

Release Bulletin Sybase® IQ 15.1 for Sun Solaris

Document ID: DC74373-01-1510-05

Last revised: May 2010

Topic	Page
1. Accessing current release bulletin information	2
2. Product summary	2
3. Special installation and migration instructions	2
4. New features and behavior changes	11
5. Product compatibilities	20
6. Restrictions	21
7. Known problems	24
8. Documentation for this version	28
9. Documentation updates and clarifications	31
10. Technical Support	58
11. Other sources of information	61
12. Accessibility features	64
13. Checklist: information for Technical Support	65

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://sybooks.sybase.com>.
- 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 15.1 for 64-bit Sun Solaris.

Refer to the *Installation and Configuration Guide for Sun Solaris* for compatible platform and operating system configurations and minimum patch levels.

Refer to the *New Features Summary Sybase IQ 15.1* document for descriptions of new features and behavior changes in Sybase IQ 15.1.

3. Special installation and migration instructions

For complete installation and migration instructions, see the *Installation and Configuration Guide for Sun Solaris*. Be sure to read the chapters “Preparing for Installation” and “Migrating Data” before installing Sybase IQ 15.1.

Before you run a new version of Sybase IQ, see “Restrictions” on page 21 for the most current requirements. This section includes installation and migration information that is new for this release or needs emphasis.

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 62.

3.1 New installer for Sybase IQ 15.x

Sybase IQ 15.x has a new procedure for installation. The Sybase IQ 15.x interactive installer is a standard installation wizard that lets you set up your software in a graphically-based environment. The interactive installer supports all product editions and installation types. A new silent install option is also available. For more information, see “Installing your software” in Chapter 1, “Preparing for Installation” in the *Installation and Configuration Guide for Sun Solaris*.

3.2 Installing Sybase IQ 15.1 ESD #3

Installation instructions for the Sybase IQ 15.1 ESD #3 GUI differ from those in the Installation and Configuration Guide: the installer includes an additional license validation step. Validating your license before you install this update will ensure that your database functions correctly after you install the software.

License validation is supported in GUI installations and Silent installations. If you install this update on a UNIX-like OS, make sure you use the same shell to start the installer (either Silent or GUI) that you use to start your current IQ server. This can be verified by ensuring that `$IQDIR15` environment variable is set.

Note Stop any servers running in the install directory before you install this update.

3.2.1 GUI installations

- 1 Start the installer. See the Installation and Configuration Guide for specific instructions.
- 2 On the Welcome screen, click Next.
- 3 Choose an installation directory, then click Next.

- 4 Validate the license for each database installed on your system:
 - Click Browse, navigate to the directory where your database is installed.
 - Choose the *.lmp* file associated with your database.

If the license fails validation, you cannot advance to the next screen. Although you can click the “Continue Without License Validation” checkbox to skip license validation, your database may fail to start after you install this upgrade.
- 5 Choose Licensed Edition, then click Next.
- 6 On the next screen, choose a location, then read and accept the license agreement. You must accept an appropriate license agreement to install Sybase IQ. You can review all available license agreements on the Sybase Web site at <http://www.sybase.com/softwarelicenses>.
- 7 When the installer prompts you for the Sybase IQ license key(s):
 - Click Browse, navigate to where you stored your license, and select it. Perform this action only if you are installing an unserved license.
 - Click “Use previously deployed license server,” then specify the server’s host name and port number. Use this option to use a previously deployed license server.
 - Click “Use previously deployed license server,” then specify the server’s host name and port number. Use this option to use a previously deployed license server.
 - Click “Continue installation without a license key” to install Sybase IQ in evaluation mode.
- 8 Click Next, then choose the appropriate product edition and license type.
- 9 Click Next, then follow the instructions on the screen.

3.2.2 Silent installations

If you plan to install Sybase IQ 15.1 ESD #3 silently, use an editor to change the database validation parameters in *SilentInstall.txt*. You must provide a path for each database you use. To add the paths, use this flag:

```
-V lmpFilePathN=<path>
```

where N is a number that determines the validation sequence; $\langle path \rangle$ is the path to the *.lmp* file you want to validate. To determine the validation sequence, assign 1 (`lmpFilePath1...`) to the first license you want to validate; 2 (`lmpFilePath2...`) to the second, and so on.

Note For validating more than three databases, you can add new lines containing `-v lmpFilePathN=<path>` to the file. Be sure to maintain the validation sequence (i.e., `lmpFilePath5` should precede `lmpFilePath6`).

To override license validation, change the

```
-V ContinueWithoutLicenseValidation
```

parameter to `true`. See the comments in *SilentInstall.txt* for other configuration options.

3.2.3 UNIX post installation tasks

If you install Sybase IQ 15.1 ESD #3 on a UNIX-like OS, you must set the `IQ-15_1` variable to run Sybase IQ Server.

For the `tcsh` or `C` (`csh`) shell, use:

```
% source $SYBASE/IQ-15_1/IQ-15_1.csh
```

For the Korn (`ksh`), Bash, or Bourne (`sh`) shells, use:

```
% . $SYBASE/IQ-15_1/IQ-15_1.sh
```

3.3 Sybase IQ 15.1 requires a license

Sybase IQ 15.1 uses the Sybase Software Asset Management (SySAM) licensing mechanism that provides system administrators with a means to enable and monitor their site's use of Sybase products.

Sybase IQ includes separate SySAM licenses for each product edition and separate licenses for any optional features available for that edition. For more information, see Chapter 2, "Licensing your software" in the *Installation and Configuration Guide for Sun Solaris*.

3.4 Updating SySAM license server

If you use a SySAM license server with Sybase IQ, Sybase IQ 15.1 on Sun Solaris requires SySAM license server version v11.5.0.0, which is distributed on the media with the product. Using an existing older version of the SySAM license server can cause the Sybase IQ server to become unresponsive.

You can check the version of your current SySAM license server by running:

```
cd $SYBASE/SYSAM-2_0/bin
./lmutil lmver lmgrd
```

If you use a SySAM license server with Sybase IQ and your SySAM license server version is earlier than v11.5.0.0, install the SySAM license server provided on the media with Sybase IQ 15.1 before you start the Sybase IQ server.

3.5 Sybase Central plug-in compatibility

The Sybase IQ 15.x plug-in for Sybase Central™ has been upgraded from version 4.3 to version 6.0 and is NOT compatible with previous versions. You can only use the 15.x plug-in with 15.x servers; you cannot use the 15.x plug-in to connect to 12.x servers.

3.6 No Sybase IQ server support on Linux 32-bit

The Sybase IQ 15.x client is supported on Red Hat 5.2 Linux for 32-bit systems, but there is no Sybase IQ 15.x server support on Linux for 32-bit systems.

3.7 SDK no longer included

Because Sybase IQ 15.x now embeds many features provided by the Software Development Kit, the SDK is no longer distributed as part of the product.

The Copy Definition utility (defncopy) is included in the Software Development Kit and is no longer shipped with Sybase IQ.

3.8 Open Client no longer included

Open Client™ is no longer installed as part of the Sybase IQ product. Sybase IQ 15.1 does provide versions of Open Client utilities that have limited functionality to enable INSERT...LOCATION:

- iqisql
- iqdsedit
- iqdscp (UNIX only)
- iqocscfg (Windows only)

3.9 Review *.iqmsgR if migration fails [CR 548182]

During migration, the 15.x server creates a message file (*.iqmsgR) as it reloads the generated schema. This file is normally deleted as part of a clean-up operation for successful migrations. If the migration fails during the reload stage, clean-up does not occur, and *.iqmsgR remains in *unload* directory. *.iqmsgR may contain information that can help solve your migration problems.

See “Migrating Data” in the *Sybase IQ Installation and Configuration Guide* for more information.

3.10 Starting Sybase Central [CR 607645]

On UNIX or Linux platforms, start Sybase Central from the `$IQDIR15/bin64` or `$IQDIR15/bin32` directory. If you try to start Sybase Central from another directory, Sybase Central may not start or function correctly.

If you try to start Sybase Central from the `$SYBASE/shared/sybcentral600` directory, for example, Sybase Central opens but will not load all plug-ins properly.

3.11 iqunload -au does not have “UNIQUE CONSTRAINT” index clause [CR 608611]

Customers who have successfully used the iqunload -au utility to migrate their 12.7 database to Sybase IQ 15.0, 15.0 ESD #1, 15.1, or 15.1 ESD #1 might be in a situation where the migrated system tables have incorrect data. This incorrect data can result in the inability to modify the schema or rows within the database. If this happens, the server may return an internal system exception. The server will also permit users to update the primary key values involved in a referential relationship in such a way that would be in RI violation.

Important!

- If you have not migrated your database, download and install 15.1 ESD #3, and use this release to perform database migration. Perform the Correcting migrated data procedure from step 3.
 - If you migrated your database with Sybase IQ 15.0, 15.0 ESD #1, 15.1, or 15.1 ESD #1, perform the Correcting migrated data procedure.
-

❖ Correcting migrated data

This procedure assumes that you migrated your database with Sybase IQ 15.0, 15.0 ESD #1, 15.1, or 15.1 ESD #1. If you plan to migrate a multiplex database, you must apply these steps to the coordinator node.

- 1 Execute the following against the migrated database using current software:

```
create temporary procedure iq_check_system_tables()
result ( Action char(255) )
begin
    if exists( select * from SYS.SYSIDX si
              join SYS.SYSIQIDX siq on (si.table_id =
              siq.table_id and si.index_id = siq.index_id)
              join SYS.SYSIDXCOL sic on (sic.table_id =
              si.table_id and sic.index_id = si.index_id)
              join SYS.SYSTAB st on (si.table_id =
              st.table_id)
              where st.server_type = 2 and siq.index_owner =
              'IQ' and si.index_category = 3 and si."unique" = 1)
    then
        select 'Database system tables require 15.1
        ESD 2 'alter database upgrade'' as Action
        else
```

```

        select 'Database system tables check was
successful.' as Action
        end if
    end;
    call iq_check_system_tables();

```

- 2 What you do next depends on the output message:

This message...	Means this...
Database system tables check was successful.	The migrated database does not contain any invalid system table information. You do not need to take any corrective action. No additional steps are required.
Database system tables require 15.1 ESD 2 'alter database upgrade'	The migrated database contains invalid entries in the system tables. Please follow the remaining steps outlined here.

- 3 Download and install Sybase IQ 15.1 ESD #3 and use this release to perform the rest of the steps below.
- 4 Make a backup copy of the *.db* and *.log* file(s) from your database.
- 5 Start the database with the new ESD software.
- 6 Connect to the database with Interactive SQL (dbisql) and execute the alter database upgrade SQL command. This will produce the output:

```

Database upgrade started
Creating system views
Creation of system views completed
Setting option values
Setting option values completed
creating variable...
**           System table update required.
**           System table update completed.
Creating migration procedures
Creation of migration procedures completed
Creating jConnect procedures
Creation of jConnect procedures completed
Database upgrade finished
29 row(s) affected
Execution time: 68.453 seconds

```

During the upgrade the system tables will be updated, the two lines prefixed with ** show the status of this update. (The ** characters will not be displayed on the console).

- 7 Repeat step 2. This should return the message:

```
Database system tables check was successful.
```

- 8 If there were no changes to the data stored in the database, then the above system table update is all that is required and the process is completed.
- 9 If the data in the database was modified after migration or you wish to run an extra verification check then the stored procedure:

```
iq_migration_ri_verification()
```

This procedure is temporary is and should be run as DBA. To load this procedure, execute this command from Interactive SQL (dbisql):

```
call sa_exec_script( 'checkpkfk.sql' )
```

After you load the script, use this command to execute the procedure:

```
call iq_migration_ri_verification()
```

This procedure will scan the database for PK/FK relationships, if any are found it will generate a script that will be used to verify the PK/FK relationships. The execution of `iq_migration_ri_verification()` will be fairly quick as it just scans the schema of the database. The generated verification script may take a while to complete.

If `iq_migration_ri_verification()` detects PK/FK relationships, the output looks like this:

```
=====
IQ PK/FK Referential Information
=====
There are 5 PK/FK relationships detected in this
database.
```

A verification script will be generated to:

```
W:/iq-15/unload/127/db/runricheck.sql
```

This script contains SQL which will perform a join on tables where there exists a primary and foreign key relationship. As a result, this script may take some time to execute.

In the event this script detects missing primary key values for which there is a foreign key value present, a log entry will be produced that contains the table name and column names that require corrective action.

To execute the generated verification script:

```
call sa_exec_script('W:/iq-
15/unload/127/db/runricheck.sql');
```

Please continue to the next step to run the runricheck.sql script.

Please continue to the next step to run the runricheck.sql script.

If `iq_migration_ri_verification()` does not detect any PK/FK relationships, then the output will look like this:

```
Database does not contain any PK/FK relationships.
No verification is required.
```

10 Load and Execute runricheck.sql.

Follow the instructions provided by `iq_migration_ri_verification()` to load the *runricheck.sql* script:

```
call sa_exec_script('W:/iq-
15/unload/127/db/runricheck.sql');
```

4. New features and behavior changes

For new features and behavior changes, see *New Features Summary Sybase IQ 15.1*, which you should read before you install Sybase IQ 15.1. This section describes changes that were not included in that book or the documentation set.

Sybase IQ supports many of the new features in SQL Anywhere® 11. For more information, see *SQL Anywhere 11 – Changes and Upgrading*.

4.1 Highlights of new features in Sybase IQ 15.1 ESD #3

Sybase IQ 15.1 ESD #3 includes new features in the areas of administration and performance.

- Sybase IQ now includes a mechanism that verifies an existing database backup. See “Verification of a database backup [CR 487181]” on page 14.
- Sybase IQ now allows you to control file system buffering on Linux, Linux IBM, and AIX platforms, in addition to Solaris and Windows platforms. See “Controlling file system buffering [CR 487762]” on page 18.

4.2 Data Definition Language (DDL) changes

This section contains new features and changes related to DDL.

4.2.1 CREATE LOCAL TEMPORARY TABLE syntax [CR 605785]

The CREATE TABLE statement now includes the LOCAL keyword for creating temporary tables. The syntax is:

```
CREATE [ { GLOBAL | LOCAL } TEMPORARY ] TABLE  
[ owner. ]table-name  
... ( column-definition [ column-constraint ] ...
```

In a procedure, use the CREATE LOCAL TEMPORARY TABLE statement, instead of the DECLARE LOCAL TEMPORARY TABLE statement, when you want to create a table that persists after the procedure completes. Local temporary tables created using the CREATE LOCAL TEMPORARY TABLE statement remain until they are either explicitly dropped, or until the connection closes.

Local temporary tables created in IF statements using CREATE LOCAL TEMPORARY TABLE also persist after the IF statement completes.

4.2.2 Behavior change declaring variables in procedures [CR 584334]

In Sybase IQ 12.7 and later releases, new procedures are invalid when there are duplicate declarations of a variable within a compound statement. This error is reported:

```
There is already a variable named '@variable_name'  
SQLCODE=-261, ODBC 3 State="42000"
```

Ensure there is only a single instance of a variable in a compound block.

4.3 Startup and connection changes

This section contains new features and changes related to startup and connection.

4.3.1 New options for mkiqdemo script

Prior to Sybase IQ 15.1, the scripts *mkiqdemo.bat* and *mkiqdemo.sh* created the IQ demo database in the current directory using a relative path, and you could not create a multiplex using this database. The scripts now accept two new parameters, `-absolute` and `-relative`. Running either script with the parameter `-absolute` changes the SQL script to create the IQ demo database with the full path name. Use `-relative` to retain the old behavior of the scripts.

4.3.2 Stop utility limitation [CR 543495]

The `stop_iq` utility on UNIX platforms fails to report running servers when the path to the server exceeds 74 characters.

As a workaround, use a `ps -ef` command to display servers with long paths. For example:

```
ps -ef|grep iqsrv

rsmithson 1133      1      0 07:04:32 ?          223:35
/sunsys1234/users/rsmithson/mybigtest1234_withdeletion
_allcol/IQ-15_1/bin64/iqsrv1

rsmithson 2046    862    0 10:02:30 pts/3      0:00 grep
iqsrv
```

4.3.3 Server startup utility switches removed [CR 537028]

These startup utility switches are removed in Sybase IQ 15.x:

- `-iqdroppls` – replaced by the `DROPLEAKS` option for `sp_iqcheckdb`
- `-iqlocalreplay` – obsolete with removal of local stores in 15.x
- `-iqnolocalreplay` – obsolete with removal of local stores in 15.x
- `-iqnomain` – obsolete with removal of local stores in 15.x

4.4 Administration and troubleshooting improvements

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

4.4.1 Change in encapsulated virtual backup behavior [CR 541342]

In versions prior to 12.6 ESD#9, no one could log onto IQ during encapsulated virtual backup and anyone connected appeared to suspend operations while the external script was called.

Encapsulated virtual backups no longer suspend database activity, including writes.

4.4.2 Verification of a database backup [CR 487181]

As of version 15.1 ESD #3, Sybase IQ includes a mechanism that verifies an existing Sybase IQ version 12.6 or later database backup using the VERIFY clause of the RESTORE SQL statement. The verification process checks the specified archive for the same errors a restore process checks, but performs no write operations. All status and error messages are written to the server log file.

Syntax

```
RESTORE DATABASE 'db_file'
FROM 'archive_device' [ FROM 'archive_device' ]...
[ CATALOG ONLY ]
[ KEY key_spec ]
[ [ RENAME dbspace-name TO 'new-dbspace-path']...
  | VERIFY [ COMPATIBLE ] ]
```

See “RESTORE statement” in Chapter 1, “SQL Statements” in *Reference: Statements and Options* for the complete RESTORE syntax, parameter descriptions, and usage.

Example

```
RESTORE DATABASE <database_name.db>
FROM '/sys1/dump/dmp1 '
FROM '/sys1/dump/dmp2 '
VERIFY
```

Usage

The VERIFY clause of the RESTORE command directs the server to validate the specified Sybase IQ database backup archives for a full, incremental, incremental since full, or virtual backup without performing any write operations.

You cannot use the RENAME clause with the VERIFY clause; an error is reported.

The backup verification process can run on a different host than the database host.

If you specify the `COMPATIBLE` clause with `VERIFY`, the compatibility of an incremental archive is checked with the existing database files. If the database files do not exist on the system on which you invoke `RESTORE...VERIFY COMPATIBLE`, an error is returned. If you specify `COMPATIBLE` while verifying a full backup, the keyword is ignored; no compatibility checks need to be made while restoring a full backup.

You must have the database and log files (`.db` and `.log`) to validate the backup of a read-only dbspace within a full backup. If you do not have these files, validate the entire backup by running `RESTORE...VERIFY` without the `READONLY dbspace` clause.

Note The verification of a backup archive is different than the database consistency checker (DBCC) verify mode (`sp_iqcheckdb 'verify...'`). `RESTORE VERIFY` validates the consistency of the backup archive to be sure it can be restored, whereas DBCC validates the consistency of the database data.

Run `sp_iqcheckdb 'verify...'` before taking a backup. If an inconsistent database is backed up, then restored from the same backup archive, the data continues to be in an inconsistent state, even if `RESTORE VERIFY` reports a successful validation.

Verification of an incremental backup

If you specify `RESTORE VERIFY` without `COMPATIBLE` for an incremental restore, Sybase IQ does not look for any dbspaces and does not perform any compatibility checks. No warning is reported, even if the files do not exist. The compatibility check is performed only when you include the `COMPATIBLE` clause.

If you specify `RESTORE VERIFY COMPATIBLE` for an incremental restore and the IQ catalog store or any of the Sybase IQ dbspaces do not exist, the compatibility check cannot be made; an error is reported and the operation fails.

During validation of an incremental backup, the `RESTORE VERIFY COMPATIBLE` process opens the Sybase IQ dbspaces in read-only mode to perform consistency checking. No dbspaces are modified.

In the case of incremental restores, if the database has been modified or the particular incremental archive is not the correct archive for the database, `RESTORE VERIFY COMPATIBLE` reports the error `Database has changed since last restore. (SQLCODE -1012008, SQLSTATE QUA08)` or `This restore cannot immediately follow the previous restore. (SQLCODE -1012009, SQLSTATE QUA09)`.

Verification progress reporting

The RESTORE VERIFY process verifies every stripe specified in the command. As the verification process checks stripes and their corresponding files, it reports progress in terms of the number of IQ blocks verified. After every 5000 IQ blocks verified, a message is written to the server log file with the percentage of the number of IQ blocks completed:

```
5000/100000 (5%) Blocks verified
```

A final message is written to the server log file when 100% of the IQ blocks are verified.

The messages RESTORE VERIFY Started, the number of IQ blocks to be verified, and RESTORE VERIFY successfully completed are also written in the server log, when the verify action starts and completes, respectively. For example:

```
I. 11/17 06:45:24. VERIFY RESTORE Started
I. 11/17 06:45:24. Total number of IQ blocks to be
verified: 764
I. 11/17 06:45:24. Total number of IQ blocks verified:
764/764 ( 100 % )
I. 11/17 06:45:24. VERIFY RESTORE Successfully Complete
```

Verification error reporting

If the verification process finds errors after which it can continue, the process continues checking the archive and logs information for the errors detected. The errors for which the verification can continue are:

```
Header of block to be restored appears to be
corrupted. (SQLCODE -10120111, SQLSTATE QUA11)
```

```
Media data appears corrupted (bad checksum).
(SQLCODE -1012012, SQLSTATE QUA12)
```

```
Media meta data appears corrupted (boundary record).
(SQLCODE -1012013, SQLSTATE QUA13)
```

```
Media meta data appear corrupted (multiple begin
boundary records). (SQLCODE -1012014, SQLSTATE QUA14)
```

If any of these errors are found and the verification process can continue to the end of the archive, this error is reported:

```
The verification of the provided archive has failed.
Please check the server log for details of the errors
thrown during verify.
```

If any error pertaining to RESTORE is found other than the errors above, the error that occurred is reported, and the verification process stops.

Permissions

Must have DBA or BACKUP authority.

See also

- “RESTORE statement” in Chapter 1, “SQL Statements” in *Reference: Statements and Options*
- “BACKUP statement” in Chapter 1, “SQL Statements” in *Reference: Statements and Options*
- Chapter 12, “Data Backup, Recovery, and Archiving” in *System Administration Guide: Volume 1*

4.4.3 Default password value change

The default password is now sql in lower case.

4.5 Data load, update, and extraction enhancements

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

4.5.1 Sybase ETL 4.8 enhancements

Sybase IQ 15.1 is certified with Sybase ETL 4.8, which includes ease of use enhancements for Sybase IQ. Sybase ETL 4.8 is now packaged separately from Sybase IQ.

Sybase ETL 4.8 allows you to bulk-load data from text files located on remote host machines into Sybase IQ 15.1, using the “DB Bulk Load - Sybase IQ” and “IQ Loader File via Load Table” components.

In earlier versions of Sybase ETL, text files were required to be available on the same host machine as the IQ server. Also, Sybase ETL was required to be installed on the same machine as Sybase IQ. In Sybase ETL version 4.8, you no longer need to install Sybase ETL and Sybase IQ on the same machine; the ETL server and Sybase IQ can communicate in a network environment that allows you to bulk-load from a remote machine in a single step.

For more information, see the Sybase ETL 4.8 online documentation at Product Manuals at <http://sybooks.sybase.com>.

4.5.2 Inserting default values

Sybase IQ supports inserting default values with this syntax:

```
INSERT [ INTO ] [ owner.]table-name [ ( column-name [, ...] ) ]
... { DEFAULT VALUES | VALUES ( [ expression | DEFAULT,... ] ) }
```

For more information, see the section “INSERT statement” in Chapter 1, “SQL statements” in *Reference: Statements and Options*.

4.5.3 SQLCODE change for 0 rows updated [CR 533932]

When an update to a Sybase IQ table affects 0 rows, Sybase IQ now sets the SQLCODE to 100 (NOTFOUND), rather than 0 (NOERROR). This behavior change provides a more accurate return code that informs you that no rows were found that match the update criteria, so no rows were updated.

4.6 Query enhancements, optimization, and changes

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

4.6.1 OS_FILE_CACHE_BUFFERING option behavior change [CR 487762]

In versions prior to Sybase IQ 15.1 ESD #3, the database option OS_FILE_CACHE_BUFFERING controls the file system buffering of both IQ Main dbspaces and IQ Temporary dbspaces. As of Sybase IQ 15.1 ESD #3, OS_FILE_CACHE_BUFFERING controls the file system buffering of IQ Main dbspaces only. The new database option OS_FILE_CACHE_BUFFERING_TEMPDB controls the file system buffering of IQ Temporary dbspaces.

See “Controlling file system buffering [CR 487762]” on page 18 for a description of the database option OS_FILE_CACHE_BUFFERING_TEMPDB.

4.7 Performance enhancements

This section contains new features and changes related to performance enhancements.

4.7.1 Controlling file system buffering [CR 487762]

As of Sybase IQ 15.1 ESD #3, you can control file system buffering on Linux, Linux IBM, and AIX platforms using the database options OS_FILE_CACHE_BUFFERING and OS_FILE_CACHE_BUFFERING_TEMPDB. This direct I/O functionality is supported in Linux kernel versions 2.6.x. Sun Solaris and Windows platforms already support direct I/O.

To enable direct I/O on Linux kernel version 2.6 and AIX, also set the environment variable `IQ_USE_DIRECTIO` to 1. Direct I/O is disabled by default in Linux kernel version 2.6 and AIX. `IQ_USE_DIRECTIO` has no effect on Sun Solaris and Windows.

Note Sybase IQ does not support direct I/O on Linux kernel version 2.4. If you set the `IQ_USE_DIRECTIO` environment variable on Linux kernel version 2.4, the Sybase IQ server does not start. The error “Error: Invalid Block I/O argument, maybe <pathname> is a directory, or it exceeds maximum file size limit for the platform, or trying to use Direct IO on unsupported OS” is reported.

In versions prior to Sybase IQ 15.1 ESD #3, the database option `OS_FILE_CACHE_BUFFERING` controls the file system buffering of both IQ Main dbspaces and IQ Temporary dbspaces. As of Sybase IQ 15.1 ESD #3, `OS_FILE_CACHE_BUFFERING` controls the file system buffering of IQ Main dbspaces only. The new database option `OS_FILE_CACHE_BUFFERING_TEMPDB` controls the file system buffering of IQ Temporary dbspaces.

To enable direct I/O for IQ Main dbspaces, set `OS_FILE_CACHE_BUFFERING` to OFF. To enable direct I/O for IQ Temporary dbspaces, set `OS_FILE_CACHE_BUFFERING_TEMPDB` to OFF. The default value of both options is OFF, which means there is no file system buffering by default.

OS_FILE_CACHE_BUFFERING_TEMPDB option

Function	Controls the use of file system buffering for IQ Temporary dbspaces.
Allowed values	ON, OFF
Default	OFF
Scope	Can be set for the PUBLIC group only. You must shut down and restart the database for the change to take effect. DBA permissions are required to set this option.

Description	<p>Setting <code>OS_FILE_CACHE_BUFFERING_TEMPDB</code> to <code>OFF</code> prevents file system buffering for IQ Temporary Store files. Turning off file system buffering saves a data copy from the file system buffer cache to the main IQ buffer cache. Usually this reduces paging caused by competition for memory between the IQ buffer manager and the file system buffer of the operating system. When <code>OS_FILE_CACHE_BUFFERING_TEMPDB</code> reduces paging, this option improves performance; however, if the IQ page size for the database is less than the block size of the file system (typically only in testing situations), performance decreases, especially during multiuser operation.</p> <p>Experiment with this option to determine the best setting for different conditions. You must restart the database for the new setting to take effect. This option has no effect on HP-UX and HP-UXi.</p> <p>In Linux, direct I/O is supported in kernel versions 2.6.x.</p> <p><code>OS_FILE_CACHE_BUFFERING</code> controls file system buffering for IQ Main Store files.</p>
See also	<p>“<code>OS_FILE_CACHE_BUFFERING</code> option” in Chapter 2, “Database Options” in the <i>Reference: Statements and Options</i></p> <p>Chapter 4, “Managing System Resources” in the <i>Performance and Tuning Guide</i></p>

4.8 Interactive SQL (dbisql) enhancements and changes

This section contains new features and changes related to Interactive SQL (dbisql).

4.8.1 Numeric values display with commas in dbisql [CR 531574]

dbisql displays numeric values with commas in the thousands place. For example, the values for `TxnID` and `LastIdle` returned by the `sp_iqconnection` stored procedure display as “1,352” and “9,375”.

5. Product compatibilities

This section describes the compatibility of Sybase IQ with other products. Certification with client and partner products is an ongoing process. See “Sybase certifications on the Web” on page 62 for instructions on accessing the most current online certification information.

5.1 Sybase products

These Sybase products are certified with Sybase IQ 15.1:

- SQL Anywhere 11.0 or later
- Adaptive Server® Enterprise 12.5.x and 15.0
- jConnect™ for JDBC™ 5.x and 6.x
- Sybase ETL 4.8

Refer to “Sybase certifications on the Web” on page 62 for instructions on accessing the latest online certification information.

6. 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 Sybase IQ 15.1 servers and databases that have been upgraded from any earlier version of Sybase IQ, unless specified otherwise.

6.1 Database name length restriction [CR 365281]

The dbbackup utility truncates the database name to 70 characters and creates a target file with a truncated name. Sybase IQ uses dbbackup when synchronizing secondary servers. Due to dbbackup restrictions, database names must be less than 70 characters long.

6.2 Restriction on number of devices for backup [CR 542986, CR 541872]

Sybase IQ supports a maximum of 32 hardware devices for backup. For faster backups, specifying one or two devices per core will help to avoid hardware and IO contention. Set the SIZE parameter on the BACKUP command to avoid creating multiple files per backup device and consider the value used in the BLOCK FACTOR clause on the BACKUP command. See Chapter 1, “SQL Statements,” in *Reference: Statements and Options*.

6.3 Dbspace 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 dbspace files on network drives owned by another server.

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

6.4 Dbspace naming in RESTORE command [CR 561366]

If the dbspace name contains a file extension such as *.iq*, *.iqtmp*, or *.iqloc*, you must enclose it in double quotation marks when specifying the name in a RESTORE command RENAME clause.

For example:

```
RENAME local1 TO '/work/local1_res.iqloc.iqloc'  
DBSPACENAME "local1_res.iqloc"
```

or

```
RENAME "test_prod2.iq" TO '/test/test_prod7.iq'
```

6.5 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 this 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 a front-end tool or an application requires scrolling, but this option is set ON, then you will see error messages like "When Force_No_Scroll_Cursors = 'ON' scrolling cursor operations are not supported."

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

6.6 Unexpected query results

In a few unusual circumstances, differences in semantics between SQL Anywhere and Sybase IQ may produce unexpected query results. These circumstances are:

- A query is issued from inside a user-defined function
- A SELECT statement has no FROM clause
- A FROM clause contains some tables that were created IN SYSTEM and others that were not created IN SYSTEM

In these circumstances, subtle differences between the semantics of SQL Anywhere and Sybase IQ may be exposed. These differences include:

- Sybase IQ treats the CHAR and VARCHAR data types as distinct and different, and SQL Anywhere treats CHAR data as if it were VARCHAR
- The behavior of the RAND function when passed an argument is deterministic in Sybase IQ and non-deterministic in SQL Anywhere

6.7 Interactive SQL restrictions

Sybase recommends that you use the Java edition of dbisql, rather than dbisqlc or iqisql, 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 you do not set this option, the output file is not created.

To set the “Show multiple result sets” option, click Tools | Options in the dbisql window, choose “SQL Anywhere,” and select the “Results” tab. Then select “Show all results sets.”

- Online help is not available for dbisql and dbisqlc. [CR 560925]
- The query plan in the dbisql Plan tab is a SQL 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:

```
% xterm -sb
```

6.8 Perl interface (DBD::SQLAnywhere) not supported by Sybase IQ [CR 574220]

The Perl interface (DBD::SQLAnywhere) is not supported by Sybase IQ.

7. Known problems

For a description of known issues in Sybase IQ 15.1, see the following sections. If there is a workaround for a problem, it is provided. Change request numbers appear in brackets, as in [CR 563004].

See also “Restrictions” on page 21 for more information.

“Documentation updates and clarifications” on page 31 contains details that were not documented in time for this release.

7.1 Sybase IQ operations

This section reports workarounds for issues with Sybase IQ operations.

7.1.1 sp_iqstatus returns error on Chinese and Japanese locale installations [CR 622928]

On Chinese and Japanese locale Sybase IQ servers, running sp_iqstatus returns an error similar to:

```
Could not execute statement.
Syntax error near '2010' on line 1
SQLCODE=-131, ODBC 3 State="42000"
Line 1, column 1
```

The workaround for this problem is:

1 Run these commands:

For Chinese locales

```
% cd $IQDIR15/res
% rm dblgzh_iq11_eucgb.res
% rm dblgzh_iq11_cp936.res
```

For Japanese locales

```
% cd $IQDIR15/res
% rm dblgja_iq11_eucjis.res
% rm dblgja_iq11_sjis.res
```

2 Restart Sybase IQ.

After this workaround, certain strings in the *.iqmsg* files and the output of *sp_iqmpxinfo*, *sp_iqstatistics*, and *sp_iqstatus* will be in English instead of Chinese or Japanese.

7.1.2 Stored procedures for information only [CR 560032]

Stored procedures *sp_iqbackupsummary* and *sp_iqrestoreaction* may give incorrect output if:

- Backup action was performed within 1 hour of DST end activity, or
- System time is changed (for example, set back in the past).

Stored procedures serve informational purposes. Sybase recommends that the DBA compare the output of the *sp_iqbackupsummary* and *sp_iqrestoreaction* procedures against the output of views *SYSIQBACKUPHISTORY* and *SYSIQBACKUPHISTORYDETAIL* to decide the restore sequence.

7.1.3 Sample programs may need modification [CR 562633]

The sample programs shipped with IQ in *\$IQDIR15/samples/sqlanywhere* may need to be modified to work in your environment. One symptom of this problem is the C ODBC sample program does not compile and "undefined reference" errors are reported.

For Java programs, download JDK 1.6 and possibly other components like *jdbcConnect*. For C programs, you need the *gnu* compiler and you need to modify the *makefiles* in order to compile them in your environment.

7.1.4 RSA certificate creation fails [CR 557702]

RSA certificate creation fails with the error `createcert` command not found, because the `createcert` utility is missing. The utility is currently available only on the AIX64, LinuxAMD64, Sun64, and Win32 platforms. Please contact Sybase, if you need this utility.

7.2 Multiplex environment

This section reports workarounds for issues within a multiplex environment.

7.2.1 Emergency shutdown on secondary servers after creating a proxy table [CR 557714]

A timing problem may cause an emergency server shutdown when a user accesses a new proxy table on a secondary server.

To avoid this problem, users should either reconnect or wait for some time and start another transaction before trying to use a newly-created proxy table.

7.3 Sybase Central

This section reports workarounds for issues with Sybase Central.

7.3.1 Interactive SQL Java option label inconsistency [CR 578634, CR 578629]

If you start Interactive SQL Java (`dbisql`) through Sybase Central and click Tools | Options, the Options window displays an option labelled “SQL Anywhere” in the left pane. This option should be labelled “Sybase IQ” and will be corrected in a later version of Sybase IQ.

If you start Interactive SQL Java (`dbisql`) using the command line, the option has the correct label “Sybase IQ.”

7.3.2 Remote client cannot connect to IQ Agent [CR 563823]

When Sybase Central running on a remote client attempts to connect to the IQ Agent running on the server host, this error may be reported:

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

The root cause of this problem may be the configuration of the network, which incorrectly resolves the hostname to loopback IP address 127.0.0.1, rather than the real IP address. You can verify the source of the problem by running the command `ping <hostname>`, which returns an IP address. If the IP address returned is 127.0.0.1, rather than the real IP address of the host, then the network configuration is the problem.

To resolve this issue, ask your system administrator to change the network configuration, so that the hostname is resolved as the real IP address. On a UNIX system, you can edit the system file `/etc/hosts`, so that the hostname is resolved to the real IP address.

7.3.3 Creating proxy tables with “Selected Columns” option [CR 559895]

Creating a proxy table in the Sybase Central wizard with the All Columns option creates the proxy table as expected. However, if you create a proxy table in the Sybase Central wizard with the Selected Columns option, even if you select all columns, the proxy table does not show the data.

The workaround is either to build the proxy table with the All Columns option, or to use `dbisql` to build the proxy table with just the selected columns that you want.

7.3.4 Sybase Central operation issues [CR 549750]

These issues with Sybase Central operations while using the Create Database Wizard and browsing for a file on remote systems are related to JRE 6.0. The workarounds are either to enter in the text field the required path name to the file, rather than selecting the file using the wizard, or to run Sybase Central on the system on which you are creating the database.

- Running Sybase Central in either Windows or Linux: you cannot traverse inside a newly created directory, while browsing on a remote system
- Running Sybase Central in Windows: double clicking on directories on a remote Linux system is non-responsive
- Running Sybase Central in either Windows or Linux: the Create New Directory button is disabled, while browsing for files on a remote Linux system
- Running Sybase Central in Linux: while navigating the C: drive on a remote Windows system, if you click the Up > Up > ... button, the listing of files for the C: drive is not correct

- Running Sybase Central on all platforms: renaming a file on a remote system does not work. After you enter the new name and press Enter, there is no response and the rename operation fails without returning an error.
- Running Sybase Central on all platforms: clicking and double clicking files and directories on a remote system does not work as expected. Sometimes Sybase Central navigates into the directory and sometimes it edits the directory name for rename. Similarly for files, sometimes Sybase Central selects the file and some times it edits the file name for rename.

8. Documentation for this version

Product packaging has changed for Sybase IQ 15.1. Along with the packaging changes, there are documentation changes. The documentation is now consolidated on two pieces of media. All the information you need to get started and to use the new features in Sybase IQ 15.1 is on the Sybase IQ Getting Started CD and the Sybase IQ SyBooks™ CD.

Table 1 lists documentation on the Sybase IQ 15.1 Getting Started CD.

Table 1: Sybase IQ 15.1 Getting Started CD contents

Part Number	Book Title
DC01014-01-1510-01	<i>Quick Start</i>
DC74373-01-1510-01	<i>Release Bulletin for Sun Solaris</i>
DC30066-01-1510-01	<i>Installation and Configuration Guide for Sun Solaris</i>

Table 2 lists Sybase IQ documentation on the Sybase IQ 15.1 SyBooks CD.

Table 2: Sybase IQ SyBooks contents

Part Number	Book Title
DC01087-01-1510-01	<i>New Features Summary Sybase IQ 15.1</i>
DC00171-01-1500-01	<i>New Features in Sybase IQ 15.0</i>
DC00462-01-1510-01	<i>Error Messages</i>
DC38159-01-1510-01	<i>Introduction to Sybase IQ</i>
DC38151-01-1510-01	<i>Reference: Building Blocks, Tables, and Procedures</i>
DC00801-01-1510-01	<i>Reference: Statements and Options</i>
DC00170-01-1510-01	<i>System Administration Guide: Volume 1</i>
DC00800-01-1510-01	<i>System Administration Guide: Volume 2</i>
DC00802-01-1510-01	<i>Using Sybase IQ Multiplex</i>
DC01034-01-1510-01	<i>User-Defined Functions Guide</i>
DC01033-01-1510-01	<i>IMSL Numerical Library User's Guide: Volume 2 of 2 C Stat Library</i>

Table 3 lists the Sybase IQ 15.1 documents included in the final stage of the release.

Table 3: Sybase IQ documents available online late July, 2009

Part Number	Book Title
DC00169-01-1510-01	<i>Sybase IQ Performance and Tuning Guide</i>
DC00168-01-1510-01	<i>Sybase IQ Utility Guide</i>
DC00172-01-1510-01	<i>Large Objects Management in Sybase IQ</i>
DC01015-01-1510-01	<i>Advanced Security in Sybase IQ</i>
DC00171-01-1500-02	<i>New Features in Sybase IQ 15.0</i>

Check the Sybase Web site Product Manuals at <http://sybooks.sybase.com> for the latest available versions of the Sybase IQ 15.1 documents.

Please note these changes to the documentation set in Sybase IQ 15.1:

- Multiplex information is now collected in *Using Sybase IQ Multiplex*.
- *Sybase IQ Troubleshooting and Recovery Guide* is no longer published. You will now find this information in *System Administration Guide: Volume 1* in Chapter 13, "System Recovery and Database Repair," and Chapter 14, "Troubleshooting Hints."
- *Sybase IQ Reference Manual* is divided into two volumes: *Reference: Building Blocks, Tables, and Procedures* and *Reference: Statements and Options*.

- *Sybase IQ System Administration Guide* is divided into two volumes: *System Administration Guide: Volume 1*, covering database administration topics, and *System Administration Guide: Volume 2*, covering programming topics.
- The information formerly in the *Encrypted Columns* book is now included in *Advanced Security in Sybase IQ*.
- The chapter “Using OLAP,” formerly in the *Performance and Tuning Guide*, is now in the *System Administration Guide: Volume 2*.

Sybase IQ and SQL Anywhere

Because Sybase IQ is an extension of SQL Anywhere Server, a component of the SQL Anywhere package, Sybase IQ supports many of the same features as SQL Anywhere Server. The IQ documentation set refers you to SQL Anywhere documentation, where appropriate. Only a subset of the functionality in the SQL Anywhere documentation is supported by Sybase IQ. For more information, refer to Appendix A, “Compatibility with Other Sybase Databases” in *Reference: Building Blocks, Tables, and Procedures*.

You may also need to refer to this SQL Anywhere 11.0.1 documentation, which is included on the SyBooks CD:

- *SQL Anywhere Server – Database Administration* describes how to run, manage, and configure SQL Anywhere databases. It describes database connections, the database server, database files, backup procedures, security, high availability, and replication with Replication Server®, as well as administration utilities and options.
- *SQL Anywhere Server – Programming* describes how to build and deploy database applications using the C, C++, Java, PHP, Perl, Python, and .NET programming languages such as Visual Basic and Visual C#. A variety of programming interfaces such as ADO.NET and ODBC are described.
- *SQL Anywhere Server – SQL Reference* provides reference information for system procedures, and the catalog (system tables and views). It also provides an explanation of the SQL Anywhere implementation of the SQL language (search conditions, syntax, data types, and functions).

The SQL Anywhere 11.0.1 documentation is also available online on the Sybase Web site Product Manuals at <http://sybooks.sybase.com> and in DocCommentXchange at http://dcx.sybase.com/dcx_home.php.

Note Some of the Sybase IQ books now contain links directly to online SQL Anywhere 11.0.1 documentation sections referenced in the Sybase IQ documentation.

Documentation included for Sybase Software Asset Management (SySAM):

- *Users Guide Sybase Software Asset Management (SySAM) 2* introduces asset management concepts and provides instructions for establishing and administering SySAM 2 licenses.
- *SySAM 2 Quick Start Guide* tells you how to get your SySAM-enabled Sybase product up and running quickly.
- *FLEXnet Licensing End User Guide* explains FLEXnet Licensing for administrators and end users and describes how to use the tools which are part of the standard FLEXnet Licensing distribution kit from Sybase.

9. Documentation updates and clarifications

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

9.1 Interactive SQL (dbisql) online help

Online help is not available for the utilities Interactive SQL (dbisqlc) and Interactive SQL Java (dbisqlj).

9.2 New Features in Sybase IQ 15.0

This section contains updates to the *New Features in Sybase IQ 15.0*.

9.2.1 Using the Sybase IQ server after forced recovery [CR 543005]

If you start a Sybase IQ server in forced recovery mode with the `start_iq -iqfrec` server switch, the server marks all space in use. After this point, no write operations are permitted on the server, until you execute either `sp_iqcheckdb 'dropleaks database'` or `sp_iqcheckdb 'dropleaks <dbspacename>'`, which recovers space that is not in use.

In a multiplex environment, you can start the coordinator with the `-iqfrec` server startup switch only in combination with the single node mode only `-iqmpx_sn` startup switch. If you try to start the coordinator normally without running `sp_iqcheckdb 'dropleaks database'` or `sp_iqcheckdb 'dropleaks dbspace IQ_SYSTEM_MAIN'`, the server cannot start. The server needs to truncate and write a new first record to the TLV log, but there is no space available to write the TLV log. If the coordinator fails to start for this reason, restart the coordinator in single node mode and execute `sp_iqcheckdb 'dropleaks dbspace IQ_SYSTEM_MAIN'`, then restart the coordinator normally.

9.3 System Administration Guide: Volume 1

This section contains updates to the *System Administration Guide: Volume 1*.

9.3.1 Working with NULLs [CR 626626]

In “Moving Data In and Out of Databases > Converting data on insertion > Working with NULLS,” the syntax and example should be “NULL” instead of “NULLS”.

The correct syntax is:

```
NULL ({BLANKS | ZEROS | literal' ['literal']...})
```

The correct example is:

```
LOAD TABLE lineitem(
    l_orderkey NULL(ZEROS) ASCII(4),
    l_partkey ASCII(3),
    l_shipdate date('MM/DD/YY'),
    l_suppkey ascii(5),
    FILLER(1))
FROM 'C:\\MILL1\\tt.t'
PREVIEW ON
```

All other instances of the word “NULLS” in this section should be “NULLs”.

9.3.2 Extracting large object data [CR 621852]

Details about extracting large object data are not documented in *System Administration Guide: Volume 1*. See *Large Objects Management in Sybase IQ*.

9.3.3 Correlating connection information [CR 620445]

The section “Correlating connection information” in Chapter 14, “Troubleshooting Hints” should be replaced with this section:

Correlating connection information

Technical Support may ask you to set the `-zr` option on the `start_iq` command in your configuration file. This server startup option sets the request logging level to track statements sent to the server. Parameters are `ALL`, `NONE`, or `SQL`. The option produces a log file named for the server with the suffix `.srvlog`.

In the Sybase IQ message file `.iqmsg`, each connection to the server is identified by a connection handle. The `.iqmsg` message file records the errors, warnings, and tracing information for each connection. This procedure tells you how to correlate the connection identifiers in the `.srvlog` and `.iqmsg` files to find relevant information.

❖ Correlating connection information between the `.srvlog` and `.iqmsg` files

- 1 In the `.iqmsg` file, locate a connection of interest. For example:

```
Connect: SA connHandle: 1000000061
```

These lines show the `.iqmsg` log file contents for this connection:

```
16:14:59. 0000000062 Connect: SA connHandle: 1000000061
      SA connID: 31 IQ connID: 0000000062 User: DBA
03/17 16:15:00. 0000000062 Cmt 12064
03/17 16:15:00. 0000000062 PostCmt 0
03/17 16:15:00. 0000000000 Disconnect: SA connHandle: 1000000061
      SA connID: 31 IQ connID: 0000000062 User: DBA
```

- 2 Isolate all of the lines for the connection by searching the `.srvlog` file for the number that follows "SA connHandle" in the `.iqmsg` file. For example, search the `.srvlog` file for "1000000061":

```
16:14:59. [,1000000061,sp_iqdbspace,48,select
str_replace(dbspaceName,'',null) into dbspaceName_literal
03/17 16:14:59. P,1000000061,[S][0]DUMMY<seq>
03/17 16:14:59. ],1000000061,sp_iqdbspace,48
03/17 16:14:59. P,1000000061,[1]ISYSIQDBFILE<seq> JNL dbf<ISYSDBFILE>
JNL ISYSDBSPACE<ISYSDBSPACE>
03/17 16:14:59. [,1000000061,sp_iqdbspace,58,execute immediate with
quotes on 'iq utilities main into iq_dbpace_temp dbpace info
'|' dbpaceName
03/17 16:14:59. P,1000000061,[S]INSERT ROWS
03/17 16:14:59. P,1000000061,[S]INSERT ROWS
03/17 16:14:59. P,1000000061,[S]INSERT ROWS
```

```

03/17 16:14:59. P,1000000061,[S]INSERT ROWS
03/17 16:14:59. P,1000000061,[S]INSERT ROWS
03/17 16:14:59. ],1000000061,sp_iqdbspace,58
03/17 16:14:59. [,1000000061,sp_iqdbspace,60,select d.dbospace_name as
DBSpaceName, min(SegType) as DBSpaceType,...
03/17 16:15:00. ],1000000061,sp_iqdbspace,60
03/17 16:15:00. P,1000000061,Work[ Sort[ GrByH[ dbf<seq> JNL
ISYSIQDBSPACE<ISYSIQDBSPACE> JNL ISYSDBSPACE<ISYSDBSPACE> JH*
iq_dbospace_temp<seq> ] ] ] : ISYSIQPARTITIONCOLUMN<seq> :
idx<seq> : tab<seq>
03/17 16:15:00. [,1000000061,sp_iqdbspace,105,drop table
dbo.iq_dbospace_temp
03/17 16:15:00. ],1000000061,sp_iqdbspace,105
03/17 16:15:00. P,1000000061,[1]Work[ Sort[ sp_iqdbspace<call> ] ]

```

The connection handle in this example is 1000000061.

9.3.4 Remove .iq files before restoring [CR 620249]

Before restoring from a full backup, the store files (by default the *.iq* files) must be removed from the directory to which you are restoring. The sixth bullet in “Data Backup, Recovery, and Archiving > Restoring your databases > Before you restore” should be replaced with:

- For a full restore, the store files (by default the *.iq* files), the catalog store (by default the *.db* file), and the transaction log (by default the *.log* file) *must not* exist in the location to which you are restoring. If any of these files exist, you must delete them or move them to a different directory before doing the full restore.

9.3.5 Using DBA user to grant group membership [CR 607427]

In Granting group membership to users section, the second sentence currently reads: “Membership in a group can be granted either by the DBA or by the group user id.”

The second sentence should read as: “Membership in a group can be granted by the DBA.” and therefore, the first step of the example should read as: “Connect to the database as a user with DBA authority.”

9.3.6 Restrictions on forced recovery [CR 610313]

The first paragraph and note in “Starting servers in forced recovery mode,” in Chapter 13, “System Recovery and Database Repair,” in the *System Administration Guide: Volume 1* are incorrect. Please disregard the following:

“If a server fails to start with an exception or an assert when opening a database, start the server with forced recovery. Forced recovery allows the server to start if the allocation map is inconsistent. In this mode, options display information about inconsistencies. You can also specify options to repair such inconsistencies.

Use forced recovery only when normal database recovery fails to restore the database to a running state.”

The correct information is:

“Follow this procedure only if you see s_buf or free list errors during failure recovery. If SQL Anywhere recovery errors occur, -iqfrec will not correct the problem.”

See also “Do not use server startup switch -f [CR 610313]” on page 56.

9.3.7 Corrections for forced recovery [CR 617597]

The first paragraph and note in “Starting servers in forced recovery mode,” in Chapter 13, “System Recovery and Database Repair,” in the *System Administration Guide: Volume 1* are incorrect. The first paragraph should read:

“Forced recovery allows the server to start if the allocation map is inconsistent. In this mode, options display information about inconsistencies. You can also specify options to repair such inconsistencies.”

Please disregard this note:

“Use forced recovery only when normal database recovery fails to restore the database to a running state.”

In the section “Forced recovery mode,” please disregard the first paragraph.

See also “-fc server startup option [CR 617597]” on page 56.

9.4 System Administration Guide: Volume 2

This section contains updates to the *System Administration Guide: Volume 2*.

9.4.1 Defining trigger conditions for events [CR 624184]

In “Chapter 6: Automating Tasks Using Schedules and Events > Understanding events > Defining trigger conditions for events”, the Sybase IQ example should be:

```

create event SecurityCheck
type ConnectFailed
handler
begin

declare num_failures int;
declare mins int;

insert into FailedConnections( log_time )
values ( current timestamp );

select count( * ) into num_failures
from FailedConnections
where log_time >= dateadd( minute, -5,
current timestamp );
if( num_failures >= 3 ) then
select datediff( minute, last_notification,
current timestamp ) into mins
from Notification;

if( mins > 30 ) then
update Notification
set last_notification = current timestamp;
call xp_sendmail( recipient='DBAdmin',
subject='Security Check', "message"=
'over 3 failed connections in last 5 minutes' )
end if
end if
end

```

9.4.2 Ranking examples [CR 619931]

In “Chapter 2: Using OLAP > Ranking Functions > Ranking examples”, ranking example 1 should be:

```

SELECT Surname, Sex, Salary, RANK() OVER (
ORDER BY Salary DESC) as RANK FROM Employees
WHERE State IN ('CA') AND DepartmentID =200
ORDER BY Salary DESC;

```

The results from the above query:

Surname	Sex	Salary	RANK
-----	---	-----	----
Savarino	F	72300.000	1
Clark	F	45000.000	2
Overbey	M	39300.000	3

In “Chapter 2: Using OLAP > Ranking Functions > Ranking examples”, ranking example 2 should be:

```
SELECT Surname, Sex, Salary, RANK() OVER (PARTITION BY
Sex
ORDER BY Salary DESC) AS RANK FROM Employees
WHERE State IN ('CA', 'AZ') AND DepartmentID IN (200,
300)
ORDER BY Sex, Salary DESC;
```

The results from the above query:

Surname	Sex	Salary	RANK
-----	---	-----	----
Savarino	F	72300.000	1
Jordan	F	51432.000	2
Clark	F	45000.000	3
Coleman	M	42300.000	1
Overbey	M	39300.000	2

9.4.3 Permissions to create external names

In the section “Sybase IQ and remote data > Working with external logins > Creating external logins” in Chapter 4, “Accessing Remote Data,” the first sentence should change from “Only the login name and the DBA account can add or modify an external login.” to “Only the DBA account can add or modify an external login.”

9.4.4 Server support of TDS password encryption [CR 530767, CR 484702]

The Sybase IQ server supports TDS password encryption. In Chapter 3, “Sybase IQ as a Data Server,” these paragraphs should be added to the end of the section “Client/server interfaces to Sybase IQ > Open Client applications and Sybase IQ > Configuring Open Client”:

“When used as a remote server, Sybase IQ supports TDS password encryption. The Sybase IQ server accepts a connection with an encrypted password sent by the client. For information on connection properties to set for password encryption, see “Client-Library Topics > Security features > Adaptive Server Enterprise security features > Security handshaking: encrypted password” in the Open Server™ 15.5 *Open Client Client-Library/C Reference Manual*.

Note Password encryption requires Open Client 15.0. TDS password encryption requires Open Client 15.0 ESD #7 or later.

To enable the Sybase IQ server to accept a jConnect connection with an encrypted password, set the jConnect ENCRYPT_PASSWORD connection property to true.”

9.5 Reference: Building Blocks, Tables, and Procedures

This section contains updates to *Reference: Building Blocks, Tables, and Procedures*.

9.5.1 String functions result data type [CR 622463]

In the “String functions” section in Chapter 4, “SQL Functions,” in the fourth paragraph the data type of the result should be LONG VARCHAR, not VARCHAR:

“Unless you supply a constant LENGTH argument to a function that produces a LONG VARCHAR result (such as SPACE or REPEAT), the default length is the maximum allowed.”

9.5.2 Location of large object function descriptions [CR 621847]

The descriptions of these functions used only with large object data are located in Chapter 4, “Function Support” in *Large Objects Management in Sybase IQ*:

- BFILE
- BYTE_LENGTH64
- BYTE_SUBSTR64
- BYTE_SUBSTR
- CHAR_LENGTH64
- SUBSTRING64

9.5.3 RANK function syntax correction [CR 619931]

The syntax for the RANK function currently states:

```
RANK () OVER ( ORDER BY expression [ ASC | DESC ] )
```

The syntax should include the [PARTITION BY] clause:

```
RANK () OVER ( [ PARTITION BY ] ORDER BY expression [ ASC | DESC ] )
```

The example should be:

```
SELECT Surname, Sex, Salary, RANK() OVER (PARTITION BY
Sex
ORDER BY Salary DESC) AS RANK FROM Employees
WHERE State IN ('CA', 'AZ') AND DepartmentID IN (200,
300)
ORDER BY Sex, Salary DESC;
```

The results from the above query:

Surname	Sex	Salary	RANK
-----	---	-----	----
Savarino	F	72300.000	1
Jordan	F	51432.000	2
Clark	F	45000.000	3
Coleman	M	42300.000	1
Overbey	M	39300.000	2

9.5.4 Number of tables per join index correction [CR 617710]

Table 6-1: Sybase IQ database object size and number limitations currently contains this row:

Number of tables per join index (number of tables that can be joined in one query block) = 32.

The row should state:

Number of tables per join index (number of tables that can be joined in one query block) = 31.

9.5.5 FIRST_VALUE function description clarification [CR 607491]

In the Usage section, the paragraph describing *window-spec* currently states: “If the *window-spec* does not contain an ORDER BY then the result is arbitrary. If there is no *window-spec*, then the result is arbitrary.”

The paragraph should state: “If the *window-spec* does not contain an ORDER BY expression, or if the ORDER BY expression is not precise enough to guarantee a unique ordering, then the result is arbitrary. If there is no *window-spec*, then the result is arbitrary.”

9.5.6 LAST_VALUE function description clarification [CR 607491]

In the Usage section, the paragraph describing *window-spec* currently reads: “If the *window-spec* does not contain an ORDER BY then the result is arbitrary. If there is no *window-spec*, then the result is arbitrary.”

The paragraph should read: “If the *window-spec* does not contain an ORDER BY expression, or if the ORDER BY expression is not precise enough to guarantee a unique ordering, then the result is arbitrary. If there is no *window-spec*, then the result is arbitrary.”

9.5.7 Binary load format corrections [CR 578116, CR 578984]

The following section replaces the section “Sybase IQ binary load format” in Chapter 3, “SQL Data Types.”

Sybase IQ binary load format

Description

Sybase IQ uses the FORMAT BINARY and BINARY column specification clauses to produce data files that can be read by the LOAD TABLE statement.

To speed data loading into Sybase IQ, customers can create data files in Sybase IQ binary format and load this data into Sybase IQ using the FORMAT BINARY syntax of the LOAD TABLE command.

You can find instructions for creating a load script using the LOAD TABLE syntax and specifying the load specification in *Reference: Statements and Options*.

Create data files with these binary formats to load into columns with the corresponding data types. In most cases, Sybase IQ uses the platform-specific binary format. These data types are exceptions that use binary formats that are specific to Sybase IQ:

- DATE
- TIME
- DATETIME
- NUMERIC

IQ binary load format and load efficiency

The Sybase IQ binary load format is a fixed width format.

In general, fixed width loads complete faster than variable width loads. When the load logic knows the length of a column and row, the data is processed more efficiently. Using delimiters to separate columns and rows that vary in width causes the load to spend time scanning the input data looking for them.

The IQ Binary Load Format is a fixed width load. The load can determine the width of each column and length of each row from information in the table definition.

Note Binary load format is endian-sensitive. This is because binary load format utilizes native binary data types to represent data.

Operating system native data types

Data for the following data types is stored in native operating system binary format and can be written to data files directly in that format. Sybase IQ reads the respective number of bytes directly into the associated data types without conversion.

- BIT (1 byte)
- TINYINT (1 byte)
- SMALLINT (2 bytes)
- INT/UNSIGNED INT (4 bytes)
- BIGINT/UNSIGNED BIGINT (8 bytes)
- FLOAT (4 bytes)
- DOUBLE (8 bytes)
- CHAR/VARCHAR (character data)
- BINARY/VARBINARY (binary data)

By default, VARCHAR and VARBINARY columns are read in as many bytes as specified by LOAD TABLE *column-spec*.

DATE

DATE column data is stored in Sybase IQ as four bytes (a 32-bit unsigned integer) representing the number of days since 0000-01-01. To convert a calendar date to the Sybase IQ binary format, use:

For a given year, month, and day:

```

year = current_year - 1;
days_in_year_0000 = 366;
binaryDateValue = (year * 365)
+ (year / 4)
- (year / 100)
+ (year / 400)
+ days_in_year_0000
+ day_of_current_year
-1;

```

For the `day_of_current_year` value in the formula above, consider this example: February 12 is day 43.

TIME

TIME data is stored as a 64-bit unsigned quantity that represents a number in microseconds (in other words, 1.0e-6 seconds). The microsecond quantity is computed using:

For a given hour, minute, second, and microsecond (usec):

```

binaryTimeValue = (hour * 3600 + minute * 60 + second +
microsecond ) * 1000000

```

TIMESTAMP

TIMESTAMP data is stored as a 64-bit unsigned integer and represents a quantity in microseconds. You can compute a binary TIMESTAMP value using:

For a given year, month, day, hour, minute, second, and microsecond:

Compute *binaryDateValue* for the date as shown above.

Compute *binaryTimeValue* for the time as shown above.

```

binaryDateTimeValue = binaryDateValue *
86400000000 + binaryTimeValue

```

NUMERIC and DECIMAL

Formats for NUMERIC and DECIMAL data types vary as a function of precision. The value must be right-padded with zeroes to the full scale of the value. The value must also be fully left-padded with zeroes, but padding happens automatically with binary programming. Once the values are padded, the decimal point is removed. For example, the value 12.34 looks like:

- NUMERIC(4,2): 1234
- NUMERIC(6,4): 123400
- NUMERIC(8,4): 00123400
- NUMERIC(12,6): 000012340000
- NUMERIC(16,8): 0000001234000000

After the value is padded and the decimal point is removed, these rules apply:

- If precision ≤ 4 , then the binary format is identical to native operating system binary format for 2 byte integer quantity.
- If precision is between 5 and 9, then the binary format is identical to native operating system binary format for a 4 byte integer quantity.
- If precision is between 10 and 18, then the binary format is identical to native operating system binary format for an 8 byte integer quantity.
- If precision ≥ 19 , then there is a special format that uses this C++ struct definition:

```
struct {
    unsigned char sign; // sign 1 for +, 0 for -
    unsigned char ndig; // # digits
    unsigned char exp; // exponent
    unsigned char erracc; // should be 0
    unsigned short digits[80];
};
```

Exponent is excess-80 form, unless the value is zero. A “zero” value is represented as:

```
sign = 1
ndig = 0
erracc = 0
exp = 0
```

The maximum exponent value is 159. The maximum number of supported digits is 288. “digits[0]” contains the least-significant digits. Digits are stored in a packed representation with two digits per “unsigned short” (2-byte) quantity. For a given “digit:”

```
lower order digit = digit[i] & 0x00FF
high order digit = digit[i] & 0xFF00
```

For example, consider the value 100 represented as NUMERIC(20). The binary layout of this value is:

```
0x0101 0x5000 0x0064 0x0000 0x0000 .....
```

```
Sign = 0x01
Number digits = 0x01
Exponent = 0x50
Erracc = 0x00
Digits = 0x0064
```

As another example, consider the value 32769:

```
0x0102 0x5000 0x0ad1 0x0003 0x0000 0x0000 ....
```

```
Sign = 0x01
Number digits = 0x02
Exponent = 0x50
Erracc = 0x00
Digits = 0x0ad1 0x0003
```

If you translate the digits into base 10, you have:

```
0x0ad1 = 2769 0x0003 = 3
```

Inserting NULL

The most expedient way to insert NULL values is to use the NULL BYTE in the input file and specify WITH NULL BYTE in the column specification of the LOAD TABLE statement. This is done by terminating each data field in the input file with 'x00' or 'x01'. Terminating a data field in the input file with 'x01' instructs the load to insert NULL into the column. For example:

```
create table d1 ( c1 date );
load table d1 ( c1 binary with null byte ) from
'filename' quotes off escapes off format binary;
```

If the content of the load input file is 000b32cb00000b32cc00, two rows will be loaded to the table. The first row will be May 7, 2009 and the second May 8, 2009. Notice that a NULL BYTE was added to the input file after each binary date. If you want NULL loaded into the first row, change the value of the NULL BYTE in the input file to 'x01'.

```
000b32cb01000b32cc00
```

As another example, to load the value 32769 into a NUMERIC(20) column, the input file contains:

```
0x0102 0x5000 0x0ad1 0x0003 0x0000 0x00
```

This includes the NULL BYTE.

To load 23456789012345678.12 into a column defined as NUMERIC(19,2), the load input file contains

```
0x0106 0x4f00 0x04b0 0x162e 0x04d2 0x1ed2 0x0d80 0x0002
0x0000 0x00
```

The digits are followed by the NULL BYTE (0x00).

Note that there are seven (numbered 0-6) unsigned shorts in the digits array of the structure that represents this numeric quantity. "digits[0]" contains the least-significant digits.

```

digits[0] = 0x04b0 (decimal 120)
digits[1] = 0x162e (decimal 5678)
digits[2] = 0x04d2 (decimal 1234)
digits[3] = 0x1ed2 (decimal 7890)
digits[4] = 0x0d80 (decimal 3456)
digits[5] = 0x0002 (decimal 2)
digits[6] = 0x0000

```

The NULL portion of the column specification indicates how to treat certain input values as NULL values, when loading into the table column. These characters can include BLANKS, ZEROS, or any other list of literals you define. When you specify a NULL value or read a NULL value from the source file, the destination column must be able to contain NULLs.

ZEROS is interpreted as follows:

- The column is set to NULL if the input data is all binary zeros (not character zeros).
- If the input data is character zero:
 - NULL(ZEROS) never causes the column to be NULL.
 - NULL('0') causes the column to be NULL. For example:

Load:

```
CREATE TABLE t1 ( c1 INT, c2 INT );
```

View the input data file, which uses big-endian byte ordering:

```
od -x data.inp
3030 3030 0000 04d2
```

Execute:

```
LOAD TABLE t1 ( c1 ASCII(4) NULL( '0000' ),
                c2 BINARY )
FROM 'data.inp'
FORMAT BINARY
QUOTES OFF
ESCAPES OFF;
```

The results:

```
SELECT * FROM t1;
c1      c2
NULL    1234
```

- If the input data is binary zero (all bits clear):
 - NULL(ZEROS) causes the column to be NULL.

- NULL('0') never causes the column to be NULL, for example:

Load:

```
CREATE TABLE t1 ( c1 INT, C2 INT );
```

VIEW the input data file, which uses big-endian byte ordering:

```
od -x data.inp
0000 0000 0000 04d2
```

Execute:

```
LOAD TABLE t1 ( c1 ASCII(4) NULL( zeros ),
                c2 BINARY )
FROM 'data.inp'
  FORMAT BINARY
  QUOTES OFF
  ESCAPES OFF;
```

The results:

```
SELECT * FROM T1;
c1      c2
NULL    1234
```

As another example, if your LOAD TABLE statement includes `col1 date('yyymmdd') null(zeros)` and the data to load is 000000, you receive an error indicating that 000000 cannot be converted to a DATE(4). To get LOAD TABLE to insert a NULL value in col1 when the data is 000000, either write the NULL clause as `null('000000')`, or modify the data to equal binary zeros and use `NULL(ZEROS)`.

Another way to load NULLs during a binary load is not to supply data for the column in the LOAD TABLE statement, if the destination column accepts null values. For example:

```
CREATE TABLE t1 ( c1 INT, c2 INT );
LOAD TABLE T1 ( c2 BINARY ) FROM 'data.inp'
  FORMAT BINARY
  QUOTES OFF
  ESCAPES OFF;

SELECT * FROM T1;
c1      c2
NULL    1234
NULL    1234
```

View the input data file, which uses big-endian byte ordering:

```
od -x data.inp
0000 04d2 0000 04d2
```

9.6 Reference: Statements and Options

This section contains updates to *Reference: Statements and Options*.

9.6.1 SUBQUERY_FLATTENING_PREFERENCE allowed values [CR 622967]

The actions of the allowed values for the SUBQUERY_FLATTENING_PREFERENCE database option in Chapter 2, “Database Options” are shown here in tabular format. This information is correct in *Reference: Statements and Options*.

Value	Action
-3	Avoid flattening both EXISTS and scalar subqueries to a join operation.
-2	Avoid flattening a scalar subquery to a join operation.
-1	Avoid flattening an EXISTS subquery to a join operation.
0	Allow the IQ optimizer to decide to flatten subqueries.
1	Ignore cost flattening EXIST, if possible.
2	Ignore cost flattening scalar, if possible.
3	Ignore cost of both EXISTS and scalar subquery.

9.6.2 @@rowcount global variable supported by Sybase IQ [CR 621197]

The section “SQL Statements > DELETE statement” currently states: “The Transact-SQL® ROWCOUNT option has no effect on DELETE operations in Sybase IQ.” This sentence should be removed.

The section “SQL Statements > UPDATE statement” currently states: “The Transact-SQL ROWCOUNT option has no effect on UPDATE operations in Sybase IQ.” This sentence should be removed.

The section “SQL Statements > SET statement [T-SQL]” currently contains this note:

Note The ROWCOUNT option has no effect on UPDATE and DELETE operations in Sybase IQ. Also note that Sybase IQ does not support the @@rowcount global variable.

This note should be changed to:

Note Sybase IQ supports the @@rowcount global variable. The SELECT, INSERT, DELETE, and UPDATE statements affect the ROWCOUNT option. The ROWCOUNT option has no effect on cursor operation, the IF statement, or creating/dropping a table or procedure.

9.6.3 SUBQUERY_CACHING_PREFERENCE option default [CR 622968]

The default value (0) for the SUBQUERY_CACHING_PREFERENCE option is not described. When you set the option to 0 (the default), the Sybase IQ optimizer chooses the appropriate caching algorithm automatically.

9.6.4 CREATE EVENT statement [CR 617597]

An error appears in the CREATE EVENT statement documentation for DiskSpace event types. The statement that the disk space event types require Windows and are not available on UNIX platforms is incorrect.

The CREATE EVENT statement includes this new example:

This statement posts a message to the server log when free disk space on the device containing the transaction log file falls below 30 percent, but executes the handler no more than once every 300 seconds.

```
CREATE EVENT LowTxnLogDiskSpace
TYPE DBDiskSpace
WHERE event_condition( 'DBFreePercent' ) < 30
AND event_condition( 'Interval' ) >= 300
HANDLER
BEGIN
message 'Disk space for Transaction Log is low.';
END;
```

9.6.5 Maximum value for LOAD TABLE SKIP clause [CR 615402]

In “SQL statements > LOAD TABLE statement,” the description of the SKIP clause in the “Usage” section should be updated to include the maximum number of rows to skip:

SKIP Defines the number of rows to skip at the beginning of the input tables for this load. The maximum number of rows to skip is $2^{31} - 1$ (2147483647). The default is 0.

9.6.6 Using the SET OPTION statement for a user group [CR 612762]

In the “Usage” section of “SQL statements > SET OPTION statement,” the third paragraph currently reads: “Specifying either a user ID or the PUBLIC user ID determines whether the option is set for an individual user, a user group represented by *userid*, or the PUBLIC user ID (the user group to which all users are a member). If no user group is specified, the option change is applied to the currently logged-on user ID that issued the SET OPTION statement.”

The paragraph should read: “Specifying either a user ID or the PUBLIC user ID determines whether the option is set for an individual user, a user group represented by *userid*, or the PUBLIC user ID (the user group to which all users are a member). If the option applies to group user ID, option settings are not inherited by members of the group—the change is applied only to the group user ID. If no user group is specified, the option change is applied to the currently logged-on user ID that issued the SET OPTION statement.”

9.6.7 Permission conditions for the GRANT statement [CR 607427]

In the “Permissions” section of “SQL Statements > GRANT statement,” the three conditions for Syntax 1 or 2 should change to the following two conditions and one of these conditions must be met:

- You are changing your own password using GRANT CONNECT
- You have DBA authority

9.6.8 Non-DBA users must use a fully-qualified name to drop an index [CR 548184]

For DROP INDEX, non-DBA users must provide a fully-qualified index name to drop an index on a base table owned by the DBA. DBA or users with the appropriate privileges can drop an index on tables that are owned by non-DBA users without using a fully-qualified name.

9.6.9 Server support of TDS password encryption [CR 530767, CR 484702]

The Sybase IQ server supports TDS password encryption. This statement should be removed from the “Usage” section of “SQL Statements > INSERT statement”:

“When used as a remote server, Sybase IQ does not support this password encryption.”

and replaced with these paragraphs:

“When used as a remote server, Sybase IQ supports TDS password encryption. The Sybase IQ server accepts a connection with an encrypted password sent by the client. For information on connection properties to set for password encryption, see “Client-Library Topics > Security features > Adaptive Server Enterprise security features > Security handshaking: encrypted password” in the Open Server 15.5 *Open Client Client-Library/C Reference Manual*.

Note Password encryption requires Open Client 15.0. TDS password encryption requires Open Client 15.0 ESD #7 or later.

To enable the Sybase IQ server to accept a jConnect connection with an encrypted password, set the jConnect ENCRYPT_PASSWORD connection property to true.”

9.6.10 Using MAX_QUERY_PARALLELISM option [CR 491483]

In the Allowed Values section, the sentence should be: “Integer less than, greater than, or equal to number of CPUs.” When the MAX_QUERY_PARALLELISM option is set, the number of parallel arms in a UNION, GROUP BY can be greater than the number of CPUs.

9.7 Using Sybase IQ Multiplex

This section contains updates to the *Using Sybase IQ Multiplex*.

9.7.1 Restoring IQ store data in a different location [CR 623585]

The procedure “Restoring IQ store data when restoring database in different location” in “Backing Up and Restoring Data in a Multiplex Environment > Backing up and restoring data in a multiplex environment > Restoring in the multiplex environment” contains errors in steps 5, 6, and 7. It should read:

- ❖ **Restoring IQ store data when restoring database in different location**
 - 1 Confirm that database home directories for each server still exist. If not, create them or restore them from file system backups.

- 2 If this is not the first time you have restored to the new location, shut down all multiplex servers running at the destination location (coordinator and secondary servers). In Sybase Central, right-click the multiplex and use the Stop Server wizard. The wizard lets you stop one or more servers. The multiplex at the original location where the backup was taken may continue running.

Note If automatic startup is enabled in your ODBC configuration, users on the same machine as the server may be set up to automatically start the server. Prevent this from happening while you are restoring the database.

- 3 After stopping all servers, confirm that the database shut down successfully. If you see an active `iqsrv15` process with the name of a multiplex server, stop the process.

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

```
% ps -ef | grep iqsrv15
fiona 434      1  1   May 19 ??  0:05 start_iq
-n myhost_myserver
-c 32MB -x tcpip(port=1234) mpxdb.db
fiona 4751    442  1 16:42:14 pts/5
0:00 grep start_iq
```

To verify on a Windows system, use Task Manager. Look on the Processes tab for `iqsrv15.exe` or find the IQ Server icon in the system tray and stop it using right-click and Shutdown.

- 4 Move files required for debugging and reconfiguring the multiplex.
 - Make a file system copy of the `.iqmsg` file. If you have message log archiving configured, see “Backing up the right files,” in *System Administration Guide: Volume 1*.
- 5 Shut down the original coordinator, then start the utility database from the coordinator server directory using the coordinator server’s name:

```
% start_iq -n coordinator_svr -c 32MB
-x tcpip(port=1234)
```

- 6 Connect to the utility database (`utility_db`).

```
% dbisql -c "eng=coordinator_svr;uid=DBA;pwd=SQL;
dbn=utility_db" -host myhost -port 1234
```

- 7 Run the RESTORE command with a new location path for the database and all of its dbspaces. To restore certain dbspace files to a different path, specify a RENAME clause. For details, see the RESTORE statement in *Reference: Statements and Options*.
- 8 Start the restored database either by reconnecting to the utility_db server and specifying the restored database file name or by stopping the server and restarting it with the restored database. If you restart the server, use the single-node and override flags (-iqmpx_sn 1 -iqmpx_ov 1).
- 9 Use DROP MULTIPLEX SERVER statement to drop all the secondary nodes. For example, DROP MULTIPLEX SERVER node_w3_skm.

Once you drop the last secondary node, the coordinator shuts down automatically, signifying conversion to simplex.
- 10 Restart the coordinator without the single node or override switch.
- 11 Recreate all the secondary nodes with the correct location path. For example:

```
CREATE MULTIPLEX SERVER node_r2_skm DATABASE
'/sunx3005/mpx_simdb.db'
HOST 'localhost' PORT 8998
ROLE READER STATUS INCLUDED
```

After you create the first secondary node, the server automatically shuts down, signifying conversion to multiplex.

Note In the preceding command, you must include the database file extension *.DB*.

- 12 When you restart the coordinator, you see a warning in the server log about the multiplex environment being invalid. This warning is generated if IQ_SYSTEM_TEMP dbspace does not contain any files, and is the case for all the secondary nodes you created in step 12. Ignore this warning for now.
- 13 Synchronize the secondary servers, following the instructions in “Synchronizing secondary servers” on page 29.
- 14 Start the secondary servers.
- 15 Connect to each secondary server and add files in IQ_SYSTEM_TEMP.
- 16 Run sp_iqmpxvalidate on the coordinator. It should report no error detected.

To restore an exact copy of the multiplex to a different location, when copies of all of the server's temporary files exist at the new location, you replace steps 9 through 17 with:

Use ALTER MULTIPLEX SERVER to alter the server name, host, port, and database path of each server.

9.7.2 LOCK TABLE unsupported on write servers [CR 555776]

Multiplex write servers do not support the LOCK TABLE command. Sybase IQ returns this message:

```
Could not execute statement.
Permission denied: Command not allowed on
Multiplex Writer servers.
(saint_iqthresholdddl.cxx 14868)
SQLCODE=-1004015, ODBC 3 State="HY000"
Line 1, column 1
```

9.7.3 Output changes when sp_iqmpxinheartbeatinfo exceeds 24-hour limit [CR 607408]

If the elapsed time exceeds 24 hours, Sybase IQ returns sp_iqmpxinheartbeatinfo output like:

```
server_id,server_name,last_positive_hb,
time_not_responding,time_until_timeout
2, 'r2_mpx_cr_srv',Jan 14 2010 11:57AM,11:59PM,11:59PM
3, 'w4_mpx_cr_srv',Jan 14 2010 11:57AM,11:59PM,11:59PM
(2 rows affected)
(return status = 0)
```

A value of 11:59PM in the time_not_responding and time_until_timeout columns means that the time has crossed the 24-hour limit.

9.8 Installation and Configuration Guide

For updates about installation and migration, including database upgrades, see “Special installation and migration instructions” on page 2.

9.9 User-Defined Functions Guide

This section contains updates to the *User-Defined Functions Guide*.

9.9.1 xIC compiler version for UDF functions on AIX [CR 609079]

Add this note to the section “Creating and executing user-defined functions > Compile and link switches for building dynamically linkable libraries > AIX switches:”

Note To compile on AIX 6.1 systems, the minimum level of the xIC compiler is 8.0.0.24.

9.9.2 Directory name \x64 changed to \bin64 for creating and executing user-defined functions [CR 619893]

In the section “Creating and executing user-defined functions > Compile and link switches for building dynamically linkable libraries > Windows switches,” the directory name \x64 is changed to \bin64 and Example 3 now reads as:

```
%VCBASE%\VC\bin\amd64\link
/LIBPATH:%VCBASE%\VC\lib\amd64
/LIBPATH:%MSSDK%\lib\bin64 kernel32.lib -manifest -DLL
-nologo
-MAP:"%OBJ_DIR%\libudfex.map_deco"
/OUT:"%OBJ_DIR%\libudfex.dll"
"%OBJ_DIR%\my_interpolate.o" "%OBJ_DIR%\my_main.o"
/DLL
-EXPORT:extfn_use_new_api -EXPORT:my_interpolate
```

9.9.3 Calling pattern for OLAP-style optimized moving window aggregate [CR 618055]

In the “OLAP-style optimized moving window aggregate” section, the calling pattern is currently documented as:

```

_start_extfn(cntxt)
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args) -- input a=1
_evaluate_extfn(cntxt, args) -- returns 1
_next_value_extfn(cntxt, args) -- input a=2
_evaluate_extfn(cntxt, args) -- returns 3
_drop_value_extfn(cntxt) -- input a=1
_next_value_extfn(cntxt, args) -- input a=3
_evaluate_extfn(cntxt, args) -- returns 5
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args) -- input a=4
_evaluate_extfn(cntxt, args) -- returns 4
_next_value_extfn(cntxt, args) -- input a=5
_evaluate_extfn(cntxt, args) -- returns 9
_drop_value_extfn(cntxt) -- input a=4
_next_value_extfn(cntxt, args) -- input a=6
_evaluate_extfn(cntxt, args) -- returns 11
_finish_extfn(cntxt)

```

The calling pattern would be called as:

```

_start_extfn(cntxt)
_reset_extfn(cntxt)
_evaluate_aggregate_extfn(cntxt, args) -- returns 1
_evaluate_aggregate_extfn(cntxt, args) -- returns 3
_drop_value_extfn(cntxt) -- input a=1
_next_value_extfn(cntxt, args) -- input a=3
_evaluate_aggregate_extfn(cntxt, args) -- returns 5
_reset_extfn(cntxt)
_next_value_extfn(cntxt, args) -- input a=4
_evaluate_aggregate_extfn(cntxt, args) -- returns 4
_next_value_extfn(cntxt, args) -- input a=5
_evaluate_aggregate_extfn(cntxt, args) -- returns 9
_drop_value_extfn(cntxt) -- input a=4
_next_value_extfn(cntxt, args) -- input a=6
_evaluate_aggregate_extfn(cntxt, args) -- returns 11
_finish_extfn(cntxt)

```

9.10 Utility Guide

This section contains updates to the *Utility Guide*.

9.10.1 Do not use server startup switch -f [CR 610313]

Table 1-1 in Chapter 1, “Running the Database Server,” contains an error. Never use `start_iq` with startup switch `-f` for IQ databases. The switch will cause database corruption, because the catalog database is not synchronized with the IQ store if the server did not shut down cleanly. Documentation for this switch will be removed from future releases. See also “Restrictions on forced recovery [CR 610313]” on page 35.

9.10.2 -fc server startup option [CR 617597]

In Chapter 1, “Running the database server”, the server startup option `-fc` is listed as a valid startup parameter. Do not use the `-fc` server startup option.

9.11 Performance and Tuning Guide

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

9.11.1 The transaction log file [CR 617597, 592023]

In Chapter 4, “Managing System Resources,” the section “The transaction log file” should read:

The transaction log file contains information that allows Sybase IQ to recover from a system failure. The transaction log is also required for auditing. The default file name extension for this file is `.log`.

To move or rename the transaction log file, use the Transaction Log utility (`dblog`). See “Transaction Log utility (`dblog`)” in Chapter 4, “Database Administration Utilities,” in the *Utility Guide*.

Warning! The Sybase IQ transaction log file is different from most relational database transaction log files. If for some reason you lose your database files, then you lose your database (unless it is the log file that is lost). However, if you have an appropriate backup, then you can reload the database.

Truncating the transaction log

Sybase IQ records in the transaction log the information necessary to recover from a system failure. Although the information logged is small for each committed transaction, the transaction log continues to grow in size. In systems with a high number of transactions that change data, over a period of time the log can grow to be very large.

When to truncate the log is really up to the DBA responsible for supporting the Sybase IQ systems, and depends on the growth profile of the log file and the operational procedures at the site.

Table 4 shows methods for truncating transaction logs in Sybase IQ.

Table 4: Truncating transaction logs

If your database is ...	Use this method ...	For details, see ...
Stopped	“The -m switch, which causes the transaction log to be truncated after each checkpoint for all databases	“Truncating the transaction log of a stopped database” on page 57
Running	The dbbackup command line utility with the -xo switch or the -r switch	<i>Utility Guide</i> > Database Administration Utilities > Backup utility (dbbackup)

Truncating the transaction log of a stopped database

Use the -m server startup switch to truncate the transaction log of a database. Note that leaving the -m server startup switch permanently set is not recommended. This switch should only be used to start Sybase IQ for a transaction log truncation. How this is done is up to the DBA, but this procedure provides a suggestion.

❖ **Truncating the transaction log of a stopped database**

- 1 Create a copy of the server switches .cfg file with a name identifying the file as the log truncation configuration setting and edit this copy of the file to add the -m switch.
- 2 Start Sybase IQ with the configuration file containing the -m option. Note that no user access or transactions should be allowed at this time.
- 3 Shut down Sybase IQ and restart using the configuration file without the -m option set.

9.11.2 UNION ALL views for faster loads [CR 619105]

In the Chapter 4, “Managing System Resources” section “Using UNION ALL views for faster load,” a new paragraph should be inserted after the first paragraph, so the section begins:

“UNION ALL views can be used to improve load performance where it is too expensive to maintain secondary indexes for all rows in a table. Sybase IQ lets you split the data into several separate base tables (for example, by date). You load data into these smaller tables. You then join the tables back together into a logical whole by means of a UNION ALL view, which you can then query.

This strategy can improve load performance, but may negatively impact the performance of some types of queries. Most types of queries have roughly similar performance against a single base table or against a UNION ALL view over smaller base tables, as long as the view definition satisfies all the constraints described in “Optimizing queries that reference UNION ALL views.” However, some types of queries, especially those involving DISTINCT or involving joins with multiple join columns, may perform significantly slower against a UNION ALL view than against a single large base table. Before choosing to use this strategy, determine whether the improvements in load performance are worth the degradation in query performance for your application.”

10. Technical Support

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you cannot resolve a problem using the manuals or online help, please have the designated person contact Sybase Technical Support or the Sybase subsidiary in your area.

Before you contact
Technical Support

Technical Support needs information about your environment 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 collect some information manually. Providing this information helps expedite the resolution of your problem.

In this list, * indicates items collected by `getiqinfo`:

- Sybase IQ version (for example, 15.1 GA or ESD level)
- Type of hardware, amount of memory, number of CPUs*
- Operating system and version (for example, Microsoft Windows 2008 Service Pack 1)*
- Operating system patch level
- Front-end tool used (for example, Business Objects Crystal Reports)

- Connectivity protocol used (for example, ODBC, JDBC, Tabular Data Stream™ (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 (if applicable)
Named *stktrc-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 type these commands and then rerun the command that produced the error:

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

The plan will be in the message log file.

If you have performance problems, set this option:

```
SET TEMPORARY OPTION Query_Plan_After_Run = 'ON'
```

This enables Technical Support to see which steps in the query processing used the time.

-
- Server logs
 - For UNIX and Linux: *IQ-15_1/logfiles/<servername>.nnnn.stderr* and *IQ-15_1/logfiles/<servername>.nnnn.srvlog**
 - On Windows platforms:
*%ALLUSERSPROFILE%\SybaseIQ\logfiles\
<servername>.nnnn.srvlog**

For example,
on Windows 2003 the server log files are located in
C:\Documents and Settings\All Users\SybaseIQ\logfiles
and
on Windows 2008 the server log files are located in
C:\ProgramData\SybaseIQ\logfiles

- Startup 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 each node (coordinator and secondary nodes)
- Screen snapshot of the problem, if possible

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 Chapter 14, “Troubleshooting Hints” in the *System Administration Guide: Volume 1*.

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, the 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.

11. Other sources of information

Use the Sybase Getting Started CD, the SyBooks CD, and the Sybase Product Manuals Web site to learn more about your product:

- The Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the SyBooks CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader, which you can download at no charge from the Adobe Web site using a link provided on the CD.
- The SyBooks CD contains product manuals and is included with your software. The Eclipse-based SyBooks browser allows you to access the manuals in an easy-to-use, HTML-based format.

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

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

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://sybooks.sybase.com>.

11.1 Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ Finding the latest information on product certifications

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click Partner Certification Report.
- 3 In the Partner Certification Report filter select a product, platform, and timeframe and then click Go.
- 4 Click a Partner 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.

11.2 Sybase EBFs and software maintenance

❖ Finding the latest information on EBFs and software maintenance

- 1 Point your Web browser to the Sybase Support Page at <http://www.sybase.com/support>.
- 2 Select EBFs/Maintenance. If prompted, type 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.

11.3 Sybase IQ 15.1 Network Clients and ODBC kits available

Sybase IQ 15.1 Network Clients and ODBC kits for development are available on the Sybase download site at <http://www.sybase.com/downloads>. A Network Client is available for each platform supported by Sybase IQ 15.1.

11.4 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 for guidelines on posting, go to Newsgroups at <http://www.sybase.com/support/newsgroups>.

11.5 IQ User's Group

The IQ User's Group provides a forum for any IQ users to exchange information about Sybase IQ. You can visit the IQUG homepage at <http://iqug.dssolutions.com/>. To have your name added to the IQ User's Group list, send e-mail to iqug-subscribe@dssolutions.com.

12. 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 15.1 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 Sybase Web site Product Manuals at <http://sybooks.sybase.com>.

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 the online help. 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.

For a Section 508 compliance statement for Sybase IQ, go to Sybase Accessibility at <http://www.sybase.com/products/accessibility>.

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 for information on the most useful configuration for your use.

13. Checklist: information for Technical Support

You can run the `getiqinfo` script to collect much of this information.

Information requested	Value
Sybase IQ version (for example 15.1 GA or ESD number)	
<code>sp_iqlmconfig</code> output	
type of hardware	
amount of memory	
number of CPUs	
operating system name and version (for example, Microsoft Windows 2008 Service Pack 1)	
operating system patch level	
front end tool used (for example, Business Objects Crystal Reports)	
connectivity protocol used (for example, ODBC, JDBC, TDS)	
Open Client version	
configuration type (single node or multiplex)	
message log file (<code>dbname.iqmsg</code>)	
server log files (<code>server.mnnn.srvlog</code> and <code>server.mnnn.stderr</code>)	
stack trace file (<code>stktrc-YYYYMMDD-HHNNSS_#.iq</code>)	
command or query that produced the error	
startup option settings	
connect option settings	
database option settings	
schema and indexes for the database	
<code>sp_iqstatus</code> output	

Information requested	Value
query plan: set options (Query_Plan, Query_Detail, Query_Plan_After_Run, Query_Plan_As_Html, Query_Plan_As_Html_Directory, Query_Timing), rerun command or query	
Screen snapshot of the problem, if possible	