



Introduction to SAP Sybase IQ

SAP Sybase IQ 16.0 SP2

DOCUMENT ID: DC38159-01-1602-01

LAST REVISED: November 2013

Copyright © 2013 by SAP AG or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and other countries. Please see <http://www.sap.com/corporate-en/legal/copyright/index.epx#trademark> for additional trademark information and notices.

Contents

About SAP Sybase IQ	1
Product Editions	1
Optional Features	2
Licensing	2
Related Products	2
SAP Sybase IQ Documentation Collection	3
Components	5
Tools and Utilities	5
Windows Options	6
Demo Database	7
Concepts	9
Key Principles	9
Column-Wise Data Structures	9
Bit-Wise Indexing	9
Architecture	10
Simplex Server Architecture	10
Multiplex Server Architecture	11
SAP Sybase IQ Database	12
Stores	12
IQ Log Files	14
Server Spaces	16
Scalability	19
Memory Use	19
Data Loads	19
MAIN_RESERVED_DBSPACE_MB Option	20
Page Size	20
Processing Threads	21
The Process Threading Model	21
Tuning Options	22
Disk Space	22
Intermediate Versioning	22

Contents

- Column-Based Indexes23
- Query Optimizer23
- Schema Design23
- UNION ALL Views23
 - UNION ALL Views for Faster Loads24
- Glossary31**
- Index35**

About SAP Sybase IQ

SAP® Sybase® IQ is a high-performance decision-support server designed specifically for mission-critical business intelligence, analytics, and data warehousing. Component Integration Services within SAP Sybase IQ provide direct access to relational and non-relational databases on mainframe, UNIX, or Windows servers.

Product Editions

SAP Sybase IQ is available in several product editions.

Product Edition	Description
Enterprise Edition	Supports all options and features. Optional features require additional licenses that are purchased separately.
Small Business Edition	Supports most of the functionality available in the Enterprise Edition, but does not support multiplex or any optional features.
Single Application Server	Provides the same functionality as the Small Business Edition, but is restricted to use in a one- or two-machine environment.
Evaluation Edition	Supports all options and features available in the Enterprise Edition for a 30-day evaluation period.
Express Edition	Supports all options and features available in the Evaluation Edition. Unlike the Evaluation Edition, the Express Edition does not expire, but does restrict the size of all IQ main dbspaces, permit other incremental license options, and cannot be deployed in a production environment.

Optional Features

Optional features extend the utility of the Enterprise Edition.

Option	Description
Unstructured Data Analytics	Supports Binary Large Object (BLOB) and Character Large Object (CLOB) storage and retrieval. Supports Binary Large Object (BLOB) and Character Large Object (CLOB) storage and retrieval.
Advanced Security	Supports these optional security features: <ul style="list-style-type: none">• Column encryption• FIPS network encryption• Kerberos connection authentication
Multiplex Grid	Lets you add additional nodes to a multiplex environment. SAP Sybase IQ requires an additional license to start secondary multiplex nodes (readers/writers).
Very Large Database Management	Lets you add multiple table spaces and dbspaces to logically partition data into manageable subsets. SAP Sybase IQ requires an additional license when the server creates or starts with two or more IQ user dbspaces.

Licensing

All product editions except the Evaluation Edition require a license. Optional features are sold and licensed separately.

The Evaluation Edition is an unlicensed server, which provides full access to all features and options available in the Enterprise Edition. To run the Evaluation Edition beyond the thirty-day evaluation period, however, you must purchase and install the appropriate license.

Related Products

SAP Sybase IQ is part of a large group of database products that provide a complete analytics infrastructure.

- SAP HANA® – The SAP HANA database is an in-memory database that combines transactional data processing, analytical data processing, and application logic processing functionality in memory.

See *Products > Database > Database Solutions > SAP HANA Database* on *sap.com* for more information.

- SAP SQL Anywhere® – SAP SQL Anywhere is a comprehensive data management package for server, desktop, mobile, and remote office environments. SAP Sybase IQ is an extension of SAP SQL Anywhere, and incorporates many of its features.

See *Products > Database > Database Solutions > SAP Sybase SQL Anywhere* on *sap.com* for more information.

- SAP Sybase PowerDesigner® – PowerDesigner is the industry-leading modeling and metadata management solution for data, information, and enterprise architecture.

See *Products > Database > Database Solutions > SAP Sybase PowerDesigner* on *sap.com* for more information.

SAP Sybase IQ Documentation Collection

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

New Release Information

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

Getting Started

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

Administration

Document Name	Contents
<i>Administration: Database</i>	Database setup and troubleshooting.

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

Reference

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

Optional Features

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

Document Name	Contents
<i>Administration: User Management and Security</i>	Contains Advanced Security option information: <ul style="list-style-type: none"> LDAP User Authentication configuration Implement Kerberos authentication

Licensing

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

Components

SAP Sybase IQ includes tools and utilities that help you work with the server.

Tools and Utilities

SAP Sybase IQ includes utilities that help you perform routine management tasks.

Table 1. SAP Sybase IQ Utilities

Utility	Description
Sybase Control Center	<p>SAP Control Center (SCC) is a Web-based administrative console for Sybase products. At least one SCC server is required to deploy SAP Control Center.</p> <p>For information on managing SAP Sybase IQ with SCC, see the SAP Control Center for SAP Sybase IQ online help in SCC or at http://sybooks.sybase.com/sybooks/sybooks.xhtml?prodID=10680.</p>
Interactive SQL	<p>Interactive SQL is a tool included with SAP Sybase IQ that lets you execute SQL statements, build scripts, and display database data.</p> <p>See the <i>Utility Guide</i> for more information about Interactive SQL.</p>
Command line utilities	<p>Command line utilities perform database administration tasks.</p> <p>See the <i>Utility Guide</i> for more information.</p>

Note: Available utilities depend on your operating system and installation options.

Windows Options

On Windows, the Sybase IQ Program Group provides quick access to common utilities. Available options differ for Server and Client installations.

Server Installations

Table 2. SAP Sybase IQ Server Options

Option	Description
Interactive SQL	Starts the Java-based Interactive SQL utility for sending SQL statements to a database.
ODBC Administrator	Manages ODBC connections to your databases for 32-bit or 64-bit platforms.
Raw Device Access	Adds or removes users and raw devices from the registry.
Sample Applications and Projects	Displays a list of sample applications and projects.
Start SAP Sybase IQ Demo Database	Starts the database server running the demo database.
SAP Sybase IQ ReadMe File	Lists last-minute changes to documentation and software and special material that needs emphasis.
Sybase Service Manager	Lets you configure, modify, or delete SAP Sybase IQ as a Windows service.
Sybase on the Web	Opens a Web page that introduces Sybase online resources.

Client Installations

Table 3. SAP Sybase IQ Client Options

Option	Description
Sample Applications and Projects	Displays a list of sample applications and projects. See the <i>Programming</i> guide.
Sybase on the Web	Opens a Web page that introduces Sybase online resources.
ODBC Administrator	Manages ODBC connections to your databases for 32-bit or 64-bit platforms.
Interactive SQL	Starts the Java-based Interactive SQL utility for sending SQL statements to a database.

Demo Database

Many of the examples throughout the documentation use the demo database (`iqdemo.db`) as a data source. The demo database is installed as part of the SAP Sybase IQ Server Suite and resides on the server. Client access is provided by tools installed as part of the SAP Sybase IQ Client Suite.

Table Names and Owners

The demo database includes data and utility tables.

Tables in the demo database are delineated by ownership. Most of the SQL examples in this document require access to `GROUPO` tables as a minimum.

GROUPO Tables

`GROUPO` tables contain internal information about a fictional company that sells athletic clothing. Sample data includes information about the company (employees, departments, and financial data) as well as product information (products) and sales information (sales orders, customers, and contacts).

Table 4. GroupO Table Names

Name	Description
Contacts	Customer contacts and sales leads.
Customers	Customer names and addresses.
Departments	Company departments, such as manager and name.
Employees	Employee information, such as name, salary, and location.
FinancialCodes	Each expense and revenue item has a financial code.
FinancialData	Quarter-by-quarter financial information about the company.
Products	Product information, such as price and quantity available.
SalesOrderItems	Sales order items. Each order consists of one or more items. Information about sales order items is held in a separate table.
SalesOrders	Individual sale orders, including customer ID, OrderDate, Financial-Code, Region, and SalesRepresentative.

DBA-Owned Tables

Tables owned by the DBA include utility and sample data tables. Access requires either the `SELECT ANY TABLE` system privilege or `SELECT` privilege on the DBA-owned table.

Table 5. DBA-Owned Table Names

Table	Description
iq_dummy	<p>iq_dummy is a one-row, one-column utility table that you can use to extract information from the database. For example, running the NOW() function against iq_dummy returns the current date and time:</p> <pre>SELECT NOW() FROM iq_dummy</pre> <p>Use of the DUMMY system table in SAP Sybase IQ is implied for all queries that do not have a FROM clause.</p>
empl	Sample employee table that includes dept_id, start_date, name, and salary columns.
sale	Sample sales table that includes prod_id, month_num, rep_id, and sales columns.

SYSOPTION - DEFAULTS Table

SYSOPTIONDEFAULTS is a utility table owned by DBO that contains all SAP Sybase IQ option names and values. You can query this table to see all default option values.

Note: The demo database is case-insensitive. This means that case is not considered in comparison and string operations. For example, you can type user IDs and passwords in either uppercase or lowercase when using the demo database. Note that, unlike the demo database, any SAP Sybase IQ databases you create are case-sensitive by default.

Concepts

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

Key Principles

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

Column-Wise Data Structures

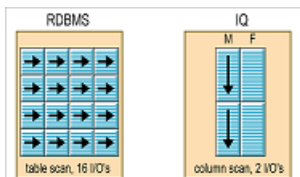
Store data column-wise rather than row-wise.

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"
```



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 Sybase IQ indexes are based on the cardinality of the data and the way the data will be used. SAP Sybase 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:


Concepts

- 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: These 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					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

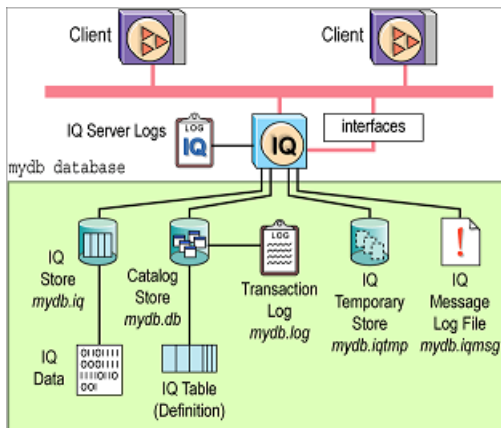
Architecture

SAP Sybase IQ supports both simplex and multiplex architecture.

Simplex Server Architecture

Simplex is a single instance of an SAP Sybase 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).

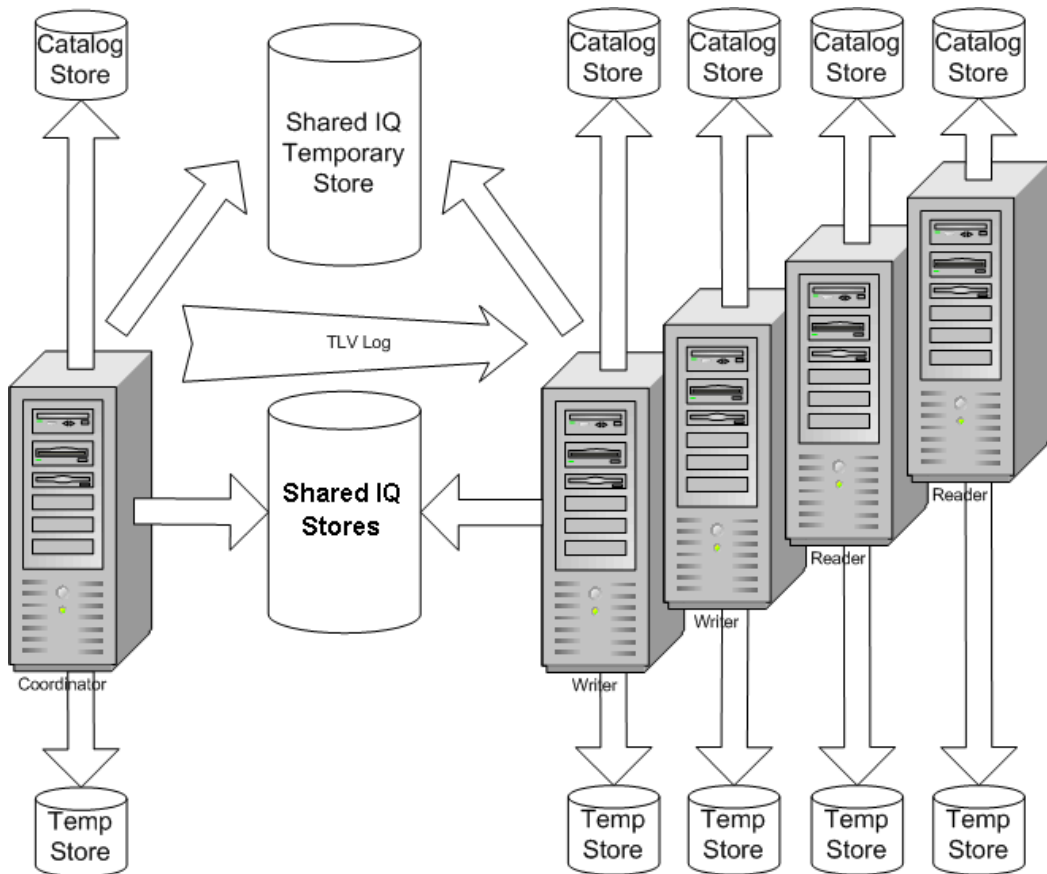


Multiplex Server Architecture

Multiplex is multiple instances of SAP Sybase IQ servers running on multiple machines. SAP Sybase 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 DLL operations and communicates information about new table versions to the secondary servers.

Concepts



SAP Sybase IQ Database

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

Stores

There are five stores in the SAP Sybase IQ database.

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 Sybase IQ database. It is always created on a file system and consists of one object for the SAP Sybase 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. Data from the system tables may be browsed using the Interactive SQL or Sybase Control Center.

The catalog store resembles a 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 Sybase IQ database.

For a complete listing of system tables and views, see *Reference: Building Blocks, Tables, and Procedures*.

IQ Main Store (for data)

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

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

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

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

The IQ main store consists of at least two dbspaces:

- **IQ_SYSTEM_MAIN** – the first dbspace created when creating a database, it contains system metadata. It should not 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` controls where tables and indexes are stored.

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.iqttmp` (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. If a file system, it should have the `.iqtmp` file extension.

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** commands 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 *Administration: In-Memory Row-Level Versioning*.

IQ Shared Temporary Store (for temporary data in a multiplex)

The IQ shared temporary store is required for distributed query processing (DQP).

The IQ shared temporary store contains temporary structures shared among nodes for DQP, and may also contain temporary user objects or local nonversioning temporary objects.

When you create a database, SAP Sybase IQ creates a single dbspace for the IQ shared temporary store, `IQ_SHARED_TEMP`. Initially, this dbspace contains no files. Before you can use DQP, you must add one or more 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, improve performance and support DQP.

IQ Log Files

There are four log files in the SAP Sybase 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.

Message Log

The IQ message file is created when an SAP Sybase 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 it the file wraps around and overwrites the oldest records.

The IQ message file captures:

- SAP Sybase 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.

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 Sybase 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 commands associated with SAP Sybase 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

Server Logs

The SAP Sybase 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`

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.

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

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.

Server Spaces

There are two server spaces in the SAP Sybase IQ database.

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.

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.

Scalability

SAP Sybase 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 Sybase 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 Sybase IQ database, but is especially important as your database grows to the multiterabyte scale.

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

Memory Use

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

SAP Sybase 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 Sybase 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.

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.

Scalability

- 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.

MAIN_RESERVED_DBSPACE_MB Option

Controls the amount of space SAP Sybase IQ reserves in the IQ main store.

Allowed Values

Integer greater than or equal to 200, in megabytes

Default

200; SAP Sybase IQ actually reserves a maximum of 50% and a minimum of 1% of the last read-write file in `IQ_SYSTEM_MAIN`

Scope

Option can be set at the database (PUBLIC) level only.

Requires the SET ANY SYSTEM OPTION system privilege to set this option. Takes effect immediately.

Remarks

MAIN_RESERVED_DBSPACE_MB controls the amount of space SAP Sybase IQ sets aside in the IQ main store for certain small but critical data structures used during release savepoint, commit, and checkpoint operations. For a production database, set this value between 200MB and 1GB, or at least 20 percent of `IQ_SYSTEM_MAIN` size. The larger your IQ page size and number of concurrent connections, the more reserved space you need.

Reserved space size is calculated as a maximum of 50 percent and a minimum of 1 percent of the last read-write file in `IQ_SYSTEM_MAIN`.

SAP Sybase IQ ignores the **MAIN_RESERVED_DBSPACE_MB** option if the actual dbspace size is less than twice the size of the **MAIN_RESERVED_DBSPACE_MB** value. In dbspaces less than 100MB (such as the demo database), half the usable space may be reserved.

Page Size

When you create your SAP Sybase 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.

Processing Threads

SAP Sybase 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.

The Process Threading Model

SAP Sybase IQ uses operating system kernel threads for best performance. By default, SAP Sybase 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 Sybase 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 Sybase 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 Sybase IQ threads. It is also possible that **-iqmt** is set too low for the number of connections.

SAP Sybase 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` command to set the maximum number of threads for a single user. The `SET OPTION`

`MAX_IQ_THREADS_PER_TEAM` command 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 a particular operation consumes. For example, you can set this option before issuing an **INSERT**, **LOAD**, **BACKUP DATABASE**, or **RESTORE DATABASE** command.

Tuning Options

Tuning options that provide faster query execution.

Disk Space

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

The most important factors in managing disk I/O for a SAP Sybase 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.

Intermediate Versioning

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

In particular, SAP Sybase 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.

Column-Based Indexes

SAP Sybase 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.

Query Optimizer

The SAP Sybase 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 Sybase IQ to choose the best plan for most queries and most databases, including the largest ones.

Schema Design

SAP Sybase 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.

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 Sybase 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.

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

This strategy can improve load performance, but may negatively impact the performance of some types of queries. Most types of queries have roughly similar performance against a single base table or against a UNION ALL view over smaller base tables, as long as the view definition satisfies all 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. UNION ALL operators are not fully parallel in this release. Their use may limit query parallelism.

UNION Operation

Combines the results of two or more select statements.

Quick Links:

Go to Parameters on page 25

Go to Examples on page 25

Go to Usage on page 25

Go to Standards on page 26

Go to Permissions on page 26

Syntax

```
select-without-order-by
... UNION [ ALL ] select-without-order-by
... [ UNION [ ALL ] select-without-order-by ]...
... [ ORDER BY integer [ ASC | DESC ] [, ...] ]
```

Parameters

(back to top) on page 24

- **ALL** – the results of UNION ALL are the combined results of the component **SELECT** statements. The results of UNION are the same as UNION ALL, except that duplicate rows are eliminated. Eliminating duplicates requires extra processing, so UNION ALL should be used instead of UNION where possible.
- **ORDER BY** – only integers are allowed in the order by list. These integers specify the position of the columns to be sorted.

Examples

(back to top) on page 24

- **Example 1** – list all distinct surnames of employees and customers:

```
SELECT Surname
FROM Employees
UNION
SELECT Surname
FROM Customers
```

Usage

(back to top) on page 24

The results of several **SELECT** statements can be combined into a larger result using a UNION clause. The component **SELECT** statements must each have the same number of items in the select list, and cannot contain an ORDER BY clause. See *FROM Clause*.

If corresponding items in two select lists have different data types, SAP Sybase IQ chooses a data type for the corresponding column in the result, and automatically converts the columns in each component **SELECT** statement appropriately.

The column names displayed are the same column names that display for the first **SELECT** statement.

Note: When **SELECT** statements include constant values and UNION ALL views but omit the FROM clause, use `iq_dummy` to avoid errors. See *FROM Clause* for details.

Standards

(*back to top*) on page 24

- SQL—ISO/ANSI SQL compliant.
- SAP Sybase Database product—Supported by SAP Adaptive Server® Enterprise, which also supports a **COMPUTE** clause.

Permissions

(*back to top*) on page 24

Requires SELECT privilege for each component of the **SELECT** statements.

Queries Referencing UNION ALL Views

To adjust performance for queries that reference UNION ALL views, set the JOIN_PREFERENCE option, which affects joins between UNION ALL views.

All partitions in a UNION ALL view must have a complete set of indexes defined for optimization to work. Queries with DISTINCT will tend to run more slowly using a UNION ALL view than a base table.

SAP Sybase IQ includes optimizations for UNION ALL views, including:

- 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.

JOIN_PREFERENCE Option

Controls the choice of algorithms when processing joins.

Allowed Values

Value	Action
0	Let the optimizer choose

Value	Action
1	Prefer sort-merge
2	Prefer nested-loop
3	Prefer nested-loop push-down
4	Prefer hash
5	Prefer hash push-down
6	Prefer asymmetric sort-merge join
7	Prefer sort-merge push-down
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
-1	Avoid sort-merge
-2	Avoid nested-loop
-3	Avoid nested-loop push-down
-4	Avoid hash
-5	Avoid hash push-down
-6	Avoid asymmetric sort-merge join
-7	Avoid sort-merge push-down
-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

Value	Action
-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

Default

0

Scope

Option can be set at the database (PUBLIC) or user level. When set at the database level, the value becomes the default for any new user, but has no impact on existing users. When set at the user level, overrides the PUBLIC value for that user only. No system privilege is required to set option for self. System privilege is required to set at database level or at user level for any user other than self.

Requires the SET ANY PUBLIC OPTION system privilege to set this option. Can be set temporary for an individual connection or for the PUBLIC role. Takes effect immediately.

Remarks

For joins within a query, the SAP Sybase IQ optimizer has a choice of several algorithms for processing the join. **JOIN_PREFERENCE** allows you to override the optimizer's cost-based decision when choosing the algorithm to use. It does not override internal rules that determine whether an algorithm is legal within the query engine. If you set it to any nonzero value, every join in a query is affected; you cannot use it to selectively modify one join out of several in a query, but join condition hint strings can do so.

This option is normally used for internal testing or tuning of report queries, and only experienced DBAs should use it.

Simple equality join predicates can be tagged with a predicate hint that allows a join preference to be specified for just that one join. If the same join has more than one join condition with a local join preference, and if those hints are not the same value, then all local preferences are ignored for that join. Local join preferences do not affect the join order chosen by the optimizer.

This example requests a hash join:

```
AND (T.X = 10 * R.x, 'J:4')
```

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;
```


Glossary

SAP Sybase IQ terms and concepts.

- **Adaptive Server Enterprise** – Adaptive Server Enterprise is a high-performance relational database management system. You can use SAP Sybase IQ to query data in Adaptive Server Enterprise databases.
- **Catalog store** – A catalog store is the portion of each SAP Sybase 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 Sybase 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 Sybase IQ server. The profile is primarily used to simplify user connections to a server. SAP Sybase 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 Sybase 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 *SYSDBFIL*E 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 Sybase IQ data off of disks and take the disks offline without any downtime.
- **Free list** – A free list is a structure that SAP Sybase 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 Sybase IQ server. IQ databases are specially indexed to take advantage of the query speed of SAP Sybase 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), and the RLV store (for in-memory data store for concurrent row-level table updates). It also generates an IQ message log file.

- **IQ main store** – The IQ main store is the portion of each SAP Sybase 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` and 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 Sybase 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 Sybase IQ database is stored in the catalog store.
- **Multiplex** – A powerful feature in SAP Sybase IQ that provides application scalability through a clustered server configuration. SAP Sybase 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 *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.
- **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, Adaptive Server Enterprise databases, and non-Sybase databases. Conversely, you can also create proxy tables that enable you to query your SAP Sybase IQ, SAP SQL Anywhere, and 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 *Administration: In-Memory Row-Level Versioning*.
- **SAP SQL Anywhere** – SAP SQL Anywhere is a transaction-processing relational database management system which 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 Sybase IQ is an extension of SAP SQL Anywhere, and supports many of the same features.

- **SAP SQL Anywhere database** – Every SAP Sybase 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 Sybase 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 Sybase IQ is referred to as a dbspace.

Index

A

About SAP Sybase IQ 1
Adaptive Server Enterprise 31
Advanced Security 1

C

Catalog Store 31
catalog store transaction log 15
component integration services (CIS) 31
Concepts 9
connection profile 31
connectivity sample applications 6

D

Database, demo
 iqdemo 7
databases
 very large 19
Databases
 demo 7
dba tables 7
dbfile 31
dbisql
 See Interactive SQL
dbspace 31
dbspces 17
demo database
 starting 6
Demo database
 DBA-owned tables 7
 GROUPO tables 7
 iq_dummy table 7
 iqdemo 7
Demo Database 7
documentation 3

E

Enterprise Edition 1
Evaluation Edition 1

F

Features, optional
 Advanced Security 1

Multiplex Grid 1
Unstructured Data Analytics 1
Very Large Database Management 1
free list 31
FROM Clause 26

G

glossary 31–33
GROUPO tables 7

I

Interactive SQL 5
IQ
 See SAP Sybase IQ about
IQ database 31
IQ message Log 32
IQ store
 reserving space 20
IQ Store 31
IQ temporary store 32
IQ transaction log 32
iq_dummy table 7
iqdemo.db file
 about 7

J

java sample applications 6
JOIN_PREFERENCE 26
JOIN_PREFERENCE option 26

L

licensing 1
lightweight processes 21

M

MAIN_RESERVED_DBSPACE_MB option 20
make files 6
manuals 3

Index

memory

- lightweight processes 21
- multithreading 21
- process threading model 21

metadata 32

multiplex 32

Multiplex Grid 1

multithreading

- performance impact 21

O

object 32

ODBC sample applications 6

Optional features

- Advanced Security 1
- Multiplex Grid 1
- Unstructured Data Analytics 1
- Very Large Database Management 1

Options

- JOIN_PREFERENCE 26

out-of-space conditions

- preventing 20

P

partition key 32

Partitioned Table 26

persistence 16

process threading model 21

Product editions

- Enterprise Edition 1
- Evaluation Edition 1
- Single Application Server 1
- Small Business Edition 1

products, related

- PowerDesigner 2
- SAP SQL Anywhere 2
- WorkSpace 2

program group 6

proxy database 32

Pushdown Join 26

R

range partition 32

Raw Device Access 6

readme file 6

real-time updating 14

Related products

- PowerDesigner 2
- SAP SQL Anywhere 2

Resource Use

- Loading with UNION ALL 24

Resource Use Options 22

RLV 33

RLV store 14

row-level versioning 14

S

sample applications 6

SAP Sybase IQ

- components 5
- Concepts 9
- demo database 7
- documentation 3
- related products 2
- Service Manager 6

Schema Design

- UNION ALL 29

Service Manager 6

shared temporary store 14

Single Application Server 1

Small Business Edition 1

SQL Anywhere 33

store 33

Sybase Control Center 5

synchronization 33

System Resources

- Resource Use Options 22

T

table partitions 33

tables, iqdemo

- dba tables 7
- GROUPO tables 7
- iq_dummy table 7

tablespace 33

U

UNION ALL

- Loading 24
- Rules 26
- View Performance 29
- Views 26

UNION operation 24
Unstructured Data Analytics 1
user manuals 3

V

Very Large Database Management 1
VLDB
 managing 19

