

Release Bulletin Adaptive Server® IQ 12.4.0 for Digital UNIX

Document ID: 74950-01-1240-01

Last revised: September 10, 1999

| Topic | Page |
|---|------|
| 1. Product summary | 1 |
| 2. Restrictions | 3 |
| 3. Installation Instructions | 5 |
| 4. Documentation for this version | 5 |
| 5. Special migration instructions | 6 |
| 6. Changed functionality in this version | 7 |
| 7. Known problems | 9 |
| 8. Product compatibilities | 14 |
| 9. Documentation updates and clarifications | 14 |
| 10. Technical Support | 49 |
| 11. Other sources of information | 50 |

1. Product summary

Enclosed is Adaptive Server™ IQ version 12.4.0, which is compatible with the following platform and operating system configurations:

- Digital UNIX V4.0d
- Digital UNIX V4.0e

Copyright 1989-1999 by Sybase, Inc. All rights reserved. Sybase, the Sybase logo, Data Workbench, InfoMaker, PowerBuilder, Powersoft, SQL Advantage, SQL Debug, Transact-SQL, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise, Adaptive Server Enterprise Monitor, AnswerBase, Backup Server, ClearConnect, Client-Library, DB-Library, dbQueue, DirectConnect, Distribution Agent, Embedded SQL, Enterprise Client/Server, Enterprise Connect, InformationConnect, KnowledgeBase, MainframeConnect, MAP, Net-Gateway, NetImpact, Net-Library, ObjectConnect, OmniConnect, OmniSQL Access Module, Open Client, Open ClientConnect, Open Client/Server, Open Gateway, Open Server, Open ServerConnect, PC DB-Net, PowerDesigner, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, RW-Library, Secure SQL Server, Security Guardian, SQL Anywhere, SQL Remote, SQL Server, SQL Server Manager, SQL Toolset, Sybase Central, Sybase IQ, Sybase SQL Desktop, Sybase SQL Workgroup, SyBooks, System 10, System 11, Watcom SQL, Web.SQL, WorkGroup SQL Server, XA-Library, XA-Server, and XP Server are trademarks of Sybase, Inc. Other product names used herein may be trademarks or registered trademarks of Sybase or other companies. 1/99

- Digital Unix V4.0f

Note The product name Digital UNIX has recently been changed to Tru64 UNIX. However the Adaptive Server IQ documentation still uses the old Digital UNIX name.

The following operating system command shows the level of your base system software:

```
% sizer -v
Tru64 UNIX V4.0D (Rev. 878);
Wed Jul 15 12:31:49 EDT 1998
```

Adaptive Server IQ 12.4.0 has been built to take advantage of certain performance features found in the Digital Alpha EV5.6 (21164A) and later chip architecture.

1.1 Required Operating System Patches

You must install the patches listed in this section to run Adaptive Server IQ on Digital UNIX.

- Digital UNIX V4.0d requires the March 1998 patch kit as the absolute minimum.

Normally, it is best to install the latest patch kit. However, as of Adaptive Server IQ 12.4.0 certification, the latest patch kits contain a problem in the pthreads library which causes a one to two times performance slowdown in Adaptive Server IQ, most notably in loads and multi-user activity. For this reason, you should avoid using those patches and use the previous patch kits. The following table shows you which kits to avoid (the middle column) and which you should use (the last column).

| For this operating system version | Do NOT use this patch kit | Use this patch kit |
|-----------------------------------|---------------------------|------------------------|
| DUNIX 4.0D | duv40das0004-19990723 | duv40das00003-19990208 |
| DUNIX 4.0E | duv40eas0002-19990617 | duv40eas00001-19990202 |
| DUNIX 4.0F | duv40fas0001-19990609 | None required |

This problem should be resolved in the next release of these patch kits.

To obtain patches, download them from the website at <http://www.service.digital.com/> or contact your Digital representative.

Note These patches require that you rebuild your kernel. You must halt your system and then boot from the console after the new kernel is put into place.

When installing patches, log in as the “root” (superuser) and follow the directions provided with the patch. All tasks for adjusting the operating system configuration must be performed as “root”.

Failure to install the appropriate patches before starting Adaptive Server IQ requires reinstalling Adaptive Server IQ.

2. Restrictions

Read this section! Your system may produce unexpected results if you ignore the restrictions and other instructions listed below.

If you are upgrading from version 11.x to 12.4.0, see chapter 3 of the *Adaptive Server IQ Installation and Feature Guide* for further restrictions.

Converting 12.0.x
databases to 12.4.0

Before you can use a database created with earlier versions of Adaptive Server IQ, you must run the script *upgrasiq.sql*, located in the *\$ASDIR/scripts* directory. For example, to update a database named *old.db*:

```
% dbisql -c "uid=dba;pwd=sql;dbn=old;eng=servername"  
$ASDIR/scripts/upgrasiq.sql
```

If you do not run this script, errors like the following may occur when you use certain front end tools with the ODBC driver to import or link to a table:

```
The operation failed. There are too many indexes  
on table 'amadbdbo_midtown'.  
Delete some of the indexes on the table  
and try the operation again.
```

You must run *upgrasiq.sql* once for each 12.0.x database to upgrade it to 12.4.0.

Note For Adaptive Server IQ versions 12.0.3, 12.03.1, and 12.4.0: Due to a timing related issue, the database server process will sometimes hang if an IQ database is started on the command line. For a workaround, see the *readme* file included with the software.

Insert into table from remote SQL database not supported

Adaptive Server IQ for Digital UNIX does not include Open Client Libraries. Current communications libraries (up to and including EBF 8263) do not support native kernel-threading. Therefore, using current communications libraries, you will be unable to insert data directly from Adaptive Server Enterprise using the **INSERT INTO *tablename* LOCATION** command. Place such data in an ASCII file and load it using the **LOAD** command.

CREATE DATABASE restrictions

For this release, to obtain the best performance, Sybase recommends the following minimum IQ page sizes:

- 32 KB (**IQ PAGE SIZE 32768**) for a database containing up to 10 million rows.
- 64 KB (**IQ PAGE SIZE 65536**) for a database containing up to 100 million rows. Note that this is the default IQ page size.
- 128 KB (**IQ PAGE SIZE 131072**) for databases with more than 100 million rows.
- Do not create any databases with an **IQ PAGE SIZE** of less than 16KB.

Query restrictions

- By default Adaptive Server IQ cursors are scrollable, meaning that Adaptive Server IQ keeps all the query results in a buffer so that you can scroll backwards. If the query returns more than a few thousand rows of output, you can improve performance by issuing the following command before running the query:

```
SET TEMPORARY OPTION Force_No_Scroll_Cursors = 'ON'
```

- Adaptive Server IQ does not support Transact-SQL style outer joins on expressions. The workaround is to use ANSI style outer joins instead.

For example, statements containing clauses like the following are unsupported:

```
SUBSTRING(COL1 ...) *= SUBSTRING(COL2 ...)
```

The following outer join format is supported:

```
FROM t1 LEFT OUTER JOIN t2 ON (SUBSTRING(COL1 ...) =
```

SUBSTRING(COL2 . . .)

3. Installation Instructions

For complete installation instructions, see *Adaptive Server IQ Installation and Feature Guide for Compaq Digital UNIX*.

3.1 Setting the LD_LIBRARY_PATH Environment Variable

You must set the LD_LIBRARY_PATH environment variable before running Adaptive Server IQ and utilities.

For instructions on setting the LD_LIBRARY_PATH environment variable, see the *Adaptive Server IQ Installation and Feature Guide*.

4. Documentation for this version

The table below lists the documentation for Adaptive Server IQ version 12.4.0.

All documentation for Version 12.4.0 is available in one of two formats:

* indicates document found on the Adaptive Server IQ Technical Library CD. Hard copy documentation can be ordered separately.

** indicates document may be provided in hard copy.

Table 1: Current Adaptive Server IQ documentation

| Part Number | Book Title |
|------------------|---|
| 74950-01-1240-01 | <i>Release Bulletin Adaptive Server IQ for Digital Unix**</i> |
| 34359-01-1240-01 | <i>Adaptive Server IQ Installation and Feature Guide for Digital Unix**</i> |
| 38152-01-1200-01 | <i>Adaptive Server IQ Administration and Performance Guide*</i> |
| 38151-01-1200-01 | <i>Adaptive Server IQ Reference*</i> |
| 38159-01-1200-01 | <i>Introduction to Adaptive Server IQ*</i> |

If you are also using Adaptive Server IQ Multiplex, please refer to the following documentation for further instructions:

| | |
|------------------|---|
| 35008-01-1240-01 | <i>Adaptive Server IQ Multiplex Installation and Feature Guide for Digital UNIX</i> |
|------------------|---|

Depending on how you use Adaptive Server IQ, you may also need to refer to the documentation for Adaptive Server Anywhere. Refer to the Version 6.02 edition of Adaptive Server Anywhere documentation on the Sybase Technical Library Web site, for the most current information that applies to Adaptive Server IQ. Older editions of the *Adaptive Server Anywhere User's Guide* and *Adaptive Server Anywhere Programming Interfaces Guide* are included on the Adaptive Server IQ Technical Library CD and in the printed documentation set.

4.1 Accessing Current Release Bulletin Information

A more recent version of this Release Bulletin *may* be available on the World Wide Web. To check for critical product or document information added after the release of the product CD, use the Sybase Technical Library Product Manual Web site.

To access release bulletins at the Technical Library Product Manual Web site:

- 1 Go to support.sybase.com.
- 2 Click the Manuals tab.
- 3 From the drop-down list choose Adaptive Server IQ.
- 4 In the window on the right, under Platform-Specific Collections, choose the appropriate version link.
- 5 In the window on the right, choose the release bulletin for your platform.

5. Special migration instructions

See the Adaptive Server IQ Installation and Feature Guide for instructions on migrating to Adaptive Server IQ 12.4.0 from earlier versions.

If you participated in the Adaptive Server IQ Beta Program and have not yet upgraded to any GA (General Availability) version, please call technical Support before upgrading to Version 12.4.0

The “Restrictions” in this *Release Bulletin* lists late-breaking requirements. Please read this section before running a new version of Adaptive Server IQ.

5.1 Obtaining query plans

In version 11.x, you could output the query plan using the command **IQ SET QUERYINFO ON**. In Adaptive Server IQ 12.0, run the following command to output the query plan:

```
SET TEMPORARY OPTION Query_Plan = 'on'
```

The plan will be in the *.IQMSG* file.

If you want additional detail or are sending the plan to Technical Support you can use:

```
SET TEMPORARY OPTION Query_Detail = 'ON'
```

6. Changed functionality in this version

The following sections describe changes since the Adaptive Server IQ documentation set was updated.

6.1 Changed functionality in Adaptive Server IQ 12.4.0

This section summarizes new features and other changes to Adaptive Server IQ 12.4.0. Details of most of these features are provided in “Documentation updates and clarifications” on page 14.

- **Improved handling of out-of-space conditions** If you run out of disk space, Adaptive Server IQ now provides better assurance that you will be able to add space where you need it.
 - Adaptive Server IQ now reserves space so you can issue a **create dbspace** command. You can control the amount of space reserved by setting two new options, **Reserved_Main_DBSpace_MB** and **Reserved_Temp_DBSpace_MB**. For details on these options, see “New options for reserving space” on page 34.
 - The error message now specifies the amount of space you need to add in megabytes, and matches the syntax in **create dbspace**.

It is extremely unlikely that you will ever need to bring down the server to add space. For revised documentation on what to do if you run out of space, see “Insufficient disk space” on page 35.

- **Forced recovery and leaked space recovery** Adaptive Server IQ 12.4.0 adds better support for recovering from crashed and potentially corrupt databases. It also adds support for recovering leaked storage space. For details of these features, see “Forced recovery and leaked space recovery” on page 36.

- **Improved stored procedure output** Stored procedures now display output in units that are easier to understand.
- **Minimum password length** Database administrators can specify a minimum password length, to discourage easily discovered passwords.
For more information, see “MIN_PASSWORD_LENGTH option” on page 42
- **ODBC 3.51** The ODBC driver has been updated to ODBC 3.51. This version of ODBC includes support for Unicode applications.
For more information, see “ODBC conformance” in the online edition of the *Adaptive Server Anywhere Programming Interfaces Guide*, Version 6.02.
- **Control of allowed JOIN syntax** In previous releases, some multi-table queries have been allowed that have ambiguous join clauses. In the present release, you can set an option to disallow such queries.
For more information, see “EXTENDED_JOIN_SYNTAX option” on page 28.
- **Zero-length data storage** If the length of a CHAR or VARCHAR cell is zero and the cell is not NULL, Adaptive Server IQ creates a zero-length cell, not a NULL. This change reverts to the original behavior in Adaptive Server IQ 12, as documented Chapter 5 of the *Adaptive Server IQ Administration and Performance Guide*.

6.2 Changed functionality in Adaptive Server IQ 12.03.1

This section summarizes new features and other changes to Adaptive Server IQ added in Version 12.03.1 and not yet included in the documentation.

- **sp_iqstatus displays IQ Page Size** See “SP_IQSTATUS now displays IQ Page Size” on page 42 for more information.
- **LOAD_MEMORY_MB option** This option replaces several older options for adjusting memory use during loads. For more information see “LOAD_MEMORY_MB option” on page 30.
- **JOIN_PREFERENCE option** Two new values were added for the **JOIN_PREFERENCE** option. For more information see “JOIN_PREFERENCE Option” on page 31.
- **Command-line options for server caching** Two new server switches override the database cache parameters you set with the **SET OPTION** command:

- **-iqmc** sets the size of the main buffer cache
- **-iqtc** sets the size of the temporary buffer cache
See “Additions to the start_asiq or asiqsrv12 command-line options” on page 15 for details.
- **stop_asiq utility** You can stop the server using the **stop_asiq** utility. For more information see “stop_asiq utility” on page 18.

6.3 Changed functionality in Adaptive Server IQ 12

Adaptive Server IQ 12 takes a giant step forward from earlier versions. Among the most important new features it includes are:

- Transactional database capabilities with table-level versioning
- The ability to update the database concurrently with queries by multiple users
- A new, more efficient database format
- More intelligent index loading, allowing faster loads for many indexes
- Syntactic compatibility with Adaptive Server Anywhere, allowing Anywhere users to build on their existing knowledge base as they begin to use Adaptive Server IQ
- Support for Transact-SQL
- Additional query and view support
- Improved front end support

See the *Adaptive Server IQ Installation and Feature Guide* Chapter 3, “Migrating Data from Prior Versions,” for important notes about how new features change database creation and connection.

7. Known problems

For a description of known problems in Adaptive Server IQ version 12.4.0, see the following sections. If there is a workaround for a problem, it is provided.

See also “Restrictions” for more information.

“Documentation updates and clarifications” contains details that were not documented in time for this release.

7.1 Data definition

This section reports problems with data definition.

7.1.1 Temporary tables in procedures

When you include an automatically created temporary table in a procedure, the table should be dropped automatically when the procedure completes.

In Adaptive Server IQ 12.4.0, however, the table is not dropped. As a result, it becomes visible outside the procedure, and a new instance of the table is created each time the procedure executes. These tables instances are not dropped until you disconnect.

This problem will be corrected in a future release. As a workaround in this release, put an explicit

```
drop table #temp_table_name
```

at the end of the procedure.

For example, the following procedure has been corrected so that the automatically created temporary table, #temp0, is explicitly dropped and will not be replicated.

```
create procedure foo
begin
  select * into #temp0 from table
  drop table #temp0 /* this line fixes the problem*/
end
```

7.2 Adaptive Server IQ Queries

This section reports problems with Adaptive Server IQ queries.

7.2.1 ANY, SOME, and ALL subquery support

Adaptive Server IQ does not yet support subqueries that use the ANY, ALL or SOME keywords. For example:

- > ALL
- >= ALL
- < ALL
- <= ALL

- = ALL
- != ALL

If you use an unsupported query in this group, Adaptive Server IQ returns an error like the following:

```
Feature, ANY, not yet implemented
```

Queries of this type can always be expressed in terms of IN subqueries or scalar subqueries using MIN and MAX set functions.

7.2.2 Large IN subqueries

There is a known performance limitation in Adaptive Server IQ version 12 that affects IN subqueries that return more than 1 million distinct values; such queries should be rewritten in term of correlated EXISTS subqueries until this server limitation is addressed.

7.3 Adaptive Server IQ Operations

This section reports problems with Adaptive Server IQ operations.

7.3.1 Output to file in DBISQL

There are known problems in using the output to *file* feature from a DBISQL session that will cause a server to abort. Do not use this feature in this release. Instead, put **> filename** on the end of a select statement.

The cases that cause the abort are:

- Using the feature when the IQ option **Force_No_Scroll_Cursors** is set to ON
- Performing a rollback command following the **output to** command, as in

```
select * from ....;  
output to file;  
rollback;
```

7.3.2 Changing length of DBISQL column values

By default, the maximum length of column values displayed by DBISQL is 30 characters. This may be inadequate for displaying output of stored procedures such as **sp_iqstatus**.

To avoid truncated output, increase the length by setting the **truncation_length** option as follows:

```
SET OPTION DBO.TRUNCATION_LENGTH = 80
```

Alternatively, from the DBISQL menu select Command→Options and enter a higher value for Limit Display Columns and/or Limit Output Columns.

7.3.3 User-defined variable issue

User-defined variables will core dump if used in IQ queries.

7.3.4 DBSPAWN ERROR when starting a server

When attempting to start a server, you may get the following message:

```
DBSPAWN ERROR -96 -- database engine already running
```

This means that dbspawn is finding the shared memory segment of a previously started server, and is unable to create a shared memory segment.

To resolve this issue, subsequent servers should be started with the parameter **-hs** which will turn off shared memory. For example:

```
start_asiq @kent.cfg kent.db -hs
```

All servers can be started with shared memory “off”. Note that if shared memory is turned off, then you will have to connect to the server using TCP/IP, instead of using the default shared memory connection.

7.3.5 Unsupported terminal types cause DBISQL error

If you set the terminal type to “dumb” or “unknown”, then start DBISQL, Adaptive Server IQ returns an error. For example:

```
% setenv TERM dumb % dbisql
Error at line 1
Unable to initialize screen routines
```

To avoid problems, use an xterm window to run DBISQL on UNIX systems. For example, you can start an xterm window with a scroll bar as follows:

```
% xterm -sb
```

7.3.6 Adding a raw disk dbspace

Do not specify the optional **SIZE** clause when adding a raw disk dbspace. Adaptive Server IQ will correctly calculate the size of the raw disk without it.

7.4 Sybase Central

This section reports problems with the Adaptive Server IQ plug-in for Sybase Central.

7.4.1 Problems with Add User-defined Data Type wizard

When you create a user-defined data type, the Adaptive Server IQ plug-in allows you to specify a default value incompatible with the user-defined base data type. For example, if you specify base type integer, you may insert a character string 'test' as default value.

In the Set Properties screen, the Add User-defined Data Type wizard lets you select, "Do not allow null values." Currently, you may select this option, but after you finish and open properties on the data type, it displays Yes next to "Allows null."

7.4.2 DSE not installed on Windows client systems

On Microsoft Windows 95 or 98 client systems, running the Directory Services Editor (DSE) utility from Sybase Central returns the following error:

```
Error Executing dsedit. Check
that the applications location
is in your Search path.
```

The Directory Services Editor is not installed with the Adaptive Server IQ Client Components on Windows 95 or Windows 98 systems. If you do not install Open Client, you cannot run Open Client utilities.

7.4.3 Add Service utility installed only with Server Components

Adaptive Server IQ NT Service Manager is not installed with the Client Components. Do not use the Add Service utility on Windows NT unless you have the Adaptive Server IQ Server Components running on the same system. If you attempt to do so, Sybase Central returns the following error:

```
Error Executing ASIQ Service. Check
that the applications location
is in your Search path.
```

If you do not install Open Client, you cannot run Open Client utilities.

7.4.4 Notification message setting omitted from Index Properties

When you use the Add Index Wizard to create a new index, the Choose IQ Index Type screen lets you specify the number of records that should be added before sending a notification message. The properties screen for the index does not display this setting, however.

7.4.5 Data Type column in Table Editor retains focus

In rare situations in the Table Editor, the focus may stay on the Data Type column when tabbing through the Table Editor.

8. Product compatibilities

See the *Adaptive Server IQ Installation and Feature Guide for Digital UNIX* for a list of client application tools that have been certified with Adaptive Server IQ version 12.4.0.

9. Documentation updates and clarifications

This section contains information omitted from documentation and new information that needs emphasis. It is organized into the following categories:

- Startup, shutdown, and connection
- Data definition (DDL)
- Data manipulation (DML)
- Backup and system administration
- Client application
- Help file

9.1 Startup, shutdown, and connection

9.1.1 Server startup requirements

The following clarification should be added to the Adaptive Server IQ Administration and Performance Guide.

You should always use the **start_asiq** utility to start the server on UNIX platforms. If you do not, among the tasks you must do which the utility normally does for you are:

- Remove all limits, and then set limits on the stack size and descriptors. To do so, go to the C shell and issue these commands:

```
% unlimited
% limit stacksize 8192
% limit descriptors 4096
```

Note Note that unlimited affects soft limits only. If you have any hard limits, you must change them by setting kernel parameters.

- Set all server parameters appropriately in the asiqsrv12 command.

9.1.2 Additions to the start_asiq or asiqsrv12 command-line options

Two server startup switches for database caching are new as of version 12.03.1. The following details about these new options will be added to the next update of the *Adaptive Server IQ Administration and Performance Guide* and the *Adaptive Server IQ Reference*.

Two parameters for the **start_asiq** (on UNIX) or the **asiqsrv12** command line override the database cache parameters that are set using the **SET OPTION** command. If the user has not set the parameters, the defaults are 8MB for the main cache size and 4MB for the temporary cache size. The following new server switches override values of the database options.

| Switch | Description |
|----------------------------------|--|
| -iqmc <i>number_of_MB</i> | Specify main cache size in MB. (Overrides default.) |
| -iqtc <i>number_of_MB</i> | Specify temporary cache size in MB. (Overrides default.) |

Two other new command-line options, **-iqfrec** and **-iqdroplks**, let you force database recovery, and recover leaked space. See “Forced recovery and leaked space recovery” on page 36 for details.

9.1.3 Specifying server switches

The range of permissible values for the **-iqsmem** switch was listed incorrectly in Chapters 2 and 12 of the *Adaptive Server IQ Administration and Performance Guide*, and in Chapter 2 of the *Adaptive Server IQ Reference*. The correct range is platform-specific. For Digital UNIX systems the range is up to 28,000 MB. The **start_asiq** utility does not set this switch on Digital UNIX systems.

The **-iqmt** switch is not set by **start_asiq** on Digital UNIX systems. The setting listed in the *Adaptive Server IQ Administration and Performance Guide* and the *Adaptive Server IQ Installation and Feature Guide* is incorrect. The default value is calculated from the number of connections and the number of CPUs, and is usually adequate.

9.1.4 -gm command line option

The description of the server parameter **-gm** in the *Adaptive Server IQ Reference Manual* should read as follows.

-gm num Limit the number of connections to the server that can be active at one time. If this number is greater than the number that is allowed under licensing constraints, it has no effect. The value should approximate the number of users expected to connect to the server.

9.1.5 -gn command-line option

The **-gn** server parameter should be added to Chapter 2 of the *Adaptive Server IQ Reference Manual* and to “Controlling performance from the command line” in Chapter 2 of the *Adaptive Server IQ Administration and Performance Guide*.

| | |
|-------------|--|
| Function | Set the number of execution threads for the Catalog Store. |
| Syntax | asqsrsv12 -gn integer... |
| Applies to | All operating systems and servers. |
| Description | <p>Set the number of execution threads that will be used for the Catalog Store and connectivity while running with multiple users.</p> <p>On Windows NT you need to specify this parameter in the asqsrsv12 command. To calculate its value use the following formula:</p> $gn_value = gm_value - ((2 * num_CPUs) + 10)$ <p>Specify a minimum of 25.</p> <p>On UNIX platforms, the start_asiq utility sets this parameter. See the <i>Adaptive Server IQ Installation and Feature Guide</i> for your platform for more information.</p> |

9.1.6 Using -v switch on 64-bit platforms

The following note should be added to Chapter 2, “The Database Server,” after the description of the **-v** server switch.

Note In order to display the version on 64-bit platforms, you must do the following:

- Run **start_asiq -v** instead, which will set up the correct paths, environment, iq parameters, etc... Anything you pass to **start_asiq** will be passed to **asiqsrv12**. End users should *always* use **start_asiq**.
 - Add **\$ASDIR/lib/iqstubs** to the beginning of the library path. However, remember, once you set this you will not be able to use any executables in **\$SYBASE**, as they will now pick up the wrong libraries.
-

9.1.7 Confirming connections

In Chapter 2 of the *Adaptive Server IQ Administration and Performance Guide*, in the section “Introduction to connections”, the second paragraph incorrectly states that **select db.name()** can be used to confirm a connection. The correct syntax is:

```
select db_name()
```

to display the current database, or

```
select db_name([database_id])
```

to display any database you specify.

9.1.8 Using ODBC connectivity with UNIX

On UNIX systems, Adaptive Server IQ installation installs only the ODBC driver, and not the driver manager. If you are using an ODBC application that uses *libodbc.so* (*libodbc.so.1*) or *libodbcinst.so* (*libodbcinst.so.1*), simply create symbolic links for these that point to *\$SYBASE/asiq12/lib/dbodbc6.so.1*. If you are creating a custom ODBC application, you can link directly to *dbodbc6.so*.

References to ODBC functions are resolved at run time. On UNIX, ODBC data sources are held in a file named *.odbc.ini*. Edit this file with any text editor to specify data sources. For details, see “Using ODBC data sources on UNIX,” in Chapter 2 of the *Adaptive Server IQ Administration and Performance Guide*.

If Adaptive Server IQ does not detect the presence of an ODBC driver manager, it will use `~/odbc.ini` for data source information. Otherwise, it will query the driver manager for data source information.

9.1.9 Using a .odbc.ini file

The following corrections apply to “Using ODBC data sources on UNIX,” in Chapter 2 of the *Adaptive Server IQ Administration and Performance Guide*.

When creating a `.odbc.ini` file on any UNIX system, you must use the long form of each identifier, as follows:

```
[My Data Source]
EngineName=myserver
CommLinks=tcpip
UserID=dba
Password=sql
```

The database server looks for the `.odbc.ini` file in:

- 1 The directory specified by the ODBCHOME environment variable
- 2 The directory specified by the HOME environment variables
- 3 The path

The database server ignores the ODBC_HOME, ODBC_INI and ODBCINI environment variables.

9.1.10 stop_asiq utility

The `stop_asiq` utility is new for version 12.03.1. Use this command to shut down an Adaptive Server IQ server and close all user connections to it.

When you issue the `stop_asiq` command, Adaptive Server IQ lists Adaptive Server IQ processes for your user ID and asks if you want to stop them. For example:

```
express@janed> stop_asiq

Checking system for ASIQ 12 Servers ...
The following 1 process(es) were found.
UID  PID  PPID  C   STIME TTY      TIME CMD
janed 6331   1    1 14:42:21 pts/1    0:19 asiqsrv12
@/express1/users/janed/sybase/asiq12/demo/asiqdemo.cfg
/express1/user
Do you want to stop the above process(es) <Y/N>?
```

Y

-

If you type Y (yes), the following message displays:

-
Shutting down asiqsrv12 (server shutdown).

If you type N, you are returned to the system prompt and IQ does not shut down the server. If no running servers were started by your user ID, Adaptive Server IQ displays information about processes run by other users, as follows:

express@janed> stop_asiq
Checking system for ASIQ 12 Servers ...
There are no ASIQ 12 Servers on this system owned by
'janed'
There were 6 other ASIQ 12 processes(s) found
UID PID PPID C STIME TTY TIME CMD
aharring 13870 13869 0 Jun 07 ? 1:20 asiqsrv12
-c 16m -gc 6000 -gd all -gr 6000 -gm 10 -gp 4096 -ti
4400 -tl 300 asi
teds 16255 5843 0 Jun 07 pts/27 0:03 asiqsrv12 -
gd all -tl 120 -gm 10 -n writer -gp 4096 -x
tcpip{port=6907}
jamesfay 3683 1 0 23:32:14 ? 14:33 asiqsrv12
-gc 6000 -gr 6000 -gm 10 -n express_daily_49765 -c 16M
-gp 4096 -gd a
redisch 5486 1 4 10:10:40 pts/6 348:20 asiqsrv12
-c 16m -gc 6000 -gd all -gr 6000 -gm 10 -gp 4096 -ti
4400 -tl 300 -n
maryc 16800 1 0 Jun 07 ? 30:10 asiqsrv12
@mary.cfg asiqdemo.db -o
/express1/users/maryc/12031/asiq12/logfiles/
ambler 26982 1 0 Jun 09 pts/18 13:11 asiqsrv12
@asiqdemo.cfg -N 6

Be sure to check with users before shutting down their servers.

9.1.11 Addition to STOP DATABASE statement

The following information should be added to the **STOP DATABASE** statement in the *Adaptive Server IQ Reference*.

When you issue

```
STOP DATABASE database-name
```

the *database-name* is the name specified in the **-n** parameter when the database is started, or in the **DBN (DatabaseName)** connection parameter. This name is typically the file name of the database file that holds the Catalog Store, without the *.db* extension.

9.1.12 Error in DBSTOP example

The *Introduction to Adaptive Server IQ* documentation on the Stop utility is incomplete. In the section “Shutting down a database server,” in Chapter 3, the UNIX system example shows no command parameters.

When you stop the server with the **DBSTOP** command, you need to specify the same parameters as when you started the server. Using a configuration file to start the server ensures that you will be able to find these parameters when you need them.

9.1.13 Disconnect details omitted

The following information was omitted from the Introduction to Adaptive Server IQ: It should be added in Chapter 4 after the section titled “Viewing connected users”:

Disconnecting users

If there are multiple connected users, Sybase Central lists them when you attempt to disconnect.

When you choose Tools —> Disconnect —> Adaptive Server IQ, a Filter Objects dialog box appears. Each database name is qualified by a server name, for example:

```
myserver.asiqdemo (DBA)
```

Select the desired database to disconnect and click OK.

9.1.14 Configuration files do not accept quotes

Do not use either single or double quotes when specifying server switches in configuration files.

The parameter string in the sample configuration file shown in Chapter 2 of the *Adaptive Server IQ Administration and Performance Guide* should be corrected to:

```
-n Elora
-c 16M
-x tcpip(port=2367)
-gm 10
-gp 4096
path\mydb.db
```

9.1.15 -Z switch must be uppercase

In Chapter 2 of the *Adaptive Server IQ Administration and Performance Guide*, in the seventh bulleted item under “What to do if you can't connect to an Adaptive Server IQ database, the **-z** command-line option should be **-Z**. This switch must always be specified uppercase.

9.2 Data definition (DDL)

9.2.1 Change to CREATE DATABASE statement

The defaults and minimums for the **IQ SIZE** and **TEMPORARY SIZE** parameters of **CREATE DATABASE**, for operating system files only, were stated incorrectly in the *Adaptive Server IQ Reference*. The default and minimum value depend on **IQ PAGE SIZE**.

Table 2: Default and minimum sizes of IQ and Temporary Stores

| IQ PAGE SIZE | IQ SIZE default | TEMPORARY SIZE default | Minimum explicit IQ SIZE | Minimum explicit TEMPORARY SIZE |
|---------------------|------------------------|-------------------------------|---------------------------------|--|
| 1024 | 512000 | 256000 | 1MB | 1MB |
| 2048 | 512000 | 256000 | 1MB | 1MB |
| 4096 | 512000 | 256000 | 1MB | 1MB |
| 8192 | 512000 | 256000 | 1MB | 1MB |
| 16384 | 1024000 | 512000 | 1MB | 1MB |
| 32768 | 2048000 | 1024000 | 2MB | 1MB |
| 65536 | 4096000 | 2048000 | 4MB | 2MB |
| 131072 | 8192000 | 4096000 | 8MB | 4MB |
| 262144 | 16384000 | 8192000 | 16MB | 8MB |
| 524288 | 32768000 | 16384000 | 32MB | 16MB |

9.2.2 MESSAGE PATH

In the MESSAGE PATH clause of CREATE DATABASE you must specify an operating system file. The message file cannot be on a raw partition. This is a correction to the *Adaptive Server IQ Reference*.

9.2.3 Low_Disk functions as High_Group

LD (Low_Disk) has been changed to function like HG (High_Group). It is no longer supported by a separate index and therefore does not offer any benefit beyond using HG. LD is still a valid syntax and may be used for compatibility with older versions.

9.2.4 Column Limit Correction

In Chapter 13 of the Adaptive Server IQ Reference the limit is incorrectly specified as 999 columns. The correct limit is 10,000 columns.

9.2.5 Joint Virtual Table (JVT) definition

A Join Virtual Table is a denormalized table which looks like a regular table; it has a name, columns, rows, and indexes. Unlike a regular table, however, a JVT cannot be created, modified or deleted by the user. It is created by Adaptive Server IQ as a result of a Create Join Index for internal processing purposes. It is deleted when the user does a Drop Join Index.

Error messages relating to join virtual tables appear only when a user tries to use or modify a JVT directly.

9.2.6 IQ database file paths must be unique

When you create a database or a dbspace, the path for the Temporary Store must differ from the path for the IQ Store. If your **CREATE DATABASE** or **CREATE DBSPACE** command specifies the identical path and filename for these two stores, you receive an error. This requirement was omitted from the *Adaptive Server IQ Reference* and the *Adaptive Server IQ Administration and Performance Guide*.

You can cause a unique path in any of these ways:

- Specify a different extension for each file (for example, *mydb.iq* and *mydb.iqtmp*)
- Specify a different file name (for example, *mydb.iq* and *mytmp.iq*)

- Specify a different pathname (for example, */iqfiles/main/iq* and */iqfiles/temp/iq* or different raw partitions
- Omit **TEMPORARY PATH** when you create the database. In this case, the temporary store is created in the same path as the Catalog Store, with the default name and extension *dbname.iqtmp* where *dbname* is the database name.

Warning! On UNIX platforms, you must be careful not to specify filenames that are links to the same file. IQ cannot detect where linked files point to. If the filenames in the command differ but they point to the same file, your database will be corrupted.

In Sybase Central, the Create Database utility fails if you do not provide extensions for file names. The *Introduction to Adaptive Server IQ* fails to mention this and also does not mention that file extensions are needed when operating system files are used.

9.2.7 Error documenting IQ PATH

In Chapter 3 of the *Adaptive Server IQ Administration and Performance Guide*, the example shown in the section titled “Specifying an IQ PATH” should be corrected. The example shows where Adaptive Server IQ puts the message and log files when you do not specify the Temporary and Message PATHS. It should read:

The Temporary Store is in ***/disk1/iqdata/company.iqtmp***

The IQ message log file is in ***/disk1/iqdata/company.iqmsg***

9.2.8 Error in raw partition limit

The *Adaptive Server IQ Reference* contains an error in the **CREATE DATABASE** statement description of the **IQ PATH** clause. The statement that raw partitions cannot be larger than 2GB is incorrect. The current limit is 128GB per file or dbspace, with a limit of 2047 dbspaces.

9.2.9 Error on CREATE DBSPACE

The **CREATE DBSPACE** command may occasionally return the following message:

```
1999-02-05 16:09:45 0002 [20152]:
You have run out of IQ STORE dbspace
```

```
in database /tech1/iq/cdrdb.db.  
In another session, please issue a CREATE DBSPACE ...  
IQ STORE command  
and add a dbspace of at least 1000 blocks.
```

After commit and checkpoint messages, the following error displays:

```
1999-02-05 16:12:53 0001  
Exception Thrown from hos_ion.cxx:170,  
Err# 4, tid 102 origtid 102  
1999-02-05 16:12:53 0001  
O/S Err#: 2, ErrID: 522 (hos_ioexception)  
1999-02-05 16:12:53 0001 File does not exist.  
File: /tech1/iq/cdrdb.iq2  
-- (hos_ion.cxx 170)
```

After many intervening commit and checkpoint messages, the following message displays:

```
The DBA has added 1 IQ STORE dbspaces  
to database /tech1/iq/cdrdb.db.  
Adaptive Server IQ (TM) is no longer  
waiting for more dbspace.  
1999-02-05 16:14:52 0002 [20902]: Insert completed.  
Index 'CDR.DBA.CDR_FE_FP_1f',
```

These messages are only informational. This is normal behavior as the server verifies that the dbspace does not yet exist, and no action is required. Such messages will be suppressed in a future release.

9.2.10 SIZE clause of CREATE DBSPACE

This clarification affects the *Adaptive Server IQ Reference* and *Adaptive Server IQ Administration and Performance Guide*.

You can only specify **SIZE** for the IQ Store and IQ Temporary Store, not for the Catalog Store.

9.2.11 Addition to DROP DATABASE statement

The following information should be added to the **DROP DATABASE** statement in the *Adaptive Server IQ Reference*.

The database must be stopped before you can drop it. If the connection parameter **AUTOSTOP=no** is used, you may need to issue a **STOP DATABASE** statement.

9.2.12 Recommended index types

In the *Adaptive Server IQ Administration and Performance Guide*, Chapter 4, “Adaptive Server IQ Indexes,” in Table 1–2, Query type/index, the recommended index types for **COUNT** and range predicate are incorrect. The correct types are:

| Type of Query Usage | Recommended Index Type |
|---|------------------------|
| As COUNT argument | LF or HG |
| In range predicate in WHERE clause (>, <, >=, <=, BETWEEN) | LF or HNG |

Note While **HNG** is recommended, in certain cases **LF** or **HG** is faster, and is often used in place of **HNG**. **HNG** tends to give consistent performance, while the performance of **LF** or **HG** with ranges depends on the size of the range selected.

9.2.13 Changes to “Using join indexes”

The section “Using join indexes” in the *Adaptive Server IQ Administration and Performance Guide*, Chapter 4, “Adaptive Server IQ Indexes, has several errors, described below.

Privileges needed to create a join index

The book states that you must be the owner of a table or the DBA to create, alter, or synchronize a join index that includes that table.

If you are not the DBA, you need to be owner of the table *and have RESOURCE authority* in order to create a join index.

One-to-many relationship

The following changes apply to the subsection “One-to-many relationship.”

The first sentence of the third paragraph should read, “If the join column is made up of more than one column, the combination of the values must be unique on the “one” side.

The example described in the third and fourth paragraphs is changed to match the *asiqdemo* database, and include sample rows that were omitted. The warning that follows these paragraphs is also changed. The corrected paragraphs read as follows:

If the join column is made up of more than one column, the combination of the values must be unique. For example, in the *asiqdemo* database, the *id* in the *customer* table and the *cust_id* in the *sales_order* table each contain a customer ID. The *customer* table contains one row for each customer and, therefore, has a unique value in the *id* column in each row. The *sales_order* table contains one row for each transaction a customer has made. Presumably, there are many transactions for each customer, so there are multiple rows in the *sales_order* table with the same value in the *cust_id* column.

So, if you join *customer.id* to *sales_order.cust_id*, the join relationship is one-to-many. As you can see in the following example, for every row in *customer*, there are potentially many matching rows in *sales_order*.

```
select sales_order.id, sales_order.cust_id,
       customer.lname
from sales_order, customer
where sales_order.cust_id = customer
```

```
id   cust_id   id           lname
      2583,101,101,'Devlin'
2001,101,101,'Devlin'
2005,101,101,'Devlin'
2125,101,101,'Devlin'
2206,101,101,'Devlin'
2279,101,101,'Devlin'
2295,101,101,'Devlin'
2002,102,102,'Reiser'
2142,102,102,'Reiser'
2318,102,102,'Reiser'
2338,102,102,'Reiser'
2449,102,102,'Reiser'
2562,102,102,'Reiser'
2585,102,102,'Reiser'
2340,103,103,'Niedringhaus'
2451,103,103,'Niedringhaus'
2564,103,103,'Niedringhaus'
2587,103,103,'Niedringhaus'
2003,103,103,'Niedringhaus'
2178,103,103,'Niedringhaus'
2207,103,103,'Niedringhaus'
2307,103,103,'Niedringhaus'
```

Warning! If the one-to-many relationship is incorrect, the join cannot be synchronized until you remove the extra rows from the "one" table. If you try to synchronize, you get a Duplicate Row error, and the transaction rolls back.

Creating star joins

The following should be added to Chapter 4 of the *Adaptive Server IQ Administration and Performance Guide* just after the figure that shows the *sales_order* table in a star join.

You can create this table using the following commands:

```
CREATE TABLE "DBA"."sales_order"

(

    "id"            integer NOT NULL,

    "cust_id"       integer NOT NULL
    REFERENCES "DBA"."customer" ("id")
    UNENFORCED,

    "order_date"    datetime NOT NULL,

    "fin_code_id"   char(2) NULL
    REFERENCES "DBA"."fin_code" ("code")
    UNENFORCED,

    "region"        char(7) NULL,

    "sales_rep"     integer NOT NULL
    REFERENCES "DBA"."employee" ("emp_id")
    UNENFORCED,

    PRIMARY KEY ("id"),

);
```

9.3 Error in DISK_STRIPING default

The *Adaptive Server IQ Reference* contains an error in the General Database Options table in Chapter 5, “Database Options.” The **DISK_STRIPING** option default should be ON.

9.4 Data manipulation (DML)

9.4.1 DISK_STRIPING default

The *Adaptive Server IQ Reference* contains an error in the General Database Options table in Chapter 5, “Database Options.” The **DISK_STRIPING** option default is ON for Windows NT, and OFF for UNIX.

9.4.2 Correction to Introduction

The first item under “Adaptive Server IQ benefits” in Chapter 1, *Introduction to Adaptive Server IQ*, should read:

- Intelligent query processing: Adaptive Server IQ uses index-only access plans to process only the data needed to satisfy any type of query

9.4.3 STRIP load option clarification

The following clarification should be added to the *Adaptive Server IQ Administration and Performance Guide*.

The **STRIP** load option does not apply to ASCII fixed-width inserts. For example, the **STRIP** option in the following statement is ignored:

```
load table dba.foo
  (col1 ascii(3), col2 ascii(3))
  from foo_data
  quotes off escapes off strip off
```

9.4.4 EXTENDED_JOIN_SYNTAX option

This option is new in Version 12.4.0 and should be added to the *Adaptive Server IQ Reference Manual*.

| | |
|----------------|---|
| Function | Controls whether queries with an ambiguous syntax for multi-table joins are allowed, or reported as an error. |
| Allowed values | ON, OFF |

| | |
|-------------|---|
| Default | ON |
| Description | <p>This option reports a syntax error for those queries containing outer joins that have ambiguous syntax due to the presence of duplicate correlation names on a null-supplying table.</p> <p>The following join clause illustrates the kind of query that is reported.</p> <pre>(R left outer join T , T join S on (C1))</pre> <p>where C1 is a condition. If the option is set to ON, this query is interpreted as follows.</p> <pre>(R left outer join T on (C1)) join S on (C2)</pre> <p>where C1 and C2 are conditions.</p> |

9.4.5 Support for joins between stores or databases

This section clarifies current support for joins between stores or between databases.

Any joins within a given IQ database are supported. This means that you can join any system or user tables in the Catalog Store with any tables in the IQ Store, in any order.

Joins of IQ tables with tables in an Adaptive Server Enterprise database are supported under the following conditions:

- On UNIX platforms, the IQ database must be the remote database, and the Adaptive Server Enterprise database must be the local database. In other words, on ASE you must add the IQ server a remote server, and define the IQ tables as proxy tables.
- On Windows NT, the IQ database can be either the local database or the remote database.

- In order to join a local ASE table with a remote IQ 12 table, the ASE version must be 11.9.2, and you must use the server class cd ASAnywhere.

Note Adaptive Server Enterprise 11.9.2 introduces new server classes to be used for accessing remote databases from the ASE 11.9.2 system. You must use server class ASAnywhere to access Adaptive Server IQ 12.0x or Adaptive Server Anywhere 6.x from ASE 11.9.2. You use server class ASIQ to access Adaptive Server IQ 11.5.1, 11.5.2, ... 11.5.x (that is, 11.5.1 or later as stated in *New Functionality in Sybase Adaptive Server Enterprise 11.9.2* is not correct for ASIQ 12).

- When you join a local IQ table with any remote table, the local IQ table must appear first in the FROM clause. This means that the local IQ table is the outermost table in the join.
- The CHAR data type is incompatible between Adaptive Server Anywhere and Adaptive Server IQ when the database is built with BLANK PADDING OFF. If you want to perform cross-database joins between ASA and ASIQ tables using character data as the join key, use the VARCHAR data type or use CHAR with BLANK PADDING ON.

9.4.6 New and changed general database options

The following changes in general database options will be added to the next update of the *Adaptive Server IQ Administration and Performance Guide* and the *Adaptive Server IQ Reference*.

LOAD_MEMORY_MB option

This version includes a new option to simplify user control of heap memory allocation during a load.

| | |
|----------------|---|
| Function | Specifies an upper bound on the amount of heap memory subsequent loads can use. |
| Allowed Values | 0 to 500 |
| Default | 0 (zero) |

Description This option specifies an upper bound (in MB) on the amount of heap memory subsequent loads can use. The default setting, 0 (zero), means that there is no upper bound, and Adaptive Server IQ can use as much heap memory as necessary to perform the load. A non-zero value means that the user has set an upper bound. This option is typically used for LOAD statements, but affects all operations where loads occur, including SYNCHRONIZE and DELETE operations.

Setting a non-zero value for this option is typically done to avoid overallocating the physical memory on the machine and exhausting virtual memory during a load operation. Given a user-specified upper limit, the load process will attempt to ensure that the load can succeed even if it means incurring degradation in load performance. However, every load requires some minimal amount of heap memory. If the user-specified limit is lower than this amount, the user will receive the following error:

The load user approximately <x> MB, but only <y> MB was specified

where <x> is a rough approximation of the minimum amount of heap memory needed and <y> is the upper limit set by the user. If this error occurs, then the user must increase or turn off the Load_Memory_MB option and try again.

We are removing the following options and replacing them with the new LOAD_MEMORY_MB option:

- INSERT_NUMBER_PRODUCERS
- INSERT_NUMBER_MT_BUFFERS
- INSERT_NUMBER_RAW_IO_BUFFERS
- INSERT_NUMBER_ROWS_PER_BUFFER
- INSERT_NUMBER_ROWS_PER_RAW_IO_BUFFER

JOIN_PREFERENCE Option

This version includes two new option values, 6 and -6 to prefer or avoid using a join index.

| | |
|-----------------------|--|
| Function | Controls the choice of algorithms when processing joins. |
| Allowed Values | -6 to 6 |
| Default | 0 |

Description

For joins within a query, the IQ optimizer has a choice of several algorithms for processing the join. This option allows you to override the optimizer's cost-based decision when choosing the algorithm to use. It does not override internal rules that determine whether an algorithm is legal within the query engine. If you set it to any non-zero value, it affects every join in a query; it cannot be used to selectively modify one join out of several in a query.

This option is normally used for internal testing, and only experienced DBA's should use it. The following table describes the valid values for this option and their action.

| Value | Action |
|-------|------------------------------|
| 0 | Let the optimizer choose |
| 1 | Prefer sort/merge |
| 2 | Prefer nested loop |
| 3 | Prefer nested loop push-down |
| 4 | Prefer hash |
| 5 | Prefer hash push-down |
| 6 | Prefer join index |
| -1 | Avoid sort/merge |
| -2 | Avoid nested loop |
| -3 | Avoid nested loop push-down |
| -4 | Avoid hash |
| -5 | Avoid hash push-down |
| -6 | Avoid join index |

9.4.7 Corrections to INSERT LOCATION

These corrections apply to the *Adaptive Server IQ Reference Manual* and the *Adaptive Server IQ Administration and Performance Guide*.

If you plan to use the **INSERT LOCATION** command (syntax 3 of the INSERT command), the Adaptive Server Enterprise server you are connecting to must exist in the *interfaces* file on the local machine.

Also, the syntax in the *Adaptive Server IQ Reference Manual* is incorrect. The *servername.dbname* must be enclosed in single quotes, and the **SELECT** statement must be in braces. An example of correct syntax is:

```
INSERT CUSTOMERS LOCATION 'BOSTON.PUBS2'
```

```
{ SELECT * FROM CUSTOMERS }
```

Note This syntax is not currently supported on Digital UNIX.

9.4.8 String function REPEAT is supported

The string function REPEAT is supported in Adaptive Server IQ version 12.x. It is documented as follows:

REPEAT function
[String]

Function

Returns a string composed of integer-expression instances of string-expression, concatenated together.

Syntax

REPEAT (string-expression, numeric-expression)

Parameters

string-expression The string-expression to be repeated.

integer-expression The number of times the string-expression will be repeated.

Examples

The statement

```
SELECT REPEAT( 'repeat', 3 )
```

returns the value repeatrepeatrepeat.

Standards and compatibility

SQL/92 Vendor extension.

Sybase REPEAT is not supported in Adaptive Server Enterprise, but REPLICATE provides the same capabilities.

9.4.9 Correction to CHAR function

The *Adaptive Server IQ Reference Manual* incorrectly lists the **CHAR** function as accepting a *string-expr* argument. The correct syntax is

```
CHAR ( integer-expr )
```

9.4.10 Number(*) function not supported

The **NUMBER(*)** function is not supported and should be deleted from the *Adaptive Server IQ Reference Manual*.

9.4.11 Using ISNULL() and COALESCE()

ISNULL() and COALESCE() can be used to convert NULL values into something else. If these are used with an outer join, the production of this expression gets pushed below the outer join. This behavior is acceptable on the row-preserving side of the join, but on the side that NULLs are added to, it can produce incorrect results. The ISNULL is processed in the vertical cursor, then NULLs are added in the outer join.

9.5 Backup and system administration

9.5.1 New options for reserving space

Two new options let you control the amount of space reserved for adding more disk space.

RESERVED_MAIN_DBSPACE_MB

| | |
|----------------|--|
| Function | Controls the amount of space Adaptive Server IQ reserves for adding dbspaces to the main IQ Store. |
| Allowed Values | Integer greater than zero, in megabytes |
| Default | 1 |
| Description | <p>This option lets you control the amount of space Adaptive Server IQ sets aside space in your main IQ Store, so that if you run out of disk space there you can add a new dbspace.</p> <p>Adaptive Server IQ sets aside 1 MB by default. This value is usually sufficient to run the DDL commands.</p> <p>You do not need to set aside room to hold the new dbspace. This option only provides space for executing related DDL commands.</p> |

RESERVED_TEMP_DBSPACE_MB

| | |
|----------------|---|
| Function | Controls the amount of space Adaptive Server IQ reserves for adding dbspaces to the Temporary IQ Store. |
| Allowed Values | Integer greater than zero, in megabytes |

| | |
|-------------|---|
| Default | 1 |
| Description | <p>This option lets you control the amount of space Adaptive Server IQ sets aside space in your temporary IQ Store, so that if you run out of disk space there you can add a new dbspace.</p> <p>Adaptive Server IQ sets aside 1 MB by default. This value is usually sufficient to run the DDL commands.</p> <p>You do not need to set aside room to hold the new dbspace. This option only provides space for executing related DDL commands.</p> |

9.5.2 Change to error message

The following error message has been changed.

| | |
|-------------------|--|
| Full message text | You have run out of '%1' dbspace in database '%2'. In another session, please issue a CREATE DBSPACE ... '%3' command and add a dbspace of at least '%4' MB. |
|-------------------|--|

| Item | Value |
|-------------|--|
| SQLCode | 1009131L |
| Constant | EMSG_IQSTORE_OUTOFDISK_HEADER |
| SQLState | QSB31 |
| ODBC State | 200152 |
| Parameter 1 | IQ STORE or IQ TEMPORARY STORE |
| Parameter 2 | Name of the database that needs more space |
| Parameter 3 | IQ STORE or IQ TEMPORARY STORE |
| Parameter 4 | Minimum number of megabytes to add |

9.5.3 Insufficient disk space

The following changes apply to the section “Insufficient disk space” in “Appendix A, Troubleshooting Hints” of the *Adaptive Server IQ Administration and Performance Guide*.

Replace the second and third Action items with the following text.

| | |
|---------|---|
| Actions | <ul style="list-style-type: none"> Try to connect to the database from a new connection. If this works, you know that the database server is running, even though the query is waiting. Then run sp_iqstatus to get more information. Check the sp_iqstatus output for the following two lines: |
|---------|---|

```
Main IQ Blocks Used: ,10188 of 12288, 82%, Max Block#: 134840
```

Temporary IQ Blocks Used: ,163 of 6144, 2%, Max Block#: 97

If the percentage of blocks used is in the nineties, you need to add more disk space with the **CREATE DBSPACE** command. In this example, 82% of the Main IQ Blocks and 2% of the Temporary IQ Blocks are used, so more space will soon be needed in the Main IQ Store.

| | |
|--------------------------------|--|
| Effect of checkpoints | <p>Insert the following text at the start of the section “Effect of checkpoints on out of disk space conditions.”</p> <p>If Adaptive Server IQ has already run out of space when a checkpoint is requested, the checkpoint command fails with the error:</p> <pre style="margin-left: 40px;">You have run out of space during the CHECKPOINT operation.</pre> <p>You must add additional dbspace before any new checkpoints can proceed.</p> |
| Adding the wrong type of space | <p>Assume, for example, the temporary dbspace has run out of space, but you accidentally add a main dbspace by omitting the temporary keyword in the create dbspace command. Your create dbspace command hangs, waiting for you to add space to the first dbspace.</p> <p>To continue, connect to the database from a new connection and create the needed temporary dbspace. Once this is done, the other create dbspace (for main) completes and all waiting connections resume running.</p> |

9.5.4 Forced recovery and leaked space recovery

This section describes parameters, options, and procedures that allow forced recovery and leaked space recovery. This information will be added to the *Adaptive Server IQ Administration and Performance Guide* and the *Adaptive Server IQ Reference Manual* in the next update.

| | |
|-----------------|--|
| Forced recovery | <p>Under certain conditions, previous versions of Adaptive Server IQ 12 could fail to recover a database after a system crash. Adaptive Server IQ 12.4.0 resolves this problem by adding a forced recovery option. The forced recovery option returns the database to its last known consistent state. Forced recovery should only be used when normal database recovery fails to restore the database to a running state.</p> <p>Normal database recovery differs from forced database recovery in these ways:</p> <ul style="list-style-type: none"> • Forced recovery marks all storage within the database as in use. In order to recover a potentially corrupt free list (allocation map) all storage within the database is marked as in use. You can return the storage allocation to its correct value by using the server startup parameter -iqdroplks and the sp_iqcheckdb storage procedure. See details below. |
|-----------------|--|

- **Incremental backups are disabled** After the database is opened in forced recovery mode, incremental backups are disabled. The next backup must be a full backup. Doing a full backup reenables incrementals.
- **Forced recovery affects all databases opened** The forced recovery parameter applies to all opens of the database while the server is up. Therefore, after the database is opened, the DBA needs to bring the server back down, and then restart the server without the forced recovery flag, to be sure that subsequent opens run in regular mode. Repeated opens of the database with forced recovery on do not harm the database, but could be confusing to the DBA because each forced recovery open marks all the storage within the database as in use.
- **Use the -gd switch during forced recovery.** Sybase recommends that you start the IQ server with the **-gd** switch set to restrict access to the server to the DBA. Forced recovery operates in all modes, but restricted access gives the DBA greater control over inadvertent opens of the database.

Leaked space recovery

Either through system failure or as a result of opening a database with forced recovery, a database's allocation map may not reflect the true allocation of its usage. When this occurs, we say that it has "leaked" storage. Adaptive Server IQ 12.4.0 resolves this problem by adding the ability to recover leaked storage.

When leaked storage is being recovered, other transactions that alter the allocation map are shut out. Such operations include checkpoints, and commands that modify the database.

Two new server startup parameters, **-iqfrec** and **-iqdroplks**, and one additional run-time setting, **dbcc_option**, support the new forced recovery and leaked storage recovery feature.

New server startup parameters

The following table describes the new server startup parameters.

| Parameter | Scope | Description |
|---------------------------------|----------|--|
| -iqfrec <i>dbname</i> | database | Marks the specified database as in use and restores database to its last known consistent state. |
| -iqdroplks <i>dbname</i> | database | Allows the sp_iqcheckdb stored procedure to recover leaked storage within the specified database. |

To verify that the data is not corrupt and set the database storage to its actual value, you start the server with the **-iqdroplks** switch and connect to the database. You then set the option **dbcc_option** and run the **sp_iqcheckdb** stored procedure. Depending on results, you may need to reset this option and rerun the procedure. See the discussion below for details.

New set option

A new set option, **dbcc_option**, controls the actions performed by **sp_iqcheckdb**. This option applies only to the database you are connected to when you set the option.

dbcc_option option

Function

Overrides the default operation of **sp_iqcheckdb**

Allowed values

0 - 3

Default

0

Description

The value of this option determines the processing that **sp_iqcheckdb** does.

| Value | Action |
|--------------|---|
| 0 | <p>Runs default action:</p> <ul style="list-style-type: none"> If -iqdroplks was specified during server startup, runs <i>CheckStorage Fix</i>. Checks for leaked blocks and corrupt database pages by walking all the block maps in the system and reading every database page. <i>If no error is detected, resets database free list to calculated allocation map</i>. If an error is found, it is reported in <i>.iqmsg</i> file, but free list is not altered. If -iqdroplks was <i>not</i> specified during server startup, runs <i>GatherStatistics</i>. Produces report. If severe error found, server may terminate. <p>In this read-only mode, sp_iqcheckdb does not lock the database to prevent other writers. Therefore, it may incorrectly report that leaks were detected, because it did not see new versions of objects outside its transactional scope.</p> |
| 1 | <p>Runs <i>CheckAllocation</i>. Checks for leaked blocks by walking all block maps in the system. Runs very fast (about 1 second per GB of data). Produces a report of findings. Does not reset the free list.</p> <p>This option provides a fast method of determining if the database has potentially leaked any storage.</p> |

| Value | Action |
|-------|---|
| 2 | <p>Runs <i>CheckStorage</i>. Checks for leaked blocks and corrupt database pages by walking all the block maps in the system and reading every database page. Runs about 50 times slower than option 1. Produces a report of findings. Does not reset the free list.</p> <p>This option is the same as the default option when -iqdroplks is passed to the server except that it runs in read-only mode.</p> |
| 3 | <p>Runs <i>CheckAllocation Fix</i>. Server <i>must</i> have been started with -iqdroplks switch. Checks for leaked blocks by walking all block maps in the system. Runs very fast (about 1 second per GB of data). Produces a report of findings. <i>If no error is detected, resets database free list to calculated allocation map.</i></p> <p>Sybase recommends that you use this option to recover the database free list only when the default option cannot do so due to errors encountered during processing, and no backed up version is available to restore.</p> |

Running
sp_iqcheckdb

In order to recover leaked storage within a database, first start the server with the **-iqdroplks** switch in the **asiqsrv12** command.

Next, connect to your database and issue the command:

```
sp_iqcheckdb
```

The stored procedure reads all storage within the database. On successful completion, it updates the database free list to reflect the true storage allocation for the database. It then generates a report listing the working and actions it has performed.

If it finds an error, **sp_iqcheckdb** reports the name of the object and the type of error found. It does not update the free list if any errors are detected.

Because it reads the entire database, **sp_iqcheckdb** may take a long time to run. The length of time depends on the size of the database and the size of the machine it executes on. Typically, **sp_iqcheckdb** can process between 20GB and 100GB per hour.

The **dbcc_option** settings of 1 and 3 provide a fast way to check for leaked storage within the system. They do this by walking the various block maps, or object directories, that make up the database. The underlying database pages that make up the actual tables and indexes are not read. Therefore, successful completion of **sp_iqcheckdb** using option 1 or 3 does not guarantee absolutely that the database is not corrupt.

The **dbcc_option** settings of 0 and 3, when combined with the server option **-iqdroppls**, update the free list if no errors are detected. In order to perform this function, write transactions are prevented before and during the running of **sp_iqcheckdb**. The stored procedure ensures this by taking the appropriate locks during its execution. Any write transactions are blocked while **sp_iqcheckdb** is running.

If it detects transactions that are not committed or not checkpointed, **sp_iqcheckdb** may refuse to recover leaked blocks. If this occurs, issue a **checkpoint** command and rerun **sp_iqcheckdb**. If **sp_iqcheckdb** still refuses to run, other users with active write transactions are connected to the database.

❖ **To recover leaked space:**

In the event that the default option (`dbcc_option = 0`) cannot recover the free list, and a previous backup is not available, use the following procedure to try to recover the database.

1 Start the server with the **-iqdroppls** switch in the `start_asiq` command (on UNIX) or `asiqsv12` command.

2 Set `dbcc_option` to 3, as a temporary option:

```
SET TEMPORARY OPTION dbcc_option = 3
```

3 Run the stored procedure:

```
sp_iqcheckdb
```

Note If this procedure fails, it is likely that the database is corrupt and beyond repair.

4 Set `dbcc_option` to 2, as a temporary option:

```
SET TEMPORARY OPTION dbcc_option = 2
```

5 Run the stored procedure again:

```
sp_iqcheckdb
```

6 From the report generated, drop the objects reporting errors.

7 With `dbcc_option` still set to 2, rerun the stored procedure to ensure no errors are present:

```
sp_iqcheckdb
```

8 Reset `dbcc_option` to the default value, 0:

```
SET OPTION dbcc_option = 0
```

Example

Assume that the DBA cannot successfully open and connect to database *foo*, because of reported IQ errors during database open and recovery. To force recovery and correct leaked space, follow the steps below.

Note Do not confuse an inability to connect to a database with an IQ server-level error while IQ is trying to open a database.

- 1 Start the database server with the **-iqfrec** switch:

```
asiqsrv12 -iqfrec foo ... -gd dba ... foo.db
```

or on UNIX

```
start_asiq -iqfrec foo ... -gd dba ... foo.db
```

- 2 Connect to the database (*foo*).

The *.iqmsg* file reports that the database was opened in forced recovery mode.

- 3 Bring down the server as you would normally. (If you use **dbstop** to bring down the server, be sure to include **-iqfrec** in the **dbstop** command.)

- 4 Start up the server again with the **-iqdroplks** switch:

```
asiqsrv12 -iqdroplks foo ... -gd dba ... foo.db
```

- 5 Connect to the database. It will be fully allocated.

- 6 To correct the leaks created by the forced recovery open, run **sp_iqcheckdb** on *foo*.

9.5.5 Improved output in stored procedures

Several stored procedures now display output in units that are easier to understand. The following table describes the new column names. The stored procedure **sp_iqindexsize** displays all of these columns except Nblocks. The stored procedures **sp_iqdbsize**, **sp_iqtablesize**, and **sp_iqjoinindexsize** display all of these columns except Info.

Table 3: New Stored Procedure Columns

| Column name | Description |
|------------------|--|
| Kbytes | Physical object size in KB |
| Pages | Number of IQ pages needed to hold the object in memory |
| Compressed Pages | Number of IQ pages when the object is compressed (on disk) |

| Column name | Description |
|-------------|--|
| Nblocks | Number of IQ blocks |
| Info | Component of the IQ index for which the Kbytes, Pages, and Compressed Pages are being reported. The components vary by index type. For example, the default (FP) index includes BARRAY and Bitmap (BM) components. The Low_Fast (LF) index includes Btree (BT) and Bitmap (BM) components. |

9.5.6 SP_IQSTATUS now displays IQ Page Size

The `sp_iqstatus` stored procedure now displays the IQ page size in addition to the block size. For example, Adaptive Server IQ used to display:

```
Block Size: 512/2bpc
```

It now displays:

```
Page Size: 1024/512blksz/2bpc
```

9.5.7 MIN_PASSWORD_LENGTH option

The **MIN_PASSWORD_LENGTH** option is new in this version.

| | |
|----------------|---|
| Function | Sets the minimum length for new passwords in the database. |
| Allowed values | Integer, greater than or equal to zero. The value is in bytes. For single-byte character sets, this is the same as the number of characters. |
| Scope | Can be set for the PUBLIC group. Takes effect immediately. DBA permissions are required to set this option. |
| Default | 0 characters |
| Description | This option allows the database administrator to impose a minimum length on all new passwords for greater security. Existing passwords are not affected. |
| Example | <ul style="list-style-type: none"> Set the minimum length for new passwords to 6 bytes. <pre>SET OPTION PUBLIC.MIN_PASSWORD_LENGTH = 6</pre> |

9.5.8 Transaction Log utility

The following information on the Transaction Log utility should be added to the *Adaptive Server IQ Administration and Performance Guide* and the *Adaptive Server IQ Reference*.

When you use the **RESTORE** statement to move and/or rename a database, you can rename all of the files except the transaction log. Transactions continue to be written to the old log file name, in the location where the Catalog Store file (the *.db* file) is located after the database is restored.

When you rename or move all other files in the database, it is preferable to do the same for the log file. To move or rename the log file, you use the Transaction Log utility (**DBLOG**). You should run this utility:

- After using **RESTORE** with a new database name
- After using **RESTORE** with the **RENAME** option

Note The database server must not be running on that database when the transaction log filename is changed. If you try to use this utility on a running database, you get an error message.

You can access the Transaction Log utility from the system command line, using the **DBLOG** command-line utility.

The DBLOG command-line utility

Syntax **dblog** [*switches*] *database-file*

Switches

| Switch | Description |
|------------------------------|----------------------------------|
| -m <i>mirror-name</i> | Set transaction log mirror name. |
| -t <i>log-name</i> | Set the transaction log name |

Description

The DBLOG command line utility allows you to display or change the name of the transaction log or transaction log mirror associated with a database. You can also stop a database from maintaining a transaction log or mirror, or start maintaining a transaction log or mirror.

Transaction log utility options

Set the name of the transaction log mirror file (-m) This option sets a filename for a new transaction log mirror. If the database is not currently using a transaction log mirror, it starts using one. If the database is already using a transaction log mirror, it changes to using the new file as its transaction log mirror. Most Adaptive Server IQ databases do not use a transaction log mirror, so this switch is rarely used.

Set the name of the transaction log file (-t) This option sets a filename, including an optional directory path, for a new transaction log. If the database is not currently using a transaction log, it starts using one. If the database is already using a transaction log, it changes to using the new file as its transaction log.

9.5.9 Error message for buffer cache settings

The following paragraph describes a change in behavior as of Version 12.4.0. It should be added to the section “Do not exceed physical memory” in Chapter 12, *Managing System Resources*,” in the *Adaptive Server IQ Administration and Performance Guide*.

If you set buffer cache sizes higher than your system will accommodate, Adaptive Server IQ will not successfully open the database. The new cache options for start up **-iqmc** and **-iqtc** should now be used to open the database and reset the defaults.

9.5.10 Setting Prefetch_Buffer_Limit option

The **SET** option **PREFETCH_BUFFER_LIMIT** defines the number of cache pages available to Adaptive Server IQ for use in prefetching (the read ahead of database pages). This option has a default value of 20, which can degrade multi-user performance. Sybase recommends that you set this option to 0 for multi-user applications. This option was omitted from Chapter 12, “Managing System Resources,” of the *Adaptive Server IQ Administration and Performance Guide*. For more information, see the *Adaptive Server IQ Reference*.

9.5.11 GRANT CONNECT for existing user ID

The following clarification should be added to Chapter 10, “Managing User IDs and Permissions” in *Adaptive Server IQ Administration and Performance Guide*, and to the **GRANT** statement in *Adaptive Server IQ Reference*.

If you have DBA authority, you can change the password of any existing user with the following command:

```
GRANT CONNECT TO userid IDENTIFIED BY password
```

The same command can also be used to add a new user. For this reason, if you inadvertently enter the user ID of an existing user when you mean to add a new user, you are actually changing the password of the existing user. You do not receive a warning because this behavior is considered normal. This behavior differs from pre-Version 12 Adaptive Server IQ.

To avoid this situation, use the system procedures *sp_addlogin* and *sp_adduser* to add users. These procedures give you an error if you try to add an existing user ID, as in Adaptive Server Enterprise, and pre-Version 12 Adaptive Server IQ.

9.5.12 Dropping users may delete tables

The following warning should be added to “Revoking user permissions” in Chapter 10, “Managing User IDs and Permissions” in *Adaptive Server IQ Administration and Performance Guide*, and to the **REVOKE** statement in *Adaptive Server IQ Reference*.

Warning! If you revoke a user's connect privileges, any database objects owned by that user are deleted without warning. Likewise, if you use the stored procedure *sp_dropuser* to drop a user, database objects owned by that user are dropped without warning. To avoid this problem, remove objects owned by a user or assign them to another user before issuing **REVOKE CONNECT** or **sp_dropuser**.

Note Procedures like *sp_dropuser* provide minimal compatibility with Adaptive Server Enterprise stored procedures. If you are accustomed to Adaptive Server Enterprise (or Adaptive Server IQ 11.x) stored procedures, you should compare their text with Adaptive Server IQ 12 procedures before using the procedure in **dbisql**. To compare, use the command

```
sp_helptext sp_name_in_question
```

9.5.13 Documentation on Data Backup and Recovery

The chapter on backup and data recovery in the *Adaptive Server Anywhere User's Guide* has been superseded by new information in Chapter 11, “Backup and Data Recovery” of the *Adaptive Server IQ Administration and Performance Guide*. The Adaptive Server Anywhere data backup and recovery information should not be used.

9.5.14 Changes to BACKUP statement

In the *Adaptive Server IQ Reference*, the **SIZE** and **STACKER** option descriptions in the **BACKUP** statement should read:

SIZE option Specifies maximum tape or file capacity (some platforms do not reliably detect end-of-tape markers). No volume used on the corresponding device should be shorter than this value. This value applies to both tape and disk files but not 3rd party devices. Units are KB so, for example, for a 3.5 GB tape you specify 3500000. Default for disk files is 2GB. No default exists for tape on UNIX platforms.

For tape devices on Windows NT:

- The value of **SIZE** must be a multiple of 64.
- If you specify a **SIZE** that is not a multiple of 64, it is automatically rounded down to a multiple of 64.
- If you do not specify **SIZE** explicitly, it is automatically set to 1.5GB.

STACKER option Specifies that the device is automatically loaded, and specifies the number of tapes it is loaded with. This value is not the tape position in the stacker, which could be zero. When **ATTENDED** is **OFF** and **STACKER** is **ON**, Adaptive Server IQ will wait for a pre-determined amount of time waiting for the next tape to be auto-loaded. The number of tapes supplied along with the **SIZE** option will be used to determine whether there is enough space to store the backed up data. Do not use this option with 3rd part media management devices.

9.5.15 Cleaning up after abnormal exit

Killing processes may result in semaphores or shared memory being left behind instead of being cleaned up automatically. To eliminate unneeded semaphores, you should periodically run the **-ipcs** command to check the status of semaphores and shared memory.

The **ipcs -a** command lists the ID numbers, owners, and create times of semaphores and shared memory segments. When all Adaptive Server IQ instances are started by the same user (as Sybase recommends), you can search the **OWNER** column for that username. Identify shared memory segments and semaphores that have a very old create date but have not been used recently.

After verifying with the owner that these are indeed unused, run **ipcrm** command to remove them. Use the **-m** parameter to specify the memory segment id and the **-s** command to specify the semaphore id number, in the following format:

```
ipcrm -m mid1 -m mid2 ... -s sid1 -s sid2 ...
```

For example:

```
% ipcrm -m 40965 -s 5130 -s36682
```

9.5.16 Monitoring server activity

It may be helpful, especially for new users, to monitor server activity. When you start a server with the **start_asiq** utility, server activity is logged in an ASCII text file placed in the directory defined by `$ASLOGDIR`. (If `$ASLOGDIR` is not defined, it defaults to `$ASDIR/logfiles`.)

The log file name has this format:

```
your_server_name.###.srvlog
```

Each time you start the server, the number is incremented. For example, your directory may look like this:

```
demo.001.srvlog  demo.002.srvlog
janedemo.001.srvlog
```

For information about your most recent session, choose the log with the largest number for the desired server. Issue a **tail -f** command to view the log contents. For example:

```
% tail -f demo.002.srvlog
```

When you run **start_asiq**, specify the **-Z** option to enhance the log file with additional information about connections. This will help new users or those troubleshooting connection problems.

To check if a particular server is running, log into the machine where it was started, and issue the command:

```
% ps -eaf | grep asiqsrv12
maryc 24836 25554 0 Feb 09 - 17:36
asiqsrv12 -c 16m -gc 6000 -gd all
-gr 6000 -gm 10 -gp 4096 -ti 4400
-tl 300 -iqmt 450 -iqsmem 2560
@fnma.cfg asiqdemo.db
janed 28932 38122 0 11:39:24 - 2:10
asiqsrv12 -c 16m -gc 6000 -gr 6000
-gm 10 -gp 4096 -ti 4400
-tl 300 -iqsmem 2560 -n janedemo -gd all
-iqmt 256 -x tcpip(port=1872)
```

In addition, the command **% stop_asiq** will display all Adaptive Server IQ processes running.

9.6 Client applications

9.6.1 ODBC AutoPreCommit omitted

The ODBC AutoPreCommit option was omitted from the *Adaptive Server IQ Reference*. Turning this option ON causes each statement to do a COMMIT *before* execution (as opposed to a COMMIT *after* execution for the AutoCommit option). The default for AutoPreCommit is OFF.

Set the AutoPreCommit option in either the Windows (NT/95) registry or the .ini file (Windows 3.1). For example:

```
[Sample DSN]
UID=DBA
PWD=SQL
AutoPreCommit=y
```

9.6.2 Using PC client applications

The following notes apply to PC client applications certified with Adaptive Server IQ:

- With BrioQuery, each query requires you to connect to the database. Be sure to close the query after processing to ensure that the connection to Adaptive Server IQ is closed. If you leave multiple queries open, you could consume more connections than you realize, eventually preventing other users from connecting to Adaptive Server IQ (since the number of configured connections would be exceeded).
- With Business Objects, you should change Password_Encryption = 1 to Password_Encryption = 0.

9.6.3 Creating attribute tables for PowerBuilder

In order to create attribute tables for PowerBuilder properly, you need to run the *iqpb.sql* script located in the Server directory on the PowerBuilder product CD against Sybase Adaptive Server IQ v.12.0.x using the provided Adaptive Server IQ ODBC driver. Earlier versions of PowerBuilder (6.5 and below) do not include the script. In this case, you may obtain the script from the Sybase FTP website. Contact your PowerBuilder support representative if you need additional information or PowerBuilder 7.0.

9.7 Help files

9.7.1 Adaptive Server IQ plug-in help reflects Multiplex support

Adaptive Server IQ Multiplex 12.4.0 is a separate product from Adaptive Server IQ 12.4.0. If you have not purchased or installed Adaptive Server IQ Multiplex, the functionality described in the online help topic Managing Multiplexes is not available.

10. Technical Support

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

Before you contact
Technical Support

Before contacting Technical Support, collect the following:

- operating system platform (For example, Sun Solaris 2.6 (SPARC))
- front end tool or connectivity protocol used (For example, Brio Query)
- configuration type (single user or multi-user)
- message log file

File named *dbname.iqmsg* in the directory, located by default where you started the database server.

- stack trace file

File named *stktrc.iq* in the directory where you started the database server.

- command or query that produced the error
- query plan

Enter the following commands and then rerun the command that produced the error

```
SET TEMPORARY OPTION Query_Plan = 'ON'  
SET TEMPORARY OPTION Query_Detail = 'ON'
```

The plan will be in the message log file.

- startup option settings
- connect option settings
- database option settings
- schema and indexes for the database

-
- Output from **sp_iqstatus** procedure

You may find additional help from the Sybase online support database, MySupport. MySupport lets you search through closed support cases, latest software bulletins, resolved and known problems, using a view customized for your needs. You can even open a technical support case online.

MySupport can be used from most Internet browsers. Open the MySupport home page, mysupport.sybase.com, and follow the instructions provided there to sign up for and use this free service.

11. Other sources of information

Use the Sybase Technical Library (available on CD and the web) resources to learn more about your product:

- Technical Library CD contains product manuals and technical documents and is included with your software. The DynaText browser (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 Technical Library.

- Technical Library Web site includes the Product Manuals site, which is an HTML version of the Technical Library CD that you can access using a standard Web browser. In addition, you'll find links to the Technical Documents Web site (formerly known as Tech Info Library), the Solved Cases page, and Sybase/Powersoft newsgroups.

To access the Technical Library Web site, go to

support.sybase.com

, scroll down to Support Services, and select a link under the Technical Library heading.

11.1 Sybase Certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ **For the latest information on product certifications and/or the EBF Rollups:**

- 1 Point your Web browser to Technical Documents at the following Web site:

techinfo.sybase.com

- 2 In the Browse section, click on the What's Hot entry.
- 3 Explore your area of interest: Hot Docs covering various topics, or Hot Links to Technical News, Certification Reports, Partner Certifications, and so on.

❖ **If you are a registered SupportPlus user:**

- 1 Point your Web browser to Technical Documents at the following Web site:

techinfo.sybase.com

- 2 In the Browse section, click on the What's Hot entry.
- 3 Click on the EBF Rollups entry.

You can research EBFs using Technical Documents, and you can download EBFs using Electronic Software Distribution (ESD).

- 4 Follow the instructions associated with the SupportPlusSM Online Services entries.

❖ **If you are not a registered SupportPlus user and you want to become one:**

You can register by following the instructions on the Web.

To use SupportPlus, you need:

- 1 A Web browser that supports the Secure Sockets Layer (SSL), such as Netscape Navigator 1.2 or later
- 2 An active support license
- 3 A named technical support contact
- 4 Your user ID and password

❖ **Whether or not you are a registered SupportPlus user**

You can use Sybase's Technical Documents. Certification Reports are among the features documented at this site

- 1 Point your Web browser to Technical Documents at the following Web site
techinfo.sybase.com
- 2 In the Browse section, click on the What's Hot entry.
- 3 Click on the topic that interests you.



Free Manuals Download Website

<http://myh66.com>

<http://usermanuals.us>

<http://www.somanuals.com>

<http://www.4manuals.cc>

<http://www.manual-lib.com>

<http://www.404manual.com>

<http://www.luxmanual.com>

<http://aubethermostatmanual.com>

Golf course search by state

<http://golfingnear.com>

Email search by domain

<http://emailbydomain.com>

Auto manuals search

<http://auto.somanuals.com>

TV manuals search

<http://tv.somanuals.com>