SYBASE®

An **SAP** Company

Quick Start

Sybase IQ 15.4

UNIX/Linux

Contents

Contents

Getting Started

The Sybase® IQ Quick Start shows you how to deploy a single server, convert the server to multiplex, and run a query that distributes processing across all nodes in the multiplex.

Tasks in this guide are sequential and progressive. Complete the tasks in one sequence before you progress to the next. The time required to complete all tasks is about 20 minutes.

1. *Create a Single Server*

   The first set of tasks describes how to deploy the demo database in a single-server environment.

2. *Convert the Server to Multiplex*

   Convert the demo database to multiplex.

3. *Distribute the Workload*

   In a multiplex environment, the query optimizer breaks complex queries into fragments and distributes the fragments to different nodes for processing.

Create a Single Server

The first set of tasks describes how to deploy the demo database in a single-server environment.

1. *Setting the Environment Variables*

   Sybase IQ environment variables point to different directories and file locations. Normally, the installer automatically sets many of these variables, but if you work in a new terminal or console, you must manually set these variables.

2. *Creating the Demo Database*

   Sybase IQ includes a demo database that includes sample data about a fictional company that sells athletic clothing. Scripts that create and load the demo database are installed as part of Sybase IQ Server Suite.

3. *Running a Sample Query*

   Interactive SQL is a utility that ships with Sybase IQ. Use Interactive SQL to connect to the server and run a sample query.

See also

Setting the Environment Variables

Sybase IQ environment variables point to different directories and file locations. Normally, the installer automatically sets many of these variables, but if you work in a new terminal or console, you must manually set these variables.

1. Open a console or terminal.
2. Set the $SYBASE variable to the parent installation directory.

Shell	Command
tcsh or C (csh)	`setenv SYBASE <directory name>`
Korn (ksh), Bash, or Bourne (sh)	`SYBASE=<directory name>;` `export SYBASE`

3. Set the $IQDIR15 variable to the Sybase IQ product directory.

Shell	Command
tcsh or C (csh)	`source $SYBASE/IQ-15_4/` `IQ-15_4.csh`
Korn (ksh), Bash, or Bourne (sh)	`$SYBASE/IQ-15_4/IQ-15_4.sh`

Note: For additional information about environment variables, see *Reference: Building Blocks, Tables, and Procedures > File Locations and Installation Settings > Environment Variables.*

Creating the Demo Database

Sybase IQ includes a demo database that includes sample data about a fictional company that sells athletic clothing. Scripts that create and load the demo database are installed as part of Sybase IQ Server Suite.

These instructions differ from those in the *Installation and Configuration Guide*. This task creates the demo database with an absolute path in a temporary directory named `myiqdemo`.

1. Change to the `$SYBASE` directory.
2. Create a subdirectory named `myiqdemo`.

   This is the temporary directory for the demo database. The full path to this directory is: .
   `$SYBASE/myiqdemo`
3. Change to `myiqdemo`.
4. Create a subdirectory named `w1`.

   This directory is used later for the multiplex conversion. The full path to this directory is:

```
$SYBASE/myiqdemo/w1
```

5. Enter:

```
$IQDIR15/demo/mkiqdemo.sh -absolute
```

The **-absolute** switch creates the demo database with an absolute path, which is required to convert the demo database to multiplex. For other options, enter:

```
$IQDIR15/demo/mkiqdemo.sh -help
```

6. Start the demo database:

```
start_iq @iqdemo.cfg iqdemo.db
```

Running a Sample Query

Interactive SQL is a utility that ships with Sybase IQ. Use Interactive SQL to connect to the server and run a sample query.

1. Enter:

   dbisql

   Interactive SQL starts and opens a connection dialog.

2. Use these connection parameters:

Tab Name	Field	Value
Identifica-tion	**Authentication**	**Database**
	User ID	DBA or dba (case-insensitive)
	Password	sql (case-sensitive)
	Action	**Connect to a running database on this computer**

3. In the SQL Statements window, enter:

```
SELECT
SalesOrders.ID,
SalesOrders.OrderDate,
Customers.CompanyName

FROM SalesOrders
KEY JOIN Customers
WHERE OrderDate > '1994/01/01'
ORDER BY OrderDate
```

Other query options are available. See *Installation and Configuration > Testing Installed Products > Demo Database > Table Names* for a list of tables and owners.

4. Press **F9** to run the query.

5. Click **File > Exit** to exit Interactive SQL.

Convert the Server to Multiplex

Convert the demo database to multiplex.

Although there are command line equivalents, all server-related tasks in this sequence use Sybase Central, which is an administrative console for Sybase products. Sybase Central lets you administer your server in a graphical environment.

Using Sybase Central to convert the demo database to multiplex lets you see the changes to your processing environment as they occur.

1. *Starting the Sybase IQ Agent*

    Sybase IQ Agent provides a communication channel between Sybase IQ and Sybase Central. Sybase Central requires an active Sybase IQ Agent to perform server-related tasks.

2. *Starting Sybase Central*

    Start Sybase Central and connect to your server.

3. *Defining the Multiplex Server*

    Use the Create Server wizard to convert a single server to multiplex.

4. *Loading Data into the Multiplex*

    Use the coordinator and writer to load the multiplex to make the best use of multiplex resources.

See also
- *Create a Single Server* on page 1
- *Distribute the Workload* on page 10

Starting the Sybase IQ Agent

Sybase IQ Agent provides a communication channel between Sybase IQ and Sybase Central. Sybase Central requires an active Sybase IQ Agent to perform server-related tasks.

1. In a console, change to $IQDIR15/bin64, and enter:

    ```
    S99SybaseIQAgent15
    ```

    This command starts the Sybase IQ Agent and sets the port to **1099**.

2. Verify that the agent is running:

    ```
    stop_iq -agent
    ```

Note: Sybase IQ writes the agent log file to $IQDIR15/logfiles/ SybaseIQAgentNNNN.MMM.log, where NNNN is the port number and MMM is a sequential number. If the environment variable $IQLOGDIR15 is set, the agent log files are in the $IQLOGDIR15 directory.

Starting Sybase Central

Start Sybase Central and connect to your server.

1. Change to:

   `$IQDIR15/bin64`

2. Enter:

   `scjview`

3. Close the **Tips** and **Welcome** windows, if necessary.
4. From the **View** menu, choose **Folders**.
5. From the **Connections** menu, choose **Connect with Sybase IQ**.
6. On the **Connect** dialog, enter:

Tab Name	Field	Value
Identifica-tion	**Authentication**	**Database**
	User ID	DBA or dba (case-insensitive)
	Password	sql (case-sensitive)
	Action	**Connect to a running database on this computer**

Defining the Multiplex Server

Use the Create Server wizard to convert a single server to multiplex.

The Quick Start uses the placeholder *<localhost>* to identify the host. When you see this placeholder in any instruction, substitute your host name.

1. In **Folders**, right-click the *<localhost>*_**iqdemo** server icon, choose **Agent > Properties**.
2. Update the **Agent Properties**, as necessary:

Field Name	Value
Username	DBA or dba (case-insensitive)
Password	sql (case-sensitive)
Port	1099 (default)
Timeout	5 (default in minutes)

3. Click **OK** to return to Sybase Central.
4. In **Folders**, right-click *<localhost>*_**iqdemo** server, choose **Convert to Multiplex**.

   The Multiplex Server Conversion wizard includes several parts:

Title	Description
Node Configuration	Configuration settings for a new multiplex node
Database Path	An absolute path to a database file on the new multiplex secondary server
Temporary Dbspace Configuration	Size and path to a temporary dbspace file
Review of Settings	Summarizes multiplex configuration settings

5. Follow the instructions on the screen.

See also
- *Loading Data into the Multiplex* on page 9

Sample Conversion Values
The Server Creation wizard derives some default field values from the single server you are converting. Refer to these sample conversion values for the Quick Start.

Node Configuration
The Node Configuration dialog sets the multiplex node configuration properties.

Field Name	Description	
Multiplex Server Name	The wizard derives the default value for this field from the name of the single server you are converting: `<localhost>_iqdemo` Change to: `<localhost>_iqdemo_mpx`	
Server Name	This is the name of the first secondary server in your multiplex. This value must be different from server you are converting, which becomes the coordinator server for the multiplex. Use this value: `<localhost>_iqdemo_w1`	
Host Information	For the Quick Start, use the default Host Name and Port Number provided by the wizard. Using these values lets you set up the multiplex with multiple servers on the same host.	
Agent	For the Quick Start use these values to bind the multiplex to the Sybase IQ Agent:	
	Port	`1099` (default)
	Username	`DBA` (case-insensitive)
	Password	`sql` (case-sensitive)

Field Name	Description
Mode	Change from **Reader** to **Writer**. You will use this node later to load data.
Create administrative shell scripts.	Generates scripts for starting, stopping, and synchronizing servers.

Database Path

Database Path dialog sets an absolute path to database stores. The *<pathname>* placeholder refers to the path. When you see this placeholder, substitute the full path name.

Field	Description
Database Path	For the Quick Start, use the `<pathname>`/myiqdemo/w1 subdirectory and identify the secondary server database file as `iqdemo_w1.db`.
	Although you can browse to the path, you must enter the database file name (`iqdemo_w1.db`).
	The full field value should read:
	`<pathname>`/myiqdemo/w1/iqdemo_w1.db
Coordinator's configuration options	Select the **Use Coordinator's configuration options** to copy the `params.cfg` options for the existing database to the new multiplex server's database directory.

Temporary Dbspace Configuration

Temporary Dbspace Configuration dialog sets the path to the temporary dbspace file. As in the previous section, substitute the full path for the *<pathname>* placeholder.

Field	Description
Temporary dbspace file path	For the Quick Start, set the path to `<pathname>`/myiqdemo/w1/iqdemo_w1.iqtmp.
	Although you can browse to the path, you must enter the database file name (`iqdemo_w1.iqtmp`).
	The full field value should read:
	`<pathname>`/myiqdemo/w1/iqdemo_w1.iqtmp
RAW device option	Unselect the **This is a RAW device** option to display the default File size and Reserve size values. For the Quick Start, do not change these values.

Review of Settings
Review the conversion values.

Multiplex Views

Converting a single server to multiplex adds different views to Sybase Central.

Folders View
The new multiplex becomes part of the Multiplex folder. Expand the new Multiplex folder to see the objects in the multiplex.

View	Description
Multiplex Server	The name of the new multiplex is: *<localhost>_iqdemo_mpx*
Coordinator	The coordinator manages multiplex resources. The name of the coordinator is: *<localhost>_iqdemo*
Writer	The writer can run read-only and read-write operations against shared IQ objects. The name of the writer is: *<localhost>_iqdemo_w1*

Multiplex Properties
Multiplex properties appear in the right pane. Select the multiplex server or one of its member objects in the Folders view to display its properties.

View	Description
Servers	The Servers tab identifies the servers in the multiplex: • *<localhost>_iqdemo* (coordinator) • *<localhost>_iqdemo_w1* (writer)
Topology	The Topology tab provides a graphical view of the multiplex. Right-click any object in the multiplex topology to display command options for that object.
Performance Monitor	The Performance Monitor tab lets you monitor performance statistics at multiplex and server levels. At the multiplex level, you can monitor one statistic at a time across multiple servers. At the server level, you can simultaneously monitor as many as 10 statistics.

Loading Data into the Multiplex

Use the coordinator and writer to load the multiplex to make the best use of multiplex resources.

Because we want to load the same data we used to create the demo database, the sample LOAD scripts include a **TRUNCATE** statement, which deletes all rows in the target table before loading the data.

1. In **Folders**, right-click **<localhost>_iqdemo > Open Interactive SQL**.

   This starts Interactive SQL with a connection to the coordinator.

2. In Interactive SQL, enter:

```
TRUNCATE TABLE Contacts;
LOAD TABLE Contacts (ID, Surname,GivenName,
Title, Street, City, State, Country, PostalCode,
Phone, Fax, CustomerID)
USING FILE '../IQ-15_4/demo/adata/contact.dat'
ROW DELIMITED BY '|'
ESCAPES OFF;
COMMIT;
```

   Enter the full path to the data file on a single line.

3. Press **F9** to load the Contacts table on the coordinator.

4. Close Interactive SQL.

5. In **Folders**, right-click **<localhost>_iqdemo_w1 > Open Interactive SQL**.

   This starts Interactive SQL with a connection to the writer.

6. In Interactive SQL, enter:

```
TRUNCATE TABLE SalesOrderItems;
LOAD TABLE SalesOrderItems (ID, LineID,
ProductID, Quantity, ShipDate)
USING FILE '../../IQ-15_4/demo/adata/sales_oi.dat'
ROW DELIMITED BY '|'
ESCAPES OFF;
COMMIT;
```

   Enter the full path to the data file on the single line.

7. Press **F9** to load the SalesOrderItems Table on the writer.

8. Close Interactive SQL.

See also

- *Defining the Multiplex Server* on page 5

Distribute the Workload

In a multiplex environment, the query optimizer breaks complex queries into fragments and distributes the fragments to different nodes for processing.

Distributing a query involves a leader node and one or more worker nodes. A query originates on the leader node. A worker node can be any node in the multiplex that can accept some of the query processing work. Any multiplex node type (reader, writer, or coordinator) can be a leader or a worker node.

The query optimizer determines whether to distribute the query. To distribute the query, the leader node breaks the query into parallel fragments and assigns the fragments to the worker nodes. The workers process the fragments and return the intermediate results. Queries that are unlikely to scale appropriately or degrade performance are not distributed; they are executed on a single node in the multiplex.

This part of the Quick Start shows how to change the multiplex to accommodate a large volume of data and run a distributed query.

1. *Adding Shared Temporary Storage*

   Distributed Query Processing (DQP) requires temporary storage that all multiplex nodes can share. This task creates a resource on the coordinator that provides temporary storage for the coordinator and writer.

2. *Increasing the iq_main File Size*

   Before you run a distributed query, increase the file size in user dbspace iq_main.

3. *Running a Distributed Query*

   The query engine is more likely to distribute a query that processes a large number of rows than those that require fewer resources.

4. *Reviewing the Query Plan*

   Check the multiplex node directories to review the query plan. The comprehensive query plan is in the leader node directory, and each distributed work unit has an HTML file generated by the node that received it.

See also
*   *Convert the Server to Multiplex* on page 4

Adding Shared Temporary Storage

Distributed Query Processing (DQP) requires temporary storage that all multiplex nodes can share. This task creates a resource on the coordinator that provides temporary storage for the coordinator and writer.

1. In **Folders**, expand **<localhost>_iqdemo_mpx > iqdemo > Dbspaces**.

2. In the **Dbspaces** folder, right-click **IQ_SHARED_TEMP**, choose **New > File**.
3. On the Add File(s) dialog, click **Add**.
4. Update the File Detail(s) dialog:

Option	Value
Logical File Name	Use `shared_temp` as the logical file name.
Physical File Path	Set the path to the `$SYBASE/myiqdemo` subdirectory and identify the secondary server database file as `shared_temp.iqtmp`. You can unselect **This is a raw device** to browse to the path, but you must manually enter the file name. The full field value should read: `$SYBASE/myiqdemo/shared_temp.iqtmp`
Raw device option	Unselect **This is a raw device** to display the default File size and Reserve size values.
File Sizes	Update with these values:

	File Size	300MB
	Reserve Size	90MB

5. Click **OK** to return to the Add File(s) dialog, then click **OK > Next > Finish**.
6. On the Files tab, right-click **shared_temp**, and choose **Properties**.
7. Change the Mode to **Read/Write**.
8. Click **Apply > OK**.

Increasing the iq_main File Size

Before you run a distributed query, increase the file size in user dbspace iq_main.

1. In **Folders**, expand ***<localhost>*_iqdemo_mpx > iqdemo > Dbspaces**.
2. In the Dbspaces folder, click **iq_main**.
3. On the **Files** tab, right-click **iq_main**, and choose **Properties**.
4. On the Properties sheet, change the **Modify File Size** to 200.
5. Click **Apply > OK**.

Running a Distributed Query

The query engine is more likely to distribute a query that processes a large number of rows than those that require fewer resources.

This task includes two scripts that load and query a significant amount of sample data (10 million rows). The first script creates and loads a Fact table and a Dimension table. The second script runs a query that fetches 3 rows, and generates a query plan in the coordinator and write server directories.

Note: Depending on your system configuration and amount of data, the Sybase IQ query engine may choose a query plan that does not distribute to the write server.

1. Choose **Tools > Sybase IQ 15 > Open Interactive SQL**.
2. On the **Connect** dialog, click **Advanced**, enter:

Tab Name	Field	Value	
Identifica-tion	**Authentication**	**Database**	
	User ID	`DBA` or `dba` (case-insensitive)	
	Password	`sql` (case-sensitive)	
	Action	**Connect to a running database on this computer**	
	Server name	`<localhost>_iqdemo`	
	Database name	`iqdemo`	
Network	**TCP/IP protocol op-tions**	**Host**	*<localhost>*
		Port	**2638**

3. In Interactive SQL, click **File > Open**.
4. Browse to `$IQDIR15/demo`, and choose `dqpdata.sql`
5. Press **F9** to create and load the Dimension table and Fact table.
6. Click **File > Open**.
7. Browse to `$IQDIR15/demo`, and choose `dqpquery.sql`
8. Press **F9** to run the script that runs the query and generates the query plan.

Reviewing the Query Plan

Check the multiplex node directories to review the query plan. The comprehensive query plan is in the leader node directory, and each distributed work unit has an HTML file generated by the node that received it.

To open the query plan in a Web browser, double-click the leader node directory. The query plan is represented visually by a structure called the query tree.

Figure 1: Sybase IQ Sample Query Plan
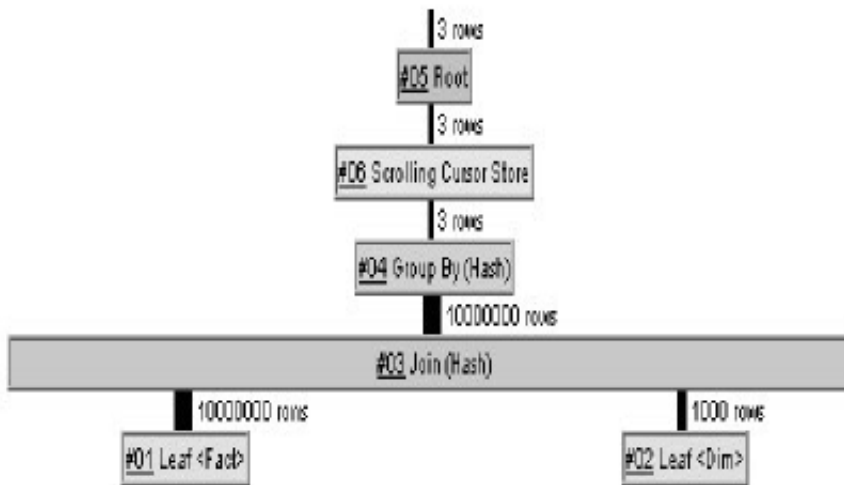
Query Tree

A query tree represents the query's data flow, and the query tree consists of nodes that represent a stage of work. Each node has a name and a data flow operator (DFO) number. All nodes in the tree are hyperlinked to more detailed information.

The lowest nodes on the tree are leaf nodes. Each leaf node represents a table or a prejoin index set in the query. At the top of the plan is the root of the operator tree. Information flows up from the tables and through any operators representing joins, sorts, filters, stores, aggregation, and subqueries.

The three vertical bars that connect the query operators indicate distribution. For example, between the group by node and the join, the three bars indicate that this operation occurs over parallel threads and is distributed over multiple servers. The operation between the join and leaf #01 is also distributed.

Note: For additional information about query plans, see the *Performance and Tuning Guide*.

Index

Index