SYBASE®

New Features Guide

Sybase IQ

15.0

Contents

About This Book

Subject   This book describes new features in Sybase® IQ 15.0.

Audience   This book is for users of previous versions of Sybase IQ who want to find out what is new and different in this release of the software.

Related Sybase IQ documents   The Sybase IQ 15.x documentation set includes:

- *Release Bulletin* provides information about last-minute changes to the product and documentation.

- *Installation and Configuration Guide* provides platform-specific instructions on installing, migrating to a new version, and configuring Sybase IQ for a particular platform.

- *Advanced Security in Sybase IQ* covers the use of user encrypted columns within the Sybase IQ data repository. You need a separate license to install this product option.

- *Error Messages* lists Sybase IQ error messages referenced by Sybase error code, SQLCode, and SQLState, and SQL preprocessor errors and warnings.

- *IMSL Numerical Library User's Guide: Volume 2 of 2 C Stat Library* contains a concise description of the IMSL C Stat Library time series C functions. This book is only available to RAP – The Trading Edition™ Enterprise users.

- *Introduction to Sybase IQ* includes hands-on exercises for those unfamiliar with Sybase IQ or with the Sybase Central™ database management tool.

- *Large Objects Management in Sybase IQ* explains storage and retrieval of Binary Large Objects (BLOBs) and Character Large Objects (CLOBs) within the Sybase IQ data repository. You need a separate license to install this product option.

- *New Features in Sybase IQ 15.0* documents new features and behavior changes for version 15.0.

- *New Features Summary Sybase IQ 15.1* summarizes new features and behavior changes for the current version.

- *Performance and Tuning Guide* describes query optimization, design, and tuning issues for very large databases.

- *Quick Start* lists steps to build and query the demo database provided with Sybase IQ for validating the Sybase IQ software installation. Includes information on converting the demo database to multiplex.

- *Reference Manual* – Includes two reference guides to Sybase IQ:

  - *Reference: Building Blocks, Tables, and Procedures* describes SQL, stored procedures, data types, and system tables that Sybase IQ supports.

  - *Reference: Statements and Options* describes the SQL statements and options that Sybase IQ supports.

- *System Administration Guide* – Includes two volumes:

  - *System Administration Guide: Volume 1* describes startup, connections, database creation, population and indexing, versioning, collations, system backup and recovery, troubleshooting, and database repair.

  - *System Administration Guide: Volume 2* describes writing and running procedures and batches, programming with OLAP, accessing remote data, setting up IQ as an Open Server, scheduling and event handling, programming with XML, and debugging.

- *User-Defined Functions Guide* provides information about the user-defined functions, their parameters, and possible usage scenarios.

- *Using Sybase IQ Multiplex* tells how to use multiplex capability, designed to manage large query loads across multiple nodes.

- *Utility Guide* provides Sybase IQ utility program reference material, such as available syntax, parameters, and options.

Other sources of information

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

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

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

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

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

- The Sybase Product Manuals Web site is an online version of the SyBooks CD that you can access using a standard Web browser. In addition to product manuals, you will find links to EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

  To access the Sybase Product Manuals Web site, go to Product Manuals at http://sybooks.sybase.com.

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://certification.sybase.com/ucr/search.do.

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.

❖ **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 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.

Sybase EBFs and software maintenance

❖ **Finding the latest information on EBFs and software maintenance**

1   Point your Web browser to the Sybase Support Page at http://www.sybase.com/support.

2   Select EBFs/Maintenance. If prompted, enter your MySybase user name and password.

3   Select a product.

4   Specify a time frame and click Go. A list of EBF/Maintenance releases is displayed.

    Padlock icons indicate that you do not have download authorization for certain EBF/Maintenance releases because you are not registered as a Technical Support Contact. If you have not registered, but have valid information provided by your Sybase representative or through your support contract, click Edit Roles to add the "Technical Support Contact" role to your MySybase profile.

5   Click the Info icon to display the EBF/Maintenance report, or click the product description to download the software.

Syntax conventions   This documentation uses the following syntax conventions in syntax descriptions:

•   **Keywords**   SQL keywords are shown in UPPER CASE. However, SQL keywords are case insensitive, so you can enter keywords in any case you wish; SELECT is the same as Select which is the same as select.

•   **Placeholders**   Items that must be replaced with appropriate identifiers or expressions are shown in *italics*.

•   **Continuation**   Lines beginning with ... are a continuation of the statements from the previous line.

- **Repeating items**   Lists of repeating items are shown with an element of the list followed by an ellipsis (three dots). One or more list elements are allowed. If more than one is specified, they must be separated by commas.

- **Optional portions**   Optional portions of a statement are enclosed by square brackets. For example:

```
RELEASE SAVEPOINT [ savepoint-name ]
```

It indicates that the *savepoint-name* is optional. The square brackets should not be typed.

- **Options**   When none or only one of a list of items must be chosen, the items are separated by vertical bars and the list enclosed in square brackets. For example:

```
[ ASC | DESC ]
```

It indicates that you can choose one of ASC, DESC, or neither. The square brackets should not be typed.

- **Alternatives**   When precisely one of the options must be chosen, the alternatives are enclosed in curly braces. For example:

```
QUOTES { ON | OFF }
```

It indicates that exactly one of ON or OFF must be provided. The braces should not be typed.

Typographic conventions

Table 1 lists the typographic conventions used in this documentation.

Table 1: Typographic conventions

Item	Description
Code	SQL and program code is displayed in a mono-spaced (fixed-width) font.
User entry	Text entered by the user is shown in bold serif type.
emphasis	Emphasized words are shown in italic.
file names	File names are shown in italic.
database objects	Names of database objects, such as tables and procedures, are shown in bold, san-serif type in print, and in italic online.

The demo database

Sybase IQ includes scripts to create a demo database (*iqdemo.db*). Many of the queries and code samples in this document use the demo database as a data source.

The demo database contains internal information about a small company (employees, departments, and financial data), as well as product (products), and sales information (sales orders, customers, and contacts).

See the Sybase IQ installation guide for your platform or talk to your system administrator for more information about the demo database.

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 15.0 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.

The online help for the Sybase IQ plug-in for Sybase Central, which you can navigate using a screen reader, describes accessibility features, including Sybase Central keyboard shortcuts.

Configuring your accessibility tool

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 UPPER CASE TEXT as initials, and Mixed Case Text as words. You might find it helpful to configure your tool to announce syntax conventions. Consult the documentation for your tool.

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.

If you need help

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.

C H A P T E R   1

New Features in Sybase IQ 15.0

About this chapter

This chapter describes the new features introduced in Sybase IQ 15.0. Features are organized according to the themes shown in the chapter table of contents.

Note For new features documented in 12.7 ESD releases, see the latest version of the *Release Bulletin* in the Sybooks Online Help at http://infocenter.sybase.com/help/index.jsp.

Contents

High performance analytics

Enhancements to loads and queries improve product performance. For a list of new options and new option settings that may provide additional performance benefits if appropriate for your environment, see "Database option changes" on page 66.

Direct loading of data from clients

Sybase IQ now supports bulk loading of remote data via the LOAD TABLE USING CLIENT FILE statement. LOAD TABLE USING FILE loads data on the local server, as does the LOAD TABLE FROM syntax, which is deprecated in this release. LOAD TABLE USING FILE also replaces the deprecated utility iq_bcp.

Both server and client need to be at Sybase IQ 15.0 level. A combination of 15.0 server and 12.7 client returns a File I/O Error.

For syntax, see LOAD TABLE statement in Chapter 1, "SQL Statements," in *Reference: Statements and Options*.

Improved loading for large single (fact) tables

In order to meet the challenges of the exponential growth of information and the demand for real-time access to data, Sybase IQ has significantly improved the performance of loading High_Group (HG) indexes and Containment (also called WORD) (WD) indexes, while transactions are still allowed to access the tables being loaded. This load performance improvement for HG and WD indexes affects:

- INSERT...SELECT

- INSERT...LOCATION

- LOAD

- UPDATE

- CREATE INDEX

- Updatable cursors

The LOAD TABLE statement now performs parallel loads of HG and WD indexes, thus executing faster than in previous releases.

IQ also supports loading partitioned tables. Table partitioning is a separately licensed Sybase IQ feature.

INSERT...LOCATION supports 2GB large object (LOB) data

By default, if a remote data column contains over 2GB, Sybase IQ silently truncates the column value to 2GB.

Users must be specifically licensed to use the Large Objects Management functionality. For details on the Large Objects Management option, see *Large Objects Management in Sybase IQ*.

LOAD TABLE QUOTES clause example

The LOAD TABLE STRIP ON keyword has been deprecated in this release. The following example shows the results of a query using the LOAD TABLE STRIP RTRIM keyword instead of the STRIP ON keyword.

LOAD TABLE QUOTES option example   Consider a table defined as:

```
CREATE TABLE t1 (c1 INT, c2 VARCHAR(20), c3 VARCHAR(20))
```

with the following input data:

```
1, apple , fruit1 ,
2, "banana" , "fruit2",
3, " pear ", " fruit3 ",
```

The result of loading this data is displayed by running the following query:

```
SELECT c1, c2, c3, LENGTH(c2), LENGTH(c3) FROM t1
```

Given the values of the QUOTES and STRIP options used during the LOAD TABLE command, the following table displays the result of the query above, with each result enclosed by '<' and '>':

LOAD TABLE options		Results of SELECT c1, c2, c3, LENGTH(c2), LENGTH(c3) FROM t1				
QUOTES	**STRIP**	**c1**	**c2**	**c3**	**length(c2)**	**length(c3)**
ON	RTRIM	<1>	<apple>	<fruit1>	<5>	<6>
		<2>	<banana>	<fruit2>	<6>	<6>
		<3>	< pear >	< fruit3 >	<6>	<8>
ON	OFF	<1>	<apple >	<fruit1 >	<6>	<7>
		<2>	<banana>	<fruit2>	<6>	<6>
		<3>	< pear >	< fruit3 >	<6>	<8>
OFF	RTRIM	<1>	< apple>	< fruit1>	<6>	<7>
		<2>	< "banana">	< "fruit2">	<9>	<9>
		<3>	< " pear ">	< " fruit3 ">	<9>	<11>
OFF	OFF	<1>	< apple >	< fruit1 >	<7>	<8>
		<2>	< "banana" >	< "fruit2">	<10>	<9>
		<3>	< " pear ">	< " fruit3 ">	<9>	<11>

Notes on the results:

- With QUOTES ON and STRIP RTRIM, both leading space and trailing space for the non-enclosed field c2 row 1 are trimmed.

- With QUOTES ON and STRIP OFF, only the leading space for the non-enclosed field c2 row 1 is trimmed.

- With QUOTES OFF and STRIP RTRIM, only the trailing space for the non-enclosed field c2 row 1 is trimmed.

- With QUOTES OFF and STRIP OFF, neither leading space nor trailing space for the non-enclosed field c2 row 1 is trimmed.

- With QUOTES ON and STRIP RTRIM, both leading space and trailing space within quotes for the enclosed fields c2 and c3 row 3 are NOT trimmed.

Loading partitioned tables

There are a few special considerations for loading partitioned tables:

- When loading a partitioned table, the best performance is achieved when the partitioning column is the first column in the column list of the statement.

  List partitioning columns before large object (LOB) columns in the SELECT statement clause of an INSERT...LOCATION statement and load data from a primary file, if possible. The data in the primary file should be rearranged using a pre-load process, if possible.

- The START ROW ID clause of the LOAD TABLE and INSERT statements is not allowed on a partitioned table. The following error is reported and a rollback is performed on the load operation:

  ```
  "Option START ROW ID not allowed on a partitioned
  table."
  ```

  (SQLCODE -1009416L, SQLSTATE QCB14, Sybase error code 21054)

- Partial width inserts are not recommended, as the START ROW ID clause of the INSERT statement is not supported on a partitioned table.

  Be sure to include the partition key column of the partitioning columns of the table in the column list of the load operation and leave all non-specified columns as NULL, if you do perform a partial width insert. If the partition key column is omitted from the column list, the following error is reported:

  ```
  "Operation not allowed — partition key column %2 not
  specified."
  ```

(SQLCODE -1009418L, SQLSTATE QCB16, Sybase error code 21056)

where %2 is the name of the partition key.

- The APPEND_LOAD database option behaves differently for partitioned and non-partitioned tables. Row ID ranges are assigned to each partition in a partitioned table. For partitioned tables, when APPEND_LOAD is ON, new rows are appended at the end of the appropriate partition. When APPEND_LOAD is OFF, the load reuses the first available row IDs and space from deleted rows of the appropriate partition.

  For non-partitioned tables, when APPEND_LOAD is ON, new rows are added after the maximum row ID that is at the end of the table rows. When APPEND_LOAD is OFF, the load reuses the deleted row IDs. With non-partitioned tables, you can also control where rows are inserted by using the LOAD or INSERT START ROW ID clause to specify the row at which to start inserting.

- An attempt to update the contents of a partitioning column results in the following error:

      "Updating partition key column on a partitioned
      table is not allowed."

  (SQLCODE -1009417L, SQLSTATE QCB15, Sybase error code 21055)

Improved multiplex configuration and concurrent load support

Multiplex functionality is significantly updated and enhanced for Sybase IQ 15.0.

A new book, *Using Sybase IQ Multiplex*, consolidates all Sybase IQ multiplex information formerly in the IQ documentation set and the new multiplex features documentation Revision 01 of the *New Features Guide*.

Using Sybase IQ Multiplex includes these topics:

Chapter/Appendix	Description	Notes
Chapter 1, "Multiplex Overview"	Introduces multiplex functionality.	Formerly Chapter 3, "Multiplex Overview," in Revision 01 of *New Features Guide*
Chapter 2, "Managing Multiplex Servers"	Describes how to create, configure, and administer multiplex servers	Formerly Chapter 4, "Managing Multiplex Servers," in first revision of *New Features Guide*
Chapter 3, "Running Multiplex Transactions"	Describes behavior of DML and DDL operations on multiplex servers	Formerly Chapter 5, "Running Multiplex Transactions," in first revision of *New Features Guide*.
Chapter 4, "Monitoring with Sybase Central"	Describes the enhanced graphical user interface in the IQ Plug-in that supports multiplex visualization and performance monitoring	Formerly Chapter 6, "Performance Monitoring with Sybase Central" in first revision of *New Features Guide*.
Chapter 5, "Backing Up and Restoring Data in a Multiplex Environment"	Describes how to back up Sybase IQ data in a multiplex.	Updated information previously in *Sybase IQ System Administration Guide*.
Appendix A, "Multiplex Reference"	Lists the syntax for SQL statements, system objects, and options that support multiplex capability.	Multiplex-specific subset of syntax formerly in Appendix A, "New and Changed Syntax Reference" in first revision of *New Features Guide*.

Improved query performance

Enhanced support for parallel queries and the ability to choose subquery predicate processing behavior enable experienced DBAs to tune query performance.

Improved tokenization

Tokenization can now be applied to columns with more than 64K distinct values.

When you create a permanent table in a Sybase IQ database, IQ stores all column values in a default index. This default index, called an **FP** (fast projection) index, optimizes projections and enables certain kind of search conditions to be evaluated.

Each column has one FP index, and each FP is an array of n fixed-length entries where n is the number of rows in the table. Each column value is stored sequentially in ascending RecordID order.

With a small number of distinct or unique values, such as a state, date, or month field, an optimized form of the FP can be created that will reduce the number of disk pages required, dramatically reducing both the storage required for a column and I/O costs for projection.

These optimized FP indexes have two pieces: 1. a lookup table where each distinct value in the column appears exactly once and 2. the logical array of column cell values, except instead to storing the actual cell values each element of the logical array is a key into the lookup table to where the cell value is stored.

The sp_iqindexmetadata stored procedure generates a report describing a specified index or indexes belonging to a specified owner or table. The output allows easy checking of whether a given index is a 1-byte, 2-byte, 3-byte, or flat style FP index. For details, see "sp_iqindexmetadata procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.

FP(1) index

When the number of values is within 256, a 1-Byte FP index is created. The FP index starts with this form. The actual key value of each row is stored in the lookup table and the FP entry (lookup key) contains the index into the lookup table. For a 1-Byte FP index, each FP entry occupies 1 byte.

FP(2) index

When the number of values exceeds 256, but is less than or equal to 65536, the same lookup table grows. Each FP entry (lookup key) contains the index into the lookup table but occupies 2 bytes instead of 1.

To avoid the overhead of converting 1-byte entries into 2-byte entries, users can specify the IQ UNIQUE value to be greater than 256 and less than or equal to 65536 at table creation time.

FP(3) index

When the distinct count exceeds 65536, you can create a 3-byte FP index. The FP(3) index is structurally similar to FP(1) and FP(2) indexes, with the following key differences:

• The maximum size of the FP(3) lookup table is 16777216, not 65536.

• The FP(3) index buffer storage contains lookup keys of 3 bytes each.

The 3-byte index stores values in a column (column data), as long as the distinct count does not exceed 16777216. Users can create a 3-byte index on columns only if the size of column data is greater than 3 bytes. Therefore, you cannot create an FP(3) index on columns with data types BIT, TINYINT, SMALLINT, CHAR(<=3), VARCHAR(<=3), BINARY(<=3) and VARBINARY(<=3). Sybase IQ also does not support FP(3) indexes for LONG VARCHAR and LONG VARBINARY data types.

To create an FP(3) index, either the MINIMIZE_STORAGE must be set ON, or the column must have been created with an IQ UNIQUE constraint value between 65537 and 16777216, including these two values. An FP(3) may also be created after rollover from an FP(2) index, once the unique count exceeds 65536, but only for data sizes shown in Table 1-1 and excluding data types BIT, TINYINT, SMALLINT, CHAR(<=3), VARCHAR(<=3), BINARY(<=3) and VARBINARY(<=3). See "Using IQ UNIQUE constraint on columns" in Chapter 9, "Ensuring Data Integrity," in the *System Administration Guide: Volume 1* and the "MINIMIZE_STORAGE option" in Chapter 2, "Database Options," in the *Reference: Statements and Options*.

Behavior changes   There are some differences in the behavior of FP indexes from earlier releases.

Table 1-1: Sybase IQ fast projection (FP) indexes

Distinct count	Column data size = 1 byte	Column data size = 2 bytes	Column data size = 3 bytes	Column data size >3 bytes
<257	FP(1)	FP(1)	FP(1)	FP(1)
257 - 65536	—	FP(2)	FP(2)	FP(2)
65537-16777216	—	—	Flat	FP(3)
>16777216	—	—	—	Flat

Like FP(1) and FP(2) indexes, the FP(3) index is not supported for columns whose data type is wider than 255 bytes or less than four bytes wide.

The creation of an FP(3) index, overflow, or a forced transition to an FP(3) index is permitted only if the space used by the lookup table is less than the current value of the FP_LOOKUP_SIZE option and less than the portion of the main cache specified by the current setting of FP_LOOKUP_SIZE_PPM.

The maximum number of lookup pages used in Sybase IQ is controlled by the FP_LOOKUP_SIZE option and the FP_LOOKUP_SIZE_PPM option, whichever is lower. See "FP_LOOKUP_SIZE option" and "FP_LOOKUP_SIZE_PPM option" in Chapter 2, "Database Options," in *Reference: Statements and Options*.

Table 1-2 calculates the maximum number of distinct values that can be supported in an FP(3) index based on the following formula:

```
FP_LOOKUP_SIZE / (Column-Data size + Cardinality size)
```

Cardinality size is the space reserved to store cardinality of all individual data in the lookup store. Cardinality size can have a value of either 4 or 8 bytes. In this example, it has a maximum value of 8 bytes.

Table 1-2: Maximum unique values in FP(3)

FP_LOOKUP_SIZE (MB)	Column data type width (bytes)					
	4	8	32	64	128	255
1 MB	87381	65536	26214	14563	7710	3986
4 MB	349525	262144	104857	58254	30840	15947
8 MB	699050	524288	209715	116508	61680	31895
16 MB	1398101	1048576	419430	233016	123361	63791
32 MB	2796202	2097152	838860	466033	246723	127583
64 MB	5592405	4194304	1677721	932067	493447	255166
128 MB	11184810	8388608	3355443	1864135	986895	510333
256 MB	16777216	16777216	6710886	3728270	1973790	1020667

Notes
The values illustrated in Table 1-2 are estimates for the number of unique values in a column for the given value of option FP_LOOKUP_SIZE; actual values may vary. Such variations are possible because counts can be stored as 4 bytes or 8 bytes.

Table 1-2 is based on the condition that the value of FP_LOOKUP_SIZE is less than or equal to the value of FP_LOOKUP_SIZE_PPM.

Configuring FP(3) indexes

You may need to adjust the temporary cache size when configuring 3-byte indexes. You can set values using the server startup command line parameter -iqtc or using the sa_server_option system procedure temp_cache_memory_mb option as follows:

```
CALL sa_server_option('temp_cache_memory_mb', value)
```

The enumerated FP indexes use a hash object to manage the values represented in the column. The size of the hash object used with a 3-byte FP can get large, depending on the number of distinct values and the width of the column. With a large enough temporary cache allocation, increasing the value of the option HASH_PINNABLE_CACHE_PERCENT above the default value of 20 percent can improve performance by allowing the entire hash object to remain in the cache.

Cache usage   In order to maximize the use of FP(3) indexes, set the FP_LOOKUP_SIZE option a value larger than the default of 16MB. Refer to Table 1-2 for maximum distinct counts allowed on a column for an FP(3) index. Table 1-2 also contains examples, with values less than 16777216, where a rollover to a Flat FP occurs for a smaller unique count than the expected 16777216.

Loads   Columns with a 3-byte index have additional cache requirement for loading data. Set FP_LOOKUP_SIZE to an appropriate value before loading columns with 3-bytes indexes.

If a scarcity of pinned buffers occurs, Sybase IQ returns a warning in the *.iqmsg* file, which also contains notification of possible thrashing:

```
Warning: Hash Insert forced buffer unpinning detected
for FP Index
Warning: Hash Insert thrashing detected for FP Index
```

Flat FP index

When the number of distinct values exceeds 16777216, no lookup table is created. Each FP entry contains an actual column cell value.

If MINIMIZE_STORAGE is on, you can avoid the overhead of converting lookup FP entries into flat style. When the distinct row count for a particular field increases beyond 16777216, then an FP(3) index is automatically converted to a flat style FP index. Specify the IQ UNIQUE value to be greater than 16777216 at table creation time to create flat style FP.

Note When you create a table with the DATE data type, a 2-byte FP index is created on the DATE field, which is independent of the settings in database option MINIMIZE_STORAGE.

If you want to create a 3-byte FP or flat style FP index on the DATE field, use the following IQ UNIQUE values when creating the table:

- For a 3-byte FP index— IQ UNIQUE should be between 65537 and 16777216.

- For a flat style FP index— IQ UNIQUE should be higher than 16777216.

Enhanced use of HG indexes

Use multicolumn HG indexes to enhance the performance of ORDER BY queries with reference to multiple columns. This change is transparent to users, but improves query performance.

Queries with multiple columns in the ORDER BY clause run faster using multicolumn HG indexes. For example, if the user has multicolumn index HG(x,y,z) on table T, then this index is used for ordered projection:

```
SELECT abs (x) FROM T
ORDER BY x, y
```

In the above example, the HG index vertically projects *x* and *y* in sorted order.

If the ROWID() function is in the SELECT list expressions, multicolumn HG indexes are also used. For example:

```
SELECT rowid()+x, z FROM T
ORDER BY x,y,z
```

If ROWID() is present at the end of an ORDER BY list, and if the columns of that list—except for ROWID()— exist within the index, and the ordering keys match the leading HG columns in order, multicolumn indexes are used for the query. For example:

```
SELECT z,y FROM T
ORDER BY x,y,z,ROWID()
```

Enhanced parallelism within queries

Earlier versions of Sybase IQ had limited ability to utilize many CPUs while running a single query. Sybase IQ 15.0 significantly increases the types of operators that are automatically parallelized by the Sybase IQ query engine. This feature is enabled by default and requires no change in query syntax or specific tuning.

Improved subquery performance

Subquery flattening is an optimization technique in which the Sybase IQ optimizer rewrites a query containing a subquery into a query that uses a join. The new database options SUBQUERY_FLATTENING_PREFERENCE and SUBQUERY_FLATTENING_PERCENT control under what circumstances the optimizer chooses to use this optimization.

The FLATTEN_SUBQUERIES option has been deprecated in Sybase IQ version 15.0.

Using caching methods

A correlated subquery contains references to one or more tables outside of the subquery and is re-executed each time the value in the referenced column changes. Sybase IQ 15.0 allows users to choose caching methods for executing the correlated subquery.

See "SUBQUERY_CACHING_PREFERENCE option" in Chapter 2, "Database Options," in *Reference: Statements and Options*.

Reduce temp space use for queries

Sybase IQ 15.0 has improved the in-memory compression of some transient structures (hashes, sorts) for many queries improving performance and decreasing resource use.

Database collation improvements

Support of database collations has changed in Sybase IQ 15.0 to take advantage of the space efficiency and speed of the SQL Anywhere Collation Algorithm.

- The new database option SORT_COLLATION allows implicit use of the SORTKEY function on ORDER BY expressions. When the value of this option is set to a valid collation name or collation ID, any string expression in the ORDER BY clause is treated as if the SORTKEY function has been invoked.

  See "SORT_COLLATION option" in Chapter 2, "Database Options," in *Reference: Statements and Options*.

- The SORTKEY function now uses the International Components for Unicode (ICU) library, instead of the Sybase Unicode Infrastructure Library (Unilib®) and has new syntax.

  For more information, see "SORTKEY function [String]" in Chapter 4, "SQL Functions," of *Reference: Building Blocks, Tables, and Procedures*.

- The new CREATE DATABASE parameter COLLATION supports specification of a collation for a database.

  See CREATE DATABASE statement in Chapter 1, "SQL Statements," in *Reference: Statements and Options*.

- The CP874toUTF8 utility converts data in the CP874 character set into UTF8 collation, a collation supported by Sybase IQ for the Thai language. The CP874toUTF8 utility now calls the International Components for Unicode (ICU) library to perform data conversion. You can also load data in the CP874 character set without converting it to UTF8 using this utility.

  See "CP874toUTF8 utility" in Chapter 3, "Database Administration Utilities," of*Utility Guide*.

See "Understanding collations" at http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dbadmin_en11/supplied-collations-choosing-natlang.html in *SQL Anywhere Server - SQL Usage*.

Information lifecycle management

Partitioning and dbspace management allow administrative operations (data placement, dbcc, backup, restore) to be performed at a finer granularity than at the table or database level.

Tablespace is a well known term for storage space management in databases. In Sybase IQ 15.0, the term dbspace management is synonymous with tablespace management.

Benefits of dbspace management and table partitioning

In Sybase IQ 15.0, dbspace management and range table partitions:

- Provide data placement capability
- Provide hierarchical storage management by supporting relocation of less critical data to cheaper storage

Dbspace management and range table partitions improve maintainability and availability by:

- Supporting dbspace read-only (RO) vs. read-write (RW) status
- Supporting dbspace online vs. offline status
- Shortening backup/restore by allowing backing up or restoring one or more RO dbspaces and/or files, or all RW files
- Supporting data validation on a dbspace or a table partition target. See "sp_iqcheckdb procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.
- Allowing Sybase IQ servers to start with unavailable and/or non-usable dbspaces marked "offline" except for the catalog store and the Sybase IQ system dbspace.

Range partitioning

Partitioning is a scheme of dividing large objects into subobjects, for example:

- Storage space is partitioned into tablespaces
- Tables can be partitioned into table partitions

Table objects can be table partitions, columns, indexes, IQ base tables, join indexes, unique constraints, primary keys or foreign keys.

A **table partition** is a collection of rows that is a subset of a user-created table. A given row cannot be placed in two different partitions. Each partition can be placed in its own dbspace and managed individually. A partition shares its parent table's logical attributes:

- Column definition with the same integrity constraints and defaults

- The same referential integrity constraints

- The same unique and primary key constraints

- The same check constraints

The table creator chooses whether to partition a table, how to partition it, and the number of partitions. The table creator defines the **partition key**, a table column that determines how a table should be partitioned. For syntax see CREATE TABLE statement and ALTER TABLE statement in Chapter 1, "SQL Statements," in *Reference: Statements and Options*.

A fundamental administration concept of partitioning is the ability to make a subobject read-only. Once a subobject is set read-only, validated, and backed up, it needs minimal maintenance.

Partitioning improves administration and runtime operations, particularly backup, restore, and database validation, by organizing storage and data according to business requirements.

Sybase IQ 15.0 **range partitioning** divides a table into logical partitions based on the values of a single table column. Only base tables can be partitioned; global temporary tables or declared local temporary tables cannot. All rows of a table partition are physically colocated, and the user must name each partition.

Sybase IQ supports a maximum of 1024 partitions for range partitioning.

The most common form of range partitioning is to partition the table by date; for example, June_2009, July_2009, and so on. A range table partition may be assigned to a dbspace.

DDL operations on partitions

You can drop, rename, partition, unpartition, merge, split, and move table partitions, with the following restrictions:

Operation	Restriction
Drop	You cannot drop a partition key column or the last partition of a partitioned table.
Rename	N/A
Partition an unpartitioned table	N/A
Merge two adjacent partitions	Both partitions must reside in the same dbspace. No data movement is required.

Operation	Restriction
Split a partition	All rows must belong to one of the two partitions after splitting. Split partition must be on same dbspace as original so that no data movement is required.
Move a partition to a new dbspace.	All rows of the partition are moved to data pages in the new dbspace. CREATE permission in the new dbspace is required.

Partitioned tables cannot participate in a join index.

DML operations on partitions

You can perform DML operations including load, insert, delete, truncate, and truncate table partition for range partitioning. Update is supported except updating the partition key column.

Sybase IQ returns an error for DML operations on a read-only table or read-only table partition. Insert and load statements or insert by updatable cursor generate an error and operations roll back, if the given row does not fit into the specified range of partitions.

The START ROW ID option for load on a partitioned table is disallowed, and partial width loads and inserts are not supported for the partitioned table. The APPEND_LOAD option for loads on a partitioned table appends new rows to the end of the appropriate partition.

For more considerations, see "Loading partitioned tables" on page 4.

Configurable tablespaces

A **tablespace** is a unit of storage within the database that you can administer as a logical subset of the total storage. Individual objects and subobjects may be allocated to individual tablespaces. Each database includes multiple tablespaces.

Dbspaces, dbfiles, and stores

In Sybase IQ 15.0, a **dbspace** is a tablespace that consists of one or more operating system files.

The meaning of the term varies according to the product version you are using. Sybase IQ 12.7 implemented one-to-one mapping between a dbspace and a database file. With DEFAULT_DISK_STRIPING option 'ON', Sybase IQ automatically distributed data across all available dbspaces for optimal performance and ease of administration. Users could not specify in which dbspace a table or index should be created.

The term **dbfile**, with a corresponding logical filename and physical file path, refers to each operating system file.

Each dbspace name, dbfile name, and physical file path must be unique. The dbfile name can be the same as the dbspace name.

A **store** is one or more dbspaces that store persistent or temporary data for a special purpose. Sybase IQ has three stores:

- The **catalog store** contains the SYSTEM dbspace and up to twelve additional catalog dbspaces.

- The **IQ main store** contains the IQ_SYSTEM_MAIN dbspace and other user dbspaces.

- The **IQ temporary store** contains the IQ_SYSTEM_TEMP dbspace.

Types of dbspaces

There are six types of dbspaces, each designed to store a particular type of Sybase IQ data:

Dbspace type	Data stored	Files contained by dbspace	Number of dbspaces
The SYSTEM dbspace	System tables, views, stored procedures, SQL Anywhere tables, and function definitions	One	One or more
Other catalog dbspaces	SQL Anywhere tables	One	One or more

Dbspace type	Data stored	Files contained by dbspace	Number of dbspaces
IQ_SYSTEM_MAIN	IQ database structures including IQ rollforward/rollback data for each committed transaction and each active checkpointed transaction, the incremental backup metadata, and database space and identity metadata. IQ user objects may be stored here but Sybase recommends that you put them in other main dbspaces.	One or more	One or more
Other main dbspaces (also called user dbspaces)	IQ objects such as tables, indexes, join indexes, and table metadata.	One or more	One or more
IQ_SYSTEM_TEMP	Set of 1 to n temporary dbfiles that define a single temporary dbspace for a standalone database or multiplex node	One or more	One
IQ_SYSTEM_MSG	External file that logs messages about database activity	One per multiplex node	One

Catalog store

The SYSTEM dbspace   The IQ catalog dbspace named SYSTEM contains metadata for your IQ database, stored in the same format as tables in a SQL Anywhere relational database system. SQL Anywhere is a relational database system that can exist with or without IQ. You may have SQL Anywhere style tables in your catalog store along with your IQ tables. Each catalog dbspace contains exactly one file.

Other catalog dbspaces   You may create SQL Anywhere tables in a separate dbspace from the SYSTEM dbspace.

IQ_SYSTEM_MAIN dbspace

The IQ_SYSTEM_MAIN dbspace is created at database creation or when you upgrade an older IQ database to Sybase IQ 15.0. IQ_SYSTEM_MAIN is a special dbspace that contains structures necessary for the database to open: the IQ checkpoint log, IQ rollforward/rollback data for each committed transaction and each active checkpointed transaction, the incremental backup metadata, and database space and identity metadata. IQ_SYSTEM_MAIN is always online when the database is open.

Additional files may be added to IQ_SYSTEM_MAIN. In Sybase IQ 15.0, IQ_SYSTEM_MAIN would normally be relatively small compared to the total main IQ store space and may be created as a file system file, even if other dbspaces and files are raw devices.

For guidelines for the sizing of IQ_SYSTEM_MAIN, see Table 1-3 on page 21. Do not include reserve space in the total database size. Multiplex users should also see "Updating the IQ_SYSTEM_MAIN dbspace" in Chapter 3, "Running Multiplex Transactions," in *Using Sybase IQ Multiplex*

The IQ_SYSTEM_MAIN dbspace cannot be made read-only.

Other user main dbspaces

The best practice is to *avoid* placing user tables or indexes in IQ_SYSTEM_MAIN. The administrator may allow user tables to be created in IQ_SYSTEM_MAIN, especially if these tables are small, very important tables. However, the recommended method is that immediately after creating the database, the administrator creates a second main dbspace (a user main dbspace), revokes CREATE privilege IN DBSPACE IQ_SYSTEM_MAIN from PUBLIC, grants CREATE IN DBSPACE for the new main dbspace to selected users or PUBLIC, and sets PUBLIC.DEFAULT_DBSPACE to the new user main dbspace.

For example:

```
CREATE DBSPACE user_main USING FILE user_main
'user_main1' SIZE 10000;
GRANT CREATE ON user_main TO PUBLIC;
REVOKE CREATE ON IQ_SYSTEM_MAIN FROM PUBLIC;
SET OPTION PUBLIC.DEFAULT_DBSPACE = 'user_main';
```

IQ temporary dbspace

A single dbspace for the IQ temporary store, IQ_SYSTEM_TEMP, is created when you create a database or upgrade an older IQ database Sybase IQ 15.0

Each IQ dbspace may contain any number of files. The only limit is that the total number of IQ files is 16384.

IQ message file dbspace

IQ_SYSTEM_MSG is a system dbspace that points to the file path of the database IQ message log file. IQ_SYSTEM_MSG is not considered a store because it doesn't store any data.

The IQ_SYSTEM_MSG dbspace has one file per multiplex node. By default, the physical file name for the message file on a simplex server or a coordinator of a multiplex is *<dbname>.iqmsg*. The physical file name for the IQ message file on a secondary node in a multiplex is *<servername>.iqmsg*.

IQ_SYSTEM_MSG is not an IQ store dbspace, so ALTER commands such as READONLY and OFFLINE do not apply to IQ_SYSTEM_MSG.

To change the path, use the command ALTER DBSPACE IQ_SYSTEM_MSG RENAME '*filepath*'. The name change takes effect when the server is restarted.

You can also change the logical filename with ALTER DBSPACE IQ_SYSTEM_MSG ALTER FILE *filename* RENAME TO *newfilename*.

Sizing guidelines for main and temporary stores

Several changes in Sybase IQ 15.0 architecture affect data storage:

• The IQ_SYSTEM_MAIN dbspace holds all of the database metadata other than IQ table metadata. IQ table metadata is stored in the table's dbspace and the TLV log. If a node is down, the multiplex needs to store versions in order to synchronize them when the node comes back up. These versions may use large amounts of space.

• Approximately 20 percent of the IQ_SYSTEM_MAIN dbspace is now used for preallocated free list space and not available for user data.

• Because Sybase IQ 15.0 performs more operations in parallel, it uses more temporary space than previous releases.

Three factors influence the space required for the IQ_SYSTEM_MAIN store:

• Versioning – the volume of versions maintained varies

- Nature of data and indexes

- Dynamic nature of the data – the capacity to load more data at any time

While Sybase can offer general guidelines, the combination of these factors makes each database's requirements unique. For a development or report server with a total size under 500GB, an IQ_SYSTEM_MAIN file of 10 to 20GB may suffice. For a production database, see Table 1-3 for size guidelines.

*Table 1-3: Size guidelines for IQ_SYSTEM_MAIN and
IQ_SYSTEM_TEMP in production databases*

Task	Guideline	Notes
Loading empty schema from iqunload -n output or for a small test database	10GB main, 5GB temporary	CREATE DATABASE sizes are in MB. The server must be at 12.7 ESD #5 or higher to use iqunload -n.
Creating new production database	• If you are migrating a database, and use a raw device for your current IQ_SYSTEM_MAIN, assign a new unused raw device of your standard size. • Total size of IQ_SYSTEM_MAIN should be at least 1/100 total database size, with a minimum 100GB main and 100GB reserve. • If using raw disks for IQ_SYSTEM_MAIN, use multiple raw disks whenever possible. Multiple raw disks enable Sybase IQ to stripe the data across devices, which improves performance. • Only use file system files for IQ dbspaces in production if the file system is fault-tolerant and implemented by a high-performance, redundant disk array (for example, RAID 5). For single-server systems, you can use a local file system, but multiplex systems require a cluster file system, ideally on a Storage Area Network device.	Omit ms_size if specifying a raw device. Always set the main reserve to 20 percent of IQ_SYSTEM_MAIN size. To set the main reserve, use the database option MAIN_RESERVED_DBSPACE_MB.
Creating main store for a multiplex	Double the space recommended for a simplex database, or at least 200GB main and 200GB reserve dbspace.	

Examples   **Example 1**   In CREATE DATABASE syntax, default size units are in MB, not
GB. The following statement creates a database with 100GB
IQ_SYSTEM_MAIN with 100GB reserve (for future expansion):

```
CREATE DATABASE 'test.db'
IQ PATH 'test.iq'
IQ SIZE 100000
IQ RESERVE 100000
TEMPORARY PATH 'test.iqtmp'
TEMPORARY SIZE 5000
```

Example 2   MAIN_RESERVED_DBSPACE_MB lets you control the amount of
space Sybase IQ sets aside in your IQ main store for certain small but critical
data structures used during release savepoint, commit, and checkpoint
operations. Sybase recommends that you set the
MAIN_RESERVED_DBSPACE_MB option value to 20 percent of the
IQ_SYSTEM_MAIN SIZE. See "IQ main store and IQ temporary store space
management" on page 45 for more about MAIN_RESERVED_DBSPACE_MB.
For example, if IQ_SYSTEM_MAIN is 100GB, set it to 20GB, as follows:

```
SET OPTION PUBLIC.MAIN_RESERVED_DBSPACE_MB = 20000
```

Example 3   You can specify the IQ_SYSTEM_MAIN size in the database
migration command. The -ms_size parameter requires a value in MB, not GB.
Omit -ms_size if specifying a raw device. For a raw device, you must specify
an unused raw partition. For more about migration, see the *Installation and
Configuration Guide*.

For example, this statement creates an IQ_SYSTEM_MAIN on a raw device:

```
iqunload -au -ms_filename /dev/rdsk/c1t0d1 -c
"UID=DBA;PWD=SQL;DBF=latest.db"
```

For example, this statement creates an IQ_SYSTEM_MAIN on a raw device:

```
iqunload -au -ms_filename \\\\.\\PhysicalDrive1 -c
"UID=DBA;PWD=SQL;DBF=latest.db"
```

Dbfiles and output files

The SYSDBFILE view shows all the dbfiles in your database, including the
catalog dbspace file, the IQ message file, dbfiles in the IQ main and temporary
dbspaces, the transaction log file, and the SA temporary file.

Files that are not dbfiles do not appear in the SYSDBFILE view. These include files that may be generated on server startup, such as the console log (specified by the -o switch) and the SQL log (specified by -zo). These log files do appear as database properties or server properties and may be examined by stored procedures such as sa_db_properties( ) or the system function db_property( ). For syntax, see "Version information in the message log" in Chapter 1, "Overview of Sybase IQ System Administration" in *System Administration Guide: Volume 1*.

Naming dbspaces

You can rename a user dbspace and dbfile names, but you cannot rename or drop catalog dbspace SYSTEM, IQ main dbspace IQ_SYSTEM_MAIN, IQ temporary dbspace IQ_SYSTEM_TEMP, and IQ message dbspace IQ_SYSTEM_MSG.

You can rename the dbfile names in IQ_SYSTEM_MAIN and IQ_SYSTEM_TEMP, and you can change the dbfile name of IQ_SYSTEM_MSG files, but you cannot change the dbfile names in SYSTEM. You cannot use ALTER DBSPACE RENAME TO in order to rename dbspaces IQ_SYSTEM_MAIN or IQ_SYSTEM_TEMP, IQ_SYSTEM_MSG, or SYSTEM.

The dbspace of a table is implicitly or explicitly specified. For base tables, the value of the DEFAULT_DBSPACE option implicitly determines the dbspace location, or the location may be specified explicitly using the CREATE TABLE IN *dbspace_name* clause. Base tables are typically created in a dbspace in the IQ main store, but may also be created without IQ indexes in a dbspace in the catalog store.

For global temporary tables, specify the IN SYSTEM clause to explicitly create a SA global temporary table. IQ temporary tables are created in IQ_SYSTEM_TEMP by default.

The dbspace of indexes and table subobjects is always specified, either implicitly based on the table dbspace or explicitly with an IN clause. Indexes and partitions are implicitly created in the table dbspace. The partitioned subsets of columns of a partitioned table are implicitly created in the dbspace of the corresponding partitions.

All indexes and table subobjects must be in the same store as the table. For example, an index on a base table in the IQ main store must also be in the IQ main store.

Dbfile attributes and operations

A dbfile has read-write or read-only status. A dbfile is read-write when it is added, and its runtime read-write status depends on both the read-write status of the dbspace and of the dbfile. The administrator can alter the read-write/read-only status of a dbfile, but cannot alter the online/offline status of a dbfile.

Operations that can be performed on dbfiles include add, drop, rename logical name, and rename the file path. See ALTER DBSPACE statement in Chapter 1, "SQL Statements," in *Reference: Statements and Options*.

Dbspace attributes and operations

A dbspace may have three types of online/offline status: online, offline, or dynamically offline. Dynamically offline means that the dbspace is marked offline in memory, as opposed to marked offline in the catalog. If a database starts and one or more dbfiles cannot be opened, the database starts but the dbspace is marked dynamically offline.An administrator can use ALTER DBSPACE ONLINE to bring a dbspace back online after fixing a problem, but this only changes the dbspace's in-memory status.

In addition to online, offline, or dynamically offline status, a dbspace also has read-write or read-only status. When created, a dbspace is online and read-write.

A dbspace also has striping attributes. An administrator may specify whether striping is on and the stripe size.

Operations that can be performed on dbspaces include add, drop, and rename. See CREATE DBSPACE statement, ALTER DBSPACE statement, and DROP statement in Chapter 1, "SQL Statements," in *Reference: Statements and Options*.

For multiplex dbspaces, see "Updating dbspaces in multiplex" in Chapter 3, "Running Multiplex Transactions," in *Using Sybase IQ Multiplex*

Read-only and read-write dbspaces and files

For a read-only dbspace, the administrator has the following capabilities:

- Add a file

- Rename the file path of a dbfile in the dbspace (requires main dbspaces are offline)

- Drop an empty file

- Rename the dbspace or dbfile in the dbspace

A file is read-only when either the file status is read-only or the file status is read-write, but the owning dbspace status is read-only. Altering a dbspace to read-only does not alter the catalog status of its associated files to read-only, but does make the file(s) read-only at the operating system level. In other words, the file's catalog read-only or read-write status remains the same, but data in the file cannot be modified.

To create a table object in a dbspace or to move a table object to a new dbspace requires CREATE permission in the dbspace. The CREATE permission in a dbspace can be granted/revoked to/from a user or a group. Any member in a group inherits CREATE permission from the group. By default, CREATE permission on IQ_SYSTEM_MAIN, IQ_SYSTEM_TEMP, and SYSTEM is granted to PUBLIC. For other IQ main dbspaces, the system administrator must explicitly grant CREATE permission on the dbspace to a group/user before they can create or move objects into that dbspace. For example, if a table is to be placed on a new IQ main dbspace, the user must have CREATE permission on that dbspace. Users must also have RESOURCE permission to create objects.

A dbspace and its associated files can have individual RO or RW status, for example:

Object	Status	Effective status	Table's dbspace	Table's status
dbspace1	RW	RW	dbspace1	RW
- file1	RO	RO		
- file2	RW	RW		
dbspace2	RO	RO	dbspace2	RO
- file1	RO	RO		
- file2	RW	RO		
dbspace3	RW	RO	dbspace3	RO
- file1	RO	RO		
- file2	RO	RO		

A table or join index is read-only when assigned to a RO dbspace. A table partition is read-only when the partition is assigned to a RO dbspace. No data modifications such as insert, delete, update, load, truncate table, and insert/delete/update through an updatable cursor are allowed to a RO table or RO table partition. No DDL operations such as ALTER TABLE add/drop column, create/drop index are allowed on a RO table or RO table partition.

Attempts to write to a read-only dbspace are detected when the modified pages are flushed to disk. Pages modified during an INSERT...VALUES statement are not written to the database until the next command that is *not* an INSERT...VALUES statement. (INSERT...VALUES is the only command that behaves this way.) Sybase IQ returns an error for DML operations on a read-only table or read-only table partition.

Operations to join indexes, including create join index, drop join index, and synchronize join index, fail if any of the join tables are RO.

The following table lists allowed dbspace configuration transitions.

Table 1-4: Allowed dbspace configuration transformations

State	Alter Type	Allowed	Allowed
		User Main	IQ_SYSTEM_MAIN, IQ_SYSTEM_TEMP
Online DBSpace			
	Alter DBSpace Offline	Yes	No
	Alter DBSpace Online	No	No for temp, yes for IQ_SYSTEM_MAIN
	Alter DBSpace RO	Yes, if it is RW	No
	Alter DBSpace RW	Yes, if it is RO	No
	Alter Striping Parameters	Yes	Yes, on single node and multiplex coordinator
	Rename DBSpace	Yes	No
	Add File	Yes	Yes
	Drop File	Yes, if empty	Yes, if empty and RO
	Alter File RO	Yes, if it is RW	Yes, if it is RW and not the last RW dbfile
	Alter File RW	Yes, if it is RO	Yes, if it is RO
	Alter File Size	Yes	Yes
	Alter File Rename Logical Name	Yes	Yes
	Alter File Rename Path	No	No
Offline DBSpace			
	Alter DBSpace Offline	No	NA
	Alter DBSpace Online	Yes	NA
	Alter DBSpace RO	No	NA

State	Alter Type	Allowed	Allowed
	Alter DBSpace RW	No	NA
	Alter Striping Parameters	Yes	NA
	Rename DBSpace	Yes	NA
	Add File	No	NA
	Drop File	Yes, if empty	NA
	Alter File RO	Yes	NA
	Alter File RW	Yes	NA
	Alter File Size	No	NA
	Alter File Rename Logical Name	Yes	NA
	Alter File Rename Path	Yes	NA
Dynamically Offline DBSpace			
	Alter DBSpace Offline	Yes, if RO	NA
	Alter DBSpace Online	Yes	NA
	Alter DBSpace RO	Yes, if RW	NA
	Alter DBSpace RW	No	NA
	Alter Striping Parameters	Yes	NA
	Rename DBSpace	Yes	NA
	Add File	No	NA
	Drop File	Yes, if empty	NA
	Alter File RO	Yes	NA
	Alter File RW	Yes	NA
	Alter File Size	No	NA
	Alter File Rename Logical Name	Yes	NA
	Alter File Rename Path	No	NA
Read-only DBSpace			
	Alter DBSpace Offline	Yes, if online	NA
	Alter DBSpace Online	Yes, if offline	NA
	Alter DBSpace RO	No	NA
	Alter DBSpace RW	Yes, if online	NA
	Alter Striping Parameters	Yes	NA
	Rename DBSpace	Yes	NA
	Add File	Yes	NA

State	Alter Type	Allowed	Allowed
	Drop File	Yes, if empty	NA
	Alter File RO	Yes, if RW	NA
	Alter File RW	Yes, if RO	NA
	Alter File Size	No	NA
	Alter File Rename Logical Name	Yes	NA
	Alter File Rename Path	Yes, if offline	NA
read-write DBSpace			
	Alter DBSpace Offline	No	No
	Alter DBSpace Online	Yes, if dynamically offline	No
	Alter DBSpace RO	Yes	No
	Alter DBSpace RW	No	No
	Alter Striping Parameters	Yes	Yes
	Rename DBSpace	Yes	No
	Add File	Yes	Yes
	Drop File	Yes, if empty	Yes, if empty and RO
	Alter File RO	Yes, if RW	Yes, if RW
	Alter File RW	Yes, if RO	Yes, if RO
	Alter File Size	Yes, if RW	Yes, if RW
	Alter File Rename Logical Name	Yes	Yes
	Alter File Rename Path	No	No
Read-only File			
	Alter File RO	No	No
	Alter File RW	Yes	Yes
	Alter File Size	No	No
	Alter File Rename Logical Name	Yes	Yes
	Alter File Rename Path	Yes, if offline	No
Read-write File			
	Alter File RO	Yes	Yes
	Alter File RW	No	No
	Alter File Size	Yes	Yes

State	Alter Type	Allowed	Allowed
	Alter File Rename Logical Name	Yes	Yes
	Alter File Rename Path	No	No

Notes

- Dynamically offline means the dbspace is marked offline in memory, as opposed to marked offline in the catalog

- A read-only IQ_SYSTEM_MAIN dbfile can be dynamically offline

- For IQ_SYSTEM_MSG, the only modification that is permitted is to rename the path, which is done using the command ALTER DBSPACE IQ_SYSTEM_MSG RENAME 'filepath'

Object placement for non-partitioned tables

You can specify a dbspace for a table object (including columns of any data type, indexes, primary and foreign keys, unique constraints, join indexes, and non-partitioned tables) at object creation or when you move the object. You must have CREATE privilege in the dbspace. For example:

```
CREATE TABLE tab1 (
        col1 INT IN dsp1,
        col2 VARCHAR(20),
        col3 CLOB IN dsp2,
        col4 DATE,
        col5 BIGINT,
        PRIMARY KEY (col5) IN dsp4) IN dsp3;
CREATE DATE INDEX col4_date ON tab1(col4) IN dsp5;
```

Resulting data allocation is as follows:

Dbspace name	Data
dsp1	FP index for col1
dsp2	FP index for col3
dsp3	FP indexes for col2, col4, and col5
dsp4	Primary key (HG on col5)
dsp5	DATE index col4_date

Object placement for partitioned tables

For a partitioned table, you can place each table partition in an individual dbspace. You may also place each column for a table partition in an individual dbspace. In general, however, individual dbspaces are recommended only for BLOB or CLOB columns or columns of CHAR, VARCHAR or VARBINARY greater than 255 bytes.   For example:

```
CREATE TABLE tab2(
    col1 INT IQ UNIQUE(65500),
    col2 VARCHAR(20),
    col3 CLOB PARTITION (p1 IN dsp11, p2 IN dsp12,
        p3 IN dsp13),
    col4 DATE,
    col5 BIGINT,
    col6 VARCHAR(500) PARTITION (P1 IN dsp21,
        p2 IN dsp22),
    PRIMARY KEY (col5) IN dsp2) IN dsp1
    PARTITION BY RANGE (c4)
    (p1 VALUES <= ('2006/03/31') IN dsp31,
    p2 VALUES <= ('2006/06/30') IN dsp32,
    p3 VALUES <= ('2006/09/30') IN dsp33
    );

CREATE DATE INDEX c4_date ON foo(c4) IN Dsp3;
```

Resulting data allocation is as follows:

Partition	Dbspace	Data
p1	dsp11	FP indexes for col3 (CLOB data)
	dsp21	FP index for col6 (VARCHAR(500) data)
	dsp31	FP indexes for col1, col2, col4, and col5
p2	dsp12	FP for col3 (CLOB data)
	dsp22	FP for col6 (VARCHAR(500) data)
	dsp32	FP indexes for col1, col2, col4, and col5
p3	dsp13	FP index for col3 (CLOB data)
	dsp33	FP indexes for col1, col2, col4, col5, and col6 (varchar(500) data)
Non-partitioned	dsp1	Lookup store for col1 and other share data (for all partitions)

Partition	Dbspace	Data
Non-partitioned	dsp2	Primary key HG on col5 (for all partitions)
Non-partitioned	dsp3	DATE index col4_date (for all partitions)

Starting databases

In order for a database to start, all files of IQ_SYSTEM_MAIN, all files of IQ_SYSTEM_TEMP, and the catalog file SYSTEM must be available. A database can be started skipping other user and catalog dbspaces that cannot be fully opened. All read-write files of IQ main store dbspaces other than the IQ_SYSTEM MAIN dbspace and any catalog dbspace other than SYSTEM that cannot successfully open all files of the dbspace, logs an error and marks the dbspace dynamically offline (marked offline in memory, as opposed to marking it offline in the catalog). If all files of IQ_SYSTEM_TEMP cannot be opened, the database will not start unless the -iqnotemp startup parameter is used.

Sybase IQ consistency checks the commit_id in each dbspace file header against the value in the system tables ISYSDBFILE and ISYSIQDBSPACE and marks any file or dbspace that does not match offline as above.

A dbspace that has been marked offline at start time may be brought online via the ALTER DBSPACE ONLINE statement, assuming that the problem has been corrected and the dbspace can be opened.

A table object that resides in an offline dbspace is unavailable. Any DDL or DML request except ALTER DBSPACE ONLINE to any table object in an offline dbspace generates an error. Note that after you make a dbspace offline, there may still be data pages in the buffer cache. In the case of a very small table, the entire table may be in memory in the buffer cache and temporarily available, even if the dbspace is offline.

Privileges for dbspace management

To create a table object in a dbspace or to move a table object to a new dbspace requires CREATE permission in the dbspace. The CREATE permission in a dbspace can be granted or revoked for a user or group. Any member in a group inherits CREATE permission from the group. By default, the CREATE permission on IQ_SYSTEM_MAIN, IQ_SYSTEM_TEMP, and SYSTEM is granted to PUBLIC. For other IQ main dbspaces, the system administrator must explicitly grant CREATE permission on the dbspace to a group/user before they can create or move objects into that dbspace.

SQL syntax for dbspace management and partitioning

Chapter 1, "SQL Statements," in *Reference: Statements and Options* includes the syntax and objects that support range partitions and configurable tablespaces. See the following sections:

- ALTER DBSPACE statement
- ALTER INDEX statement
- ALTER TABLE statement
- BACKUP statement
- COMMENT statement
- CREATE DBSPACE statement
- CREATE INDEX statement
- CREATE JOIN INDEX statement
- CREATE TABLE statement
- DROP statement
- GRANT statement
- RESTORE statement
- REVOKE statement
- TRUNCATE TABLE statement

Database options for dbspace management and partitioning

Chapter 2, "Database Options," in *Reference: Statements and Options* describes the new database options that support range partitions and configurable tablespaces. See the following sections:

- "DEFAULT_DBSPACE option"
- "DEFAULT_DISK_STRIPING option"
- "DEFAULT_KB_PER_STRIPE option"

Stored procedures for dbspace management and partitioning

Chapter 7, "System Procedures," includes the new and changed stored procedures that support range partitions and configurable tablespaces. For new procdures, see:

- "sp_iqbackupdetails procedure"
- "sp_iqbackupsummary procedure"
- "sp_iqdbspace procedure"
- "sp_iqdbspaceinfo procedure"
- "sp_iqdbspaceobjectinfo procedure"
- "sp_iqemptyfile procedure"
- "sp_iqfile procedure"
- "sp_iqobjectinfo procedure"
- "sp_iqrestoreaction procedure"

Some existing procedures have new versions and new behavior:

- The stored procedure sp_iqdbspace displays detailed information about each IQ dbspace. The procedure sp_iqfile now displays the information formerly provided by sp_iqdbspace in versions prior to 15.0.

- A new version of the sp_iqdbspaceinfo stored procedure now displays the size of each object and sub-object used in a specified table or join index and has new keywords.

- The following procedures now display a isPartitioned output column: sp_iqtable and sp_iqcolumn.

- The following procedures now display a dbspace_id output column: sp_iqtable and sp_iqjoinindex.

  This example displays the new dbspace_id and isPartitioned columns:

```
call sp_iqtable ('Departments')

Table_name   Table_type  Table_owner  Server_type  location  dbspace_id
Departments  BASE        GROUPO       IQ           Main      16384
```

```
isPartitioned  Remarks                                 table_constraints
    N          contains the names and heads of the     (NULL)
               various departments in the sporting
               goods company
```

• The dbspace name column has a different name for the following procedures:

Procedure name	Old dbspace column name	New dbspace column name
sp_iqdbspace	Name	DBSpaceName
sp_iqindexinfo	dbspace_name	DbspaceName
sp_iqspaceinfo	dbspace_name	DbspaceName

• The sp_iqcheckdb procedure has several changes. See "sp_iqcheckdb procedure" in *Reference: Building Blocks, Tables, and Procedures*.

Database consistency checking for dbspace management and partitioning

Changes to the database consistency checking stored procedure sp_iqcheckdb support range partitions and configurable tablespaces. See "sp_iqcheckdb procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.

Secure business intelligence

Sybase IQ 15.0 security enhancements include changes to network encryption support and user login management.

Support of FIPS encryption, Kerberos authentication, and column encryption is included in the separately licensed Sybase IQ Advanced Security Option. Support of IPv6 is available without a separate license.

FIPS support

Sybase IQ includes enhancements to Federal Information Processing Standards (FIPS) approved encryption technology. FIPS is supported on all platforms supported by Sybase IQ 15.0, except Microsoft Windows for 64-bit systems.

The main impact of FIPS support for Sybase IQ is that encryption can be nondeterministic, which is now the default behavior. A nondeterministic algorithm is one in which the same input yields different output values each time. This means that when a string is encrypted with a key, the encrypted string is different each time. The algorithm, however, is still able to decrypt the nondeterministic result using the key. This feature makes analyzing the encryption algorithm more difficult and the encryption more secure.

Support of FIPS is part of the separately licensed Sybase IQ Advanced Security Option.

Both RSA and FIPS security are included with Sybase IQ. RSA encryption requires no separate libraries, but FIPS requires two optional libraries: *dbfips11.dll* and *sbgse2.dll*. The library *sbgse2.dll* is provided by Certicom. Both security models require certificates. The rsaserver certificate has been renamed from *rsaserver.crt* to *rsaserver.id*.

FIPS also requires the following registry setting, which is set automatically by the Sybase IQ installation utility:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Certicom\libsb]
"expectedtag"=hex:5b,0f,4f,a6,e2,4a,ef,3b,44,07,05,2e,
b0,49,02,71,1f,d9,91,b6
```

For more information on using FIPS and RSA encryption, see "Transport-layer security" at http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dbadmin_en11/da-transport-layer-security.html and "Keeping your data secure" at http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dbadmin_en11/da-security.html in *SQL Anywhere Server – Database Administration*.

Kerberos authentication support

Sybase IQ now supports Kerberos authentication. This login feature allows you to maintain a single user ID and password for both database connections and operating system and/or network logins. You can use your Kerberos credentials to connect to the database without specifying a user ID or password.

Kerberos authentication is part of the separately licensed Sybase IQ Advanced Security Option.

See "Kerberos authentication" at
http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dba
dmin_en11/da-kerberos-authentication.html in *SQL Anywhere Server –
Database Administration*.

IPv6 support

Sybase IQ now supports Internet Protocol version 6 (IPv6), which contains
addressing and control information to route packets over the Internet. IPv6
supports 2^{128} unique IP addresses, which is a substantial increase over the
number of addresses supported by its predecessor IPv4. Sybase IQ supports
both IPv4 and IPv6 addresses anywhere you can specify an IP address on the
client or server.

ODBC classes support the use of IPv6 addresses for remote data access. JDBC
classes do not support the use of IPv6 addresses for remote data access.

See "Using the TCP/IP protocol" at
http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dba
dmin_en11/da-using-tcpip.html in SQL Anywhere Server – Database
Administration.

CREATE DATABASE ENCRYPTED clause extension

The syntax for the ENCRYPTED clause of the CREATE DATABASE statement
has been extended to allow you to specify SIMPLE as an encryption type.
Additionally, you can specify the encryption key and the algorithm in any
order.

You can now specify the algorithms AES256 and AES256_FIPS in the CREATE
DATABASE ENCRYPTED clause. See CREATE DATABASE statement in
Chapter 1, "New Features in Sybase IQ 15.0," in *Reference: Statements and
Options*.

Login management changes

Sybase IQ defines the rules to be followed when establishing a user's database connection in a database object called a **login policy**. A login policy is a named object in the database that consists of a set of options. Each login policy is associated with a set of options called login policy options. See CREATE LOGIN POLICY statement Chapter 1, "SQL Statements," in *Reference: Statements and Options*.

You must have DBA privileges to create new login policies or assign an existing login policy to a user. Login policies cannot be inherited through the user group hierarchy. For the SQL command syntax to manage policies, see "SQL syntax for login management" on page 37.

Each new database is created with a default login policy, called the root policy. You can modify the option values for the root login policy, but you cannot drop the policy. When a user account is created without specifying its login policy, the user becomes part of the root login policy. Any options that are not explicitly set when creating a login policy inherit their values from the root policy.

See "Managing login policies overview" at http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dbadmin_en11/da-permissi-s-4686947.htmlin *SQL Anywhere Server - Database Administration*.

You can execute login management commands on any multiplex server and they get automatically propagated to all servers in the multiplex. As recommended for any DDL for performance reasons, these commands should be executed on the coordinator. For details, see "Multiplex login management," in Chapter 1, "Managing Multiplex Servers," in *Using Sybase IQ Multiplex*.

Migrating databases to Sybase IQ 15.0 removes existing login management settings and replaces certain stored procedures and system tables with new ones listed in "New and changed system tables" on page 55 and "New and changed stored procedures" on page 62. To recreate login management settings after migration, use the SQL syntax in the following section.

SQL syntax for login management

See the following sections in Chapter 1, "New Features in Sybase IQ 15.0,"in *Reference: Statements and Options*:

- ALTER LOGIN POLICY statement

- ALTER USER statement

- COMMENT statement

- CREATE LOGIN POLICY statement

- CREATE USER statement

- DROP LOGIN POLICY statement

- "DROP USER statement

System tables for login management

System views display system table columns in readable format. See "SYSLOGINPOLICY system view," "SYSLOGINPOLICYOPTION system view," and "SYSUSER system view" in Chapter 8, "System Views," of *Reference: Building Blocks, Tables, and Procedures*.

System procedures for login management

Sybase IQ 15.0 login policies support setting maximum connections and password expiration for subsets of users, such as guest accounts, rather than on a per user basis. See Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*. The following new and changed system procedures support login policies:

- "sa_get_user_status system procedure"

- "sp_expireallpasswords procedure"

- "sp_iqaddlogin procedure"

- "sp_iqcopyloginpolicy procedure"

- "sp_iqmodifyadmin procedure"

- "sp_iqmodifylogin procedure"

- "sp_iqpassword procedure"

Improved operational management

Sybase IQ 15.0 provides improved operational management through new and changed query and system diagnostics and monitoring utilities, license management using SySAM 2.0, new message log management features, and improved management of IQ main store and IQ temporary store space.

Improved multiplex high availability

Multiplex functionality is undergoing significant changes for Sybase IQ 15.0. For more information, see "New multiplex features in Sybase IQ 15.0" in Chapter 1, "Multiplex Overview," in *Using Sybase IQ Multiplex.*

Query and system diagnostics and monitoring utilities

The following changes that affect system diagnostics and monitoring are described in this section:

- A new stored procedure, sp_iqstatistics, returns information about available in-memory statistics.

- A new stored procedure, sa_dependent_views, provides information about dependent views for a table or view. Modifications to the ALTER VIEW, ALTER FUNCTION, and CREATE FUNCTION SQL statements support the new procedure.

New sp_iqstatistics procedure

Function   Returns serial number, name, description, value, and unit specifier for each available statistic or a specified statistic. For syntax, see "sp_iqstatistics procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.

New sa_dependent_views procedure support

Function   Returns the list of all dependent views for a given table or view. For syntax, see "sa_dependent_views procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.

SQL statement support for sa_dependent_views

Modifications to the ALTER VIEW, ALTER FUNCTION, and CREATE FUNCTION SQL statements support the new sa_dependent_views procedure.

License management through SySAM 2

For additional information about license management, see the *Sybase Software Asset Management (SySAM) 2 Users Guide*. For syntax, see "sp_iqlmconfig procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.

Message log management

The new message log management feature allows the server to archive the message log file and maintain all of the information written to the message log, preventing the log file size from growing indefinitely. Message log management is controlled by the new server properties IQMsgMaxSize and IQMsgNumFiles or the new server startup switches -iqmsgsz and -iqmsgnum.

Note The IQMSG_LENGTH_MB database option has been deprecated and must be removed from existing scripts and code. Attempts to use the IQMSG_LENGTH_MB option will result in an error.

Message logging

An IQ message log file exists for each database. The default name of this file is *dbname.iqmsg*. The IQ message log file is created when the database is started for the first time after the database is created. The default location of the message log file is the same directory as the catalog store, although the database creator may specify another location, in addition to specifying a different file name.

By default, Sybase IQ logs the following types of messages in the message log file:

- Error messages

- Status messages

- Insert notification messages

- Query plans

You can examine this file as you would any other text file.

Message log file management

By default, the message log file continues to grow in size without limit and exists until you drop the database. To control the size of the message log file, you can set a limit on the size of the file and enable either message log wrapping or log archiving.

You can delete, rename, or copy the message file at any time after stopping the database.

Message log management is controlled by either the server properties IQMsgMaxSize and IQMsgNumFiles or the server startup switches -iqmsgsz and -iqmsgnum:

- IQMsgMaxSize or -iqmsgsz sets an upper limit in megabytes (MB) on the active message log size. Allowed values are integers between 0 and 2047 (inclusive). The default is 0, which means there is no limit on the size of the message log file.

- IQMsgNumFiles or -iqmsgnum sets the number of message log archives. Allowed values are integers between 0 and 64 (inclusive). The default is 0, which means that messages are wrapped in the main message log file and there is no archiving.

The value of the server property takes precedence over the corresponding server startup switch. When the server starts, the values of the -iqmsgsz and -iqmsgnum server switches are written in the server log file.

For information on setting the IQMsgMaxSize and the IQMsgNumFiles server properties, see "sa_server_option system procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.

For information on setting the -iqmsgsz and -iqmsgnum server startup switches, see "Starting the database server" in Chapter 1, "Running the Database Server," in the *Utility Guide*.

Message log wrapping

When you enable message log wrapping, as soon as the log file reaches the maximum size specified in the IQMsgMaxSize server property or the -iqmsgsz server startup switch, new messages are written at the beginning of the file. Existing messages are overwritten, line-by-line.

Message log wrapping is enabled by setting IQMsgMaxSize or -iqmsgsz to a value greater than zero (the default value of zero indicates there is no limit on the size of the message log file) and setting the IQMsgNumFiles server property or the -iqmsgnum server startup switch to zero (the default).

When wrapping is enabled, the tag `<next msg insertion place>` in the message log file tells you where new messages are being placed. Additional tags at the beginning and end of the file remind you that log wrapping is enabled, and that the last message in the file may not be the most recent one.

Message log archiving

You can maintain all of the information written to the message log file without allowing the file size to grow indefinitely by enabling message log archiving. A message log archive is a file in which the contents of the active *.iqmsg* message log file is saved.

Enabling message log archiving

Message log archiving is enabled by setting the IQMsgMaxSize server property using the sa_server_option system stored procedure or the -iqmsgsz server startup switch to a value greater than zero and setting the IQMsgNumFiles server property or the -iqmsgnum server startup switch to the number of message log archives. The value of the server property takes precedence over the value of its corresponding server switch.

Message log archive names

The names of the *.iqmsg* message log archives follow the pattern *logname.iqmsg.n*, where *n* is a positive integer greater than zero and *logname* is the name of the message path as specified in the CREATE DATABASE statement or is the database name with the suffix *.iqmsg*. The archives are created as read-only files in the same directory as the message log file.

When the message log file *logname.iqmsg* is full and the number of message log archives is less than the number specified in IQMsgNumFiles or -iqmsgnum, the server renames the current message log to create a new archive. Then a new message log is created with the name *logname.iqmsg*.

For example, if -iqmsgnum is equal to 5, the message log archives are created in the following order: *logname.iqmsg.1*, *logname.iqmsg.2*, *logname.iqmsg.3*, *logname.iqmsg.4*, and *logname.iqmsg.5*.

When the message log file *logname.iqmsg* is full and the maximum number of message log archives already exists, the oldest archive (with file extension *.1*) is deleted before the current message log is archived.

For example, when -iqmsgnum is 5 and *logname.iqmsg.5* exists:

- The file *logname.iqmsg.1* is deleted

- The files *logname.iqmsg.2* to *logname.iqmsg.5* are renamed to *logname.iqmsg.1* to *logname.iqmsg.4*,respectively

- The active message log file is renamed to *logname.iqmsg.5*

• A new message log file *logname.iqmsg* is created

Using this method, the server always keeps the most recent message logs, when the value of IQMsgNumFiles (or -iqmsgnum) is greater than zero.

New database server startup switches

Two new server command-line switches control the size of the message log and the number of archives of old message logs: -iqmsgnum and -iqmsgsz. For information, see "Starting the database server" in Chapter 1, "Running the Database Server," in the *Utility Guide*.

New sa_server_option system procedure options

The sa_server_option system procedure has two new options that can be used to reset message log size and number of log archives, while the database server is running. For details, see "sa_server_option system procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.

Message log management errors

Message logging stops, if the disk becomes full during message logging. The error `"Disk Full!!! Message logging stopped."` is written in the server log.

Message logging stops, if the following errors occur during message log management:

• File rename error: the server cannot archive the current active log or cannot rename any existing log archive

• File creation error: the server cannot create the new message log

• File deletion error: the server cannot delete the old archive

• Disk I/O error

The server automatically resumes message logging when the error condition is resolved. A message is written to the server log when message logging stops and when logging resumes.

For example, if renaming the file fails, the message `"Renaming of <filename> failed. Message log could not be archived. Message logging stopped."` is written in the server log. When logging resumes, the message `"Message logging resumed."` is written in the server log.

If the server fails to delete or rename a file because the file does not exist, the error is ignored and the log management process continues.

Backing up the iqmsg file and log archives

Backing up the message log file *dbname.iqmsg* and the message log archives is a good idea, even though these files are not required for a restore.

If problems occur during a restore, the *.iqmsg* file contains information that proves that the database was shut down before the backup started. The message log files may be useful in diagnosing the cause of the database failure from which you are recovering. Be sure to make a copy before restoring, for use in later analysis.

If IQ message log wrapping is enabled (the IQMsgMaxSize server option or the -iqmsgsz server startup switch is not equal to zero, and IQMsgNumFiles server option or the -iqmsgnum server startup switch is zero), you will probably want to back up the *.iqmsg* file so that all messages are accessible in the event you need them for diagnostic purposes.

If message log archiving is enabled (the IQMsgMaxSize server option or the -iqmsgsz server startup switch is not equal to zero and the IQMsgNumFiles or the -iqmsgnum server startup switch is not equal to zero), the server automatically backs up the message log archives. The maximum amount of message log that is archived is 128GB, which is sufficient in most cases.

Note Backing up the message log archives *is* required before a server restart. After the server restarts, the existing log archives are ignored and a new archive is created when the *dbname.iqmsg* file is full. To preserve the old archive logs, back up the files before restarting the server.

See also
- For more information on message log management, see "Message log file management" in Chapter 1, "Overview of Sybase IQ System Administration," of the *System Administration Guide: Volume 1*.

- For information on setting server startup switches, see "Starting the database server" in Chapter 1, "Running the Database Server," in the *Utility Guide*.

- For information on setting server properties, see "sa_server_option system procedure" in Chapter 7, "System Procedures," in the *Reference: Building Blocks, Tables, and Procedures*.

IQ main store and IQ temporary store space management

The IQ main store and IQ temporary store space management has undergone some improvements in Sybase IQ 15.0.

The user with DBA authority can now limit the amount of space used per connection. In addition, when IQ runs out of space in IQ main store or the IQ temporary store, it no longer suspends the transaction that ran out of space until new space is added. In IQ 15.0, the transaction that runs out of space in the IQ main store or the IQ temporary store will fail and gets rolled back.

Quota management for IQ temporary store

The new database option MAX_TEMP_SPACE_PER_CONNECTION limits the amount of IQ temporary store space used per connection and tracks temporary store usage for all Data Manipulation Language (DML) statements, in addition to queries. For syntax, see "MAX_TEMP_SPACE_PER_CONNECTION option" in Chapter 2, "Database Options," in *Reference: Statements and Options*. MAX_TEMP_SPACE_PER_CONNECTION monitors and limits the actual run time temporary store usage by the statement. If the connection exceeds the quota set by the MAX_TEMP_SPACE_PER_CONNECTION option, an error is returned and the current statement rolls back.

The default value of the existing QUERY_TEMP_SPACE_LIMIT database option has changed from 2GB to 0, which means there is no limit on temporary store usage by queries. To limit the temporary store usage per connection, the DBA can set the MAX_TEMP_SPACE_PER_CONNECTION option for all DML statements, including queries.

When a Sybase IQ database is upgraded from an earlier release to version 15.0, the MAX_TEMP_SPACE_PER_CONNECTION database option is set to the default value of 0. You can use sp_iqcheckoptions to find the default and current values of options before and after upgrading, to help determine if the new option settings are appropriate for the upgraded database.

Out of dbspace handling

Sybase IQ 15.0 introduces a new feature so that the IQ server no longer waits for additional space on an out-of-dbspace condition, but instead rolls back either the entire transaction or rolls back to a savepoint. If there is not enough temporary or main dbspace available for a buffer or dbspace allocation request, then the statement making the request rolls back. At this point, the DBA can add more space to a dbspace using the ALTER DBSPACE or the ALTER FILE command.

Out of dbspace error
messages

If the entire transaction is rolled back on an out-of-dbspace condition, the
following new error message is reported:

```
"%1 -- Transaction rolled back"
```

Item	Value
SQLCode	-1285L
Constant	IQ_TRANSACTION_ROLLBACK
SQLState	40W09
ODBC 2 State	ERROR
ODBC 3 State	ERROR
Sybase Error Code	2973
Severity Code	16
Parameter 1	error because of which the transaction rolled back

Probable cause: This transaction has been rolled back because the server
encountered an error during a critical operation.

If a buffer or dbspace allocation request fails because there is no space in the
dbspace, the following new error message is logged in the *iqmsg* message file:

```
"You have run out of space in %2 DBSpace. %1"
```

Item	Value
SQLCode	-1009170L
Constant	EMSG_OUT_OF_DBSPACE
SQLState	QSB66
ODBC 2 State	ERROR
ODBC 3 State	ERROR
Sybase Error Code	20223
Severity Code	14
Parameter 1	location of the exception
Parameter 2	name of dbspace

This error message replaces the error message "You have run out of { IQ
STORE | IQ TEMPORARY STORE } dbspace in database <***dbname***>.
In another session, please issue a CREATE DBSPACE ... { IQ
STORE | IQ TEMPORARY STORE } command and add a dbspace of
at least ***nn*** MB."

If a buffer allocation request finds a dirty buffer, but the buffer manager cannot flush the buffer due to an out-of-space condition, the following new error message is returned and the current statement rolls back:

```
"%2: All buffer cache pages are in use, ask your DBA to
increase the size of the buffer cache. %1"
```

Item	Value
SQLCode	-1009031L
Constant	EMSG_BUFMAN_ALLSLOTSLOCKED
SQLState	QSA31
ODBC 2 State	ERROR
ODBC 3 State	ERROR
Sybase Error Code	20052
Severity Code	14
Parameter 1	location of the exception
Parameter 2	the particular buffer cache throwing the exception

Space usage monitoring example

The following code creates a timer based event that monitors space usage to help avoid unexpected rollbacks, which may occur in out of space situations on non-privileged operations.

```
CREATE EVENT DBSpaceLogger
SCHEDULE START TIME '00:00:01' EVERY 300 SECONDS
HANDLER
BEGIN
DECLARE DBSpaceName VARCHAR(128);
DECLARE Usage SMALLINT;
DECLARE cursor_1 CURSOR FOR
SELECT DBSpaceName, Usage
FROM sp_iqdbspace()
WHERE Usage > 0
ORDER BY Usage
FOR READ ONLY;

OPEN cursor_1;
idx1: LOOP
FETCH cursor_1 INTO DBSpaceName, Usage;
IF SQLCODE <> 0 THEN LEAVE idx1 END IF;
IF Usage >= 70 AND Usage < 80 THEN
call dbo.sp_iqlogtoiqmsg('Information: DBSpace' +
DBSpaceName + '''s usage is more than 70%');
```

```
ELSEIF Usage >= 80 AND Usage < 90 THEN
call dbo.sp_iqlogtoiqmsg('Warning: DBSpace ' +
DBSpaceName + '''s usage is more than 80%');
ELSEIF Usage >= 90 AND Usage < 100 THEN
call dbo.sp_iqlogtoiqmsg('Critical Warning: DBSpace
' + DBSpaceName + '''s usage is more than 90%');
END IF;
END LOOP;
CLOSE cursor_1;
END;
```

The DBSpaceLogger event is created in the sample iqdemo database.

Behavior Changes in Sybase IQ 15.0

About this chapter

This chapter describes the behavior changes introduced in Sybase IQ 15.0 that are not described in the new features chapter.

Note For changes documented in 12.7 ESD releases, see the latest version of the *Release Bulletin* in the Product Manuals at http://www.sybase.com/support/manuals/.

Contents

Deprecated features

This section lists deprecated features supported in the current software that will not be supported in release 15.0 of Sybase IQ. Changes in these features may affect existing applications.

Note Lists of deprecated features may be incomplete and are subject to change.

The following SQL statements and clauses are deprecated:

- IQ local store for store type in the CREATE DBSPACE command.

- RELOCATE clause in the ALTER DBSPACE command.

- The STRIP ON clause of the LOAD TABLE statement has been deprecated. To continue stripping trailing blanks, use STRIP RTRIM instead. For details, see LOAD TABLE statement in Chapter 1, "SQL Statements," in *Reference: Statements and Options*.

The following database options are deprecated:

Database options deprecated in Sybase IQ 15.0

Option name	Option name
ANSI_INTEGER_OVERFLOW	MIN_SMPDJ_OR_HPDJ_INDIRECT_SIZE
AUTO_COMMIT	MIN_SMPDJ_OR_HPDJ_FILTERED_PPM
AUTO_REFETCH	MIN_SMPDJ_OR_HPDJ_SIZE
AUTOMATIC_TIMESTAMP	MIN_SMPDJ_OR_HPDJ_SIZE
BELL	MPX_GLOBAL_TABLE_PRIV
CHAR_OEM_TRANSLATION	MPX_LOCAL_SPEC_PRIV
COMMAND_DELIMITER	MPX_OPTIONS
COMMIT_ON_EXIT	NULLS
CONVERT_HG_TO_1242	ON_ERROR
DESCRIBE_JAVA_FORMAT	OPTIMISTIC_WAIT_FOR_COMMIT
DISK_STRIPING	OPTIMIZATION_LOGGING
DISK_STRIPING_PACKED	OUT_OF_DISK_MESSAGE_REPEAT
DIVIDE_BY_ZERO_ERROR	OUT_OF_DISK_WAIT_TIME
ECHO	OS_OPTION_CRASH
ENABLED_ORDERED_PUSHDOWN_INSERTION	OUTPUT_FORMAT
ENABLE_THREAD_ALLOWANCE	OUTPUT_LENGTH
FLATTEN_SUBQUERIES	OUTPUT_NULLS
FLOAT_AS_DOUBLE	PARALLEL_GBH_ENABLED
HEADINGS	PARALLEL_GBH_MIN_ROWS_PER_UNIT

Database options deprecated in Sybase IQ 15.0

Option name	Option name
INPUT_FORMAT	PARALLEL_GBH_UNITS
IO_EXERCISER_PARAM_FILE	PERCENT_AS_COMMENT
IQMSG_LENGTH_MB	QUERY_PLAN_ON_OPEN
ISQL_COMMAND_TIMING	QUIET
ISQL_ESCAPE_CHARACTER	RETURN_JAVA_AS_STRING
ISQL_FIELD_SEPARATOR	RI_TRIGGER_TIME
ISQL_LOG	SCREEN_FORMAT
ISQL_PLAN	SORT_PHASE1_HELPERS
ISQL_PLAN_CURSOR_SENSITIVITY	STATISTICS
ISQL_PLAN_CURSOR_WRITABILITY	SQLCONNECT
ISQL_QUOTE	SQLSTART
JAVA_HEAP_SIZE	TEMP_KB_PER_STRIPE
JAVA_INPUT_OUTPUT	THREAD_COUNT
JAVA_NAMESPACE_SIZE	THREAD_STACK
JAVA_PAGE_BUFFER_SIZE	THREAD_SWAPS
LOCAL_KB_PER_STRIPE	TRUNCATE_WITH_AUTO_COMMIT
LOCAL_RESERVED_DBSPACE_MB	TRUNCATION_LENGTH
LOG_DETAILED_PLANS	TRUNCATE_DATE_VALUES
LOG_MAX_REQUESTS	TRUNCATION_LENGTH
MAIN_KB_PER_STRIPE	TSQL_HEX_CONSTANT
MAX_WORK_TABLE_HASH_SIZE	UUID_HAS_HYPHENS
MIN_NLPDJ_FILTERED_PPM	MAIN_CACHE_MEMORY_MB
MIN_NLPDJ_TABLE_SIZE	TEMP_CACHE_MEMORY_MB

Additional notes on deprecated database options:

- The IQMSG_LENGTH_MB database option has been deprecated and must be removed from existing scripts and code. Attempts to use the IQMSG_LENGTH_MB option return an error. See "Message log management" on page 40 for information on the new message log management feature.

- The FLATTEN_SUBQUERIES option has been deprecated. For information on new subquery flattening options, see "Improved query performance" on page 6.

- Deprecated options related to disk striping in releases prior to Sybase IQ 15.0: DISK_STRIPING, DISK_STRIPING_PACKED, MAIN_KB_PER_STRIPE, TEMP_KB_PER_STRIPE, LOCAL_KB_PER_STRIPE. To align your schema with storage, see "Migrating Data," in the *Installation and Configuration Guide*.

- MAIN_CACHE_MEMORY_MB and TEMP_CACHE_MEMORY_MB are removed. The IQ main and IQ temporary cache size settings are now server options only. If databases need to run with different cache sizes, change the option on the server using the stored procedure sa_server_option main_cache_memory_mb or temp_cache_memory_mb options. For example, instead of:

  ```
  execute {set option
  "PUBLIC".MAIN_CACHE_MEMORY_MB = 200}
  ```

  use this:

  ```
  execute {call sa_server_option
  ('main_cache_memory_mb',200)}
  ```

  This command changes the option on a running engine before starting a database but cannot change the cache size on a running database. If two databases need to be run with different cache sizes, then set the option before starting each database.

  If you set the cache sizes using the server startup -iqmc or -iqtc switches, the specified values are used for all databases started on that server unless the values are changed with an sa_server_option option. If the cache sizes are not specified, the default values apply.

  For more information on setting IQ main and IQ temporary cache sizes and the use of the sa_server_option stored procedure, see "Setting IQ main and IQ temporary cache sizes" on page 78.

The following server startup option is deprecated:

- -iqdroplks

  To rebuild the free list at the dbspace level, use sp_iqcheckdb command with appropriate arguments instead.

See "Compatibility views" at http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dbreference_en11/views-s-5195671.htmlin *SQL Anywhere Server - SQL Reference*.

The following stored procedure and stored procedure mode are deprecated:

- sp_iqrelocate

To empty out all objects in a file within a dbspace, use the sp_iqemptyfile procedure. To relocate files, use ALTER TABLE MOVE and sp_iqemptyfile.

• sp_iqcheckdb repair mode

A new mode dropleaks allows you to reset the allocation map for the entire database or a specified dbspace target.

The iq_bcp utility is deprecated. Sybase IQ 15.0 has a true client-side loader that supports the full functionality of the IQ LOAD TABLE command, including the ability to load BCP formatted text files. For more information, see "Direct loading of data from clients" on page 1.

Deprecated collations

The following collations that were available in 12.6 and 12.7 are deprecated in Sybase IQ 15.0:

• 437EBCDIC – Code Page 437, EBCDIC ordering

• 437ESP – Code Page 437, Spanish

• 437LATIN1 – Code Page 437, Latin 1, Western

• 819DAN – Code Page 819, Danish

• 819ESP – Code Page 819, Spanish

• 819ISL – Code Page 819, Icelandic

• 819LATIN1 – Code Page 819, Latin 1, Western

• 819NOR – Code Page 819, Norwegian

• 819SVE – Code Page 819, Swedish/Finnish

• 850DAN – Code Page 850, Danish

• 850ESP – Code Page 850, Spanish

• 850ISL – Code Page 850, Icelandic

• 850LATIN1 – Code Page 850, Latin 1, Western

• 850NOR – Code Page 850, Norwegian

• 850SVE – Code Page 850, Swedish/Finnish

• 852CYR – Code Page 852, Cyrillic

• 852LATIN2 – Code Page 852, Latin 2, Central/Eastern European

- 852POL – Code Page 852, Polish

- 855CYR – Code Page 855, Cyrillic

- 856HEB – Code Page 856, Hebrew

- 857TRK – Code Page 857, Turkish

- 860LATIN1 – Code Page 860, Latin 1, Western

- 861ISL – Code Page 861, Icelandic

- 862HEB – Code Page 862, Hebrew

- 863LATIN1 – Code Page 863, Latin 1, Western

- 865NOR – Code Page 865, Norwegian

- 866RUS – Code Page 866, Russian

- 869ELL – Code Page 869, Greek

- 920TRK – Code Page 920, Turkish, ISO 8859-9

- 950TWN – Code Page 950, Traditional Chinese, Big 5 Encoding

- 1252DEU – Code Page 1252, Windows Specialty German, Umlaut chars not equal

- C – Standard C collation

- internal – Code Page 850, Multilingual

- SJIS – Japanese Shift-JIS Encoding

- SJIS2 – Japanese Shift-JIS Encoding, ASE-compatible

- UTF8 – UTF-8, 8-bit multibyte encoding for Unicode, legacy ordering

- WIN_LATIN1 – Code Page 1252, Windows Latin 1, Western, ISO8859-1 with extensions

- WIN_LATIN5 – Code Page 1254, Windows Latin 5, Turkish, ISO8859-9 with extensions

You can use the iqunload utility to migrate to Sybase IQ 15.0 an existing 12.6 or 12.7 database that was created with a deprecated collation. For details about iqunload, see the *Installation and Configuration Guide*.

The deprecated collations are no longer visible in the Sybase Central list of supported collations. Although the SQL syntax still allows you to create new databases with the deprecated collations, creating new databases using these deprecated collations is not recommended.

SQL Anywhere changes

The following features are no longer supported or are deprecated, and may affect existing applications. For a complete list of SQL Anywhere changes, see *SQL Anywhere 11 – Changes and Upgrading*.

* ANSI_INTEGER_OVERFLOW option — If you use the ANSI_INTEGER_OVERFLOW option, an overflow always results in a SQLSTATE = 22003 - overflow error. The settings of this option were previously ignored and the default behavior was always to overflow.

System tables and views

This section contains information on new and changed system tables and views.

 The catalog has been redesigned in Sybase IQ 15.0:

* All system tables are new.

* Access to system tables is now only through system views. (For descriptions of corresponding system views, remove the "I" from each system table name and see "System views" on page 60.)

* To support existing applications, compatibility views that match Sybase IQ 12.7 are provided for some system tables. These compatibility views should not be used for new applications.

New and changed system tables

Upgrading to Sybase IQ 15.0 adds the following system tables:

* ISYSDBSPACE

* ISYSDBSPACEPERM

* ISYSIQBACKUPHISTORY

* ISYSIQBACKUPHISTORYDETAIL

* ISYSIQDBSPACE

* ISYSIQMPXLOGINPOLICYOPTION

- ISYSIQMPXSERVER

- ISYSIQPARTITIONCOLUMN

- ISYSLOGINPOLICY

- ISYSLOGINPOLICYOPTION

- ISYSPARTITION

- ISYSPARTITIONKEY

- ISYSPARTITIONSCHEME

- ISYSSUBPARTITIONKEY

- ISYSUSER

See "ISYSIQMPXLOGINPOLICYOPTION system table" and "SYSIQMPXSERVER system view" in Chapter A, "Multiplex Reference,"in *Using Sybase IQ Multiplex*.

Upgrading to Sybase IQ 15.0 removes the following system tables:

- IQ_USER_LOGIN_INTO_TABLE

- IQ_SYSTEM_LOGIN_INFO_TABLE

- IQ_LOGIN_PER_SERVER

- IQ_MPX_SERVER_PROPERTIES

- IQ_MPX_INFO

- IQ_MPX_STATUS

- IQ_MPX_VERSIONLIST

The system tables listed in Table 2-1 are replaced. Compatibility views with the same names as the replaced tables support older versions of Sybase IQ that supported these tables.

Table 2-1: Tables with compatibility views

Removed/deprecated system table/ New compatibility view	Replacement system tables	Corresponding system views
SYSFILE	ISYSDBSPACE and ISYSDBFILE	SYSDBSPACE and SYSDBFILE
SYSINDEX	ISYSIDX	SYSIDX
SYSIQCOLUMN	ISYSIQTABCOL	SYSIQTABCOL
SYSIQFILE	ISYSIQDBFILE	SYSIQDBFILE
SYSIQINDEX	ISYSIQIDX	SYSIQIDX

Removed/deprecated system table/ New compatibility view	Replacement system tables	Corresponding system views
SYSIQJOININDEX	ISYSIQJOINIDX	SYSIQJOINIDX
SYSIQTABLE	ISYSIQTAB	SYSIQTAB

Changes to system tables and view columns

In the SYS.ISYSTAB and SYS.ISYSIDX system tables, a new column called dbspace_id replaces the file_id column, which is deprecated in Sybase IQ 15.0. The file_id column remains in these tables and their associated compatibility views.

The following table lists other differences in system table and view columns between Sybase IQ versions 12.7 and 15.0.

System table and view column differences

System table name	Sybase IQ 12.7		Sybase IQ 15.0		
	Column name	Data type	Column name	Data type	Possible application impact
SYSCATALOG	check	text	check	smallint	N
SYSCOLUMN	width	smallint	width	unsigned int	Y
	unused				N
	unused2				N
	format_str				N
	remote_name				N
	remote_type				N
SYSCONSTRAINT	table_id*				Y
	column_id*				Y
	index_id*				Y
	fkey_id*				Y
SYSINDEXES	indid	numeric	indid	bigint	N
	fname	text	fname	char	Y
SYSINFO	page_size	int	page_size	varchar	Y

System table and view column differences

System table name	Sybase IQ 12.7		Sybase IQ 15.0		Possible application impact
	Column name	**Data type**	**Column name**	**Data type**	
	default_collation	char	default_collation	varchar	Y
	classes_version				Y
SYSIQCOLUMN	info_location	hs_vdorecid	info_location	smallint	N
	info_recid_size				N
	info_location_size	unsigned int	info_location_size	smallint	N
SYSIQFILE	file_name	char	file_name	text	Y
SYSIQINDEX	info_location	hs_vdorecid	info_location	smallint	N
	info_recid_size	unsigned int	info_recid_size	smallint	N
	info_location_size	unsigned int	info_location_size	smallint	N
SYSIQINFO	last_full_backup				N
	last_incr_backup				N
SYSJAR	create_time**				N
	remarks				N
SYSJARCOMPONENT	create_time**				N
	remarks				N
SYSJAVACLASS	replaced_by				N
	type_id				N
	class_descriptor				N
	create_time**				N
	remarks				N
SYSPROCPARMS	length	smallint	length	unsigned int	Y
	remarks				Y

System table and view column differences

| System table name | Sybase IQ 12.7 | | Sybase IQ 15.0 | | Possible application impact |
	Column name	Data type	Column name	Data type	
SYSREMOTEUSER	log_send	numeric	log_send	unsigned bigint	N
	log_sent	numeric	log_sent	unsigned bigint	N
	confirm_sent	numeric	confirm_sent	unsigned bigint	N
	log_received	numeric	log_received	unsigned bigint	N
	confirm_received	numeric	confirm_received	unsigned bigint	N
SYSREMOTEUSERS	log_send	numeric	log_send	unsigned bigint	N
	log_sent	numeric	log_sent	unsigned bigint	N
	confirm_sent	numeric	confirm_sent	unsigned bigint	N
	log_received	numeric	log_received	unsigned bigint	N
	confirm_received	numeric	confirm_received	unsigned bigint	N
SYSSCHEDULE	event_id	int	event_id	unsigned int	Y
SYSTABLEPERM	ttable_id				Y
SYSUSERTYPE		smallint		unsigned int	Y
	format_str				N
	super_type_id				N

Notes:

* Replaced by ref_object_id: the object ID of the column, table, or index to which the constraint applies

** The create time is available in SYSOBJECT.create_time

System views

❖ **To view detailed system information views and definitions (Sybase Central):**

1 Connect to the database as a user with DBA authority.

2 Right-click the database and choose Configure Owner/Container Filtering.

3 Click SYS and then click OK.

4 In the left pane, double-click Views.

5 In the left pane, click a view, and in the right pane, click the SQL tab.

6 Click the Data tab to view detailed information about the selected view.

The following table lists system views. For syntax, see Chapter 8, "System Views," of *Reference: Building Blocks, Tables, and Procedures*

System View name
SYSCOLUMNS system view
SYSFKEY system view
SYSIQBACKUPHISTORY system view
SYSIQBACKUPHISTORYDETAIL system view
SYSIQDBFILE system view
SYSIQDBSPACE system view
SYSIQIDX system view
SYSIQINFO system view
SYSIQJOINIDX system view
SYSIQJOINIXCOLUMN system view
SYSIQJOINIXTABLE system view
SYSIQPARTITIONCOLUMN system view
SYSIQTAB system view
SYSLOGINPOLICY system view
SYSLOGINPOLICYOPTION system view
SYSPARTITION system view
SYSPARTITIONKEY system view
SYSPARTITIONSCHEME system view
SYSSSYNCPROFILE system view
SYSUSER system view

Additional system views are common to both Sybase IQ and SQL Anywhere.

See "System views" at
http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dbr
eference_en11/rf-system-views.html in *SQL Anywhere Server - SQL Reference*.

For deprecated system views, see "Deprecated features" on page 50. These
views are visible in the product, but replaced by other views and therefore no
longer documented.

Compatibility views

Sybase IQ provides some deprecated views for compatibility with its older
versions. See "To view detailed system information views and definitions
(Sybase Central):" on page 60 for instructions on how to display columns and
SQL syntax for these views. Table 2-1 on page 56 lists compatibility views and
the current views that you should use instead.

See "Compatibility views" at
http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dbr
eference_en11/views-s-5195671.html in *SQL Anywhere Server - SQL Reference*.

For information about deprecated and compatibility columns in current views,
see "System views" on page 60.

Consolidated views

Consolidated views provide data in a form more frequently required by users.
For example, consolidated views often provide commonly-needed joins.
Consolidated views differ from system views in that they are not just a
straightforward view of raw data in a underlying system table(s). For example,
many of the columns in the system views are unintelligible ID values, whereas
in the consolidated views, they are readable names.

Consolidated views such as SYSCATALOG and SYSINDEXES are common to
both Sybase IQ and SQL Anywhere.

See "Consolidated views" at
http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dbreference
_en11/views-s-4117677.html in *SQL Anywhere Server - SQL Reference*.

Data Definition Language (DDL) changes

This section contains behavior changes related to Data Definition Language (DDL).

DDL syntax changes to support new features are described in "High performance analytics" on page 1 and "Information lifecycle management" on page 13. Please see these sections for additional DDL changes.

Backup and restore changes

There are four significant changes to backup and restore:

- An IQ full, incremental-since-full or incremental backup may be restricted to only the set of read-write files in the database using the READWRITE FILES ONLY keywords. The read-write files must be IQ dbspaces. For syntax and details, see BACKUP statement in Chapter 1, "SQL Statements," in *Reference: Statements and Options*

- An IQ backup may back up a set of read-only dbspaces and/or read-only files. The read-only dbspaces/files must be IQ dbspaces.

- A restore of a backup may, optionally, restore only a subset of the dbspaces and/or files. The subset must be READONLY dbspaces and/or files. For syntax, see RESTORE statement in Chapter 1, "SQL Statements," in *Reference: Statements and Options*.

- System tables now contain backup history. See "SYSIQBACKUPHISTORY system view" and "SYSIQBACKUPHISTORYDETAIL system view" in Chapter 8, "System Views," of *Reference: Building Blocks, Tables, and Procedures*.

New and changed stored procedures

Upgrading to Sybase IQ 15.0 adds the following stored procedures:

- sa_get_user_status

- sp_expireallpasswords

- sp_iqbackupdetails

- sp_iqbackupsummary

- sp_iqcardinality_analysis

- sp_iqcopyloginpolicy
- sp_iqdbspaceobjectinfo
- sp_iqemptyfile
- sp_iqfile
- sp_iqlmconfig
- sp_iqobjectinfo
- sp_iqpassword
- sp_iqrestoreaction
- sp_iqstatistics

  For details, see "New sp_iqstatistics procedure" on page 39.

New procedures are documented in Chapter 7, "System Procedures," in
Reference: Building Blocks, Tables, and Procedures.

Upgrading to Sybase IQ 15.0 removes the following stored procedures:

- sp_iq_post_process_login
- sp_iq_process_login
- sp_iqlistexpiredpasswords
- sp_iqlistlockedusers
- sp_iqlistpasswordexpirations
- sp_iqlocklogin
- sp_iqrelocate
- sp_iqretryselectinto

The following changes affect system procedures in Sybase IQ 15.0.

- The sp_iqcheckdb procedure has several changes. The repair mode is
  removed. A new mode dropleaks allows you to reset the allocation map for
  the entire database or a specified dbspace target. The main and local targets
  are removed. You can now check a dbspace target and check a subset of
  the rows of a table that belong to a specified partition. For details, see
  "sp_iqcheckdb procedure" in Chapter 7, "System Procedures," in
  Reference: Building Blocks, Tables, and Procedures.

- The stored procedure sp_iqdbspace is a new procedure that displays detailed information about each IQ dbspace. The new procedure sp_iqfile now displays the information formerly provided by sp_iqdbspace. For details, see "sp_iqdbspace procedure"and "sp_iqfile procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.

- A new version of the sp_iqdbspaceinfo stored procedure now displays the size of each object and sub-object used in a specified table or join index and has new keywords. For details, see "sp_iqdbspaceinfo procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.

- The following procedures now display a isPartitioned output column: sp_iqtable and sp_iqcolumn. For sample output, see "Stored procedures for dbspace management and partitioning" on page 33.

- The following procedures now display a dbspace_id output column: sp_iqtable and sp_iqjoinindex. For sample output, see "Stored procedures for dbspace management and partitioning" on page 33.

- The dbspace name column has a different name for the following procedures:

Procedure name	Old dbspace column name	New dbspace column name
sp_iqdbspace	Name	DBSpaceName
sp_iqindexinfo	dbspace_name	DbspaceName
sp_iqspaceinfo	dbspace_name	DbspaceName

  For sample output, see "sp_iqdbspace procedure"and "sp_iqdbspaceinfo procedure" in Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures*.

- The sp_iqprocparm procedure returns nine columns instead of ten. The remarks column is removed.

- Sybase IQ 15.0 no longer has store type IQ local store. The local keyword has been removed from the sp_iqindexinfo and sp_iqspaceinfo stored procedures. LOCAL is no longer a valid value for the location output column of the sp_iqcolumn, sp_iqconstraint, and sp_iqtable procedures. LOCAL is no longer a valid value for the SegmentType output column of sp_iqdbspace.

The following stored procedures now support 3-byte indexes:

- sp_iqcheckdb

- sp_iqcolumn

- sp_iqindexadvice

- sp_iqindexmetadata

- sp_iqindexsize

- sp_iqindex

- sp_iqindexfragmentation

- sp_rebuildindex

- sp_iqrowdensity

Changes to identifiers

In Sybase IQ 15.0, rules for identifiers have changed as follows:

- Table names can no longer contain double quote characters.

- User names and database names can no longer contain double quote characters, single quote, and semicolon characters.

- User names and database names may no longer start with a space or end with a space.

- Dbspace names are no longer case sensitive in a CASE RESPECT database. For example, the second command below now fails with the error "Item 'main2' already exists":

```
CREATE DBSPACE MAIN2 AS 'main_file2' IQ STORE 10;
CREATE DBSPACE main2 AS 'main_file3' IQ STORE 10;
```

Database collation behavior changes

The following is a list of database collation related features that are no longer supported or deprecated and that may affect existing applications.

- The creation of custom collations is no longer supported.

  If you are rebuilding a database with a custom collation, the collation is preserved if you rebuild in a single step. If you choose to unload the database and then load the schema and data into a database that you create, then you must use one of the supplied collations.

- The Collation utility (dbcollat) is not longer supported.

- The SQLLOCALE environment variable is no longer supported and has been replaced by the SALANG and SACHARSET environment variables.

- The ASLANG environment variable has been renamed SALANG.

- The ASCHARSET environment variable has been renamed SACHARSET.

- The SYSCOLLATION, SYSCOLLATIONMAPPINGS, and SYSINFO system tables have been deprecated. Collation mapping information is now stored as database properties.

  See "Compatibility views" at http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1 /dbreference_en11/views-s-5195671.htmlin *SQL Anywhere Server - Database Administration*.

- The SORTKEY function now uses the International Components for Unicode (ICU) library, instead of the Sybase Unicode Infrastructure Library (Unilib) and has new syntax. Sort key values created using a version of Sybase IQ prior to 15.0 do not contain the same values created using version 15.0 and higher. You should regenerate any sort key values in your database that were generated using a version of Sybase IQ prior to 15.0.

  See "SORTKEY function [String]" in Chapter 4, "SQL Functions," of *Reference: Building Blocks, Tables, and Procedures*.

- The CP874toUTF8 utility now calls the International Components for Unicode (ICU) library to perform data conversion. You can also load data in the CP874 character set without converting it to UTF8 using this utility. See "CP874toUTF8 utility" in Chapter 3, "Database Administration Utilities," in the *Utility Guide*.

For a list of collations deprecated in Sybase IQ 15.0, see "Deprecated collations" on page 53. For more information on changes to database collations, see "Database collation improvements" on page 12.

Database option changes

This section includes pointers to new options and details on database option changes.

For database options related to multiplex capability, see Appendix A, "Multiplex Reference," of *Using Sybase IQ Multiplex*.

New database options for Sybase IQ 15.0

The following sections in Chapter 2, "Database Options," in *Reference: Statements and Options* describe new options:

- "DEFAULT_DBSPACE option"
- "DEFAULT_DISK_STRIPING option"
- "DEFAULT_KB_PER_STRIPE option"
- "FP_LOOKUP_SIZE option"
- "FP_LOOKUP_SIZE_PPM option"
- "MAX_TEMP_SPACE_PER_CONNECTION option"
- "SORT_COLLATION option"
- "SUBQUERY_CACHING_PREFERENCE option"
- "SUBQUERY_FLATTENING_PERCENT option"
- "SUBQUERY_FLATTENING_PREFERENCE option"

The following sections in Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex* describe new options:

- "MPX_AUTOEXCLUDE_TIMEOUT option"
- "MPX_HEARTBEAT_FREQUENCY option"
- "MPX_IDLE_CONNECTION_TIMEOUT option"
- "MPX_MAX_CONNECTION_POOL_SIZE option"
- "MPX_MAX_UNUSED_POOL_SIZE option"

Changed option names and values

Three database options have changed as follows:

- DEFAULT_HAVING_SELECTIVITY is now DEFAULT_HAVING_SELECTIVITY_PPM

- DEFAULT_LIKE_MATCH_SELECTIVITY is now
  DEFAULT_LIKE_MATCH_SELECTIVITY_PPM

- DEFAULT_LIKE_RANGE_SELECTIVITY is now
  DEFAULT_LIKE_RANGE_SELECTIVITY_PPM

The meanings of these parameters is the same. Instead of "percent" they are now "parts-per-million" parameters, where 15.0 1% = 10,000. In IQ 12.7, 1% = a parameter setting of "1". The default values for DEFAULT_LIKE_MATCH_SELECTIVITY_PPM and DEFAULT_LIKE_RANGE_SELECTIVITY_PPM parameters have changed accordingly. They were "15" and are now "150000" (150,000, or 15%). The DEFAULT_HAVING_SELECTIVITY_PPM parameter still has a default of zero (meaning to let the optimizer choose).

Changed default values of database options

The default values of some database options have changed in Sybase IQ 15.0:

- The default value for the STRING_RTRUNCATION option changed from OFF to ON. This value changes for newly created databases.

- The default value of the QUERY_TEMP_SPACE_LIMIT database option has changed from 2GB to 0, which means there is no limit on temporary store usage by queries. To limit the temporary store usage per connection, the DBA can set the new MAX_TEMP_SPACE_PER_CONNECTION database option for all DML statements, including queries. For more information, see "Quota management for IQ temporary store" on page 45.

- The default value of the TEMP_SPACE_LIMIT_CHECK database option has changed to ON from OFF.

Changed default behavior of database options

The default behavior of the following database options has changed in Sybase IQ 15.0:

- The low limit of the MAIN_RESERVED_DBSPACE_MB and the TEMP_RESERVED_DBSPACE_MB database options has changed from 0 to 200MB. Reserved space size is now calculated as a maximum of 50% and a minimum of 1% of the last read-write file in IQ_SYSTEM_MAIN or IQ_SYSTEM_TEMP. A change in the value of MAIN_RESERVED_DBSPACE_MB or TEMP_RESERVED_DBSPACE_MB takes effect immediately. The server does *not* need to be restarted in order to change reserved space size.

- The default behavior of the ANSI_INTEGER_OVERFLOW database option changed in Sybase IQ 15.0. For details, see "SQL Anywhere changes" on page 55.

- The FORCE_DROP database option is no longer allowed on a secondary node. If a force drop is attempted on a secondary node, an error is returned. FORCE_DROP is now a temporary option, so that the value of the option does not get propagated to secondary nodes at synchronization.

Removed database options

See "Deprecated features" on page 50 for options that are removed from Sybase IQ 15.0.

Start-up and connection changes

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

Start and stop utility changes

The start and stop utilities, start_asiq and stop_asiq, are renamed to start_iq and stop_iq. The Windows 32-bit platform does not support the stop_iq utility.

Default upper limit of catalog store cache size (-ch) increased [CR 540471]

The default for the switch -ch (upper limit of catalog store cache size) doubled from 128MB to 256MB. See "Starting the database server," in Chapter 1, "Running the Database Server," in the *Utility Guide*.

New version number in start_iq [CR 494427]

The start_iq -v command now returns a four-digit IQ build number. For example:

Sybase IQ 12.7 output:

```
$ start_iq -v
Sybase IQ/12.7.0/080707/P/ESD 5/Enterprise Linux64 -
amd64,Opteron64,X86_64/2.4.21-27.0.1.ELsmp/64bit/2008-
07-07 14:27:26
```

Sybase IQ 15.0 output:

```
$ start_iq -v
Sybase IQ/15.0.0.5120/090225/P/GA/Linux ES release 4 -
x86_64 - 2.6.9-67.0.4.ELsmp/64bit/2009-02-25 00:22:48
```

File and directory name changes

Table 2-2 lists files and directories that are renamed in Sybase IQ 15.0.

Table 2-2: File and directory name changes

12.7 name	15.0 name
ASIQ-12_7.sh	IQ-15_0.sh
ASIQ-12_7/	IQ-15_0/
asiqdemo.*	iqdemo.*
ASIQservice.exe (Windows only)	SybaseIQService15.exe
asiqsrv12 (UNIX only)	iqsrv15
asiqsrv12.exe (Windows only)	iqsrv15.exe
dscp (UNIX only)	iqdscp
dsedit	iqdsedit
IQAgent12.jar	IQAgent15.jar
IQHelpen12.jar	IQHelpen15.jar
IQPlugin12.jar	IQPlugin15.jar
isql	iqisql
S99SybaseIQAgent12 (UNIX only)	S99SybaseIQAgent15
SybaseIQAgent12.exe (Windows only)	SybaseIQAgent15.exe
win32/ (Windows only)	bin32/
x64/ (Windows only)	bin64/

> **Note** Use start_iq to start Sybase IQ servers on UNIX. iqsrv15.exe is not
> supported for starting UNIX servers.

Server log files default to four digits

The default server log file name changed from *<servername>*.999.srvlog to
<servername>.9999.srvlog.

Request log file analysis

The format of the output in the request log file (generated by setting the -zr
server startup switch) has changed in Sybase IQ 15.0. The stored procedures
sa_get_request_profile and sa_get_request_times can be used to read the -zr log
file and summarize the results.

The procedure sa_get_request_profile analyzes the request log to determine the
execution times of similar statements and summarizes the results in the global
temporary table satmp_request_profile. For example:

```
call sa_get_request_profile('/sys1/users/jones/iqreqs1_zr.log');
select * from satmp_request_profile;
```

The procedure sa_get_request_times analyzes the request log to determine
statement execution times and summarizes the results in the global temporary
table satmp_request_time. For example:

```
call sa_get_request_times('/sys1/users/jones/iqreqs1_zr.log');
select * from satmp_request_time;
```

See "Request logging" at
http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1/dbu
sage_en11/performance-s-5753386.htmlin *SQL Anywhere Server - Database
Administration*.

Windows registry changes

Windows registry entries that contained the obsolete name Adaptive Server IQ
are updated to Sybase IQ. You must update Windows utilities that rely on these
registry entries to reflect the new name. Some of the main entries are shown in

Table 2-3: Windows registry entry changes

12.7 name	15.0 name
Software\Adaptive Server IQ	Software\Sybase IQ
ODBC\Adaptive Server IQ	ODBC\Sybase IQ

Demo database changes

Sybase IQ 15.0 ships a script and data that let you create your own demo database. The new database, iqdemo, supports an imaginary sporting goods company.

You can still create the demo database used in 12.6 or 12.7 by specifying the parameter -old when you invoke *mkiqdemo.sh* or *mkiqdemo.bat*. Table 2-4 lists the former names and new names of demo tables so that you can update scripts or applications.

Table 2-4: Demo database table names

12.7 name	15.0 name	15.0 owner
contact	Contacts	GROUPO
customer	Customers	GROUPO
department	Departments	GROUPO
emp1	emp1	DBA
employee	Employees	GROUPO
fin_code	FinancialCodes	GROUPO
fin_data	FinancialData	GROUPO
iq_dummy	iq_dummy	DBA
product	Products	GROUPO
sale	sale	DBA
sales_order_items	SalesOrderItems	GROUPO
sales_order	SalesOrders	GROUPO
SYSOPTIONDEFAULTS		DBA

In Sybase IQ 15.0, the demo database password has changed from 'SQL' to 'sql'.

❖ **Creating the demo database on Windows**

The file *%ALLUSERSPROFILE%\SybaseIQ\demo\mkiqdemo.sql* contains SQL statements that create the database and *%ALLUSERSPROFILE%\SybaseIQ\demo\demodata* contains its data.

• In the directory *%ALLUSERSPROFILE%\SybaseIQ\demo*, run the batch file *mkiqdemo.bat* and respond to the prompts.

❖ **Creating the demo database on UNIX**

The file *$IQDIR15/demo/mkiqdemo.sql* contains SQL statements that create the database and *$IQDIR15/demo/demodata* contains its data.

• In the directory *$IQDIR15/demo*, run the script *mkiqdemo.sh* and respond to the prompts.

Sybase IQ 15.0 does not include Megaphone Telco Demo.

Environment variable name changes

The following Sybase IQ 15.0 environment variable names have changed. Note that certain Sybase IQ 15.0 environment variable names now include the current release number, enabling you to run multiple versions of Sybase IQ on the same machine.

Table 2-5: Environment variable name changes

12.7 name	15.0 name
ASDIR	IQDIR15
ASIQTIMEOUT	IQTIMEOUT
ASIQPORT	IQPORT
ASLOGDIR	IQLOGDIR15
ASTMP	IQTMP15

Character set conversion

The -ct database server option for enabling and disabling character set conversion is no longer supported.

Character set conversion is always enabled for the database server, but if the database server determines that it is not required, then it is not used. You can disable character set conversion from the client by specifying CharSet=none in your connection string.

Localization environment variables renamed

The following environment variables are renamed for this release:

Table 2-6: Environmental variables renamed

Former name	New name
ASLANG	SALANG
ASCHARSET	SACHARSET

Simplified TDS tools iqdsedit and iqisql

Sybase IQ 15.0 includes simplified versions of the Tabular Data Stream (TDS) tools dsedit (Directory Services Editor) and isql (Interactive SQL), called iqdsedit and iqisql.

These tools provide support for a limited set of functions for creating interfaces files and testing connections for INSERT...LOCATION functionality and are for users who do not have any Open Client functionality from another Sybase product such as TDS. To use advanced features, like encrypted connections, you need to use the full versions of dsedit and isql included with SDK, ASE, or Replication Server.

The simplified versions of these tools, iqdsedit and iqisql, are located in *$IQDIR15/bin64* on UNIX and in *%SYBASE%\%IQDIR15%\bin64* on Windows (*%SYBASE%\%IQDIR15%\bin32* on Windows 32-bit systems).

Query enhancements and changes

This section contains behavior changes related to queries.

Disjunction of subquery predicates

Previous IQ versions supported only limited disjunction of subqueries. Sybase IQ 15.0 removes these restrictions and supports arbitrary combination of ANSI-SQL-89 style subqueries in WHERE or HAVING clauses, with two exceptions.

The following are not supported:

• Disjunction of subqueries in CASE statements

Disjunction of subqueries in OUTER JOIN conditions.

Examples **Example 1**   Disjunction of uncorrelated EXISTS and IN subqueries:

```
SELECT COUNT(*)
FROM supplier
WHERE s_suppkey IN (SELECT MAX(l_suppkey)
                          FROM lineitem
                          GROUP BY l_linenumber)
OR EXISTS (SELECT p_brand
           FROM part
           WHERE p_brand = 'Brand#43');
```

Example 2   Disjunction of correlated EXISTS subqueries:

```
SELECT COUNT(*)
FROM supplier S
WHERE EXISTS (SELECT l_suppkey
              FROM lineitem
              WHERE l_suppkey = S.s_suppkey )
OR EXISTS     (SELECT p_brand
              FROM part
              WHERE p.partkey = S.s_suppkey and
p_brand = 'Brand#43');
```

Example 3   Disjunction of correlated/uncorrelated quantified comparison subqueries:

```
SELECT COUNT(*)
FROM lineitem L
WHERE l_quantity > (SELECT MAX(s_acctbal)
                          FROM supplier
                          WHERE L.l_suppkey
= s_suppkey and s_nationkey = 10)
OR l_partkey >= ANY (SELECT MAX(p_partkey)
                     FROM part
                     GROUP BY p_mfgr);
```

Data load, update, and extraction changes

This section discusses behavior changes related to data load, update, and extraction facilities.

Load performance improvements

See the section "Improved loading for large single (fact) tables" on page 2 for details on LOAD TABLE syntax changes and information on the increased limit on the length of large object data (LOB) that can be retrieved from a remote database using INSERT...LOCATION.

FP conversion

In Sybase IQ 12.7, adding a new distinct value via LOAD, INSERT, or UPDATE to a FP(1) column of char(2), binary(2), or small integer, converts FP (1) to FP(2). In version 15.0, with a default main cache setting of 32 MB and a default FP_LOOKUP_SIZE_PPM setting of 2500, adding a new distinct value via LOAD, INSERT, or UPDATE to a FP(1) column of char(2), binary(2), or small integer, converts FP(1) to flat FP.

SQLCODE change for 0 rows updated [CR 533932]

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

Non-blocking named pipes [CR 406712]

Non-blocking named pipes have been implemented for input for LOAD TABLE and output for EXTRACT and BFILE. Statements waiting on pipe I/O can now be interrupted using CTRL-C.

For more information, see "Extraction with named pipes" in CChapter 7, "Moving Data In and Out of Databases," of*System Administration Guide: Volume 1*.

Installation and migration changes

See *Installation and Configuration Guide* for your platform for new migration steps you must take before installing Sybase IQ 15.0, and for details about the following subsections.

Network Clients for all platforms

Sybase IQ 15.0 now provides a Network Client for all supported server platforms and includes the platform-specific network client with each server purchased.

SDK no longer installed

Sybase IQ 15.0 no longer installs the Software Development Kit.

Sybase Central enhancements

The Sybase IQ 15.0 plug-in for Sybase Central is upgraded from Sybase Central version 4.3 to version 6.0.

Limits, memory, and disk use changes

This section contains behavior changes related to limits, memory, and disk use.

Cache size defaults increased

The default cache sizes are increased as follows:

- Default temporary cache: 24MB

- Default main cache: 32MB

These default cache sizes are suitable for the demo database only and are too small for any production application.

Reserved space size calculation

Reserved space size is now calculated as a maximum of 50% and a minimum of 1% of the last read-write file in IQ_SYSTEM_MAIN or IQ_SYSTEM_TEMP. A change in the value of MAIN_RESERVED_DBSPACE_MB or TEMP_RESERVED_DBSPACE_MB takes effect immediately. The server does *not* need to be restarted in order to change reserved space size.

The low limit of the MAIN_RESERVED_DBSPACE_MB and the TEMP_RESERVED_DBSPACE_MB database options has changed from 0 to 200MB.

Setting IQ main and IQ temporary cache sizes

The MAIN_CACHE_MEMORY_MB and TEMP_CACHE_MEMORY_MB database options used to set IQ main and IQ temporary cache sizes are removed in Sybase IQ 15.0. The IQ main and IQ temporary cache size settings are now server options only. Server options may be specified on the command line or in a configuration file when starting the server. Server options apply to all databases started or created by the server.

Many server options, for example, the request logging options -zr, -zo, -zs, -zn, -zt, -zl and -zp, the console output options -o and -on, and the idle and liveness timeout values -ti and -tl, may be dynamically changed for a running server using the sa_server_option stored procedure.

Server option settings never persist. They exist in memory only while the server is running. When the server is restarted, the options must be specified again using the start parameters or by calling sa_server_option.

The IQ main and IQ temporary cache sizes set using the -iqmc and -iqtc server options may now be changed using the sa_server_option stored procedure as follows:

```
sa_server_option 'main_cache_memory_mb', value;
sa_server_option 'temp_cache_memory_mb', value;
```

Unlike the other server options listed above, the IQ main and IQ temporary cache settings are used only when starting a database. Therefore, to have an effect, a database must be started after calling sa_server_option to change the values.

Using the stored procedure sa_server_option to set the IQ main and IQ temporary cache sizes is useful mainly in a test environment, where many tests with different databases or different cache settings are run using the same server without shutting the server down. A production database should always be started using a configuration file to specify appropriate values for the number of IQ threads, IQ thread stack size, number of connections, and IQ cache sizes.

Documentation changes

This section describes changes related to the Sybase IQ 15.0 documentation set.

New Features Guide

The *New Features Guide* for Sybase IQ 15.0 contains new syntax and details for all new features. Refer back to the 15.1 core documentation for syntax that is unchanged. For example, if a new clause is added to a SQL Statement, see *New Features Guide* for that clause and the 15.1 *Reference: Statements and Options* for the complete statement and its other clauses.

In prior releases, the *New Features Guide* contained a brief summary linked to the details in core documentation.

Reference Manual

This section describes changes to the *Sybase IQ Reference Manual*.

- The *Reference Manual* is now divided into two volumes:

  - *Reference: Building Blocks, Tables, and Procedures* describes SQL, stored procedures, data types, and system tables that Sybase IQ supports.

- *Reference: Statements and Options* describes the SQL statements and options that Sybase IQ supports.

- The system tables diagrams are no longer included in the "System Tables" chapter.

Administration Guide

This section describes changes to the *Sybase IQ System Administration Guide*.

- The *Sybase IQ System Administration Guide* is now divided into two volumes:

  - *System Administration Guide: Volume 1* describes startup, connections, database creation, population and indexing, security, data integrity, transactions, versioning, collations, backup and archiving, and troubleshooting (formerly in its own book).

  - *System Administration Guide: Volume 2* covers topics for application programming such as procedures and batches, remote data access, events, OLAP and XML.

- Appendix A, "XML in the Database," is removed.

  See "XML in the Database" at http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1 /dbusage_en11/ug-sqlxml.html in *SQL Anywhere Server - SQL Usage*.

- Appendix B, "Data Access Using JDBC," is removed.

  See "SQL Anywhere JDBC Driver" at http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere.11.0.1 /dbprogramming_en11/pg-jdbc.html in *SQL Anywhere Server - SQL Usage*.

New and Changed Syntax

About this chapter

This chapter lists the SQL syntax, system object, and utility changes that support 15.0 functionality.

SQL statements

Unless otherwise noted, see Chapter 1, "SQL Statements," in *Reference: Statements and Options* for syntax and examples.

The following statements are new or changed in Sybase IQ 15.0:

- ALTER DBSPACE statement

  ALTER DBSPACE supports configurable tablespaces and partitioning. Operations that can be performed on dbfiles include add, drop, rename logical name, and rename the file path. See "Naming dbspaces" on page 23.

- ALTER FUNCTION statement

  New ALTER FUNCTION syntax supports the sa_dependent_views procedure.

- ALTER INDEX statement

  ALTER INDEX statement supports configurable tablespaces and partitioning.

- ALTER LOGIN POLICY statement

  New syntax supports login policies described in "Login management changes" on page 37.

- ALTER MULTIPLEX RENAME statement

  The ALTER MULTIPLEX RENAME and ALTER MULTIPLEX SERVER statements support multiplex configuration outside of Sybase Central. For syntax and examples, see ALTER MULTIPLEX RENAME statement in Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- ALTER MULTIPLEX SERVER statement

  For syntax and examples, see ALTER MULTIPLEX SERVER statement in Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- ALTER TABLE statement

  ALTER TABLE statement supports configurable tablespaces and partitioning.

- ALTER USER statement

  ALTER USER statement supports login management. See "Login management changes" on page 37.

- ALTER VIEW statement

  New ALTER VIEW syntax supports the sa_dependent_views procedure.

- BACKUP statement

  There are four significant changes to backup and restore:

  - An IQ full, incremental-since-full or incremental backup may be restricted to only the set of read-write files in the database using the READWRITE FILES ONLY keywords. The read-write files must be IQ dbspaces. For syntax and details, see "BACKUP statement" in Chapter 1, "SQL Statements," in *Reference: Statements and Options*

  - An IQ backup may back up a set of read-only dbspaces and/or read-only files. The read-only dbspaces/files must be IQ dbspaces.

- A restore of a backup may, optionally, restore only a subset of the dbspaces and/or files. The subset must be READONLY dbspaces and/or files. For syntax, see "RESTORE statement" in Chapter 1, "SQL Statements," in *Reference: Statements and Options*.

- System tables now contain backup history. See "SYSIQBACKUPHISTORY system view" and "SYSIQBACKUPHISTORYDETAIL system view" in Chapter 8, "System Views," of *Reference: Building Blocks, Tables, and Procedures*.

- COMMENT statement

  The COMMENT statement now supports:

  - Tablespaces and partitioning described in "Information lifecycle management" on page 13.

  - The clause LOGIN POLICY *policy-name*. See "Login management changes" on page 37

  - The clause KERBEROS LOGIN *"client-Kerberos-principal"*.

- CREATE DATABASE statement

  The CREATE DATABASE statement now supports:

  - Range partitioning and tablespace management. See "Sizing guidelines for main and temporary stores" on page 20.

  - Database collation improvements. See "Database collation improvements" on page 12

  - Message log archiving. See "Message log archiving" on page 42.

  - Encryption enhancements. See "CREATE DATABASE ENCRYPTED clause extension" on page 36.

  - Changes to identifiers such as table, user, dbspace and database names. See "Changes to identifiers" on page 65.

- CREATE DBSPACE statement

  The CREATE DBSPACE statement now supports:

  - Range partitioning and tablespace management. See "Dbspace attributes and operations" on page 24.

  - Rollback instead of waiting when a server is out of space. See "Out of dbspace handling" on page 45.

Using IQ local store for store type in the CREATE DBSPACE command is deprecated.

• CREATE FUNCTION statement

A new stored procedure, sa_dependent_views, provides information about dependent views for a table or view. Modifications to the ALTER VIEW, ALTER FUNCTION, and CREATE FUNCTION SQL statements support the new procedure.

For syntax and examples, see CREATE FUNCTION statement in Chapter 1, "SQL Statements," in *Reference: Statements and Options*.

• CREATE INDEX statement

CREATE INDEX syntax is unchanged for this release. If you omit the IN clause, which signals index placement, Sybase IQ creates the index in the dbspace where the table is created. You must have DBA authority or RESOURCE authority and CREATE privilege in the specified dbspace to create an index.

Sybase IQ has significantly improved the performance of loading High_Group (HG) indexes and Containment (also called WORD) (WD) indexes, while transactions are still allowed to access the tables being loaded. This load performance improvement for HG and WD indexes affects CREATE INDEX.

• CREATE JOIN INDEX statement

CREATE JOIN INDEX now supports:

  • Range partitions and configurable tablespaces

  • Read-only and read-write dbspaces and files.

• CREATE LOGIN POLICY statement

New syntax supports login policies described in "Login management changes" on page 37.

• CREATE MULTIPLEX SERVER statement

You can now configure your multiplex manually using DDL statements for creating, altering and dropping multiplex server definitions. See Chapter 2, "Managing Multiplex Servers," and Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

• CREATE TABLE statement

CREATE TABLE now supports range partitions and configurable tablespaces. See "Object placement for non-partitioned tables" on page 29 and "Object placement for partitioned tables" on page 29.

- CREATE USER statement

  New syntax supports login policies described in "Login management changes" on page 37.

- DROP DBSPACE statement

  New syntax supports operations described in "Dbspace attributes and operations" on page 24

- DROP LOGIN POLICYstatement

  New syntax supports login policies described in "Login management changes" on page 37.

- DROP MULTIPLEX SERVER statement

  You can now configure your multiplex manually using DDL statements for creating, altering and dropping multiplex server definitions. See Chapter 2, "Managing Multiplex Servers," and Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- DROP USER statement

  New syntax supports login policies described in "Login management changes" on page 37.

- GRANT statement

  New syntax supports login policies described in "Login management changes" on page 37and configurable tablespaces and partitioning described in "Read-only and read-write dbspaces and files" on page 24.

- LOAD TABLE statement

  The LOAD TABLE syntax includes the new STRIP clause keyword RTRIM. The keyword RTRIM replaces the keyword ON, which has been deprecated.

  When bulk loading large objects, the USING CLIENT FILE clause applies to both primary and secondary files. (If you have the Large Objects Management Option, see *Large Objects Management in Sybase IQ* for details.)

- RESTORE statement (see BACKUP statement)

- REVOKE statement

  See GRANT statement.

- TRUNCATE TABLE statement

  The partition clause specifies which partition to truncate. It does not affect data in other partitions.

SQL functions

The SORTKEY function now uses the International Components for Unicode (ICU) library, instead of the Sybase Unicode Infrastructure Library (Unilib®) and has new syntax.

For syntax, see "Alphabetical list of functions," in Chapter 4, "SQL Functions," in *Reference: Building Blocks, Tables, and Procedures*.

Database options

The following options are new or changed for Sybase IQ 15.0. Unless otherwise noted, see Chapter 2, "SQL Options," in *Reference: Statements and Options* for syntax and examples. "Deprecated features" on page 50 lists options deprecated in Sybase IQ 15.0.

- DEFAULT_DBSPACE option

- DEFAULT_DISK_STRIPING option

- DEFAULT_KB_PER_STRIPE option

- FP_LOOKUP_SIZE option

- FP_LOOKUP_SIZE_PPM option

- MAX_TEMP_SPACE_PER_CONNECTION option

- MPX_AUTOEXCLUDE_TIMEOUT option

  For syntax and examples, see Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- MPX_HEARTBEAT_FREQUENCY option

For syntax and examples, see Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- MPX_IDLE_CONNECTION_TIMEOUT option

    For syntax and examples, see Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- MPX_MAX_CONNECTION_POOL_SIZE option

    For syntax and examples, see Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- MPX_MAX_UNUSED_POOL_SIZE option

    For syntax and examples, see Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- SORT_COLLATION option

- SUBQUERY_CACHING_PREFERENCE option

- SUBQUERY_FLATTENING_PERCENT option

- SUBQUERY_FLATTENING_PREFERENCE option

System tables

The following tables are new or changed in Sybase IQ 15.0. Unless otherwise noted, see Chapter 9, "System Tables," in *Reference: Building Blocks, Tables, and Procedures* for syntax and examples.

- ISYSIQINFO system table

- ISYSIQMPXLOGINPOLICYOPTION system table

    For syntax and examples, see Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- ISYSIQMPXSERVER system table

    For syntax and examples, see Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

System views

The following views are new or changed in Sybase IQ 15.0. Unless otherwise noted, see Chapter 8, "System Views," in *Reference: Building Blocks, Tables, and Procedures* for syntax and examples.

- SYSCOLUMNS system view
- SYSFKEY system view
- SYSIQBACKUPHISTORY system view
- SYSIQBACKUPHISTORYDETAIL system view
- SYSIQDBFILE system view
- SYSIQDBSPACE system view
- SYSIQIDX system view
- SYSIQINFO system view
- SYSIQJOINIDX system view
- SYSIQJOINIXCOLUMN system view
- SYSIQJOINIXTABLE system view
- SYSIQMPXLOGINPOLICYOPTION system view

  For syntax and examples, see Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- SYSIQMPXSERVER system view

  For syntax and examples, see Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

- SYSIQPARTITIONCOLUMN system view
- SYSIQTAB system view
- SYSLOGINPOLICY system view
- SYSLOGINPOLICYOPTION system view
- SYSPARTITION system view
- SYSPARTITIONKEY system view
- SYSPARTITIONSCHEME system view
- SYSSUBPARTITIONKEY system view
- SYSUSER system view

System procedures

The following system procedures are new or changed in Sybase IQ 15.0. Unless otherwise noted, see Chapter 7, "System Procedures," in *Reference: Building Blocks, Tables, and Procedures* for syntax and examples.

- sa_dependent_views procedure

  Returns the list of all dependent views for a given table or view.

- sa_get_user_status procedure

  For details see the *SQL Anywhere Server SQL Reference* section sa_get_user_status procedure at http://infocenter.sybase.com/help/topic/com.sybase.help.sqlanywhere/dbreference_en11/sa-get-user-status.html.

- sa_server_option system procedure has two new options to reset message log size and number of log archives while the database server is running.

- sp_expireallpasswords procedure

- sp_iqaddlogin procedure

- sp_iqbackupdetails procedure

- sp_iqbackupsummary procedure

- sp_iqcheckdb procedure has a new dropleaks mode and syntax to support table partitions.

- sp_iqcolumn procedure

- sp_iqconnection procedure

- sp_iqcopyloginpolicy procedure

- sp_iqdbspace procedure

- sp_iqdbspaceobjectinfo

- sp_iqemptyfile procedure

- sp_iqfile procedure

- sp_iqindexinfo procedure

- sp_iqlmconfig procedure

- sp_iqmodifyadmin procedure

- sp_iqmodifylogin procedure
- sp_iqmpxinfo procedure
- sp_iqpassword procedure
- sp_iqmpxincconnpoolinfo procedure
- sp_iqmpxincheartbeatinfo procedure
- sp_iqobjectinfo procedure
- sp_iqrestoreaction procedure
- sp_iqstatistics procedure
- sp_iqtransaction procedure

Startup and database administration utilities

The following utilities are new or changed in Sybase IQ 15.0.

- CP874toUTF8 utility

    See "CP874toUTF8 utility" in Chapter 3, "Database Administration Utilities," in *Utility Guide* for changes.

- start_iq utility

    The new switches -iqmpx_failover, -iqmpx_ov, -iqmpx_reclaimwriterfreelist, and -iqmpx_sn are documented in "Server startup utility (start_iq)" in Appendix A, "Multiplex Reference," in *Using Sybase IQ Multiplex*.

    The new switches -iqmsgnum and -iqmsgsz are documented in "Starting the database server," in Chapter 1, "Running the Database Server," in *Utility Guide*.

- The default cache size for the switch -ch has been increased from 128mb to 256mb., and is documented in "Starting the database server," in Chapter 1, "Running the Database Server," in *Utility Guide*.

Index

Numerics

1-byte indexes   6, 7
2-byte indexes   6, 7
32-bit network clients   77
3-byte indexes   8
64-bit network clients   77

A

Administration Guide
    changes in this release   80
ALTER DBSPACE statement   81
ALTER FUNCTION statement   81
ALTER INDEX statement   81, 82
ALTER TABLE MOVE statement   52
ALTER TABLE statement   82
altering
    functions   81
analyzing output   71
APPEND_LOAD option
    partitioned tables   4
archiving message log   42
ASDIR environment variable   73
ASIQ-12_7.sh
    name change   70
ASIQPORT environment variable   73
ASIQTIMEOUT environment variable   73
ASLOGDIR environment variable   73
ASTMP environment variable   73

B

backup operations
    summary   89
backups
    iqmsg file   44
    message log   44
    message log archives   44
behavior changes
    UPDATE return code   76
bulk loading   1
bulk loading client data   1

C

cache default sizes   77
cache size
    setting with sa_server_option   52, 78
collation
    changes in 15.0   65
    CP874toUTF8 utility   13, 66
    dbcollat   66
    improvements in 15.0   12
    SORT_COLLATION option   87
    SORTKEY function   66, 86
    system tables   66
columns
    SYSCOLUMNS system view   88
COMMENT ON LOGIN POLICY statement   83
COMMENT statement   83
compatibility views
    SYSFILE   61
connections
    maximum   36
    number of   89
consolidated views
    SYSCATALOG   61
    SYSINDEX   61
conventions
    documentation   x, xi
    syntax   x
    typographic   xi
CP874toUTF8 utility   66
CREATE DATABASE ENCRYPTED clause   36

Z