SYBASE®

An **SAP®** Company

New Features Summary

SAP Sybase IQ 16.0 SP01

Contents

Contents

Read This First: Helpful Hints Before You Start Using SAP Sybase IQ 16

Although the SAP® Sybase® IQ 16 New Features Summary describes all new SAP Sybase IQ functionality, some features may require additional action on your part to take advantage of the new architecture.

Load Performance Configuration Issues

Customers upgrading from a previous release, for example, may need to change some initial compatibility options or rebuild wide columns to accommodate different datatypes. The new load engine provides better performance, but requires changes to the default memory allocation to use all available hardware resources efficiently.

This topic highlights migration issues and features which impact load performance. See the appropriate topics in the SAP Sybase IQ core documentation set for details. *Migration (Linux and UNIX)* and *Migration (Windows)* are new administration manuals in 16.0, describing the steps for upgrading your database, steps to install maintenance releases, and information on upgrading to the role-based security model. *Administration: Load Management* is also new in 16.0, describing data import and export procedures.

NBit

Continuous `NBit` dictionary compression replaces 1, 2, and 3 byte dictionary compression as the default column storage mechanism in 16.0. All datatypes except LOB (character and binary) and BIT datatypes can be `NBit` columns.

The `IQ UNIQUE` column constraint determines whether a column loads as `Flat FP` or `NBit FP`. An `IQ UNIQUE` n value set to 0 loads the column as `Flat FP`. An n value greater than 0 but less than the `FP_NBIT_AUTOSIZE_LIMIT` creates a `NBit` column initially sized to *n*. Columns without an `IQ UNIQUE` constraint implicitly load as `NBit` up to the auto-size limit.

Using `IQ UNIQUE` with an *n* value less than the auto-size limit is not necessary. The load engine automatically sizes all low or medium cardinality columns as `NBit`. Use `IQ UNIQUE` in cases where you want to load the column as `Flat FP` or when you want to load a column as `NBit` when the number of distinct values exceeds the value of the **FP_NBIT_AUTOSIZE_LIMIT** option.

Loads and Large Memory

Large memory represents the maximum amount of memory that SAP Sybase IQ can dynamically request from the OS for temporary use. Because some load operations may require more large memory than the 2GB default provides, adjust the startup options that

control large and cache memory allocation based on the total amount of available physical memory.

As a general rule, large memory requirements represent one third of the total available physical memory allocated to SAP Sybase IQ. To ensure adequate memory for the main and temporary IQ stores, set the **–iqlm**, **–iqtc**, and **–iqmc** startup parameters so that each parameter receives one third of all available physical memory allocated to SAP Sybase IQ.

In most cases, you should allocate 80% of total physical memory to SAP Sybase IQ to prevent SAP Sybase IQ processes from being swapped out. Adjust actual memory allocation to accommodate other processes running on the same system. For example, on a machine with 32 cores and 128GB of total available physical memory, you would allocate 100GB (approximately 80% of the 128GB total) to SAP Sybase IQ processes. Following the general rule, you would set the **–iqlm**, **–iqtc**, and **–iqmc** parameters to 33GB each.

See *-iqlm iqsrv16 Server Option* and *-iqmc iqsrv16 Server Option* in the *Utility Guide*.

Index Changes

Changes to FP and HG indexes take advantage of the new column compression mechanism and improve load performance.

Index	Description
New Fast Projection (FP) Indexes	Take advantage of the new continuous NBit dictionary compression, which replaces 1, 2, and 3 byte dictionary compression. FP(1), FP(2), and FP(3) indexes roll over to NBit(8), NBit(16), and NBit(24) respectively. If FP_NBIT_IQ15_COMPATIBILITY='OFF', IQ UNIQUE constraints applied to the column determine whether the column loads as Flat FP or NBit. See *Fast Projection (FP) Index* in *Administration: Database*.

Index	Description
New tiered HG index structure	Decouples load performance from HG index size. In 15.x, load throughput could degrade as the amount of data in an HG index increased. As the index grew, loading the same amount of data could take more time. The new tiered structure decouples load performance from the HG index size to increase throughput. The CREATE_HG_WITH_EXACT_DISTINCTS option determines whether newly created HG indexes are tiered or non-tiered. This option is ON in all new 16.0 databases and all 16.0 databases migrated from 15.x. To take advantage of the new structure, set this option to OFF. See *CREATE_HG_WITH_EXACT_DISTINCTS Option* in *Reference: Statements and Options*. Use **sp_iqrebuildindex** to convert non-tired HG indexes to tiered HG and vice-versa.

Stored Procedures

New stored procedures return information about column indexes and constraints.

Procedure	Description
sp_iqindexmetadata	Returns details about column indexes, including the index types (Flat FP, NBit, HG, and tiered HG), distinct counts, IQ UNIQUE *n* value, and NBit dictionary size. See *sp_iqindexmetadata Procedure* in *Reference: Building Blocks, Tables, and Procedures*.
sp_iqcolumnmetadata	Returns FP index metadata for one or more user tables or all tables in the database. See *sp_iqcolumnmetadata Procedure* in *Reference: Building Blocks, Tables, and Procedures*.
sp_iqindexrebuildwidedata	Identifies wide columns that you must rebuild before they are available for read/write activities. Output includes statements that you can use with **sp_iqrebuildindex** to rebuild the columns. See *sp_iqindexrebuildwidedata Procedure* in *Reference: Building Blocks, Tables, and Procedures*.

Procedure	Description
sp_iqrebuildindex	Rebuilds FP indexes (`Flat FP as NBit`, or `NBit as Flat FP`) and HG indexes (single HG as tiered HG, or tiered HG as single HG). Before you can insert or update new data, you must rebuild all columns greater than 255 bytes wide. The `index_clause` can reset `IQ UNIQUE` n to an explicit value from 0 (to recast an `NBit` column to `Flat FP`) up to the limits defined in the `FP_NBIT_AUTOSIZE_LIMIT` and `FP_NBIT_LOOKUP_MB` options. **sp_iqrebuildindex** also enables read-write access to columns that contain large object (LOB) data. LOB columns migrated from 15.x databases are read-only until you run **sp_iqrebuildindex**. The estimated cardinality for `NBit` columns with an `IQ UNIQUE` value below or equal to the `FP_NBIT_AUTOSIZE_LIMIT` is stored as 0 regardless of the `FP_NBIT_IQ15_COMPAT-IBILITY` setting. This affects the value returned from **sp_iqindexmetadata**. See *sp_iqrebuildindex Procedure* in *Reference: Building Blocks, Tables, and Procedures*.

Database Options
Some database options are not enabled to take advantage of 16.0 features. Maintaining limited compatibility after a database upgrade provides some flexibility to transition existing applications.

Option	Description
FP_NBIT_IQ15_COMPATIBILI-TY	Provides tokenized `FP` support similar to that available in 15.x. This option is ON by default in all 16.0 databases upgraded from 15.x and OFF in all newly created 16.0 databases. • If this option is ON, the database engine uses the `MINI-MIZE_STORAGE`, `FP_LOOKUP_SIZE`, and `FP_LOOKUP_SIZE_PPM` options to optimize column storage. These options are ignored in 16.0. • If this option is OFF, the database engine ignores 15.x options and columns conform to SAP Sybase IQ `NBit` storage options. Set this option to OFF to take advantage of `NBit` column compression.

Option	Description
CREATE_HG_WITH_EX-ACT_DISTINCTS	Determines whether new HG indexes explicitly created with a **CREATE INDEX** command, or implicitly creating or altering a table with a PRIMARY KEY or a FOREIGN KEY declaration, are tiered or non-tiered. This option is ON 16.0 databases upgraded from 15.x and all newly created 16.0 databases. If this option is ON, all new HG indexes are non-tiered. To take advantage of the new tiered HG index structure, set this option to OFF. Use **sp_iqrebuildindex** to convert non-tiered HG indexes to tiered HG and vice-versa.
CRE-ATE_HG_AND_FORCE_PHYSI-CAL_DELETE	Governs 16.0 delete behavior for tiered HG indexes. This option determines whether SAP Sybase IQ performs a physical delete immediately or defers the delete to a point later in the load. CREATE_HG_AND_FORCE_PHYSICAL_DELETE is ON by default, which instructs SAP Sybase IQ to perform physical deletes.
REVERT_TO_V15_OPTIMIZER	REVERT_TO_V15_OPTIMIZER forces the query optimizer to mimic SAP Sybase IQ 15.x behavior. RE-VERT_TO_V15_OPTIMIZER='ON' by default in all 16.0 databases upgraded from 15.x. REVERT_TO_V15_OPTI-MIZER='OFF' by default in all newly created SAP Sybase IQ 16.0 databases. If you plan to use SAP Sybase IQ hash partitioning features, set the REVERT_TO_V15_OPTIMIZER ='OFF' in databases upgraded from 15.x to 16.0.

See also

New Features

New features were introduced in SAP Sybase IQ 16.0.

Sybase Control Center

Sybase Control Center is a Web-based administrative console that replaces Sybase Central.

Sybase Central is not supported in SAP Sybase IQ 16.0 and is not available as an installation option. Sybase Control Center replaces Sybase Central.

Sybase Control Center provides a single comprehensive Web administration console for managing real-time performance, status, and availability monitoring of large-scale Sybase enterprise servers. Sybase Control Center combines a rich client administrative console, agents, common services, and tools for managing Sybase products. It provides historical monitoring, threshold-based alerts and notifications, and intelligent tools for identifying performance and usage trends. Sybase Control Center (SCC) architecture includes an SCC server and SCC agent. A single SCC server can monitor up to 100 SAP Sybase IQ servers.

Use Sybase Control Center to administer SAP Sybase IQ:

- Simplex servers
- Multiplex servers
- Logical servers
- Remote servers
- External logins
- Databases
- Tables
- Views
- Materialized views
- Domains
- Text configuration objects
- Text indexes
- Sequence generators
- Spatial support
- Security and user management
- Dbspaces
- DB Files
- Events
- External Environments

- Funtions
- Procedures
- Web Services

Note: Sybase Control Center for SAP Sybase IQ (the IQMAP 3.2.7 plugin for SCC 3.2.8) is included with the SAP Sybase IQ installer for these platforms:

- Win64
- Solaris Sparc
- Solaris x64
- LinuxAMD64

Sybase Control Center for SAP Sybase IQ is not supported on:

- AIX64
- HP64
- IBM LinuxPPC64

To install Sybase Control Center for SAP Sybase IQ on Win32 and LinuxAMD32, you must install SCC 3.2.8 and IQMAP 3.2.7 via the SCC 3.2.8 DVD.

See also

Concealment of Sensitive Information in Output Requests

Sensitive information - passwords and encryption keys - is now hidden when statements are printed to request level logs, are logged by diagnostic tracing, or when expressions containing sensitive information are used as column names.

Additionally, output of the **REWRITE** function, connection_property('LastStatement') and any statements recorded in event tracing is sanitized. Any expression that represents a password, key, or sensitive data is hidden.

For example, encrypt('a', 'abc', 'AES') now outputs as encrypt('***', '***', AES).

The sensitive parameters in the built-in procedures ENCRYPT and DECRYPT are now hidden.

Sensitive parameters are now hidden in these system procedures:

- sa_verify_password
- sp_password
- xp_startmail
- xp_startsmtp

Passwords and keys are now hidden in these statements:

- GRANT CONNECT
- CREATE DATABASE
- START DATABASE
- DROP DATABASE
- CREATE EXTERNLOGIN
- SET TEMPORARY OPTION secure_feature_key = <key>

Connection Blocking

Connection blocking blocks a statement until a server write-lock becomes available. The new connection blocking behavior enables competing write-transactions to establish their snapshot versions more efficiently and collaboratively, minimizing lock contention during table writes to the IQ main store and RLV store.

In earlier releases, when several clients attempted to perform write-operations against the same table, the server used a queuing mechanism to queue up waiting write requests while the table was locked by the write-transaction. Client applications used retry logic to send write-lock requests until the table was available for locking. However, the transaction snapshot version was created at transaction creation time, prior to obtaining the table write lock. It was possible, therefore, for a competing write-transaction to obtain the lock when the server put your transaction to sleep. In this scenario, when the server woke your transaction, the server noted that your version was older than the transaction committed by the other user, causing a future-version error and an inability to write.

Connection blocking is the mechanism where the server puts the connection to sleep and wakes it when the other connection releases the table write- lock. When you enable connection blocking, any transaction attempting to obtain a lock conflicting with an existing lock held by another transaction waits until either the conflicting lock is released, or until the blocking timeout value is reached. If the lock isn't released when the blocking timeout value is reached, then your waiting transaction receives an error.

When you disable connection blocking, your connection attempting to obtain the lock held by another connection receives an error immediately. The server rolls back the current operation.

Two database options support connection blocking:

- **BLOCKING database option –** Controls the behavior in response to locking conflicts. If the blocking option is set to ON, any transaction attempting to obtain a lock that conflicts with an existing lock held by another transaction waits until every conflicting lock is released or until the blocking timeout value is reached. If the lock is not released within the blocking timeout period (in milliseconds), then an error is returned for the waiting transaction.
- **BLOCKING_TIMEOUT database option –** Controls how long (in milliseconds) a transaction waits to obtain a lock.

In release 16.0, transaction snapshot versioning is delayed until your transaction obtains the table write-lock. This guarantees your transaction is blocked if another writer exists. Delayed transaction snapshot versioning is enabled by default, and requires no configuration.

See these topics in the SAP Sybase IQ 16.0 documentation:

*   *Administration: Database > Manage Transactions and Versioning > Transaction Blocking*
*   *Administration: In-Memory Row-Level Versioning > Manage Blocking in the RLV Store*
*   *Reference: Statements and Options > Database Options > Alphabetical List of Options > BLOCKING Option*
*   *Reference: Statements and Options > Database Options > Alphabetical List of Options > BLOCKING_TIMEOUT Option*

See also
*   *Database Option Changes* on page 47
*   *RLV Data Store for Write-Optimized Storage* on page 13

Global Transaction Resiliency

DML read-write transactions on multiplex writer nodes now survive temporary communication failures between coordinator and writer nodes and temporary failure of the coordinator due to server failure, shutdown or failover. Prior to SAP Sybase IQ 16.0 such failures guaranteed failure of the global transaction on the writer node and required the user to roll back the entire transaction.

When a user connects to a writer node and executes read-write DML commands against shared objects, the writer starts a global transaction on the coordinator. The transaction starts on an internal internode communication (INC) connection from writer to coordinator.

For example, INSERT or LOAD commands on shared database objects are global transactions. If a failure occurs, the global transaction and corresponding INC connection is suspended.

If the temporary failure resolves within a user-defined timeout period, the global transaction continues as if there was no failure. The user can commit, roll back, or continue the transaction. Use the MPX_LIVENESS_TIMEOUT option to set the timeout period, which defaults to an hour (default value 3600).

If the failure persists longer than the user-defined timeout period, the global transaction cannot resume and the user must roll back the whole transaction.

If there is a DML command actively executing while the failure happens, the command behavior depends on the user-defined timeout and the command type.

To check connection status (active or suspended), use the **sp_iqconnection** system procedure on a writer node or **sp_iqmpxsuspendedconninfo** system procedure on a coordinator. Run

sp_iqmpxincstatistics for a snapshot of the aggregate statistics of the INC status since server startup.

See these new and updated topics in the SAP Sybase IQ 16.0 documentation:

* *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures*
* *Reference: Statements and Options > Database Options > Alphabetical List of Options > MPX_LIVENESS_TIMEOUT Option*
* *Administration: Multiplex > Manage Transactions > Troubleshoot Transactions*

This feature does not affect transactions initiated on the coordinator.

Communication Failure or Coordinator Failure and Restart During Global Transaction

If internode communication (INC) fails or the coordinator fails or is shut down during a writer-initiated global transaction, transactions suspend and resume automatically if the INC is restored before a user-specified timeout expires.

Delays in command execution may indicate INC suspend and resume operations. If INC is interrupted, the coordinator suspends a global transaction for an hour. The transaction resumes successfully as soon as INC is restored. If the timeout value elapses, the transaction fails. Set the **MPX_LIVENESS_TIMEOUT** database option to change the timeout period.

The following cases describe the behavior of writer nodes.

Table 1. Communication to Coordinator Resumes Before Timeout

Writer Command Status	Command Behavior	Result
Actively executing command	Command suspends, except for ROLLBACK, which executes locally on writer.	Command succeeds.
New DML command	Command suspends and resumes, except for ROLLBACK and ROLLBACK TO SAVEPOINT, which execute locally on the writer.	If communication is restored, resumed commands succeed.

Table 2. Communication Failure Exceeds Timeout

Writer Command Status	Command Behavior	Result
Suspended DML command on connection	The suspended command fails and returns an error about the non-recoverable state of the transaction.	You must roll back the transaction. Rollback happens automatically if the suspended command is COMMIT or ROLLBACK to SAVEPOINT.
No suspended DML command on connection	The next command returns an error about the non-recoverable state of the transaction.	You must roll back the transaction.

To check connection status, use the **sp_iqconnection** system procedure on a writer node or the **sp_iqmpxsuspendedconninfo** system procedure on a coordinator.

Run **sp_iqmpxincstatistics** for a snapshot of the aggregate statistics of the INC status since server startup.

Note: If a global transaction initiated from a writer node modifies both global and local persistent objects (for example, an SA base table and an IQ base table), and the coordinator fails during commit, global object changes may be committed while local object changes are lost. This is consistent with a scenario that updates both local and proxy tables in the same transaction, where "best effort" is used to commit both local and global components of a transaction.

Login Redirection

SAP Sybase IQ 16.0 provides load balancing when a user tries to log into an overloaded node by redirecting the attempted login to a node that is less loaded in the same logical server.

Use the LOGIN_REDIRECTION logical server policy to enable redirection. You can now define your own logical server policies and assign them to one or more logical servers to specify server behavior. Connection parameters, logical server policies, login policies, and user privileges determine the nodes available to a particular user for processing.

For example, the following creates a logical server policy lsp1 where login redirection is enabled and associates the new policy with logical server ls1:

```
CREATE LS POLICY lsp1 LOGIN_REDIRECTION=ON
ALTER LOGICAL SERVER ls1 POLICY lsp1
```

To specify the target logical server, connect using the LOGICALSERVER parameter. To specify the target server role instead, connect with the NODETYPE parameter. For example, to connect an application executing INSERT or UPDATE statements to a member node with the writer role, use NODETYPE=WRITER.

You can specify the current coordinator without knowing the host name by using the built-in logical server, COORDINATOR.

If login redirection is enabled and a connection is allowed, SAP Sybase IQ redirects connections when:

- The initial connection node is not a member of the target logical server.
- The initial connection node is a member of the target logical server but has a role other than that requested.
- The initial node is a member of the target logical server and has the requested role, but the user has reached the limit of maximum connections on the current logical server member node.

If the initial request fails, SAP Sybase IQ returns a list of available nodes to the user. If available member nodes of the requested role are unavailable, the connection fails.

See also
- *Connection Changes* on page 46
- *Logical Server Policy Option Changes* on page 55
- *SQL Statement Changes* on page 62

RLV Data Store for Write-Optimized Storage

Use the new row-level versioning (RLV) data store in your simplex database to perform row-level updates, inserts, and deletes, in real-time. When a table is registered for storage in the RLV data store, multiple users can write to different rows of the same table concurrently.

For information on architecture, configuration, blocking, locking, monitoring, merging, and troubleshooting row-level versioning, see the new manual *Administration: In-Memory Row-Level Versioning*.

SAP Sybase IQ 16.0 introduces a new data store:

- **RLV Store –** The in-memory data store optimized for high-performance row-level updates. The RLV store acts as a staging area for write events. If a table is enabled for row-level versioning, then all **LOAD TABLE**, **INSERT**, **UPDATE**, and **DELETE** events write directly to the RLV store. In-memory data in the RLV store is periodically merged into the IQ main store. You can set parameters for automatic merges, and you can merge on-demand.

The IQ main store is optimized for high performance queries, rather than table writes. The RLV store adds high-performance write capability to your database. You can use row-level versioning in conjunction with your existing insert mechanisms such as bulk loading. For example, your applications can make real-time updates during business hours, while overnight a bulk load operation bypasses the RLV store and loads updated data directly to the IQ main store.

You do not require a separate license to use the RLV store. However, row-level versioning requires:

- A simplex database
- A single RLV dbspace
- You have enabled row-level snapshot versioning
- You have enabled RLV storage on the specified table, or on all base tables

If a table is RLV-enabled, SAP Sybase IQ table DDL behavior is unchanged. DDL events on a table (adding a column, for example) require an exclusive lock on that table.

See these topics in the SAP Sybase IQ 16.0 documentation:

- *Administration: In-Memory Row-Level Versioning > About In-Memory Row-Level Versioning*
- *Administration: In-Memory Row-Level Versioning > Appendix: SQL Reference*

See also
- *Database Option Changes* on page 47
- *Tables and Views Changes* on page 78
- *SQL Statement Changes* on page 62
- *Connection Blocking* on page 9

Scale-Up and Scale-Out Performance Features

Scale up features take advantage of multiple cores and high parallelism in a single servers to provide high load and query performance. Multiplex and logical servers scale-out to multiple servers to improve query performance.

Fully Parallel Bulk Load

New load engine effectively scales data loads to make use of all additional hardware.

SAP Sybase IQ 16.0 uses all available cores in the box for better throughput, regardless of the number of columns you load. The higher number of cores, the better load performance you get. Due to high parallelism, you may encounter larger demands on your I/O system and you may need to increase your storage bandwidth to realize improvements.

Tiered High Group Index Structure

New tiered High Group (HG) index structure decouples load performance from HG index size.

In SAP Sybase IQ 15, load throughput could degrade as the amount of data in an HG index increased. As the index grew, loading the same amount of data could take more time. The new tiered structure decouples load performance from the HG index size to increase throughput.

The CREATE_HG_WITH_EXACT_DISTINCTS option determines whether newly created HG indexes are tiered or non-tiered. If this option is ON, all new HG indexes are non-tiered. To take advantage of the new structure, set this option to OFF. Use **sp_iqrebuildindex** to convert non-tiered HG indexes to tiered HG and vice-versa .

Additional Information
See these topics in the SAP Sybase IQ documentation:

- **Database Migration –** *Migration > Database Upgrades*
- **Database Administration –** *Administration: Database > Index Data > Index Types Comparison > High_Group (HG) Index*
- **Database Options –** *Reference: Statements and Options > Database Options > Alphabetical List of Options >CREATE_HG_WITH_EXACT_DISTINCTS*
- **Stored Procedures –**

    - *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqrebuildindex Procedure*
    - *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqindexmetadata Procedure*

NBit Dictionary Compression

Continuous NBit dictionary compression replaces 1, 2, 3 byte dictionary compression. All datatypes except LOB (character and binary) and BIT datatypes may be NBit columns.

NBit dictionary compression provides more efficient compression and greater performance than 1, 2, 3 byte dictionary compression. The IQ UNIQUE column constraint explicitly determines whether the column loads as Flat FP or NBit FP. An IQ UNIQUE *n* value set to 0 specifies a Flat FP. An *n* value greater than 0 but less than FP_NBIT_AUTOSIZE_LIMIT creates a NBit column initially sized to *n*.

Columns without an IQ UNIQUE constraint implicitly load as NBit up to the limits defined in the FP_NBIT_AUTOSIZE_LIMIT. Use IQ UNIQUE in cases where you want to load the column as Flat FP or when you want to load a column as NBit when the number of distinct values exceeds the auto-size limit.

Note: BIT, BLOB,and CLOB data types do not support NBit dictionary compression. If FP_NBIT_IQ15_COMPATIBILITY='OFF', a non-zero IQ UNIQUE column

specification in a CREATE TABLE or ALTER TABLE statement that includes these data types returns an error.

NBIT Levels and corresponding IQ UNIQUE cutoffs

SAP Sybase IQ 15 supported three levels of tokenization: FP(1), FP(2) and FP(3). SAP Sybase IQ 16.0 supports 31 possible NBit levels.

NBit Value	IQ UNIQUE
1	2
2	4
3	8
4	16
5	32
6	64
7	128
8	256 [1]
9	512
10	1024
11	2048
12	4096
13	8192
14	16384
15	32768
16	65536 [2]
17	131972
18	262144
19	524288
20	1048576
21	2097152
22	4194304
23	8388608

NBit Value	IQ UNIQUE
24	16777216 [3]
25	33554432
26	67108864
27	134217728
28	268435456
29	536870912
30	1073741824
31	2147483647
[1] Equivalent to FP(1) [2] Equivalent to FP(2) [3] Equivalent to FP(3)	

Additional Information
See these topics in the SAP Sybase IQ documentation:

- **Database Migration –** *Migration > Database Upgrades*

- **Database Administration –** *Administration: Database > Index Data > Index Types Comparison > Fast Projection (FP) Index*

- **Statements –**
  - *Reference: Statements and Options > SQL Statements > CREATE TABLE Statement*
  - *Reference: Statements and Options > SQL Statements > ALTER TABLE Statement*

- **Database Options –**
  - *Reference: Statements and Options > Database Options > Alphabetical List of Options > FP_NBIT_AUTOSIZE_LIMIT*
  - *Reference: Statements and Options > Database Options > Alphabetical List of Options > FP_NBIT_ENFORCE_LIMITS*
  - *Reference: Statements and Options > Database Options > Alphabetical List of Options > FP_NBIT_IQ15_COMPATIBILITY*
  - *Reference: Statements and Options > Database Options > Alphabetical List of Options > FP_NBIT_LOOKUP_MB*
  - *Reference: Statements and Options > Database Options > Alphabetical List of Options > FP_NBIT_ROLLOVER_MAX_MB*

- **Stored Procedures –**

  - *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqrebuildindex Procedure*
  - *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqcardinality_analysis Procedure*
  - *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqindexmetadata Procedure*

Hash Partitioning

Hash Partitioning reduces resource usage and internode communication.

To take advantage of hash partitioning, you need to create a new hash partitioned table, choose an appropriate high cardinality column as a partitioning key, and reload the data from an existing table into the newly created hash partitioned table. With hash-partitioning, queries on a large number of rows (of the order of 1 billion rows) use less temp space and perform better with a large number of cores. In a multiplex environment, hash-partitioning significantly reduces shared temp bandwidth requirements, allowing for better scale-out across multiple servers.

Data Affinity

Data affinity accesses cached data to boost query performance in multiplex servers.

For optimum scale-out performance, Sybase recommends that you have 4-8 multiplex nodes with total memory across the cluster that can hold most of the working set data for the queries. Affinity is automatically enabled for multiplex databases. No user intervention or maintenance is required.

Non-Partitioned Query Scale-Out

For non-partitioned data, new `JOIN` and `GROUP BY` algorithms reduce the amount of intermediate results that are exchanged during distributed query execution.

These new features improve query performance in cases limited by shared temp or interconnect bandwidth. The optimizer automatically selects these new operators on a cost basis. No user intervention is required but some limitations apply:

- The optimizer has a limited cost mode for the new algorithms. It may not use them in all cases where they may provide benefit or may use them sub-optimal situations
- New grouping algorithms are limited to non-null data

Cache Ejection

Caches are handled more efficiently with multiple writers in SAP Sybase IQ Multiplex.

Each server in the Multiplex now has to re-read only a fraction of pages to determine what data pages have changed since the last read. Only the data pages that have changed are re-read and

others are just read from the cache. This reduces I/O and thereby gains in overall load and queries.

Login Redirection

Logical servers can redirect connections at login time.

Login Redirection lets you dynamically change the logical server group by adding or removing multiplex nodes to rebalance the workload as connections are redirected to servers in the logical server group at any given time. All redirection logic is contained in the client libraries and the server. There is no need to update existing applications. Applications will try to connect to any server in the multiplex, and if login redirection is turned on for the login policy in question, the connection will be redirected to the appropriate node in the logical server group.

LDAPUA

`LDAP` user authentication (`LDAPUA`) allows client applications to send user names and passwords to SAP Sybase IQ for authentication by an `LDAP` server.

Authentication using the `LDAP` server allows you to use server-wide passwords instead of SAP Sybase IQ or application-specific passwords. This feature requires an enterprise-wide, central `LDAP` server.

Security Management

New features pertaining to security, user management, and LDAP user authentication.

Role-Based Security

Role-based security, also called role-based access control (RBAC), allows the breakdown of privileged operations into fine-grained sets that can be individually granted to users.

It provides granular access control in a flexible and easy to use manner and enforces the separation of duties principle by making it possible to assign disjointed sets of privileged tasks to separate autonomous entities. Finally, it allows full control over which system privileges and roles can be granted to other entities.

Role-based security is based on the concepts of system and object-level privileges and roles. A privilege controls the ability to perform a specific task and can be granted directly to a user. A role is an entity to which system privileges and object-level privileges can be associated. Granting system and object-level privileges to a role and then granting the role to users allows users to inherit the privileges of the role. You can also grant roles to other roles to create a hierarchical security structure.

Each system privilege or role has designated administrators whose responsibility it is to control which users are granted the system privilege or role, which users can themselves act as administrators of the system privilege or role, or which users can do both.

See these topics:

*   *Administration: User Management and Security > Security Management*
*   *Administration: User Management and Security > Security Management > Plan and Implement Role-Based Security*
*   *Administration: User Management and Security > Security Management > Roles*
*   *Administration: User Management and Security > Security Management > Privileges*
*   *Administration: User Management and Security > Security Management > Users*

Ability to Impersonate Another User

The SET USER system privilege allows a user to assume the exact roles and system privileges of (impersonate) another user. Once begun, impersonation of another user remains in effect until it is terminated or the current session ends.

To grant the SET USER system privilege:

*   Each grantee (user doing the impersonating) must be an existing user or role.
*   Each target user (user being impersonated) must be an existing user with a login password.
*   The grantor (user granting the system privilege) must have administrative rights to the SET USER system privilege.
*   Each grantee must have been granted at least all roles and system privileges with same or higher administrative options as the target user.

    Note: Failure to meet the final condition does not prevent the system privilege from being granted to the user; it does, however, prevent the user from executing the SETUSER command to successfully impersonate another user.

Note: This functionality is not currently implemented in Sybase Control Center.

See these topics:

*   *Administration: User Management and Security > Security Management > Impersonation*

See also
*   *SQL Statement Changes* on page 62

Minimum Number of Role Administrators

The ability to manage roles is a privilege granted to select users and roles.

Through the process of dropping roles and users or revoking roles and system privileges, you might have a scenario where there are no users or roles left with sufficient system privilege to manage the remaining roles. This configurable option ensures that no role can be revoked or dropped if doing so reduces the remaining number of role administrators below a set value.

See these topics:

*   *Administration: User Management and Security > Security Management > Roles > User-Defined Roles > Role and Global Role Administrators*

See also
*   *Database Option Changes* on page 47

User Management

Enhancements to user management now support dual control password change and the elimination of an potential lockdown scenario due to failed login attempts.

Dual Control of User Password Management

The dual control option provides an added level of security for users resetting passwords for other users. It ensures that no single user knows or controls the password of another user.

This option requires that two users participate when resetting the password of a user other than themselves. Each user defines a portion of the password. The reset password then becomes the union of the two passwords.

Note: This functionality is not currently implemented in Sybase Control Center.

See these topics:

*   *Administration: User Management and Security > Security Management > Manage Passwords*

See also
*   *SQL Statement Changes* on page 62
*   *Login Policy Option Changes* on page 56

Automatic Unlocking of User Accounts

User accounts are locked if users exceed the maximum failed login attempts limit. Once locked, the account must be manually unlocked by a user granted the MANAGE ANY USER system permission.

However, since users with the MANAGE ANY USER system privilege are also subject to lockdown due to failed login attempts, you might have a scenario whereby all users with the necessary system privilege to unlock accounts are themselves locked out. The automatic unlocking of user accounts feature prevents this lock-down scenario from occurring.

See these topics :

*   *Administration: User Management and Security > Security Management > Login Policies*

See also
*   *Login Policy Option Changes* on page 56

SAP Sybase IQ LDAP Authentication

SAP Sybase IQ can be integrated into any existing enterprise-wide directory access framework based on Lightweight Directory Access Protocol (LDAP), a widely accepted international standard.

Integration with LDAP user authentication supports:

- Authentication using searched distinguished name (DN)
- Failover to a secondary LDAP server for high availability
- Automatic failback to previously failed servers
- Integration with OpenLDAP third-party libraries
- Secure communication with LDAP servers
- Efficient design for frequent, short-lived connections
- Extensibility to multiple domains and multiple LDAP servers

See these topics in the SAP Sybase IQ 16.0 Administration: User Management and Security documentation:

- See *Administration: User Management and Security > Advanced Security Options in SAP Sybase IQ > LDAP Authentication Support in SAP Sybase IQ*

See also
- *Database Option Changes* on page 47
- *Login Policy Option Changes* on page 56
- *SQL Statement Changes* on page 62
- *Tables and Views Changes* on page 78
- *Stored Procedure Changes* on page 70
- *Utility Option Changes* on page 81

Transport Layer Security for Multiplex

SAP Sybase IQ supports encrypted communication connections on multiplex.

Use the -ec server option to specify FIPS-certified RSA encryption on packets transmitted to and from clients.

See these topics in the SAP Sybase IQ 16.0 documentation:

- *Administration: User Management and Security > Security Management > Transport Layer Security > Setting up transport-layer security*
- *Utility Guide > start_iq Database Server Startup Utility > start_iq Server Options > -ec iqsrv16 database server option*

Running Privileged System Procedures

SAP Sybase IQ provides increased security when running privileged system procedures.

For new databases, system procedures that perform authorized tasks now execute with the privileges of the person running the procedure rather than the privileges of the creator of the procedure.

This increased security can be enabled on migrated databases, with some considerations. See *Migration > Database Upgrades > Changes to System Procedures that Perform Privileged Operations*

See also
* *Privileged System Procedure Execution* on page 45

Shared System Temporary Store

A multiplex configuration with shared temporary storage can use the IQ_SHARED_TEMP dbspace as a shared system temporary store instead of requiring a separate local store for each secondary server. The shared system temporary store simplifies multiplex configuration, improves performance, and supports distributed query processing.

On multiplex systems:

* When you set the logical server policy option TEMP_DATA_IN_SHARED_TEMP ON, SAP Sybase IQ creates all temporary objects on the IQ_SHARED_TEMP dbspace. You must restart secondary nodes after setting this option or after adding a read-write file to the shared temporary store. (If the shared temporary store contains no read-write file, or if you do not restart secondary nodes, data instead writes to IQ_SYSTEM_TEMP.)
* Temporary user objects (such as tables or table indexes) that you create using the **IN IQ_SYSTEM_TEMP** clause go in either IQ_SYSTEM_TEMP or IQ_SHARED_TEMP, depending on the value of the logical server option TEMP_DATA_IN_SHARED_TEMP:
  * If TEMP_DATA_IN_SHARED_TEMP is 'OFF', objects go in IQ_SYSTEM_TEMP.
  * If TEMP_DATA_IN_SHARED_TEMP is set 'ON', objects go in IQ_SHARED_TEMP.

  SAP Sybase IQ does not support creating temporary user objects using the **IN IQ_SHARED_TEMP** clause.
* The **WITH STOP SERVER** clause automatically shuts down all servers in the logical server. These statements support **WITH STOP SERVER**:
  * **ALTER LOGICAL SERVER**
  * **ALTER LS POLICY**
  * **CREATE LOGICAL SERVER**
  * **DROP LOGICAL SERVER**

- If you use **ALTER LS POLICY ... WITH STOP SERVER** to change the TEMP_DATA_IN_SHARED_TEMP option 'ON|OFF', all servers in that logical server shut down automatically. You must restart the servers to force the logical server to place temporary data in the store specified by the TEMP_DATA_IN_SHARED_TEMP option.
- If you use **ALTER LS POLICY** to set TEMP_DATA_IN_SHARED_TEMP 'OFF', the logical server starts placing temporary data in the SYSTEM temporary area after the next normal server startup.
- You can also change the TEMP_DATA_IN_SHARED_TEMP value indirectly using **CREATE LOGICAL SERVER**, **ALTER LOGICAL SERVER**, or **DROP LOGICAL SERVER** statements and the **WITH STOP SERVER** clause.

See also

Table Partitioning

Table partitioning can improve performance by dividing large tables into smaller, more manageable storage objects. SAP Sybase IQ 16.0 now supports range, hash, and hash-range partitioning schemes.

- Range partitioning divides large tables by a range of partition-key values established for each partition. As part of an information life cycle management strategy, range partitioning can shorten backup and restore times; provide a finer level of granularity for data validation; and support tiered storage.
- Hash partitioning maps data to partitions based on *partition-key* values processed by an internal hashing function. Hash partitioning can enhance join performance on large tables and distributed queries (DQP).
- Hash-range partitioning provides the benefits of hash partitioning and range partitioning. Hash partitioning provides the best distributed query performance; range sub-partitioning enhances administrative tasks as part of an information life cycle management strategy.

Note: Range-partitions and composite partitioning schemes, like hash-range partitions, require the separately licensed VLDB Management option.

Additional Information

- *Administration: Database > Manage SAP Sybase IQ Database Objects > Table Partitions*
- *Administration: Load Management > Bulk Loads with the LOAD TABLE Statement > Considerations for Partitioned Table Loads*
- *Reference: Statements and Options > SQL Statements > ALTER TABLE Statement*
- *Reference: Statements and Options > SQL Statements > CREATE TABLE Statement*

System Secure Feature Keys

The system secure feature key now allows you to create customized secure feature keys that are assigned to a specific users, limiting users' access to only the features secured by the administrator for that key.

Customized secure feature keys are managed by system procedures.

See these topics in the SAP Sybase IQ 16.0 documentation:

• *Administration: User Management and Security > Security Management > Data Security*.

New Features

Performance Improvements

Distributed query processing and micro query performance has improved in SAP Sybase IQ 16.0. The new Cache Ejection feature improves performance by eliminating unnecessary cache page reads.

Cache Ejection

This enhancement prevents the buffer manager from performing unnecessary rereads of cached pages from disk, improving performance through a reduction of disk I/O and internode communication.

A cache ejection policy determines how a system identifies and replaces stale cache data. Each server in a multiplex system has its own cache.

DQP Performance Improvements

Distributed Query Processing (DQP) performance is improved in SAP Sybase IQ 16.0.

Performance Improvement	Description
Affinity-based work allocation	For DQP eligible queries, SAP Sybase IQ identifies preferred nodes for scans of specific row ranges and partition IDs. Once a query establishes an affinity for a piece of data, subsequent queries that need the same data use the same node for their work. Use the database options **AFFINITY_AUTOEXCLUDE_TIMEOUT** and **CACHE_AFFINITY_PERCENT** to control allocation.
Asymmetric sort merge join algorithm	The new *asymmetric sort merge join* algorithm improves performance in distributed queries that join and sort high-cardinality data. Advanced DBAs can control algorithm choice using the **JOIN_PREFERENCE** option.
Early aggregation	An optimization where distincting, grouping, and aggregates are pushed into a sort, rather than sorting all rows first, then performing any aggregation outside the sort as a second step. Advanced DBAs can control algorithm choice using the **AGGREGATION_PREFERENCE** option.

Performance Improvement	Description
Option to use DQP over the network	DQP over the network keeps data in memory in the temporary cache instead of in the IQ_SHARED_TEMP dbspace. The option that provides the best performance may vary by query, and depends on the performance of the network relative to the I/O system. Use the **DQP_ENABLED** logical server policy option to control DQP processing. The **DQP_ENABLED** and **DQP_ENABLED_OVER_NETWORK** database options control DQP at the connection level.

See also
* *Query Plan Changes* on page 58

Micro Query Performance Improvement

In previous releases, the QUERY_PLAN option was ON by default, meaning that every query generated a query plan regardless of the size of the query. In SAP Sybase IQ 16.0, use the new QUERY_PLAN_MIN_TIME option when QUERY_PLAN is ON to prevent the server from generating query plans for micro queries, resulting in a database server performance improvement.

See these topics in the SAP Sybase IQ 16.0 documentation:

* *Performance and Tuning Guide > Optimizing Queries and Deletions > Planning Queries > Using Query Plans > Preventing Query Plan Generation for Micro Queries*
* *Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN Option*
* *Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_MIN_TIME Option*

See also
* *Database Option Changes* on page 47

Discontinued Features

If you are upgrading from an earlier release of SAP Sybase IQ 16.0, note the features and functionality discontinued in this release. Discontinued features and functionality are not supported by SAP.

JDBC™ 3 Connectivity

SAP Sybase IQ no longer supports JDBC 3 connectivity.

SAP Sybase IQ 16.0 supports JDBC 4 and supplies the `jodbc4.jar` file in `$IQDIR16/java` on UNIX or `%IQDIR16%\Java` on Windows. The `jodbc.jar` file is no longer supplied.

SAP Sybase IQ 16.0 also supports jConnect™ for JDBC 7.0.

JDBC-Based Server Classes

JDBC-based server classes are no longer supported and generate errors.

Discontinued Item	Description
JDBC-Based Server Classes	All support for JDBC-based server classes is discontinued. If your application uses any of the following, update your application to use the corresponding odbc server class: • asejdbc • iqjdbc • sajdbc

LOAD TABLE Statement: Discontinued Syntax

Several clauses of the LOAD TABLE statement, deprecated in previous releases, are discontinued in SAP Sybase IQ 16.0.

The STRIP ON clause of the LOAD TABLE statement is discontinued. Using the STRIP ON clause will result in an error. Use STRIP RTRIM instead.

The BLOCK FACTOR clause of the LOAD TABLE statement, deprecated in release 15.2, is discontinued. Using the BLOCK FACTOR clause will result in an error.

The BLOCK SIZE clause of the LOAD TABLE statement, deprecated in release 15.2, is discontinued. Using the BLOCK SIZE clause will result in an error.

The UNLOAD FORMAT clause of the LOAD TABLE statement, deprecated in release 15.2, is discontinued. Using the UNLOAD FORMAT clause will result in an error.

Time Series Analysis

The time series analysis feature has been removed from SAP Sybase IQ 16.0. The Visual Numerics Inc. (VNI) third-party external libraries `IMSL C Stat` and `C Math` have been removed from the product.

These time series manuals have been removed from the SAP Sybase IQ documentation collection:

* *Time Series Guide*
* *IMSL C Numerical Library User Guide Volume 2 of 2: C Stat Library*

Time series functions are no longer bundled with the server. Users can still continue to build their own UDFs.

See also
* *Discontinued Time Series SQL Functions* on page 32

Discontinued Database Server Startup Options

The **start_iq** options -cw, -cm, and -ec are no longer supported.

Discontinued Item	Description
-cw iqsrv16 database server startup option (deprecated)	Deprecated in previous releases, support for this **start_iq** option has been removed in release 16.0.
-cm iqsrv16 database server startup option	Support for this **start_iq** option has been removed in release 16.0.
-ec iqsrv16 database server startup option parameter TLS_TYPE=ECC	Support for the ECC algorithm has been removed in release 16.0. **TLS_TYPE=ECC** is no longer supported.

Discontinued Indexes

JOIN, LD, FP(1), FP(2), and FP(3) indexes are not supported in SAP Sybase IQ 16.0.

Discontinued Item	Description
JOIN Indexes	SAP Sybase IQ no longer supports `JOIN` indexes. All references to this item have been removed from the documentation.
	Customers who are migrating a database from a previous edition must drop all `JOIN` indexes before upgrading to SAP Sybase IQ 16.0.
LD Indexes	`LD` (LOWDISK) indexes were replaced by `HG` indexes in release 12.0. If you have a very old database, note that SAP Sybase IQ no longer supports `LD` indexes. All references to this item have been removed from the documentation.
	Customers who are migrating a database from a previous edition must drop all `LD` indexes before upgrading to SAP Sybase IQ 16.0.
FP(1), FP(2), FP(3) indexes	SAP Sybase IQ no longer supports `FP(1)`, `FP(2)`, `FP(3)` indexes.
	Continuous `NBit` dictionary compression replaces 1, 2, 3 byte dictionary compression. `FP(1)`, `FP(2)`, `FP(3)` indexes rollover to `NBit` 8, 16, 24 respectively.

Note: Failing to drop all `LD` or `JOIN` indexes before an ALTER DATABASE UPGRADE in SAP Sybase IQ 16.0 causes the upgrade to fail, and the server will return a message that indicates that a database upgrade is not possible when these objects are in use.

To recover, shutdown the 16.0 server, open the database with SAP Sybase IQ 15.x, and drop all `LD` and all `JOIN` indexes before proceeding with the migration.

Discontinued Login Policy Option

The dqp_enabled login policy option is not a valid login policy option in SAP Sybase IQ 16.0.

On upgrade, if you had configured this option for any login policy, these options will be lost.

After upgrading, configure the dqp_enabled logical server policy option for the logical servers where Distributed Query Processing (DQP) is required. You can to choose whether DQP should use the network for intermediate result sets, or use shared temporary storage.

For details, see *DQP_ENABLED Option* and *DQP_ENABLED_OVER_NETWORK Option* in *Administration: Multiplex*.

Discontinued Time Series SQL Functions

Time series functions are not supported in SAP Sybase IQ 16.0.

Discontinued Item	Description
TS_ARMA_AR Function [Aggregate]	Discontinued SQL function. Not supported.
TS_ARMA_CONST Function [Aggregate]	Discontinued SQL function. Not supported.
TS_ARMA_MA Function [Aggregate]	Discontinued SQL function. Not supported.
TS_AUTOCORRELATION Function [Aggregate]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA Function [Aggregate]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA_OUTLIER Function [Aggregate]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA_RESULT_AIC Function [Scalar]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA_RESULT_AICC [Scalar]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA_RESULT_BIC Function [Scalar]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA_RESULT_FORECAST_ERROR Function [Scalar]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA_RESULT_FORECAST_VALUE Function [Scalar]	Discontinued SQL function. Not supported.

Discontinued Item	Description
TS_AUTO_ARIMA_RE-SULT_MODEL_P Func-tion [Scalar]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA_RE-SULT_MODEL_Q Func-tion [Scalar]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA_RE-SULT_MODEL_S Func-tion [Scalar]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA_RE-SULT_MODEL_D Func-tion [Scalar]	Discontinued SQL function. Not supported.
TS_AUTO_ARIMA_RE-SULT_RESIDUAL_SIG-MA Function [Scalar]	Discontinued SQL function. Not supported.
TS_AUTO_UNI_AR Func-tion [Aggregate]	Discontinued SQL function. Not supported.
TS_BOX_COX_XFORM Function [Aggregate]	Discontinued SQL function. Not supported.
TS_DIFFERENCE Func-tion [Aggregate]	Discontinued SQL function. Not supported.
TS_DOUBLE_ARRAY Function [Scalar]	Discontinued SQL function. Not supported.
TS_ESTIMATE_MISS-ING Function [Aggregate]	Discontinued SQL function. Not supported.
TS_GARCH Function [Ag-gregate]	Discontinued SQL function. Not supported.
TS_GARCH_RESULT_A Function [Scalar]	Discontinued SQL function. Not supported.
TS_GARCH_RE-SULT_AIC Function [Sca-lar]	Discontinued SQL function. Not supported.
TS_GARCH_RE-SULT_USER [Scalar]	Discontinued SQL function. Not supported.
TS_INT_ARRAY Function [Scalar]	Discontinued SQL function. Not supported.

Discontinued Item	Description
TS_LACK_OF_FIT Function [Aggregate]	Discontinued SQL function. Not supported.
TS_LACK_OF_FIT_P Function [Aggregate]	Discontinued SQL function. Not supported.
TS_MAX_ARMA_AR Function [Aggregate]	Discontinued SQL function. Not supported.
TS_MAX_AR-MA_CONST Function [Aggregate]	Discontinued SQL function. Not supported.
TS_MAX_ARMA_LIKE-LIHOOD Function [Aggregate]	Discontinued SQL function. Not supported.
TS_MAX_ARMA_MA Function [Aggregate]	Discontinued SQL function. Not supported.
TS_OUTLIER_IDENTIFI-CATION Function [Aggregate]	Discontinued SQL function. Not supported.
TS_PARTIAL_AUTO-CORRELATION Function [Aggregate]	Discontinued SQL function. Not supported.
TS_VWAP Function [Aggregate]	Discontinued SQL function. Not supported.

See also
* *SQL Function Changes* on page 61

Discontinued TLS and HTTPS Certificate and Certificate_Password Protocol Options

TLS and HTTPS `certificate` and `certificate_password` protocol options were renamed to `identity` and `identity_password`, respectively, in an earlier release.

Attempting to use the discontinued names `certificate` and `certificate_password` in SAP Sybase IQ 16.0 results in an error.

Discontinued Tables and Views

The system tables ISYSIQJOINIDX and ISYSIQJOINIXCOLUMN are no longer supported. Their corresponding system views SYSIQJOINIDX and SYSIQJOINIXCOLUMN are no longer supported.

System Tables

System Table	Description
ISYSIQJOINIDX	SAP Sybase IQ no longer supports this table.
ISYSIQJOINIXCOLUMN	SAP Sybase IQ no longer supports this table.

System Views

System View	Description
SYSIQJOINIDX	SAP Sybase IQ no longer supports this view.
SYSIQJOINIXCOLUMN	SAP Sybase IQ no longer supports this view.

See also

- *Tables and Views Changes* on page 78

Discontinued Tools and Utilities

Sybase Central, the SQL Anywhere debugger, iqdsedit, and iqisql are discontinued in SAP Sybase IQ 16.0.

Debugger

The SQL Anywhere debugger has been discontinued in SAP Sybase IQ 16.0. The debugger was used during the development of SQL stored procedures.

A debugging tool will be available in a future release.

iqdsedit Utility

All support for the iqdsedit utility has been discontinued. Use the DSEdit utility (dsedit.exe) instead.

iqisql Utility

All support for the iqisql utility has been discontinued.

Use the isql Interactive SQL utility instead.

Sybase Central

The Sybase Central graphical administration tool has been discontinued. Sybase Control Center replaces it.

Sybase Control Center is included with your SAP Sybase IQ 16.0 installation.

Deprecated Features

If you are upgrading from an earlier release of SAP Sybase IQ, note the features and functionality deprecated in release 16.0. Deprecated features and functionality are supported for backward compatibility, but are not recommended. Support will be removed in an upcoming release.

Deprecated Database Options

Some database options, and the **DQP_ENABLED** login policy option, are deprecated or discontinued in SAP Sybase IQ 16.0.

Deprecated Item	Description
APPEND_LOAD	The **APPEND_LOAD** option is deprecated. This option still exists but the setting will be ignored.
DQP_ENABLED login policy option	Discontinued. Replaced by **DQP_ENABLED** logical server policy option.
LARGE_DOUBLES_ACCUMULATOR	Discontinued. Replaced by **FLOATING_POINT_ACCUMULATOR** option.
MAX_WARNINGS	Since SAP Sybase IQ no longer supports **JOIN INDEXES**, the **MAX_WARNINGS** option is no longer required.

Deprecated Database Server Startup Options

The **host** parameter of **start_iq** option -x is deprecated in SAP Sybase IQ 16.0.

Deprecated Item	Description
host parameter for -x iqsrv16 database server startup option	The host network connection parameter of the -x option is deprecated in release 16.0.

Deprecated Stored Procedures

Several stored procedures are deprecated in SAP Sybase IQ 16.0.

Deprecated Item	Description
sp_addgroup	With the implementation of role-based security, SAP Sybase IQ no longer supports groups. However, the procedure remains, for backward compatibility.
sp_changegroup	With the implementation of role-based security, SAP Sybase IQ no longer supports groups. However, the procedure remains, for backward compatibility.
sp_dropgroup	With the implementation of role-based security, SAP Sybase IQ no longer supports groups. However, the procedure remains, for backward compatibility.
sp_iqcardinality_analysis	**sp_iqcardinality_analysis** no longer returns an index type value or index recommendation. Users are advised to run Index Advisor for suggestions about additional column indexes. **sp_iqcardinality_analysis** is deprecated and will be removed in a future release.

See also
- *Stored Procedure Changes* on page 70

Deprecated System Tables and Views

Several system tables and views are deprecated in SAP Sybase IQ 16.0.

Table	Description
ISYSGROUP	Table no longer used in SAP Sybase IQ 16.0; however, SYSGROUP system view remains for backward compatibility.
ISYSUSERAUTHORITY	Table no longer used in SAP Sybase IQ 16.0; however, SYSUSERAUTHORITY system view remains for backward compatibility.

Deprecated User Management Functionality

Authorities and groups are deprecated in SAP Sybase IQ 16.0 and are replaced by role-based security. However, role-based security is backwards compatible with authorities and groups.

Deprecated Item	Description
Authorities	Authorities have become user-defined roles. During the upgrade of an existing database to SAP Sybase IQ 16.0, each deprecated SAP Sybase IQ 15.x authority is replaced with an equivalent user-defined compatibility role, which contains an equivalent set of privileged tasks.
Groups	Groups have become user-extended roles. During the upgrade of an existing database to SAP Sybase IQ 16.0, each existing group in the database is replaced with an equivalent user extended role, using the same name. All members of the original group are automatically granted membership in the equivalent new role. Any authorities granted to the original group are automatically replaced with the equivalent compatibility role and granted to the new user-extended role.

Deprecated Environment Variables

Several environment variables are deprecated in SAP Sybase IQ 16.0.

The following SYBASE_JRE6 Java environment variables are deprecated:

- SYBASE_JRE6_64
- SYBASE_JRE6
- SYBASE_JRE6_32
- SYBASE_JRE5_64

See also
- *JRE and Java Runtime Environment Variable Changes* on page 54

Deprecated Features

Behavior Changes

Several behavior changes have been introduced in SAP Sybase IQ 16.0.

Backward Compatibility: Changes to Default Behavior

The default behavior of SAP Sybase IQ has changed in some areas.

See also

Constraint Violation Checks

The order in which the database server evaluates check constraints and unique constraints has changed in SAP Sybase IQ 16.0. This behavior change occurs only for **LOAD** or **INSERT** statements containing the clause IGNORE CONSTRAINT CHECK <x>, UNIQUE <y>.

Check constraints are now evaluated before unique constraints. This behavior change may lead to different load results compared to 15.4, depending on the constraints specified, and the data being loaded.

Consider this simple **LOAD TABLE** example comparing 15.4 constraint evaluation behavior with 16.0:

1. Create the table:
```
CREATE TABLE t1( c1 int, c2 int, primary key(c1), CHECK( c1 <
c2 ) );
```
2. Load values:

```
LOAD TABLE t1 { 4,3 }, {4,5 } FROM 'sample_data.dat' IGNORE
CONSTRAINT CHECK 0, UNIQUE 0
```

In 15.4, the database server detects unique constraint violations before check constraint violations:

- HG index detects a unique violation on { 4, 5 }, so row 2 is removed.
- This leaves row 1 which contains { 4, 3 }. This violates the check constraint, so row 1 is removed.
- The result is zero rows in the table.

In 16.0, check constraint violations are detected first, resulting in a different outcome:

- Row { 4,3 } is discarded due to the check constraint violation.
- HG index only sees one row so there is no unique constraint violation.
- The resulting row { 4, 5 } does not violate the check constraint.
- The result is one row in the table: { 4, 5 }.

A related behavior change affects column defaults in statements containing the clause IGNORE CONSTRAINT CHECK. In SAP Sybase IQ 16.0, column default violations are never ignored.

Constraint Violation Checks: Message Log and Row Log

In previous versions, the row ID listed in msg.log and row.log indicated the actual row ID. In SAP Sybase IQ 16.0, the row ID represents the *logical row ID* of the specified constraint violation.

Database Options: Changed Defaults

If you are upgrading from an earlier release of SAP Sybase IQ, note that default values of some database options changed in release 16.0.

JOIN_SIMPLIFICATION_THRESHOLD Option

The default value for the JOIN_SIMPLIFICATION_THRESHOLD database option changed to 12 in SAP Sybase IQ 16.0.

MIN_PASSWORD_LENGTH Option

The default value for the MIN_PASSWORD_LENGTH database option changed to 3 in SAP Sybase IQ 16.0.

In previous releases, any length or an empty password were permitted.

QUERY_PLAN Option

The default value for the QUERY_PLAN database option changed to OFF in SAP Sybase IQ 16.0.

In previous releases, the default was ON.

QUERY_DETAIL Option

The default value for the QUERY_DETAIL database option changed to ON in SAP Sybase IQ 16.0.

In previous releases, the default was OFF.

QUERY_PLAN_AFTER_RUN Option

The default value for the QUERY_PLAN_AFTER_RUN database option changed to ON in SAP Sybase IQ 16.0.

In previous releases, the default was OFF.

QUERY_TIMING Option

The default value for the QUERY_TIMING database option changed to ON in SAP Sybase IQ 16.0.

In previous releases, the default was OFF.

Reserved Words

Reserved words in SAP Sybase IQ 16.0.

Item	Description
Object Names	Unlike previous releases, reserved words in SAP Sybase IQ 16.0 cannot be used as object names, unless enclosed in brackets or double quotes.
	A SAP Sybase IQ 15 database could contain tables, columns, and other objects named *row*. In SAP Sybase IQ 16.0, *row* is a reserved word and cannot be used as an object name.
	To use a reserved word as an object name, enclose the object name in brackets (regardless of the QUOTED_IDENTIFIER setting) or double quotes (if QUOTED_IDENTIFIER='ON' [default]):
	```// QUOTED_IDENTIFIER ON | OFF\nselect * from [row];\nalter table row2 rename [row] to\ncol_row;\n\n// QUOTED_IDENTIFIER='ON'\nselect "row" from row2;\nalter table "row" rename rownew;```
array	New reserved word in SAP Sybase IQ 16.0.
json	New reserved word in SAP Sybase IQ 16.0.

Item	Description
row	New reserved word in SAP Sybase IQ 16.0.
rowtype	New reserved word in SAP Sybase IQ 16.0.
unnest	New reserved word in SAP Sybase IQ 16.0.
varray	New reserved word in SAP Sybase IQ 16.0.

Backward Compatibility: Migration Considerations

If you are upgrading from an earlier release of SAP Sybase IQ, several new features and functionality are backward compatible.

See also

GRANT/REVOKE Syntax

The syntax to grant and revoke system privileges in the role-based security model has changed significantly.

However, support for the previous syntax continues. One-to-one mappings between the authority-based and role-based models and group-based and role-based syntax assist in the transition to the new role-based syntax.

See:

- *Migration > Upgrading to Role-Based Security > Changes to the GRANT Statement Syntax*.

- *Migration > Upgrading to Role-Based Security > Changes to the REVOKE Statement Syntax.*

Role-Based Security in Stored Procedures

The enhanced role-based security model of SAP Sybase IQ 16.0 retains backward compatibility with authority-based syntax.

All user-defined stored procedures, functions, and queries created on SAP Sybase IQ 15.x databases will continue to run after they are upgraded.

Privileged System Procedure Execution

As part of role-based security, the way in which privileged system procedures run has changed.

In pre-16.0 databases, privileged system procedures ran with the privileges of its owner, referred to as the SYSTEM PROCEDURE DEFINER model. In new 16.0 databases, privileged system procedures run with the privileges of the person executing it, referred to as the SYSTEM PROCEDURE INVOKER model.

To minimize potential loss of functionality as a result of this behavior change, after upgrading using the default options, all pre-16.0 system procedures continue to run using the SYSTEM PROCEDURE DEFINER model. Any system procedures introduced in 16.0 or later run using the SYSTEM PROCEDURE INVOKER model.

See:

- *Migration > Upgrading to Role-Based Security > Changes to System Procedures that Perform Privileged Operations*
- *Migration > Upgrading to Role-Based Security > Managing System Procedure Execution*

Connection Changes

SAP Sybase IQ 16.0 enhances database connections.

Table 3. Connection Parameter Changes

Parameter	Description
LogicalServer [LS]	Specifies the target logical server, enabling redirection of a connection from one multiplex node to another. When unspecified, The logical server defaults to the setting of the *default_logical_server* option in the user login policy. See *Administration: Database > Connect to Servers and Databases > Connection and Communication Parameters > Connection Parameters > LogicalServer Connection Parameter [LS]*
NodeType	Connects an application to a logical server member node that has a specific role. See *Administration: Database > Connect to Servers and Databases > Connection and Communication Parameters > Connection Parameters > NodeType Connection Parameter*
Redirect	Controls login redirection at the connection level. See *Administration: Database > Connect to Servers and Databases > Connection and Communication Parameters > Connection Parameters > Redirect Connection Parameter*

See also
- *Backward Compatibility: Changes to Default Behavior* on page 41
- *Backward Compatibility: Migration Considerations* on page 44
- *Database Option Changes* on page 47
- *JRE and Java Runtime Environment Variable Changes* on page 54
- *Logical Server Policy Option Changes* on page 55
- *Login Policy Option Changes* on page 56
- *ODBC Driver Changes* on page 57
- *Query Plan Changes* on page 58
- *SQL Function Changes* on page 61
- *SQL Statement Changes* on page 62
- *Stored Procedure Changes* on page 70

Database Option Changes

SAP Sybase IQ 16.0 includes database option changes.

Option	Description
AFFINITY_AUTOEXCLUDE_TIMEOUT	The amount of time before SAP Sybase IQ removes a shut down node from the affinity map and reassigns its partitions to other nodes.
	Allowed values are 0 to 10080 minutes (1 week). The default is 10 minutes.
AGGREGATION_PREFERENCE	The allowed value range is now -6 to 6.
	These values are new in SAP Sybase IQ 16.0:
	• **4** – prefer aggregation with a distinct/grouping sort.
	• **5** – prefer aggregation with a sort if grouping columns include all the partitioning keys of a hash partitioned table.
	• **6** – prefer aggregation with a hash if grouping columns include all the partitioning keys of a hash partitioned table.
	• **-4** – avoid aggregation with a distinct/grouping sort.
	• **-5** – avoid aggregation with a sort if grouping columns include all the partitioning keys of a hash partitioned table.
	• **-6** – avoid aggregation with a hash if grouping columns include all the partitioning keys of a hash partitioned table.
ALLOW_SNAPSHOT_VERSIONING	New option. Applies to all base tables in the database (as opposed to RLV-enabled tables only). Restricts table versioning for all base tables to either table-level or row-level snapshot versioning. This option does not apply to the IQ catalog store.
	The default is YES.

Option	Description
ANSI_SUBSTRING [TSQL]	The default changed to ON.
	Note: Default values in recent releases: • **16.0** – ON • **15.4** – OFF • **15.3** – OFF • **15.0 - 15.2 ESD3** – ON
APPEND_LOAD	The **APPEND_LOAD** option is deprecated. This option still exists but the setting will be ignored.
BASE_TABLES_IN_RLV	New option. Registers new tables in the RLV store, enabling row-level versioning for IQ main store tables. RLV-enabled tables are eligible for multiple writer concurrent access. You can override this setting at the table level using the **CREATE TABLE** statement. Has no effect on IQ main store tables created before you set this option ON. Has no effect on tables from an upgraded database.
BLOCKING	New allowed value: ON. This setting causes any transaction attempting to obtain a lock that conflicts with an existing lock held by another transaction to wait until every conflicting lock is released, or until the blocking_timeout threshold is reached.
	BLOCKING is not supported on secondary nodes of a multiplex. It is supported on the multiplex coordinator node, and in simplex implementations.
BLOCKING_TIMEOUT	New option. Controls how long (in milliseconds) a transaction waits to obtain a lock.
	BLOCKING_TIMEOUT is not supported on secondary nodes of a multiplex. It is supported on the multiplex coordinator node, and in simplex implementations.
CACHE_AFFINITY_PERCENT	Allowed values are 0 -100.
	Maximum percentage of main cache to use for affinity. Non-affinity data can use this area if insufficient affinity data exists.
CREATE_HG_WITH_EXACT_DIS-TINCTS	Determines whether newly created HG indexes are tiered or non-tiered.This option is ON in all new 16.0 databases and all 16.0 databases migrated from 15.x.To take advantage of the new structure, set this option to OFF.

Option	Description
CREATE_HG_AND_FORCE_PHYSI-CAL_DELETE	Governs 16.0 delete behavior for tiered HG indexes. This option determines whether SAP Sybase IQ performs a physical delete immediately or defers the delete to a point later in the load. **CREATE_HG_AND_FORCE_PHYSICAL_DELETE** is ON by default, which instructs SAP Sybase IQ to perform physical deletes.
DQP_ENABLED	Set the temporary database option **DQP_ENABLED** OFF to disable DQP for the current connection. Set the option ON (the default value) to enable DQP for the current connection, but only when DQP is enabled by the policy option for the logical server of the current connection.
DQP_ENABLED_OVER_NETWORK	Set the temporary database option **DQP_ENA-BLED_OVER_NETWORK** to ON to enable DQP over the network for the current connection. The **DQP_ENABLED** database option must also be set to ON and the DQP_EN-ABLED logical server policy option must not be set to 1. If the DQP_ENABLED logical server policy option is set to 2, then this database option has no affect and all queries will use DQP over the network.
ENABLE_ASYNC_IO	New option. Allows a DBA to enable or disable the asynchronous IO used by the RLV persistence log for row-level versioning.
FLOATING_POINT_ACCUMULATOR	New option. Replaces **LARGE_DOUBLES_ACCUMULA-TOR**.
FP_LOOKUP_SIZE	Sets the number of lookup pages and reserves cache memory for FP indexes in SAP Sybase IQ databases where the **FP_NBIT_IQ15_COMPATIBILITY** option is ON.
FP_LOOKUP_SIZE_PPM	Controls the amount of main cache allocated to FP indexes in SAP Sybase IQ databases where the **FP_NBIT_IQ15_COMPATIBILITY** option is ON.
FP_NBIT_AUTOSIZE_LIMIT	New option. Limits the number of distinct values that an NBit column can load implicitly. Columns implicitly load as NBit up to the auto-size limits. Columns with an IQ UNIQUE n value set to 0 loads as Flat FP. Columns with an n value greater than 0 but less than the auto-size limit sets the NBit limit to n.
FP_NBIT_ENFORCE_LIMITS	New option. If this option is ON, and an NBit column exceeds explicit or implicit sizing limits, the operation throws an error and rolls back.

Option	Description
FP_NBIT_IQ15_ COMPATIBILITY	New option. Provides tokenized FP support similar to that available in SAP Sybase IQ 15. Defaults to OFF in a newly created 16.0 database. Set to ON in SAP Sybase IQ databases upgraded from 15.x.
FP_NBIT_LOOKUP_MB	New option. Limits the dictionary size for implicit NBit columns. Columns that exceed these limits rollover to Flat FP.
FP_NBIT_ROLLOVER_MAX_MB	New option. Sets the dictionary size (values and counts) for implicit NBit rollovers. If this option is ON, columns that exceed this limit rollover to Flat FP.

Option	Description
JOIN_PREFERENCE	The allowed value range is now -12 to 12. These values changed in SAP Sybase IQ 16.0: • **6** – previously the action was *prefer prejoin*. The new action is *prefer asymmetric sort merge join*. • **-6** – previously the action was *avoid prejoin*. The new action is *avoid asymmetric sort merge join*. These values are new in SAP Sybase IQ 16.0: • **8** – prefer asymmetric sort merge push down join. • **9** – prefer partitioned hash join if the join keys include all the partition keys of a hash partitioned table. • **10** – prefer partitioned hash-push down join if the join keys include all the partition keys of a hash partitioned table. • **11** – prefer partitioned sort-merge join if the join keys include all the partition keys of a hash partitioned table. • **12** – prefer partitioned sort-merge push-down join if the join keys include all the partition keys of a hash partitioned table. • **-8** – avoid asymmetric sort merge push down join. • **-9** – avoid partitioned hash join if the join keys include all the partition keys of a hash partitioned table. • **-10** – avoid partitioned hash-push down join if the join keys include all the partition keys of a hash partitioned table. • **-11** – avoid partitioned sort-merge join if the join keys include all the partition keys of a hash partitioned table. • **-12** – avoid partitioned sort-merge push-down join if the join keys include all the partition keys of a hash partitioned table.
JOIN_SIMPLIFICATION_THRESHOLD	New option - Controls the minimum number of tables being joined together before any join optimizer simplifications are applied.
LOG_DEADLOCKS	New option - Controls whether deadlock reporting is turned on or off.
LOGIN_MODE	New option. Add LDAPUA to the LOGIN_MODE option to use LDAP user authentication.
MAX_WARNINGS	Since SAP Sybase IQ no longer supports **JOIN INDEXES**, the **MAX_WARNINGS** option is no longer required.

Option	Description
MIN_ROLE_ADMINS	New option. Specifies the minimum number of role administrators for each role. Ensures that no role can be dropped if doing so reduces the remaining number of role administrators below a set value. Value range is 1 (default) – 10.
MINIMIZE_STORAGE	Behavior change. Minimizes disk space for columns in 16.0 databases where the **FP_NBIT_IQ15_COMPATIBILITY** option is ON. Ignored where the **FP_NBIT_IQ15_COMPATIBILITY** option is OFF.
NOTIFY_MODULUS	Behavior change – Default value for new SAP Sybase IQ 16.0 database changed from 100000 to 0. Default value remains unchanged for an upgraded pre-16.0 database.
QUERY_DETAIL	Default changed to ON.
QUERY_PLAN	Default changed to OFF.
QUERY_PLAN_AFTER_RUN	Default changed to ON.
QUERY_PLAN_MIN_TIME	New option. Specifies a threshold for query execution in microseconds. The query plan is generated only if query execution time exceeds the threshold. **QUERY_PLAN** must be ON.
QUERY_TIMING	Default changed to ON.
REVERT_TO_V15_OPTIMIZER	New option. Forces the query optimizer behaves as it did in 15.4. Forces the query optimizer to mimic 15.x behavior. This option is OFF in all newly created 16.0 databases. Set to ON in SAP Sybase IQ databases upgraded from 15.x. To take advatage of new DQP and algorithms and hash partitioning features, set this option to OFF in databases upgraded from 15.x.
ROUND_TO_EVEN	New option. When **ROUND_TO_EVEN** option is set to ON, the **ROUND** function rounds half to the nearest even number. When the option is set to OFF (the default), the **ROUND** function rounds half away from zero.
RV_AUTO_MERGE_EVAL_INTERVAL	New option. Configures the evaluation period used to determine when an automated background merge of the row-level versioned (RLV) and IQ main stores should occur. The default is 15 minutes.

Option	Description
RV_MERGE_NODE_MEMSIZE	New option. An automated merge of the row-level versioned (RLV) store and IQ main stores occurs based on the merge thresholds, including **RV_MERGE_NODE_MEM-SIZE**. When this node threshold is exceeded, a merge will be triggered. The default is 75%.
RV_MERGE_TABLE_MEMPERCENT	New option. An automated merge of the row-level versioned (RLV) store and IQ main stores occurs based on the merge thresholds, including **RV_MERGE_TABLE_MEM-PERCENT**. If this table threshold is exceeded, a merge will be triggered for the specific table. The default is 0%.
RV_MERGE_TABLE_NUMROWS	New option. An automated merge of the row-level versioned (RLV) store and IQ main stores occurs based on the merge thresholds, including **RV_MERGE_TABLE_NUM-ROWS**. If this table threshold is exceeded, a merge will be triggered for the specific table. The default is 10000000.
RV_RESERVED_DBSPACE_MB	New option. A portion of the RLV store must be reserved for memory used by data structures during critical operations. The default is The minimum of 50 Mb or half the size of the RLV dbspace.
SNAPSHOT_VERSIONING	New option. Applies to RLV-enabled tables only (as opposed to all base tables in the database). Controls whether RLV-enabled tables are accessed using single-writer table-level versioning, or multiple writer row-level versioning. This option does not apply to the IQ catalog store. The default is TLV.
TRUSTED_CERTIFICATE_FILE	New option. For TLS connections from SAP Sybase IQ to other servers. Specifies the file name containing the certificate of the certificate authorities (CA) that are trusted. The trusted CAs are the signers of the certificates used by: • The external LDAP directory server for use with LDAP User Authentication • The server certificate in use by the SAP Sybase IQ multiplex server for INC and MIPC connections.

See also
• *Backward Compatibility: Changes to Default Behavior* on page 41

JRE and Java Runtime Environment Variable Changes

SAP Sybase IQ 16.0 requires JRE_7. The SYBASE_JRE Java environment variables have changed accordingly.

The SYBASE_JRE7_64, SYBASE_JRE7_32 environment variable specifies the location of the Java Runtime Environment used by Sybase Control Center. On startup, Sybase Control Center checks SCC_JAVA_HOME for Java version definition. If SCC_JAVA_HOME is undefined, Sybase Control Center checks for installed JREs in this order:

- SYBASE_JRE7_64
- SYBASE_JRE7_32

The following SYBASE_JRE6 Java environment variables are deprecated:

- SYBASE_JRE6_64
- SYBASE_JRE6
- SYBASE_JRE6_32
- SYBASE_JRE5_64

See also

Logical Server Policy Option Changes

Changes affected logical server policy options in SAP Sybase IQ 16.0.

Option	Description
DQP_ENABLED	New option. When set to 0, query processing is not distributed. When set to 1 (the default), query processing is distributed as long as a writable shared temporary file exists. When set to 2, query processing is distributed over the network, and the shared temporary store is not used.
LOGIN_REDIRECTION	New option. When ON, enables login redirection for logical servers governed by specified login policy. When OFF (the default), disables login redirection.
REDIRECTION_WAIT-ERS_THRESHOLD	New option. Specifies the number of connections that can queue before SAP Sybase IQ redirects a connection to this logical server. Can be any integer value (defaults to 5).
TEMP_DA-TA_IN_SHARED_TEMP	When ON, all temporary table data and eligible scratch data writes to the shared temporary store, provided that the shared temporary store has at least one read-write file added. You must restart all multiplex nodes after setting this option or after adding a read-write file to the shared temporary store. (If the shared temporary store contains no read-write file, or if you do not restart nodes, data instead writes to IQ_SYSTEM_TEMP.) When OFF (the default), all temporary table data and scratch data writes to the local temporary store.

See *Reference: Statements and Options > SQL Statements > ALTER LS POLICY Statement*.

See also
- *Backward Compatibility: Changes to Default Behavior* on page 41
- *Backward Compatibility: Migration Considerations* on page 44
- *Connection Changes* on page 46

Login Policy Option Changes

There are changes to some login policy options in SAP Sybase IQ 16.0.

Option	Description
AUTO_UNLOCK_TIME	New option. The time period after which locked accounts not granted the MANAGE ANY USER system privilege are automatically unlocked. This option can be defined in any login policy, including the root login policy.
ROOT_AUTO_UNLOCK_TIME	New option. The time period after which locked accounts granted the MANAGE ANY USER system privilege are automatically unlocked. This option can be defined in the root login policy only.
CHANGE_PASSWORD_DU-AL_CONTROL	New option. Requires input from two users, each granted the CHANGE PASSWORD system privilege, to change the password of another user.
LDAP_PRIMARY_SERVER	New option. Specifies the name of the primary LDAP server.
LDAP_SECONDARY_SERVER	New option. Specifies the name of the secondary LDAP server.
LDAP_AUTO_FAILBACK_PERI-OD	New option. Specifies the time period, in minutes, after which automatic failback to the primary server is attempted.
LDAP_FAILOVER_TO_STD	New option. Permits authentication with standard authentication when authentication with the LDAP server fails due to system resources, network outage, connection timeouts, or similar system failures. However, it does not permit an actual authentication failure returned from an LDAP server to fail over to standard authentication.

Option	Description
LDAP_REFRESH_DN	New option. Updates the `ldap_refresh_dn` value in the `ISYSLOGINPOLICYOPTION` system table with the current time, stored in Coordinated Universal Time (UTC).
	Each time a user authenticates with LDAP, if the value of the option `ldap_refresh_dn` in `ISYSLOGINPOLI-CYOPTION` is more recent than the `user_dn` value in `ISY-SUSER`, a search for a new user DN occurs. The `user_dn` value is then updated with the new user DN and the `user_dn_changed_at` value is again updated to the current time.

See *Reference: Statements and Options > SQL Statements > ALTER LOGIN POLICY Statement*.

See also
- *Backward Compatibility: Changes to Default Behavior* on page 41
- *Backward Compatibility: Migration Considerations* on page 44
- *Connection Changes* on page 46
- *Database Option Changes* on page 47
- *JRE and Java Runtime Environment Variable Changes* on page 54
- *Logical Server Policy Option Changes* on page 55
- *ODBC Driver Changes* on page 57
- *Query Plan Changes* on page 58
- *SQL Function Changes* on page 61
- *SQL Statement Changes* on page 62
- *Stored Procedure Changes* on page 70
- *Tables and Views Changes* on page 78
- *Utility Option Changes* on page 81

ODBC Driver Changes

The name of the ODBC driver (DBODBC16.DLL) changed in SAP Sybase IQ 16.0.

In 15.4, the driver name was *SQL Anywhere*. In 16.0, the driver name is *Sybase IQ*.

If you have an ODBC application that references the driver name directly, you will need to update the driver name in the application. If your ODBC application references the datasource name, no application changes are required.

See also
- *Backward Compatibility: Changes to Default Behavior* on page 41
- *Backward Compatibility: Migration Considerations* on page 44
- *Connection Changes* on page 46
- *Database Option Changes* on page 47
- *JRE and Java Runtime Environment Variable Changes* on page 54
- *Logical Server Policy Option Changes* on page 55
- *Login Policy Option Changes* on page 56
- *Query Plan Changes* on page 58
- *SQL Function Changes* on page 61
- *SQL Statement Changes* on page 62
- *Stored Procedure Changes* on page 70
- *Tables and Views Changes* on page 78
- *Utility Option Changes* on page 81

Query Plan Changes

New load execution plans and query plan nodes in SAP Sybase IQ 16.0.

Load Execution Plans

The Interactive SQL Query Plan Viewer now supports SAP Sybase IQ load execution plans. Load execution plans detail the steps that the database engine uses to insert data into a table.

Load plans use the same database options as query execution plans. Text versions of the plan are written to the *.iqmsg* file. HTML versions include a Data Flow Object (DFO) tree that identifies the number of rows processed at each stage of the load. Different SQL statements generate different DFO trees and the same statement may generate different trees for different kind of tables (un-partitioned, range partitioned, hash partitioned, hash-range partitioned, etc.).

Load execution plans support these SQL statements:

- LOAD..INTO
- INSERT..LOCATION
- INSERT..VALUES
- INSERT..INTO
- CREATE INDEX
- ALTER TABLE ADD column (with default value)
- DECLARE CURSOR FOR UPDATE
- PARALLEL IQ statements

To generate a load plan, set the appropriate QUERY_PLAN options, then use the Interactive SQL Query Plan Viewer to create the plan.

Note: Use query plans only to evaluate the efficiency of a particular query or load. Running SAP Sybase IQ with the QUERY_PLAN option set to ON can significantly impact performance, particularly as the volume of **INSERT...VALUE** statements increase.

New query plan nodes

Node Change	Description
Distinct Sort (inserter)	New node indicating early aggregation algorithm is used.
Distinct Sort (retriever)	New node indicating early aggregation algorithm is used.
Grouping Sort (inserter)	New node indicating early aggregation algorithm is used.
Grouping Sort (retriever)	New node indicating early aggregation algorithm is used.
Join (Asymmetric Sort-Merge)	New node indicating an asymmetric sort-merge join is used.
Join (Asymmetric Sort-Merge PushDown)	New node indicating an asymmetric sort-merge join is used.
Order By (Full Rewind)	New node indicating an asymmetric sort-merge join is used. Displays on the right branch of a **Join (Asymmetric Sort-Merge)** node.
Order By (Per-Work-Unit)	New node indicating an operator that performs a sort on a work-unit worth of data. For example, it displays on the left branch under a **Join (Asymmetric Sort-Merge)** node.

Additional Information

- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_DETAIL Option
- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_NAME Option
- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN Option
- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_AFTER_RUN Option
- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_AS_HTML Option

- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_AS_HTML_DIRECTORY Option
- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_MIN_TIME
- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_TEXT_ACCESS Option
- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_TEXT_CACHING Option
- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_ROWS_RETURNED_LIMIT Option
- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_TEMP_SPACE_LIMIT Option
- Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_TIMING Option

See also

SQL Function Changes

Function changes for SAP Sybase IQ 16.0.

Table 4. System Functions

Function	Description
sp_has_role	New function. Returns an integer value indicating whether the invoking user has been granted a specified system privilege or user-defined role. When used for permission checking within user-defined stored procedures, this function can display an error message when a user fails a permission check.
	See *Reference: Building Blocks, Tables, and Procedures > SQL Functions > Alphabetical List of Functions> sp_has_role Function [System]*.

See also
- *Backward Compatibility: Changes to Default Behavior* on page 41
- *Backward Compatibility: Migration Considerations* on page 44
- *Connection Changes* on page 46
- *Database Option Changes* on page 47
- *JRE and Java Runtime Environment Variable Changes* on page 54
- *Logical Server Policy Option Changes* on page 55
- *Login Policy Option Changes* on page 56
- *ODBC Driver Changes* on page 57
- *Query Plan Changes* on page 58
- *SQL Statement Changes* on page 62
- *Stored Procedure Changes* on page 70
- *Tables and Views Changes* on page 78
- *Utility Option Changes* on page 81
- *Discontinued Time Series SQL Functions* on page 32

SQL Statement Changes

Syntax changes for SAP Sybase IQ 16.0.

Statement	Description
ALTER DATA-BASE	New clauses: • RESTART prevents the database from automatically restarting after an upgrade. • SYSTEM PROCEDURE AS DEFINER defines whether to execute system procedures that perform privileged tasks with the privileges of the invoker (the person calling the procedure) or the definer (the owner of the procedure). 　• OFF means all system procedures execute with the privileges of the invoker. 　• ON (default), or not specified means pre-16.0 system procedures execute with the privileges of the definer and 16.0 or later system procedures execute with the privileges of the invoker. See *Reference: Statements and Options > SQL Statements > ALTER DATA-BASE*
ALTER LDAP SERVER	New statement to modify the defined properties for the LDAP server. See *Reference: Statements and Options > SQL Statements > ALTER LDAP SERVER Statement.*
ALTER LOGICAL SERVER	New clauses: • POLICY clause associates a logical server with a user-defined logical server policy. • WITH STOP SERVER clause automatically shuts down all servers in the logical server. AUTO is now a reserved logical server name. See *Reference: Statements and Options > SQL Statements > ALTER LOGI-CAL SERVER Statement.*

Statement	Description
ALTER LOGIN POLICY	New clauses: • DEFAULT_LOGICAL_SERVER sets the target logical server context if you omit LogicalServer from the connection string. • LDAP_PRIMARY_SERVER, LDAP_SECONDARY_SERVER, LDAP_AUTO_FAILBACK_PERIOD, LDAP_FAILOVER_TO_STD, and LDAP_REFRESH_ON allow the defining of SAP Sybase IQ LDAP user authentication related properties in login policies. • ROOT_AUTO_LOCK_TIME (root login policy only) and AUTO_UN-LOCK_TIME let you define user account management related properties in login policies. • CHANGE_PASSWORD_DUAL_CONTROL requires input from two users, each granted the CHANGE PASSWORD system privilege, to change the password of another user. See *Reference: Statements and Options > SQL Statements > ALTER LOGIN POLICY*
ALTER LS POLICY	You can alter user-created policies to control the behavior of associated logical servers in your multiplex. Enable login redirection, specify which nodes in a logical server are available for processing queries, the placement of temporary table data, and how many connections can queue before redirection. New options DQP_ENABLED, LOGIN_REDIRECTION, REDIRECTION_WAIT-ERS_THRESHOLD, and TEMP_DATA_IN_SHARED_TEMP added. New WITH STOP SERVER clause automatically shuts down all servers in the logical server. This statement can also alter the root ls policy to change the value of different ls policy options. See *Reference: Statements and Options > SQL Statements > ALTER LS POL-ICY Statement.*
ALTER ROLE	New statement lets you migrates a compatibility role to a user-defined system role, then automatically drops the compatibility role. See *Reference: Statements and Options > SQL Statements > ALTER ROLE Statement.*
ALTER TABLE	New syntax extensions enable you to either register the table with the RLV store for real-time in-memory updates, or disable RLV storage. New syntax supports changing a table owner. See *Reference: Statements and Options > SQL Statements > ALTER TABLE Statement.*
ALTER USER	New REFRESH_DN clause lets you clear the saved DN and timestamp for a user, which is used during SAP Sybase IQ LDAP user authentication. See *Reference: Statements and Options > SQL Statements > ALTER USER.*

Statement	Description
COMMENT ON	New LDAP_SERVER clause lets you create a comment on a SAP Sybase IQ LDAP server object.
	See *Reference: Statements and Options > SQL Statements > COMMENT.*
COMMENT ON LS POLICY	New clause allows comment to document the purpose or guidelines for using new user-defined logical server policies.
	See *Reference: Statements and Options > SQL Statements > COMMENT ON LS POLICY Statement.*
CREATE DATA-BASE	New SYSTEM PROCEDURE AS DEFINER clause defines whether to execute system procedures that perform privileged tasks with the privileges of the invoker (the person calling the procedure) or the definer (the owner of the procedure).
	• OFF (default), or not specified means all system procedures execute with the privileges of the invoker.
	• ON means pre-16.0 system procedures execute with the privileges of the definer and 16.0 or later system procedures execute with the privileges of the invoker.
CREATE DBSPACE	New syntax extension for creating an RLV store dbspace.
	See *Reference: Statements and Options > SQL Statements > CREATE DBSPACE Statement.*
CREATE LDAP SERVER	New statement to create a new LDAP server for LDAP user authentication.
	See *Reference: Statements and Options > SQL Statements > CREATE LDAP SERVER Statement.*
CREATE LOGICAL SERVER	New POLICY clause associates a logical server with a user-defined logical server policy. AUTO is now a reserved logical server name.
	New WITH STOP SERVER clause automatically shuts down all servers in the logical server.
	See *Reference: Statements and Options > SQL Statements > CREATE LOGI-CAL SERVER Statement.*

Statement	Description
CREATE LOGIN POLICY	New clauses: • DEFAULT_LOGICAL_SERVER sets the target logical server context if you omit LogicalServer from the connection string. • LDAP_PRIMARY_SERVER, LDAP_SECONDARY_SERVER, LDAP_AUTO_FAILBACK_PERIOD, LDAP_FAILOVER_TO_STD, and LDAP_REFRESH_ON allow the defining of SAP Sybase IQ LDAP user authentication related properties in login policies. • ROOT_AUTO_LOCK_TIME (root login policy only) and AUTO_UN-LOCK_TIME let you define user account management related properties in login policies. • CHANGE_PASSWORD_DUAL_CONTROL requires input from two users, each granted the CHANGE PASSWORD system privilege, to change the password of another user. See *Reference: Statements and Options > SQL Statements > CREATE LOGIN POLICY*.
CREATE LS POLICY	New statement lets you create your own policies to control the behavior of associated logical servers in your multiplex. Enable login redirection. Specify which nodes in a logical server are available for processing queries, the placement of temporary table data, and how many connections can queue before redirection. New options DQP_ENABLED, LOGIN_REDIRECTION, REDIRECTION_WAITERS_THRESHOLD, and TEMP_DATA_IN_SHARED_TEMP added. See *Reference: Statements and Options > SQL Statements > CREATE LS POLICY Statement*.
CREATE ROLE	New statement lets you create a new user-defined role, extends an existing user to act as a role, or manages role administrators on a role. See *Reference: Statements and Options > SQL Statements > CREATE ROLE Statement*.

Statement	Description

CREATE TABLE	New syntax enables you to register the table with the RLV store.
	
	``
	CREATE [GLOBAL TEMPORARY] TABLE <table-name> {
	      ENABLE | DISABLE } RLV STORE
	``
	
	New syntax extensions support range, hash, and hash-range partitioning schemes:
	
	``
	partitioning-scheme:
	      { range-partitioning-scheme
	        | hash-partitioning-scheme
	        | hash-range-partitioning-scheme
	       }
	``
	
	Hash partitioning maps data to partitions based on the composite value of a set of partition key columns and an internal hash function. Hash partitioning distributes data to partitions based on composite value of a set of partition key columns, which can enhance query join performance on a large table. Hash-range partitioning is a composite partitioning scheme that distributes data to partitions based on the composite value of a set of partition key columns, an internal hash function, and range sub-partition. Hash-range partitioning maps data to sub-partitions based on a range of sub-partition key values.
	
	See *Reference: Statements and Options > SQL Statements > CREATE TABLE Statement*.

DECLARE LOCAL TEMPORARY TABLE	New syntax enables you to register the table with the RLV store. See *Reference: Statements and Options > SQL Statements > DECLARE LOCAL TEMPORARY TABLE Statement*.
DROP LDAP SERVER	New statement lets you remove an LDAP server from the database after verifying that it is not in a READY or ACTIVE state. See *Reference: Statements and Options > SQL Statements > DROP LDAP SERVER*.
DROP LOGICAL SERVER	New WITH STOP SERVER clause automatically shuts down all servers in the logical server. See *Reference: Statements and Options > SQL Statements > DROP LOGICAL SERVER*
DROP LS POLICY	New statement lets you delete the policy if a user-defined policy is not currently used for any logical server. You cannot drop the root policy. See *Reference: Statements and Options > SQL Statements > DROP LS POLICY Statement*.
DROP ROLE	New statement lets you remove a user-defined role from the database at any time as long as all dependent roles will are left with the minimum required number of administrator users with active passwords. See *Reference: Statements and Options > SQL Statements > DROP ROLE*.

Statement	Description

INSERT	New INSERT...VALUES support for (DEFAULT), DEFAULT VALUES or VALUES() clauses that insert rows with all default values. Assuming that there are 3 columns in table t2, these examples are semantically equivalent:
	
	`INSERT INTO t2 values (DEFAULT, DEFAULT, DEFAULT);`
	
	`INSERT INTO t2 DEFAULT VALUES;`
	
	`INSERT INTO t2() VALUES();`
	
	INSERT...VALUES also supports multiple rows. The following example inserts 3 rows into table t1:
	
	`CREATE TABLE t1(c1 varchar(30));`
	`INSERT INTO t1 VALUES  ('morning'),('afternoon'),`
	`    ('evening');`
	
	SAP Sybase IQ treats all load/inserts as full-width inserts. Columns not explicitly specified on the load/insert statement, the value loaded will either be the column DEFAULT value for the column (if defined) or (if no DEFAULT value is defined).
	
	INSERT no longer supports the START ROW ID option. A LOAD or INSERT statement can still include this option, but SAP Sybase IQ will ignore the value.
	
	INSERT/SELECT from an empty table that has a data type mismatch with the target table returns an error rather than not reporting the mismatch.
	
	You can use the statement to insert data into an RLV-enabled table.
	
	See *Reference: Statements and Options > SQL Statements > INSERT Statement*.

Statement	Description
LOAD TABLE	SAP Sybase IQ treats all load/inserts as full-width inserts. Columns not explicitly specified on the **LOAD**/**INSERT** statement, the value loaded will either be the DEFAULT value for the column (if defined) or NULL (if the column has no DEFAULT value).
	LOAD TABLE no longer supports the START ROW ID option. **LOAD** or **INSERT** statements can still include START ROW ID, but SAP Sybase IQ will ignore the value.
	The load engine no longer supports files in which input rows span file boundaries. A partial input row that occurs at the end of any file will now be treated according to the ON PARTIAL INPUT ROW setting.
	LOAD TABLE no longer supports UNLOAD FORMAT syntax.
	You can use the statement to load a file into an RLV-enabled table.
	The default value of the **LOAD TABLE** NOTIFY clause was 100000 prior to release 16.0. In SAP Sybase IQ 16.0 the default has changed to 0, meaning no notifications are printed by default.
	See *Reference: Statements and Options > SQL Statements > LOAD TABLE Statement*.
INSERT..SELECT	Prior to SAP Sybase IQ 16.0, attempting to use **INSERT..SELECT** on only NULL values when the source and target datatypes mismatched inserted the NULL values even when explicit conversion was required. Now SAP Sybase IQ will properly enforce that explicit conversion is required and correctly return the error `Unable to implicitly convert column <name> from datatype <type>`.
	For example:
	``` CREATE TABLE t_int(c1 INT null); CREATE TABLE t_char(c1 CHAR(5) null); INSERT INTO t_int values (NULL); INSERT INTO t_char select * from t_int; ```
	Prior to 16.0, a NULL value would be inserted into t_char. Starting in 16.0, the required conversion error will occur. To use explicit conversion use CAST or CONVERT, for example:
	``` INSERT INTO t_char select cast(c1 as char(5)) from t_int; ```
	See *Reference: Statements and Options > SQL Statements > INSERT Statement*.

Statement	Description
GRANT/REVOKE *database object permissions*	New clauses. You can grant and revoke the TRUNCATE and LOAD object permissions to users and roles. See *Reference: Statements and Options > SQL Statements > GRANT Database Object Permissions.* See *Reference: Statements and Options > SQL Statements > REVOKE Database Object Permissions.*
GRANT/REVOKE *system privilege*	New statements let you grant and revoke of specific system privileges to specific users, with or without administrative rights. See *Reference: Statements and Options > SQL Statements > GRANT System Privilege Statement.* See *Reference: Statements and Options > SQL Statements > REVOKE System Privilege Statement.*
GRANT/REVOKE CHANGE PASS-WORD	New statement that lets you allow users to manage passwords for other users and administer the CHANGE PASSWORD system privilege. See *Reference: Statements and Options > SQL Statements > GRANT CHANGE PASSWORD Statement.* See *Reference: Statements and Options > SQL Statements > REVOKE CHANGE PASSWORD Statement.*
GRANT/REVOKE SET USER	New statements that allow you to grant and revoke the ability for one user to impersonate another user and administer the SET USER system privilege. See *Reference: Statements and Options > SQL Statements > GRANT SET USER Statement.* See *Reference: Statements and Options > SQL Statements > REVOKE SET USER Statement.*
GRANT/REVOKE ROLE	New statements that allow you to grant and revoke roles to users or other roles, with or without administrative rights. See *Reference: Statements and Options > SQL Statements > GRANT ROLE Statement.* See *Reference: Statements and Options > SQL Statements > REVOKE ROLE Statement.*
SELECT	New FOR JSON clause lets you execute an SQL statement against your database and return the results as a JSON document.
TRUNCATE TA-BLE	New syntax extension let you TRUNCATE partitions in hash partitioned and hash-range partitioned tables. See *Reference: Statements and Options > SQL Statements > TRUNCATE TABLE Statement.*

Statement	Description
VALIDATE LDAP SERVER	New statement allows the validation of proposed changes to the settings of existing LDAP servers before applying them. See *Reference: Statements and Options > SQL Statements > VALIDATE LDAP SERVER Statement.*

See also

Stored Procedure Changes

Catalog stored procedure and system stored procedure changes for SAP Sybase IQ 16.0.

Table 5. Catalog Stored Procedures

Procedure	Description
sa_get_ldapserver_status	New procedure. Allows you to determine the current status of the LDAP server configuration object. See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sa_get_ldapserver_status System Procedure.*

Procedure	Description
sa_get_user_status	New output information. For LDAP user authentication, output includes the distinguished name (DN) for the user and the date and time the distinguished name was found. For dual control password management, output includes current status of a dual control password change, and if active, the user IDs of the two administrators performing the change.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sa_get_user_status System Procedure*.
sa_make_object	Behavior change. Previously, the **sa_make_object** system procedure required RESOURCE authority and could be used to create objects owned by other users, even though RESOURCE authority didn't actually give a user permissions to create objects owned by others. For example, if a user with RESOURCE authority attempted to execute the equivalent **CREATE** statement, the statement would fail.
	Now, when using **sa_make_object**, you must have all appropriate privileges required to create objects owned by self or others. The actual privileges you need depend on the creation operation you are performing.
sa_report_deadlocks	New procedure. Retrieves information about deadlocks from an internal buffer created by the database server.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sa_report_deadlocks System Procedure*.
sa_server_option	Behavior change. New parameters added: `rlv_auto_merge` and `rlv_memory_mb`.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sa_server_option System Procedure*.
sp_auth_sys_role_info	New procedure. Generates a report that maps legacy authorities to corresponding system roles and role IDs. This procedure returns a row for each authority.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sp_auth_sys_role_info System Procedure*.

Procedure	Description
sp_alter_secure_feature_key	New procedure. Alters a previously-defined secure feature key by modifying the authorization key and/or the feature list.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sp_alter_secure_feature_key system procedure.*
sp_create_secure_feature_key	New procedure. Creates a new secure feature key.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sp_create_secure_feature_key system procedure.*
sp_displayroles	New procedure. Displays all roles granted to a user-defined role or a user, or displays the entire hierarchical tree of roles.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sp_displayroles System Procedure.*
sp_drop_secure_feature_key	New procedure. Deletes a secure feature key.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sp_drop_secure_feature_key system procedure.*
sp_list_secure_feature_keys	New procedure. Returns a list of defined secure feature keys.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sp_list_secure_feature_keys system procedure.*
sp_objectpermission	New procedure. Generates a report on object permissions granted to the specified role or user name or object permissions granted on the specified object or dbspace.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures> sp_objectpermission System Procedure.*
sp_proc_priv	New procedure. Generates a report of the minimum system privileges required to run a stored procedure and pass the permission check for the procedure.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sp_proc_priv system procedure.*

Procedure	Description
sp_sys_priv_role_info	New procedure. Generates a report to map system privileges to corresponding system roles and role IDs. A singe row is returned for each system privilege.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sp_sys_priv_role_info System Procedure.*
sp_use_secure_feature_key	New procedure. Enables an existing secure feature key.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of Catalog Stored Procedures > sp_use_secure_feature_keys system procedure.*

Table 6. System Stored Procedures

Procedure	Description
sp_iqcardinality_analysis	**sp_iqcardinality_analysis** no longer returns an index type value or index recommendation. Users are advised to Run Index Advisor for suggestions about additional column indexes. **sp_iqcardinality_analysis** is deprecated and will be removed in a future release.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqcardinality_analysis.*
sp_iqcolumn	Behavior change. **sp_iqcolumn** no longer includes est_cardinality as part of the output.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqcolumn.*
sp_iqcolumnmetadata	New procedure. Returns index metadata for all columns in one or more tables.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqcolumnmetadata.*
sp_iqconnection	New columns added to display the name of the underlying INC connection and connection suspension status.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqconnection.*

Procedure	Description
sp_iqdbsize	New columns `RLVLogBlocks` and `RLVLogBytes` added.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqdbsize*.
sp_iqdbspace	The `DBSpace Type` column now displays RLV.
	Block Type `R` (RLV Free List Manager) is added to the BlkTypes block type identifiers.
	The `Stripingon` column is now always F.
	The `StripeSize` column is now always 1.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqdbspace*.
sp_iqdbspaceinfo	This stored procedure is not supported for RLV dbspaces. The procedure returns no results if you specify an RLV dbspace.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqdbspaceinfo*.
sp_iqdbspaceobjectinfo	This stored procedure is not supported for RLV dbspaces. The procedure returns no results if you specify an object located in an RLV dbspace.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqdbspaceobjectinfo*.
sp_iqemptyfile	This stored procedure is not supported for RLV dbspaces. The procedure returns an error message f you specify a file located in an RLV dbspace.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqemptyfile*.

Procedure	Description
sp_iqfile	The `SegmentType` column now displays RLV. Block Type `R` (RLV Free List Manager) is added to the BlkTypes block type identifiers. The `RWMode` column is now always RW. The `StripeSize` column is now always 1. See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqfile*.
sp_iqindexmetadata	Behavior change. Returns details about column indexes, column constraints, and NBit dictionary size. New HG index output columns include Cardinality Range Min - Max, Estimated Cardinality, Accuracy of Cardinality, and Force Physical Delete. See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqindexmetadata*.
sp_iqindexrebuildwidedata()	New procedure. Identifies wide columns in migrated databases that you must rebuild before they are available for read/write activities. See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqindexrebuildwidedata()*.
sp_iqlocks	Behavior change. Support for write-intent locks and row-level locks for RLV. See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqlocks*.
sp_iqmergerlvstore	New procedure. Triggers a merge of the row-level versioned (RLV) store with the IQ main store. See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqmergerlvstore*.
sp_iqmpxincstatistics	New procedure. Displays a snapshot of the aggregate statistics of internode communication (INC) status since server startup as of the moment of execution. See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures >sp_iqmpxincstatistics*.

Procedure	Description
sp_iqmpxsuspendedconninfo	New procedure. Shows details about currently suspended connections and transactions on the coordinator node.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqmpxsuspendedconninfo* .
sp_iqrebuildindex	Behavior changes. Rebuilds FP indexes (Flat FP as NBit, or NBit as Flat FP); HG indexes (single HG as tiered HG, or tiered HG as single HG); LOB data and columns greater than 255 bytes wide.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqrebuildindex.*
sp_iqrlvmemory	New procedure. Monitors RLV store memory usage per table.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqrlvmemory.*
sp_iqrowdensity	Behavior change. Returns 1.0 for density in all cases.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqrowdensity.*
sp_iqspaceused	New columns `rlvLogKB` and `rlvLogKBUsed` added.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqspaceused.*
sp_iqstatistics	Behavior change. `OperationsActiveLoadTableStatements` statistic now returns correct value.
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqstatistics.*
sp_iqstatus	New output rows: • RLV memory limit • RLV memory used
	See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqstatus.*

Procedure	Description
sp_iqtable	New PartitionType column indicates table partition type (Range, Hash, Hash-Range, None). New `isRLV` column indicates if the table is RLV-enabled. See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures >*.
sp_iqtablesize	New columns `RlvLogPages` and `RlvLogKBytes` added. See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqtablesize.*
sp_iqtransaction	New `VersioningType`, `Blocking`, and `Blocking Timeout` columns. See *Reference: Building Blocks, Tables, and Procedures > System Procedures > Alphabetical List of System Stored Procedures > sp_iqtransaction.*

See also
- *Backward Compatibility: Changes to Default Behavior* on page 41
- *Backward Compatibility: Migration Considerations* on page 44
- *Connection Changes* on page 46
- *Database Option Changes* on page 47
- *JRE and Java Runtime Environment Variable Changes* on page 54
- *Logical Server Policy Option Changes* on page 55
- *Login Policy Option Changes* on page 56
- *ODBC Driver Changes* on page 57
- *Query Plan Changes* on page 58
- *SQL Function Changes* on page 61
- *SQL Statement Changes* on page 62
- *Tables and Views Changes* on page 78
- *Utility Option Changes* on page 81
- *Deprecated Stored Procedures* on page 38

Tables and Views Changes

SAP Sybase IQ 16.0 includes new and changed system tables and views.

New System Tables

System Table	Contains
ISYSIQPARTITION-COLUMN	One row for each column of a partition described in ISYSPARTITION view in a partitioned table described in the ISYSPARTITIONSCHEME view.
	See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > System Views > SYSPARTITION System View* for details.
ISYSLDAPSERVER	Contains attributes for the LDAP server.
	See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > Alphabetical List of System Views > SYSLDAPSERVER*.
ISYSIQRLVMERGE-HISTORY	New table logging RLV merge events from the RLV store to the IQ main store. A log entry is added for each row-level versioning (RLV) enabled-table each time a merge between the RLV store and the IQ main store begins. Log entries are updated when the merge is complete.
	See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > System Views > SYSIQRLVMERGEHISTORY System View* for details.
ISYSIQRVLOG	New table tracking RLV log streams.
	See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > System Views > SYSIQRVLOG System View* for details.
ISYSCERTIFICATE	New table storing certificates in text PEM-format.
	See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > System Views > SYSCERTIFICATE System View* for details.
ISYSPARTITION	One row for each range subpartition of a hash-range partitioned table. Because Hash partitions are logical rather than physical objects, `ISSY-PARTITION` will not contain any rows for hash partitions of hash partitioned table or hash-range partitioned table.
	See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > System Views > SYSPARTITION System View* for details.

System Table	Contains
ISYSPARTITION-KEY[*]	One row for each column of a partition key. See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > System Views > SYSPARTITIONKEY System View* for details.
ISYSPARTITION-SCHEME	One row for each hash partitioned table and each hash-range partitioned table. The hash function and the number of hash partitions per table is fixed for the database and is not stored as part of this table. See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > System Views > SYSPARTITIONSCHEME System View* for details.
ISYSROLEGRAN-TEXT	New table containing syntax extensions pertaining to the SET USER and CHANGE PASSWORD system privileges. See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > Alphabetical List of System Views > SYSROLEGRAN-TEXT.*
ISYSROLEGRANT	New table containing one row for each grant of a system or user-defined role. See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > Alphabetical List of System Views > SYSROLEGRANT.*
ISYSSUBPARTI-TIONKEY[*]	One row for each column of a subpartition key. See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > System Views > SYSSUBPARTITIONKEY System View* for details.

Note: [*]ISYSPARTITIONKEY and ISYSSUBPARTITIONKEY share identical schemas.

Changed System Tables

System Table	Description of Change
ISYSIQDBSPACE	New column is_rlv_store indicating if the dbspace is a RLV store dbspace. See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > System Views > SYSIQDBSPACE System View* for details.

System Table	Description of Change
ISYSIQIDX	Removed columns: • `link_table_id` • `link_column_id`
ISYSIQTAB	New column `is_rlv` indicating if RLV storage is enabled on this table. Removed column `join_id`. See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > System Views > SYSIQTAB System View* for details.
ISYSIQTABCOL	Removed columns: • `link_table_id` • `link_column_id` • `has_original` • `original_not_null` • `original_unique` Added column `is_nbit`
ISYSTABLEPERM	Column `loadauth` and `truncateauth` added to indicate current grant status of the object privileges. See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > Alphabetical List of System Views > SYSTABLEPERM.*
ISYSUSER	Columns for `user_dn` and `user_dn_cached_at` added for LDAP user authentication. Column `user_type` added to identify user type for RBAC. Column `dual_password` added for the first and second parts of the dual password for a user. See *Reference: Building Blocks, Tables, and Procedures > System Tables and Views > Alphabetical List of System Views > SYSUSER.*

See also

Utility Option Changes

SAP Sybase IQ 16.0 includes option changes for the **isql** and **start_iq** utilities.

Table 7. isql Changes

Option	Description
--URP	New option for the Interactive SQL utility (**isql**). Allows you to specify a remote password and provides TDS login redirection for logical servers. See *Utility Guide > isql Interactive SQL Utility.*.

Table 8. start_iq Changes

Utility	Description
-al	New option for the **start_iq** utility, **-al** allows you to extend authentication to select users using standard authentication. This may be necessary when the *login_mode* database options are restricted to SAP Sybase IQ LDAP user authentication (LDAPUA) only, but no user or login policy exists that permits LDAPUA authentication. See *start_iq > -al iqsrv16 Server Option* or *start_iq > -al iqsrv16 Database Option* under *start_iq Database Server Startup Utility* in the *Utility Guide* .

Utility	Description
-iqlm	New option. The **–iqlm** startup option specifies the maximum amount of memory that SAP Sybase IQ can dynamically request from OS for temporary use. Set **–iqlm** as a switch as part of the command or configuration file that starts the server. Default large memory is 2048MB. Large memory requirements represent one third of the total available physical memory. To ensure adequate memory for the main and temporary IQ stores, set the **–iqlm**, **–iqtc**, and **–iqmc** startup parameters so that each parameter receives one third of all available memory. See *Utility Guide > start_iq Database Server Startup Utility > start_iq Server Options > -iqlm iqsrv16 Server Option*.
–iqmc	Behavior change. Default memory for **–iqmc** is 64MB. Large memory requirements are one third of all available physical memory. To ensure adequate memory for the main store, set the **–iqmc** startup parameter to one third of available physical memory. See *Utility Guide > start_iq Database Server Startup Utility > start_iq Server Options > -iqmc iqsrv16 Server Option*.
-iqrlvmem	New option. Server startup argument to govern the RLV store memory usage. Only used if RLV storage is enabled. The default is 2GB.
–iqtc	Behavior change. Default memory for **–iqtc** is 64MB. Large memory requirements are one third of all available physical memory. To ensure adequate memory for the IQ temporary store cache, set the **–iqtc** startup parameter to one third of available physical memory. See *Utility Guide > start_iq Database Server Startup Utility > start_iq Server Options > -iqtc iqsrv16 Server Option*.
-sf	New feature set. `manage_keys` added to the `manage_security` feature set for the **start_iq -sf** option. Allows you to specify which features you want to secure for databases running on the database server.

See also
- *Backward Compatibility: Changes to Default Behavior* on page 41
- *Backward Compatibility: Migration Considerations* on page 44
- *Connection Changes* on page 46
- *Database Option Changes* on page 47
- *JRE and Java Runtime Environment Variable Changes* on page 54
- *Logical Server Policy Option Changes* on page 55

- *Login Policy Option Changes* on page 56
- *ODBC Driver Changes* on page 57
- *Query Plan Changes* on page 58
- *SQL Function Changes* on page 61
- *SQL Statement Changes* on page 62
- *Stored Procedure Changes* on page 70
- *Tables and Views Changes* on page 78

Behavior Changes

Administration Tool Changes

Sybase Control Center Changes

The SAP Sybase IQ plug-in for Sybase Control Center (SCC) 3.2.7 provides complete administration and management functionality. SCC replaces the Sybase Central administration tool.

For documentation on administering SAP Sybase IQ with SCC, see the *Sybase Control Center for SAP Sybase IQ* online help at *http://sybooks.sybase.com/sybooks/sybooks.xhtml? prodID=10680*.

Brief descriptions of new and enhanced features in SCC 3.2.7:

Table 9. Sybase Control Center Changes

New/Enhanced Feature	Description
External Environments	External environment administration – In previous releases, only the Java external environment was supported. In SCC 3.2.7, you can administer these external environments: ODBC, ESQL, .NET, Java, Perl, PHP.
Tables	Table administration – Perform table management tasks such as creating/deleting a base table and moving a table to another dbspace. • Creating/deleting a base table • Creating/deleting a global temporary table • Creating/deleting a proxy table • Modifying table properties • Moving a table to another dbspace • Managing table columns • Managing table constraints • Managing table indexes • Managing table partitions • Managing IQ catalog store (system store) table triggers • Generating table DDL commands • Viewing table data in the Execute SQL window
SAP Sybase IQ server and SCC agent logs	View and filter logs, copy and paste from log snapshots. The log displays as a static snapshot that does not refresh automatically.

New/Enhanced Feature	Description
Domains	Domain administration – Create, view and update domains. The Administration Console tree lists domains under schema objects.
Views	View administration – Perform view management tasks such as creating a view and recompiling a materialized view. • Creating a view • Modifying a view • Disabling a view • Recompiling and enabling a view • Deleting a view • Generating view DDL commands • Managing triggers on an IQ catalog store (system store) view • Listing materialized views • Creating a materialized view • Modifying a materialized view • Validating materialized view data • Refreshing materialized view data • Truncating materialized view data • Setting a clustered index on a materialized view • Disabling a materialized view • Recompiling a materialized view • Generating materialized view DDL commands • Managing triggers on a materialized view
Web Services	Web service administration – Perform web service management tasks such as creating, modifying, and deleting web services. • Creating a Web service • Enabling a Web service • Disabling a Web service • Modifying a Web service • Deleting a Web service • Generating Web service DDL commands

New/Enhanced Feature	Description

Logical Servers	Logical server support has been enhanced as follows:
	
	• When you authenticate a multiplex server resource, SCC automatically generates one logical server resource per logical server, if one does not already exist. SCC sets the resource name to the multiplex name followed by a colon and the logical server name. For example, `myMPX1:MyLS1`.
	• Create, edit, and delete logical server policies to control logical server access and context.
	• Monitor logical servers from the Perspective Resources View, Heat Chart, or Administration Console.
	• For better load balancing, SAP Sybase IQ can now redirect new multiplex logins to a node with a lighter load in the same logical server.
	• Two generated administration scripts include new connection parameters in the connection string:
	  • stop_server.sh (.bat) includes LS=SERVER
	  • sync_server.sh (.bat) includes `LS=COORDINATOR`

Login Mappings	Login mappings – Map a Windows user profile or a Kerberos principal to an existing user in the database to maintain a single user ID for database connections, operating system, and network logins.
Transport layer security for multiplex	Transport layer security – Configure transport layer security with RSA encryption in a multiplex environment.
Global Transaction Resiliency	Global transaction resiliency – Monitor the suspended, resumed, and rolled back status of INC (internode communication) connections for global transactions.
Row-level Versioning	Row-level Versioning (RLV) – Use the new RLV data store in your simplex database to perform row-level updates, inserts, and deletes, in real-time. When a table is enabled for storage in the RLV data store, multiple users can write to different rows of the table concurrently.
Remote Access	Remote Access – Access data from other data sources by creating Remote Server definitions, Proxy Tables, Remote Procedures and External Logins.
LDAP Server	LDAP server administration - Create, modify, and delete LDAP servers.
Role-Based Security	Role-based security administration – Perform user and role management for role-based users, user-extended roles, and standalone roles. Tasks include create, modify and delete, as well as grant and revoke system privileges and object permissions to users and roles.

New/Enhanced Feature	Description
Triggers on Views	View trigger administration – Create, modify, and delete triggers on views.
Sequence Generator	Sequence Generator administration - Create, modify, and delete sequence generators.

See also
- *Discontinued Tools and Utilities* on page 35
- *Sybase Control Center* on page 7

Documentation Collection Changes

The SAP Sybase IQ 16.0 documentation collection contains new documentation, redesigned manuals, and retitled manuals. This topic helps you locate content in the 16.0 collection.

Table 10. Documentation Collection Changes in SAP Sybase IQ 16.0

Location in 15.4 Collection	New Location in 16.0 Collection

• Sybase Central online help	• Sybase Control Center online help, available from:
	   • The SCC GUI
	   • The SAP Sybase IQ documentation collection
	   • The SCC documentation collection at *http://sybooks.sybase.com/sybooks/sybooks.xhtml?prodID=10680.*

• *System Administration Guide: Volume 1* • *System Administration Guide: Volume 2*	• *Administration: Database* • *Administration: Backup, Restore, and Data Recovery* • *Administration: Globalization* • *Administration: User Management and Security* • *Administration: Load Management* • *Programming*
• *Advanced Security*	• *Administration: User Management and Security > Advanced Security Options in Sybase IQ*
• *Using Sybase IQ Multiplex*	• *Administration: Multiplex*

Location in 15.4 Collection	New Location in 16.0 Collection
• *Installation and Configuration Guide for HP-UX > Database Upgrades* • *Installation and Configuration Guide for AIX > Database Upgrades* • *Installation and Configuration Guide for Linux > Database Upgrades* • *Installation and Configuration Guide for Solaris > Database Upgrades* • *Installation and Configuration Guide for Windows > Database Upgrades*	• *Migration (Linux and UNIX)* • *Migration (Windows)*

SAP Sybase IQ Documentation Collection

Refer to this summary to locate information about specific subject areas.

New Release Information

Document Name	Contents
Installation and Configuration	Installation and configuration procedures.
Release Bulletins	Late-breaking product information.
Guide to Licensed Options	Features you can buy separately.
New Features Summary	Release-specific feature summaries.

Getting Started

Document Name	Contents
Introduction to SAP Sybase IQ	Hands-on introduction to SAP Sybase IQ.
Quick Start (UNIX/Linux), Quick Start (Windows)	Steps to create and query an IQ demo database.
Migration (UNIX/Linux), Migration (Windows)	Steps to install maintenance releases, steps for upgrading your database, and information on upgrading to the role-based security model.

Administration

Document Name	Contents
Administration: Database	Database setup and troubleshooting.

Document Name	Contents
Administration: Backup, Restore, and Data Recovery	Data preservation strategies.
Administration: Globalization	Locales, collations and character set configuration.
Administration: Load Management	Data import and export procedures.
Administration: User Management and Security	User and Security implementation and administration.
Administration: In-Memory Row-Level Versioning	Row-level versioning configuration and administration.
Administration: Spatial Data	IQ catalog store spatial data administration.

Reference

Document Name	Contents
Reference: Statements and Options	Syntax and parameters for SQL statements and options.
Reference: Building Blocks, Tables, and Procedures	SQL, functions, procedures, tables, and views.
Error Messages	Error and warning messages.
Utility Guide	Command-line utility reference.
Performance and Tuning Guide	Database, system, and query tuning options.
Programming	Developer reference for building and deploying database applications.

Optional Features

Document Name	Contents
Unstructured Data Analytics	Binary Large Object (BLOB) and Character Large Object (CLOB) storage and retrieval.
User-Defined Functions Guide	C/C++ interface for user-defined functions.
Administration: Multiplex	Multiplex setup and administration.

Document Name	Contents
Administration: User Management and Security	Contains Advanced Security option information: • LDAP User Authentication configuration • Implement Kerberos authentication

Licensing

Document Name	Contents
Sybase Software Asset Management (SySAM) 2	License generation, options, and management.
SySAM 2 Quick Start Guide	SPDC and SAP Service Marketplace license generation.
FLEXnet Licensing End User Guide	FLEXnet Licensing utilities.