

Release Bulletin Adaptive Server Enterprise Version 12.5.1 for Sun Solaris

Document ID: DC73430-01-1251-02

Last revised: October 15, 2003

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	5
3. Special installation instructions	6
3.1 Known installation issues	6
4. Special upgrade instructions	7
4.1 Change in versioning of SDK components	8
4.2 Adaptive Server 12.5.1 backward compatibility	8
5. Product and platform interoperability	9
6. New functionality in this version	10
6.2 sysconfig is not Section 508 compliant	10
6.3 PowerTransfer	10
7. SDK compatibility issues	11
7.1 Moving Open Client/Server applications to version 12.5.1	11
8. Known problems	12
8.1 Highlighted known LDAP-related problems	12
8.2 alter table and triggers with if update() clauses	13
8.3 Truncation error using select with power() function and numeric argument	13

Copyright 1989-2003 by Sybase, Inc. All rights reserved. Sybase, the Sybase logo, AccellaTrade, 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-FORMS, APT-Translator, APT-Library, AvantGo, AvantGo Application Alerts, AvantGo Mobile Delivery, AvantGo Mobile Document Viewer, AvantGo Mobile Inspection, AvantGo Mobile Marketing Channel, AvantGo Mobile Sales, AvantGo Mobile Pharma, AvantGo Pylon, AvantGo Pylon Application Server, AvantGo Pylon Conduit, AvantGo Pylon PM 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, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, Distribution Director, e-ADK, E-Anywhere, 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, 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 AvantGo, My AvantGo Media Channel, My AvantGo Mobile Marketing, 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 Library, Open ServerConnect, Open Solutions, Optima++, PB-Gen, PC APT Execute, 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, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RW-DisplayLib, S-Designer, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILLS, smart.partners, 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 CFT, 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 MPT, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, Syber-Assist, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, TradeForce, TransactSQL, Translation Toolkit, UltraLite.NET, UNIBOM, Unibit, Unimult, Unisep, Unistring, URK Runtime Kit for UniCode, Viewer, Visual Components, VisualSpeller, 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. 03/03

Topic	Page
8.4 Highlighted known ASE Replicator problems	13
8.5 Highlighted known XML problems	18
8.6 Default values not always set correctly in the Job Scheduler templete creation wizard	20
8.7 SQL derived tables cannot have unnamed columns	21
8.9 Reconfiguring from utf8 with external sort order	21
8.10 Versions strings for Open Client/Server are not correct	22
8.11 Backup Server sybmultbuf mismatch error	23
8.12 Queries containing a select statement with a group by statement	23
8.13 Job Scheduler Agent cannot terminate a job on request	23
8.14 Job timeout in Job Scheduler not currently available	24
8.15 XML identifiers that cannot be used as element names	24
8.16 sybmigrate error messages	25
8.17 XA might hang when used with Open Client/Server	25
8.18 Enhanced Full-Text Search installation	25
8.19 Errors in queries with derived tables	25
8.20 SQL Advantage queries hanging	26
8.21 Views must have defined column names	26
8.22 Migrating between single-byte and multibyte charactersets	26
8.23 Migration of stored procedures	26
8.24 cis connect timeout and enable SNMP do not work	26
8.25 XP Server returns incorrect results	27
8.26 Network license daemon support	27
8.27 bcp and abstract data types	27
8.28 Referencing Java data in a nonrecovered database	28
8.29 Configuring connectivity for clients accessing wide data	28
8.30 Upgrading check constraints	28
8.31 Handling multibyte character sets during migration	29
8.32 Adaptive Server plug-in default character set and language	29
8.33 Dynamic and host variables not allowed in DDL statements	29
9. Product compatibilities	30
9.1 Replication Server compatibility	31
10. Documentation updates and clarifications	31
10.1 Utility Guide	32
10.2 Transact-SQL User's Guide	34
10.3 Reference Manual: Building Blocks	34

Topic	Page
10.4 Reference Manual: Commands	35
10.5 System Administration Guide	38
10.6 Performance and Tuning Guide	39
11. Technical support	39
12. Other sources of information	40
12.1 Sybase certifications on the Web	40
12.2 Sybase EBFs and software maintenance	41

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.

- v **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 12.5.1. 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.

This release bulletin provides the latest information about Adaptive Server and its known problems, as well as documentation updates and clarifications.

2.1 Installation kit

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, XP Pro, and Windows 2003 computers.
- *Release Bulletin for Adaptive Server Enterprise Version 12.5.1 for Sun Solaris* (this document).
- *Installation Guide for Sybase Adaptive Server Enterprise Version 12.5.1 for Sun Solaris*.

2.2 Server components

The Server CD includes the following components:

- Adaptive Server 12.5.1
- Adaptive Server Enterprise Monitor™ Server 12.5.1
- Monitor Client Library 12.5.1
- Historical Server 12.5.1
- Software Developer’s Kit (SDK) 12.5.1 ASE Edition which includes the following:
 - Open Client™ 12.5.1 ASE Edition
 - Embedded SQL™/C 12.5.1 ASE Edition
 - Embedded SQL/Cobol 12.5.1 ASE Edition
- SQL Server Monitor™ API 12.5
- SQL Remote™ 7.0.2
- Language Modules 12.5.1
- jConnect for JDBC 5.5
- Java utilities:
 - jisql 2.0
 - Ribo 2.0
- Sybase Central™ 4.x

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 ASE Edition, which includes the following:
 - Open Client 12.5.1 ASE Edition
 - Embedded SQL/C 12.5.1 ASE Edition
 - Embedded SQL/Cobol 12.5.1 ASE Edition
- Open Client Runtime 12.5.1
- ODBC 4.2
- OLE DB 2.7
- jConnect 5.5 and related documentation
- Java utilities:
 - jisql 2.0
 - Ribo 2.0
- InfoMaker® 9.0
- Language Modules 12.5.1
- XA Interface Library for ASE Distributed Data Manager 12.5.1
- Physical Architect™ 8.0
- SQL Remote 7.0.2
- Monitor Client Library 12.5.1
- PowerDesigner 9.5.2

Note Java Central is installed with PowerDynamo, not Adaptive Server. The only plug-in that comes with Adaptive Server is the Adaptive Server Enterprise plug-in for Sybase Central Java Edition 12.5.1.

2.4 Operating system updates

Solaris 8 requires the following operating system patches to run Adaptive Server 12.5.0.3 components, for both 32-bit and 64-bit versions:

- 108434-04
- 108435-04
- 108528-12
- 108827-12
- 108991-18
- 109210-12
- 109384-02
- 111177-06
- 110648-13
- 112138-01

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

Contact your operating system provider for any patches recommended for your installation. 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.

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 list all currently installed patches and display the operating system version level, enter:

```
showrev -p
```

3. Special installation instructions

This section provides information for installing Adaptive Server version 12.5 and later that is not included or should be corrected in the installation guide.

3.1 Known installation issues

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

3.1.1 Job Scheduler templates scripts must be downloaded from Sybase CodeXchange Web site

To install the Job Scheduler templates, you must download three installation scripts from the Sybase CodeXchange Web site at <http://www.codexchange.sybase.com/>. These files are used to install template stored procedures to target ASEs and to install template XML into the Job Scheduler database.

v **Downloading Job Scheduler templates installation files**

- 1 From the CodeXchange home page at <http://www.codexchange.sybase.com/>, go to: Projects | ASE. Scroll to the bottom of the page to the table listing Subprojects. Select Job Scheduler.
- 2 Go to Project Tools | News.
- 3 View the news item dated October 1, 2003, called JS Template Files Missing from 12.5.1 Release CDs.
- 4 Follow the instructions on the View news item page to download the files.

3.1.2 InstallShield and inaccessible NFS mounted devices

[CR #334942] If you have an inaccessible NFS mounted device, InstallShield may hang while displaying:

```
Performace install checks. Please wait ...
```

If this happens, the console from which InstallShield was executed displays a message similar to:

```
NFS3 server <some server> not responding still trying
```

Workaround: kill InstallShield. Unmount the inaccessible NFS device. Retry the installation.

4. Special upgrade instructions

This section provides information for upgrading to Adaptive Server version 12.5.1 that is not included or should be corrected in the installation guide.

4.1 Change in versioning of SDK components

Starting with version 12.0, the full SDK (Software Developer Kit) product has been shipped as one of the components of Adaptive Server. In addition, it is shipped as its own point product.

With 12.5.1, the SDK point product will be shipped approximately two months after the Adaptive Server 12.5.1 release. To differentiate between the two SDK releases (one as a point product and the other as an Adaptive Server component), the SDK component files have the words “ASE Edition” embedded in their version strings. For example, 'bcp -v' currently displays:

```
Sybase CTBCP Utility/12.5.1 ASE Edition/P-EBF-...
```

Once the SDK point product is released, you may need to install this product over your current installation to use new fixes to the SDK products (isql, bcp, CT-Library, ESQL/C & ESQL/COBOL, and so on). After the “ASE Edition” of SDK has been overlaid with the point product release, the version number will revert to the previously existing standard. For example 'bcp -v' will display:

```
Sybase CTBCP Utility/12.5.1/P-  
EBF11382/DRV.12.5.1/SPARC/Solaris 2.8/BUILD1251-...
```

As Adaptive Server ESDs typically do not contain files from the SDK component, this regular versioning scheme will now resume. The SDK component is a fully functioning product and has the same features and functionality as the point product.

4.2 Adaptive Server 12.5.1 backward compatibility

Before using Adaptive Server 12.5.1 features, make sure you read the Technote “Adaptive Server 12.5.1 Backward Compatibility” at the Sybase Web site at <http://www.sybase.com/support/techdocs>.

You can 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 12.5.1 features.

4.3 HA and upgrading files

- If HA is installed, you must re-run the *installhasvss* script after you have run *installmaster*.
- *sybha* is owned by “root.” Its ownership must be restored to “sybase” for the copy to execute correctly.

4.4 Enhanced Full-Text Search Japanese environment setup

Enhanced Full-Text Search needs specific a Japanese environment setup on Japanese Solaris 2.8. Enhanced Full-Text Search needs to be set to this configuration:

```
language=japanese
charset=sjis
vdkLanguage=japanb
vdkCharset=sjis
```

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 ClientServer, and Adaptive Server

Adaptive Server 12.5.1	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	X	N/A	X	N/A	N/A
HP Tru64 UNIX 5.x	X	X	X	X	X	X
HP HP-UX 11.x	X	X	X	X	X	X
IBM RS 6000 AIX 4.3.x	X	X	X	X	X	X
IBM RS 6000 5.1	X	X	X	X	X	X
Red Hat Enterprise 2.1	N/A	X	N/A	X	X	N/A
Silicon Graphics IRIX	N/A	X	N/A	X	X	N/A
Sun Solaris	X	X	X	X	X	X
Microsoft Windows	X	X	X	X	X	X

Legend

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

Note 1 Open Client/Server 12.5.1 is currently scheduled for release in December 2003 and Replication Server 12.6 is currently scheduled for release in November 2003 (dates may be subject to change.)

Note 2 12.5.1 is a shelf-replacement for the 12.5 versions of Adaptive Server and Open Client/Server. Therefore, the 12.5 version of either product is compatible with 12.5.1, although Sybase strongly recommends that customers upgrade to 12.5.1 as soon as possible.

Note 3 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 12.5.1 that is not included in the *What's New in Adaptive Server 12.5.1*.

6.1 TCP_NODELAY is turned on by default

TCP_NODELAY is turned on by default, with changes made by the user to the configuration file.

6.2 sysconfig is not Section 508 compliant

sysconfig is not Section 508 compliant.

If you are unable to use sysconfig, see the *System Administration Guide* for alternative methods of configuring Adaptive Server.

6.3 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.

7. SDK compatibility issues

7.1 Moving Open Client/Server applications to version 12.5.1

The Adaptive Server Edition Software Developer Kit (SDK) is shipped as part of Adaptive Server version 12.5.1. The 12.5.1 SDK point product is scheduled for release before the end of 2003. It contains new features and bug fixes not included in the Adaptive Server Edition of SDK. In addition, the Adaptive Server Edition does not contain fixes to known bugs since SDK 12.5 ESD #10.

Below is a list of features and major bug fixes that will be included in the SDK 12.5.1 point product, but not in the Adaptive Server Edition of SDK.

The following table lists features not available for the 12.5.1 SDK.

CR number	Description
321477	CTISQL: Makes the display of date/time from isql shorter to avoid wasting space.
322200	KERBEROS: Adds Kerberos to 64-bit Sun Solaris.
325311	KERBEROS: Adds Kerberos security driver with MIT GSS library.
331054	KERBEROS: Adds the ability to load Cybersafe or MIT libraries upon request
325728	KERBEROS: Adds support on Linux

The following CRs are known problems that occur in the Adaptive Server Edition of SDK:

CR number	Description
315862	CTLIB: Reads from unallocated memory causes crash in ct_get_data().
320690	CTLIB: Connection may hang when notification is turned on.
322466	CTLIB: Reentrant libraries return CS_TIMED_OUT instead of CS_QUIET for a zero timeout.
323380	CTLIB: ct_poll(CS_NO_LIMIT) may timeout.
323377	CTLIB: Full asynchronous I/O broken in non-reentrant libraries.
323378	CTLIB: FULL asynchronous I/O causes "Socket I/O possible" in the threaded libraries.

If your application is likely to experience any of the CRs listed above, or requires one of the new features listed in the first table, Sybase advises you to wait for the release of the 12.5.1 SDK point product. In the meantime, you may continue to run your 12.5 Open Client/Server applications from your existing Adaptive Server version 12.5 install location against the newer Adaptive Server 12.5.1 installation.

If the CRs listed above are not likely to adversely impact your applications, you may upgrade your Open Client applications to version 12.5.1. However, Sybase strongly recommends that you install the 12.5.1 SDK point product release as soon as it is available.

7.1.1 Notes for installing version 12.5.1 SDK

Sybase will release Open Server 12.5.1 before the end of 2003, at the same time as the SDK 12.5.1 point product.

If you continue to use your current 12.5 SDK installation, be sure to install Adaptive Server 12.5.1 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.

8. Known problems

This section describes known problems and workarounds for Adaptive Server.

8.1 Highlighted known LDAP-related problems

This section discusses Lightweight Directory Access Protocol (LDAP) issues with Adaptive Server.

8.1.1 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 will be using.

8.1.2 *langinstall* and *charset* utilities do not support LDAP

[CR #233080] The *langinstall* and *charset* utilities are built with DB-Library™, and therefore do not support the use of LDAP.

Workaround: There is no workaround for this issue.

8.2 *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.

8.3 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.

8.4 Highlighted known ASE Replicator problems

This section describes known problems and workarounds for ASE Replicator.

8.4.1 Subscription status incorrect after abnormal shutdown

[CR #270782] The subscription status returned by *sp_helprep* and *sp_helpsub* is incorrect under certain conditions.

This problem occurs after the ASE Replicator process is shut down abnormally, such as by a power failure on the host machine or a process kill command, while a subscription is resumed (*up*), and the remote server associated with that subscription is still down when ASE Replicator is restarted. When ASE Replicator is restarted, the subscription status is actually 0 (*down*), but *sp_helprep* and *sp_helpsub* return the subscription status as 1 (*up*).

Workaround: This is an error only in the status information returned. After the remote server has successfully restarted, resume the subscription normally by using `sp_resumesub`.

8.4.2 No errors with invalid database connection name

[CR #s 270729, 271653, and 271676] The following ASE Replicator procedures do not return an error message when you invoke them with an invalid or null database connection name:

- `sp_helprep`
- `sp_resumeprimaryconn`
- `sp_resumereplicateconn`
- `sp_suspendprimaryconn`
- `sp_suspendreplicateconn`

Workaround: There is no workaround for this issue.

8.4.3 Incorrect error message with invalid database connection name

[CR #269484] The following ASE Replicator procedures return an incorrect error message when you invoke them with an invalid or null database connection name, or an invalid or null connection configuration parameter name:

- `sp_configprimaryconn`
- `sp_configreplicateconn`

Workaround: There is no workaround for this issue.

8.4.4 Incorrect error message with invalid subscription name

[CR #271677] When you invoke `sp_helprep` with the status keyword and an invalid subscription name (`sub=sub_name`), it returns an incorrect error message.

Workaround: There is no workaround for this issue.

8.4.5 Statistics not returned by some commands

[CR #270080] The following ASE Replicator procedures do not return a statistics result set:

- `sp_helprep stats`
- `sp_helprep stats, pri_conns`
- `sp_helprep stats, subs`
- `sp_helprep stats, sub=subname`
- `sp_helpreplicateconn stats`
- `sp_helpprimaryconn stats`

Workaround: To get statistics for a primary connection, use the following command:

```
sp_helpprimaryconn stats
```

To get statistics for a replicate connection, use the following command:

```
sp_helprep stats, rep_conns
```

8.4.6 UNIX permission problems with RUN script creation

[CR #279616] When *aserep.sh* creates the *RUN* script, it sets permissions on the *RUN* script file to allow execution. If *aserep.sh* encounters a problem setting permissions on the *RUN* script file, the following error message is returned:

```
Component message: Problem setting permissions;  
exitValue = 1
```

Workaround: To solve this problem:

- 1 Log in to the operating system with a user ID that has authority to set execute permissions in the ASE Replicator instance subdirectory.
- 2 Execute the *aserep.sh* script to initialize the ASE Replicator instance.

See Chapter 2, “Setting Up and Starting ASE Replicator,” in the *ASE Replicator User’s Guide* for more information on executing the *aserep.sh* script to initialize the ASE Replicator instance.

8.4.7 Publishing multiple tables with the same name not supported

[CR #279148] Adaptive Server allows non-unique table names in a database, so a specific table must be identified by a qualified object name. ASE Replicator does not fully support qualified object names when identifying primary tables to be published.

If you use owner-qualification in `sp_addprimaryart`, the correct table (identified by its qualified object name) is published, but you cannot publish additional tables that have the same name in that database.

Workaround: There is no workaround for this issue.

8.4.8 Some nested quotes not allowed

[CR #269224] When parameters passed to ASE Replicator command procedures contain nested quotes, outer quote characters must be double quotes and inner (nested) quote characters must be single quotes. For example:

```
sp_addreplicateart $subname, $pubobj, $subobj,  
@where_clause = "where coll = '3' "
```

Workaround: Always use double quote characters outside single quote characters when you use strings with nested quotes in ASE Replicator command procedure parameters.

8.4.9 Java object replication not supported

[CR #s 270775 and 286817] ASE Replicator does not support replicating Java objects.

Workaround: There is no workaround for this issue.

8.4.10 Some commands fail on case-insensitive servers

[CR #270776] ASE Replicator command procedures that deal with articles may fail on case-insensitive servers.

Workaround: Use the exact case of the database object name returned by catalog stored procedures on the case-insensitive server.

8.4.11 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_addreplicateart` and `sp_addreplicateconn` return some of their error messages incorrectly. This problem does not affect the function of these procedures.

Note Although the error messages are returned incorrectly, the errors are logged correctly in the ASE Replicator log files.

Workaround: There is no workaround for this issue.

8.4.12 Roman-8 character set issues

[CR #291478] The ASE Replicator process uses jConnect for JDBC for all of its communication with Adaptive Server. Therefore, the Adaptive Server default character set must map to a character set supported by jConnect. If the Adaptive Server default character set does not map to a character set supported by jConnect (such as Roman-8, for example), ASE Replicator cannot connect to the Adaptive Server.

When this problem occurs, the following error message appears in the *system.log* file:

```
Message: JZ0IB: The server's default charset of roman8
does not map to an encoding that is available in the
client Java environment.
```

Workaround: Use an Adaptive Server default character set that maps to a character set supported by jConnect (such as UTF-8, for example).

8.4.13 Java out-of-memory problems

[CR #273087] In some circumstances, ASE Replicator may shut down with a Java `OutOfMemory` error. This problem can occur when:

- The value of the `queue_size` parameter is set too high for either a database connection or the ASE Replicator general configuration parameter. Sybase recommends that you use the default values.
- A problem occurs on one or more connections where data gets backed up in the queue. Depending on the size of the data, the `queue_size` value for the affected connection, and the amount of memory allocated to the Java virtual machine (VM), a Java `OutOfMemory` exception can occur.

Workaround: If you encounter a Java `OutOfMemory` error, you can do either or both of the following:

- Reduce the `queue_size` for each database connection and for the ASE Replicator general configuration.
- Increase the amount of memory allocated to the Java VM by editing the `-Xmx` parameter in the `aserep` start-up script. The default is 64MB.

You can estimate the amount of memory to allocate to the Java VM for a particular `queue_size` value by evaluating the average size of a row of data in any table at each database connection, and factoring that value with the amount of memory available on the machine.

- For each primary database connection:
 - Determine the average size (in bytes) of a row of table data replicated.
 - Multiply that value by the `queue_size` value for the connection, and add an overhead factor of 25 percent of the data size.
 - Multiply that number by the number of subscriptions that subscribe to data in a publication on the primary connection.
- Find the average (or maximum) for all primary connections and multiply that number by the ASE Replicator `queue_size` setting. Add an overhead factor of 25 percent.
- The sum of these numbers, plus approximately 32MB, should give you an estimate of the amount of memory required by the Java VM.

Note This amount of memory is required only in a worst-case scenario, in which the queues are filled with data.

8.5 Highlighted known XML problems

8.5.1 Use of XPath queries contain wild cards that fetch an attribute

[CR #331064] The `xmlextract` and `xmltest` built-in functions return incorrect results for an XPath query that specifies an attribute reference to the right of a wild card and an element reference.

For example, the following should return “99,” but instead returns an empty result:

```
select xmlextract("/a/*/c/@at", "<a><b> <c at='99' /></b></a>")
```

Workaround: Specify a descendant-path query:

```
select xmlextract("/a/*/c//c/@at", "<a><b> <c at='99' /></b></a>")
```

8.5.2 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.

8.5.3 XML samples: *jcs.jar* and *sqlx.jar* interplay issues

[CR #332553] The files *jcs.jar* and *sqlx.jar* cannot be installed in the same database.

Adaptive Server 12.0 and Adaptive Server 12.5 provided sample Java classes for XML processing. They were described in Chapters 8 and 9 of *Java in Adaptive Server*, and their source code and JAR files were provided in the directory `$$SYBASE $$SYBASE_ASE/sample`. The JAR file for these sample Java classes was named *jcs.jar*.

Adaptive Server version 12.5.1 also includes a set of Java-based SQLX mapping classes that perform mappings between SQL and XML. The source code and JAR file for those SQLX mapping classes are in the `$$SYBASE/$$SYBASE_ASE/sample` directory of Adaptive Server version 12.5.1.

The JAR file for these classes is named *sqlx.jar*.

One of the classes provided in *jcs.jar* file is jXml, a class that provides a set of Java wrapper methods on DOM processing for a stored XML document. These sample Java classes for XML processing are not included in Adaptive Server 12.5.1, since they have been superseded by the facilities of the XQL processor of Adaptive Server 12.5 and the native XML processor of Adaptive Server 12.5.1.

Workaround: For applications that continue to use the jXml class, *jcs.jar* is available in the following location:

```
$$SYBASE/$$SYBASE_ASE/sample/JavaSqlXml/SQLX-examples/jcs.jar
```

Applications that use the jXML class or other facilities of the sample Java classes for XML must include that file in their CLASSPATH for client applications, and install it in the database for server applications.

jcs.jar for the jXml class and *sqlx.jar* cannot be installed in the same database, because the files contain several utility classes with the same names.

Applications that use those utility classes continue to run without change when the *jcs.jar* is replaced with the *sqlx.jar*. Therefore, you must install *jcs.jar* in a different database than *sqlx.jar*.

8.5.4 New reserved words

[CR #333291] The following words have been added to the reserved word list for Adaptive Server 12.5.1:

- xmlextract
- xmlparse
- xmltest

Workaround: Use quoted identifiers:

```
setquoted identifiers on
create table "exmextract" ("xmlparse" int)
```

8.6 Default values not always set correctly in the Job Scheduler templete creation wizard

[CR #334601] When you select a template that provides a dropdown list with possible values, the first value on the list is selected as the default. Double check the value selected to make sure it is the value you want. In some cases, the value selected is not the correct default.

Workaround: There is no workaround for this issue

8.7 SQL derived tables cannot have unnamed columns

[CR #333873] SQL derived tables cannot have unnamed columns.

For example:

```
select * from (select sum(advance) from titles) d1
```

causes an error message.

Workaround: Rewrite the query using derived column list or the “AS” clause.

For example:

```
select * from (select sum(advance) from titles) d1(a1)
```

or:

```
select * from (select sum(advance) as a1 from titles) d1
```

8.8 OpenClient SDK header files and libraries

[CR #334381] Adaptive Server version 12.5.1 is distributed without the Open Client SDK header files and libraries.

Workaround: Client-side utilities including isql, bcp, defncpy, and the jConnect jdbc driver are distributed with this release. You can use the jConnect jdbc driver to develop applications.

8.9 Reconfiguring from utf8 with external sort order

[CR #333614] Before Adaptive Server version 12.5.1, the utf8 character set could be used only with the binary sort order. However, Adaptive Server release 12.5.1 supports non-binary sort orders with the utf8 character set.

If you use utf8 and then configure an external sort order (see Table 2), you cannot reconfigure Adaptive Server with another character set without first reverting to the binary sort order with utf8. After reverting to binary sort order (which has an ID of 50), you can then safely change to any other character set.

If you change the character set without first going to the binary sort order with utf8, recovery of the master database hangs and Adaptive Server displays error message 1569:

```
Unable to get a conversion handler for the built-in sort
function.
```

If this occurs, you must download and start Adaptive Server version 12.5.1 ESD #1 or later to successfully bring up a server (see the Sybase Web site at <http://www.sybase.com> for information about downloading Adaptive Server ESDs). Check with Sybase support for the availability of Adaptive Server version 12.5.1 ESD #1.

Alternatively, you can re-create the master device and load from database dumps.

Table 2 lists the external sort orders for utf8 and their IDs.

Table 2: External sort orders

Sort order	Sort order ID
big5bin	194
cp932bin	192
cyrdict	140
eucjisbn	192
euckscbn	161
gb2312bn	137
gbpinyinbn	163
rusdict	165
sjisbin	179

`$$SYBASE/collate/unicode` also lists the sort orders and their IDs

8.10 Versions strings for Open Client/Server are not correct

Problem: Open Client version strings are incorrect. The current string reads:

```
Sybase CTISQL Utility/12.5.1 ASE Edition/P-EBF11358; SSConfig \
VIRELTYPE/DRV.12.5.1/SPARC/Solaris 2.8/BUILD1251004/64bit/OPT/Mon Aug \
11 10:11:22 2003
Sybase CTBCP Utility/12.5.1 ASE Edition/P-EBF11358; SSConfig \
VIRELTYPE/DRV.12.5.1/SPARC/Solaris 2.8/BUILD1251-004/64bit/OPT/Mon Aug \
11 10:11:22 2003
```

The version string should read:

```
Sybase CTISQL Utility/12.5.1 ASE Edition/P-EBF11358 \
/DRV.12.5.1/SPARC/Solaris 2.8/BUILD1251-004/64bit/OPT/Mon Aug \
11 10:11:22 2003
Sybase CTBCP Utility/12.5.1 ASE Edition/P-EBF11358 \
/DRV.12.5.1/SPARC/Solaris 2.8/BUILD1251-004/64bit/OPT/Mon Aug \
11 10:11:22 2003
```

Workaround: The strings contain a value name rather than the value itself. This has no effect on Open Client/Server functionality and you can disregard the different version strings.

8.11 Backup Server *sybmultbuf* mismatch error

[CR #333539] In some circumstances, the *sybmultbuf* entry in the `$SYBASE_ASE/install/RUN_backupserver` script does not point to the upgraded release area after an upgrade from pre-12.5 Adaptive Server to Adaptive Server 12.5.1. This results in a backupserver *sybmultbuf* mismatch. The error message you see when using backupserver may be similar to:

```
Dec 25 13:50:59 2003: A00: SYBMULTBUF ERROR: Unsupported
revision number between backupserver and sybmultbuf.
Got revision number=19, but expected revision
number=18.
```

Workaround: Manually edit the `$SYBASE_ASE/install/RUN_backupserver` script so the `-M` option for *sybmultbuf* is pointing to the newly upgraded area, as well as the Backup Server.

8.12 Queries containing a select statement with a group by statement

[CR #333077] Queries containing a select with a group by with a case expression using a local variable, may return wrong results.

Workaround: Remove the use of the local variable in group by.

8.13 Job Scheduler Agent cannot terminate a job on request

[CR #326255] The Job Scheduler Agent cannot terminate a job on request.

Workaround:

- 1 Retrieve the `spid` and target server name from the stored procedures used in the job you want to terminate. The GUI displays this information.

List all running jobs by scheduled job name:

```
sp_sjobhelp @name='sjname=<sjob name>',
@option='all_users,running'
go
```

The following values are returned:

- `runid` – the xid.
- `server` – the target server name.
- `job start` – when the job started.
- `user_run` – the name of the user running the job.

List all running jobs by job name:

```
sp_sjobhelp @name='jname=<job name>',  
@option='all_users,running'  
go
```

The following values are returned:

- `runid` – the xid
 - `server` – the target server name
 - `job start` – when the job started
 - `user_run` – the name of the user running the job
- 2 Log on to the target server.
 - 3 Use `sp_who` to check that the job you want to terminate has the `spid` associated with the `user_run` you expect.
 - 4 Terminate the job. Enter:

```
kill spid
```

8.14 Job timeout in Job Scheduler not currently available

[CR #332317] Job timeout functionality is not included in Adaptive Server 12.5.1. There is no workaround other than to use a stopwatch to determine the timeout, and use the workaround included with “Job Scheduler Agent cannot terminate a job on request” on page 23.

8.15 XML identifiers that cannot be used as element names

[CR #332511] You cannot use the following identifiers as element names or attribute names in XPath queries:

- `and`
- `not`
- `or`

- union
- to
- eq
- ne
- lt
- le
- it
- ge

Workaround: Revise the XML documents you want to query to avoid using these identifiers as element or attribute names.

8.16 *sybmigrate* error messages

[CR #332051] When evoking *sybmigrate*, warning messages display in the log.

Workaround: These messages are harmless and can be ignored.

8.17 XA might hang when used with Open Client/Server

[CR #331225] Applications using the XA protocol might hang when linked with the Open Client Server SDK that is provided as part of this release.

Workaround: Use the Open Client Server 12.5.1 SDK point product available after Adaptive Server 12.5.1 GA.

8.18 Enhanced Full-Text Search installation

[CR #331127] When installing Enhanced Full-Text Search with InstallShield, the text server entry is not added into *syssservers*.

Workaround: Manually add the entry:

```
sp_addserver <textservername>, sds
```

8.19 Errors in queries with derived tables

[CR #329454] Incorrect results may be returned for a query having more than one derived table with group by aggregates in the derived table expression.

Workaround: Do not run these types of queries.

8.20 SQL Advantage queries hanging

[CR #327303] Killing a server process associated with a SQL Advantage program does not notify the SQL Advantage client of the process termination. Subsequent queries issued by SQL Advantage hang.

Workaround: Do not kill server processes associated with a SQL Advantage program.

8.21 Views must have defined column names

[CR #314976] ANSI requirements specify that a view must have defined column names.

Workaround: You can suppress this requirement by using trace flag 1702.

8.22 Migrating between single-byte and multibyte charactersets

[CR #295842] sybmigrate does not support migration of data from single-byte (iso_1, roman8) to multibyte (utf8), even with -Jutf8.

Workaround: There is no workaround for this issue.

8.23 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.

8.24 *cis connect timeout* and *enable SNMP* do not work

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

Workaround: Do not use these configuration parameters.

8.25 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.

8.26 Network license daemon support

Problem: Adaptive Server version 12.5.1 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 12.5.1 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.

8.27 bcp and abstract data types

[CR #268302] Adaptive Server allows you to use bcp in to copy data containing Java objects, but it performs no load-time checking to ensure that the objects are consistent with the classes that are installed in the server. Any checking is delayed until the objects are actually accessed by an application. To be sure that all objects are compatible with their corresponding classes, Sybase recommends that you run a select query to access these objects on the server immediately after they have been copied. If there are any incompatibilities, they are reported via a Java serialization error (10707).

Workaround: There is no workaround for this issue.

8.28 Referencing Java data in a nonrecovered database

[CR #274071] Referencing Java data in a database that is not recovered or operational causes a “process (0xf000f) infected with 11” condition, which terminates the database connection.

Workaround: Verify that the database is recovered and available for use by issuing:

```
use database database_name
```

8.29 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.

8.30 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.

8.31 Handling multibyte character sets during migration

[CR #271839] 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.

8.32 Adaptive Server plug-in default character set and language

[CR #270906] The Adaptive Server plug-in default character set is *iso_1*, and the default language is *us_english*. If you are running an Adaptive Server that is using a different language or character set, you may encounter compatibility problems.

Workaround: Set the character set and language values in the Adaptive Server plug-in Connect dialog, so that they correspond to the configuration of the Adaptive Server.

8.33 Dynamic and host variables not allowed in DDL statements

[CR #234217] Dynamic and host variables in create view statements were allowed in earlier versions, but worked incorrectly when:

- create view statements were replicated, or
- Views containing dynamic and host variables were upgraded from text in the *syscomments* table.

Dynamic and host variables are no longer allowed in create view statements.

Workaround: To allow dynamic and host variables with this incorrect behavior in create view statements, start Adaptive Server with trace flag 108:

```
dataserver -dd_master -T108
```

9. 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.

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 12.0
- 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
- Security Guardian™ 11.1.1
- SQL Remote 6.0.2
- 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.

9.1 Replication Server compatibility

Replication Server
version 12.1
compatibility

Replication Server version 12.1:

- Does not replicate extended data such as wide rows or wide columns, without truncating it.
- Does not replicate Unicode data.
- Does not use SSL when connecting to Adaptive Server, although Replication Server can still replicate data from and to an Adaptive Server that uses SSL.
- Does not use an LDAP server to retrieve server information, but instead relies on traditional interfaces files for this information. Replication Server can, however, still replicate data from and to an Adaptive Server that uses an LDAP server for this purpose.

DirectConnect and
OpenSwitch
compatibility

Direct Connect and OpenSwitch do not currently support the following Adaptive Server version 12.5 and later features:

- Extended limits, such as wide rows and wide columns
- Unicode
- SSL
- LDAP

10. Documentation updates and clarifications

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

10.1 Utility Guide

10.1.1 *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
```

10.1.2 *ddlgen* update

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.

- `-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`.

10.1.3 Omission from *symigrate*

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

`symigrate` 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`.

10.1.4 Utilities moved

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

- `certauth`
- `certpk12`
- `certreq`

10.1.5 *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.

10.2 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)
```

10.3 Reference Manual: Building Blocks

10.3.1 `@@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.”

10.3.2 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).

10.4 Reference Manual: Commands

10.4.1 dbcc check parameter

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

Syntax	<pre>dbcc upgrade_object({dbid dbname}, { compiled_object_name "view" "procedure" "default" "rule" "default" "trigger" }, ["force" "check"])</pre>
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 dbcc <i>upgrade_object</i> with check option, you must drop and re-create the <i>compiled_object</i>.

10.4.2 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.

10.4.3 *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.”

10.4.4 *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.

10.4.5 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: %1!", @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

```

10.5 System Administration Guide

[CR #322843] 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”

[CR #331120] 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"|

```

10.6 Performance and Tuning Guide

10.6.1 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.

10.6.2 Multiple temp databases

The information on multiple tempdbs was inadvertently included in the *Performance and Tuning Guide*. This information should be in the *System Administration Guide* and will be moved there on 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. 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.

12. 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, you will find links to EBFs/Updates, 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/>.

12.1 Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

v **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.

- v **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.

12.2 Sybase EBFs and software maintenance

- v **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. Enter user name and password information, if prompted (for existing Web accounts) or create a new account (a free service).
 - 3 Select a product.
 - 4 Specify a time frame and click Go.
 - 5 Click the Info icon to display the EBF/Maintenance report, or click the product description to download the software.

