

Release Bulletin Sybase® IQ 12.6 for IBM RISC System/6000 AIX

Document ID: DC74600-01-1260-07

Last revised: June 2007

Topic	Page
1. Accessing current release bulletin information	2
2. Product summary	2
3. Package contents	5
4. Special installation and migration instructions	6
5. New features and behavior changes	17
6. Product compatibilities	37
7. Restrictions	39
8. Known problems	42
9. Documentation for this version	50
10. Documentation updates and clarifications	52
11. Technical Support	100
12. Other sources of information	102
13. Accessibility features	105
14. Checklist: information for Technical Support	107

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 product release, use the Sybase Product Manuals Web site.

❖ **Accessing release bulletins at the Sybase Product Manuals Web site**

- 1 Go to Product Manuals at <http://www.sybase.com/support/manuals/>.
- 2 Select a product and language and click Go.
- 3 Select a product version from the Document Set list.
- 4 Select the Release Bulletins link.
- 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® IQ version 12.6 with multiplex capability, which is compatible with the IBM RISC System/6000 on these operating system versions:

- AIX 64-bit 5L version 5.1
- AIX 64-bit 5L version 5.2
- AIX 64-bit 5L version 5.3

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

```
# oslevel
5.1.0.0
# oslevel
5.2.0.0
```

Refer to the *New Features in Sybase IQ 12.6* document for descriptions of the new features and behavior changes in Sybase IQ 12.6.

Note Sybase IQ is also called Adaptive Server IQ. You will sometimes see the name Adaptive Server IQ in documentation and in the product.

2.1 Required RS6000 patches

Sybase IQ version 12.6 requires the following maintenance levels:

Table 1: Maintenance levels for AIX

AIX Version	Maintenance level
5.3	5300-02
5.2	5200-02
5.1	5100-05

Use the following operating system command to verify the maintenance level of your system:

```
oslevel -r
```

For example, on AIX 5.3 systems, this command should return:

```
5300-02
```

Note Sybase recommends that you keep your system current with the latest patches. AIX 5.3 Technology Level 4, however, causes Sybase IQ to become unresponsive during shutdown. This issue, CR 431685, has been resolved with AIX 5.3 Technology Level 5. Refer to the IBM Web site for the latest patches.

2.2 Updated software components

Sybase IQ 12.6 installs software components updated to the latest versions available at release time. For example, the Sybase IQ 12.6 installation:

- Includes the Sybase jConnect JDBC Driver, Version 5.5. This component is optional, but strongly recommended. If you install the jConnect driver, Java classes installed into a database can make JDBC calls to access and modify data.

Sybase jConnect JDBC Driver, Version 6.0 is also included with this release and is optional. For more information, see the *Sybase IQ Installation and Configuration Guide for IBM RISC System/6000 AIX*.

- Updates the Java Runtime Environment (JRE) to version 1.4.2. This updated version is required to use the latest Sybase Central plug-in for Sybase IQ.

- Installs Open Client Software Developer's Kit 12.5.1.

Note Before installing the software, Sybase strongly recommends that you check the online support Web site for software updates to these components. If a software update (ESD or EBF) has been released, it contains bug fixes made after this product shipped. You must download the latest update and install it after installing IQ from the product CD. For more information, see "Finding the latest information on EBFs and software maintenance" on page 103.

2.3 Network Client supported platforms

The Sybase IQ Network Client is available for Windows and Linux. Both the Windows and the Linux versions of the Sybase IQ Network Client operate with the Sybase IQ server on any supported platform.

The Sybase IQ Network Client for Windows is included with Sybase IQ 12.6 on all platforms. The Sybase IQ Network Client for Linux can be downloaded or ordered separately.

- **Linux** – The Sybase IQ Network Client for Linux can be installed from the Network Client for Linux CD on the following configurations:
 - Red Hat Enterprise Linux 2.1 x86, Advanced Server or Workstation, with kernel version 2.4.9-e.40smp and glibc version 2.2.4-32.8, or
 - Red Hat Enterprise Linux 3.0 IA32, Advanced Server or Workstation Edition, distribution of kernel version 2.4.21-4.0.1.ELhugemem #1 SMP and glibc version 2.3.2-95.30, or
 - Red Hat Enterprise Linux 4.0, Advanced Server or Workstation Edition, distribution of kernel version 2.6.9-11.ELsmp #1 SMP and glibc version 2.3.4-2.9, or
 - SuSE Linux Enterprise Server 8.0, or
 - SuSE Linux Enterprise Server 9.0, with kernel version 2.6.5-7.244-pseries64 #1 SMP and glibc version 2.3.3.-98.61, or
 - SuSE Linux Enterprise Server (SLES) 10.0, with kernel version 2.6.16.21-0.25-ppc64 #1 SMP and glibc version 2.4-31.5.
- **Windows** – The Sybase IQ Network Client for Windows can be installed from the Network Client for Windows CD on a Windows 98, Windows NT SP6, Windows ME (Millennium Edition), Windows 2000 SP2, Windows 2003, or Windows XP SP2 system.

The Sybase IQ installation for UNIX and Linux platforms includes ODBC drivers needed for the client and DBISQL.

3. Package contents

You receive the following CDs:

- Sybase IQ 12.6 CD for your edition and platform
 - Edition may be Enterprise or Developer's Kit
 - AIX 64 is your platform
- Sybase IQ 12.6 Network Client for Windows CD

Note Sybase IQ 12.6 Network Client for Linux can be ordered separately.

- Sybase IQ 12.6 Sybooks CD
- Getting Started Sybase IQ 12.6 CD
- Sybase IQ 12.6 Megaphone Telco Demo CD

You also receive a booklet, *Getting Started*. This booklet tells you how to use the Getting Started CD, which contains Release Bulletins and Installation and Configuration Guides for all Sybase IQ 12.6 platforms.

3.1 Megaphone Telco Demo

The Megaphone Telco Demo, shipped on its own CD, is a larger demo than the *asiqdemo* sample database. The Megaphone Telco Demo shows how fast Sybase IQ can be as the size of the data grows. It features sample queries, a star schema, and a fact table with over one million rows.

4. Special installation and migration instructions

For complete installation and migration instructions, see the *Sybase IQ Installation and Configuration Guide for IBM RISC System/6000 AIX*. Be sure to read the section “Before you install” before installing Sybase IQ 12.6.

Note Sybase IQ 12.6 enforces check constraints. If you have existing check constraints, see the section “Before you install” in the chapter “Installing Sybase IQ” in the *Sybase IQ Installation and Configuration Guide* before installing.

“Before you install” also contains important information on upgrading LONG BINARY columns.

You *must* upgrade all existing databases after installing Sybase IQ 12.6, by running ALTER DATABASE UPGRADE. See the section “Upgrading databases to Version 12.6” for more information about ALTER DATABASE UPGRADE.

Before you run a new version of Sybase IQ, see “Restrictions” on page 39 for the most current requirements. This section includes installation and migration information that is new for this release or needs emphasis.

Subsequent sections include installation and migration information that needs emphasis or is new for this release.

Sybase strongly recommends that you check the online support Web site for software updates to these components before you install the software. If a software update (ESD or EBF) has been released, it contains bug fixes made after this product shipped. You must download the latest update and install it after installing IQ from the product CD. See “Sybase EBFs and software maintenance” on page 103.

4.1 Correction to ODBC connection variables [CR 447344]

The second paragraph in the section “Storing connection information” in Chapter 4, “Configuring Sybase IQ” of the *Sybase IQ Installation and Configuration Guide* should read as follows:

To connect with ODBC data sources, the location of your *.odbc.ini* file must be referenced by one of the following variables. Sybase IQ searches the directories specified by the variables below in the following order:

- \$ODBCINI – must contain the exact full pathname of the *.odbc.ini* file.
- \$HOME

- current directory
- \$PATH

4.2 Correction to ODBC driver reference [CR 446798]

The following paragraph appeared erroneously in the section “Installing ODBC Drivers” in Chapter 4, “Configuring Sybase IQ” of the *Sybase IQ Installation and Configuration Guide*.

If you use ODBC with UNIX or Linux, see “Using ODBC without the driver manager” in Chapter 4 of the *Adaptive Server Anywhere Programming Interfaces Guide* to ensure that you use the correct driver.

The *Programming Interfaces Guide* does not exist in the Adaptive Server Anywhere documentation set. The relevant section is now “Linking ODBC applications on UNIX” in Chapter 7 of the *Adaptive Server Anywhere Programming Guide*. This section is cross-referenced earlier in the same section of the *Sybase IQ Installation and Configuration Guide*. Another reference is unnecessary.

4.3 Correction to odbc.ini file format [CR 446795]

The section before the first example in “Storing connection information” in Chapter 4, “Configuring Sybase IQ” of the *Sybase IQ Installation and Configuration Guide* should read as follows:

Each entry in the *.odbc.ini* file should have the following format:

[an_entry_name]

Driver — the driver path

Userid — the user ID

Password — the password

EngineName — the desired engine

CommLinks — tcpip(port=engine_port_number)

AutoStop — no (Required parameter - must be set to no)

DatabaseName — the database name

DatabaseFile — the desired database with path. Used with embedded databases.

4.4 Correction to data source name utility [CR 446793]

An error appeared in Chapter 4, “Configuring Sybase IQ,” in the *Sybase IQ Installation and Configuration Guide*. The first sentence of the third paragraph in “Creating ODBC data sources” should read as follows:

You can also use the cross-platform iqdsn utility to create data sources.

4.5 Correction to ODBC driver names [CR 446702]

The first paragraph in the section “Installing ODBC drivers” in Chapter 4, “Configuring Sybase IQ” of the *Sybase IQ Installation and Configuration Guide* should read as follows:

When you install Sybase IQ on your UNIX or Linux server, the installation procedure also installs the ODBC driver, which can be directly accessed by applications. If you are using an ODBC application that uses *libodbc.so* (*libodbc.so.1*) or *libodbcinst.so* (*libodbcinst.so.1*), simply create symbolic links that point to *\$SYBASE/ASIQ-12.6/lib/libdbodbc9.so.1* for single threaded or *\$SYBASE/ASIQ-12.6/lib/libdbodbc9_r.so.1*. If you are creating an ODBC application, you can link directly to *libdbodbc9.so* for non-threaded applications and *libdbodbc9_r.so* for threaded applications. References to ODBC functions are resolved at run time.

4.6 Migrating across hardware platforms [CR 445754]

This section was omitted from Chapter 3, “Migrating Data from Previous Versions,” in the *Sybase IQ Installation and Configuration Guide*.

Sybase IQ supports migrating your database from one platform to another, as long as both have the same endian structure.

Platforms with big-endian structure are:

- AIX64
- HP-UX64 PA-RISC
- HP-UX64 Itanium
- IBM Linux on POWER
- SunOS64

Platforms with little-endian structure are:

- Linux32**

- Linux64
- Windows 32
- WinAMD64
- SunAMD64

Sybase IQ 12.6 ESD #2 and higher releases support migration between Windows and Linux.

IMPORTANT!

****** If you created your Sybase IQ database on a Linux 32-bit version prior to Sybase IQ 12.6 ESD #2, you must first install IQ 12.6 ESD #2 for Linux 32-bit systems and create a new data backup before migrating to another platform.

❖ **Migrating a database from one platform to another (same endian structure)**

- 1 Back up the database.
- 2 Shut down the Sybase IQ server.
- 3 Install the Sybase IQ server on the new platform. Your migration can take place on the same or a different machine.
- 4 Start the Sybase IQ server on the new hardware platform.
- 5 Connect to the utility database, *utility_db*.
- 6 Restore the database from the backup you created in Step 1.
- 7 Shut down the server and restart it against the restored database. If the current version of Sybase IQ is higher than the version on which you were previously running, you need to upgrade databases, and therefore restart the server in a way that restricts user connections. Sybase recommends using two server start-up options:
 - Use -gd DBA so that only users with DBA authority can start and stop databases.
 - Use -gm 1 to allow a single connection plus one DBA connection above the limit so that a DBA can connect and drop others in an emergency.

An alternate way to restrict connections is to specify

```
sa_server_option 'disable_connections', 'ON'
```

on the connection where you intend to perform the upgrade and

```
sa_server_option 'disable_connections', 'OFF'
```

on the same connection after upgrading. *The disadvantage is that this method precludes emergency access from another DBA connection.*

- 8 Start Interactive SQL and issue the database upgrade statement. For example:

```
ALTER DATABASE UPGRADE
```

If the database was created with the Java options off, append the keywords `JAVA OFF JCONNECT OFF` to the preceding command.

For more information, see “Upgrading non-multiplex databases,” in Chapter 3, “Migrating Data,” in the *Sybase IQ Installation and Configuration Guide*.

4.7 Moving data in one endian format to a system with a different endian format [CR 445754]

The following information was omitted from *Sybase IQ Installation and Configuration Guide* Chapter 3, “Migrating Data from Previous Versions,” and from the *Sybase IQ Release Bulletin*.

This section documents a procedure for moving data from a database in big-endian format to a database in little-endian format. This procedure moves table definitions but does not include migration of database objects, such as stored procedures or events, which you must recreate.

For example, Sybase IQ databases built on Sun64 SPARC systems store binary data in big-endian (Most Significant Byte first) format. Because Sun Solaris x64 is a little-endian system, Sybase IQ databases built on Sun64 SPARC cannot be upgraded with `ALTER DATABASE UPGRADE` to run on Sun Solaris x64 systems.

To move data for each database across hardware platforms of different endian structures, you must:

- Copy the database schema from the source platform (tables, indexes, etc.).
- Create a new database on the target platform.
- Perform a binary data dump from the source database.
- Load data into the new target database.

The following steps describe this process in detail.

❖ **Moving data from big-endian to little-endian systems or the reverse**

Note Before you begin, make sure that you have a process for capturing your database and table schema.

Check operating system documentation for the maximum file size for your system. For example, an extract file on Sun Solaris x64 has a maximum size of 512GB.

The following example loads a table named *lineitem* and identifies one extract file on UFS (file system) called *lineitem_binary.inp*.

- 1 Activate the extract utility:

```
SET TEMPORARY OPTION Temp_Extract_Name1 =
'lineitem_binary.inp'
SET TEMPORARY OPTION Temp_Extract_Name2 = ''
```

- 2 Set up a binary extract of the *lineitem* table:

```
SET TEMPORARY OPTION Temp_Extract_Binary = 'on'
SET TEMPORARY OPTION Temp_Extract_Swap = 'off'
```

- 3 Place output in the file *lineitem_binary.inp*:

```
SELECT * FROM lineitem
```

- 4 Turn off the extract utility:

```
SET TEMPORARY OPTION Temp_Extract_Name1 = ''
```

- 5 Create a duplicate of your database on the target system.

- 6 Assuming table *lineitem* as defined below, load the *lineitem* table as follows:

```
LOAD TABLE lineitem
( l_orderkey      BINARY WITH NULL BYTE,
  l_partkey       BINARY WITH NULL BYTE,
  l_suppkey       BINARY WITH NULL BYTE,
  l_linenumbers   BINARY WITH NULL BYTE,
  l_quantity      BINARY WITH NULL BYTE,
  l_extendedprice BINARY WITH NULL BYTE,
  l_discount      BINARY WITH NULL BYTE,
  l_tax           BINARY WITH NULL BYTE,
  l_returnflag    BINARY WITH NULL BYTE,
  l_linestatus    BINARY WITH NULL BYTE,
  l_shipdate      BINARY WITH NULL BYTE,
  l_commitdate    BINARY WITH NULL BYTE,
  l_receiptdate   BINARY WITH NULL BYTE,
```

```

l_shipinstruct      BINARY WITH NULL BYTE,
l_shipmode          BINARY WITH NULL BYTE,
l_comment           BINARY WITH NULL BYTE )
FROM 'C:\\mydata\\lineitem_binary.inp'
FORMAT BINARY
STRIP OFF
QUOTES OFF
ESCAPES OFF
PREVIEW ON
BYTE ORDER HIGH
COMMIT

```

Note particularly two clauses:

- BINARY WITH NULL BYTE is required when loading a binary file.
- BYTE ORDER HIGH specifies the byte order from the system where the data *originated*. The source database in this example is a big-endian platform; therefore, this data requires byte order HIGH. (Little-endian databases require byte order LOW.)

When loading a multiplex database, *use absolute (fully-qualified) paths in all filenames*. Do not use relative pathnames.

4.8 Database upgrades and procedure identifier limits [CR 421618]

The catalog `proc_ids` (procedure identifiers) are assigned sequentially and unused `proc_ids` are not reused. Over time, as procedures are dropped and created, databases created prior to IQ 12.6 may eventually reach the maximum `proc_id` limit of 32767, causing `CREATE PROCEDURE` to return an “Item already exists” error in IQ 12.6.

For databases created with a version prior to IQ 12.6 GA, the maximum `proc_id` for procedures is 32767, even if the database has been upgraded to IQ 12.6 or higher. This limit does not apply to databases created with IQ 12.6 and higher.

To determine if your database has a 32767 maximum `proc_id` limit, you can run `sp_columns sysprocedure`. If the data type for the `proc_id` column is `smallint`, the maximum `proc_id` of 32767 applies. To determine the current maximum `proc_id` value in use for your database, run the following query:

```
SELECT MAX(proc_id) FROM sys.sysprocedure
```

In IQ 12.6 ESD #7 and higher a check has been added to ensure that for databases created prior to IQ 12.6, the maximum `proc_id` is at a level that allows `ALTER DATABASE UPGRADE` to complete. If the maximum `proc_id` is above that level, then `ALTER DATABASE UPGRADE` does not run and instead returns the message “Database upgrade not possible”.

To resolve this issue for databases created prior to IQ 12.6, the `ALTER DATABASE UPGRADE` command supports a `PROCEDURE ON` clause in 12.6 ESD #7 and higher which will compact the `proc_ids` by recreating all stored procedures. The syntax is `ALTER DATABASE UPGRADE PROCEDURE ON`. The `PROCEDURE ON` clause is ignored for databases created in 12.6 and above.

Please note that using `ALTER DATABASE UPGRADE PROCEDURE ON` recreates all procedures without comments. If you want the comments back in the procedures after running the command, you must run `ALTER PROCEDURE <procedure_name>` with your source code for the procedures that contain comments. The `sp_helptext '<owner>.<procname>'` command can be used to save the text of procedures with comments before running `ALTER DATABASE UPGRADE PROCEDURE ON`.

Sybase recommends that you make a copy of the `.DB` and `.LOG` files for the database immediately before running `ALTER DATABASE UPGRADE PROCEDURE ON` as a backup. Because only the catalog is modified during an `ALTER DATABASE UPGRADE` command, a full backup is unnecessary.

4.9 Restrict connections during upgrades [CR 419328]

When you start a server to perform an `ALTER DATABASE UPGRADE` operation, always specify the `-gm 1` server start-up switch to allow only one user to connect and the `-gd` switch to restrict access to the DBA. Never allow other users to connect when `ALTER DATABASE UPGRADE` is running.

4.10 Setting the SYBASE_OCS environment variable [CR 409969]

If you have Sybase 15.0 products, such as Adaptive Server Enterprise 15, install Sybase IQ 12.6 in a separate directory instead of the same `$SYBASE` directory, and set the `SYBASE_OCS` environment variable correctly for each product.

OCS 15.0 (Open Client /Server), requires changes in the environment that conflict with OCS 12.5. Because Sybase IQ 12.6 uses OCS 12.5, installing it in the same `$SYBASE` directory may cause problems for Sybase products that use OCS 15.0, such as Adaptive Server Enterprise 15.

4.11 DROP PROCEDURE on multiplex after upgrade [CR 398236]

Attempting to drop or alter a procedure on a query server after upgrading from 12.5 to 12.6 returns a “permission denied” error.

To DROP or ALTER PROCEDURE on a multiplex upgraded to Sybase IQ 12.6, perform the following steps:

- 1 On the write server, recreate the procedure. This overwrites the query server’s version. For example:

```
CREATE PROCEDURE quarterly_results
```

- 2 On the write server, drop the procedure, for example:

```
DROP PROCEDURE quarterly_results
```

4.12 Data Source utility should be iqdsn not dbdsn [CR 393049]

References to dbdsn should be iqdsn in “Creating ODBC sources,” Chapter 4 of the *Sybase IQ Installation and Configuration Guide*.

4.13 Steps for preserving check constraints [CR 392330]

In the *Sybase IQ Installation and Configuration Guide*, Chapter 1, section “Before you install,” the steps for preserving check constraints are out of order. “Install Sybase IQ” should be step 1, not step 3. Remaining steps follow in order. After installation, the special stored procedures are in the scripts subdirectory of ASDIR.

4.14 Avoiding port number conflicts [CR 358218]

To avoid product conflicts, change the IQ port number in *default.cfg* and other configuration files (for example, *asiqdemo.cfg*) if Adaptive Server Anywhere is installed on the same system as Sybase IQ. Both products use the default port 2638.

First, add the following line to *\$ASDIR/scripts/default.cfg* with an unused port number, for example:

```
-x tcpip{port=4444}
```

Next, look for a port number definition in each configuration file. For example, *\$ASDIR/demo/asiqdemo.cfg* contains the following line:

```
-x tcpip{port=2638}
```

Edit the line and replace the default port number with the new one, for example:

```
-x tcpip{port=4444}
```

Save each file when finished.

If Adaptive Server Anywhere is on the same subnet as Sybase IQ, the server names must be unique.

4.15 Upgrading databases to Version 12.6

You must upgrade all existing databases after installing Sybase IQ 12.6. Running ALTER DATABASE UPGRADE makes the databases compatible with the new version of Sybase IQ, and lets you take advantage of performance optimizations and enhancements. Upgrading carries out the following tasks:

- Adds new system tables.
- Adds new columns to system tables and updates existing system table columns.
- Makes several changes to multiplex system tables.
- Adds several database options, and changes defaults or ranges for others. Database options provide a degree of tuning for overall database behavior or for individual user behavior. The list of available options is kept in the system tables, and needs to be upgraded for the options to become available. Options removed from Sybase IQ 12.6 are not removed from existing databases.
- Adds and updates system stored procedures. Multiplex stored procedures, while documented in *Sybase IQ Reference Manual*, are intended for use by Sybase Central for managing the multiplex, or are used internally by IQ during multiplex operation.

Sybase strongly recommends that you perform a full backup of your databases before upgrading to version 12.6.

See the chapter “Migrating Data from Prior Versions” in your *Sybase IQ Installation and Configuration Guide* for other migration issues. See *New Features in Sybase IQ 12.6* for more information on new system procedures, new options, and other new features and behavior changes.

Note Sybase IQ 12.6 enforces previously unenforced column and table CHECK constraints on updates, inserts, and loads of new data. Existing data in databases created with previous versions of Sybase IQ may contain unsupported constraints that now generate errors. To avoid these errors, see “Preserving check constraints before database upgrade” in your *Sybase IQ Installation and Configuration Guide*.

4.16 Upgrading database options

Sybase IQ 12.6 has many new and modified database options. For information on these options, see *New Features in Sybase IQ 12.6*.

If you have explicitly set database options that adjust performance and want to be sure that the new settings are appropriate for your environment, you may find it helpful to run `sp_iqcheckoptions`, which displays

- Database options whose value is changed from the default
- Current value of these options
- Default value for the connected user

It is a good idea to run `sp_iqcheckoptions` and capture its output before and after running `ALTER DATABASE UPGRADE`.

4.17 Migrating from versions prior to 12.5

You need to upgrade your database to Version 12.4.3 or Version 12.5 before you can upgrade to Version 12.6. Sybase IQ 12.6 supports upgrading from Version 12.4.3 to Version 12.6.

If you are migrating from any Adaptive Server IQ or Sybase IQ release prior to Version 12.5, be sure to read the Release Bulletins for all interim versions for important migration information, as well as new features in those versions.

5. New features and behavior changes

For new features and behavior changes, see *New Features in Sybase IQ 12.6*, which you should read before installing Sybase IQ 12.6. The following sections describe changes that were not included in that book or the documentation set.

5.1 Data Definition Language (DDL) changes

This section contains new features and changes related to DDL.

5.1.1 Database created with JAVA ON now returns correct error [CR 469250]

ESD #10 corrects a problem where Sybase IQ returned the same error (Error -845) for two different situations.

Now a database created with JAVA ON correctly reports a different error for incorrect dotted references than a database created with JAVA OFF.

In a database created with JAVA OFF, the following statement reports error -845, `SQLSTATE_INVALID_COLUMN_QUALIFICATION`, "Owner used in a qualified column reference does not match correlation name."

```
CREATE TABLE dot ( a int, b int );
INSERT dot VALUES ( 1, 1 );
COMMIT;
SELECT * FROM dot WHERE
dot..a = 100;
```

In a database created with JAVA ON, the statement reports error -706, `SQLSTATE_OMNI_SERVER_NOT_CAPABLE`, "Remote server does not have the ability to support this statement."

5.1.2 Unique table names on a connection [CR 455684]

In Sybase IQ 12.6 ESD #10, an attempt to create a base table or a global temporary table will fail if a local temporary table of the same name exists on that connection, as the new table cannot be uniquely identified by *owner.table*.

You can, however, create a local temporary table with the same name as an existing base table or global temporary table. References to the table name access the local temporary table, as local temporary tables are resolved first.

For example, consider the following sequence:

```
CREATE TABLE t1 (c1 int);
INSERT t1 VALUES (9);

DECLARE LOCAL TEMPORARY TABLE t1 (c1 int);
INSERT t1 VALUES (8);

SELECT * FROM t1;
```

The result returned is 8. Any reference to t1 refers to the local temporary table t1 until the local temporary table is dropped by the connection.

5.1.3 Creating dbspaces on raw devices in Sybase Central [CR 441162]

In the Create DBSpace wizard, on the Path screen, the raw device option is selected by default. Sybase recommends that you create main stores on raw devices. Be careful to deselect the raw device option when it is unnecessary. For example, raw devices disallow size values. If you inadvertently leave the raw device option selected and specify a size, Sybase IQ attempts to allocate the entire device. The raw device option disables path verification, so you should also deselect the option before specifying a file system name.

5.1.4 Moving the message file in Sybase Central [CR 405842]

Sybase Central now provides a way to specify the location of the *.iqmsg* file for a query server if you do not wish it to default to the same directory as the catalog database. On the Temporary DB Information page, select the Override Default File Paths check box. In the Specify the Path to the Message Files pop-up box, type a new path to the *.iqmsg* file. You may then create the query server.

5.1.5 MAX_IQ_THREADS_PER_CONNECTION minimum value changed [CR 395718]

As of 12.6 ESD #5, the minimum value of the MAX_IQ_THREADS_PER_CONNECTION database option is 3, rather than 2. This value is documented in Chapter 2 “Database Options” in Table 2-1 “General database options” and in “MAX_IQ_THREADS_PER_CONNECTION option” in the section “Alphabetical list of options.”

5.1.6 IQ UNIQUE field added to Sybase Central [CR 382663]

You can now add IQ UNIQUE constraints on the Sybase Central column properties page.

5.1.7 Domain names are case insensitive [CR 372849]

In Sybase IQ 12.5, user-defined data types (domains) were case sensitive in case sensitive databases. For example, you could create domains *HIRE_DATE* and *hire_date*. In Sybase IQ 12.6, domain names are case insensitive. Sybase IQ 12.6 returns an error if you attempt to create a domain with the same name as an existing domain except for case.

5.2 Multiplex enhancements

This section contains new features and changes related to multiplex.

5.2.1 Multiplex permission changes (ESD #3)

Starting with ESD #3, Sybase IQ 12.6 restricts permissions by default.

Prior to ESD #3, Sybase IQ allowed query servers to create and drop users, groups, and group memberships and to change permissions on objects created by write servers. This caused conflicts with permissions set by write servers.

Now only write servers can:

- Modify TABLE permissions (select, insert, delete, update, references and alter) on tables owned by the write server.
- Modify EXECUTE permission on stored procedures and functions owned by the write server. For related information, see “Multiplex DDL restrictions [CR 367460]” on page 64.
- Execute GRANT/REVOKE of DBA, RESOURCE, GROUP, or MEMBERSHIP.
- Create a user (GRANT CONNECT TO... where the user does not exist).
- Drop a user (REVOKE CONNECT FROM...).

The following objects, when created on a write server, are owned by that server and cannot be dropped or altered on a query server by default:

- Tables
- Views
- Indexes
- Data types
- Messages
- Constraints

- Procedures
- Functions

The following, when created, altered, or dropped on a write server, propagate to query servers:

- IQ base tables, IQ global temporary tables, and indexes on either IQ base or global temporary tables
- IQ referential integrity constraints and IQ check constraints
- User-defined data types (domains)
- Messages
- Users and groups
- Permissions
- Views
- Stored procedures and functions

The following are permitted on query servers:

- CREATE, ALTER, and DELETE EVENT
- Changing a user password via GRANT CONNECT TO... IDENTIFIED BY... when the user already exists

The database options MPX_GLOBAL_TABLE_PRIV and MPX_LOCAL_SPEC_PRIV, described in the following sections, let you override the new permission restrictions.

Note Sybase strongly recommends that you create users domains and messages on the write server only, to avoid static collisions. For details, see “Setting multiplex permissions [CR 404004].”

New database option MPX_GLOBAL_TABLE_PRIV

Function	Lets a query server grant and revoke permissions on objects created by the write server.
Allowed values	ON, OFF
Scope	DBA permissions are required to set this option. Can be set only for the PUBLIC group. Takes effect immediately.
Default	OFF

Description To enable this option, set it ON. Setting MPX_GLOBAL_TABLE_PRIV ON allows grant and revoke of table and execute permissions of write server objects on a query server.

See also “New database option MPX_LOCAL_SPEC_PRIV.”

Note This option was omitted from the chapter “Database Options” in the *Sybase IQ Reference Manual*.

New database option MPX_LOCAL_SPEC_PRIV

Function Lets a query server create and drop users, groups, and group memberships.

Allowed values 0 to 63

A bit mask indicating the corresponding special privileges to be granted and revoked on a query server:

- 0x01 = DBA
- 0x02 = create user via GRANT CONNECT and drop user via REVOKE CONNECT
- 0x04 = RESOURCE
- 0x10 = GROUP
- 0x20 = MEMBERSHIP

To combine two or more privileges, add the bit masks in hexadecimal (base 16), then convert to decimal to determine the value for the option. For example, to combine RESOURCE and GROUP privileges on a database, use the formula $4 + 16 \text{ (10 in base 16)} = 20$ and set MPX_LOCAL_SPEC_PRIV to 20.

To allow all privileges, you must set all bits. To do this, set MPX_LOCAL_SPEC_PRIV to 63.

Scope Can be set only for the PUBLIC group. Takes effect immediately.

Default 0

Description To enable the MPX_LOCAL_SPEC_PRIV option, set it to the appropriate value between 1 and 63 (See Allowed values.) DBA permissions are required to set this option. This option takes effect immediately.

See also

“New database option MPX_GLOBAL_TABLE_PRIV” on page 21.

Note This option was omitted from the chapter “Database Options” in the *Sybase IQ Reference Manual*.

5.2.2 Mixed-mode multiplex

Any multiplex where all servers are not at the same version is mixed-mode multiplex. Upgrading Sybase IQ from version 12.5 to 12.6, for example, typically results in a mixed-mode multiplex for a short period.

Please see “Before You Install” in the *Sybase IQ Installation and Configuration Guide* Chapter 1 for details about running multiple IQ Agents on a mixed-mode multiplex.

The following operations are not recommended on mixed-mode multiplexes:

- Creating a local store

Creating a local store succeeds on a 12.6 query server but not on the write server and is, therefore, lost at the next synchronize.

- Creating a main store

The create operation succeeds but displays a warning. The space cannot be deleted explicitly from any 12.6 query server. The query servers that still define the deleted main store will continue running, but once stopped will not restart without override flags, and a file_id mismatch between write and query servers causes subsequent main store creates to fail.

5.2.3 System tables and multiplex dbspace management [CR 375326]

Multiplex configuration information is found in the system table SYSIQFILE and in the following special system tables:

- IQ_MPX_INFO
- IQ_MPX_STATUS
- IQ_MPX_VERSIONLIST

For details, see Chapter 9, “System Tables,” in the *Sybase IQ Reference Manual*.

The following rules govern dbspace management in the multiplex:

- Each multiplex server has a set of rows in SYSIQFILE. This set includes:

- One row for each shared main dbspace
- One row for each temporary dbspace of the server
- One row for the .message file (.iqmsg)
- If it is a query server, one row for each local dbspace of the server
- The server_name column of SYSIQFILE identifies the server that owns the row. Each row that belongs to the write server has the empty string, '', as the server_name. Each row that belongs to a query server names that query server in the server_name column.

Do not confuse owning a row in SYSIQFILE with owning the dbspace. For example, a shared main dbspace named *main2* is owned by the write server, but each query server owns a row in SYSIQFILE for dbspace *main2* with the query server's server_name.

- Filenames may use either absolute or relative paths. On UNIX, they may also be UNIX links rather than actual files. The most flexible setup is to use relative paths for NFS files and, on UNIX, relative paths to links for raw devices.
- Filenames used for a shared main dbspace may differ across servers, either because the device names differ on different servers or because the DBA uses links with different names for each server.
- SYSFILE contains one row for each physical IQ dbspace file. Since there is one row in SYSIQFILE for each main dbspace for each server, there is a many-to-one relationship between the rows in SYSIQFILE and the rows in SYSFILE for shared main dbspace. All temporary and local dbspaces have exactly one row in SYSFILE and one row in SYSIQFILE.
- The query server rows in SYSIQFILE for a shared main dbspace are called **aliases**. The term alias may also refer to any row in SYSIQFILE that belongs to another server.

5.2.4 System procedures for managing multiplex dbspaces [CR 375326]

The following system procedures simplify multiplex dbspace management:

Table 2: Multiplex stored procedures

Procedure	Purpose
sp_iqmpxcountribremote	Returns a count of dbremote connections for a multiplex database.
sp_iqmpxgetconnversion	Displays the version number for a specified connection.
sp_iqmpxreplacewritserver	Converts the query server on which it runs into the new write server for the multiplex.
sp_iqmpxvalidate	Detects inconsistencies in the multiplex configuration information.
sp_iqmpxversioninfo	Displays server type (W for write server, Q for query server, S for single-node mode) and synchronization status.
sp_iqmpxcfg_<servername>	Sets up specified query server for SQL Remote replication.

For details, see Chapter 10, “System Procedures,” in the *Sybase IQ Reference Manual* and “sp_mpxcfg_<servername> procedure” on page 89.

5.3 Start-up and connection changes

This section contains new features and changes related to start-up and connection.

5.3.1 Thread stack memory error message change for ESD #9 [CR 449762]

The error on server start-up when thread stack memory exceeds machine resources now includes the message “malloc failed”.

Prior to ESD #9, the error displays as follows:

```
Sybase IQ fatal error: Insufficient heap memory
for stacks: try smaller -iqmt or -igtss values
```

As of ESD #9, the error displays as follows:

```
Sybase IQ fatal error: malloc failed Insufficient
heap memory for stacks: try smaller -iqmt or -igtss
values
```


5.3.2 TDS connects through default database [CR 449141]

Sybase IQ now allows connections to continue through to the default database even if the Open Client login server name does not match the name of the default database, provided that the connection string does not start a database (that is, there is no DBF parameter) and the server is running only one database. If both of these conditions are met, then the Open Client connection to the default database is allowed.

5.3.3 Reconnecting through *DBISQL* after restore [CR 415872]

Sybase IQ now requires the DBF parameter and database filename in order to connect to a database under certain circumstances. This situation occurs when you use DBISQLC or DBISQL and you have restored that database from backup while connected to *utility_db*.

For example, include the DBF parameter as follows:

```
CONNECT USING
'uid=DBA;pwd=SQL;dbf=node1/users/fiona/mydb.db;
links=tcpip{host=serv1;port=1234};eng=serv1_asiqdemo'
```

Prior to ESD #5, you could connect to a restored database using the following syntax:

```
CONNECT DATABASE mydb USER DBA IDENTIFIED BY SQL
```

The preceding command now returns a “specified database not found” error.

Another way to avoid the error is to enter a START DATABASE command while connected to *utility_db*, for example:

```
START DATABASE mydb
```

Use this method when connecting via DBISQL (Java).

5.3.4 Performance limitations for Java Virtual Machine [CR 404659]

Java applications running in IQ run slower than when run outside in a Sun JVM. Despite this limitation, Sybase recommends that you tune your applications by increasing the available memory for IQ JVM use with the database options `JAVA_HEAP_SIZE` and `JAVA_NAMESPACE_SIZE`.

5.3.5 LOGIN_PROCEDURE database options reset by ODBC [CR 393543]

The following option settings are explicitly set by the Sybase jConnect driver and the iAnywhere ODBC driver in accordance with the ODBC specification:

- `Time_format = 'hh:nn:ss'`
- `Timestamp_format = 'yyyy-mm-dd hh:nn:ss.ssssss'`
- `Date_format = 'yyyy-mm-dd'`
- `Date_order = 'ymd'`
- `Isolation_level = 0`

These options will overwrite settings by the `LOGIN_PROCEDURE` database option. Because these option settings are mandated by the ODBC specification, ODBC applications, including DBISQL applications, must explicitly set these options if they want different behavior. This could be done using the ODBC connection parameter `InitString`, for example:

```
iqdsn -wu foo -c
"uid=dba;pwd=sql;eng=foo;InitString='SET OPTION
PUBLIC.DATE_ORDER = 'DMY ''''"
```

5.3.6 New parameter for IQ Agent start-up [CR 391890]

In ESD #4, an optional host parameter was added to the start-up command on UNIX for the IQ Agent, `S99SybaseIQAgent1260`.

To start the agent using the host name returned by `uname -n` enter:

```
S99SybaseIQAgent1260 -host
```

To start the agent using the host's alias enter:

```
S99SybaseIQAgent1260 -host <foo>
```

where *foo* is an alias present in the `/etc/hosts` file.

5.3.7 Search Network check box defaults to OFF [CR 389092]

In the DBISQL Connect dialog, on the Database tab, the check box Search network for database servers appears. In Sybase IQ 12.6, this check box is deselected by default. You must enable this check box in order for DBISQL to locate remote servers.

5.3.8 Overriding IQ Agent Port Number

On UNIX, you can override the IQ Agent port number on the start-up command line; for example:

```
$ASDIR/bin/S99SybaseIQAgent1260 -port nnnn
```

On Windows, use the Service Manager to override the IQ Agent port number, as follows:

- 1 Select the name Sybase IQ Agent 1260.
- 2 Click the Properties icon or choose Action > Properties from the menu bar.
- 3 Click Stop under Service status to stop the agent.
- 4 In the Start Parameters text box, type `-port nnnn` where *nnnn* is the port number.
- 5 Click Start to restart the agent.

If the agent fails to start on Windows, check the event log for diagnostic information.

5.3.9 New `start_asiq` utility switches

The following command switches were added to IQ 12.6:

- `-iqnotemp` lets you start a multiplex database without a temporary dbspace by creating a temporary file in place of the dbspace.

The allowed value is *size*, the file size in MB.

For example, to create a 10MB temporary file:

```
-iqnotemp 10
```

- `-iqnomain` means start the query server without opening the shared IQ Main Store.

The allowed value, 1, means start without the store.

- `-iqnolocalreplay` specifies that IQ *not* replay the local TLV log after synchronizing the multiplex.

The allowed value, 1, means do not replay the log.

These switches were omitted from the *Sybase IQ Utility Guide*.

5.4 Administration and troubleshooting improvements

This section contains new features and changes related to administration and troubleshooting.

5.4.1 Displaying inconsistent indexes with sp_iqcheckdb [CR 402706]

The sp_iqcheckdb system stored procedure was enhanced in ESD #4.4 so that allocation mode now displays indexes that are inconsistent due to multiply-owned blocks.

The following keywords print block numbers for affected database blocks to the IQ message file:

- dupleaks — leaked blocks
- dumpdups — duplicate blocks
- dumpunallocs — unallocated blocks

For example, the following performs allocation checking for the entire database and reports block numbers for any leaked blocks:

```
sp_iqcheckdb 'allocation database dupleaks'
```

5.4.2 sp_iqversionuse procedure added [CR 401646]

Function	Displays version usage for the IQ Main store.
Syntax	call dbo.sp_iqversionuse ()
Description	The sp_iqversionuse system stored procedure helps troubleshoot situations where the databases uses excessive storage space due to multiple table versions. If out-of-space conditions occur or sp_iqstatus shows a high percentage of main blocks in use, run sp_iqversionuse to find out which versions are being used and the amount of space that can be recovered by releasing versions. Produces a row for each user of a version. Run sp_iqversionuse first on the write server to determine which versions should be released and the amount of space in KB to be released when the version is no longer in use. Connection IDs are displayed in the IQConn column for users connected to the write server. Version usage due to query servers is displayed as the query server name with connection ID 0. Run sp_iqversionuse on multiplex query servers to determine individual connections to query servers. Users from other servers are not displayed on a query server.

The amount of space is expressed as a range because the actual amount typically depends on which other versions are released. The actual amount of space released can be anywhere between the values of MinKBRelease and MaxKBRelease. The oldest version will always have MinKBRelease equal to MaxKBRelease.

WasReported indicates if version usage information has been sent from the query server to the write server. WasReported will be 0 initially on a write server for new versions. WasReported will change to 1 once SQL Remote replicates version usage information back to the write server. If WasReported is 0 for an extended period, SQL Remote might be stopped.

Table 3: Columns in sp_iqversionuse

Column name	Description
VersionID	The version identifier
Server	The server to which users of this version are connected
IQConnID	The connection ID using this version
WasReported	Indicates whether the server has received usage information for this version
MinKBRelease	The minimum amount of space that will be returned once this version is no longer in use
MaxKBRelease	The maximum amount of space that will be returned once this version is no longer in use

Example

In this example, the oldest version 42648 is in use by connection 108 on the write server (*mpxw*). Committing or rolling back the transaction on connection 108 will release 7.9MB of space. Version 42686 is in use by query server (*mpxq*) according to output from the write server. Using the query server output, the actual connection is connection 31. The actual amount of space returned from releasing version 42686 will depend on whether 42648 is released first.

WasReported is 0 for versions 42715 and 42728 on the write server because these are new versions that have not yet been replicated by SQL Remote. Since version 42728 does not appear on the query server output, it has not yet been used by the query server.

The following output is returned when sp_iqversionuse executes on the write server *mpxw*:

```
call dbo.sp_iqversionuse
```

VersionID	Server	IQConn	WasReported	MinKBRelease	MaxKBRelease
42648	'mpxw'	108	1	7920	7920

VersionID	Server	IQConn	WasReported	MinKBRelease	MaxKBRelease
42686	'mpxq'	0	1	7920	304
42702	'mpxq'	0	1	0	688
42715	'mpxq'	0	0	0	688
42728	'mpxq'	0	0	0	688

The following output is returned when `sp_iqversionuse` executes on the query server (*mpxq*):

```
call dbo.sp_iqversionuse
```

VersionID	Server	IQConn	WasReported	MinKBRelease	MaxKBRelease
42686	'mpxq'	31	1	0	0
42715	'mpxq'	00	1	0	0

5.4.3 Specifying index type for `sp_iqcheckdb` [CR 353538]

The new parameter *indextype* specifies an index type for all of the `sp_iqcheckdb` modes. The syntax is:

```
sp_iqcheckdb 'mode target [...] [ resources resource-percent ]'
```

where

mode:

```
{ allocation | check | verify } | repair
```

target:

```
[ main | local | indextype index-type [...] ] database | database resetclocks |  
{ [ indextype index-type ] [...] table table-name | index index-name [...] }
```

index-type::

```
{ FP | CMP | LF | HG | HNG | WD | DATE | TIME | DTTM }
```

Example

Verify the FP and HG indexes in the table t1 and the LF indexes in the table t2.

```
sp_iqcheckdb 'verify indextype FP indextype HG table t1  
indextype LF table t2'
```

5.5 Data load, update, and extraction enhancements

This section contains new features and changes related to data load, update, and extraction.

5.5.1 Inserting ASE DATE and TIME type data [CR 442348]

12.6 ESD #9 supports inserting data from an Adaptive Server Enterprise database column of data type DATE or TIME using the INSERT...LOCATION syntax of the INSERT statement.

5.5.2 Procedures added for disabling LOB compression [CR 352758]

In ESD #4, two stored procedures were added to control the compression (done when writing database buffers to disk) for columns of data type LOB. This functionality saves CPU cycles because certain data formats stored in a LOB column (e.g., jpg files) are already compressed and gain nothing from additional compression.

sp_iqsetcompression procedure

Function	Sets compression for LOB/CLOB data types (long binary and long varchar, respectively).
Syntax	call dbo.sp_iqsetcompression (owner, table, column, on_off_flag)
Permissions	Requires DBA authority.
Description	The sp_iqsetcompression system stored procedure lets you control compression of LOB/CLOB audiotape columns. The compression setting only applies to IQ base tables. A side effect of sp_iqsetcompression is that a COMMIT occurs after the compression setting is changed.

Table 4: Column in sp_iqsetcompression

Name	Data type	Description
<i>owner</i>	char(128)	Owner of the table for which you are setting compression.
<i>table</i>	char(128)	Table for which you are setting compression.
<i>column</i>	char(128)	Column for which you are setting compression.
<i>on_off_flag</i>	char(3)	Compression setting: ON enables compression, OFF disables compression.

Example

For this example, assume the following table definition:

```
CREATE TABLE USR.pixTable (picID INT NOT NULL,
picJPG LONG BINARY NOT NULL);
```

To turn off compression on the LOB column picJPG, call the `sp_iqsetcompression` procedure using the following command (you must have DBA permission):

```
CALL sp_iqsetcompression('USR', 'pixTable', 'picJPG',
'OFF') ;
```

This command returns no rows.

sp_iqshowcompression procedure

Function	Displays compression settings for columns of LOB/CLOB data types (long binary and long varchar, respectively).
Syntax	call dbo.sp_iqshowcompression (<i>owner, table, column</i>)
Permissions	Requires DBA authority.
Description	Returns the column name and compression setting. Compression setting values are 'ON' (compression enabled) or 'OFF' (compression disabled).

Table 5: Columns in sp_iqshowcompression

Name	Data type	Description
<i>owner</i>	char(128)	Owner of the table for which you are setting compression.
<i>table</i>	char(128)	Table for which you are setting compression.
<i>column</i>	char(128)	Column for which you are setting compression.

Example For this example, assume the following table definition:

```
CREATE TABLE USR.pixTable (picID INT NOT NULL,
picJPG LONG BINARY NOT NULL);
```

To check the compression status of the columns in the pixTable table, call the `sp_iqshowcompression` procedure using the following command (you must have DBA permission):

```
CALL sp_iqshowcompression('USR', 'pixTable',
'picJPG') ;
```

This command returns one row:

```
'picJPG', 'ON'
```


5.6 Query enhancements, optimization, and changes

This section contains new features and changes related to queries and optimization.

5.6.1 Function with VARCHAR result exceeds length limit [CR 406638]

In Sybase IQ 12.6, the maximum size of CHAR data changed from 255 bytes to 32K, as described in *New Features in Sybase IQ 12.6*. Unless a constant expression is the supplied LENGTH argument to a function that produces a VARCHAR result (such as SPACE or REPEAT), the default length is the maximum allowed.

Sybase IQ 12.6 queries containing one or more of such functions may return one of the following errors:

```
ASA Error -1009080: Key doesn't fit on a single
database page: 65560(4, 1)
```

```
ASA Error -1009119: Record size too large for
database page size.
```

For example:

```
SELECT COUNT(*) FROM test1 a WHERE
(a.col1 + SPACE(4-LENGTH(a.col1)) +
a.col2 + space(2- LENGTH(a.col2)))
IN (SELECT (b.col3) FROM test1 b);
```

To avoid such errors, cast the function result with an appropriate maximum length, for example:

```
SELECT COUNT(*) FROM test1 a WHERE
(a.col1 + CAST(SPACE(4-LENGTH(a.col1))
AS VARCHAR(4)) + a.col2 +
CAST(SPACE(2-LENGTH (a.col2))
AS VARCHAR(4))) IN (SELECT (b.col3) FROM test1 b);
```

The errors are more likely with an IQ page size of 64K or a multibyte collation.

5.6.2 RAND function results changed [CR 397838]

The values returned by RAND vary depending on whether you use a FROM clause or not and whether the referenced table was created in SYSTEM or in an IQ store.

If RAND is called with a constant or no argument in a query containing only tables in IQ stores, Sybase IQ evaluates the function once per query.

If RAND is called with a FROM clause and an argument in a query containing only tables in IQ stores, the function returns an arbitrary but repeatable value.

When no argument is called, RAND is a non-deterministic function. Successive calls to RAND may return different values. The query optimizer does not cache the results of the RAND function.

The following example will be added to the *Sybase IQ Reference Manual*:

The following statement returns an arbitrary 5% sampling of a table:

```
SELECT AVG(table1.number_of_cars),
       AVG(table1.number_of_tvs)
FROM table1 WHERE RAND(ROWID(table1)) < .05
and table1.income < 50000;
```

5.6.3 ENABLED_ORDERED_PUSHDOWN_INSERTION option added [CR 392905]

Function	The ENABLE_ORDERED_PUSHDOWN_INSERTION option was added for ESD #5. This option controls how the query optimizer adds in the semijoin predicates for push-down joins selected by the join optimizer.
Allowed values	ON, OFF
Default	OFF
Scope	Can be set temporary, for an individual connection, or for the PUBLIC group. Takes effect immediately.
Description	Change this option only if Sybase Support instructs you to do so. If OFF (the default), this option reverts the optimizer to the behavior seen in Sybase IQ 12.6 GA. If ON, this option allows the insertion of semijoin predicates in projection-savings order.

5.6.4 Database options for push-down join optimizer tables [CR 368261]

New database options provide a lower bound, expressed in parts-per-million, for the tables into which the optimizer is willing to consider using push-down joins. These options only affect the optimizer when the target table is very large. Under normal circumstances, you should not need to change these options.

- MIN_SMPDJ_OR_HPDJ_FILTERED_PPM
UNSIGNED INT Default: 2500

Specifies the minimum percentage of rows that must remain after all simple local predicates (expressed in parts-per-million) before the optimizer will consider using either the hash push-down join (HPDJ) or sort-merge push-down join (SMPDJ) algorithms. The default is equivalent to a selectivity of 0.0025, or one quarter of one percent of the table.

- MIN_NLPDJ_FILTERED_PPM

UNSIGNED INT Default: 2500

Specifies the minimum percentage of rows remaining after all simple local predicates (expressed in parts-per-million) before the optimizer will consider using the nested-loop push-down join (NLPDJ) algorithm.

5.6.5 New database option DEFAULT_HAVING_SELECTIVITY

Function	Provides default selectivity estimates to the optimizer for most HAVING clauses.
Allowed values	0 to 100
Scope	Can be set for an individual connection or the PUBLIC group. Takes effect immediately.
Default	0
Description	DEFAULT_HAVING_SELECTIVITY sets the selectivity for HAVING clauses, overriding optimizer estimates. A HAVING clause filters the results of a GROUP BY clause or a query with a select list consisting solely of aggregate functions. When DEFAULT_HAVING_SELECTIVITY is set to the default of 0, the optimizer estimates how many rows will be filtered by the HAVING clause. Sometimes the IQ optimizer does not have sufficient information to choose an accurate selectivity, and in these cases will choose a generic estimate of 40%. DEFAULT_HAVING_SELECTIVITY allows a user to replace the optimizer estimate for all HAVING predicates in a query. Users can also specify the selectivity of individual HAVING clauses in the query, as described in the section “User-supplied estimates” of the “Search conditions” section in Chapter 3, “SQL Language Elements.”
See also	Chapter 3, “Improving Query Performance” in the <i>Sybase IQ Performance and Tuning Guide</i>

Note This option was omitted from the chapter “Database Options” in the *Sybase IQ Reference Manual*.

5.7 Security enhancements

This section contains new features and changes related to security.

5.7.1 JDK/JRE patch level updated for Daylight Saving Time change [CR 448643]

The JDK/JRE patch level has been updated to 1.4.2_SR6 for IBM AIX due to new Daylight Saving Time start and end dates resulting from the United States Energy Policy Act.

5.7.2 GROUP_MEMBER function added [CR 393686]

In ESD #4.1, the following new function was added:

GROUP_MEMBER function [System]

Function	Identifies whether the user belongs to the specified group.
Syntax	GROUP_MEMBER (<i>group-name-string-expression</i> [, <i>user-name-string-expression</i>])
Parameters	group-name-string-expression Identifies the group to be considered. user-name-string-expression Identifies the user to be considered. If not supplied, then the current user name is assumed.
Return value	• 0 Returns 0 if the group does not exist, if the user does not exist, or if the user does not belong to the specified group. • 1 Returns an integer other than 0 if the user is a member of the specified group.

5.7.3 Login expiration [CR 392592]

A userid/password created using `sp_iqaddlogin` and set to expire in one day is now defined as valid all day tomorrow and invalid on the following day. In other words, a login created today and set to expire in n days will not be usable once the date changes to the $(n+1)$ th day.

5.7.4 New output from `sp_iqlistpasswordexpirations` [CR 369558]

The `sp_iqlistpasswordexpirations` stored procedure output has changed in this release. The column `Days_till_Expiration` now is the number of days until the password expires. If the password has expired, the column returns “Expired”. If the password does not expire, the column returns zero (0). A new column, `Password_Expiration_Interval`, holds the former contents of `Days_till_Expiration`, the total number of days the password is valid on its creation date. A zero (0) means that the password does not expire.

5.8 Sybase IQ documentation on CD

In Sybase IQ 12.6, a SyBooks CD replaces the Technical Library CD. The SyBooks browser software runs on Windows and Linux platforms. Users with non-Linux UNIX platforms must use Acrobat Reader to open PDF files on the SyBooks CD.

The Sybase Product Manuals Web site is an online version of the SyBooks CD that you can access using a standard Web browser. To access the Sybase Product Manuals Web site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

6. Product compatibilities

This section describes the compatibility of Sybase IQ with other products. Certification with client and partner products is an ongoing process. For the latest information on versions of Sybase products and client application tools that have been certified with Sybase IQ 12.6, refer to the Sybase Web site. See the section “Sybase certifications on the Web” on page 102 for instructions on accessing online certification information.

6.1 Sybase products

The following Sybase products are certified with Sybase IQ 12.6:

- Adaptive Server Anywhere
- Adaptive Server Enterprise
- InfoMaker
- jConnect for JDBC
- Open Client

- PowerBuilder
- PowerDesigner
- Replication Server

6.2 Client applications

The following third party client application tools have been certified with Sybase IQ 12.6. An asterisk (*) indicates that certification is ongoing. Refer to the section “Sybase certifications on the Web” on page 102 for instructions on accessing the latest online certification information.

- Ascential Software DataStage *
- Business Objects
- Business Objects Crystal Reports
- Business Objects Data Integrator
- Cognos PowerPlay
- ColdFusion MX
- Informatica PowerCenter
- Microsoft Access
- Microsoft SQL Server
- Microsoft Data Transformation Services*
- MicroStrategy*
- SAS*
- SPSS Software Clementine

6.3 EMC Symmetrix and Hitachi certification

EMC Symmetrix and Hitachi are certified with Sybase IQ and are supported on all platforms. For details about the EMC Symmetrix certification, refer to the white paper “Configuring and Performance Tuning EMC Symmetrix for Sybase IQ Multiplex” available on the Sybase Web site.

7. Restrictions

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

The information in this section applies to Version 12.6 servers and databases that have been upgraded from any earlier version of Sybase IQ, unless specified otherwise.

7.1 Resolving static and dynamic collisions [CR 405314]

A static collision is a database state that prevents a DDL statement from executing, even if the statement is retried. For example, attempting to drop a user data type fails if that data type is in use in a table definition. This situation is static, because retrying the statement never succeeds, even after all users are disconnected or after the server is restarted. In an interactive setting, a static collision results in a SQL error.

A dynamic collision is a database server state that prevents a DDL statement from executing at the moment, but eventually clears on its own and allows the statement to succeed, if the statement is retried. For example, altering a table fails if the table is in use in a query by another user. This situation is dynamic, because retrying the statement succeeds once other operations complete, after all users are disconnected, or after the server is restarted. In an interactive setting, a dynamic collision results in a SQL error.

In a multiplex environment, the query server resolves a static collision that occurs while executing a DDL statement propagated from the write server, by renaming objects that conflict. The object (a table or domain) gets a new name made from the original name with the character “~” (tilde) appended, followed optionally by a sequence number. This creates a unique name that is unlikely to look like a name users would choose.

For information on objects that have been renamed due to static collisions, run the stored procedure `sp_iqmpxrenameinfo`. This procedure returns both the old name and the new name of the object.

To resolve a dynamic collision during DDL execution, the query server finds the connections responsible and disconnects them.

The disconnect is logged in the `.iqmsg` file. Here is an example of a message that appears in the `.iqmsg` when a dynamic collision occurs:

```
Shared IQ Store update DDL statement:  
drop table DBA.gtt44  
Disposition: SQLSTATE:42W21 --
```

```
dropped 1 connection(s) for table:  
DBA.gtt44 Retry successful
```

Users can avoid collisions by managing domain and table names carefully across the members of the multiplex. To avoid the inconvenience of disconnects, DDL at the write server that affects normal query operation on the multiplex servers should be scheduled to avoid active users.

7.2 CREATE SCHEMA prohibited on active multiplex [CR 366937]

The CREATE SCHEMA command is not valid on an active multiplex. It returns the error:

```
create schema command not allowed on an active multiplex  
server
```

7.3 Dbospace management and file placement

When you allocate file system files for dbspaces (System, IQ Main or IQ Temporary), do not place the files on a file system that is shared over a local area network. Violating this rule can lead to poor I/O performance and other problems, including overloading the local area network. On UNIX and Linux platforms, avoid Network File System (NFS) mounted file systems. On Windows, do not place dbospace files on network drives owned by another server.

Dbospace management should be performed by a single database administrator on a single connection to avoid conflicts.

7.4 Query restrictions

By default Sybase IQ cursors are scrollable, meaning that Sybase IQ keeps all the query results in a buffer so that you can scroll backwards. If the query returns very large numbers (millions) of rows of output, you can improve performance by issuing the following command before running the query:

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

Some front-end applications make use of scrolling cursor operations, however, and require this option to be set "OFF".

If scrolling cursors are never used in your application, you should make this a permanent public option. It will use less memory and make a modest improvement in query performance.

7.5 User-defined functions and query results

In very few cases, differences in semantics between Adaptive Server Anywhere and Sybase IQ can produce different results for a query if it is issued in a user-defined function. For example, IQ treats the CHAR and VARCHAR data types as distinct and different, while Anywhere treats CHAR data as if it were VARCHAR.

7.6 Interactive SQL restrictions

Sybase recommends that you use the Java edition of DBISQL, rather than DBISQLC or isql, unless instructed otherwise for specific tasks. Note that although dbisqlc is supported, dbisqlc does not contain all the features of the Java version dbisql.

- Do not use the Import option in the DBISQL Data menu (or Command > Options > Input Format in DBISQLC). This option is not supported for use with IQ databases. Use the LOAD TABLE statement or the INSERT statement to load data into IQ tables.
- If you run DBISQL (Interactive SQL Java) with the -q (quiet mode) option and the data extraction commands (primarily setting the option TEMP_EXTRACT_NAME1 to an output file) are in a command file, you must first set and make permanent the DBISQL option “Show multiple result sets.” If this option is not set, the output file is not created.

To set the “Show multiple result sets” option, click Tools > Options in the DBISQL window, then check the box “Show multiple result sets” and click “Make permanent.”

- In DBISQL and DBISQLC, if you click the Help button or the Online Books button, the error “Cannot open the file” or “Cannot open help” is returned. Online help is not available for DBISQL and DBISQLC. [CR 370040]
- The query plan in the DBISQL Plan tab is an Adaptive Server Anywhere style query plan. For a Sybase IQ query plan, refer to the IQ *.iqmsg* file.
- If you set the terminal type on UNIX and Linux systems to “dumb” or “unknown” and then start DBISQLC, Sybase IQ returns an error. For example:

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

To avoid problems, run DBISQL (Interactive SQL Java) instead, or use an xterm window to run DBISQLC on UNIX and Linux systems. For example, you can start an xterm window with a scroll bar as follows:

```
% xterm -sb
```

7.7 SQL command line length is restricted

A current restriction limits the SQL command line length to the Catalog page size (not IQ page size) of the connected database. Because the default Catalog page size is 4KB, this is only a problem when connected to a database such as `utility_db`, which has a page size of 1024. This restriction may cause RESTORE commands that reference numerous dbspaces to fail.

To avoid the problem, make sure the length of SQL command lines is less than the Catalog page size.

8. Known problems

For a description of known issues in Sybase IQ 12.6, see the following sections. If there is a workaround for a problem, it is provided. Change request numbers appear in brackets, for example [CR 235449].

See also “Restrictions” on page 39 for more information.

“Documentation updates and clarifications” on page 52 contains details that were not documented in time for this release.

8.1 Sybase IQ operations

This section reports workarounds for issues with Sybase IQ operations.

8.1.1 Time conversions in the Java JDK/JRE

The Java Run-time Environment (JRE) contains library functions for time conversion from Coordinated Universal Time (also called GMT) to local time, including any Daylight Saving Time compensation. Patches to the JDK/JRE will be issued as necessary by Sun and Sybase, Inc. to update these functions to reflect Daylight Savings Time changes caused by the United States Energy Policy Act of 2006.

8.1.2 CIS returns “No Suitable Driver” error [CR 414501]

When using Component Integration Services (CIS) in certain geographic regions, connection attempts return the error No Suitable Driver. Java Development Kits used with Sybase IQ 12.6 support only time zone codes shown in Table 6 and Table 7.

- For databases using default JDK 1.1.8:

Substitute JST for unsupported time zone KST, which gives the same GMT+9 time, as follows:

```
setenv TZ JST
```

- For databases using JDK 1.3:

```
setenv TZ Asia/Seoul
```

Set the time zone environment variable to a supported setting, start the server, and CIS works as expected. To ensure that the correct setting is always used, you can set the time zone in a star-up script such as *.cshrc* or *.login*.

See Table 6 and Table 7 for the appropriate time zone code and settings.

Table 6: Time zone settings for JDK 1.1.8

Time zone setting	Time zone code	Time zone setting	Time zone code
0	GMT	16	SST
1	UTC	17	NST
2	ECT	18	MIT
3	EET	19	HST
4	ART	20	AST
5	EAT	21	PST
6	MET	22	PNT
7	NET	23	MST
8	PLT	24	CST
9	IST	25	EST
10	BST	26	IET
11	VST	27	PRT
12	CTT	28	CNT
13	JST	29	AGT
14	ACT	30	BET
15	AET	31	CAT

Table 7: Time zone settings for JDK 1.3

Time zone setting	Time zone code	Time zone setting	Time zone code
0	Pacific/Niue	161	Europe/Belgrade
1	Pacific/Apia	162	Europe/Paris
2	MIT	163	ECT
3	Pacific/Pago_Pago	164	Africa/Bujumbura
4	Pacific/Tahiti	165	Africa/Gaborone
5	Pacific/Fakaofu	166	Africa/Lubumbashi
6	Pacific/Honolulu	167	Africa/Maseru
7	HST	168	Africa/Blantyre
8	America/Adak	169	Africa/Maputo
9	Pacific/Rarotonga	170	Africa/Kigali
10	Pacific/Marquesas	171	Africa/Khartoum
11	Pacific/Gambier	172	Africa/Mbabane
12	America/Anchorage	173	Africa/Lusaka
13	AST	174	Africa/Harare
14	Pacific/Pitcairn	175	CAT
15	America/Vancouver	176	Africa/Johannesburg
16	America/Tijuana	177	Europe/Sofia
17	America/Los_Angeles	178	Europe/Minsk
18	PST	179	Asia/Nicosia
19	America/Dawson_Creek	180	Europe/Tallinn
20	America/Phoenix	181	Africa/Cairo
21	PNT	182	ART
22	America/Edmonton	183	Europe/Helsinki
23	America/Mazatlan	184	Europe/Athens
24	America/Denver	185	Asia/Jerusalem
25	MST	186	Asia/Amman
26	America/Belize	187	Asia/Beirut
27	America/Regina	188	Europe/Vilnius
28	Pacific/Galapagos	189	Europe/Riga
29	America/Guatemala	190	Europe/Chisinau
30	America/Tegucigalpa	191	Europe/Bucharest
31	America/El_Salvador	192	Europe/Kaliningrad
32	America/Costa_Rica	193	Asia/Damascus
33	America/Winnipeg	194	Europe/Kiev
34	Pacific/Easter	195	Europe/Istanbul
35	America/Mexico_City	196	EET

Time zone setting	Time zone code	Time zone setting	Time zone code
36	America/Chicago	197	Asia/Bahrain
37	CST	198	Africa/Djibouti
38	America/Porto_Acre	199	Africa/Asmera
39	America/Bogota	200	Africa/Addis_Ababa
40	America/Guayaquil	201	EAT
41	America/Jamaica	202	Africa/Nairobi
42	America/Cayman	203	Indian/Comoro
43	America/Managua	204	Asia/Kuwait
44	America/Panama	205	Indian/Antananarivo
45	America/Lima	206	Asia/Qatar
46	America/Indianapolis	207	Africa/Mogadishu
47	IET	208	Africa/Dar_es_Salaam
48	America/Nassau	209	Africa/Kampala
49	America/Montreal	210	Asia/Aden
50	America/Havana	211	Indian/Mayotte
51	America/Port-au-Prince	212	Asia/Riyadh
52	America/Grand_Turk	213	Asia/Baghdad
53	America/New_York	214	Europe/Simferopol
54	EST	215	Europe/Moscow
55	America/Antigua	216	Asia/Tehran
56	America/Anguilla	217	MET
57	America/Curacao	218	Asia/Dubai
58	America/Aruba	219	Indian/Mauritius
59	America/Barbados	220	Asia/Muscat
60	America/La_Paz	221	Indian/Reunion
61	America/Manaus	222	Indian/Mahe
62	America/Dominica	223	Asia/Yerevan
63	America/Santo_Domingo	224	NET
64	America/Grenada	225	Asia/Baku
65	America/Guadeloupe	226	Asia/Aqtau
66	America/Guyana	227	Europe/Samara
67	America/St_Kitts	228	Asia/Kabul
68	America/St_Lucia	229	Indian/Kerguelen
69	America/Martinique	230	Asia/Tbilisi
70	America/Montserrat	231	Indian/Chagos
71	America/Puerto_Rico	232	Indian/Maldives
72	PRT	233	Asia/Dushanbe

Time zone setting	Time zone code	Time zone setting	Time zone code
73	America/Port_of_Spain	234	Asia/Ashkhabad
74	America/St_Vincent	235	Asia/Tashkent
75	America/Tortola	236	Asia/Karachi
76	America/St_Thomas	237	PLT
77	America/Caracas	238	Asia/Bishkek
78	Antarctica/Palmer	239	Asia/Aqtobe
79	Atlantic/Bermuda	240	Asia/Yekaterinburg
80	America/Cuiaba	241	Asia/Calcutta
81	America/Halifax	242	IST
82	Atlantic/Stanley	243	Asia/Katmandu
83	America/Thule	244	Antarctica/Mawson
84	America/Asuncion	245	Asia/Thimbu
85	America/Santiago	246	Asia/Colombo
86	America/St_Johns	247	Asia/Dacca
87	CNT	248	BST
88	America/Fortaleza	249	Asia/Almaty
89	America/Cayenne	250	Asia/Novosibirsk
90	America/Paramaribo	251	Indian/Cocos
91	America/Montevideo	252	Asia/Rangoon
92	America/Buenos_Aires	253	Indian/Christmas
93	AGT	254	Asia/Jakarta
94	America/Godthab	255	Asia/Phnom_Penh
95	America/Miquelon	256	Asia/Vientiane
96	America/Sao_Paulo	257	Asia/Saigon
97	BET	258	VST
98	America/Noronha	259	Asia/Bangkok
99	Atlantic/South_Georgia	260	Asia/Krasnoyarsk
100	Atlantic/Jan_Mayen	261	Antarctica/Casey
101	Atlantic/Cape_Verde	262	Australia/Perth
102	America/Scoresbysund	263	Asia/Brunei
103	Atlantic/Azores	264	Asia/Hong_Kong
104	Africa/Ouagadougou	265	Asia/Ujung_Pandang
105	Africa/Abidjan	266	Asia/Macao
106	Africa/Accra	267	Asia/Kuala_Lumpur
107	Africa/Banjul	268	Asia/Manila
108	Africa/Conakry	269	Asia/Singapore
109	Africa/Bissau	270	Asia/Taipei

Time zone setting	Time zone code	Time zone setting	Time zone code
110	Atlantic/Reykjavik	271	Asia/Shanghai
111	Africa/Monrovia	272	CTT
112	Africa/Casablanca	273	Asia/Ulan_Bator
113	Africa/Timbuktu	274	Asia/Irkutsk
114	Africa/Nouakchott	275	Asia/Jayapura
115	Atlantic/St_Helena	276	Asia/Pyongyang
116	Africa/Freetown	277	Asia/Seoul
117	Africa/Dakar	278	Pacific/Palau
118	Africa/Sao_Tome	279	Asia/Tokyo
119	Africa/Lome	280	JST
120	GMT	281	Asia/Yakutsk
121	UTC	282	Australia/Darwin
122	Atlantic/Faeroe	283	ACT
123	Atlantic/Canary	284	Australia/Adelaide
124	Europe/Dublin	285	Australia/Broken_Hill
125	Europe/Lisbon	286	Australia/Hobart
126	Europe/London	287	Antarctica/DumontDURville
127	Africa/Luanda	288	Pacific/Truk
128	Africa/Porto-Novo	289	Pacific/Guam
129	Africa/Bangui	290	Pacific/Saipan
130	Africa/Kinshasa	291	Pacific/Port_Moresby
131	Africa/Douala	292	Australia/Brisbane
132	Africa/Libreville	293	Asia/Vladivostok
133	Africa/Malabo	294	Australia/Sydney
134	Africa/Niamey	295	AET
135	Africa/Lagos	296	Australia/Lord_Howe
136	Africa/Ndjamena	297	Pacific/Ponape
137	Africa/Tunis	298	Pacific/Efate
138	Africa/Algiers	299	Pacific/Guadalcanal
139	Europe/Andorra	300	SST
140	Europe/Tirane	301	Pacific/Noumea
141	Europe/Vienna	302	Asia/Magadan
142	Europe/Brussels	303	Pacific/Norfolk
143	Europe/Zurich	304	Pacific/Kosrae
144	Europe/Prague	305	Pacific/Tarawa
145	Europe/Berlin	306	Pacific/Majuro

Time zone setting	Time zone code	Time zone setting	Time zone code
146	Europe/Copenhagen	307	Pacific/Nauru
147	Europe/Madrid	308	Pacific/Funafuti
148	Europe/Gibraltar	309	Pacific/Wake
149	Europe/Budapest	310	Pacific/Wallis
150	Europe/Rome	311	Pacific/Fiji
151	Europe/Vaduz	312	Antarctica/McMurdo
152	Europe/Luxembourg	313	Asia/Kamchatka
153	Africa/Tripoli	314	Pacific/Aucklandoo
154	Europe/Monaco	315	NST
155	Europe/Malta	316	Pacific/Chatham
156	Africa/Windhoek	317	Pacific/Enderbury
157	Europe/Amsterdam	318	Pacific/Tongatapu
158	Europe/Oslo	319	Asia/Anadyr
159	Europe/Warsaw	320	Pacific/Kiritimati
160	Europe/Stockholm		

8.1.3 Nested stored procedure with dynamic result set returns errors [CR 401334]

A data truncation error or conversion error occurs when a procedure calls another procedure with a dynamic result set and all of the following are true:

- The IQ server is version 12.5
- DBISQL Java connects through iAnywhere JDBC driver
- DBISQL Java version is higher than 7.04.

The problem doesn't happen if DBSIQL Java connects through jConnect5 or the ODBC driver or if Sybase IQ 12.6 is used with DBISQL 9.0.1.

There are several ways to avoid the problem:

- Connect DBISQL Java through jConnect5 or ODBC driver.
- Use Sybase IQ 12.6 with DBISQL version 9.0.1 or use Sybase IQ 12.5 with DBISQL version 7.0.4
- Add a statement like the following to the start of the procedure, to keep the server from adding a result set:

```
IF 1 = 0 THEN
SELECT 1 AS a FROM nosuchtable;
END IF;
```


8.1.4 Connecting to IQ 12.6 using ODBC DSN for 12.5 [CR 376618]

You cannot connect to a Sybase IQ 12.6 server using an ODBC Data Source Name (DSN) created for a 12.5 server, even if you specify the same server name, port number, and database name.

Use the ODBC Data Source Administrator to remove DSNs created in 12.5 and add new DSNs.

8.1.5 IQ plug-in on non-Linux machines cannot connect to IQ Agent on Linux [CR 362017]

When creating a new database, if you specify a remote Linux host as the IQ Agent host server, the Create Database Server wizard fails with an error like the following:

```
Unable to connect to the IQ Agent on host "<hostname>"
```

An error referencing host 127.0.0.1 may also appear in the server log.

To prevent the problem:

- 1 On the Linux machine, use the `ypmatch` command to locate its IP address. For example, for system *maynardqa*:

```
%ypmatch maynardqa hosts
10.47.81.70 maynardqa
```

- 2 Edit the `/etc/hosts` file on the Linux machine and replace the IP address of 127.0.0.1 with the true IP address of the Linux host.

8.2 Data definition

This section reports workarounds for issues with data definition.

8.2.1 Dropping dbspaces that contain pages from TLV log [CR 405730]

If a write server is started in single-node mode, a user may be unable to drop some of the dbspaces if they contain pages belonging to the TLV (table version) log. In single-node mode, IQ does not update the TLV log or explicitly relocate pages when a dbspace is changed to relocate mode.

To prevent this situation, start the database in multiplex mode before dropping the dbspace. When the write server starts in multiplex after being run in single-node mode, it truncates the TLV log and the dbspace can be dropped.

8.2.2 Dropping the last local dbspace on query server [CR 363338]

When a dbspace is created, IQ needs to save some checkpoint information for the transaction control. If the dbspace is the only one for the type of store (IQ Local Store, in this case), the checkpoint data is saved in the very same dbspace. This makes the dbspace “in use”, even though, from a user’s point of view, nothing has been done yet. The dbspace is therefore protected from being dropped.

To release the checkpoint information saved in the dbspace, a DBA needs to issue a CHECKPOINT SQL statement to force IQ to flush out the checkpoint data. If there are multiple connections to the server, only the CHECKPOINT statement that is executed *after* the completion of all the transactions that temporally overlap with the transaction of creating the dbspace will flush the checkpoint data. Therefore a DBA may have to issue CHECKPOINT several times before he or she can successfully drop the dbspace.

A stored procedure, `sp_iqdbspace`, reports the current status of all the IQ dbspaces, including whether or not a dbspace is ready to be dropped. Please see “Dropping a dbspace,” in Chapter 5, “Working with Database Objects” in *Sybase IQ System Administration Guide* for details on the `sp_iqdbspace` returned results.

If the IQ Local Store to be dropped is not empty, you must do the following before you can drop it:

- 1 Alter the database to relocate mode. For details, see “Altering dbspaces” in Chapter 6, “Managing Dbspaces” in *Introduction to Sybase IQ*.
- 2 Run `sp_iqrelocate` on the database. For details, see Chapter 9, “System Procedures” in *Sybase IQ Reference Manual*.

9. Documentation for this version

All documentation for Version 12.6 is included with your product shipment on CD:

- The Sybase IQ Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the SyBooks CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader, which you can download at no charge from the Adobe Web site using a link provided on the CD.

The Sybase IQ Getting Started CD also contains release bulletins and installation guides in an HTML format that is specialized for accessibility. You can navigate the HTML with an adaptive technology such as a screen reader, or view it with a screen enlarger.

- The Sybase IQ SyBooks CD contains product manuals and is included with your software. The Eclipse-based SyBooks browser allows you to access the manuals in an easy-to-use, HTML-based format.

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

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

Note The SyBooks browser software runs on Windows and Linux platforms. Users with non-Linux UNIX platforms must use Acrobat Reader to open PDF files on the SyBooks CD.

See Table 8 for a list of documentation on the Sybase IQ 12.6 Sybooks CD.

Table 8: Sybase IQ Sybooks contents

Part Number	Book Title
DC00171-01-1260-01	<i>New Features in Sybase IQ 12.6</i>
DC38159-01-1260-02	<i>Introduction to Sybase IQ</i>
DC00170-01-1260-02	<i>Sybase IQ System Administration Guide</i>
DC00169-01-1260-02	<i>Sybase IQ Performance and Tuning Guide</i>
DC38151-01-1260-02	<i>Sybase IQ Reference Manual</i>
DC36584-01-1260-02	<i>Sybase IQ Troubleshooting and Error Messages Guide</i>
DC00168-01-1260-01	<i>Sybase IQ Utility Guide</i>
DC00172-01-1260-01	<i>Large Objects Management in Sybase IQ</i>

A hard copy edition of the Sybase IQ documentation set, 48200-01-1260-01, can be ordered separately.

You may also need to refer to the documentation for Adaptive Server Anywhere:

- *Adaptive Server Anywhere Programming Guide*

Intended for application developers writing programs that directly access the ODBC, Embedded SQL™, or Open Client™ interfaces, this book describes how to develop applications for Adaptive Server Anywhere.

- *Adaptive Server Anywhere Database Administration Guide*
Intended for all users, this book covers material related to running, managing, and configuring databases and database servers.
- *Adaptive Server Anywhere Error Messages*
This book lists all Adaptive Server Anywhere error messages with diagnostic information.
- *Adaptive Server Anywhere SQL Reference Manual*
Intended for all users, this book provides a complete reference for the SQL language used by Adaptive Server Anywhere. It also describes the Adaptive Server Anywhere system tables and procedures.

You can also refer to the Adaptive Server Anywhere documentation in the SQL Anywhere Studio 9.0.1 collection on the Sybase Product Manuals Web site. To access this site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

10. Documentation updates and clarifications

This section contains information omitted from documentation and new information that needs emphasis.

See *New Features in Sybase IQ 12.6* for a description of changes to the Sybase IQ documentation set in version 12.6.

10.1 Customer issues

The *CustomerIssues.txt* file lists customer issues resolved in Sybase IQ 12.6. Issues resolved in a subsequent software update (ESD or EBF) are documented in the cover letter for that update. *CustomerIssues.txt* is installed in the following directories:

- `$SYBASE/ASIQ-12.6` on UNIX
- `%SYBASE%/ASIQ-12.6` on Windows

For more information, see “Sybase EBFs and software maintenance” on page 103.

10.2 Interactive SQL (DBISQL) online help

Online help is not available for the utilities Interactive SQL (DBISQLC) and Interactive SQL Java (DBISQL).

10.3 Introduction to Sybase IQ

This section contains updates to the *Introduction to Sybase IQ*.

10.3.1 Reserving space for dbspace expansion [CR 394000]

In the Chapter 6 “Managing Dbspaces” section “Creating dbspaces,” the sentence “If desired, specify the size in megabytes of free list space to reserve, so that the dbspace can be increased in size in the future.” should say “If desired, specify the size in megabytes of space to reserve, so that the dbspace can be increased in size in the future.” The words “free list” should be removed from the sentence.

10.3.2 Fast Launcher option [CR 386628]

The Fast Launcher option for Sybase Central is only available on Windows. This information was omitted from the *Introduction to Sybase IQ*.

10.4 Installation and Configuration Guide

For updates about installation and migration, including database upgrades, see “Special installation and migration instructions” on page 6.

10.5 Troubleshooting and Error Messages Guide

This section contains updates to the *Sybase IQ Troubleshooting and Error Messages Guide*.

10.5.1 “Not enough space” error when synchronizing multiplex [CR 374292]

A multiplex synchronize may fail with the following error in the Agent log file:

```
java.io.IOException: Not enough space
```

This indicates an out of memory condition, not lack of disk space or dbspace. This was omitted from Chapter 1, *Sybase IQ Troubleshooting and Error Messages Guide*.

10.5.2 Database start-up after recovery [CR 351590]

When performing forced recovery or leaked blocks recovery, you must start the database with the “.db” extension, not “.DB”. For example:

```
start_asiq -n my_db_server -x 'tcpip{port=7934}'  
-gd dba -iqdroppls my_db /work/database/my_db.db
```

This information was omitted from “Repairing allocation problems” in Chapter 2, *Sybase IQ Troubleshooting and Error Messages Guide*.

10.5.3 Error message descriptions

The error message descriptions in Chapter 6, “Database Error Messages” have not been fully updated for version 12.6. For additional information on error messages, see *Adaptive Server Anywhere Error Messages*.

10.6 System Administration Guide

This section contains updates to the *Sybase IQ System Administration Guide*.

10.6.1 Files located relative to Catalog Store [CR 469000]

The information in “Example 2 — Moving the Catalog Store” in the “Moving database files” section of the “The RESTORE statement” section under “Restoring your databases” in Chapter 14, “Backup and Data Recovery” is not correct regarding the directory in which a file is restored.

- In the second bullet, the sentences “Files originally created relative to the IQ_SYSTEM_MAIN dbspace, which holds the main IQ Store file, are restored relative to the main IQ Store file.” will be replaced with the following:

“Files originally created relative to the Catalog Store are restored relative to the Catalog Store.”
- The sentence “Relative pathnames in the RENAME clause work as they do when you create a database or dbspace: the main IQ Store dbspace, Temporary Store dbspaces, and Message Log are restored relative to the location of *db_file* (the Catalog Store); user-created IQ Store dbspaces are restored relative to the directory that holds the main IQ dbspace.” will be replaced with the following:

Relative pathnames in the RENAME clause work as they do when you create a database or dbspace: the main IQ Store dbspace, Temporary Store dbspaces, and Message Log are restored relative to the location of *db_file* (the Catalog Store); user-created IQ Store dbspaces are restored relative to the directory that holds the Catalog Store.”

10.6.2 Unique table names on a connection [CR 455684]

In Sybase IQ 12.6 ESD #10, an attempt to create a base table or a global temporary table will fail, if a local temporary table of the same name exists on that connection, as the new table cannot be uniquely identified by *owner.table*.

The information in the section “Unique table names on a connection [CR 455684]” on page 17 will be added in Chapter 5, “Working with Database Objects,” to the “Types of tables” section in “Creating tables” in the “Working with tables” section.

10.6.3 LOAD TABLE FORMAT syntax correction [CR 445737]

In the section “Bulk loading data using the LOAD TABLE statement” in Chapter 7, “Moving Data In and Out of Databases,” the syntax for the FORMAT option of the LOAD TABLE statement should not contain single quotes. The correct syntax is:

```
LOAD TABLE [ owner.table-name
[ ( load-specification, ... ) ]
FROM 'filename-string', ...
[ FORMAT { ascii | binary } ]
... [ DELIMITED BY string ]
... [ STRIP { ON | OFF } ]
... [ QUOTES { ON | OFF } ]
... [ ESCAPES { ON | OFF } ]
... [ ESCAPE CHARACTER character ]
... [ WITH CHECKPOINT ON|OFF ]
... [ load-options ]
```

10.6.4 Group permissions cannot be revoked for a specific user [CR 444928]

When users are assigned membership in a group, they inherit all the permissions on tables, views, and procedures associated with that group. If you do not want a specific user to access a particular table, view, or procedure, then do not make that user a member of a group that has permissions on that object. You cannot revoke permissions for a specific user within a group.

10.6.5 Inserting ASE DATE and TIME type data [CR 442348]

12.6 ESD #9 supports inserting data from an Adaptive Server Enterprise database column of data type DATE or TIME using the INSERT...LOCATION syntax of the INSERT statement.

Add the following rows to Table 7-18: DATE/TIME data types in Chapter 7, “Moving Data In and Out of Databases.”

Adaptive Server Enterprise Datatype	Sybase IQ Datatype	Notes
date	date	You can insert data from an ASE database column of data type date using INSERT...LOCATION.
time	time	The Sybase IQ data type is the Time of day, containing hour, minute, second, and fraction of a second. The fraction is stored to 6 decimal places. A time value requires 8 bytes of storage. The Adaptive Server Enterprise data type is between 00:00:00:000 and 23:59:59:999. You can use either military time or 12AM for noon and 12PM for midnight. A time value must contain either a colon or the AM or PM signifier. AM or PM may be in either uppercase or lowercase. A time value requires 4 bytes of storage. You can insert data from an ASE database column of data type time using INSERT...LOCATION.

10.6.6 TIME to DATE conversion unsupported for INSERT and UPDATE [CR 442348]

In Table 7-8: IQ conversions for INSERT and UPDATE in Chapter 7, “Moving Data In and Out of Databases,” the conversion in the row From: time To: dt (date) should be changed to “U” (unsupported) from “E” (explicit).

10.6.7 ESCAPE_CHARACTER option for system use only [CR 440545]

The ESCAPE_CHARACTER option is reserved for system use only. Do not change the setting of this option.

In Chapter 5, “Working with Database Options,” the paragraph following Table 5-2 should not refer to the ESCAPE_CHARACTER option. The paragraph should say:

“On Windows systems, when you specify device names that include a backslash, you must double the backslash to keep the system from mistaking a backslash/letter combination for an escape sequence such as tab or newline command.”

10.6.8 Multiplex servers and remote data access [CR 426772]

This section was omitted from the *Sybase IQ System Administration Guide*. It should appear after the section “Using remote procedure calls (RPCs)” in Chapter 16, “Accessing Remote Data”.

Sybase IQ supports remote data access among servers in a multiplex configuration in all directions:

- Write server to query server
- Query server to write server
- Query server to query server

For example, the write server may be defined as a remote server to a query server and proxy tables may be defined on the query server for tables on the write server.

Consider the following when using remote data access among multiplex servers:

- Use server class ASAODBC or ASAJDBC for IQ servers
- Use proxy table names that differ from local table names to avoid multiplex static collisions. For example, if a query server local store has a table named *employee*, then the proxy table name should not be *employee*.
- Add remote definitions created on the query servers to the `sp_mpxcfg_<servername>` stored procedure so that they persist after a multiplex SYNCHRONIZE. If you do not add the definitions to the `sp_iqmpxcfg_<servername>` procedure, they disappear after SYNCHRONIZE.

For more information on the `sp_mpxcfg_<servername>` procedure, please see “Synchronizing query servers” in Chapter 5, “Working with Database Objects” in the *Sybase IQ System Administration Guide* and

- A remote server definition that points to the current server returns an expected server definition is circular error if attempting to use that remote server definition.

10.6.9 sp_iqcommandstats not supported [CR 416986]

Table 1-2, Stored Procedures for the IQ Store, incorrectly listed `sp_iqcommandstats` in the *Sybase IQ System Administration Guide*. This procedure is no longer supported.

10.6.10 Using multiplex with Interactive SQL [CR 405319]

A note in the section “Connection shortcuts in Sybase Central”, in Chapter 3, “Sybase IQ Connections” incorrectly states that you must use `dbisql`, not `isql` to access multiplex databases. Both utilities are supported.

10.6.11 Creating query servers [CR 405319]

The following sentence was omitted from the section “Creating query servers” in Chapter 5, “Working with Database Objects”: In order to have different server and database names in the interfaces file, you must give the `.DB` file a different name from the server name when you create the query server.

10.6.12 Setting multiplex permissions [CR 404004]

Sybase strongly recommends that you create users, domains and messages on the write server only; otherwise static collisions could lead to unintended results.

User names on the write server are global. For example, suppose that the DBA creates a user name `chris` for user Christopher Jones that exists on a query server only and owns objects on that query server.

Subsequently, the user `chris` is created for user Christine Smith on the write server. Now there is only one user `chris`, who is Christine Smith, and Christine Smith now owns objects on the query server formally owned by Christopher Jones.

A similar problem exists for domains and messages because both have a flat global name space. Even though they are owned by a user, there is no way to qualify the name with the owner name.

After a static collision, the original query server message or domain is renamed. However, any query server object that references the original query server message will now reference the new message. A query server object that references the original query server domain will either reference the original or the new domain, depending on whether it references it in a definition (e.g. a table) or dynamically (e.g. in a stored procedure).

Sometimes it may be best to set permissions differently on query servers. Consider these factors:

- Synchronizing copies user and permission definitions from the write server. The stored procedure `sp_mpxcfg_name` (where *name* is the query server name) may be used to reset the query server values to the desired state as part of the synchronize. If the query server has an IQ Local Store, the user privileges will be re-established as part of the synchronize. You can use the `-iqlocalreplay` server switch to override this.
- In the event of a disaster, the Catalog Store of a query server may be used to recreate the write server. In that case, differences on the query server leave the new write server in a different state from before the disaster.
- Setting permission on the query server is not permanent. The query server will continue to see GRANT and REVOKE commands from the write server. If the write server resets the permission, the change propagates to the query server.

10.6.13 LOAD TABLE statement requires column specifications [CR 403963] [453567]

The simple LOAD TABLE example in Chapter 7, “Moving Data In and Out of Databases” in the *Sybase IQ System Administration Guide* is incorrect, because it contains no column specifications. The following example is correct:

```
LOAD TABLE department
( dept_id, dept_name, dept_head_id )
FROM '/d1/MILL1/dept.txt'
```

10.6.14 SYSFILE file_name column after a RESTORE [CR 403817]

Add the following paragraph to the section “Restore accommodates dbspace changes” in the section “Restoring your databases” in Chapter 14, “Backup and Data Recovery”:

Note that the `file_name` column in the SYSFILE system table for the SYSTEM dbspace is not updated during a restore. For the SYSTEM dbspace, the `file_name` column always reflects the name when the database was created. The filename of the SYSTEM dbspace is the name of the database file.

10.6.15 Restoring multiplex databases [CR 398979, CR 375326]

The following procedure replaces “Restoring the multiplex,” in Chapter 14, *Sybase IQ System Administration Guide*.

❖ **Restoring the multiplex**

To restore the multiplex, you must first restore the Catalog Store and IQ Store, then restore any IQ Local Store(s) on query servers.

- 1 Confirm with Technical Support that a restore operation is needed.

If you have trouble opening your database on a query server, try doing a synchronize operation first. For details, see “Synchronizing Query Servers in Chapter 5, “Working with Database Objects,” *Sybase IQ System Administration Guide*.

- 2 Confirm that database home directories for each server still exist. If not, create them or restore them from file system backups:
- 3 Shut down every server in the multiplex (write server and query servers) using the Stop Multiplex command in Sybase Central.

Note If automatic start-up is enabled in your ODBC configuration, users on the same machine as the server may be able to start the server automatically and you will need to prevent this from happening while you are restoring the database.

- 4 After stopping all servers, verify that the database shut down successfully. If you see an active `asqsrsv12` process with the server name of a server in the multiplex (in a start-up parameter) you should stop it.

To verify on a UNIX system, use the `ps` command. For example:

```
% ps -ef | grep asqsrsv12
fiona 434      1  1   May 19 ??  0:05 asqsrsv12
-n myhost_myserver -c 32MB -x tcpip(port=1234)
mpxdb.db
fiona 4751    442  1 16:42:14 pts/5    0:00 grep asqsrsv12
```

To verify on a Windows system, use Task Manager. Look on the Processes tab for `asqsrsv12.exe` or find the IQ Server icon in the system tray and stop it using right-click and Shutdown.

- 5 Move files required for debugging and reconfiguring the multiplex.
 - Make a file system copy of the `.iqmsg` file.
 - On each server, preserve any `dbname.iqtmp` dbspace files, to reconfigure the multiplex. If the IQ Temporary store is damaged, use the start the server with the `-iqnotemp` switch to drop and recreate the temporary store dbspaces. For more information, see *Sybase IQ Release Bulletin*.

Delete the following files from the write server:

```
<database_home>/<dbname>.db
```

```
<database_home>/<dbname>.log
```

If a query server is damaged, however, drop it and re-add it after RESTORE but do not synchronize it.

- 6 Start the utility database from the write server directory using the write server's server name:

```
% start_asiq -n Alcott_Server1 -c 32MB
-x tcpip(port=1234)
```

To restore the database to a different home directory or a different machine from the one that created the backup, you must start the utility database using the multiplex override switch (-iqmpx_ov)

- 7 Connect to the utility database (*utility_db*).

```
% dbisql -c
"eng=Alcott_Server1;uid=DBA;pwd=SQL;dbn=utility_db"
```

- 8 Run the RESTORE command. Moving the database (restoring to a different machine) requires the RENAME clause. For details, see the *Sybase IQ Reference Manual*.
- 9 Shut down the write server.
- 10 Make sure that the temporary dbspaces exist as before, on raw devices or as files of the correct length. See "Backing up the IQ Store and Catalog Store," *Sybase IQ System Administration Guide*. For information on starting without the IQ Temporary Store, see *Sybase IQ Release Bulletin*.
- 11 Start the write server and, *if restoring to the same location*, synchronize the multiplex. For more information, see "Synchronizing query servers," in Chapter 5, "Working with Database Objects," *Sybase IQ System Administration Guide*. When restoring to a new location, you must start the write server using the multiplex override switch (-iqmpx_ov 1), and you must not synchronize.
- 12 If restoring to a new location, you must connect to the write server using Interactive SQL (not Sybase Central) and drop the query servers after starting the multiplex. This removes the absolute paths to query servers in their former locations from the system tables. (If you do not know the names of the query servers, open the Multiplex container in Sybase Central to list them. Ignore any warnings that occur in this situation.)

To remove each query server, run the `sp_iqmpxdropqueryserver` stored procedure in DBISQL or DBISQLC for each query server. For example, to drop a query server named `iqmpx_qs`, enter:

```
sp_iqmpxdropqueryserver ('iqmpx_qs')
```

If you wish to make the database multiplex, you may create new query server(s).

10.6.16 Specifying an IQ PATH [CR 385894]

On Windows, you must double the backslashes in SQL commands that specify a raw device. For example:

```
CREATE DBSPACE main2 AS '\\\\.\H:' IQ STORE;
```

This information was omitted from Chapter 5, section “Specifying an IQ PATH” in the *Sybase IQ System Administration Guide* and from CREATE DBSPACE and CREATE DATABASE syntax in Chapter 6, *Sybase IQ Reference Manual*.

10.6.17 ASE DATETIME accuracy [CR 385183]

Table 7-18 in the *Sybase IQ System Administration Guide* incorrectly documents ASE datetime accuracy. ASE datetime granularity is 1/300th of a second.

10.6.18 CREATE and DROP JOIN INDEX for multiplex [CR 382199]

CREATE and DROP JOIN INDEX operations return an error instead of propagating from write server to query server. You must perform them in single-node mode on the write server, then synchronize query servers.

10.6.19 IMAGE, TEXT, and LONG VARCHAR data types [CR 382185]

The TEXT, IMAGE, and LONG VARCHAR data types were erroneously listed as unsupported in the *Sybase IQ System Administration Guide*. The IMAGE and TEXT data types are supported via the IQ Binary Large Object (BLOB) and Character Large Object (CLOB) data types.

10.6.20 Using INSERT LOCATION with ODBC [CR 380217]

The IQ INSERT LOCATION syntax uses '{' and '}' characters, which represent a start and end escape sequence in the ODBC standard. Because the INSERT LOCATION select statement passed within the '{ }' characters is not a valid ODBC escape sequence, ODBC may return a syntax error like [ODBC Driver] Syntax error or access violation.

To avoid the ODBC syntax error, use one of the following methods:

- Turn off ODBC escape processing when sending an INSERT LOCATION statement by using the SQL_NOSCAN attribute setting in the ODBC program. For example:

```
SQLSetStmtAttr( stmt, SQL_NOSCAN, TRUE );
```

- Create an IQ stored procedure to accept string input that will be a valid INSERT LOCATION statement with different characters used in place of '{ }' that are not part of the INSERT LOCATION statement. The stored procedure would replace the alternate characters with '{ }' and execute the INSERT LOCATION statement.

For example, to create an IQ stored procedure that replaces '[' with '{':

```
create procedure sp_ins(string1 varchar(1024))
begin
    declare @cmdline varchar(1024);
    set @cmdline = replace(string1, '[', '{' );
    set @cmdline = replace(@cmdline, ']', '}' );
    execute immediate @cmdline;
end
```

Then the ODBC application would send the following to IQ:

```
call sp_ins
('INSERT into t1(c1) LOCATION 'host.db' [select c1
from t2]')
```

The stored procedure would convert the string to:

```
INSERT into t1(c1) LOCATION 'host.db' {select c1
from t1}
```

and execute that command.

10.6.21 Cannot update join index base table [CR 376130]

You cannot update a base table that is part of any join index. *Sybase IQ Reference Manual* and *Sybase IQ System Administration Guide* do not emphasize this restriction and error message “-1000102 Cannot update table %2 because it is defined in one or more join indexes” is undocumented.

10.6.22 ALTER INDEX command restriction [CR 375585]

The ALTER INDEX command returns an “index not found” error if you try to alter an index in a local temporary table. If you try to alter an index that is automatically created, such as a default index, IQ returns a “cannot alter index” error.

Only indexes in base tables or global temporary tables with owner type USER can be altered.

This information was omitted from Chapters 5 and 6 of the *Sybase IQ System Administration Guide*, from Chapter 6 or the *Sybase IQ Reference Manual*, and from Chapters 3 through 6 of the *Sybase IQ Troubleshooting and Error Messages Guide*.

10.6.23 Creating main dbspaces in a multiplex [CR 370316]

When creating a main dbspace on a write server, you must create aliases for the query servers before you synchronize them or they cannot open the new file.

This was omitted from Chapter 5, “Working with Database Objects,” in the *Sybase IQ System Administration Guide*.

10.6.24 Multiplex DDL restrictions [CR 367460]

Sybase IQ will not propagate CREATE, ALTER, and DROP statements for procedure and function definitions from the write server to query servers in a multiplex. You must use one of the following methods to keep procedure and function definitions consistent across multiplex servers.

- After executing the DDL statements to modify procedures and functions at the write server, execute the same statements on each query server. Create a SQL script for this approach to ensure consistent definitions.
- After executing the DDL statement(s) to modify procedures and functions at the write server, synchronize each query server.

ALTER and DROP statements for procedures and functions can fail if the existing definition is in use by another connection to the server. The user should schedule maintenance of procedures and functions for periods when the existing definitions will not be used.

This information was omitted from the *Sybase IQ System Administration Guide*.

Note If you have upgraded 12.5 databases to 12.6, see “DROP PROCEDURE on multiplex after upgrade [CR 398236]” on page 14.

10.6.25 Automatic login logout [CR 359391]

The following example shows how you can prevent a user from connecting after five failed login attempts.

This example uses a ConnectFailed event handler to total failed connect attempts. The totals are stored in table dba.event_table. The first time a user fails a connection attempt, a row will be inserted to dba.event_table for that user. On subsequent failed connections, the count in that user's row will be updated. A ConnectFailed event cannot prevent a user from continuing to try to connect. The login procedure must be used to deny access if the allowed number of failed login attempts is exceeded. This example tracks failed login attempts for all users including those who may be database administrators.

```
--First, create the table to hold the user
--information and insert a row for each use.

CREATE TABLE dba.event_table ( username
CHAR(128) NOT NULL, failed_login_attempts
INTEGER, PRIMARY KEY(username) );

--Create the event handler that will
--increment the number of failed login
--attempts.

CREATE EVENT ev_badlogin type ConnectFailed
handler
BEGIN
    DECLARE uid CHAR(128);
    DECLARE xx INTEGER;

    SET uid = event_parameter ('User')

    IF EXISTS(SELECT * FROM dba.event_table
WHERE ucase( uid) = ucase(username)) THEN
```

```
-- The user is already in the table.

UPDATE event_table
SET failed_login_attempts =
failed_login_attempts+1
    WHERE ucase(username) = ucase(uid);

ELSE

-- Insert the user for the first time.

INSERT dba.event_table
VALUES (ucase(uid), 1);

END IF;

SELECT failed_login_attempts INTO xx
FROM dba.event_table
WHERE ucase(username)=ucase(uid);

-- It is not possible to stop the user from
-- attempting to connect after 5 tries.
-- Instead, send a message to the server
-- console to notify the database
-- administrator that a user has exceeded
-- the allowable connect attempts.

IF xx > 5 THEN

RAISERROR 17001 uid + 'has had more
than 5 failed login attempts.'

END IF;

END

--The login procedure follows:

CREATE PROCEDURE dba.check_logins()

BEGIN

    DECLARE xx integer;
    DECLARE uid char(128);

-- See if the connected user is in the
-- event_table and proceed accordingly.

IF EXISTS(SELECT * from dba.event_table
WHERE ucase(username) =
ucase(current user)) THEN

SELECT failed_login_attempts INTO xx
FROM dba.event_table
    WHERE ucase(username) =
```

```
ucase(current user);

IF ( xx >= 5 ) THEN

RAISERROR 17010 current user + 'has been
locked out by the Database Administrator.'

ELSE

-- The user has connected.
-- Remove the user's row from
-- dba.event_table and call
-- the default login procedure
-- for the database.

SET uid=current user;

DELETE DBA.event_table
WHERE ucase(username) = ucase(uid);

CALL sp_login_environment();

    end if;

ELSE

-- The user is not in dba.event_table,
-- but has connected. Call the default
-- login procedure for the database.

CALL sp_login_environment();

END IF;

END;

--The following enables all users to run
--the login stored procedure and set the
--login procedure option.

GRANT EXECUTE ON dba.check_logins to PUBLIC;
SET OPTION PUBLIC.Login_Procedure =
'dba.check_logins';
```

To enable a user that has exceeded the number of allowed failed connection attempts to connect, a database administrator must delete the row for that user from dba.event_table.

This example was omitted from the *Sybase IQ System Administration Guide*, Chapter 12, “Managing User IDs and Permissions.”

10.6.26 Backing up query servers with IQ Local Stores

The following information was omitted from Chapter 14, “Backup and Data Recovery.” Use steps 2 through 5 when restoring either full or incremental backups. Do not use this procedure to restore to a different query server.

❖ Backing up and restoring query servers with IQ Local Stores

Enter commands in the following procedure without line breaks.

- 1 Connect to the query server using DBISQL or DBISQLC. Enter the following command:

```
BACKUP DATABASE FULL to 'full_path_to_backup_file'
```

- 2 Stop the query server.
- 3 In the write server’s directory, enter the following from a Command Prompt:

```
sync_qnode SQL query_server_name
```

- 4 In the query server’s directory, enter the following command:

```
dbbackup -y -d <full_path_of_the_query_server_dir>
-c "uid=DBA;pwd=SQL;eng=<write_server_name>;
links=tcip{host=<write_server_host>;
port=<write_server_port>}"
```

Omit spaces or line breaks from the -c parameter string. Line breaks in the preceding example are for readability only.

- 5 Delete the log file. For example:

```
rm <database>.log
```

If the log file or database name on the query server is different from the write server’s files, use the dblog utility to set the transaction log file name. You may use the relative path of the query server’s database file. For more about dblog, see Chapter 3 in the *Sybase IQ Utility Guide*. For example:

```
dblog -r -t all_types.log all_types.db
```

- 6 Start the query server without the database as:

```
start_asiq @params.cfg -n query_server_name -x
'tcpip{port=query_server_port}'
```

- 7 Connect to the query server using DBISQLC or DBISQL and specify the utility database in the connect string. Enter the following command:

```
RESTORE DATABASE '<database_name>' FROM
'<full_path_to_backup_file>'
```

- 8 Stop the query server that was started without the database.
- 9 Restart the query server via Sybase Central.

10.6.27 Correction to VerifyServerName examples

In Chapter 4, “Connection and Communication Parameters, the examples for the VerifyServerName [Verify] parameter are incorrect. They should read as follows, with a semicolon instead of a comma between “NONE” and “VERIFY”:

On the Network tab, check the TCP/IP check box and type in the text box:

```
host=ip_address:port#;DOBROADCAST=NONE;VERIFY=NO
```

For example:

```
host=123.456.77.888:2222;DOBROADCAST=NONE;VERIFY=NO
```

10.6.28 Extraction with named pipes

Any UNIX process issuing a read request on a named pipe will wait forever until the process writing data to the pipe either sends data or an EOF (end-of-file). Each time the reading process receives data, it issues another read. If the writing process stops sending data and fails to send an EOF, the read process will wait in the kernel and cannot be interrupted from IQ.

Make sure that any process writing data into a named pipe always finishes with an EOF. If an IQ connection becomes unresponsive while writing out to a named pipe, try dumping the data out of the pipe. For example, issue the following commands from another thread:

```
cat NamedPipeFile > /dev/null
```

10.6.29 Restoring to a raw device

In Chapter 14, “Backup and Data Recovery,” in “Moving database files,” the table name in the sample SELECT statement for restoring to a raw device should be SYS.SYSIQINFO, not SYS.SYSINFO.

10.6.30 Clarification for Chapter 16, “Accessing Remote Data”

In the Overview section, the list of external data sources should be divided into two groups:

- Sybase IQ, Adaptive Server Anywhere, and Adaptive Server Enterprise are accessible from both Windows and UNIX.
- Oracle, IBM DB2, Microsoft SQL Server, and Other ODBC data sources are available only from Windows.

In the Server classes section, the list of ODBC-based server classes should be divided into two groups:

- asaodbc works on all platforms of remote servers.
- aseodbc, db2odbc, msodbc, oraodbc, and odbc can only be used on Windows.

The sections “Loading remote data without ODBC access” and “Querying data without ODBC access” should read “Loading remote data without native classes” and “Querying data without native classes.”

Information was omitted from “Creating remote servers” under “ODBC-based server example.” After the line “ODBC-based connections may only be used for IQ on 32-bit systems,” add “On 64-bit systems, other methods of access are available, described in the sections ‘Loading remote data without native classes’ and ‘Querying remote data without native classes’.”

10.6.31 JDBC connection example code location

The directory *ASIQ-12_6/Samples/ASA/JavaSQL/manual-examples* is the location of the JDBC connection example source code referenced in “External connection example code” in the section “Establishing JDBC connections” in Appendix B, “Data Access Using JDBC.” The source code files are *not* located in the *java* directory.

10.6.32 JDBC and debugger documentation

Appendix B, “Data Access Using JDBC,” and Appendix C, “Debugging Logic in the Database,” have not been fully updated for version 12.6.

10.7 Reference Manual

This section contains updates to the *Sybase IQ Reference Manual*.

10.7.1 Index limit correction [CR 471124]

Table 8-1 in Chapter 8, “Physical Limitations,” shows a limit of 32,767 indexes per table. The correct limit is 2^{32} (approximately 4,000,000) indexes per table.

Corrections to the *Sybase IQ Reference Manual* will be available in the next GA release.

10.7.2 GARRAY_FILL_FACTOR_PERCENT clarifications [CR 469543]

The function and description of the GARRAY_FILL_FACTOR_PERCENT option have been clarified in Sybase IQ 12.6 ESD #10. Corrections to the *Sybase IQ Reference Manual* will be available in the next GA release.

Function	Specifies the percent of space on each HG garray page to reserve for incremental inserts.
Allowed values	0 – 1000
Default	25
Scope	DBA permissions are not required to set this option. Can be set temporary for an individual connection or for the PUBLIC group. Takes effect immediately.
Description	An HG index can reserve some storage on a per-group basis (where group is defined as a group of rows with identical values). Reserving space consumes additional disk space but can help the performance of incremental inserts into the HG index. If you plan to do future incremental inserts into an HG index, and those new rows have values that are already present in the index, a nonzero value for this option may improve incremental insert performance. If you do not plan to incrementally update the index, you can reduce this option to save disk space.

10.7.3 Files located relative to Catalog Store [CR 469000]

The information in the “Usage” section of the “CREATE DBSPACE statement” section in Chapter 6, “SQL Statements” contains incorrect information regarding the directory in which a dbspace is created, if an explicit directory is not specified.

The sentences “A filename without an explicit directory is created in the same directory as the initial dbspace of that store. Any relative directory is relative to that initial dbspace.” will be replaced with the following:

“A filename without an explicit directory is created in the same directory as the Catalog Store of the database. Any relative directory is relative to the Catalog Store.”

10.7.4 TOP_NSORT_CUTOFF_PAGES option [468580]

This option, added in 12.6 ESD #8, was not documented.

Function	Sets the result size threshold for TOP N algorithm selection.
Allowed values	1 – 1000
Default	1
Description	The TOP_NSORT_CUTOFF_PAGES option sets the threshold, measured in pages, where evaluation of a query that contains both a TOP clause and ORDER BY clause switches algorithms from ordered list-based processing to sort-based processing. Ordered list processing performs better in cases where the TOP N value is smaller than the number of result rows. Sort-based processing performs better for large TOP N values. In some cases, increasing TOP_NSORT_CUTOFF_PAGES can improve performance by avoiding sort-based processing.
See also	“SELECT Statement” in <i>Sybase IQ Reference Manual</i> .

10.7.5 Clarify TEMP_EXTRACT_DIRECTORY option [CR 468430]

The TEMP_EXTRACT_DIRECTORY option controls not only whether a user is allowed to use the data extraction facility but the directory into which temp extract files are placed. It can also override a directory path specified in the TEMP_EXTRACT_NAME options.

If the TEMP_EXTRACT_DIRECTORY option is set to the string FORBIDDEN (case insensitive) for a user, then that user is not allowed to perform data extracts. An attempt by this user to use the data extraction facility results in an error: “You do not have permission to perform Extracts”.

If TEMP_EXTRACT_DIRECTORY is set to FORBIDDEN for the PUBLIC group, then no one can run data extraction.

If TEMP_EXTRACT_DIRECTORY is set to a valid directory path, temp extract files are placed in that directory, overriding a path specified in the TEMP_EXTRACT_NAME options.

If TEMP_EXTRACT_DIRECTORY is set to an invalid directory path, an error occurs: “Files does not exist File: <invalid path>”

If TEMP_EXTRACT_DIRECTORY is blank, then temp extract files are placed in directories according to their specification in TEMP_EXTRACT_NAMEn. If no path is specified as part of TEMP_EXTRACT_NAMEn, the extract files are by default placed in the server startup directory.

The Sybase IQ manuals will be updated to reflect these changes in the next GA release.

10.7.6 Determining version using sp_iqstatus output [CR 463635]

To display space that can be reclaimed by dropping connections, use sp_iqstatus and add the results from the two returned rows:

```
(DBA)> select * from sp_iqstatus() where name like
'%Versions:%'
Execution time: 6.25 seconds
Name                               Value
-----
Other Versions: 2 = 1968Mb
Active Txn Versions: 1 = C:2175Mb/D:2850Mb

(First 2 rows)
```

The above example output shows that 1 active write transaction created 2175MB and destroyed 2850 MB of data. The total data consumed in transactions and not yet released is 4818MB, or 1968MB + 2850MB = 4818MB.

If sp_iqstatus shows a high percentage of main blocks in use on a multiplex server, run sp_iqversionuse to find out which versions are being used and the amount of space that can be recovered by releasing versions. See sp_iqversionuse in the *Sybase IQ Reference Manual*.

You can also display version information by using sp_iqtransaction to see total amount of MainTableKBCreated and MainTableKBDropped over transactions. See sp_iqtransaction in the *Sybase IQ Reference Manual* for details.

10.7.7 Correction to ASA NUMERIC data type default precision [CR 461147]

The default precision of Adaptive Server Anywhere NUMERIC and DECIMAL data types is incorrect in the “Compatibility” section of the “Numeric data types” section in Chapter 4, “SQL Data Types”. The default precision of the Adaptive Server Anywhere NUMERIC and DECIMAL data types is 30, not 3.

10.7.8 Using correct syntax when creating stored procedures [CR 456307]

There are two ways to create stored procedures: T-SQL and SQL/92. BEGIN TRANSACTION, for example, is T-SQL specific when using CREATE PROCEDURE syntax. Do not mix syntax when creating stored procedures.

10.7.9 Unique table names on a connection [CR 455684]

In Sybase IQ 12.6 ESD #10, an attempt to create a base table or a global temporary table will fail, if a local temporary table of the same name exists on that connection, as the new table cannot be uniquely identified by *owner.table*.

The information in the section “Unique table names on a connection [CR 455684]” on page 17 will be added to the following chapters and sections:

- In Chapter 6, “SQL Statements,” the “Usage” section in “DECLARE LOCAL TEMPORARY TABLE statement”
- In Chapter 6, “SQL Statements,” the “Usage” section in “CREATE TABLE statement”

10.7.10 VARCHAR and trailing blanks correction [CR 451561] [CR 450225]

Appendix A, “Compatibility with Other Sybase Databases” incorrectly states in the section “Character data types” that Adaptive Server Enterprise trims trailing blank spaces from VARCHAR values, but Sybase IQ does not. The correct statement is:

Adaptive Server Enterprise trims trailing blank spaces from VARCHAR values. Sybase IQ trims trailing blanks from VARCHAR values depending on the form of the data and the operation. For details, see “Character data types” in Chapter 4, “SQL Data Types.”

The following information has been added to the section “Character data types” in Chapter 4, “SQL Data types”:

VARCHAR data and trailing blanks

Data inserted via INSERT, UPDATE, or LOAD TABLE can be in one of the following forms:

- Enclosed in quotes
- Not enclosed in quotes
- Binary

For a column of data type VARCHAR, trailing blanks within the data being inserted are handled as follows:

- 1 For data enclosed in quotes, trailing blanks are never trimmed.
- 2 For data not enclosed in quotes:
 - Trailing blanks are always trimmed on insert and update.
 - For a LOAD statement, you can use the STRIP ON/OFF LOAD option to specify whether to have the trailing blanks trimmed. The STRIP ON/OFF option applies only to variable-length non-binary data. For example, assume the following schema:

```
CREATE TABLE t( c1 VARCHAR(3) );
LOAD TABLE t( c1 ',' ) ... STRIP ON           // trailing blanks trimmed

LOAD TABLE t( c1 ',' ) ... STRIP OFF          // trailing blanks not trimmed

LOAD TABLE t( c1 ASCII(3) ) ... STRIP ON      // trailing blanks not trimmed
LOAD TABLE t( c1 ASCII(3) ) ... STRIP OFF     // trailing blanks trimmed

LOAD TABLE t( c1 BINARY ) ... STRIP ON        // trailing blanks trimmed
LOAD TABLE t( c1 BINARY ) ... STRIP OFF       // trailing blanks trimmed
```

- 3 For binary data, trailing blanks are always trimmed.

You should not depend on the existence of trailing blanks in VARCHAR columns, when you write your applications. If an application relies on trailing blanks, a CHAR column should be used instead of a VARCHAR column.

10.7.11 Example revised for sa_checkpoint_execute procedure [CR 450192]

The example for the sa_checkpoint_execute system procedure is revised as follows:

Assuming you have created a subdirectory named *backup*, the following statement issues a checkpoint, copies all of the asiqdemo database files to the backup subdirectory, and completes the checkpoint:

```
sa_checkpoint_execute 'cp asiqdemo.* backup/'
```

10.7.12 LOAD TABLE FORMAT syntax correction [CR 445737]

In the section “LOAD TABLE statement” in Chapter 6, “SQL Statements,” the syntax for the FORMAT option should not contain single quotes. The correct syntax is:

```
LOAD [ INTO ] TABLE [ owner.]table-name
... ( load-specification [, ...] )
... FROM { 'filename-string' | filename-variable } [, ...]
```

```

... [ CHECK CONSTRAINTS { ON | OFF } ]
... [ DEFAULTS { ON | OFF } ]
... QUOTES OFF
... ESCAPES OFF
... [ FORMAT { ascii | binary } ]
...

```

10.7.13 OFF not an option for CLOSE_ON_ENDTRANS [CR 444249]

In Chapter 6, “SQL Statements,” “Set statement [T-SQL],” OFF is not an option for CLOSE_ON_ENDTRANS and has been removed from table 6-14.

In the subsequent bulleted list, CLOSE_ON_ENDTRANS should read as follows:

SET CLOSE_ON_ENDTRANS { ON } When CLOSE_ON_ENDTRANS is set to ON (the default and only allowable value), cursors are closed at the end of a transaction. With the option set ON, CLOSE_ON_ENDTRANS provides Transact-SQL compatible behavior.

10.7.14 ESCAPE_CHARACTER option for system use only [CR 440545]

The ESCAPE_CHARACTER option is reserved for system use only. Do not change the setting of this option.

Chapter 2, “Database Options” requires the following changes:

- Table 2-3: Remove the value and default of the ESCAPE_CHARACTER option.
- Add the following sentence to the note that follows Table 2-3:

The ESCAPE_CHARACTER option is reserved for system use only. Do not change the setting of this option.

10.7.15 More information requested on LOAD_MEMORY_MB [CR 426783]

The amount of virtual memory used by load command is a function of the total number of bytes (as defined by the table schema) for all columns being loaded. For example, consider the following schema:

```
CREATE table xx(c1 integer, c2 varchar(300), c2 double)
```

The total number of bytes is 312:

integer	4 bytes
varchar(300)	300 bytes
double	8 bytes

```
-----
total          312 bytes
```

Using this example, the load would require approximately 140MB:

```
312 * 45 * 10000 => 140MB
```

The amount of virtual memory used can become quite large if many columns (such as in a very wide table) are loaded at once. The wider the table, the more the load memory. The more users doing loads, the more heap/load memory is allocated outside IQ.

There are several courses of action you can take if you encounter the following error:

```
"All available virtual memory has been used ..."
```

You can set an upper limit on the amount of virtual memory a LOAD command can use by setting `LOAD_MEMORY_MB` to a non-zero value, with 2000MB the maximum allowed value.

You can also adjust `BLOCK FACTOR` or `BLOCK SIZE` LOAD command options. These command options default to 10000 and 500000, respectively, but you can set them to any number. Setting them lower forces the load to use less virtual memory.

You can also resort to loading a subset of the columns at a time, which is referred to as a partial-width load.

10.7.16 Subquery support in CREATE VIEW statement [CR 426715]

In Chapter 6, “SQL Statements,” the CREATE VIEW statement omits a subquery restriction. In the “Usage” section, the AS option should read as follows:

AS The SELECT statement on which the view is based must not contain an ORDER BY clause, a subquery in the SELECT list, or a TOP or FIRST qualification. It may have a GROUP BY clause and may be a UNION.

10.7.17 NUMBER(*) syntax errors [CR 426096]

Though the documentation for the NUMBER function stated “Use the NUMBER function only in a select list or a SET clause of an UPDATE statement. A syntax error is generated if you use NUMBER in any other type of statement,” it was requested that the statements be listed. This section now contains the following note:

Note A syntax error is generated if you use NUMBER in a DELETE statement, WHERE clause, HAVING clause, ORDER BY clause, subquery, query involving aggregation, any constraint, GROUP BY, DISTINCT, a query containing UNION ALL, or a derived table.

10.7.18 TEMP_EXTRACT_SIZE extract file size limits [CR 424817]

Extract file size limits listed for TEMP_EXTRACT_SIZE are default values, rather than upper limits.

When large file systems, such as JFS2, support file size larger than the default value, set TEMP_EXTRACT_SIZE to the value that the file system allows. For example, to support 1TB set option:

```
TEMP_EXTRACT_SIZE1 = 1000000000 KB
```

10.7.19 IDENTITY_INSERT option setting corrections [CR 418760, CR 439496, CR 456339]

To turn off the IDENTITY_INSERT database option, you must set it to an empty string, not a blank space. This was incorrectly documented in Chapter 2, “Database Options,” on page 26 (Table 2-1) and page 73, and in Chapter 6, “SQL Statements,” on page 492.

The following is incorrect:

```
SET TEMPORARY OPTION IDENTITY_INSERT = ' '
```

The correct syntax is:

```
SET TEMPORARY OPTION IDENTITY_INSERT = ''
```

The following changes apply to Chapter 2, “Database Options.”

The following note was omitted from the Scope section.

Note If you set a user level option for the current option, the corresponding temporary option is also set. For details, see “Scope and duration of database options” on page 17.

The following note was omitted from the Example section.

To illustrate the effect of user level options on temporary options (see note above), if you are connected to the database as DBA, and issue:

```
SET OPTION IDENTITY_INSERT = 'customer'
```

the value for the option is set to customer for the user DBA and temporary for the current connection. Other users who subsequently connect to the database as DBA find their option value for IDENTITY_INSERT is customer also.

10.7.20 ASIQTIMEOUT environment variable omitted [CR 418745]

The following new feature for Sybase IQ 12.6 was omitted from Chapter 1, section “Environment Variables,” in the *Sybase IQ Reference Manual*.

Setting

ASIQTIMEOUT=nnn

Operating system

Optional but recommended in multiplex environments.

Description

The IQ Agent waits indefinitely for a process to complete. Setting a wait time is recommended when creating or synchronizing query servers for a multiplex with a very large catalog store. Large catalog stores extend the time needed for the dbbackup part of synchronization, and increasing the wait time accommodates a larger synchronize.

This variable overrides the default wait time (zero, no timeout). The argument *nnn* is the number of minutes for the IQ Agent to wait. For example:

- To wait 45 minutes (Korn or Bourne shell):

```
ASIQTIMEOUT=45
export ASIQTIMEOUT
```

- To wait an hour (C shell):

```
setenv ASIQTIMEOUT 60
```

To overwrite the current setting, use the following server start-up option, where *nnn* is the number of minutes to wait:

```
-Dasigttimeout=nnn
```

10.7.21 ALTER TABLE and modifying columns [CR 408589]

In Chapter 3, “SQL Statements,” the ALTER TABLE statement description does not define “column-modification.”

ALTER column-name column-modification Change the definition of a column. The permitted modifications are as follows:

- **SET DEFAULT default-value** Change the default value of an existing column in a table. You can also use the MODIFY clause for this task, but ALTER is SQL/92 compliant, and MODIFY is not. Modifying a default value does not change any existing values in the table.
- **DROP DEFAULT** Remove the default value of an existing column in a table. You can also use the MODIFY clause for this task, but ALTER is SQL/92 compliant, and MODIFY is not. Dropping a default does not change any existing values in the table.
- **ADD** Add a named constraint or a CHECK condition to the column. The new constraint or condition applies only to operations on the table after its definition. The existing values in the table are not validated to confirm that they satisfy the new constraint or condition.
- **CONSTRAINT column-constraint-name** The optional column constraint name allows you to modify or drop individual constraints at a later time, rather than having to modify the entire column constraint.
- **SET COMPUTE (expression)** Change the expression associated with a computed column. The values in the column are recalculated when the statement is executed, and the statement fails if the new expression is invalid.
- **DROP COMPUTE** Change a column from being a computed column to being a non-computed column. This statement does not change any existing values in the table.

Another ALTER clause is incorrect and should read as follows:

| **ALTER CONSTRAINT** *constraint-name* **CHECK (new-condition)**

10.7.22 Valid password identifiers [CR 405855]

The following information should be added to GRANT statement section in Chapter 6, “SQL Statements.”

Passwords have a maximum length of 255 bytes.

The following are invalid for database user IDs and passwords:

- Names that begin with white space or single or double quotes

- Names that end with white space
- Names that contain semicolons

10.7.23 SYSDATABASE file_name column after a RESTORE [CR 403817]

Add the following item to the list of RESTORE issues in the Usage section of the “RESTORE statement” section in Chapter 6, “SQL Statements”:

- The file_name column in the SYSDATABASE system table for the SYSTEM dbspace is not updated during a restore. For the SYSTEM dbspace, the file_name column always reflects the name when the database was created. The filename of the SYSTEM dbspace is the name of the database file.

10.7.24 Database object size and number limitations [CR 403652]

The following corrections apply to Chapter 8, “Physical Limitations.”

- Database size: The maximum database size is approximately the number of files times the file size on a particular platform, depending on the maximum disk configuration. Refer to your operating system documentation for kernel parameters that affect the maximum number of files.
- Number of tables per database: 4,293,918,719
- Number of rows per table: 2^{48}
- Field size:
 - 255 bytes for BINARY, VARBINARY
 - 32,767 for CHAR, VARCHAR
 - Up to 512TB for 128KB pages or 1PB for 512KB pages for LONG BINARY, LONG VARCHAR

10.7.25 Selects from ASE database objects with fully qualified names [CR 402123]

A clarification is needed for the section “Data manipulation language” in Appendix A, “Compatibility with Other Sybase Databases.” The section is unclear about results of queries using fully qualified names in Adaptive Server Enterprise SELECT statements, such as a FROM clause with **<database name>.<owner>.<table name>**. Sybase IQ ignores the qualifiers in such statements. For example, Sybase IQ interprets the query `SELECT * FROM XXX..TEST` as `SELECT * FROM TEST`.

10.7.26 Global variable example correction [CR 397174]

In the Chapter 3 “SQL Language Elements” section “Global variables” under the section “Variables,” the data type in the CREATE PROCEDURE statement in the example should be VARCHAR (100):

```
CREATE PROCEDURE VersionProc ( OUT ver
                               VARCHAR ( 100) )
BEGIN
    SELECT @@version
    INTO ver;
END
```

10.7.27 MAX_IQ_THREADS_PER_CONNECTION minimum value changed [CR 395718]

As of 12.6 ESD #5, the minimum value of the MAX_IQ_THREADS_PER_CONNECTION database option is now 3, rather than 2. This value is documented in Chapter 2 “Database Options” in Table 2-1 “General database options” and in “MAX_IQ_THREADS_PER_CONNECTION option” in the section “Alphabetical list of options” in the *Sybase IQ Reference Manual*.

10.7.28 Reserving space for dbspace expansion [CR 394000]

In the Chapter 6 “SQL Statements” section “CREATE DATABASE statement,” the descriptions of the IQ RESERVE and TEMPORARY RESERVE clauses should say “Specifies the size in megabytes of space to reserve...” rather than “Specifies the size in megabytes of free list space to reserve...” The words “free list” should be removed from the sentence.

In the Chapter 6 “SQL Statements” section “CREATE DBSPACE statement,” the description of the RESERVE clause should say “Specifies the size in megabytes of space to reserve...” rather than “Specifies the size in megabytes of free list space to reserve...” The words “free list” should be removed from the sentence.

10.7.29 JDBC connections set QUOTED_IDENTIFIER option ON [CR 393934]

The following line should be added to the description of the QUOTED_IDENTIFIER option (TSQL) in Chapter 2, “Database Options.”

The JDBC driver also turns QUOTED_IDENTIFIER to ON.

10.7.30 FLOAT_AS_DOUBLE option and REAL data type join columns [CR 387270]

If a join column is a REAL data type, you must set the database option `FLOAT_AS_DOUBLE` to OFF when creating join indexes, or an error occurs. Issues may also result when using inexact numerics for join columns. This information was omitted from Chapter 2, “Database Options,” and Chapter 6, “SQL Statements,” *Sybase IQ Reference Manual*.

10.7.31 Multiplex system procedures [CR 385880]

Chapter 9, “System Procedures,” in the *Sybase IQ Reference Manual* contains errors in syntax for multiplex system procedures.

The syntax for `sp_iqmakempx` should read:

```
call sp_iqmakempx('host', 'servername',
'dbname.db', 'nnnn ')
```

where *nnnn* is the TCP port number where the write server should listen.

Syntax for the following stored procedures also lists internal names (such as *L_host_name*) instead of parameters (such as *host*):

The syntax for `sp_iqmpxaddremoteusers` should read:

```
call sp_iqmpxaddremoteusers('servername', 'qname')
```

Permissions should state that user must have DBA authority.

The syntax for `sp_iqmpxaliasdbspace` should read:

```
call sp_iqmpxaliasdbspace( dbspace, 'server1',
'path', 'server2', 'offset' )
```

Function should state that it adds an alias of an existing dbspace on a new server to SYSIQFILE.

Permissions should state that user must have DBA authority.

The syntax for `sp_iqmpxcreatepublication` should read:

```
sp_iqmpxcreatepublication('newserver')
```

Table 9: Parameters in `sp_iqmpxcreatepublication`

Name	Data type	Description
<i>new server</i>	varchar(30)	Name of server that requires a publication.

Permissions should state that user must have DBA authority.

The syntax for `sp_iqmpxcreatequeryserver` should read:

call sp_iqmpxcreatequeryserver('host', 'servername', 'path', 'nnnn')

The syntax for sp_iqmpxdropqueryserver should read:

call sp_iqmpxdropqueryserver('servername')

Permissions should state that user must have DBA authority.

The syntax for sp_iqmpxdropserverdbspaces should read:

call sp_iqmpxdropserverdbspaces ('servername')

Permissions should state that user must have DBA authority.

The syntax for sp_iqmpxdumptvlog should read:

sp_iqmpxdumptvlog('struc_type', 'row_order')

This procedure displays the following log information:

Table 10: Columns returned by sp_iqmpxdumptvlog

Parameter	Value	Description
<i>struc_type</i>	local or main	Type of IQ store
<i>row_order</i>	desc or asc	Descending or ascending row order

The syntax for sp_iqmpxexcludeserver should read:

call sp_iqmpxexcludeserver('servername', 'status')

Permissions should state that user must have DBA authority.

The syntax for sp_iqmpxmakeclean should read:

call sp_iqmpxmakeclean()

Permissions should state that user must have DBA authority.

The syntax for sp_iqmpxprotectexec should read:

call sp_iqmpxprotectexec('command')

The syntax for sp_iqmpxreplacewriteserver should read:

call sp_iqmpxreplacewriteserver('servername')

Usage should state that the DBA must run it on a query server.

The syntax for sp_iqmpxresetquerysubscription should read:

call sp_iqmpxresetquerysubscription('servername')

The syntax for `sp_iqmpxretryexec` should read:

```
call sp_iqmpxretryexec('cmd' , 'msg' )
```

The syntax for `sp_iqmpxsetpublisher` should read:

```
call sp_iqmpxsetpublisher('servername')
```

The syntax for `sp_iqmpxsubscribeuser` should read:

```
call sp_iqmpxsubscribeuser('user ' , 'path' , 'perm')
```

The syntax for `sp_iqmpxunsubscribeuser` should read:

```
call sp_iqmpxunsubscribeuser('user ')
```

The syntax for `sp_iqmpxvalidate` should read:

```
call dbo.sp_iqmpxvalidate('show_msgs')
```

The syntax for `sp_iqmpxversionfetch` should read:

```
call sp_iqmpxversionfetch( 'CatalogID ' , 'VersionID ' , 'OAVID ' ,  
'ServerType' , 'CatalogSync' , 'WCatalogID ' 'WVersionID ' )
```

10.7.32 Maximum number of rows [CR 380487]

The maximum number of rows in a table is incorrect in the *Sybase IQ Reference Manual* chapter, “Physical Limitations”. The former limit of 127 billion has changed to 256 trillion.

10.7.33 Cannot BCP into IQ table [CR 379792]

The second bullet in “BCP support in loading,” Appendix A, should read as follows:

IQ supports BCP indirectly. You can perform a BCP into an Anywhere table and then transfer the contents to IQ; however, the transfer of rows from Anywhere to IQ is executed one row at a time. IQ does not support BLKLIB, so BCP, which uses Open Client's Bulk-Library, doesn't work in load mode. Both Sybase IQ and Adaptive Server Enterprise BCP format support a blank when 1 digit is in the date.

10.7.34 LOCK TABLE statement [CR 379695]

The `LOCK TABLE` statement is allowed but undocumented in Sybase IQ 12.6. This statement prevents other concurrent transactions from accessing or modifying a table. Syntax is:

```
LOCK TABLE table-name [WITH HOLD] IN {SHARE | EXCLUSIVE}
```

MODE

Parameters are:

table-name - Must be a base table, not a view.

WITH HOLD — Holds the lock until the end of the connection. Otherwise, releases the lock when current transaction is rolled back or ends.

SHARE mode — Other transactions cannot modify the table but have read access. Requires SELECT privileges.

EXCLUSIVE mode — Other transactions cannot execute queries, updates, or any other transaction against the table. Requires DBA authority or ownership of the table.

This statement blocks or delays other transactions that may require access to the table. LOCK TABLE allows direct control over concurrency at a table level, independent of isolation level.

10.7.35 ROWID examples incorrect [CR 379498]

The ROWID examples in the *Sybase IQ Reference Manual* are incorrect. Output column headers should read “rowid(product)” instead of “rowid(product.id)”.

10.7.36 STRING_RTRUNCATION option setting correction [CR 371360]

In Chapter 2, “Database Options,” the description for the option STRING_RTRUNCATION is incorrect. It should read:

Determines whether an error is raised when an INSERT or UPDATE truncates a CHAR or VARCHAR string.

10.7.37 Compatibility of datetime and time values from ASE

A DATETIME or TIME value retrieved from an Adaptive Server Enterprise database using INSERT...LOCATION can have a different value due to the datetime precision of Open Client.

For example, the DATETIME value in the Adaptive Server Enterprise database is ‘2004-11-08 10:37:22.823’ as retrieved using INSERT...LOCATION is ‘2004-11-08 10:37:22.823333’.

10.7.38 DROP QUERY SERVER or DROP DATABASE on multiplex requires path

If you use Interactive SQL instead of Sybase Central to drop databases or query servers, always provide an explicit path. For example, if you drop the write server's database before dropping the query servers on the same machine, the following may return an error:

```
DROP DATABASE 'mydbname'
```

To avoid the error, specify the full database path, for example:

```
DROP DATABASE '/s1/mpx/wsrvr/mydbname.db'
```

10.7.39 SYSEVENT table source column

The source column, type LONG VARCHAR, was omitted from the SYSEVENT system table. This column contains the original source for the event handler if the PRESERVE_SOURCE_FORMAT option is ON. It is used to maintain the appearance of the original text.

10.7.40 SYSPROCEDURE columns omitted

Add the following columns to the section “SYSPROCEDURE system table” in Chapter 10, “System Tables.”

```
source LONG VARCHAR,  
avg_num_rows FLOAT,  
avg_costs FLOAT,  
stats LONG BINARY,
```

source This column contains the original source for the procedure if the PRESERVE_SOURCE_FORMAT option is ON. It is used to maintain the appearance of the original text. For more information, see “PRESERVE_SOURCE_FORMAT option [database]” in Chapter 2, “Database Options.”

avg_num_rows Information collected for use in query optimization when the procedure appears in the FROM clause.

avg_cost Information collected for use in query optimization when the procedure appears in the FROM clause.

stats Information collected for use in query optimization when the procedure appears in the FROM clause.

10.7.41 Updatable cursor row limitation with ODBC

There is a limitation regarding updatable cursors and ODBC. A maximum of 65535 rows or records can be updated, deleted, or inserted at a time using the following ODBC functions:

- `SQLSetPos` `SQL_UPDATE`, `SQL_DELETE`, and `SQL_ADD`
- `SQLBulkOperations` `SQL_ADD`, `SQL_UPDATE_BY_BOOKMARK`, and `SQL_DELETE_BY_BOOKMARK`

There is an implementation-specific limitation to the maximum value in the statement attribute that controls the number of effected rows to the largest value of an `UNSIGNED SMALL INT`, which is 65535.

```
SQLSetStmtAttr(HANDLE, SQL_ATTR_ROW_ARRAY_SIZE,  
VALUE, 0)
```

This information should be added to the section “Updatable cursor limitations” in the “Usage” section for the `DECLARE CURSOR` statement description in the “SQL Statements” chapter.

10.7.42 Double quotes within identifiers

Chapter 3, “SQL Language Elements,” incorrectly states under “Identifiers” that double quotes are not allowed within identifiers. Double quotes are allowed.

10.7.43 Additional notes on `MONITOR_OUTPUT_DIRECTORY` option

The `MONITOR_OUTPUT_DIRECTORY` option controls placement of output files for the IQ buffer cache monitor. All monitor output files are used for the duration of the monitor runs, which cannot exceed the lifetime of the connection. The output file still exists after the monitor run stops. A connection can run up to two performance monitors simultaneously, one for main cache and one for temp cache. A connection can run a monitor any number of times, successively.

The `MONITOR_OUTPUT_DIRECTORY` option controls the directory in which the monitor output files are created, regardless of what is being monitored or what monitor mode is used.

The chapter “Database Options” in the *Sybase IQ Reference Manual* contains an example of setting the MONITOR_OUTPUT_DIRECTORY option. In this example, the output directory string is set to both “/tmp” and “tmp/”. The trailing slash (“/”) is correct and is supported by the interface. The example illustrates that the buffer cache monitor does not require a permanent table; a temporary table can be used.

10.7.44 sp_mpxcfg_<servername> procedure

The following information was omitted from Chapter 9, “System Procedures.”

Function	Sets up query server named <i>servername</i> for SQL Remote replication.
Syntax	call “DBA”.sp_mpxcfg_<servername>(' ')
Description	Sybase IQ calls this procedure when synchronizing query servers. This procedure in turn runs specified procedure or procedures on the named query server. When finished, this procedure returns the following message in the server log: Query server auto-configuration complete. If the query server is already configured or if you run sp_mpxcfg_<servername> on a write server, the procedure does nothing.
Permissions	Must have DBA authority.

10.8 New Features in Sybase IQ 12.6

This section contains updates to the *New Features in Sybase IQ 12.6*.

10.8.1 Perl interface (DBD::ASAny) not supported by Sybase IQ [CR 443439]

Under “Miscellaneous enhancements,” the paragraph stating “The new DBD::ASAny driver for the Perl DBI module allows you to access and modify Adaptive Server Anywhere databases from Perl scripts,” is incorrect.

Perl interface (DBD:ASAny) not supported by Sybase IQ.

10.8.2 Enhanced checkpoint processing [CR 394506]

IQ 12.6 ESD #4.6 includes an enhancement to allow checkpoint processing to occur even when the database is experiencing a high rate of commits. The change allows fair access for checkpoints to occur in the presence of operations that suspend checkpoints such as commits. While it is still true that checkpoints, commits, and rollbacks are mutually exclusive, the change allows checkpoints to occur between commits and rollbacks, even when the database is experiencing a high commit rate.

Since checkpoints have opportunity to run more frequently in a high commit rate scenario, this reduces the number of transactions that checkpoint must apply at one time, thereby shortening the checkpoint time. This enhancement applies to both explicit checkpoints started by issuing a CHECKPOINT command and implicit checkpoints based on the -gc server start-up parameter or CHECKPOINT_TIME option setting. The new checkpoint behavior is automatically enabled when running 12.6 ESD #4.6 and is required.

10.8.3 Dropping join indexes on a write server [CR 385317]

You can no longer drop a join index on a write server if the FORCE_DROP option is set ON. To force drop such a join index, you must first start the server in single-node mode. If you drop an object when FORCE_DROP = ON, the FORCE_DROP setting on query servers is unaffected, and you need only restart the write server after the drop completes.

10.8.4 Managing thread allocation [CR 384081]

Sybase IQ 12.6 featured a new algorithm to avoid running out of threads. When the algorithm is used (the default), the IQ optimizer assigns a thread quota to each table that contains invariant predicates, based on the row counts after high selectivity filters.

In ESD #3, the ENABLE_THREAD_ALLOWANCE option was added to let users control when the optimizer uses the algorithm. The option may be set temporary, for an individual connection, or for the PUBLIC group. It may be set ON or OFF, and takes effect immediately. The default is now OFF.

10.8.5 LOAD_ZEROLENGTH_ASNULL option added [CR 383943]

The LOAD_ZEROLENGTH_ASNULL database option added to 12.6 ESD #3 lets you specify LOAD statement behavior under the following conditions:

- Inserting a zero-length data value into a column of data type CHAR, VARCHAR, LONG VARCHAR, BINARY, VARBINARY, or LONG BINARY, and
- A NULL column-spec (e.g. NULL(ZEROS) or NULL(BLANKS)) is also given for that same column

Setting option LOAD_ZEROLENGTH_ASNULL 'ON' loads a zero-length value as NULL when the above conditions are met.

Setting option LOAD_ZEROLENGTH_ASNULL 'OFF' (the default) causes LOAD to load a zero-length value as zero-length subject to the setting of option NON_ANSI_NULL_VARCHAR.

10.8.6 HG_DELETE_METHOD option added [CR 381564]

Sybase IQ 12.6 included performance enhancements for delete batch processing. The HG_DELETE_METHOD option lets you specify which of three algorithms to apply to the delete. If no algorithm is specified, the 12.6 cost model considers many factors, including I/O costs, CPU costs, available resources, index metadata, parallelism, and predicates available from the query. See the *Sybase IQ Reference Manual* for details about HG_DELETE_METHOD.

Note Costing currently does not consider the effects of range predicates on the large delete. This can cause mid delete to be chosen in cases where large delete would be faster. You can use HG_DELETE_METHOD to force the use of the large delete algorithm.

10.8.7 CONVERSION_ERROR option behavior change [CR 380625]

When the CONVERSION_ERROR option is set OFF, each thread doing data conversion for a LOAD statement now writes at most one warning message to the *.iqmsg* file.

In order to write all data conversion warning messages to the *.iqmsg* file, you must set the CONVERSION_ERROR OFF, then set the DDL_OPTIONS2 option to 8. For example:

```
SET TEMPORARY OPTION CONVERSION_ERROR='OFF'  
SET TEMPORARY OPTION DDL_OPTIONS2=8
```

10.8.8 Temporary tables in stored procedures [CR 373124]

If a procedure dynamically creates and then selects the same temporary table within a stored procedure, you must use the EXECUTE IMMEDIATE WITH RESULT SET ON syntax to avoid “Column not found” errors. For example:

```
CREATE PROCEDURE p1 (IN @t varchar(30))
BEGIN
EXECUTE IMMEDIATE
'SELECT * INTO #resultSet FROM ' || @t;
EXECUTE IMMEDIATE WITH RESULT SET ON
'SELECT* FROM #resultSet';
END
```

This was omitted from Chapter 8, “Using Procedures and Batches,” *Sybase IQ System Administration Guide*.

10.8.9 Location of Java sample programs

The Java sample programs are now installed in the directory *ASIQ-12_6/Samples/ASA/JavaSQL*, not the directory *samples/asa/java* as stated in the “XML and Java support” section.

10.9 Utility Guide

This section contains updates to the *Sybase IQ Utility Guide*.

10.9.1 Effect of start_asiq -ti on connection to local server using shared memory [CR 422369]

The start_asiq -ti server switch, which specifies the client idle time before shutdown, has no effect on connections to a local server using shared memory.

This information has been added to the -ti switch description in the Chapter 1, “Running the Database Server” section “Starting the database server” under “Usage.”

10.9.2 Truncating the transaction log on a running server [CR 404902]

The following information was omitted from the *Sybase IQ Utility Guide*.

The Backup utility (dbbackup) makes a copy of the transaction log of a running IQ database. This utility lets you truncate the transaction log, freeing disk space and improving recovery speed, without having to stop and restart your server.

Note To back up an entire Sybase IQ database, always use the BACKUP statement, not dbbackup. BACKUP backs up all database files, and is the only way to back up the Catalog Store. For details, see “BACKUP statement,” in the *Sybase IQ Reference Manual*.

Syntax

dbbackup [options] target-directory

Parameters

The following table lists the available options for the dbbackup utility.

Table 11: Options for dbbackup

Option	Description
@data	Read options from the specified environment variable or configuration file
-c "keyword=value; ..."	Supply database connection parameters
-l (lowercase L) file	Live backup of the transaction log to a file
-q	Quiet mode — do not print messages
-r	Copy the old transaction log to a new name and start a new empty log
-xo filename	Truncate (delete and restart) the transaction log

Usage

The Backup utility allows you to back up the transaction log while other applications or users are using the database. Backup filenames are the same as the database filenames.

If you have adequate disk space, use -r to preserve the existing log file under a new name and start a new empty log. If disk space is limited, use -xo instead to truncate the existing log.

Exit codes are 0 (success) or non-zero (failure).

Options

@data Reads in options from the specified environment variable or configuration file. If both exist with the same name, Sybase IQ uses the environment variable.

For more information about configuration files, see your *Sybase IQ Installation and Configuration Guide*.

To protect passwords or other information in the configuration file, you can use the File Hiding utility (dbfhide) to obfuscate configuration file contents.

Connection parameters (-c) If the connection parameters are not specified, connection parameters from the SQL CONNECT environment variable are used, if set. The user ID must have DBA authority or REMOTE DBA authority.

For a description of the connection parameters, see Chapter 4, “Connection and Communication Parameters” in *Sybase IQ System Administration Guide*.

Live backup (-l (lowercase L)) Enables a secondary system to be brought up rapidly in the event of server failure. A live backup does not terminate, but continues while the server runs. It runs until the primary server becomes unavailable. At that point, it shuts down, but the backed up log file is intact and can be used to bring a secondary system up quickly.

The live backup of the transaction log is always the same length or shorter than the active transaction log. When a live backup is running and another backup restarts the transaction log (dbbackup -x), the live backup automatically truncates the live backup log and restarts the live backup at the beginning of the new transaction log.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from a command prompt.

Rename and start new transaction log (-r) Forces a checkpoint and the following three steps to occur:

- Copies and saves the current working transaction log to the directory specified in the command.
- Keeps the current transaction log in its current directory, but renames it using the format *yymmddxx.log*, where *xx* are sequential characters starting at *AA* and running to *ZZ*, and *yymmdd* represents the current year, month, and day. This file is then no longer the current transaction log.
- Generates a new transaction log file that contains no transactions. The new file has the name of the former current transaction log and becomes the current transaction log.

Back up the transaction log file only (-t) This can be used as an incremental backup since the transaction log can be applied to the most recently backed up copy of the database file(s).

Delete and restart the transaction log without a backup (-xo) Delete the current transaction log and start a new one. This operation does not carry out a backup; its purpose is to free up disk space.

target-directory The directory to which the backup files are copied. If the directory does not exist, Sybase IQ creates it. The parent directory must exist.

Examples

The following Windows command backs up the transaction log from the *asiqdemo* database running on the *sample_server* server into the directory *asiqbakup*, connecting as user ID DBA with password SQL:

```
dbbackup -c
"eng=sample_server;dbn=asiqdemo;uid=DBA;pwd=SQL"
c:\sample\asiqbakup
```

The following command uses the *-xo* option of *dbbackup* to delete the transaction log and start a new one with the same name, and the *-l* option to perform these actions while connected to the live database:

```
dbbackup -l -xo -c
"eng=sample_server;dbn=asiqdemo;uid=DBA;pwd=SQL"
```

10.9.3 Default thread stack size and number of threads [CR 381625]

The *Sybase IQ Utility Guide* Chapter 1 specified incorrect values for default thread stack size and number of threads. The thread stack size default (*-iqtss*) for 64-bit systems is 350KB. The total number of threads (*-iqmt* plus *-gn*) must not exceed 4096 on 64-bit platforms or 2048 on 32-bit platforms.

The calculation for the number of threads IQ creates is 60 per CPU for the first four CPUs and 50 per CPU for the rest, plus connection threads.

For example, on a system with 12 CPUs and 10 connections:

$$60 * 4 + 50 * (numCPUs - 4) + 2 * (numConnections + 2) + 1 = 655$$
10.9.4 Requirement for Interactive SQL [CR 375382]

In order for DBISQL to function correctly, \$HOME must exist and must be writable by the user.

10.9.5 Switches for the start_asiq utility

The following utility switches were omitted from Table 1-1 in the *Sybase IQ Utility Guide*:

Table 12: Server start-up options for start_asiq

Switch	Description
-cc {+ -}	Enable/disable page collection for cache warming
-cr {+ -}	Enable/disable cache warming
-cs	Display cache sizing statistics
-cv {+ -}	Enable/disable cache warming status messages
-f	Force database to start without transaction log
-gf	Disable firing of triggers
-iqnolocalreplay 1	Do not replay the TLV log after synchronizing a multiplex.
-iqnomain 1	Start a query server with a local store without opening shared IQ Main Store.
-iqnotemp size	Create a temporary file in place of the defined temporary dbspace. The argument to the switch specifies the file size in MB an argument to the switch.
-oe filename	Set filename to log start-up errors, fatal errors, and assertions
-startdir dirname	Start the server in the specified directory. If you use this parameter on Windows, IQ assumes that each server is being started by the Sybase IQ Agent, and starts the server in the background

If there is a problem starting the server, start_asiq returns a non-zero value. If you did not specify a log file after the -o switch on start-up, the error is written to the first one of the following that is defined:

- \$ASLOGDIR/<servername>.xxxx.srvlog
- \$ASDIR/logfiles/<servername>.xxxx.srvlog
- \$ASLOGDIR/start_asiq.log
- \$ASDIR/logfiles/start_asiq.log
- the Systems applications log file

10.9.6 Switches for the stop_asiq utility

Syntax

Switches were omitted from Chapter 1, section “The stop_asiq utility (UNIX only).” The title should read “(UNIX and Linux only).” Updated information is as follows:

To stop the server, run the stop_asiq utility, using the following command format:

```
stop_asiq [ -agent | -cleanup ] [ -stop [one | all ] ] [ -help ]
```


Switches

Table 13: Switches for stop_asiq utility

Parameter	Purpose
-agent	Stops the IQ Agent on UNIX or Linux systems
-cleanup	Removes the orphan Sybase IQ process on Linux
-help	Displays stop_asiq syntax and switches
-stop [one all]	Removes user interaction with stop_asiq. Assumes “Y” response to all questions.

Stopping servers in cron or at jobs

To use stop_asiq in a cron or at job, specify the utility with the appropriate -stop option:

```
stop_asiq -stop one
```

Setting -stop one shuts down a single server, when exactly one running server was started by the user ID that starts the cron or at job. This prevents accidentally shutting down the wrong server if several are running.

```
stop_asiq -stop all
```

Setting -stop all shuts down all servers that were started by the user ID that starts the server. This command is compatible with Sybase IQ 12.5 if you have installed ESD #8 or higher.

You can specify both options on the same command, for example:

```
stop_asiq -agent -stop all
```

Note You must specify the full pathname to the stop_asiq executable in the cron statement.

10.9.7 dbremote -k parameter

Table 1-4 “Recommended dbremote options for multiplex servers” in Chapter 1, “Running the Database Server” incorrectly lists the -k parameter. The dbremote -k parameter is no longer supported by Sybase IQ.

10.10 Performance and Tuning Guide

This section contains updates to the *Sybase IQ Performance and Tuning Guide*.

10.10.1 Restricting concurrent queries with -iqgovern [CR 432651]

In Chapter 5, “Managing System Resources,” the description for -iqgovern under “Restricting concurrent queries,” is incorrect. It has been corrected to read as follows:

The -iqgovern switch lets you specify the number of concurrent queries on a particular server. This is not the same as the number of connections, which is controlled by your license. By specifying the -iqgovern switch, you can help IQ optimize paging of buffer data out to disk, and avoid overcommitting memory. The default value of -iqgovern is $(2 \times \text{the number of CPUs}) + 10$. You may need to experiment to find an ideal value. For sites with large numbers of active connections, try setting -iqgovern slightly lower.

10.10.2 Estimating temporary cache size [CR 425940]

The following was omitted from “Managing buffer caches” in Chapter 5, “Managing System Resources” in the *Sybase IQ Performance and Tuning Guide*.

For guidelines about estimating temporary cache size, see the Hardware Sizing Guide for Sybase IQ 12.6 and 12.7 at <http://www.sybase.com/detail?id=1050862>. This technical white paper describes the relationship between hardware configuration and Sybase IQ memory management.

10.10.3 Memory buffer cache [CR 391499]

Information on the memory split in section “Sybase IQ main and temp buffer caches” in Chapter 4, “Managing System Resources,” does not agree with Table 4.1.

The sentence, “The general rule of thumb for IQ, unlike most other databases, is a split of about 40% for the main buffer cache and 60% for temp buffer cache,” is correct, but the table shows percentages the other way around.

Table 4-1 has been edited to agree with the paragraph.

10.10.4 Cross reference errors [CR 388182]

The *Sybase IQ Performance and Tuning Guide* contains the following errors:

- Page 33 reference to page 515 should read 45
- Page 43 reference to page 526 should read 57

- Page 45 reference to Table 12-1 on page 520 should be Table 4-1 on page 50
- Page 48 reference to page 562 should be 87
- Page 50 reference to Tables 12-2 and 12-3 should be 4-2 and 4-3
- Page 51 reference to 526 should be 57 and 521 should be 51

10.10.5 Using dbremote -o and -k parameters when truncating transaction log

The start dbremote example in the section “Truncating the transaction log for a multiplex database” in Chapter 4, “Managing System Resources” contains incorrect parameters. The dbremote -k parameter is no longer supported by Sybase IQ and the -o parameter should be just before the output file path. The correct syntax of the start dbremote command is:

```
start dbremote -q -v -x -o
"d:\Server01\mpxdb\dbremote.log" -c
"uid=DBA;pwd=SQL;eng=Server01;dbf=
d:\Server01\mpxdb\mpxdb;
links=tcipip{port=1704;host=FIONA-PC}"
```

10.11 Large Objects Management in Sybase IQ

This section contains information on items documented in *Large Objects Management in Sybase IQ*.

10.11.1 LOAD TABLE FORMAT syntax correction [CR 445737]

In the section “Loading large object data” in Chapter 6, “Moving Large Object Data,” the syntax for the FORMAT option of the LOAD TABLE statement should not contain single quotes. The correct syntax is:

```
LOAD [ INTO ] TABLE [ owner.]table-name
... ( column-name load-column-specification [, ...] )
... FROM 'filename-string' [, ...]
... [ QUOTES { ON | OFF } ]
... ESCAPES OFF
... [ FORMAT { ascii | binary } ]
... [ DELIMITED BY 'string' ]
...
```

11. Technical Support

Before you contact 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.

Technical Support needs information about your environment in order to resolve your problem. Before contacting Technical Support, run the `getiqinfo` script to collect as much information as possible automatically. You may also need to collection some information manually. Providing this information helps expedite the resolution of your problem.

In the following list, * indicates items collected by `getiqinfo`.

- Type of hardware, amount of memory, number of CPUs*
- Operating system and version (for example, Sun Solaris 2.9)*
- Operating system patch level*
- Front end tool used (for example, Brio Query)
- Connectivity protocol used (for example, ODBC, JDBC, TDS)
- Open Client version
- Configuration type (single user or multiuser)
- Message log file (*very important*)*

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

- Stack trace file for the date and time this problem occurred

Named *stktc-YYYYMMDD-HHMMSS_#.iq*, located in the directory where you started the database server*

- Command or query that produced the error
- Query plan * (recorded in *.iqmsg* file)

Note The query plan is collected automatically by `getiqinfo`. If you collect information manually you must 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.

If you have performance problems, set the following option:

```
SET TEMPORARY OPTION Query_Plan_After_Run = 'ON'
```

This enables Technical Support to see which steps query processing steps used the time.

- Server logs
 - For UNIX and Linux, *ASIQ-12_6/logfiles/<servername>.00n.stderr* and *ASIQ-12_6/logfiles/<servername>.00n.srvlog**
 - On Windows platforms, *ASIQ-12_6\logfiles\<servername>.00n.srvlog**
- start-up and connection option settings, from the configuration file (by default, *dbname.cfg*)*
- Database option settings and output from *sa_conn_properties** (if the server is still running)
- Schema and indexes for the database
- Output from *sp_iqstatus* and *sp_iqcheckdb*
- On multiplex databases, you must execute *getiqinfo* on the write server and/or query servers, and the following information is also collected:
 - *servername.out* on query servers
 - *write_server_name/repDirs/logfiles/servername.dbrlog*

A checklist for recording this information for Technical Support is provided at the end of this release bulletin.

For more information on *getiqinfo*, see “Collecting diagnostic information using *getiqinfo*” in *Sybase IQ Troubleshooting and Error Messages Guide*.

Online support

You may find additional help from the Sybase online support Web site MySybase at <http://www.sybase.com/support/>. MySybase is a free service that allows you to search through closed support cases, latest software bulletins, and resolved and known problems, using a view customized for your needs. You can even open a Technical Support case online.

MySybase can be used from most Internet browsers. Go to the Sybase support home page at <http://www.sybase.com/support/> and click the MySybase tab for information on how to sign up for and use this free service.

12. Other sources of information

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

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

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

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

Note The SyBooks browser software runs on Windows and Linux platforms. Users with non-Linux UNIX platforms must use Acrobat Reader to open PDF files on the SyBooks CD.

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

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

- Infocenter is an online version of SyBooks that you can view using a standard Web browser. To access the Infocenter Web site, go to Sybooks Online Help at <http://infocenter.sybase.com/help/index.jsp>.

12.1 Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ Finding the latest information on product certifications

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

❖ Finding the latest information on component certifications

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

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

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

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

12.2 Sybase EBFs and software maintenance

❖ Finding the latest information on EBFs and software maintenance

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

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

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

12.3 IQ Newsgroup

Sybase newsgroups provide a means for users to exchange information over the Internet. The newsgroup for Sybase IQ is sybase.public.iq.

For information on subscribing to Sybase newsgroups, configuring your newsreader or Web browser, and guidelines for posting, go to Newsgroups at <http://www.sybase.com/support/newsgroups>.

Note Sybase newsgroups are moving to a new, more fault-tolerant environment. If you already subscribe to Sybase newsgroups, you need to replace your existing newsgroup headers with the new ones. Follow these steps.

❖ Replacing existing newsgroup information

- 1 Remove the forums.sybase.com account from your newsreader.
- 2 Add the forums.sybase.com account again.
All headers will be marked as unread.
- 3 Resubscribe to your newsgroups, following the instructions at Newsgroups at <http://www.sybase.com/support/newsgroups>.
- 4 Download headers into those newsgroups.

12.4 IQ User's Group

The IQ User's Group provides a forum for any IQ users to exchange information on IQ. To have your name added to the IQ User's Group list, send e-mail to iqug@odshp.com.

13. Accessibility features

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

Sybase IQ 12.6 and the HTML documentation have been tested for compliance with U.S. government Section 508 Accessibility requirements. Documents that comply with Section 508 generally also meet non-U.S. accessibility guidelines, such as the World Wide Web Consortium (W3C) guidelines for Web sites.

For Section 508 compliant core documentation, see your Sybooks CD or the Infocenter Web site at <http://infocenter.sybase.com/help/index.jsp>.

For Section 508 compliant installation guides and release bulletins, see the HTML files on your Getting Started CD.

The online help for this product is also provided as HTML-based JavaHelp, which you can navigate using a screen reader.

For information about accessibility support in the Sybase IQ plug-in for Sybase Central, see “Using accessibility features” in *Introduction to Sybase IQ*. The online help for this product, which you can navigate using a screen reader, also describes accessibility features, including Sybase Central keyboard shortcuts and using this product without a mouse.

For information about how Sybase supports accessibility, see Sybase Accessibility at <http://www.sybase.com/accessibility>. The Sybase Accessibility site includes links to information on Section 508 and W3C standards, and Section 508 compliance statements for Sybase IQ and other Sybase products.

Note You might need to configure your accessibility tool for optimal use. Some screen readers pronounce text based on its case; for example, they pronounce ALL UPPERCASE TEXT as initials, and MixedCase Text as words. You might find it helpful to configure your tool to announce syntax conventions, or to pronounce certain characters or punctuation such as underscore. Consult the documentation for your tool. See the following table for pronunciation guidelines for important terms.

Table 14: ABBR and ACRONYM values

Term	Element value
\$ASDIR/bin	dollar A S D I R slash bin
asiqdemo	A S I Q demo
ASIQPlugin.jar	A S I Q plug in dot jar

Term	Element value
DBISQL	D B I sequel
dbname.db	D B name dot D B
dbname.log	D B name dot log
dbname.iqtmp	D B name dot I Q T M P
Note You can apply this technique to other filenames with extensions.	
dbspace	D B space
DSEEDIT	D S edit
IDs	I deez
I/O	I O
iq_dummy	I Q underscore dummy
IQ_SYSTEM_MAIN	I Q underscore system underscore main
libtcl.cfg	L I B T C L dot C F G
start_asiq	start underscore A S I Q
start_asiq.exe	start underscore A S I Q dot e x e
SQL	sequel (not S Q L, except when explaining how to pronounce SQL, or when it's what you type, like a password. JAWS reads SQL as S Q L by default.)
®	registered trademark
™	trademark
→	right arrow

14. Checklist: information for Technical Support

You can run the `getiqinfo` script to collect much of this information.

Information requested	Value
type of hardware	
amount of memory	
number of CPUs	
operating system name and version (e.g., Sun Solaris 2.9)	
operating system patch level	
front end tool used (e.g., Brio Query)	
connectivity protocol used (e.g., ODBC, JDBC, TDS)	
Open Client version	
configuration type (single or multiuser)	
message log file (dbname.iqmsg)	
stack trace file stktrc-YYYYMMDD-HHNNSS_#.iq	
command or query that produced the error	
start-up option settings	
connect option settings	
database option settings	
schema and indexes for the database	
sp_iqstatus output	
query plan: set options (Query_Plan, Query_Detail, Query_Plan_After_Run), rerun command or query	

