

Release Bulletin Sybase® IQ 15.2 for Linux

Document ID: DC10082-01-1520-02

Last revised: December 2010

Topic	Page
1. Accessing current release bulletin information	1
2. Product summary	2
3. Special installation and migration instructions	2
4. New features and behavior changes	7
5. Product compatibilities	14
6. Restrictions	14
7. Known problems	17
8. Documentation for this version	32
9. Documentation updates and clarifications	34
10. Technical Support	53
11. Other sources of information	55
12. Accessibility features	58
13. Checklist: information for Technical Support	60

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 Documentation Web site.

❖ Accessing release bulletins at the Sybase Product Documentation Web site

- 1 Go to Product Documentation 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 15.2 for Linux.

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

Refer to *New Features Summary Sybase IQ 15.2* for descriptions of new features and behavior changes in Sybase IQ 15.2.

3. Special installation and migration instructions

For complete installation and migration instructions, see the *Installation and Configuration Guide for Linux*. Read the chapters “Preparing for Installation” and “Migrating Data” before installing Sybase IQ 15.2.

Before you run a new version of Sybase IQ, see “Restrictions” on page 14 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 57.

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 graphical installation wizard. The interactive installer supports all product editions and installation types. A new silent installation option is also available. See “Installing your software” in Chapter 1, “Preparing for Installation” in the *Installation and Configuration Guide for Linux*.

3.2 Sybase IQ 15.2 requires a license

Sybase IQ 15.2 uses the Sybase Software Asset Management (SySAM) licensing mechanism that lets system administrators 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. See Chapter 2, "Licensing your software" in the *Installation and Configuration Guide for Linux*.

3.3 Sub-capacity licensing

Sybase now offers sub-capacity licensing options for Sybase IQ Enterprise Edition. Sub-capacity licensing refers to licensing a Sybase product on a subset of the CPUs available on a physical machine.

3.3.1 Platform support

Sub-capacity licensing is supported on the following platforms:

Table 1: Sub-capacity licensing vendor support

Vendor	Product	Platform support	Virtualization Type
HP	nPar	HP IA 11.31	Physical Partition
	vPar		Virtual Partition
	Integrity Virtual Machines with Resource Manager		Virtual Machine
	Secure Resource Partitions		OS Containers
IBM	LPAR	AIX 6.1	Virtual Partition
	dLPAR		Virtual Partition
SUN	Dynamic System Domains	Solaris 10	Physical Partition
	Solaris Containers/Zones with Solaris Resource Manager		OS Partition
INTEL/ AMD	VMWare ESX Server ¹ Guest OS: Windows	VMWare ESX Server 3.5 Guest OS: Windows 2003	Virtual Machine
	VMWare ESX Server Guest OS: Linux	VMWare ESX Server 3.5 Guest OS: RH 5.3, SuSE 10	Virtual Machine
	Xen ² , DomainU: Windows	Windows 2003	Virtual Machine
	Xen, DomainU: Linux	RH 5.3, SuSE 10	Virtual Machine

¹ VMWare excludes VMWare Workstation and VMWare Server

² Xen excludes Sun Solaris x64

3.3.2 How to enable Sybase sub-capacity licensing

You must have a sub-capacity licensing agreement with Sybase before you can enable sub-capacity licensing. If you choose to use Sybase IQ in a sub-capacity environment, see the *SySAM Quick Start Guide* for instructions for generating your license key.

Keep your license server up-to-date

Although the installation media includes a copy of the most recent SySAM license server, Sybase recommends that you periodically check for license server updates on the SySAM Standalone License Server Install site at <http://www.sybase.com/sysam/server>.

3.4 SySAM Licensing checkout [CR #628594]

If Sybase IQ was licensed using a per-processor license type in previous releases, the license quantity was determined at startup. In this release, Sybase IQ periodically checks the number of processors it can use and attempts to check-out additional licenses for any increase. Sybase IQ shuts down after 30 days, if additional licenses are not made available within this period. See the *SySAM Users Guide* for details regarding license grace.

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 earlier versions. You can use the 15.x plug-in only 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 Using remote server classes for ASE [CR #615420]

Component Integration Services (CIS) connectivity from Sybase IQ to Adaptive Server® Enterprise (ASE) is available using server classes asejdbc and aseodbc. The asejdbc server class uses the Adaptive Server JDBC™ driver and the aseodbc server class uses the Adaptive Server ODBC driver. CIS connectivity through aseodbc server class is expected to provide better performance than the asejdbc class. The Adaptive Server ODBC driver requires a separate installation of SDK 15.5 ESD #5 which can be downloaded from the Sybase EBF download site. SDK 15.5 ESD #5 is not available for download at the time of the Sybase IQ 15.2 release and is expected to be available shortly afterwards. Contact Sybase Technical Support for details on release dates.

To use the aseodbc server class, you must also install:

- IQ 15.2 EBF17907 for Linux x64
- IQ 15.2 EBF17908 for Linux IBM64

Contact Sybase Technical Support to obtain the appropriate EBF.

3.8 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 SDK is available as a separately purchased option.

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

3.9 Open Client no longer included

Open Client™ is no longer installed as part of the Sybase IQ product. Sybase IQ 15.2 does provide versions of Open Client utilities that have limited functionality to enable INSERT...LOCATION:

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

3.10 Increase IQ temporary cache setting for loads

The `LOAD_MEMORY_MB` option has been deprecated. Instead of setting `LOAD_MEMORY_MB`, you must now adjust the IQ temporary cache setting. The amount to increase the IQ temporary cache is approximately the sum of all the `LOAD_MEMORY_MB` settings of all of the concurrent `LOAD TABLE` executions.

3.11 Creating demo database fails with sort error [CR #631097]

Creating the demo database may fail with an Informational stack trace for `SORT` and the error:

```
Exception Thrown from s_csort.cxx:1496, Err# 2,
tid 3676 origtid 3676
O/S Err#: 0, ErrID: 2062 (hs_csortexception);
SQLCode: -1009134, SQLState: 'QSB34', Severity: 14
[21023]: Insufficient buffers for 'Sort'.
-- (s_csort.cxx 1496)
```

The “insufficient buffers” error may be related to the small default Sybase IQ memory caches (`-iqmc 32` and `-iqtc 24`) and a high number of processors on the machine.

Workaround:

If hyperthreading is on, set the `-iqnumbercpus` server startup option to the number of physical processors. This may avoid the “insufficient buffers” error, if `-iqnumbercpus` is 64 or less.

If hyperthreading is off and the number of processors is greater than 64, edit `mkiqdemo.sh` to increase the size of the memory cache. For example, on a machine with 128 processors, change `-iqmc 32` and `-iqtc 24` to `-iqmc 100` and `-iqtc 100`.

See *System Administration Guide: Volume 1 > Troubleshooting hints > Solutions for specific conditions > Performance issues > Slow performance on a multi-CPU or hyperthreaded machine.*

3.12 Drop and recreate TEXT indexes after ESD installation [CR #620317]

After installing a Sybase IQ 15.2 ESD and upgrading the database, drop and recreate all Sybase IQ TEXT indexes created in Sybase IQ 15.2 GA. See “Drop and recreate TEXT indexes after ESD upgrade or rollback [CR #620317]” on page 18.

4. New features and behavior changes

For new features and behavior changes, see *New Features Summary Sybase IQ 15.2*, which Sybase recommends you read before you install Sybase IQ 15.2. 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 15.2 ESD #2

Sybase IQ 15.2 ESD #2 includes new features in the areas of load and TEXT index searching.

- The extended syntax HEADER SKIP allows LOAD TABLE to skip header rows at the beginning of a data file. See “LOAD TABLE HEADER SKIP syntax [CR #593620]” on page 8.
- The extended full text search capability allows fuzzy searching on a TEXT index with an n-gram term breaker. See “NGRAM TEXT index and fuzzy and non-fuzzy search” on page 10.

4.2 Data Definition Language (DDL) changes

This section contains new features and changes related to DDL.

4.2.1 LONG BINARY now maps to MS SQL VARBINARY [CR #636192]

The mapping of the columns of a proxy table to a Microsoft SQL Server table changed as of Sybase IQ 15.2 ESD #1. A LONG BINARY column now maps to a Microsoft SQL Server VARBINARY(max) column. In Sybase IQ 15.2 GA, a VARBINARY column was mapped to a Microsoft SQL Server VARBINARY(max) column.

4.3 Administration and troubleshooting improvements

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

4.3.1 Default password value change

The default password is now sql in lowercase.

4.4 Data load, update, and extraction enhancements

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

4.4.1 LOAD TABLE HEADER SKIP syntax [CR #593620]

As of Sybase IQ 15.2 ESD #2, you can use the HEADER SKIP clause to specify a number of lines at the beginning of the data file, including header rows, for LOAD TABLE to skip. All LOAD TABLE column specifications and other load options are ignored.

Syntax

```
LOAD [ INTO ] TABLE [ owner.]table-name
... ( load-specification [, ...] )
... { FROM | USING [ CLIENT ] FILE }
{ 'filename-string' | filename-variable } [, ...]
.
.
.
... [ HEADER SKIP number [ HEADER DELIMITED BY 'string' ] ]
```

Parameters

number: the number of lines to skip; greater than or equal to zero

string: the line delimiter, 1 to 4 characters in length; default is '\n'

Examples

Example 1 Ignore one header row at the beginning of the data file, where the header row is delimited by '&&'.

```
LOAD TABLE
...HEADER SKIP 1 HEADER DELIMITED by '&&'
```

Example 2 Ignore 2 header rows at the beginning of the data file, where each header row is delimited by '\n'.

```
LOAD TABLE
...HEADER SKIP 2
```


Usage

You can use the `HEADER SKIP` clause to specify a number of lines at the beginning of the data file, including header rows, for `LOAD TABLE` to skip.

- Lines are determined by a delimiter string specified in the `HEADER DELIMITED BY` clause. The default `HEADER DELIMITED BY` string is the `'\n'` character.
- The `HEADER DELIMITED BY` *string* has a maximum length of four characters. An error is returned, if the string length is greater than four or less than one.
- When a non-zero `HEADER SKIP` value is specified, all data inclusive of the `HEADER DELIMITED BY` delimiter is ignored, until the delimiter is encountered the number of times specified in the `HEADER SKIP` clause.
- All `LOAD TABLE` column specifications and other load options are ignored, until the specified number of rows has been skipped. After the specified number of rows has been skipped, the `LOAD TABLE` column specifications and other load options are applied to the remaining data.
- The “header” bytes are ignored only at the beginning of the data. When multiple files are specified in the `USING` clause, `HEADER SKIP` only ignores data starting from the first row of the first file, until it skips the specified number of header rows, even if those rows exist in subsequent files. `LOAD TABLE` does not look for headers once it starts parsing actual data.
- No error is reported, if `LOAD TABLE` processes all input data before skipping the number of rows specified by `HEADER SKIP`.

4.5 Query enhancements, optimization, and changes

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

4.5.1 NGRAM TEXT index and fuzzy and non-fuzzy search

As of Sybase IQ 15.2 ESD #2, Sybase IQ supports NGRAM TEXT index, fuzzy search, and non-fuzzy search.

NGRAM TEXT index

NGRAM TEXT index stores the text in the column by breaking the text into NGRAMs of text value N , where N is the value given by a user. You can perform a search over an NGRAM TEXT index by matching the NGRAMs of the text value in the CONTAINS clause of the query against the stored NGRAMs in the index.

NGRAM TEXT index accommodates fuzzy searching capability over the text for both European and non-European languages.

Note NGRAM TEXT index search is mainly useful when words are misspelled. Sybase IQ does not support searches like synonyms and antonyms.

NGRAM is built on TEXT indexes, so use text configuration object settings to define whether to use an NGRAM or GENERIC TEXT index.

For more information on text configuration object settings, see *SQL Anywhere 11.0.1 > SQL Anywhere Server - SQL Usage > Querying and Modifying Data > Querying Data > Text configuration objects > Text configuration object settings*.

Creating NGRAM TEXT index

For information on how to create a NGRAM TEXT index, see *SQL Anywhere 11.0.1 > SQL Anywhere Server - SQL Usage > Querying and Modifying Data > Querying Data > Types of full text searches > Tutorial: Performing a fuzzy full text search*.

Fuzzy search

Fuzzy search capability over a TEXT index is possible only if the TEXT index is of type NGRAM. The GENERIC TEXT index cannot handle the fuzzy search.

Here is an example of a fuzzy search over an NGRAM TEXT index. Create a table and a NGRAM TEXT index:

```
CREATE TEXT CONFIGURATION NGRAMTtxtcfg
  FROM default_char;
ALTER TEXT CONFIGURATION NGRAMTtxtcfg TERM BREAKER
  NGRAM;
ALTER TEXT CONFIGURATION NGRAMTtxtcfg maximum term
  length 3;
CREATE TABLE t_iq(a int, b varchar(100));
CREATE TEXT INDEX TXT_IQ on t_iq(b) CONFIGURATION
  NGRAMTtxtcfg
```

Insert this data into the table:

```
INSERT INTO t_iq values (1,'hello this is hira ');
INSERT INTO t_iq values(2, ' book he ookw worm okwo
kwor');
INSERT INTO t_iq values(3,'Michael is a good person');
INSERT INTO t_iq values(4,'hello this is evaa');
INSERT INTO t_iq values(5,'he is a bookworm');
INSERT INTO t_iq values (6,'boo ook okw kwo wor orm');
```

After inserting the data, execute this query to perform fuzzy searching over an NGRAM TEXT index:

```
SELECT * FROM t_iq WHERE CONTAINS (b, 'FUZZY "bookerm"');
```

The results of the query are:

a	b
2	book he ookw worm okwo kwor
5	he is a bookworm
6	boo ook okw kwo wor orm

Fuzzy searching can be used to search for misspellings or variations of a word. To do so, use the FUZZY operator followed by a string in double quotes to find an approximate match for the string.

Using the FUZZY operator is equivalent to breaking the string manually into substrings of length *n* and separating them with OR operators. For example, if you have a text index configured with the NGRAM term breaker and a MAXIMUM TERM LENGTH of 3, specifying 'FUZZY "500 main street" ' is equivalent to specifying '500 OR mai OR ain OR str OR tre OR ree OR eet'.

The FUZZY operator is useful in a full text search that returns a score. Many approximate matches may be returned, but usually only the matches with the highest scores are meaningful.

Example 1—This query illustrates a additional letter in the fuzzy search clause.

```
SELECT * FROM t_iq WHERE CONTAINS (b, 'FUZZY "hellow"');
```

The results of the query are:

a	b
1	hello this is hira
4	hello this is evaa

Example 2—In this query, a letter is removed from the fuzzy search clause.

```
SELECT * FROM t_iq WHERE CONTAINS(b, 'FUZZY "hlllo"');
```

The results of the query are:

a	b
1	hello this is hira
4	hello this is evaa

Note Fuzzy search does not support prefix or suffix searching. For example, the search clause cannot be “v*” or “*vis”.

For more information on fuzzy search, see *SQL Anywhere 11.0.1 > SQL Anywhere Server - SQL Usage > Querying and Modifying Data > Querying Data > Types of full text searches > Fuzzy searches*.

Non-fuzzy search

Non-fuzzy search on NGRAM breaks the term into corresponding NGRAMs and searches for the NGRAMs in the NGRAM TEXT index.

The query `CONTAINS (M.Description, 'ams') ct;` illustrates a non-fuzzy NGRAM search over a 2GRAM index, which is semantically equal to searching query `CONTAINS(M.Description, "'am ms'") ct;`

If you search for a ‘v*’ TERM on a 2GRAM index, then v followed by any alphabet is considered as a matching 2GRAM for the searching term and is output as a result.

The query `CONTAINS (M.Description, 'white whale')` ct; illustrates a non-fuzzy NGRAM search over a 3GRAM index and is semantically equal to searching query `CONTAINS (M.Description, "whi hit ite wha hal ale")`;

The difference between NGRAM fuzzy and non-fuzzy search is that fuzzy search is a disjunction over individual GRAMS. Non-fuzzy search is a conjunction over the individual GRAMS. When `GENERIC` and `NGRAM TEXT` indexes are created on the same column, then the `GENERIC TEXT` index is used for a query with non-fuzzy search and the `NGRAM TEXT` index is used for fuzzy search.

Example 3—This query illustrates non-fuzzy search after creating a `GENERIC TEXT` index on the same column.

```
SELECT * FROM t_iq WHERE CONTAINS (b, 'bookworm');
```

The results of the query are:

a	b
5	he is a bookworm

Example 4—This query illustrates fuzzy search with both `NGRAM` and `GENERIC TEXT` indexes on the same column.

```
SELECT * FROM t_iq
WHERE CONTAINS (b, 'FUZZY "bookwerm"');
```

The results of the query are:

a	b
2	book he ookw worm okwo kwor
5	he is a bookworm
6	boo ook okw kwo wor orm

Example 4—This query illustrates the behavior of a fuzzy search phrase in a non-fuzzy search clause.

```
SELECT * FROM t_iq WHERE CONTAINS (b, 'bookwerm');
```

No result is returned for this query.

4.6 Security enhancements

This section contains new features and changes related to security.

4.6.1 System procedure behavior change [CR #589250]

The SQL text of all system procedures is hidden. `sp_iqhelp` returns encrypted results for the procedure definition of any Sybase IQ system procedure.

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 56 for instructions on accessing the most current online certification information.

5.1 Sybase products

These Sybase products are certified with Sybase IQ 15.2:

- SQL Anywhere 11.0.1
- Adaptive Server Enterprise 15.5
- jConnect™ for JDBC™ 5.x and 6.x
- Sybase ETL 4.8
- Sybase Control Center (SCC) 3.1

Refer to “Sybase certifications on the Web” on page 56 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 15.2 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 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. Doing so 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. These file placement recommendations also apply to log files.

To avoid conflicts, Sybase recommends that dbspace management be performed by a single database administrator on a single connection.

Performance related to randomly accessed files, including the System, IQ Main, and IQ Temporary dbfiles, can be improved by increasing the number of disk drives devoted to those files. Performance related to sequentially accessed files, including the transaction log and message log files, can be improved by locating these files on dedicated disk drives.

Suggestions to reduce file placement impact on performance:

- Keep random disk I/O away from sequential disk I/O.
- Place the database file, temporary dbspace, and transaction log file on the same physical machine as the database server.
- Isolate Sybase IQ database I/O from I/O for proxy tables in other databases, such as Adaptive Server Enterprise.
- Place the transaction log and message log on separate disks from the IQ store, catalog store, and temporary store, and from any proxy databases such Adaptive Server Enterprise.

6.3 Dbspace naming in RESTORE command [CR #561366]

If the dbspace name contains a file extension such as *.iq*, *.iqtmp*, or *.iqloc*, you must enclose the dbspace name 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'
```

6.4 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; SQL Anywhere treats CHAR data as if it were VARCHAR.
- When the RAND function is passed an argument, the behavior is deterministic in Sybase IQ and nondeterministic in SQL Anywhere.

6.5 Interactive SQL restrictions

Sybase recommends that you use dbisql, rather than dbisqlc or iqisql, unless instructed to do otherwise, for specific tasks. dbisqlc is supported, but does not contain all the features of dbisql.

- Do not use the Import option on 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) 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 option “Show multiple result sets,” click Tools | Options in the dbisql window, choose “SQL Anywhere,” and select the “Results” tab. Then select “Show all results sets.”

- There is no online help 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) instead, or use an xterm window to run dbisqlc on UNIX and Linux systems. For example, to start an xterm window with a scroll bar, enter:

```
% xterm -sb
```

7. Known problems

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 14, and “Documentation updates and clarifications” on page 34.

7.1 SySAM licensing

Learn about known issues and apply workarounds with SySAM licensing.

7.1.1 Changes to licensing behavior [CR #628596, CR #628403]

Sybase IQ cannot be installed or fails to start in a virtual environment and reports an error similar to the following:

```
Cannot determine the total number of processor cores on
the machine from within this virtual system. The IQ_CORE
("DH") license type is restricted for use on a machine
with a maximum of 4 processor cores regardless of
whether the virtual system is limited to use fewer
processor cores. You need to run 'sysamcap' on the
physical machine (or virtual-machine/partition control
domain) and set the reported environment variable prior
to running Installer.
```

This error is raised in the following circumstances:

- 1 Sybase IQ is running on a virtual system which doesn't have access to the processor counts of the physical machine (this applies to all virtual environments except AIX LPAR and HP-UX vPar).

and

- 2 Sybase IQ is using a license type that is restricted for use to a machine with no more than a specific number of processors (for example, Sybase IQ Single Application Server Edition cannot be used on a machine with more than 16 processor cores, ASE Small Business Edition cannot be used on a machine with more than 2 processor chips).

Workaround:

Run the `sysamcap DEFAULT` command on the physical machine or control domain for the virtualization technology (for example, VMware ESX console, or XEN dom0), then set the environment variable reported prior to running the product or product installer. The process is described in the SySAM Sub-capacity Licensing section of the *SySAM Users Guide*.

7.2 Installation

Learn about known issues and apply workarounds for Sybase IQ installations.

7.2.1 Drop and recreate TEXT indexes after ESD upgrade or rollback [CR #620317]

In Sybase IQ version 15.2 ESD #1 or later, operations that access a Sybase IQ TEXT index created in 15.2 GA report the error “Incorrect TEXT Index Version (1): Index must be dropped and recreated.” Also, `sp_iqcheckdb` reports errors and inconsistent indexes related to these TEXT indexes.

Workaround: Use Sybase IQ 15.2 ESD #1 or later to drop and recreate all TEXT indexes created in 15.2 GA.

After a rollback of Sybase IQ to version 15.2 GA, an operation that accesses a TEXT index created with 15.2 ESD #1 or higher reports an error like “Non-Fatal IQ Internal Error Detected at `s_postingTbl.cxx:1503` on thread 4113.”

Workaround: Use Sybase IQ 15.2 GA to drop and recreate all IQ TEXT indexes.

7.3 Sybase IQ operations

Learn about known issues and apply workarounds for Sybase IQ operations.

7.3.1 Incorrect result from LOCATE with CLOB [CR #648637]

The LOCATE string function has an optional third parameter, which may be supplied to provide a starting offset for the LOCATE string search.

The syntax for LOCATE is:

```
LOCATE(string-expression1, string-expression2 [, numeric-expression])
```

If *string-expression1*, the string to be searched, is a CLOB data type and the *numeric-expression* provided for the offset is evaluated as a BIGINT data type, the offset may be interpreted incorrectly.

Workaround: To ensure the offset is evaluated correctly, either use an integer constant or CAST the *numeric-expression* as an integer.

This example illustrates both workarounds.

```
CREATE TABLE t1(c_clob CLOB, c_bigint BIGINT);
INSERT INTO t1 VALUES ('abc abc def', 4);
```

CAST *numeric-expression* as an integer:

```
SELECT LOCATE( c_clob, 'abc',
              CAST(LOCATE(c_clob,'abc def') as INT)) FROM t1;
SELECT LOCATE( c_clob, 'abc',
              CAST(c_bigint as INT)) FROM t1;
```

Or use an integer constant:

```
SELECT LOCATE( c_clob, 'abc', 4) FROM t1;
```

For more information on the LOCATE function, refer to *Reference: Building Blocks, Tables, and Procedures*.

7.3.2 NGRAM TEXT index CONTAINS BEFORE search [CR #642744]

An NGRAM TEXT index search with a CONTAINS... BEFORE clause may fail with the error:

```
Could not execute statement.
Syntax error - CONTAINS Near/Before first arguments
must be single term.
-- (df_Heap.cxx 6876)
SQLCODE=-1005028, ODBC 3 State="HY000"
```

The same query with a CONTAINS... NEAR clause returns a result without error.

7.3.3 Login fails when FORCE CHANGE PASSWORD = ON [CR #640915]

If FORCE CHANGE PASSWORD = ON, a login using Interactive SQL (isql) fails with the error:

```
Msg 13361, Level 16, State 0:  
SQL Anywhere Error -1152: Password has expired  
CT-LIBRARY error:  
ct_connect(): protocol specific layer: external error:  
The attempt to connect to the server failed.
```

Workaround: Log in using dbisql.

7.3.4 Server failure when running sa_text_index_vocab procedure [CR #627966]

If two connections simultaneously run the sa_text_index_vocab system procedure against the same TEXT index, the server may fail with a Sybase IQ Abort and the message caught signal 11, program abort.

Workaround: Call sa_text_index_vocab on a given TEXT index from one connection at a time.

7.3.5 PHP bound parameters on big endian platforms [CR #627872]

In Sybase IQ 15.2, PHP: Hypertext Preprocessor (PHP) drivers on big endian platforms require bound parameters to be initialized before the call to sasql_stmt_bind_param for INT and BIGINT data types.

This example shows the values properly initialized before the call to sasql_stmt_bind_param:

```
$stmt = sasql_prepare($conn, "insert into  
testdefault(c1, c2, c3, c5) values(?,?,?,?)");  
#Binding parameters with statement prepared  
$c1=22; $c2=33; $c3="col3data"; $c5="col5data";  
sasql_stmt_bind_param  
($stmt,"iiss", $c1, $c2, $c3, $c5);  
#executing statement  
sasql_stmt_execute($stmt);
```

7.3.6 SySAM does not start on Xen and VMWare [CR #625061]

In order to run Sybase IQ in Linux Xen DomainU or VMWare ESX guest OS partition, you must first create an environment variable that instructs SySAM to check for a virtual machine.

For Xen, set SYBASE_SAM_CPUINFO=Xen or SYBASE_SAM_CPUINFO=XENCHECKXEN.

For VMWare, set SYBASE_SAM_CPUINFO=VMCHECKVM.

7.3.7 Remote query error [CR #626133]

NO PASSTHRU and PARTIAL PASSTHRU modes are useful in remote queries that use Component Integration Services (CIS) functional compensation, or in remote queries that join Sybase IQ tables to SQL Anywhere tables. When you execute a remote query in NO PASSTHRU or PARTIAL PASSTHRU mode and the query execution plan contains a group-by hash, distinct hash, join hash, or sort operator, a statement may fail and issue this error:

```
-728,Update operation attempted on non-updatable  
remote query error
```

The affected statement fails, but no other statements or connections are impacted.

This error occurs only when the result computed by the operator is much larger than the available memory. Other operations in the query plan may reduce the amount of available memory. If available memory is small—for example, if fewer than 20 rows can be held in the available memory at one time—then Sybase IQ may report the above error. If the result set is very large—for example, if it contains more than 520×10^9 rows—then Sybase IQ may report the above error.

To avoid this problem, increase the amount of memory allocated to SQL Anywhere by increasing the SQL Anywhere cache, or wait for the load on the server to decrease, or rewrite the query so that it can be executed in FULL PASSTHRU mode.

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

On Sybase IQ servers that are set up for Chinese and Japanese locales, 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
```

Workaround:

- 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 procedure, certain strings in the *.iqmsg* files and the output of sp_iqmpxinfo, sp_iqstatistics, and sp_iqstatus are in English, instead of Chinese or Japanese.

7.3.9 ASE connections to Sybase IQ with BIGDATETIME [CR #622007]

When a version 15.5 GA Adaptive Server Enterprise server uses Component Integration Services (CIS) to connect to a Sybase IQ server version that has BIGDATETIME and BIGTIME capability, any chronological data type sent to the Adaptive Server server using CIS results in the error:

```
Msg 7225, Level 16, State 4:  
Line 1:  
Unknown datatype token 188 'BIGDATETIME NULL'  
encountered.  
Exited passthru mode from server 'QA_IQ15_ASECIS'.
```

Sybase IQ sends the data as BIGDATETIME and relies on Adaptive Server to convert appropriately. The BIGDATETIME data type was implemented in CIS in Adaptive Server Enterprise version 15.5 ESD #1, so this error occurs with Adaptive Server Enterprise 15.5 before any ESDs have been applied.

Workaround:

```
SET TEMPORARY OPTION RETURN_DATE_TIME_AS_STRING='ON' in
Sybase IQ for each session. The Sybase IQ server sends all chronological data
as strings and Adaptive Server completes the conversion. This workaround is
for Adaptive Server Enterprise 15.5 GA connecting to Sybase IQ in
passthrough mode.
```

You can also use this workaround in the definition of a remote stored procedure, but it is not a solution for proxy tables created on Adaptive Server returning chronological data types from Sybase IQ.

7.3.10 dbisql does not start on Linux Red Hat 5.3 [CR #571993]

The Interactive SQL utility dbisql fails to start on Linux Red Hat 5.3, and issues a message similar to:

```
Error! could not load the Java Virtual machine DLL:
/root/users/user1/050509/shared/JRE-
6_0_7_32BIT/lib/i386/client/libjvm.so
```

Workaround: Run:

```
% cd $IQDIR15/bin32
% dbisql -batch # creates below
% dbisql.sh
```

7.3.11 SELinux policy requirement for Sybase IQ 15.x on Linux PowerPC x64 [CR #571627]

You must install selinux-policy-2.4.6-25.e15 or later to use the java -version command with SELinux enabled. If the policy version is incorrect, the system displays this message:

```
Errorloading: /libjvm.so: cannot restore segment prot
after reloc: Permission denied
```

This is the preferred method to meet this policy requirement, but alternatively, you can disable SELinux.

7.3.12 Stored procedures for information only [CR #560032]

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.

7.3.13 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 Technical Support, if you need this utility.

7.3.14 FORWARD TO from Sybase IQ to ASE fails [CR #539484]

Execution of a `FORWARD TO` command from Sybase IQ to Adaptive Server Enterprise using the ASEODBC driver fails with ASA Error -660 “The identifier that starts with ‘<identifier>’ is too long. Maximum length is 28.”

Workaround:

To avoid this problem, do one of the following after creating a remote Adaptive Server Enterprise server using either server class `aseodbc` or server class `asejdbc` and before executing a SQL statement that contains a string with double quotes against the remote server using the `FORWARD TO` statement:

- 1 Change the double quotes to single quotes.
- 2 Include the database option `SET QUOTED_IDENTIFIER OFF` before the SQL statement to be executed:

```
FORWARD TO <remote server name> {  
  SET QUOTED_IDENTIFIER OFF  
  <SQL statement>  
}
```

- 3 Set the `QUOTED_IDENTIFIER` database option to `OFF` using a separate `FORWARD TO` statement:

```
FORWARD TO <remote server name> {  
  SET QUOTED_IDENTIFIER OFF }  
FORWARD TO <remote server name> {  
  SET QUOTED_IDENTIFIER OFF }
```


7.4 Multiplex environment

Learn about known issues and apply workarounds within a multiplex environment.

7.4.1 Restoring IQ store data in a different location [CR #623585, CR #633738]

There is an incorrect procedure in *Using Sybase IQ Multiplex > Backing Up and Restoring Data in a Multiplex Environment > Backing up and restoring data in a multiplex environment > Restoring in the multiplex environment*. This is the corrected procedure (steps 5, 6, 7, and 8 differ from the original version):

❖ 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. You may specify any valid identifier as a server name except the name of a registered secondary server. If you use the name of the coordinator, use that name when restoring.

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

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

```
% dbisql -c
"eng=utility_startup_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 supply a `RENAME` clause for dbspaces to move. For details, see the `RESTORE` statement in *Reference: Statements and Options*.

- 8 Start the restored database with single-node and override flags (`-iqmpx_sn 1 -iqmpx_ov 1`).

When creating a shadow multiplex (on the same network but in a different location), do not start the restored database with the coordinator server name if the original multiplex is running with that server name. Either use a unique server name or start the database with `-DoBroadcast=NO`.

- 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 Re-create 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 coordinator 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 30.
- 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.

7.4.2 Emergency shutdown on secondary servers after `CREATE TEXT INDEX` [CR #611990]

A secondary server goes into emergency shutdown, when:

- It is started with the `-sf external_library_full_text` flag, which disables loading of external libraries, and
- The user issues a `CREATE TEXT INDEX` statement on the coordinator with a text configuration that uses external libraries.

All other servers replay the DDL successfully.

To avoid the problem, start secondary nodes in the multiplex without the `-sf external_library_full_text` flag.

7.4.3 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.

7.5 Sybase Central

Learn about known issues and apply workarounds for Sybase Central.

7.5.1 Create Trigger Wizard does not list user tables [CR #631447]

The Sybase Central Create Trigger Wizard does not list user tables.

Workaround: Use Interactive SQL (dbisql) to create a trigger.

7.5.2 Editing Sybase Central password field [CR #625254]

In Sybase Central, the password field cannot be edited on the Connection Parameters step in the Create Database Wizard, unless password is the first field edited.

Workaround:

- Set the LC_ALL environment variable in the scjview script to a value other than LC_ALL=C (recommended), or
- Click Back or Next on the Connection Parameters screen.

7.5.3 Sybase Central warning fails to refresh [CR #623638]

Viewing dbspace contents requires DBA authority or EXECUTE permission to run the sp_iqdbspaceobjectinfo procedure.

The first time you try to view dbspace contents without DBA authority or EXECUTE permission, Sybase Central returns an error message. Successive tries continue to show the same error, even when you press F5 to refresh. To improve performance, Sybase Central does not refresh each time you change folders.

7.5.4 Sybase Central fails to start after ASE install [CR #618082]

After installing Adaptive Server Enterprise (ASE) in the same directory in which Sybase IQ is already installed, Sybase Central fails to start and this error is reported:

```
Unable to locate the Java Runtime Environment
($SYBASE_JRE6).
The directory $SYBROOT[/shared]/jre-6_0 is missing or
invalid.
```

Workaround:

Source *IQ-15_2.csh* for Sybase IQ, then source *SYBASE.csh* for ASE, before starting Sybase Central:

```
cd $SYBASE
source IQ-15_2.csh
source SYBASE.csh
cd $SYBASE/shared/sybcentral600
./scjview.sh &
```

7.5.5 Sybase IQ plug-in for Sybase Central not registered [CR #571188]

If Sybase IQ is installed on the same system as other products, such as Adaptive Server Enterprise (ASE), Replication Server®, or Replication Manager, the Sybase IQ plug-in for Sybase Central may not be registered automatically. If you attempt to register the Sybase IQ plug-in, Sybase Central reports this error:

```
The library "dbput11_r" could not be loaded. Please
ensure that this file exists as part of your
installation.
```

Workaround:

Source *IQ-15_2.csh* for Sybase IQ, then source *SYBASE.csh* for ASE, Replication Server, or Replication Manager, before starting Sybase Central:

```
cd $SYBASE
source $SYBASE/IQ-15_2/IQ-15_2.csh
source SYBASE.csh
setenv LD_LIBRARY_PATH_64
"$SYBASE/IQ-15_2/lib64:${LD_LIBRARY_PATH_64}"
cd $SYBASE/shared/sybcentral600
./scjview.sh &
```

7.5.6 Remote client cannot connect to Sybase IQ Agent [CR #563823]

When Sybase Central running on a remote client attempts to connect to the Sybase 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.

7.5.7 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.

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

7.5.8 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.

7.5.9 Starting Sybase Central online help on Linux on POWER [CR #390320]

When you use the Help button from a Sybase Central wizard or utility on Linux on POWER, the loading image displays the message “Loading online help,” but remains open without starting help. If the image displays in a window with a Close icon (for example, in an Exceed for Windows display), close it manually. If there is no icon, use Alt+F4 (for example, on Solaris CDE display).

Once initialized, the help system works correctly and may subsequently be invoked from wizards or utilities without problems.

8. Documentation for this version

Documentation for Sybase IQ version 15.2 is included with your product shipment. All the information you need to get started and to use the new features in Sybase IQ 15.2 is on the Sybase IQ Getting Started CD and the Sybase IQ SyBooks™ CD.

Sybase Central includes online help for Sybase IQ.

Table 2 lists documentation on the Sybase IQ 15.2 Getting Started CD.

Table 2: Sybase IQ 15.2 Getting Started CD contents

Part number	Book title
DC01014-01-1520-01	<i>Quick Start</i>
DC10082-01-1520-01	<i>Release Bulletin for Linux</i>
DC10083-01-1520-01	<i>Installation and Configuration Guide for Linux</i>

Table 3 lists Sybase IQ documentation on the Sybase IQ 15.2 SyBooks CD.

Table 3: Sybase IQ SyBooks contents

Part number	Book title
DC01087-01-1520-01	<i>New Features Summary Sybase IQ 15.2</i>
DC00462-01-1520-01	<i>Error Messages</i>
DC38159-01-1520-01	<i>Introduction to Sybase IQ</i>
DC38151-01-1520-01	<i>Reference: Building Blocks, Tables, and Procedures</i>
DC00801-01-1520-01	<i>Reference: Statements and Options</i>
DC00170-01-1520-01	<i>System Administration Guide: Volume 1</i>
DC00800-01-1520-01	<i>System Administration Guide: Volume 2</i>
DC00169-01-1520-01	<i>Performance and Tuning Guide</i>
DC00168-01-1520-01	<i>Utility Guide</i>
DC01268-01-1520-01	<i>Unstructured Data Analytics in Sybase IQ</i>
DC01015-01-1520-01	<i>Advanced Security in Sybase IQ</i>
DC00802-01-1520-01	<i>Using Sybase IQ Multiplex</i>
DC01034-01-1520-01	<i>User-Defined Functions Guide</i>
DC01033-01-1520-01	<i>IMSL Numerical Library User's Guide: Volume 2 of 2 C Stat Library</i>
DC01154-01-1520-01	<i>Time Series Guide</i>

Check Product Documentation at <http://sybooks.sybase.com> for the latest available versions of the Sybase IQ 15.2 documents.

Note these changes to the documentation set in Sybase IQ 15.2:

- The *Time Series Guide* contains the descriptions of the SQL functions for time series forecasting and analysis, formerly in *Reference: Building Blocks, Tables, and Procedures*.
- *Unstructured Data Analytics in Sybase IQ* includes the former contents of *Large Objects Management in Sybase IQ*, plus instructions about the new full text search capability. See *New Features Summary Sybase IQ 15.2*.

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 Documentation at <http://sybooks.sybase.com> and in DocCommentXchange at http://dcx.sybase.com/dcx_home.php.

Note Most 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):

- *SySAM 2 Users Guide* 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 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 New Features Summary Sybase IQ 15.2

This section contains updates to *New Features Summary Sybase IQ 15.2*.

9.1.1 ODBC Driver Manager library name on Unix and Linux platforms [CR #646312]

Sybase IQ now provides ODBC driver manager capability on all supported UNIX and Linux platforms. For applications that will only access Sybase IQ with one data source, Sybase provides the *libdbodbl1* shared object for linking. For applications that will access multiple data sources, Sybase provides the *libdbodbm11* shared object.

See “Using the SQL Anywhere ODBC Driver Manager on UNIX” in *SQL Anywhere Server - Programming* for more information about the *libdbodm* shared object.

9.1.2 CIS performance enhancements

Sybase IQ uses Component Integration Services (CIS) to query tables on remote servers. Changes in Sybase IQ 15.2 allow queries with proxy tables or SA IN SYSTEM tables to execute faster than earlier versions, when the number of rows returned from an IN SYSTEM or proxy table is small compared to the number of rows returned from Sybase IQ.

Queries that benefit from these changes include:

- SELECT, SELECT INTO, or INSERT...SELECT statements only
- A statement that contains one or more Sybase IQ tables and one or more proxy or SA IN SYSTEM tables

Queries that do not benefit from these changes include:

- References to system catalog tables
- A subquery in a CASE expression
- LOBs (columns of LONGBINARY, LONGVARCHAR, LONGNVARCHAR, LONGBITSTRING, or XML) on remote tables
- A function not supported by target node (for example, bit_or or user defined functions)
- Queries with global variables

9.2 Interactive SQL (dbisql) online help

There is no online help available for Interactive SQL (dbisql or dbisqlc).

9.3 System Administration Guide: Volume 1

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

9.3.1 Using ODBC data sources on UNIX [CR# 646317]

Sybase IQ Connections > *Using ODBC data sources on UNIX* incorrectly states that only the ODBC driver is installed on UNIX and Linux platforms. Sybase IQ 15.2 includes an ODBC driver and the iAnywhere driver manager.

For clarity, we changed the *dbodbc11.so* reference to *libdbodbc11.so*. The sentence that reads “If you are creating a custom ODBC application, you can link directly to *dbodbc11.so*” now reads “If you are creating a custom ODBC application, you can link directly to *libdbodbc11.so*.”

9.3.2 Formula for calculating microseconds [CR #644926]

Moving Data In and Out of Databases > Using binary load format > TIME should read as follows:

TIME data is stored as a 64-bit unsigned quantity that represents a number in microseconds (in other words, 1.0e-6 seconds). Compute the microsecond quantity for a given hour, minute, second, and microsecond (*usec*):

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

9.3.3 Merge certain DATE partitions before migration to 15.2 ESD #2 [CR #640727]

This section is added to *Working with Database Objects > Building your Sybase IQ databases > Range partitioning*:

In versions before 15.2 ESD #2, Sybase IQ ignored database partitions defined on a column of type DATE when the partition definition contained time portions.

For example, assume that you created this table before migrating to 15.2 ESD #2:

```
CREATE TABLE pt_date(c1 INT, c2 DATE)
IN IQ_SYSTEM_MAIN
PARTITION BY RANGE (c2)
( p1 VALUES <= ('2006-03-31' ) ,
p2 VALUES <= ('2006-06-30:01:02' ) ,
p3 VALUES <= ('2006-09-30' ) ,
p4 VALUES <= ('2006-12-31' )
)
INSERT INTO pt_date VALUES( 1, '2006-03-30');
INSERT INTO pt_date VALUES( 1, '2006-05-30:01:02');
INSERT INTO pt_date VALUES( 1, '2006-10-30');
COMMIT;
```

After migration to 15.2 ESD #2, an operation (such as SELECT) on the pt_date table returns the error “cannot convert '2006-06-30:01:02' to a date.”

To avoid such errors:

- 1 Before migration, to avoid the error, merge the partition to its next partition. This creates a new partition range and removes the former range:

```
ALTER TABLE pt_date MERGE PARTITION p2 INTO p3
```

- 2 Split the partition into the desired partition definition with the same partition criteria, but without a time portion:

```
ALTER TABLE pt_date
split PARTITION p3 INTO
(P31 VALUES <= ('2006-06-30') ,
 P32 VALUES <= ('2006-09-30') );
```

9.3.4 Troubleshooting server startup [CR #638520]

The following is added to *Troubleshooting Hints > Solutions for specific conditions > Server operational issues > Sybase IQ will not start*:

Server has been started with an incorrect path When you start a new multiplex server, the database file path must match the database file path specified when creating that server. If you use the wrong path, server startup fails with the following messages in the server log file (*.srvlog):

```
E. 08/18 07:22:19. MPX: server myserver
has been started with an incorrect catalog path
(expected path: /work/IQ-15_3/demo/mympx/iqdemo.db).
-- (st_database.cxx 7883)
I. 08/18 07:22:19. Database server shutdown due
to startup error
DBSPAWN ERROR: -82
Unable to start specified database: autostarting
database failed
```

If you see these messages, restart the server with the expected path. If you plan to use UNIX soft (symbolic) links for server paths, you must create the soft link before you run CREATE MULTIPLEX SERVER.

9.3.5 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”.

9.3.6 Null .ini file error [CR #632953]

The following is added to *Troubleshooting Hints > Solutions for specific conditions > Database connection issues > Cannot connect to a database*:

On UNIX or Linux systems using an aseodbc server class, the error message returned is:

```
Unable to connect to server
```

followed by the message:

```
Null ini file or ini file not found in $SYBASE or current  
directory
```

Make sure that your .ini file is named .odbc.ini and that the ODBCINI environment variable contains the exact full path name of the .odbc.ini file.

9.3.7 Remove .iq files before restoring [CR #620249]

In *Data Backup, Recovery, and Archiving > Restoring your databases > Before you restore*, replace the sixth bullet item with this text:

- 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, delete them or move them to a different directory before doing the full restore.

A full restore destroys all old database files and re-creates them. Manually deleting the store, catalog store, and transaction log files protects you from inadvertently performing a full restore.

9.4 System Administration Guide: Volume 2

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

9.4.1 Sybase IQ 15.x does not include iqprocs.sql file [CR #648475]

In *Using Procedures and Batches > Using IQ UTILITIES to create your own stored procedures > Requirements for using IQ UTILITIES*, references to the *iqprocs.sql* file are not valid. Sybase IQ 15.x releases do not include *iqprocs.sql*. All SQL code for procedures is encrypted and compiled into the shared library *libiqscripts15_r.so* on UNIX and *iqscripts15.dll* on Windows.

You can view the stored procedure code by using Sybase Central or by entering `sp_helptext 'owner.procname'` in Interactive SQL.

9.4.2 Correction in OLAP AVG function example result output [CR #634333]

In *Using OLAP > Analytical Functions > Windowing > Window Framing > CURRENT ROW concept* the correct result output for the OLAP AVG function example is:

- Row [1] = (10 + 50 + 100)/3
- Row [2] = (50+ 100 + 120)/3
- Row [3] = (100 + 120 + 500)/3
- Row [4] = (120 + 500 + NULL)/2
- Row [5] = (500 + NULL + NULL)/1

9.4.3 Null ini file error for aseodbc connection [CR #632953]

The following was added to *Server Classes for Remote Data Access > ODBC-based server classes*:

Note If Sybase IQ returns the error `Null ini file or ini file not found`, make sure that your `.ini` file is named `.odbc.ini` and that the `ODBCINI` environment variable contains the exact full path name of the `.odbc.ini` file.

9.4.4 Syntax correction in CREATE SERVER statement [CR #628912]

In *Accessing Remote Data > Sybase IQ and remote data > Working with remote servers > Querying data without native classes > Querying remote data using DirectConnect and proxy table from UNIX > Setting up Sybase IQ to connect to the ASE server*, the correct syntax for `CREATE SERVER` is:

```
create SERVER jones_1207 CLASS 'asejdbc'  
    USING 'jones:4101/tst_db'
```

9.5 Reference: Building Blocks, Tables, and Procedures

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

9.5.1 Correction to sp_iqemptyfile procedure [CR #639511]

System Procedures > System stored procedures > sp_iqemptyfile procedure is updated as follows:

Function – Empties a dbfile and moves the objects in the dbfile to another available read-write dbfile in the same dbspace.

Description – `sp_iqemptyfile` empties a dbfile. The dbspace must be read-only before you can execute the `sp_iqemptyfile` procedure. The procedure moves the objects in the file to another available read-write dbfile in the same dbspace. If there is no other read-write dbfile available, then Sybase IQ displays an error message.

9.5.2 Dbspace does not support NAS [CR #629367]

Currently *Physical Limitations > Size and Number Limitations > Table 6-1: Sybase IQ database object size and number limitations* contains the following two rows:

Item	Limitation
Catalog file size	Maximum is 1TB for all platforms except for Windows systems with FAT 32-file systems, which have a 4GB limit. Windows systems with NTFS support the 1TB maximum. Sybase IQ does not support creating dbspaces on NAS (Network Attached Storage) devices.
Dbspace size	Raw device: Maximum size is 4TB. Operating system file: Maximum size supported by the operating system.

The rows should be:

Item	Limitation
Catalog file size	Maximum is 1TB for all platforms except for Windows systems with FAT 32-file systems, which have a 4GB limit. Windows systems with NTFS support the 1TB maximum.
Dbspace size	Raw device: Maximum size is 4TB. Operating system file: Maximum size supported by the operating system. Sybase IQ does not support creating dbspaces on NAS (Network Attached Storage) devices.

9.5.3 Example query correction [CR #627459]

Currently *SQL Functions > System functions > Properties Available for each Database* contains the following example query:

```
select db_property ( 'PageSize') from iq_dummy
```

This query should be:

```
select db_property ( 'PageSize')
```

9.6 Reference: Statements and Options

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

9.6.1 ALTER DBSPACE Example Syntax [CR #647844]

Replace Example 3 in *Reference: Statements and Options > SQL Statements > ALTER DBSPACE statement* with:

```
ALTER DBSPACE DspHist ADD
FILE FileHist3 '/History1/data/file3' SIZE 500MB,
FILE FileHist4 '/History1/data/file4' SIZE 500
```

9.6.2 LOAD TABLE...WITH CHECKPOINT [CR #646321, CR #648651]

Replace the WITH CHECKPOINT description in *Reference: Statements and Options > SQL Statements > LOAD TABLE statement* with:

WITH CHECKPOINT This option is useful only when loading SQL Anywhere tables in a Sybase IQ database.

Use this clause to specify whether to perform a checkpoint. The default setting is OFF. If this clause is set to ON, a checkpoint is issued after successfully completing and logging the statement. If the server fails after a connection commits and before the next checkpoint, the data file used to load the table must be present for the recovery to complete successfully. However, if WITH CHECKPOINT ON is specified, and recovery is subsequently required, the data file need not be present at the time of recovery.

The data files are required, regardless of what is specified for this clause, if the database becomes corrupt and you need to use a backup and apply the current log file.

Warning! If you set the database option CONVERSION_ERROR to OFF, you may load bad data into your table without any error being reported. If you do not specify WITH CHECKPOINT ON, and the database needs to be recovered, the recovery may fail as CONVERSION_ERROR is ON (the default value) during recovery. It is recommended that you do not load tables when CONVERSION_ERROR is set to OFF and WITH CHECKPOINT ON is not specified.

See also *Reference: Statements and Options > Database Options > Alphabetical list of options > CONVERSION_ERROR option [TSQL]*.

For information regarding automatic recovery of Sybase IQ data, see *System Administration Guide: Volume 1 > System Recovery and Database Repair*.

9.6.3 RESTORE statement syntax correction [CR #633572]

In *SQL Statements > RESTORE Statement*, the RENAME clause specifies *logical-dbfile-name*, not *dbspace-name*.

9.6.4 Identifying altered views requiring manual recompile [CR #633463]

This section was added to *SQL Statements > ALTER VIEW Statement*:

Identifying and fixing invalid dependent views

Check for, and correct, any dependent views that become invalid due to changes to their underlying tables.

Under most circumstances the database server automatically recompiles views to keep them valid if the underlying tables change. However, if your table alteration removes or materially changes something referenced by the view definition, then the dependent view becomes invalid. For example, if you remove a column referenced in the view definition, then the dependent view is no longer valid. Correct the view definition and manually recompile the view.

❖ Identifying and Fixing Invalid Dependent Views

- 1 Run `sa_dependent_views` to get the list of dependent views.
- 2 Perform the DDL operation that alters the table. The server automatically disables dependent views, and attempts to recompile them once the DDL is complete.
- 3 Check that all the views listed by `sa_dependent_views` are valid. For example, perform a simple test such as `SELECT * FROM myview`.
- 4 If a view is invalid, it is likely you will need to alter the view definition to resolve the issue. Examine the view definition against the DDL change that you made and make the necessary changes. Run `ALTER VIEW RECOMPILE` to correct the view definition.
- 5 Test the corrected view to make sure it works. For example, perform a simple test such as `SELECT * FROM myview`.

For detailed information on how the database server handles view dependencies, see *SQL Anywhere Server – SQL Usage > Creating Databases > Working with database objects > Working with views > View Dependencies*.

9.6.5 Backing up read-only dbspaces [CR #631508]

The following was added to *SQL Statements > BACKUP Statement* in the paragraph preceding CRC clause:

Once you back up a read-only dbspace that you do not plan to change, you may omit it from subsequent backups, reduce time to back up, and save space. Be careful; changing such a dbspace back to read-write at any time after the backup invalidates the last read-only backup.

9.6.6 Clarification on Selective Backups [CR #630990]

Update the third example in *SQL Statements > RESTORE Statement* as follows:

A selective backup backs up either all READWRITE dbspaces or specific read-only dbspaces or dbfiles. READWRITE backups are a subtype of either full or incremental backups.

Notes:

- You can take a READONLY selective backup and restore all objects from this backup (as in the second example above).
- You can take an all-inclusive backup and restore read-only files and dbspaces selectively.
- You can take a READONLY selective backup of multiple read-only files and dbspaces and restore a subset of read-only files and dbspaces selectively. See “Permissions.”
- You can restore the read-only backup, only if the read-only files have not changed since the backup. Once the dbspace is made read-write again, the read-only backup is invalid, unless you restore the entire read-write portion of the database back to the point at which the read-only dbspace was read-only.
- Sybase recommends that you use one backup subtype, either selective or non-selective. If you must switch from a non-selective to a selective backup, or vice versa, always take a non-selective full backup before switching to the new subtype, to ensure that you have all changes.

9.6.7 ON_ERROR option [CR #624849]

The ON_ERROR option was deprecated in Sybase IQ 15.0, but reinstated in Sybase IQ 15.1.

Table 2-4 “DBISQL options” in the Sybase IQ 15.1 and Sybase IQ 15.2 documentation states that ON_ERROR is not currently supported by Sybase IQ. This is incorrect.

“Alphabetical list of options” should include this section:

ON_ERROR option [DBISQL]

Function	Controls what happens if an error is encountered while executing statements in Interactive SQL (dbisql).
Allowed values	String. See Description for allowed values.
Default	PROMPT
Description	<p>Controls what happens if an error is encountered while executing statements as follows:</p> <ul style="list-style-type: none"> • STOP – dbisql stops executing statements from the file and returns to the statement window for input. • PROMPT – dbisql prompts the user to see if he or she wants to continue. • CONTINUE – The error displays and dbisql continues executing statements. • EXIT – dbisql terminates. • NOTIFY_CONTINUE – The error is reported, and the user is prompted to press ENTER or click OK to continue. • NOTIFY_STOP – The error is reported, and the user is prompted to press ENTER or click OK to stop executing statements. • NOTIFY_EXIT – The error is reported, and the user is prompted to press ENTER or click OK to terminate Interactive SQL.

When you are executing a *.SQL* file, the values STOP and EXIT are equivalent.

9.7 Using Sybase IQ Multiplex

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

9.7.1 Database file paths for starting multiplex databases [CR #638520]

The following was added to *Multiplex Reference > SQL statement > CREATE MULTIPLEX SERVER > Usage*:

If you plan to use UNIX soft (symbolic) links for server paths, create the soft link before you run CREATE MULTIPLEX SERVER. When you start the new server, the database file path must match the database file path specified when creating that server.

9.7.2 Configuring the IQ Agent on Windows [CR #638823]

In *Managing Multiplex Servers > Using Sybase Central > Running the Sybase IQ Agent > Configuring the Sybase IQ Agent to Start Automatically > Configuring the IQ Agent on Windows*, change “Sybase IQ 15.0” to “Sybase IQ 15.2”.

9.7.3 Replacing All Dbfiles in IQ_SYSTEM_TEMP [CR #633738]

The following was added to *Running Multiplex Transactions > Running DDL Commands > Updating dbspaces in multiplex*.

If a subset of the temporary IQ dbfiles cannot be opened or are damaged, drop each damaged dbfile, then replace it.

If all temporary IQ dbfiles are damaged, follow these steps:

❖ Replacing all dbfiles in IQ_SYSTEM_TEMP

- 1 Start the database with the -iqnotemp parameter.

```
start_iq @params.cfg mpctest.db -iqnotemp 100 -n
mpxnode_c
```

The only temporary file operation allowed on a database while running with -iqnotemp is to drop one or more of the temporary file definitions.

- 2 Drop the last file in IQ_SYSTEM_TEMP:

```
ALTER DBSPACE IQ_SYSTEM_TEMP DROP FILE filename
```

- 3 Stop the database.

- 4 For MULTIPLEX:

- Start the database with no files to empty the checkpoint log:

```
start_iq @params.cfg mpctest.db -iqmpx_ov 1
-iqmpx_sn 1 -n mpxnode_c
```

- Add a “dummy” dbfile to the coordinator. You will need to drop this file and add it again in coordinator mode, because files added in single node mode have a null server ID; they are owned by the coordinator instead of the server that added the file.
- To add more space to IQ_SYSTEM_MAIN with a new file on a raw device, use syntax like the following:

Table 4:

Platform	Command
Windows:	ALTER DBSPACE IQ_SYSTEM_TEMP ADD FILE tempfile '\\.\PhysicalDrive3'
Unix:	ALTER DBSPACE IQ_SYSTEM_TEMP ADD FILE tempfile '/dev/rdsk/c4t600A0B80005A7F5D0000024'

where *tempfile* is the logical or chosen logical name of the new dbfile.

- Stop the server. Now start it without the `-iqnotemp` and `-iqmpx_ov` and `-iqmpx_sn` flags.
- Add a new dbfile in IQ_SYSTEM_TEMP. Because the server is in coordinator mode, the new temporary store dbfile is correct, has a valid server ID, and is used by the coordinator.
- Drop the “dummy” dbfile that you added to the coordinator.

9.7.4 Correction to Start command for Sybase Central on Windows [CR #632128]

In *Managing Multiplex Servers > Using Sybase Central > Starting Sybase Central on Windows*, Step 1 should read:

- 1 To start Sybase Central on Windows, run Sybase > Sybase IQ 15.2 > Sybase Central Java Edition from the Programs menu.

9.7.5 Correction to mpx_mode column description [CR #627055]

In Table A-5, `sp_iqmpxinfo` columns, in *Multiplex Reference > System procedures > sp_iqmpxinfo procedure*, the Description for the `mpx_mode` column erroneously lists ‘writer’ twice.

9.7.6 Multiplex and virtualization [CR #625061]

In *Multiplex Overview > Introduction*, the following note should be included:

Note Sybase IQ Multiplex is not supported in virtualized environments.

9.7.7 Restoring IQ store data in a different location [CR #623585, CR #633738]

See “Restoring IQ store data in a different location [CR #623585, CR #633738]” on page 25 for corrections to a procedure in *Backing Up and Restoring Data in a Multiplex Environment > Backing up and restoring data in a multiplex environment > Restoring in the multiplex environment*.

9.8 Utility Guide

This section contains updates to the *Utility Guide*.

9.8.1 Renaming transaction log with dblog [CR #634508]

The following restriction was removed from *Utility Guide > Database Administration Utilities > dblog Utility*.

- The databases cannot be participating in SQL Remote replication. For Sybase IQ, this means that dblog cannot be used in multiplex environments.

9.8.2 Server option -iqnotemp [CR #633738]

The -iqnotemp server option was updated in *Utility Guide > Running the Database Server > Starting the database server > Table 1-1*. The Description now contains the following:

Creates a temporary file in place of the defined temporary dbspace. *size* is file size in MB. This parameter causes the server to ignore all temporary IQ dbfile definitions when starting a database. You can use -iqnotemp to solve temporary dbfile problems by dropping damaged files and replacing them later.

If you start the utility database server with -iqnotemp when restoring, Sybase IQ ignores all temporary IQ file definitions in the backed up database during the restore operation, including intermediate starts. You can thus restore a database to a different temporary file topology without recreating and using old temporary file definitions.

The only temporary file operation allowed on a database started with `-iqnotemp` is `ALTER DBSPACE IQ_SYSTEM_TEMP DROP FILE filename`.

9.8.3 Show multiple results sets in Interactive SQL -nogui mode [CR #633140]

These topics were added to *Utility Guide > Using Interactive SQL*:

Showing multiple result sets (UNIX)

By default, Interactive SQL shows the first result set of the most-recently executed statement. Use this procedure to see all result sets on UNIX platforms.

❖ Showing multiple results sets (UNIX)

- 1 If running Interactive SQL as a command line program (`-nogui` mode), navigate to your `$HOME` directory and locate the file `.isqlPreferences11`.
 - a Change the line

```
<entry
key="SybaseIQ.showMultipleResultSets">0</entry>
```

to:

```
<entry
key="SybaseIQ.showMultipleResultSets">1</entry>
```
 - b Save your changes to `.isqlPreferences11`.
- 2 If using GUI mode, click `Tools > Options` in the Interactive SQL window. The Interactive SQL Options window appears.
 - a On the Sybase IQ page, select the Results tab.
 - b Choose Show All Result Sets.
 - c Click OK.

Showing multiple result sets (Windows)

By default, Interactive SQL shows the first result set of the most-recently executed statement. Use this procedure to see all result sets on Windows platforms.

Setting the Show Multiple Result Sets option in GUI mode also enables multiple result sets in -nogui mode.

❖ Showing multiple result sets (Windows)

- 1 Click Tools > Options in the Interactive SQL window. The Interactive SQL Options window appears.
- 2 On the SQL Anywhere page, select the Results tab.
- 3 Choose Show all Result Sets.
- 4 Click OK.

9.8.4 Server option -startdir [CR #631259]

The -startdir server option was added to *Utility Guide > Running the Database Server > Starting the database server > Table 1-1*.

- Syntax: -startdir dirname.
- Description: Starts the server in the specified directory. If you use this parameter on Windows, Sybase IQ assumes that each server is being started by the Sybase IQ Agent, and starts the server in the background.

9.9 Installation and Configuration Guide

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

9.9.1 New name for iqunload -n procedure [CR #612677]

The name of the “Correcting invalid SQL” procedure has been changed to “Running iqunload -n.” See *Migrating Data > Avoiding potential migration problems > Correcting invalid SQL syntax*.

9.10 Performance and Tuning Guide

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

9.10.1 Incorrect default value for MAX_HASH_ROWS option [CR #633491]

In *Optimizing Queries and Deletions > Controlling query processing > Setting query optimization options*, the MAX_HASH_ROWS option incorrectly states that the default value is 1,250,000 rows; the correct value is 2,500,000.

9.11 User-Defined Functions Guide

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

9.11.1 -O2 and -Wall compile switches on Linux [CR #640213]

In *Creating and executing user-defined functions > Compiling and linking source code to build dynamically linkable libraries > Linux Switches*, add the -O2 and -Wall switches to the list of compile switches.

Including the -O2 and -Wall switches in the list of compile UDF switches, when compiling C++ applications for building shared libraries on Linux, may decrease computation time.

9.12 Unstructured Data Analytics in Sybase IQ

This section contains updates to *Unstructured Data Analytics in Sybase IQ*.

9.12.1 BFILE LOB data extraction example [CR #621844]

Disregard the example in *Unstructured Data Analytics in Sybase IQ > Large Object Data Load and Unload > Exporting large object data > BFILE function*. The SELECT command extracts data, but the extracted data contains two extraneous columns and cannot be reloaded.

This new example extracts LOB data using the BFILE function and reloads the data.

1 Create table LobA:

```
create table LobA
  (rowid int primary key,
   col1 clob null,
   col2 blob null)
```

Assume LobA has two rows of data.

- 2 Extract the non-LOB data and the paths to the files into which the LOB data is extracted:

```
BEGIN
  SET TEMPORARY OPTION
    Temp_Extract_Name1 = LobA_data.txt';
  SELECT rowid,
    'row' + string(rowid) + '.' + 'col1',
    'row' + string(rowid) + '.' + 'col2'
  FROM LobA;
END
```

The file *LobA_data.txt* is created and contains this non-LOB data and these filenames:

```
1,row1.col1,row1.col2,
2,row2.col1,row2.col2,
```

- 3 Perform the LOB data extraction:

```
SELECT
  BFILE('row' + string(rowid) + '.' + 'col1',col1),
  BFILE('row' + string(rowid) + '.' + 'col2',col2)
FROM LobA;
```

After the extraction, there is a file for each cell of LOB data extracted. For example, if table LobA contains two rows of data with rowid values of 1 and 2, you have these files:

```
row1.col1
row1.col2
row2.col1
row2.col2
```

- 4 Reload the extracted data:

```
LOAD TABLE LobA
(rowid,
 col1 ASCII FILE ('') NULL('NULL'),
 col2 BINARY FILE ('') NULL('NULL'))
FROM LobA_data.txt'
DELIMITED BY ','
ROW DELIMITED BY '\n'
ESCAPES OFF;
```

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 automatically collect as much information as possible. You may also need to manually collect some information.

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

- Sybase IQ version (for example, 15.2 GA or ESD level)
- Type of hardware, amount of memory, number of CPUs*
- Operating system and version (for example, Red Hat Enterprise Linux 5.2 x86-64)*
- 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)
- (Important) Message log file* – by default, named *dbname.iqmsg* and located in the directory where you started the database server.
- Stack trace file for the date and time the problem occurred (if applicable), by default named *stktrc-YYYYMMDD-HHMMSS_#.iq* and located in the directory where you started the database server.*
- Command or query that produced the error

- Query plan* (recorded in *.iqmsg* file)

The query plan is collected automatically by `getiqinfo`. If you manually collect information, enter these commands, then re-run 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 is in the message log file.

If you have performance problems, set this database 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_2/logfiles/<servername>.nnnn.stderr* and *IQ-15_2/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, 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.

See *System Administration Guide: Volume 1 > Troubleshooting Hints > Reporting problems to Technical Support > Collecting diagnostic information using getiqinfo*.

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 Documentation 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 Documentation 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 Documentation Web site, go to Product Documentation at <http://sybooks.sybase.com>.

❖ **Creating a filter to narrow your search of the online documentation**

- 1 Point your Web browser to Sybooks Online Infocenter Help at <http://infocenter.sybase.com/help/index.jsp>.
- 2 Click Create/Edit search filter.
- 3 Select Search within a user-defined search filter in the Select search filter dialog box.
- 4 Click Create New Filter.
- 5 Enter a name for your search filter in the New Search Filter list dialog box. For example, `IQ_15.2_Install_Guides`.
- 6 Scroll down to your product and select the product or expand the list of documentation for this product in the Select product(s) list. For example, scroll down to Sybase IQ 15.2 and click the plus sign.
- 7 Select the documents to search and click OK. For example, select the five Installation and Configuration Guides for the `IQ_15.2_Install_Guides` filter.
- 8 Select the filter name and click OK in the Select search filter dialog box. The name of the search filter displays in the Current search filter drop-down box and is used for your next search.

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 15.2 Network Clients and ODBC kits available

Sybase IQ 15.2 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 15.2.

11.4 IQ Newsgroup

Sybase newsgroups allow 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 Users' Group

The IQ Users' Group lets Sybase IQ users exchange information about Sybase IQ. You can visit the IQUG homepage at <http://iqug.dssolutions.com/>. To have your name added to the IQ Users' Group list, send an e-mail message 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 15.2 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 Documentation 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

Run the `getiqinfo` script to automatically collect much of this information.

Information requested	Value
Sybase IQ version (for example 15.2 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, Red Hat Enterprise Linux 5.2 x86-64)	
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.nmnn.srvlog</code> and <code>server.nmnn.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	

