUXEDO	47
12. Technical support	48
13. Other sources of information	48
13.1 Sybase certifications on the Web	48
13.2 Sybase EBFs and software maintenance	49

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 12.5.2. 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.

For more information specific to jConnect™ for JDBC™, jisql, and Ribo, see the jConnect documentation.

2.1 Installation kit

Note The installation guide for Adaptive Server version 12.5.2 is the same installation guide as for Adaptive Server version 12.5.1.

The installation kit includes:

- The server CD – for contents list, see “Server components” on page 4.
- The PC-client CD – contains software client components to be installed on Windows NT, Windows 2000, Windows XP Pro, and Windows 2003 computers.
- *Release Bulletin for Adaptive Server Enterprise Version 12.5.2 for HP Tru64 UNIX* (this document).
- *Installation Guide for Adaptive Server Enterprise 12.5.1 on HP Tru64 UNIX*.

2.2 Server components

The Server CD includes the following components:

- Sybase Servers:
 - Adaptive Server 12.5.2
 - Adaptive Server Enterprise Monitor™ Server 12.5.2
 - Historical Server 12.5.2
- Language Modules 12.5.1
- Adaptive Server Administration Tools
 - Sybase Central™ 4.x
 - Adaptive Server plug-in
 - Job Scheduler Templates and Utilities
 - SQL Advantage
- Software Developer’s Kit (SDK) 12.5.1, which includes the following:
 - Connectivity:
 - Open Client™ 12.5.1
 - Embedded SQL™/C 12.5.1

- Monitor Client Library 12.5.2
- XA Interface for Adaptive Server Distributed Transaction Manager 12.5.1
- jConnect for JDBC 5.5 and related utilities and documentation
- jConnect for JDBC 6.0 and related utilities and documentation

2.3 PC-client components

The PC-client CD includes the following components, which can be installed on Windows:

- Software Development Kit (SDK) 12.5.1, which includes the following:
 - Open Client 12.5.1
 - Embedded SQL/C 12.5.1
 - Embedded SQL/Cobol 12.5.1
 - Monitor Client Library 12.5.1
 - XA Interface for Adaptive Server Distributed Transaction Manager 12.5.1
- Adaptive Server Data Providers:
 - ODBC 4.2
 - OLE DB 2.7
 - ADO.NET 1.1
- Language Modules 12.5.1
- jConnect 5.5 for JDBC and related utilities and documentation
- jConnect 6.0 for JDBC and related utilities and documentation
- Adaptive Server Administration Tools:
 - Sybase Central 4.x
 - Adaptive Server Plug-in
 - Job Scheduler Templates and Utilities
 - SQL Advantage
 - Microsoft Cluster Server Resource Type
- InfoMaker 9.0

- PowerDesigner 10.0
- PowerTransfer 10.0

2.4 Operating system updates

HP Tru64 UNIX 5.0a requires the following operating system patches to run Adaptive Server version 12.5 and later components:

- T64V50A16-C0011700-10650-20010507.tar – 5.0A with Patch Kit 2 (BL16)
- T64V50A1-C0011600-10649-20010507.tar – 5.0A with Patch Kit 1 (BL01)

You can obtain the patches from the HP Tru64 UNIX patch kit Web site at <http://ftp1.support.compaq.com/public/unix/v5.0a/>.

If your operating system requires patches, install them before you install Adaptive Server.

To determine which patches have been installed on your system, see “Viewing installed patches,” next.

For a complete list of patches, contact your operating system representative. Do not use a patch that is earlier than the version suggested for your operating system. Use the patch recommended by the operating system vendor even if it supersedes the patch listed here.

If you plan to use Sybase Failover in a high availability system, install all vendor-recommended operating system and high availability patches. For more information, contact your operating system and high availability vendors.

Viewing installed patches

To display installed patches, enter:

```
/usr/sbin/setld -i
```

To display the operating system version level, enter:

```
uname -a
```

Note The version level should be 464 or higher for Adaptive Server version 12.5 and later.

3. Special installation instructions

This section provides information for installing Adaptive Server version 12.5.2 and later that is not included or should be corrected in the installation guide.

3.1 Running dbcc checkcatalog

When any column value is accessed in a database table, Adaptive Server version 12.5.2 has been enhanced to validate the length of the row and the length of the column value within the row, using information stored within the system catalogs in that database. Any discrepancy is reported in the error log with the message:

```
Invalid column length: <length>. Value must be between  
0 and <max col length> for '<row type>' row with minimum  
row length of <min row length>.
```

This check requires accurate information within the system catalogs. It is important to verify the accuracy of this information before any database activity occurs after the upgrade to Adaptive Server version 12.5.2 to prevent any false alarms.

To verify the information, run `dbcc checkcatalog(<database name>)` for all databases. If any errors are reported, fix them by running `dbcc checkcatalog(<database name>, "fixall")`.

Dump the database after the fixes have been made. If an older dump is loaded into Adaptive Server 12.5.2, then any fixes that `dbcc checkcatalog` has applied to system catalogs is lost by the load. Rerun `dbcc checkcatalog` after the load has completed.

3.2 Installing Adaptive Server version 12.5.2 with previous releases

This release of Adaptive Server can be installed on top of an existing 12.5 or later release without having to do an upgrade. Just shutdown your servers, install the release on top of the existing release area and do not configure any servers, and then start your servers back up. There is no upgrade.

For more information about post installation tasks, see the Technote at the Sybase Web site at <http://my.sybase.com/detail?id=1013610>.

3.3 General configuration steps for Kerberos

For information about installing Kerberos security software, see the whitepaper on the Sybase Web site at <http://www.sybase.com/detail?id=1029260>. This helps you get started with using Kerberos with Adaptive Server, OpenClient, OpenServer and jConnect applications.

3.4 Known installation issues

This section describes known problems or issues that you may encounter during the installation process.

3.4.1 Some language files not installed during silent installation

[CR #347178] When installing the Sybase products using the silent mode and an option file, languages other than `us_english` are not installed in `$SYBASE/locales`.

Workaround: Replace in the newly created options file:

```
-W setupTypes.selectedSetupTypeId=custom
```

With:

```
-W setupTypes.selectedSetupTypeId=
```

3.4.2 Adding a local server name

[CR #348555] As part of the installation, you should have a local server name for the server. If you do not, `sp_sysmon "cachewizard"` does not work.

Workaround: To add a local server, enter:

```
sp_addserver <local servername>, local
```

You must restart the server for this change to take effect.

3.4.3 Windows Exceed terminals display issue

[CR #331564] When you set `DISPLAY` to an Exceed terminal on Windows platforms, some of the text may not be displayed, or cannot be seen in the GUI windows

Workaround: Perform the installation in interactive console mode. To do this, issue:

```
./setup -is:javahome JVM -console
```

4. Special upgrade instructions

This section provides information for upgrading to Adaptive Server version 12.5.2 that is not included or should be corrected in the installation guide.

4.1 Upgrading Job Scheduler

[CR #371574] Because of several changes made to sybmgmtdb, including the univarchar changes, you must run the *installjsdb* script anytime you upgrade Job Scheduler.

❖ Upgrading Job Scheduler

- 1 Disable Job Scheduler before you upgrade Adaptive Server:

```
sp_configure "enable job scheduler", 0
```

- 2 Upgrade Adaptive Server. See the installation guide for details on upgrading.

- 3 Run the *installjsdb* script:

```
isql -Usa -Psa_password -Sservername -i  
$SYBASE_ASE/scripts/installjsdb
```

Note You must have the directory with the location of the isql executable (*\$SYBASE_OCS/bin*) in your path.

The *installjsdb* script looks for the sybmgmtdb database. The script updates Job Scheduler tables and stored procedures.

- 4 Enable Job Scheduler:

```
sp_configure "enable job scheduler", 1
```

- 5 To start Job Scheduler, you can either restart the server or enter the following:

```
use sybmgmtdb  
go  
sp_js_wakeup "start_js",1  
go
```

4.2 Using standby_access mode during migration

[CR #361616] You can use the online database for standby_access command when you are performing upgrades within the same family of interim releases (IR) and ESDs of Adaptive Server, such as from version 12.0.0.7 to version 12.0.0.8.

You cannot, however, use online database for standby_access within the same family of maintenance releases, or different families of Adaptive Server, such as:

- Version 12.5 to 12.5.2 – same family, maintenance release upgrade
- Version 12.5.0.x to 12.5.2 – same family, maintenance release upgrade
- Version 12.5.1.x to 12.5.2 – same family, maintenance release upgrade
- Version 12.0.0.8 to 12.5.2 – upgrade to a different family

You can, however, use the command if you are updating Adaptive Server with a minor maintenance release, such as version 12.0.0.7 to 12.0.0.8.

You can perform dump transaction with standby_access before you upgrade:

```
dump database
dump transaction with standby_access
```

You cannot, however, then perform the following:

```
load database
load transaction
online for standby_access
```

This causes the standby_access load sequence to break.

Workaround: Perform the load sequence without using the standby_access parameter:

```
load database
load transaction
online
```

4.3 Upgrade support for Adaptive Server version 12.5.2

Sybase supports upgrades from Adaptive Server versions 11.5.x, 11.9.x, and 12.0.x to Adaptive Server version 12.5.2.

4.4 Adaptive Server 12.5.2 backward compatibility

Before using Adaptive Server 12.5.2 features, make sure you read the Technote “Adaptive Server 12.5.1 Backward Compatibility” at the Sybase Web site at <http://www.sybase.com/detail?id=1025879>.

You can downgrade to earlier 12.5.x versions of Adaptive Server, but to do so you must take the precautions described in the technical note *before* using any of the new 12.5.2 features.

5. 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 12.5.2	Adaptive Server 12.0	Open Client/Server 12.5.1	Open Client/Server 12.0	Replication Server 12.6	Replication Server 12.5	Replication Server 12.1
Apple Mac OS X	N/A	N/A	N/A	X	N/A	N/A
HP Tru64 UNIX	X	X	X	X	X	X
HP HP-UX	X	X	X	X	X	X
HP Itanium2 HP-UX	N/A	X	N/A	X	N/A	N/A
IBM AIX	X	X	X	X	X	X
Linux X86 Red Hat Enterprise	N/A	X	N/A	X	X	N/A
Linux Itanium2 Red Hat Enterprise	N/A	X	N/A	N/A	N/A	N/A
Silicon Graphics IRIX	N/A	X	N/A	X	X	N/A
Sun Solaris	X	X	X	X	X	X
Sun Solaris x86	N/A	N/A	N/A	N/A	N/A	N/A
MS Windows 32-bit	X	X	X	X	X	X

Legend

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

Sun Solaris x86	N/A	N/A	N/A	N/A	N/A	N/A
MS Windows 32-bit	X	X	X	X	X	X

Note 1 12.5.2 is a shelf replacement for the 12.5.1 versions of Adaptive Server. Therefore, the 12.5.1 version of Adaptive Server is compatible with 12.5.2, although Sybase strongly recommends that customers upgrade to 12.5.2 as soon as possible.

Note 2 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.

For the latest interoperability report on these products, see the MySybase Web site at <http://my.sybase.com/detail?id=1026087>

6. New functionality in this version

This section describes changed functionality in Adaptive Server version 12.5.2 that is not included in the *What's New in Adaptive Server 12.5.2*.

6.1 Non-clustered indexes and the UNIQUE attribute set

If the index is not unique, use upgrade nstep to make it unique.

6.2 PowerTransfer

PowerTransfer is an extension to PowerDesigner that allows the transfer of data into Adaptive Server. The source database can be any ODBC data source supported by PowerDesigner. PowerTransfer uses the Sybase bulk copy mechanism (bcp) for high-performance inserts into Adaptive Server.

6.3 HP Tru64 UNIX support for asynchronous I/O

HP Tru64 UNIX supports asynchronous I/O on file systems.

7. Changes that may affect existing applications

The following list describes changes for Adaptive Server 12.5.2 that may affect existing applications.

7.1 File location changes

7.1.1 Updated *trusted.txt* file

Previous versions of Adaptive Server included the *trusted.txt* file in `$$SYBASE/$SYBASE_ASE/certificates/`.

This version of Adaptive Server includes an updated version of the *trusted.txt* file in the `$$SYBASE/config/` directory.

7.2 Changes to commands, stored procedures, and global variables

7.2.1 Number of subqueries in a union statement

Earlier versions of Adaptive Server had a maximum value of 16 for the number of subqueries in a single side of a union statement. In Adaptive Server 12.5.2, the maximum number of subqueries within a single side of a union statement is 50.

7.2.2 *online database* command and error number 7411

In earlier versions of Adaptive Server, the *online database* command would stop if it encountered error message 7411.

online database for this release of Adaptive Server no longer considers error number 7411 to be fatal. Instead, *online database* activates the last chance threshold, and advises the user to rebuild the *systhresholds* clustered index using `sp_fixindex`.

7.2.3 Stored procedures must be run in single-user mode

Earlier versions of Adaptive Server allowed you to run the following stored procedures in multiuser mode: `sp_logdevice`, `sp_dropsegment`, and `sp_extendsegment`.

With Adaptive Server, 12.5.2, you must run these stored procedures in single-user mode.

7.2.4 Changes to the *langinstall* and *charset* utilities

The *langinstall* and *charset* utilities were originally written for DBLib but have been rewritten for CTLib, and can now work with LDAP and SSL.

With Adaptive Server 12.5.2, *charset* allows you to update the current server default character set and sort order. These updates are to the *syscharsets* entry only, and become active when you next restart the server. However, when you update the character set definition, you must:

- Specify only the character set—*do not* reference a sort order),
- Answer yes when prompted at the login for the new character set. (for example, if you enter `charset -Usa -P charset.loc utf8`)

langinstall is no longer restricted to installing messages of 255 bytes in length, and can currently manage messages up to 512 bytes in length.

7.2.5 *sp_helptext* no longer truncates trailing spaces

sp_helptext from earlier versions of Adaptive Server trimmed trailing spaces when it selected text from *syscomments*, which could cause problems when creating scripts.

In version 12.5.2, *sp_helptext* does not truncate trailing spaces when it displays text from *syscomments*.

7.2.6 Changes to global variables

Global variables in default statements

You can now use global variables in `default` clauses. For example:

```
create table foo (c1 int, c2 int default @@spid)
```

NULL values for @@servername

Earlier versions of Adaptive Server allowed you to use a NULL value for @@servername with monitoring tables. With Adaptive Server 12.5.2, @@servername cannot be NULL when you are using monitoring tables, *sp_sysmon* cache wizard, or *sp_monitor*.

If you do issue any of these commands with @@servername set to NULL, Adaptive Server issues error message 5629, “remote distributed transaction participant cannot be started.” If this occurs:

- 1 Define a local server name with sp_addserver.
- 2 Restart the server.

You must also perform these steps if you are using the Real Time Messaging Services option. See the *Real Time Messaging Service User’s Guide* for more information.

7.3 Changes to system behavior

7.3.1 status column value for new logins

In previous versions of Adaptive Server, when you created a new login, a new row was added to syslogins with the status column value set to 0.

Adaptive Server 12.5.2 sets the syslogins status column for new logins to a value of 224, which is functionally the same as setting the status column to a value of 0.

7.3.2 Starting Adaptive Server with -T4110

If you start Adaptive Server 12.5.2 with trace flag -T4110, the optimizer favors an index covering the or clause for min or max aggregates involving in (...), or, or not between clauses.

7.3.3 Adaptive Server using column statistics

Adaptive Server may not take into account column statistics for a non-optimizable in clause, which can cause it to select a sub-optimal plan for a query with in or or clauses that join multiple tables.

Start Adaptive Server with trace flag -T301 to use column statistics.

7.3.4 Proxy tables mapped to RPCs

Adaptive Server version 12.5.0.3 allowed RPCs to create temporary objects for the duration of the remote procedure call. Because of this, proxy tables mapped to RPCs could not participate in transactions, and the work done by the RPC could not be rolled back.

Adaptive Server version 12.5.2 allows proxy tables mapped to RPC to participate in transactions by default, which allows Adaptive Server to maintain data integrity with roll backs.

Start Adaptive Server with trace flag 11228 to revert to the version 12.5.0.3 behavior. However, if you use the `insert.. select` syntax with proxy tables and proxy RPCs (for example, `insert proxy_table1 select proxy_rpc`), and the proxy table and proxy RPC both reside on the same remote server, the `insert` portion of the query does not rollback.

7.3.5 Granting and revoking permissions on system tables

Adaptive Server no longer grants or revokes permissions to "public" on system catalogs in the *installmaster* and *installmodel* scripts. For more information about granting and revoking permissions on system tables, see Chapter 10, "Changes to Global Variables, Commands, and Stored Procedures" in *New Functionality in Adaptive Server 12.5.2*.

7.3.6 Cache memory pools in high availability

Cache memory pools are no longer reconfigured for optimal sizes during failover in a high-availability setup.

7.4 Changes that affect performance

7.4.1 Reducing spinlock contention

To reduce spinlock contention on the default data cache, Adaptive Server version 12.5.2 automatically changes the way it determines the number of cache partitions for the default data cache.

If both the global cache partition and local cache partition are set to "default," Adaptive Server version 12.5.2 automatically tunes the number of cache partitions for the default data cache based on the number of engines and the memory available for the default data cache. However, if either of these parameters set to a specific number, Adaptive Server uses that value as the number of local cache partitions for the default data cache and no automatic tuning is performed.

7.5 Diagnostic changes

7.5.1 optdiag includes new statistic

optdiag includes a new statistic, sampling percent last used, which indicates that statistics are gathered with a user-specified sampling percent.

7.6 Changes to Job Scheduler

7.6.1 Script and utility name updates

This version of Job Scheduler removes the following duplicate script files from different *bin* directories; *GenericTempUtil.csh*, *temprem*, *temputil*, *tempimp*, *temptojob*. This version of Job Scheduler also adds Bourne shell scripts for each of the template utilities: import, remove, template-to-job, and general util. The C shell scripts have been renamed to include the *.csh* extension.

7.6.2 Localization changes

***xml* directory adds more localized versions**

This version of Job Scheduler includes localized versions of the English (*en*), Korean (*ko*), Japanese (*ja*), and French (*fr*) language subdirectories in the *%SYBASE/JS-12_5/Templates/xml* directory.

Specifying languages for XML

The *installTemplateXml* and *installTemplateXml.bat* XML installation scripts for Job Scheduler locate XML documents based on the language code English (*en*), Korean (*ko*), simple Chinese (*zh*), Japanese (*ja*), and French (*fr*). You can now specify the language for XML, and the install script finds the appropriate XML files. If you do not specify a language, Job Scheduler uses *en* as the default.

7.6.3 Stopping jobs with Job Scheduler

Job Scheduler allows you to:

- Stop a job after waiting for a specified period of time
- Immediately stop a job

See the Job Scheduler documentation for more information.

7.6.4 Job Scheduler stored objects in unichar

Earlier versions of Job Scheduler stored information about jobs, schedules, scheduled jobs, job templates and job output and history in varchar columns.

This version of Job Scheduler stores information about jobs, schedules, scheduled jobs, job templates and job output and history in unichar. Existing data is converted from varchar to unichar during an upgrade. You may see some errors, which you can ignore, during the upgrade process.

7.7 XML and Web Services changes

7.7.1 XML and subscripts

Adaptive Server 12.5.2 makes the following changes to the way subscripts are handled in XML, to conform to the XPath standard:

- Prohibits reverse ranges, such as “//a [8 to 3]”
- Prohibits negative subscripts, such as “//a [-2]”
- Prohibits subscripts and ranges that are not integer values, such as “//a [1 to b]”

In the Adaptive Server version 12.5.2, these situations provide appropriate error messages.

7.7.2 Treatment of XPath operators, “//” and “/”

In Adaptive Server version 12.5.2, the interpretation of the “//” operator has been changed to reflect the XPath standard.

For example, the following returns the title of each book whose author’s first name is “Mary.”

```
/bookstore/book[author/first-name = "Mary"]/title
```

In Adaptive Server 12.5.1, you could also reference the first name using the “//” operator:

```
/bookstore/book[//first-name = "Mary"]/title
```

Adaptive Server 12.5.1 interpreted the leading “//” operator as a relative reference to all `first-name` elements contained in the current book element. The XPath standard, however, specifies that a leading “/” operator is an absolute reference that references every `first-name` element in the entire document. To reference the `first-name` elements contained in the current book, you must precede the “//” operator with the “.” operator to indicate the current context:

```
/bookstore/book[.//first-name = "Mary"]/title
```

Adaptive Server 12.5.2 adopts the XPath standard interpretation of the “//” operator. Queries using Adaptive Server 12.5.1 that were written in the second example’s form raise an exception error should be rewritten in form of the third example.

7.7.3 Error reporting with invalid *for xml* usage

If you use the `for xml` clause in a context in which it is not allowed (for example, `exec a_stored_proc for xml`), Adaptive Server version 12.5.2 produces error 156, “Incorrect syntax near the keyword <keyword>”.

In the previous versions of Adaptive Server this syntax produced error 152.

7.7.4 samples directory for XML

Adaptive Server version 12.5.2 enhances the XML samples and restructures some directories. The XML and Java samples are in `$$SYBASE/ASE-12_5/sample/XML` and `$$SYBASE/ASE-12_5/sample/Java` respectively.

7.7.5 XML libraries location

In Adaptive Server version 12.5.1, the libraries `jcs.jar` and `sqlx.jar` were located in `$$SYBASE/ASE-12_5/sample`.

In Adaptive Server 12.5.2, they are located in the `$$SYBASE/ASE-12_5/lib` directory.

These files are used for running Java-based XML processors, Java, and XML samples provided with the Sybase installation.

7.7.6 Web Services libraries located in `$SYBASE`

In Adaptive Server version 12.5.1, the libraries *dsparser*, *jconnect* and *jtds* were included as a part of the Web Services jar files located in `$SYBASE/WS-12_5/lib`.

With Adaptive Server version 12.5.2, these libraries are retrieved from the Adaptive Server installation directories, *jConnect-5_5* and *shared-1_0*

8. SDK compatibility issues

8.1 Moving Open Client/Server applications to version 12.5.1

Adaptive Server version 12.5.2 ships the SDK version 12.5.1 plus ESD #2.

8.1.1 Notes for installing version 12.5.1 SDK

If you continue to use your current 12.5 SDK installation, be sure to install Adaptive Server 12.5.2 in a separate directory.

Make sure that your SYBASE and load library path variables (for example, LD_LIBRARY_PATH on Solaris, LIBPATH on IBM AIX or HP Tru64 UNIX, and so on) point to the older SDK installation directory.

9. Known problems

This section describes known problems and workarounds for Adaptive Server.

9.1 Some database devices do not display from ASE Plug-in

[CR #330279] Some database devices are not displayed when ASE Plug-in issues SQL commands to retrieve the database devices information.

Workaround: The database devices retrieval SQL statements are corrected to avoid the Arithmetic overflow exception.

9.2 Setting the application name

[CR #275782] ASE Plug-in does not register the application name on Adaptive Server.

Workaround: Set the host name of the connection. Select the user name when opening LogOnDialog.

9.3 Running a pre-12.5.1 XML processor in standalone mode

[CR #331069] When a pre-12.5.1 XML processor is run in standalone mode, it requires JDK and JRE.1.2.x jar files in the CLASSPATH. Failure to have these JDK/JRE1.2.x files means that the XML processor does not start.

Workaround: Sybase does not recommend running a pre-12.5.1 Adaptive Server XML processor in standalone mode.

9.4 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-12_5/config/libtcl.cfg` file to add the ldap server you plan to use.

9.5 *alter table* and triggers with *if update()* clauses

[CR #199655] If a trigger contains an `if update()` clause, data modifications that should fire the trigger that are performed after executing `alter table add`, `alter table drop`, `alter table lock`, or `alter table modify` may cause errors in column references. Triggers on the altered table that use an `if update()` clause in the body of the trigger to reference a column may not fire, or may fire incorrectly.

Workaround: After the `alter table` operation has completed, drop and re-create all triggers on the altered table. Doing so causes the `if update()` clause in the triggers to correctly reference the new columns by their new column offsets, so the trigger code executes correctly.

9.6 Truncation error using *select with power()* function and numeric argument

[CR #236723] If you use the `select` command with the function `power()`, and use a numeric argument, the result is a truncation error.

Workaround: Convert the numeric argument to a float.

9.7 Some messages returned incorrectly on UTF-8 servers

[CR #271646] When the Adaptive Server default character set is UTF-8, the ASE Replicator procedures `sp_addreplicatart` and `sp_addreplicateconn` return some error messages incorrectly.

This problem does not affect the function of these procedures.

Note Even though the error messages are not returned correctly, the errors are logged correctly in the ASE Replicator log files.

Workaround: There is no workaround for this issue.

9.8 Cross platform *bcp* of parsed XML images

[CR #332012] One 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.

9.9 Starting the Sybase License Manager Utility

[CR #351387] If the `DISPLAY` environment variable is not set, executing `$$SYBASE/SYSAM-1_0/bin/lmgr` to start the Sybase License Manager Utility causes errors.

Workaround: You must set the `DISPLAY` environment variable by executing:

In Bourne shell:

```
set DISPLAY=<host name>:0.0
export DISPLAY
```

In C shell:

```
setenv DISPLAY <host name>:0.0
```

Where *host name* is the name of the system where you want the installer screens to display.

9.10 Memory allocation failure

[CR #350856] If there is insufficient memory in the heap class Adaptive Server hangs.

Workaround: Double the heap class size.

9.11 Migration of stored procedures

[CR #328131] `sybmigrate` does not migrate stored procedures that have a space character in their name, or belong to a procedure group.

Workaround: There is no workaround for this issue.

9.12 *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.

9.13 XP Server returns incorrect results

[CR #280073] On all systems, using an extended stored procedure (ESP) with parameters may result in wrong data being sent to XP Server. Because of this, XP Server may return incorrect results.

This happens when there is no parameter declaration list in the ESP, but one or more parameters are supplied at execution time. Before CR #280073, all parameters supplied at execution time were sent to the XP Server, causing incorrect results. CR #280073 fixes this problem by silently ignoring all parameters supplied at execution time when the parameter declaration list is empty.

This fix does not cause any behavior change when the parameter declaration list in the ESP matches the parameters supplied at execution time.

If you are using ESPs without a parameter declaration list, and you supply parameters at execution time, drop and re-create the ESP in Adaptive Server. Make sure the parameter declaration list in the ESP matches the parameter list supplied at execution time

You can use trace flag 523 to revert to pre-CR280073 behavior. This is a server-wide trace flag to be set when you restart Adaptive Server.

9.14 Network license daemon support

Adaptive Server version 12.5.2 and higher has upgraded the SySAM daemons and libraries (FlexLM) to version 8. The SySAM daemons (lmgrd and SYBASE) are backward compatible, however the libraries are not. As a result, Adaptive Server 12.5.2 cannot check out licenses from earlier versions of the SySAM license daemons.

Workaround: Sites using the Network License Model should make sure that the central SySAM license daemons are upgraded first.

9.15 Configuring connectivity for clients accessing wide data

[CR #271822] If you are using wide data (that is, data other than text and image that is greater than 255 bytes), and you are using the bulk library APIs, make sure you configure the connectivity capability for wide data. Failure to do so may result in memory corruption.

9.16 Upgrading check constraints

[CR #272941] When you upgrade to Adaptive Server 12.5.0.1 or later, check constraints created with the following criteria. They may not have been upgraded correctly.

- The check constraint was created as a named check constraint.
- Quoted identifiers were set to on while the constraint was created.
- Quote identifiers were used around the entire constraint name.
- The constraint name was also a key word, or not a legal identifier.

You see parser syntax errors after you have upgraded to Adaptive Server 12.5.0.1 and later, and when you first add data to the constrained column, or when you run `dbcc upgrade_object`.

To determine whether this problem occurs on your upgraded databases, run this command to check for syntax errors on your constraints and rules:

```
dbcc upgrade_object (dbname, rule)
```

Workaround: There are two workarounds to this problem:

- Issue `alter table drop constraint_name` to drop a constraint that was not upgraded successfully. Then issue `alter table add constraint_name`, and re-create the constraint.
- As System Administrator, manually update `syscomments.text` to add quoted identifiers to the constraints that were not successfully upgraded.

9.17 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.

10. Product compatibilities

This section lists the Sybase components that are compatible with Adaptive Server Enterprise version 12.5 and later. For information about operating system requirements, see individual component documentation.

Note DirectConnect 12.5 for Informix, MSSS, UDB does not support LDAP.

Note Component Integration Services and SSL are not supported on Adaptive Server version 12.5.2.

The following components have been tested for compatibility with Adaptive Server 12.5 and later:

- Sybase Character Sets 3.0
- Sybase Central viewer 4.x

- ECDA option for ODBC 12.5
- ECDA option for Informix 12.5
- MainframeConnect DirectConnect for OS/390 12.5
- ECDA option for Oracle 12.5
- ECDA option for Microsoft SQL Server 12.5
- ECDA option for DB2 Universal Database 12.5
- ECDA option for AS/400 12.5
- OpenSwitch 12.5

Note OpenSwitch 12.5 does not support SSL.

- jConnect for JDBC 4.5
- jConnect for JDBC 5.5
- Open Server™ 12.5
- SDK 12.5
- Open ServerConnect™ (CICS, IMS/MVS) 4.0
- Open ClientConnect™ (CICS, IMS/MVS) 4.0
- ODBC Driver for Adaptive Server 12.5
- OLE DB Provider for Adaptive Server 12.5
- InfoMaker 7.0.3
- Replication Server® 11.5.1, 12.0
- SQL Remote 7.0.02
- Physical Architect 8.0
- XA-Library for CICS/ENCINA 11.1.1
- XA-Server for Tuxedo 11.1.1

Note Backup Server 12.5 is compatible with Backup Server 11.9.2.1 and 12.0. Either server can be local or remote. Backup Server 12.5 is compatible with versions of Backup Server earlier than 11.9.2.1 *only* if the Backup Server 12.5 is the remote server.

10.1 Known compatibility issues

10.1.1 Using Job Scheduler templates

Sybase recommends upgrading to Adaptive Server 12.5.2 Job Scheduler templates and template stored procedures. The templates for this release have been improved and internationalized. To upgrade, install the new template XML from the `$Sybase/Sybase/JS-12_5/Templates/xml` directory, and the new template stored procedures from the `$Sybase/Sybase/JS-12_5/Templates/procs` directory. For more information on installing the templates, see the *Job Scheduler User's Guide*. Particular combinations of unlike template and server versions have the following constraints:

- You can:
 - Run a job on Adaptive Server 12.5.2 that was created from an Adaptive Server 12.5.1 template.
 - Install and use the Adaptive Server 12.5.2 template stored procedures with the 12.5.1 template XML.
- You cannot create a new job from:
 - Adaptive Server 12.5.1 Job Scheduler using Adaptive Server 12.5.2 template XML.
 - Adaptive Server 12.5.2 Job Scheduler using Adaptive Server 12.5.1 template XML.

10.1.2 ASIQ does not start

[CR #346774] Uninstalling Adaptive Server 12.5.1 from `$$SYBASE` where ASIQ 12.5 also exists breaks ASIQ, and ASIQ cannot to start.

Workaround: Do not install Adaptive Server and ASIQ in the same `$$SYBASE` directory.

10.1.3 Installing ECDA

[CR #347617] Installing ECDA version 12.5 over Adaptive Server version 12.5 causes Adaptive Server not to work.

Workaround: Install ECDA into a separate `$$SYBASE` release area.

11. Documentation updates and clarifications

This section describes changes and additions to the Adaptive Server version 12.5.1 documentation set.

11.1 Web Services User's Guide

The following are updates for Adaptive Server version 12.5.2 for *Adaptive Server Enterprise Web Services User's Guide*.

11.1.1 Remote server class definition setting

Include the following in the “Troubleshooting Issues” section of Chapter 7, “Troubleshooting” in the *Adaptive Server Enterprise Web Services User's Guide*.

Issue The `sp_webservices` add command may return the following error when generating proxy tables:

```
Warning: Row size (3347 bytes) could exceed row size limit, which is 1962
bytes.
Msg 208, Level 16, State 1:
Server 'JMALVARADO', Line 1:
tempdb..ws_4338e6e122cd4ef0a not found. Specify owner.objectname or uses to
check whether the object exists (sp_help may produce lots of output).
No proxy tables were created for the WSDL URL:
[http://www.xignite.com/xquotes.asmx?WSDL]
(return status = 0)
```

This error occurs because the remote server representing the Web Services Consumer has been added using `sp_addserver` with a class other than “sds.” To verify that this is so, use `sp_helpserver` in isql:

```
1> sp_helpserver ws
2> go
```

Here, `ws` is the name of the Web Services Consumer. This is the default. The remote server class is returned in the indicated column of the result:

```
name network_name class ...
---- -
ws ws null ...
```

User action Change the class of the remote server to “sds” by using `sp_dropserver` and `sp_addserver` in isql:

```
1> sp_dropserver ws_name
```

```

2> go
...
1> sp_addserver ws_name, sds, ws_name
2> go

```

Here, *ws_name* is the name chosen for the Web Services Consumer.

11.2 Using Sybase Failover in a High Availability System

The following are updates for Adaptive Server version 12.5.2 for *Using Sybase Failover in a High Availability System*

11.2.1 Configuring Adaptive Server resource groups

Step 2 of the section titled “Configuring Adaptive Server resource groups” in Chapter 11, “Active-Active Configuration for Sun Cluster 3.0” and in Chapter 12, “Active-Passive Configuration for Sun Cluster 3.0” is replaced with:

If you use another file at a different location, specify the full path for the resource extension property *Dataserver_login_file* when configuring the *SYase* resource. Create or edit a file that contains Adaptive Server login information for system administrator and the user you added for the fault monitor. The default file is

```
/$SYBASE/$SYBASE_ASE/SC-3_0/etc/ase_login_file.
```

The file consists of two lines; the first line is the login and password of the System Administrator, and the second line is the login and password of *monitor_user*. The fault monitoring program, *ase_monitor*, performs the thorough probe as user *monitor_user*.

```

login_type <tab> login_string
login_type <tab> login_string

```

Valid values for login type are “encrypted” and “normal”. If you set *login_type* to “normal”, the value of the *login_string* is in the form “login_name/password”. If you set *login_type* to “encrypted”, the value of *login_string* is the encrypted string you get from the *haisql* utility (located in */\$SYBASE/\$SYBASE_ASE/bin*). Sybase recommends usage of “encrypted” *login_type* so the sensitive information in the file is well protected. To use *haisql* to generate the encrypted login string:

- 1 Run *haisql* with no arguments to generate the encrypted string for a given *login_name* and *password*:

```

/$SYBASE/ASE-12_5/bin/haisql
Enter Username: sa

```

Enter Password:

```
TWAS8n1jSF2gBsvayUlW97861.cyTKaS1YhavBRQ2qKcJwtx.TmFBarGS2K1553WDR7g8m5
vrf86t@K4CU62HEccm4zkeexsP9E=FeuvX
```

2 Copy and then paste the encrypted string to the *ase_login_file* file.

The following is an example of the *ase_login_file* using the “encrypted” login type:

```
encrypted
TWAS8n1jSF2gBsvayUlW97861.cyTKaS1YhavBRQ2qKcJwtx.TmFBarGS2K1553WDR7g8m5vrf86t
@K4CU62HEccm4zkeexsP9E=FeuvX
encrypted
rX2S8n1jSF2gBuD0q=AXEXKCZvzGcK5K3kWnp_P+e4avf=67kYVSzy7+h640@97FSP_dlkH_oV2zi
ma5+7tUyHnsm4zmSIHIUnKSTPoTD
```

The following is an example of the *ase_login_file* file using “normal” login type:

```
normal    sa/sa_password
normal    monitor_user/monitor_user_password
```

Note The two lines of the *ase_login_file* may use different login types.

You should protect the *ase_login_file* file with proper access permissions, particularly if you are not using the encrypted login strings. Perform the following to make the file readable only to the root user after editing the file with proper *login_type* and *login_string* values:

```
chmod 400 ase_login_file
chown root ase_login_file
chgrp sys ase_login_file
```

11.3 Utility Guide

The following are updates for Adaptive Server version 12.5.2 for the *Utility Guide*.

11.3.1 bcp modes

Correction In the section, “bcp modes” in Chapter 3, “Using bcp to Transfer Data to and from Adaptive Server” of the *Utility Guide*, the description should read:

bcp in works in one of two modes:

- Slow bcp – logs each row insert that it makes, used for tables that have one or more indexes or triggers.
- Fast bcp – logs only page allocations, copying data into tables without indexes or triggers at the fastest speed possible.

Version 12.5.1 of the book erroneously reversed the descriptions.

11.3.2 ddlgen

ddlgen, the Java-based tool that generates definitions for server- and database-level objects in Adaptive Server, has expanded its list of filters to include indexes, triggers, and constraints out of table and database definitions in the DDL of database- as well as table-level objects. The valid filters are:

-F filters out indexes, triggers, and constraints out of table and database definitions in the DDL of table- and database-level objects. The valid filters are:

- *For tables* – [% | I | RI | KC | TR]
- *For databases* – [% | SGM | GRP | USR | R | D | UDD | U | V | P | XP | I | RI | KC | TR]

The filter options are:

Filter option	Filters out:
%	Everything, and retrieves the schema-only definition of a database or table.
SGM	Segments
GRP	Groups
USR	Users
R	Rules
D	Defaults
UDD	User-defined datatypes
U	User tables
V	Views
P	Stored procedures
XP	Extended stored procedures
I	Indexes
RI	Referential integrity constraints
KC	Primary- and unique-key constraints
TR	Triggers

Examples

Generates DDL for the authors table in the pubs2 database, filtering for all indexes (I), and referential integrity constraints (RI), primary and unique key constraints (KC), and triggers (TR) from the DDL of a table:

```
ddlgen -Uroy -Proyl23 -TU -Nauthors -Dpubs2 -F%
```

Alternatively, you can specify each of the filters individually:

```
ddlgen -Ulogin -Ppassword -TU -Ndbname.owner.table
-FI,RI,KC,TR
```

Note When you use the -F parameter to filter a table, then indexes, referential integrity, key constraints and triggers automatically get filtered, as they are a subset of table object.

The following generates the definition of *table_name* while filtering out foreign keys and primary-unique keys:

```
ddlgen -Ulogin -Ppassword -TU -Ntable_name -Ddbname
-FRI,KC
```

Both of these generate foreign keys for a specified user in the entire database:

```
ddlgen -Ulogin -Ppassword -TRI -N%.%.% -Ddbname
```

Or:

```
ddlgen -Ulogin -Ppassword -TRI -Ndbname%.%.%
```

Both of these generate DDL for the primary and unique keys of all the tables in a database that begin with “PK”:

```
ddlgen -Ulogin -Ppassword -TKC -Ndbname.%.%.PK%
```

Or:

```
ddlgen -Ulogin -Ppassword -TKC -N%.%.PK% -Ddbname
```

Note Although you can normally generate all indexes only for a table, the -T object type parameter with the RI and KC value allows you to generate foreign keys as well as primary and unique keys for an entire database.

The following generates schema-only definition of a database:

```
ddlgen -Ulogin -Ppassword -Sserver:port -TDB -Ndbname
-F%
```

Alternatively, you can specify each of the filters individually:

```
ddlgen -Ulogin -Ppassword -Sserver:port -TDB -Ndbname
```

-FSGM, GRP, USR, R, D, UDD, V, P, XP, I, RI, KC, TR

The following generates the database DDL skipping the compiled object:

```
ddlgen -Ulogin -Ppassword -Sserver:port -TDB -Ndbname
-FTR, D, XP, V, R
```

The following generates database definition without a table definition:

```
ddlgen -Ulogin -Ppassword -Sserver:port -TDB -Ndbname
-FU
```

ddlgen update

Servers The existing description of the -S parameter in ddlgen is:

- *-Shost_name : port_number*

specifies the host name or IP address of Adaptive Server, as well as its port number. Separate *host_name* and *port_number* with a colon, without spaces before or after it.

Note You must use the -S option because ddlgen does not connect to a default server.

The updated description is as follows:

- *-Sserver* – specifies the name of the Adaptive Server. ddlgen looks this name up in the interfaces file or LDAP configuration file. If you specify:
 - *-S [host_name:port_number]* – ddlgen uses the *host_name* and *port_number* provided, and neither interfaces nor LDAP configuration files are read or parsed.
 - *-S[server] -l* – ddlgen parses the interfaces file specified at the user location for the server name (see the -l parameter description for more information).
 - *-S[server]* – without specifying an interfaces file, ddlgen:
 - 1 Tries to read the LDAP configuration file from the standard location.
 - 2 If the LDAP file does not exist, or exists but does not contain an Adaptive Server entry, then the interfaces file is parsed at its standard location for the server name.

- 3 If the LDAP file exists, then `ddlgen` uses it to search the server name. The interfaces file is not parsed, and the LDAP configuration file is parsed.

Separate *host_name* and *port_number* with a colon, without spaces before or after it.

Note You must use the `-S` option because `ddlgen` does not connect to a default server.

Interfaces files `-I` – specifies the interfaces file name, and corresponds to `$$SYBASE/interfaces` for UNIX, and `$$SYBASE%\ini\sql.ini` for Windows. Use this optional parameter with `-S`.

Indexes You can now run the `ddlgen` utility for all indexes of all tables in a database in the following command:

```
ddlgen -Ulogin -Ppassword -Sserver:port -TI
-Ndbname.%.%.%
```

For example, this generates DDL for all indexes for all tables in the `pubs2` database:

```
ddlgen -Usa -P -SHARBOR:1955 -TI -Npubs2.%.%.%
```

11.3.3 sp_sproc_columns

There is a new column in the result set for the `sp_sproc_columns` catalog stored procedure called `mode`:

Column	Datatype	Description
<code>mode</code>	<code>varchar(20)</code>	The parameter mode information stored in <code>syscolumns</code> that contains: <ul style="list-style-type: none"> • <i>For SQL procedures</i> – in, out, or “return value”. • <i>For SQLJ procedures (Java)</i> – in, out, inout, or “return value”.

11.3.4 sqldbgr

Add the following to the description of the `sqldbgr` SQL debugger utility:

Note You do not have the ability to view `sqldbgr` version strings.

11.3.5 *dataserver* clarification

In the *dataserver* reference page in the *Utility Guide*, Example 2 is incorrect. It currently reads:

Example 2 rewrites a corrupt model database:

```
dataserver -d my_master_device -w model
```

The correct example is:

Example 2 rewrites a corrupt model database:

```
dataserver -d d_master -w model -s server_name
```

11.3.6 Omission from *sybmigrate*

The following information was omitted from the “*sybmigrate*” Chapter, in the *Utility Guide*.

sybmigrate requires allow resource limits to be set to 0.

Note If metadata already exists on the target Adaptive Server, you cannot migrate server data.

SYBASE_JRE – defines the location of the Java runtime environment. This is generally set to *\$SYBASE/shared-1_0/jre-1_3* in the Adaptive Server release area. This environment variable overrides *JAVA_HOME*.

SYBASE_JRE defaults to *\$SYBASE/shared-1_0/jre-1_3*.

11.3.7 Utilities moved

The following three utilities have moved from *\$SYBASE/\$SYBASE_ASE/bin* to *\$SYBASE/\$SYBASE_OCS/bin*:

- certauth
- certpk12
- certreq

11.3.8 *bcp* clarification

In the section “*bcp* modes” in Chapter 3, “Using *bcp* to Transfer Data to and from Adaptive Server” in the *Utility Guide*, the definitions of *bcp* in were reversed and incorrectly stated as:

bcp in works in one of two modes:

- Fast bcp – logs each row insert that it makes, used for tables that have one or more indexes or triggers.
- Slow bcp – logs only page allocations, copying data into tables without indexes or triggers at the fastest speed possible.

The correct definitions are as follows:

bcp in works in one of two modes:

- Slow bcp – logs each row insert that it makes, used for tables that have one or more indexes or triggers.
- Fast bcp – logs only page allocations, copying data into tables without indexes or triggers at the fastest speed possible.

11.4 Transact-SQL User’s Guide

In “Sorting Query results” in Chapter 3, the word “aggregates” has been dropped from this sentence:

“Adaptive Server does not allow subqueries, variables and constant expressions in the order by list.” A successful example in the following section, order by and group by, uses the aggregate function avg(price) in an order by clause:

```
order by avg(price)
```

11.5 Reference Manual: Building Blocks

11.5.1 *str_replace*

The *str_replace* function now accepts NULL in the third parameter and treats it as an attempt to replace *string_expression2* with NULL, effectively turning *str_replace* into a “string cut” operation.

For example, the following returns “abcghijklm”:

```
str_replace("abcdefghijklm", "def", NULL)
```

11.5.2 @@version_as_integer update

In Chapter 3, “Global Variables” in *Reference Manual: Building Blocks*, the definition for @@version_as_integer states:

“Returns the version of the current release of Adaptive Server as an integer.”

This information should read:

“Returns the number of the last upgrade version of the current release of Adaptive Server as an integer. For example, @@version_as_integer returns 12500 if you are running Adaptive Server version 12.5, 12.5.0.3, or 12.5.1.”

11.5.3 set_appcontext update

In the reference page for set_appcontext in Chapter 2, “Transact-SQL Functions” in *Reference Manual: Building Blocks*, the parameter for *attribute_value* is incorrectly stated as being a char(2048). The correct datatype is char(30), and should appear as follows:

attribute_value – is a row that specifies and application attribute value. It is saved as the datatype char(30).

11.6 Reference Manual: Commands

11.6.1 where clause syntax

The following syntax is no longer applicable in the where clause.

```
where [not] time_period1 overlaps time_period2
```

The following example in the where clause reference pages is also no longer applicable:

Example 9 Compares two time periods and determines whether they overlap each other. The first period begins March 16, 1994 and lasts for one month. The second period begins March 31, 1994 and lasts until December 31, 1994. The predicate returns a value of TRUE because the two periods have points in common:

```
where (date "1994-03-16", interval +"1" month) overlaps
      (date "1994-03-31", date "1994-12-31")
```

11.6.2 dbcc upgrade object has a new parameter “check”

The following description of the dbcc check command should be included in *Reference Manual: Commands*.

Syntax

```
dbcc upgrade_object( {dbid | dbname},
                    { compiled_object_name |
                    "view" | "procedure" | "default" |
```

```
"rule" | "default" | "trigger" },
["force" | "check" ])
```

Keywords and options	check – checks syntax for the specified <i>compiled_object</i> in syscomments in the specified database. Does not raise errors on occurrences of select.
Comments	<ul style="list-style-type: none"> • check is used to detect syscomments text corruption which was caused by Adaptive Server defects that occurred before the upgrade. This syscomments text corruption is serious because it causes upgrade to fail. • If any error is reported by <i>dbcc upgrade_object</i> with check option, you must drop and re-create the <i>compiled_object</i>.

11.6.3 *create index* clarification

The create index reference page in Chapter 1, “Commands” of *Reference Manual: Commands* incorrectly states that the *max_rows_per_page* and *fillfactor* parameters are exclusive.

“*max_rows_per_page* – *max_rows_per_page* and *fillfactor* are mutually exclusive. Unlike *fillfactor*, the *max_rows_per_page* value is maintained until it is changed with *sp_chgattribute*.”

It should read:

max_rows_per_page – limits the number of rows on data pages and the leaf level pages of indexes. Unlike *fillfactor*, the *max_rows_per_page* value is maintained until it is changed with *sp_chgattribute*.

11.6.4 *bcp* reference page clarification

In the *bcp* reference page in Chapter 8, “Utility Commands Reference” of the *Reference Manual: Commands*, the description of the *-b* batchsize parameter has been clarified. It used to read:

“*-b* batchsize is the number of rows per batch of data copied (default is to copy all the rows in one batch).”

The text now reads:

“*-b* batchsize is the number of rows per batch of data copied. By default, *bcp* in copies *n* rows in one batch, where *n* is equal to the batch size.”

11.6.5 *set* command option *explicit_transaction_require* not supported

The *set* command option *explicit_transaction_require* is not currently supported, and should not appear in the *Adaptive Server Reference Guide* version 12.5.1.

11.6.6 New *set* parameters *proc_output_params* and *set proc_return_status*

The following new parameters should be included in the *set* command reference section of *Reference Manual: Commands*.

Parameters

- *proc_output_params* – controls sending of output parameters that a stored procedure generates back to the client. *set proc_output_params off* suppresses sending the output parameters back to the client. The default for this parameter is on.
- *proc_return_status* – controls sending of a return status TDS token back to the client. *set proc_return_status off* suppresses sending the return status token to the client, and *isql* client does not display the (*return status = 0*) message. The default for this parameter is on.

Warning! If the client application that executes a procedure relies on the success or failure of the procedure based on the return status, then do not use the *set proc_return_status off* option.

Example 1

Suppresses the output of both parameters and the return status TDS token:

```
1> set proc_output_params OFF
2> go

1> set proc_return_status OFF
2> go

1> declare @x int
2> select @x = 2
3> exec sp_pout @x output
4> print "Value of @x returned from sproc is: %!", @x
5> go

(1 row affected)
Value of @x returned from sproc is: 3
(1 row affected)
```

In addition, you can also suppress the lines reporting the number of rows affected to generate output with no extra messages using the *set nocount on* option before running this batch.

Example 2

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: 2
```

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

Return parameters:

```
-----
                2
```

11.7 Reference Manual: Procedures**11.7.1 sp_autoformat**

`sp_autoformat` is a stored procedure introduced in Adaptive Server version 12.5.1.

Description	A utility stored procedure that produces readable result set data, <code>sp_autoformat</code> reformats the width of variable-length character data to display only non-blank characters. Trailing blanks are truncated in the output.
Syntax	<code>sp_autoformat @fulltabname [, @selectlist , @whereclause , @orderby]</code>
Parameters	<p><code>@fulltabname</code></p> <p>specifies the name of table from which data is being selected. Use owner names if the object owner is not the user running the command.</p> <p><code>@selectlist</code></p>

specifies the comma-separated list of columns to be selected in the result set. Columns in the table can be renamed using the `<name> = <column>` notation. See examples. If `@selectlist` is not provided, all columns in the table specified are output in column ID order.

@whereclause

is a search predicate, specified as a where clause, that filters out rows from the table being selected.

@orderby

is an optional order by clause that specifies the order in which the output result set is presented.

Examples

Example 1 Returns a result set from a select statement similar to `select id, colid, name from syscolumns where id = 3`, where the character columns are autoformatted:

```
1> sp_autoformat "syscolumns", "id, colid, name", "where id = 3"
2> go
```

id	colid	name
-----	-----	-----
3	1	id
3	2	number
3	3	colid
3	4	status
3	5	type
3	6	length
3	7	offset
3	8	usertype
3	9	cdefault
3	10	domain
3	11	name
3	12	printfmt
3	13	prec
3	14	scale
3	15	remote_type
3	16	remote_name
3	17	xstatus
3	18	xtype
3	19	xdbid
3	21	accessrule
3	22	status2

Example 2 Renames the output columns using the following syntax:

```
[ < AS-Name label of Column> ][*]=[ ]*<column name>
```

<AS-Name label of Column> is optional, and you can use white spaces around the = separator:

```
1> sp_autoformat syscolumns, "'Object Id' = id, 'Column Name'=name,
   'Column ID'=colid", "where id = 3"
2> go
```

Object Id	Column Name	Column ID
-----	-----	-----
3	id	1
3	number	2
3	colid	3
3	status	4
3	type	5
3	length	6
3	offset	7
3	usertype	8
3	cdefault	9
3	domain	10
3	name	11
3	printfmt	12
3	prec	13
3	scale	14
3	remote_type	15
3	remote_name	16
3	xstatus	17
3	xtype	18
3	xdbid	19
3	accessrule	21
3	status2	22

(1 row affected)

Example 3 Uses the @orderby parameter to specify an ordering in the result output:

```
sp_autoformat @fulltabname = 'syscolumns',
              @selectlist = "id, name",
              @orderby = "ORDER BY name"
```

Example 4 To generate an autoformatted result when you select from multiple tables, or if you have a complex SQL select statement with expressions in the select list, you must:

- 1 Use temporary tables to generate the result set:

The following generates the list of the columns with the highest column ID on all system catalogs:

```

select o.id, o.name, c.colid, c.name
from sysobjects o, syscolumns c
where o.id < 100 and o.id = c.id
      and c.colid = (select max(c2.colid) from syscolumns c2
                    where c2.id = c.id)
order by o.name

```

The following generates the same result set with auto-formatting of character data using a temporary table to produce readable output, and includes minor changes to provide column names in the temporary table:

```

select o.id, ObjectName = o.name, c.colid, ColumnName = c.name
into #result
from sysobjects o, syscolumns c
where o.id < 100 and o.id = c.id
      and c.colid = (select max(c2.colid) from syscolumns c2
                    where c2.id = c.id)

```

2 Use sp_autoformat on that temporary table to produce formatted output:

The order by clause in the original select statement is skipped when generating the temporary table, and is instead added to the call to sp_autoformat when generating the output result.

```

1> exec sp_autoformat @fulltablename = #result, @orderby = "order by
      ObjectName"

```

```

2> go

```

id	ObjectName	colid	ColumnName
11	sysalternates	2	altsuid
21	sysattributes	13	comments
55	syscertificates	6	suid
45	syscharsets	8	sortfile
3	syscolumns	22	status2
6	syscomments	8	status
37	sysconfigures	9	value4
17	sysconstraints	7	spare2
38	syscurconfigs	15	type
30	sysdatabases	19	status4
12	sysdepends	10	readobj
35	sysdevices	7	mirrorname
43	sysengines	12	starttime
...			

```
(1 row affected)
(return status = 0)
```

You can further process the temporary table to report only on the required output for selected tables, as shown below:

```
1> exec sp_autoformat #result, "id, 'Object Name' = ObjectName,
   'Column Name' = ColumnName", "where id < 5"
2> go
```

id	Object Name	Column Name
-----	-----	-----
1	sysobjects	loginame
2	sysindexes	crdate
3	syscolumns	status2
4	systypes	accessrule

Usage

- `sp_autoformat` looks for an object only in the current database. To use `sp_autoformat` on temporary tables, issue the procedure from `tempdb`.
- `sp_autoformat` does not validate that the columns referenced in any of the parameters actually exist in the table specified by the `@fulltablename` parameter. `sp_autoformat` fails if you reference any nonexistent columns.
- Provide only one instance of a column in the select list.

Return codes

- 0 – successful completion
- 1 – internal error, or usage error in invocation
- Other – any other errors raised by Adaptive Server during the execution of the generated SQL statement are returned back to the caller.

Restrictions

- `sp_autoformat` uses internal SQL variables to generate SQL statements that are then executed using `execute immediate`. The length of the generated SQL statement is limited to 2K bytes. Auto-formatting result sets for a large column list, or columns with long names can sometimes cause an error due to insufficient size of the buffer for the generated SQL statement.
- Quoted identifiers are not supported for either the table or column names. If you have result sets that use quoted identifiers and that need autoformatting, you must:
 - a Generate the required data in a temporary table, where the columns in the temporary table do not have any quoted identifiers.

- b Use `sp_autoformat` to produce the required output using the temporary table.
- c Rename the columns in the `@selectlist` in the desired output format.

Permissions

Any user can execute `sp_autoformat`. However, users selecting from the tables must have appropriate select privileges.

11.7.2 sp_helpdb

The `sp_helpdb` stored procedure now includes the `order` parameter. Add the following to the reference page for `sp_helpdb` in *Reference Manual: Procedures*. The full syntax of `sp_helpdb` is:

```
sp_helpdb [dbname, order]
```

The default order of the output is by `lstart`, which is the order in which the databases were created or altered. Use `device_name` along with `dbname` to display the output of `sp_helpdb` ordered by `device_name`.

Example Specifies `device_name` for the `order` parameter to display the device fragments for `mydb` in alphabetical order, overriding the default sort order of `sp_helpdb`.

```
sp_helpdb mydb, device_name
```

name	db_size	owner	dbid	created	status
mydb	4.5 MB	sa	5	Feb 27, 2003	no options set

(1 row affected)

device_fragments	size	usage	created	free kbytes
A	1.5 MB	data only	Feb 27 2003 7:50AM	1530
B	1.0 MB	log only	Feb 27 2003 7:50AM	not applicable
C	2.0 MB	data only	Feb 27 2003 7:50AM	846

11.8 System Administration Guide**11.8.1 Table information incomplete**

In the *System Administration Guide*, Chapter 12, “Auditing,” Table 12-6, Values in event and extrainfo columns, is incomplete.

Table 12-6 should include these lines for the create index entry under:

“Other Information: Index name”

11.8.2 Size information incorrect in procedure

In the *System Administration Guide*, Chapter 26, “Checking Database Consistency,” some size information is incorrect in the “Creating the dbccdb database” procedure. The procedure should read:

```
* In step 1, output:
Recommended size for dbccdb is 23MB (data = 21MB, log
= 2MB)
* In step 3, user input:
...
size = "4096"
...
size="1024"
* In step 7 user input:
use dbccdb|
go|
sp_dbcc_createws dbccdb, scanseg, scan_pubs2, scan,
"64K"|
sp_dbccvcreatews dbccdb, textseg, text_pubs2, text,
"64K"|
```

11.9 Performance and Tuning Guide

11.9.1 OAM-scans on DOL tables

OAM-pages are used when performing a table scan on a DOL-table. These OAM-pages refer to allocation-units, which must then be scanned to find the extents on the pages for the table being scanned. When the scan is started, a list of pages is generated the information in the allocation-page. Any newly allocated pages produced by another process in this allocation-unit are not seen by the scanning process as they were not allocated when the scan began. This is due to the way DOL-tables are implemented. If you want to include the newly allocated pages,

- use isolation level 3 type of scans (select with holdlock) or,
- add an index to the DOL-table which can be used to prevent a table-scan or,
- convert the DOL-table to an APL-table.

11.9.2 Multiple temporary databases

The information on multiple tempdbs was inadvertently included in the *Performance and Tuning Guide*. This information is being moved to the *System Administration Guide* in the next update.

Login bindings to multiple tempdbs

In prior versions only a user with the “sa” role could be bound to a tempdb.

As of this version of Adaptive Server, any user can also be bound to an application at the time of login. For example:

If there is a binding for login “user1” and for application isql, when “user1” logs in with application isql the “user1” binding determines the tempdb assigned to the session, not the isql binding.

11.9.3 Optimizer and Abstract Plans

Correction: In Chapter 11, “Multiple Temporary Databases,” the “Unchanged yet notable procedures” section noted the following for the sp_changedowner stored procedure:

You cannot map system databases, including tempdb to a default location. However, you can map user-created temporary databases.

This information is incorrect.

11.10 XA Interface Integration Guide for CICS, Encina, and TUXEDO

In Chapter 4, “Application Programming Guidelines in the *XA Interface Integration Guide for CICS, Encina, and TUXEDO*,” the introductory paragraph for the section, “Tightly coupled transactions,” has been revised. The new text is emphasized.

11.10.1 Tightly coupled transactions

The XA environment treats each thread or process that works on a transaction as a transaction branch. Each transaction branch is assigned a different xid and works independently of the other branches. However, all branches are committed or rolled back as a unit. *This applies to MTS/COM+ environments only, and only if you are using ctlib-based drivers that are older than version 3.6.*

12. 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.

13. Other sources of information

Use the Sybase Getting Started CD, the Sybase Technical Library CD, and the Technical Library 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 Technical Library CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader (downloadable at no charge from the Adobe Web site, using a link provided on the CD).
- The Technical Library CD contains product manuals and is included with your software. The DynaText reader (included on the Technical Library CD) allows you to access technical information about your product in an easy-to-use format.

Refer to the *Technical Library Installation Guide* in your documentation package for instructions on installing and starting the Technical Library.

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

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

13.1 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.

13.2 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.

