User Guide | PUBLIC
SAP IQ 16.2
Document Version: 1.0.0 – 2026-08-11

Introduction to SAP IQ

Content

1 Introduction to SAP IQ

This guide introduces the SAP IQ product, its concepts, and its components.

Glossary

- Glossary [page 32]

SAP IQ terminology.

2 About SAP IQ

SAP IQ is a high-performance decision-support server designed specifically for mission-critical business intelligence, analytics, and data warehousing. Component Integration Services within SAP IQ provide direct access to relational and non-relational databasess.

SAP IQ is a fully ACID-compliant database with guaranteed data integrity regardless of whether your database (or dbspace) is created on raw devices or file system files. Multiplex implementations support GPFS and NFS; simplex implementations support any file system. Both simplex and multiplex support raw devices.

2.1 Product Licensing

Available licenses are determined by product edition. Optional features are sold and licensed separately.

Whether you install a licensed or evaluation edition, you have 90-day access to all features and options. To use an option beyond the 90-day evaluation period, you must purchase and install an appropriate SAP Software Asset Management (SySAM) license. See SySAM Users Guide.

The e-mail message or Web key you receive when you purchase an SAP product provides specific product and licensing information. Before installing and configuring any product edition, know which features you are licensed to use.

2.2 Product Editions

- **SAP IQ Enterprise Edition**
  Supports all options and features. The Enterprise Edition requires a CPU license, which restricts the number of cores (not CPUs) the server can use. Optional features are sold and licensed separately.
- **SAP IQ 90-Day Trial Evaluation**
  Provides full access to all features available in the SAP IQ Enterprise Edition for a 30-day evaluation period. When the evaluation period expires, customers are required to purchase an appropriate license to continue running the server. Available as a free download here.

2.3 Available Licenses

The SAP IQ Evaluation Edition and SAP IQ Enterprise Edition provide 90-day access to all features and options. Access to optional features beyond the 90-day evaluation period requires an SAP IQ Enterprise Edition and appropriate license.

Orderable License	Actual License	Description
CPU	IQ_CORE	Restricts the number of cores (not CPUs) the server can use.
VLDB MO	IQ_VLDBMGMT	Lets you create multiple tablespaces and dbspaces beyond the defaults. Each VLDB license allows for 1 TB of storage. Additional licenses are required for each additional TB of storage in the main store.
ASO	IQ_SECURITY	Enables column and ECC Kerberos encryption options. The number of ASO licenses must match the number of cores.
MPXO	IQ_MPXNODE	Allows you to starts secondary multiplex nodes (readers/writers). The number of nodes must always be less than cores, as each multiplex server must have at least one core.
IDBA-PSO	IQ_UDF	Lets you define and run high-performance scalar and aggregate user-defined functions. Available with approved third-party libraries only.
	IQ_IDA	Allows you to build, deploy and run C/C++ V4 User-Defined Functions (UDF). The IQ_IDA license will function as both the InDatabase Analytics Option and Partner Solutions license.
UDA	IQ_UDA	Applies to TEXT indexes and enables the IQ text search functionality, which lets search unstructured and semi-structured data. An IQ_UDA also includes an IQ_LOB license.

> **Note**
>
> The DBA must explicitly authorize access to an optionally licensed feature. SAP IQ will not attempt to check-out an optional license by default. Unless the DBA grants access to an option, the option will not be available. See the `sp_iqlmconfig` stored procedure in *SAP IQ SQL Reference* for more information.

Related Information

sp_iqlmconfig Procedure

2.4 SAP IQ as Near-Line Storage

SAP Business Information Warehouse users can use SAP IQ as a near-line storage (NLS) solution.

With NLS implemented, frequently-used hot and warm data is stored in SAP HANA, while infrequently-used cold data is stored in SAP IQ, thereby reducing your in-memory footprint. For detailed information, see SAP-NLS Solution for SAP BW.

2.5 Related Products

There are a number of products that are related to SAP IQ.

Product	Description
SAP HANA	SAP HANA is an in-memory database that combines transactional data processing, analytical data processing, and application logic processing functionality in memory. See SAP HANA on the SAP Help Portal for more information.
SAP SQL Anywhere	SAP SQL Anywhere is a comprehensive data management package for server, desktop, mobile, and remote office environments. SAP IQ is an extension of SAP SQL Anywhere, and incorporates many of its features. See SAP SQL Anywhere on the SAP Help Portal for more information.
SAP PowerDesigner	PowerDesigner is the industry-leading modeling and metadata management solution for data, information, and enterprise architecture. See SAP PowerDesigner on the SAP Help Portal for more information.

Related Information

TEMP_EXTRACT_SIZEn Options

3 Documentation

SAP IQ documentation is topic-oriented and task-based. Divisions within the set are designed to address key subject areas. Refer to this summary to locate information about specific items of interest.

All documentation is available online in HTML and pdf format on the SAP IQ Help Portal.

3.1 New Release Information

See the *Release Notes: What's Changed in SAP IQ 16.2*:

- SAP IQ 16.2 Initial Release (SP00)

3.2 Getting Started

- Introduction to SAP IQ [page 4]
- SAP IQ Guide to Licensed Options
- SAP IQ Installation and Update Guide for Linux
- SAP IQ Installation and Update Guide for Windows
- Getting Started With SAP IQ

3.3 Administration

- SAP IQ Administration: Database
- SAP IQ Administration: Multiplex
- SAP IQ Administration: User Management and Security
- SAP IQ Administration: Backup, Restore, and Data Recovery
- SAP IQ Administration: Globalization
- SAP IQ Administration: Load Management
- SAP IQ Administration: In-Memory Row-Level Versioning (Deprecated)
- SAP IQ Administration: Spatial Data

Configuration

- SAP IQ 16.2 Best Practices Guide
- SAP IQ Performance and Tuning Series: Basics
- SAP IQ Performance and Tuning Series: Monitoring Views

3.4 Reference Guides

- SAP IQ SQL Reference
- SAP IQ Utility Reference
- SAP IQ Performance and Tuning Series: Basics
- SAP IQ Programming Reference
- SAP IQ Interactive SQL Guide
- SAP IQ Connection and Communication Parameters Reference

3.5 Optional Features

- SAP IQ Administration: Unstructured Data Analytics

3.6 SySAM Licensing

- SySAM Users Guide

4 Concepts

Understanding some basic terms and concepts will help you work with SAP IQ..

In this section:

4.1 Key Principles

The key principles of SAP IQ are the concepts of column-wise data structures and bit-wise indexing.

In this section:

4.1.1 Column-Wise Data Structures

Store data column-wise rather than row-wise.

Column-wise data structures are easily searched without having to scan the full rows. Vertical partitioning of the data means never performing a table scan. Only those columns necessary for a query are returned, significantly reducing I/O.

Row-wise data structures, such as those used in traditional RDBMSs, can be searched only by reading an entire page and then locating the data on the page. Since the columns being searched for are only a fraction of the total data read, there is excessive I/O.

Compare the number of I/O's generated between column-wise and row-wise data structures using the same SELECT Statement.

```
SELECT count(*) FROM customer WHERE gender = "M"
```

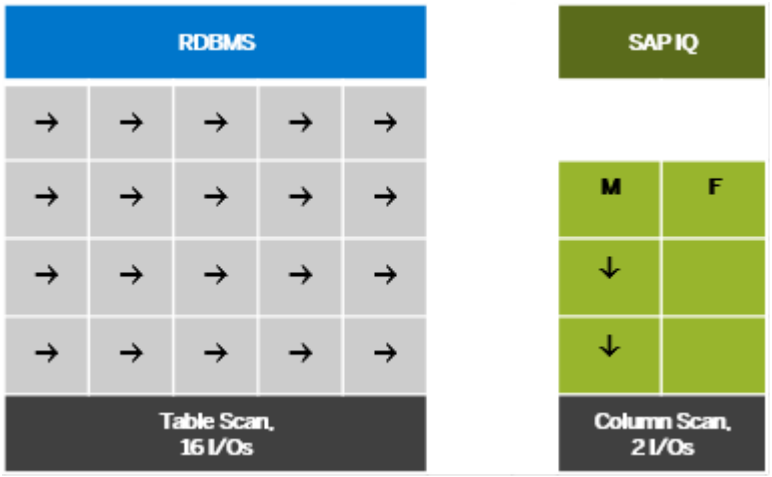

4.1.2 Bit-Wise Indexing

Bitmap indexes use bit arrays (commonly called bitmaps) and answer queries by performing bitwise logical operations on these bitmaps.

Unlike other indexing methods, SAP IQ indexes are based on the cardinality of the data and the way the data will be used. SAP IQ has indexes for both low- and high-cardinality columns.

Bit-wise indexing benefits both low- and high-cardinality columns. It translates the distinct data elements of a column into a bit mask.

Cardinality is the number of unique values in a column:

- Low-cardinality data: fewer than 1500 unique values. Examples are days of the week, gender and states.
- High-cardinality data: 1500 unique values or more. Examples are primary keys, telephone numbers, social security numbers.

Example of Low-Cardinality Bit-Wise Indexing

The table on the left represents unindexed data. The table on the right represents the same data as a bit-wise index.

> **Note**
>
> This is not an actual index, but merely a representation of how bit-wise indexing works for low-cardinality data.

Color	Green	Blue	Red
Red	0	0	1
Blue	0	1	0
Green	1	0	0

Example of High-Cardinality Bit-Wise Indexing

The data is stored vertically.

Actual Data	In Binary				SAP IQ Index			
6	0	1	1	0	0	1	1	0
9	1	0	0	1	1	0	0	1
5	0	1	0	1	0	1	0	1
11	1	0	1	1	1	0	1	1
0	0	0	0	0	0	0	0	0
1	0	0	0	1	0	0	0	1
7	0	1	1	1	0	1	1	1

4.2 Architecture

SAP IQ supports both simplex and multiplex architecture.

In this section:

Simplex Server Architecture [page 13]
Simplex is a single instance of an SAP IQ server running on a single node (machine).

Multiplex Server Architecture [page 13]
Multiplex is multiple instances of SAP IQ servers running on multiple machines. SAP IQ supports read-write transactions from multiple servers in the multiplex.

4.2.1 Simplex Server Architecture

Simplex is a single instance of an SAP IQ server running on a single node (machine).

In a simplex configuration, the files may be located on a host machine or on a network storage device (or distributed between the host and the storage device).

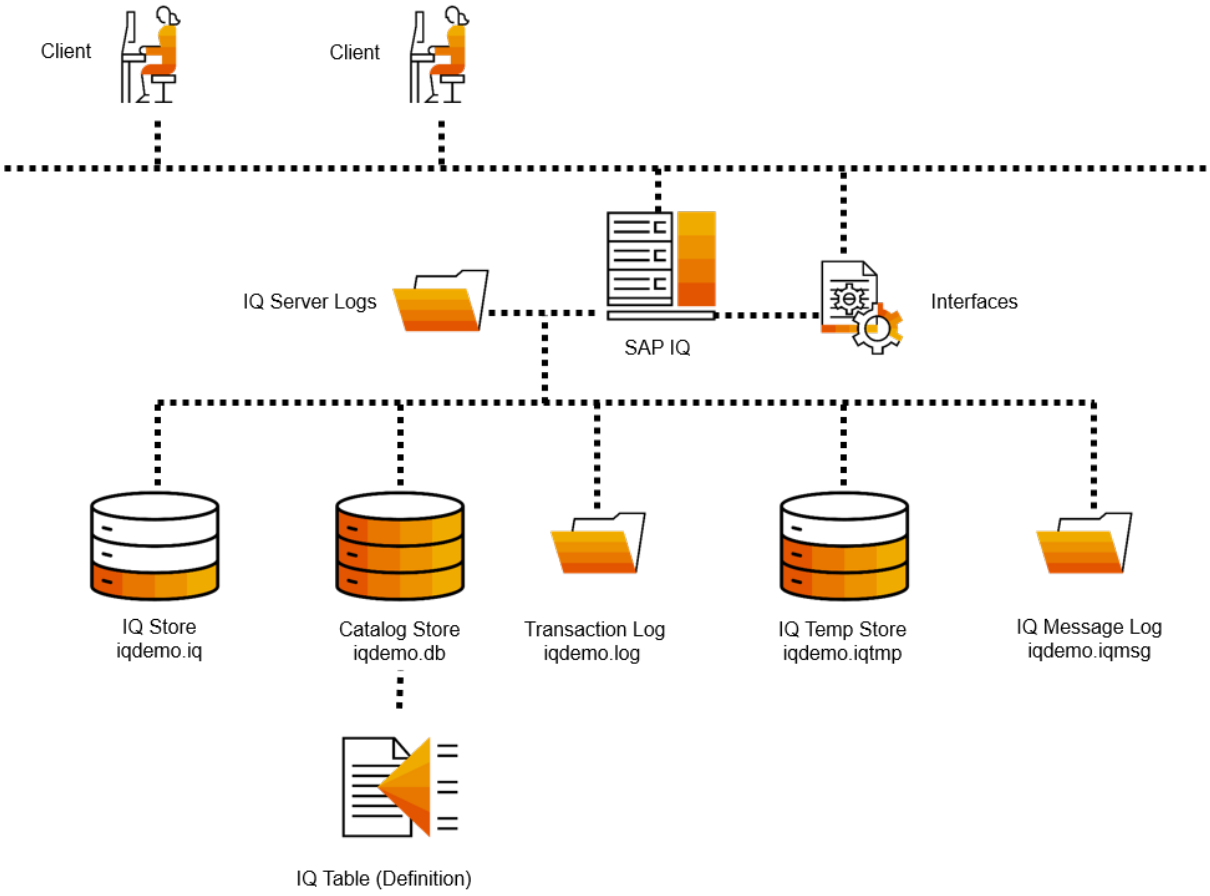

4.2.2 Multiplex Server Architecture

Multiplex is multiple instances of SAP IQ servers running on multiple machines. SAP IQ supports read-write transactions from multiple servers in the multiplex.

Each server has its own temporary store and catalog store. The shared IQ Store is common to all servers. The primary server, or coordinator, manages all global read-write transactions and maintains the global catalog. The table version log (TLV log) stores information about DML and DDL operations and communicates information about new table versions to the secondary servers.

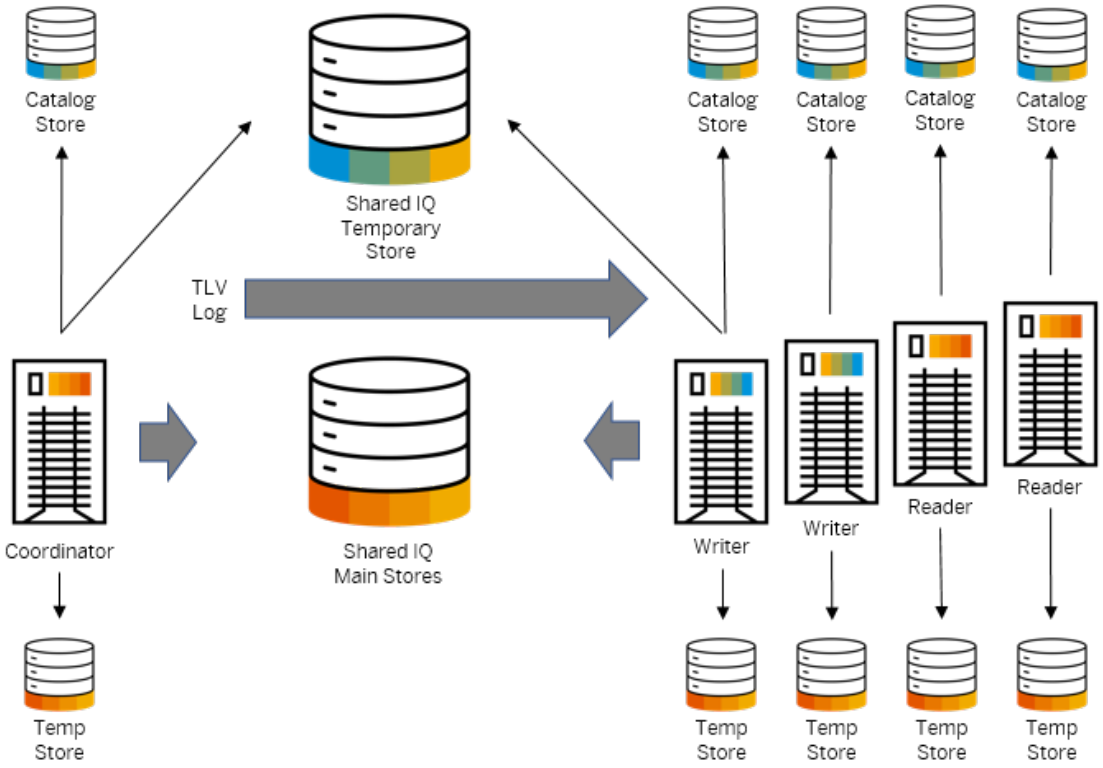

4.3 SAP IQ Database

An SAP IQ database is made up of stores, log files, and server space.

In this section:

Stores [page 14]
There are five stores in the SAP IQ database.

IQ Log Files [page 17]
There are four log files in the SAP IQ database.

Server Spaces [page 20]
There are two server spaces in the SAP IQ database.

4.3.1 Stores

There are five stores in the SAP IQ database.

In this section:

4.3.1.1 Catalog Store (for Metadata)

The catalog store is also known as the database file since it has a `.db` file extension. It may also be referred to as the catalog.

There is one catalog store per SAP IQ database. It is always created on a file system and consists of one object for the SAP IQ database. It contains all information needed to manage the database, including metadata stored in the system tables and stored procedures (system and user).

It is also known as the database file and its physical file name is: `dbname.db`.

The catalog store consists of the SYSTEM and TEMPORARY dbspaces. It is not recommended that you overload the catalog store with additional dbspaces.

All the information about tables in a database is held in the system tables within the catalog store. Browse data from the system tables using Interactive SQL.

The catalog store resembles an SAP SQL Anywhere database as it uses traditional row-based storage, typically on 4K (4096 byte) pages. The page size for the catalog is configured when you create the SAP IQ database.

For a complete listing of system tables and views, see SAP IQ SQL Reference.

4.3.1.2 IQ Main Store (for Data)

The IQ main store holds all the user data (indexes) for the database compressed on disk.

It contains the transaction log for data in the IQ main store tables and structure to manage space allocation (free list). The internal transaction log manages transactions involving IQ main store data. The free list tracks page usage with a given dbspace.

The IQ main store is a preallocated space (although it can be enlarged with advance planning). Initially, the IQ main store has one file or raw device when the database is created, but additional files may be added to the IQ main store as needed.

> **Note**
>
> Raw devices aren't available in Linux. See SAP Note 3483103 The RAW_DRIVER, and support for devices located in `/dev/raw/`, was deprecated in:
>
> - SuSE Linux Enterprise Server (SLES) 15.4
> - Red Hat Enterprise Linux 9

There's one IQ main store per database, which usually consists of many files or raw partitions. When using raw partitions, the physical name is the partition name. If setting up multiplex on multiple hosts, the IQ main store must be created using raw partitions.

The IQ main store consists of least two dbspaces:

IQ_SYSTEM_MAIN the first dbspace created when creating a database, it contains system metadata. It shouldn't be used for user data.

USER_DBSPACE created and named by DBA, it holds user tables and indexes. DBAs may create additional dbspaces, as needed. The database option Default_Dbspace changes the default dbspace where tables are created.

> **Note**
>
> You must be licensed for the IQ_VLDBMGMT option in order to create more than one user dbspace per database.

4.3.1.3 IQ Temporary Store (for Temporary Data)

The IQ temporary store is used for both loading data and for queries.

It is the work area for the database. It is used to build or modify HG and optimized FP indexes during loads and deletes on columns with those indexes. It is also used to sort data during order by queries, some group by operations and some table join processing.

The physical file name of the IQ temporary store is `dbname.iqtmp` (for file system devices). There is one IQ temporary store per database.

The IQ temporary store is initialized at the time the IQ database is created and has the logical name IQ_SYSTEM_TEMP. It contains the IQ_SYSTEM_TEMP dbspace, which can consist of several physical files. These files can be raw or file system. If a raw partition, the physical name will be the partition name.

4.3.1.4 RLV Store (for In-Memory Data)

The row-level versioning (RLV) store is an in-memory store for high-performance row-level updates. If a table is registered for RLV storage, then all `LOAD TABLE`, `INSERT`, `UPDATE`, and `DELETE` statements write directly to the RLV store.

Multiple connections can make simultaneous updates to different rows of an RLV-enabled table. In the IQ main store, only one connection can write to a table at one time.

The RLV store periodically, and automatically, merges its in-memory contents with the IQ main store, although you can change merge preferences. You can trigger a manual merge when desired.

For detailed information on using the RLV store to perform real-time concurrent updates, see *SAP IQ Administration: In-Memory Row-Level Versioning*.

Related Information

SAP IQ SQL Reference

4.3.1.5 IQ Shared Temporary Store (for Temporary Data in a Multiplex)

The IQ shared temporary store contains temporary user objects or local nonversioning temporary objects.

When you create a database, SAP IQ creates a single dbspace for the IQ shared temporary store, IQ_SHARED_TEMP. Initially, this dbspace contains no files.

A multiplex configuration with shared temporary storage can use the IQ_SHARED_TEMP dbspace as a shared system temporary store to simplify multiplex configuration, and improve performance.

4.3.2 IQ Log Files

There are four log files in the SAP IQ database.

The names of the IQ log files are derived from the name of the database with unique extensions. For example, an IQ database named "customer" would contain a catalog store named `customer.db`, an IQ message file named `customer.iqmsg`, and a catalog store transaction log file named `customer.log`. If using file systems for the IQ store and IQ temporary store, their names would be `customer.iq` and `customer.iqtmp`, respectively. For raw partitions, the dbspaces would be named by the partition name.

In this section:

RLV store logging is different, and in many ways simpler, than traditional database (disk or in-memory) logging. because data in the RLV store is transient. Data pages are not persisted to disk, and data is merged to the IQ main store on a frequent basis.

4.3.2.1 Message Log

The IQ message file is created when an SAP IQ server starts for the first time.

If the file has been deleted or renamed, the server automatically creates a new IQ message file when the server is restarted. The IQ message file will continue to grow over time. However, the DBA can set a maximum file size after which the file wraps around and overwrites the oldest records.

The IQ message file captures:

- SAP IQ server start up information
- Messages concerning the time and records for table loads, deletes and truncations
- Transaction IDs for all connections to the database
- Times of server checkpoints
- Error messages which are typically identified by "`exception thrown`"
- Status messages
- Query plans

The name of the physical file is `dbname.iqmsg`. There is one message log per database, for file system only.

4.3.2.2 Catalog Store Transaction Log

The transaction log file contains transaction for the catalog store only. It is used to roll back (undo) or roll forward changes as necessary.

The catalog store transaction log file does not contain any transactions for the warehouse data in the IQ store. The name of the physical file is `dbname.log` and it is stored in the same directory as the `.db` file. The catalog store transaction log is a small file whose growth over time is related to its activity. This file must exist as a component of the SAP IQ database.

Do not delete or alter this file. If this file is lost or damaged, technical support may be necessary to recover the database. To protect `.log` file, it may be mirrored to another directory. All transactions involving data in the catalog store are logged to this transaction log, including:

- Any create, alter or drop statements associated with SAP IQ database objects
- Adding or removing users from the database
- Adding or changing user permissions
- Adding or changing a dbspace to the database
- Any DML statement for user data that is stored in the catalog

4.3.2.3 Server Logs

The SAP IQ server engine logs messages to several log files, which are created each time the server is started.

On Windows servers, the server log files are located in the `C:\Documents and Settings\All Users\Sybase IQ\logfiles` directory.

- `iq_startup_nt.log`—records server startup parameters only
- `<server_name>.00n.srvlog`—where n is the number of times the server has been started
- Other server messages are sent to the server console window

On UNIX and Linux servers, the server log files are located in $IQDIR16`/logfiles`. These server log files have 4 digit numbers.

- `<servername>.000n.stderr`
- `<servername>.000n.srvlog`

4.3.2.4 RLV Log

RLV store logging is different, and in many ways simpler, than traditional database (disk or in-memory) logging. because data in the RLV store is transient. Data pages are not persisted to disk, and data is merged to the IQ main store on a frequent basis.

In this section:

4.3.2.4.1 Log Space Usage

RLV log space usage is reported by the `sp_iqtablesize` stored procedure.

In particular, two columns of `sp_iqtablesize` relate to the RLV log.

Column Name	Description
RlvLogPages	Number of IQ pages being used to store RLV logs for this table
RlvLogKBytes	Number of kilobytes being used to store RLV logs for this table

Related Information

sp_iqtablesize Procedure

4.3.2.4.2 Commit Log

The commit log is a specialized log stream, which speeds the analysis pass during recovery.

The commit log contains only end transaction and merge log records. During recovery a list of committed transactions can be quickly built by reading the relatively small commit log followed by the tail of the individual log streams.

4.3.3 Server Spaces

There are two server spaces in the SAP IQ database.

In this section:

Dbspaces [page 20]
A dbspace is a logical name for a container of files or raw partitions called dbfiles.

Dbfiles [page 21]
A dbfile is an operating system file contained within a dbspace.

4.3.3.1 Dbspaces

A dbspace is a logical name for a container of files or raw partitions called dbfiles.

Catalog and IQ main stores may consist of 1 or more dbspaces. The IQ temporary store consists of only one dbspace.

> **Note**
>
> You must be licensed for the IQ_VLDBMGMT option in order to create more than one user dbspace per database.

4.3.3.2 Dbfiles

A dbfile is an operating system file contained within a dbspace.

The RLV_STORE dbspace, the catalog store SYSTEM dbspace and other catalog dbspaces can contain only one dbfile each. The IQ store user dbspaces, the IQ_SYSTEM_MAIN dbspace, and the IQ_SYSTEM_TEMP dbspace can each contain multiple dbfiles.

5 Scalability

SAP IQ's patented design permits databases to scale to contain many terabytes of data.

Its index-based structure allows the database to store your data in a much smaller space than the size of the raw input data, and access it far faster than a traditional relational database. These features make SAP IQ ideal for storing and accessing very large databases (VLDBs).

Database administrators need to understand the options and features that affect performance, and follow documented guidelines. While many default settings automatically provide the greatest efficiency, you may need to experiment with certain option settings for the fastest results, based on your configuration, your loading requirements, and your queries. Setting these options appropriately is necessary for top performance in any SAP IQ database, but is especially important as your database grows to the multiterabyte scale.

This section introduces SAP IQ features that help you manage a very large database, and points you to more detailed discussion and recommendations.

In this section:

5.1 Memory Use

Allocating memory appropriately is a key factor in performance for all SAP IQ databases.

SAP IQ uses memory in its buffer caches for loads and queries. It also uses some memory for managing connections, transactions, buffers, and database objects.

SAP IQ has two buffer caches, one for the main store and one for the temporary store. The default sizes of these caches are not sufficient for a production data warehouse. You must adjust them to reflect the size of your database and tables, your mix of loads and queries, and other factors such as your operating system and other applications that can affect the amount of memory available.

5.2 Data Loads

As your database grows, it is crucial to manage data loading properly.

These features ensure that your loads can scale to meet your needs:

- Buffer manager partitioning to avoid lock contention. Buffer partitioning based on the number of CPUs is enabled by default, and can be adjusted by setting server or database options.
- Allowing sufficient memory for loads, without allocating more memory than is available on your system.
- Reserving space for data structures used during release savepoint, commit, and checkpoint operations. Reserve IQ main store space using the MAIN_RESERVED_DBSPACE_MB Option.

5.3 Page Size

When you create your SAP IQ databases, it is especially important to choose the correct IQ page size.

For very large databases, you need an IQ page size of 128KB or larger.

5.4 Processing Threads

SAP IQ uses operating system threads to process queries and loads.

The default settings of options that control thread use are usually sufficient to provide good performance. In some cases, you may need to change these settings.

In this section:

5.4.1 The Process Threading Model

SAP IQ uses operating system kernel threads for best performance. By default, SAP IQ allocates the number of threads based on the number of CPUs on the system.

Lightweight processes are underlying threads of control that are supported by the kernel. The operating system decides which lightweight processes (LWPs) should run on which processor and when. It has no knowledge about what the user threads are, but does know if they are waiting or able to run.

The operating system kernel schedules LWPs onto CPU resources. It uses their scheduling classes and priorities. Each LWP is independently dispatched by the kernel, performs independent system calls, incurs independent page faults, and runs in parallel on a multiprocessor system.

A single, highly threaded process serves all SAP IQ users. The database server assigns varying numbers of kernel threads to each user connection, based on the type of processing being done by that connection, the total number of threads available, and the various option settings.

Insufficient Threads Error

If there are insufficient threads for a query, SAP IQ generates this error:

```
Not enough server threads available for this query
```

This condition may well be temporary. When some other query finishes, threads are made available and the query may succeed the next time. If the condition persists, you may need to restart the server and specify more SAP IQ threads. It is also possible that `-iqmt` is set too low for the number of connections.

SAP IQ Options for Managing Thread Usage

- Use the server start-up option `-iqmt` to set the maximum number of threads. The default value is calculated from the number of connections and the number of CPUs and is usually adequate.
- Use the server start-up option `-iqtss` to set the stack size of the internal execution threads. The default value is generally sufficient, but may be increased if complex queries return an error indicating that the depth of the stack exceeded this limit.
- Use the `SET OPTION MAX_IQ_THREADS_PER_CONNECTION` statement to set the maximum number of threads for a single user. The `SET OPTION MAX_IQ_THREADS_PER_TEAM` statement sets the number of threads available to a team of threads, enabling you to constrain the number of threads (and thereby the amount of system resources) allocated to a single operation.
- Use these options to control the amount of resources an operation consumes. For example, you can set this option before issuing an `INSERT`, `LOAD`, `BACKUP DATABASE`, or `RESTORE DATABASE` statement.

Related Information

MAX_IQ_THREADS_PER_CONNECTION Option

MAX_IQ_THREADS_PER_TEAM Option

5.5 Disk Space

Learn about managing disk I/O for an SAP IQ system.

The most important factors in managing disk I/O for an SAP IQ system are:

- Having enough disk space for queries and loads
- Using that disk space effectively, so that the fastest I/O is available to support the processing speed of high-powered, multi-CPU systems

The `sp_iqstatus` stored procedure indicates the percentage of space used in the IQ main and temporary stores. If there is not enough temporary or main dbspace available for a buffer or dbspace allocation request, then the statement making the request rolls back. You can create a timer-based event to monitor space usage to help avoid unexpected rollbacks, which may occur in out of space situations on non-privileged operations.

Disk striping is an important means of obtaining maximum I/O performance. Disk striping distributes data randomly across multiple disk drives. You can take advantage of disk striping capabilities in your operating system or disk management software and hardware, as well as internal striping. Disk striping is enabled by default.

Related Information

sp_iqstatus Procedure

5.6 Intermediate Versioning

A key aspect of managing loads and queries in larger databases is SAP IQ's transaction-level versioning.

In particular, SAP IQ offers the ability to roll back transactions to intermediate save points, so that you may not need to repeat the entire load if a long transaction is unable to complete.

5.7 Column-Based Indexes

SAP IQ's column-based indexing structure optimizes your ability to perform selections or calculations on attributes of interest to you.

For the best performance, you need the right set of indexes for your data and queries. Your database should have an index on every column that affects performance.

5.8 Query Optimizer

The SAP IQ query optimizer evaluates every query, choosing among various processing options to produce a query plan that offers optimal performance.

The optimizer is tuned for each release of SAP IQ to choose the best plan for most queries and most databases, including the largest ones.

5.9 Schema Design

SAP IQ often works better with denormalized schemas common in data warehouse design.

In a traditional relational database, normalization improves transaction processing by removing redundancy and improving consistency. In a data warehouse, especially a very large one, denormalization improves performance when processing queries against large amounts of data.

5.10 UNION ALL Views

Tables with a large number of rows can have lengthy load times. The UNION ALL view is one way to address this issue.

SAP IQ lets you partition tables by splitting the data into several separate base tables (for example, by date). You then join them back together into a logical whole by means of a UNION ALL view.

UNION ALL views are simple to administer. If the data is partitioned by, for example, month, you can drop an entire month's worth of data by deleting a table and updating the UNION ALL view definition appropriately. You can have many view definitions for a year, a quarter, and so on, without adding extra date range predicates.

In this section:

5.10.1 UNION ALL Views for Faster Loads

`UNION ALL` views can improve load performance when it is too expensive to maintain secondary indexes for all rows in a table.

SAP IQ lets you split the data into several separate base tables (for example, by date). You load data into these smaller tables. You then join the tables back together into a logical whole by means of a `UNION ALL` view, which you can then query.

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

To create a `UNION ALL` view, choose a logical means of dividing a base table into separate physical tables. The most common division is by month. For example, to create a view including all months for the first quarter, enter:

```
CREATE VIEW
SELECT * JANUARY
UNION ALL
SELECT * FEBRUARY
UNION ALL
SELECT * MARCH
UNION ALL
```

Each month, you can load data into a single base table—JANUARY, FEBRUARY, or MARCH in this example. Next month, load data into a new table with the same columns, and the same index types.

> **Note**
>
> You cannot perform an `INSERT...SELECT` into a `UNION ALL` view.

In this section:

Related Information

UNION Statement

5.10.1.1 Queries Referencing UNION ALL Views

SAP IQ includes optimizations for `UNION ALL` views.

Optimizations include:

- Split `GROUP BY` over `UNION ALL` view
- Push-down join into `UNION ALL` view

A `UNION` can be treated as a partitioned table only if it satisfies all of the following constraints:

- It contains only one or more `UNION ALL`.
- Each arm of the `UNION` has only one table in its `FROM` clause, and that table is a physical base table.
- No arm of the `UNION` has a `DISTINCT`, a `RANK`, an aggregate function, or a `GROUP BY` clause.
- Each item in the `SELECT` clause within each arm of the `UNION` is a column.
- The sequence of data types for the columns in the `SELECT` list of the first `UNION` arm is identical to the sequence in each subsequent arm of the `UNION`.

Related Information

JOIN_PREFERENCE Option

5.10.1.2 UNION ALL View Performance

Structure queries to evaluate the `DISTINCT` operator before the `ORDER BY`, where the sort order is `ASC`.

Certain optimizations, such as pushing a `DISTINCT` operator into a `UNION ALL` view, are not applied when the `ORDER BY` is `DESC` because the optimization that evaluates `DISTINCT` below a `UNION` does not apply to `DESC` order. For example, the following query would impact performance:

```
SELECT DISTINCT state FROM testVU ORDER BY state DESC;
```

To work around this performance issue, queries should have the `DISTINCT` operator evaluated before the `ORDER BY`, where the sort order is `ASC` and the optimization can be applied:

```
SELECT c.state FROM (SELECT DISTINCT state
  FROM testVUA) c
ORDER BY c.state DESC;
```

6 Maximum Limits

Limits exist on the size of objects and the number of objects in SAP IQ databases. In most cases, computer memory and disk drive are more limiting factors.

Item	Limit
Catalog file size	Maximum is 1TB for all platforms. Windows systems with NTFS support the 1TB maximum.
Database name size	250 bytes
Database size	Maximum database size approximates the number of files times the file size on a particular platform, depending on the maximum disk configuration. Refer to your operating system documentation for kernel parameters that affect the maximum number of files.
Dbfile size	File system device: Maximum size is 100 TB Operating system file: Maximum size supported by the operating system
Dbspace size	Maximum 2000 dbfiles per dbspace.
Field size	255 bytes for `BINARY`, 32,767 bytes for `VARBINARY` 32,767 for `CHAR`, `VARCHAR` Up to 512 TB for 128 KB pages or 1 PB for 512 KB pages for `LONG BINARY`, `LONG VARCHAR`
IQ page size	Must be between 64KB and 512KB
Maximum key size	255 bytes for single-column index. 5300 bytes for multicolumn index
Maximum length of string literal	32KB

Item	Limit
Maximum length of SQL statement	The maximum length of SQL statements is limited to the amount of memory available for the IQ catalog, and to the size of the catalog stack If your SQL statements are long, increase the catalog stack size using `-gss`, and increase catalog memory cache size using `-c` or a combination of `-ch` and `-cl` When printing the SQL statement in error messages, the text is limited to the IQ catalog page size. To print long commands, you can start the server with an increased `-gp` setting, although in general SAP recommends that you use the default of `-gp` 4096.
Maximum length of variable-length FILLER column	512 bytes
Maximum number of dbfiles	The total number of files that can be opened depends on the number of unique file descriptors that an operating system can support
Maximum number of users (connected and concurrent)	1000 (100 concurrent connections per node, up to a maximum of 10 nodes)
Maximum size of temp extract file	Set by `TEMP_EXTRACT_SIZEn` option. Platform limits are: Windows: 0 – 128GB Linux: 0 – 512GB
Table_name size	128 bytes
Number of columns per table	SAP IQ supports up to 45,000 columns in a table. You may see performance degradation if you have more than 10,000 columns in a table. The limit on the number of columns per table that allow NULLs is approximately 8*(database-page-size - 30).
Column_name size	128 bytes
Number of events per database	2^31 – 1 = 2 147 483 647
Number of files per database	Operating system limit that user can adjust; for example, using `NOFILE`. Typically, 2047 files per database
Number of indexes	2^32 (~4,000,000,000) per table
Index_name size	128 bytes
Number of rows per table	Limited by table size, upper limit 2^48 -1
Number of stored procedures per database	2^32 – 1 = 4 294 967 295
Number of tables or views in a single `FROM` clause	16 – 64, depending on the query, with join optimizer turned on

Item	Limit
Number of tables or views referenced per query	512
Number of tables per database	4,293,918,719
Number of tables referenced per transaction	No limit.
Number of UNION branches per query	2048. If each branch has multiple tables in the `FROM` clause, the limit on tables per query reduces the number of `UNION` branches allowed.
Number of values in an IN list	250,000
Row size	SAP recommends a limit of half the page size
Table size	Limited by database size

Related Information

sp_iqlmconfig Procedure
-gss Database Server Option
-c Database Server Option
-ch Database Server Option
-cl Database Server Option
-gp Database Server Option
TEMP_EXTRACT_SIZEn Options

7 Glossary

SAP IQ terms and concepts.

Term	Description
SAP Adaptive Server Enterprise	SAP Adaptive Server Enterprise is a high-performance relational database management system. You can use SAP IQ to query data in SAP Adaptive Server Enterprise databases.
Catalog store	A catalog store is the portion of each SAP IQ database that contains its metadata. The catalog store contains the SYSTEM dbspace and up to 12 additional other catalog dbspaces. The default name is `dbname.db`.
Component Integration Services (CIS)	Component Integration Services (CIS) provides SAP IQ users with direct access to relational or nonrelational databases on the mainframe, UNIX, or Windows servers.
Connection Profile	Connection profiles store connection information to a running SAP IQ server. The profile is primarily used to simplify user connections to a server. SAP IQ extends connection profiles to facilitate starting servers and creating databases.
Dbfile	A dbfile is an operating system file used to store data for a SAP IQ database. Each dbfile has a corresponding logical filename and physical file path. Each dbspace name, dbfile name, and physical file path must be unique. The dbfile name can be the same as the dbspace name. The `<SYSDBFILE>` view shows all the dbfiles in your database, including the catalog dbspace file, the IQ message file, dbfiles in the IQ main and temporary dbspaces, the transaction log file, and the SA temporary file.
Dbspace	A dbspace is a logical collection of dbfiles. If a database runs out of room, you can expand it by adding additional dbspaces. Users can move SAP IQ data off of disks and take the disks offline without any downtime.
Free list	A free list is a structure that SAP IQ uses to track which blocks are in use by a dbspace.
IQ database	An IQ database is a database that you create using an SAP IQ server. IQ databases are specially indexed to take advantage of the query speed of SAP IQ. Each IQ database that you create includes these stores: • An IQ main store – for data • A catalog store – for metadata • An IQ temporary store – for temporary data • The RLV store – for in-memory data store for concurrent row-level table updates It also generates an IQ message log file.

Term	Description
IQ main store	The IQ main store is the portion of each SAP IQ database that contains the IQ_SYSTEM_MAIN dbspace and other user dbspaces. The IQ main store contains persistent database structures, such as backup metadata and rollback data for committed transactions. The IQ main store is sometimes called the IQ store. → Tip Avoid storing user tables and indexes in IQ_SYSTEM_MAIN. Instead, create additional dbspaces — called user main dbspaces — to store user tables and indexes.
IQ message log	An IQ message log file created when the first user connects to an IQ database. The default name for this file is `dbname.iqmsg`. IQ_SYSTEM_MSG is a system dbspace that points to the file path of the database IQ message log file. IQ_SYSTEM_MSG is not considered a store because it doesn't store any data. SAP IQ logs error messages, status messages and insert notification messages in this file.
IQ temporary store	The IQ temporary store contains the IQ_SYSTEM_TEMP dbspace. The IQ temporary store is the portion of each IQ database that stores temporary tables and temporary scratch space data structures. The database server uses temporary data structures to sort and process data. Data in these tables persists only as long as you are connected to the database.
IQ transaction log	The IQ transaction log records changes to the database. The transaction log includes version information, free space, and other information you can use to recover from a system failure. By default, the transaction log is created in the same directory as the catalog store. The default name for this dbfile is `dbname.log`.
Metadata	Metadata is data that describes the data in your database — for example, the size and data type of each column in a table. Metadata for each SAP IQ database is stored in the catalog store.
Multiplex	A powerful feature in SAP IQ that provides application scalability through a clustered server configuration. SAP IQ multiplex allows concurrent data loads and queries via independent data processing nodes connected to a shared data source. Each multiplex server has its own catalog store and IQ temporary store and all the servers share a common IQ store. For more information, see *SAP IQ Administration: Multiplex*.
Object	An object can be a user-created table or index. Objects are divided into persistent objects, which remain in the database over user disconnects and server restarts, and temporary objects, tables and views that only remain in the database during the current session. Permanent tables are also called base tables.
Partition key	A partition key is a table column defined by the table creator that determines how a table should be partitioned.

Term	Description
Proxy table	A proxy table is a table object that maps to a table on a remote server, and whose column attributes and index information are derived from the object at the remote location. You can use proxy tables to search data in multiple SAP SQL Anywhere servers, SAP Adaptive Server Enterprise databases, and non-SAP databases. Conversely, you can also create proxy tables that enable you to query your SAP IQ, SAP SQL Anywhere, and SAP Adaptive Server Enterprise databases.
Range partition	A range partition is logical subset of table rows based on the values of a single table column.
RLV	Row-level snapshot versioning. The RLV store is an in-memory data store for concurrent row-level table updates. If a table is not enabled for RLV storage, it uses table-level snapshot versioning. See *SAP IQ Administration: In-Memory Row-Level Versioning*.
SAP SQL Anywhere	SAP SQL Anywhere is a transaction-processing relational database management system that can be used standalone or as a network server in a multiuser client/server or three-tier environment. SAP SQL Anywhere is specifically designed to use fewer memory and disk resources than the average database management system. SAP IQ is an extension of SAP SQL Anywhere, and supports many of the same features.
SAP SQL Anywhere database	Every SAP IQ database uses an SAP SQL Anywhere database for the catalog store.
Store	A store is one or more dbspaces that store persistent or temporary data for a special purpose. SAP IQ has these stores: the catalog store, the IQ main store, the IQ temporary store, and the RLV store.
Synchronization	Synchronization brings an outdated multiplex secondary node server up to date.
Table partition	A table partition is a collection of rows that is a subset of a user-created table. A given row cannot be placed in two different partitions. Each partition can be placed in its own dbspace and managed individually.
Tablespace	A tablespace unit of storage within the database that may be administered as a logical subset of total storage. You may allocate individual objects and subobjects to individual tablespaces. A tablespace in SAP IQ is referred to as a dbspace.

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.
About the icons:

- Links with the icon  : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
  - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
  - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon  : You are leaving the documentation for that particular SAP product or service and are entering an SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.
The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

www.sap.com/contactsap