



Utility Guide

SAP Sybase IQ 16.0 SP04

DOCUMENT ID: DC00168-01-1604-01

LAST REVISED: May 2014

Copyright © 2014 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

Audience	1
CP874toUTF8 Database Administration Utility	3
dbbackup Database Administration Utility	5
dbfhide Database Administration Utility	9
Hiding the Contents of .ini Files Using dbfhide	10
dbinfo Database Administration Utility	11
dbisql Interactive SQL Utility	13
Starting Interactive SQL from the Windows Start Menu	13
Starting Interactive SQL from a Command Prompt	13
Interactive SQL Utility Syntax	13
Interactive SQL Command Line Options	14
Setting Quiet Mode Prerequisites	16
Interactive SQL Connection Parameters	17
Supplying Connection Parameters Using the Connect Assistant	17
Supplying Connection Parameters Manually	18
Connect Window Tools	21
Execute SQL Statements and Command Files	21
Using Text Completion	21
Inserting Comments	23
Look Up Tables, Columns, and Procedures	23
Indent SQL Statements	24
Clearing the SQL Statements Pane	25
Executing All SQL Statements	25
Executing Selected SQL Statements	25
Executing SQL Statements One at a Time	25
Configuring the Execute Statements Toolbar Button	26
Cancelling a SQL Statement	26
Executing Command Files	26

Executing COMMIT and ROLLBACK	
Statements	27
Customizing the Interactive SQL Interface	27
Interactive SQL Window Reference	28
Data Menu Window Reference	29
Tools Menu Window Reference	29
Interactive SQL Keyboard Shortcuts	30
Result Sets	31
Showing Multiple Results Sets (UNIX)	32
Showing Multiple Result Sets (Windows)	32
Edit Result Sets in Interactive SQL	32
Generating SQL Statements from Result Sets	36
Printing SQL Statements and Result Sets	36
Favorites List	36
Adding a .sql File to Favorites	36
Adding SQL Statements to Favorites	37
Adding a Connection to Favorites	37
Showing the Favorites	37
Opening a Favorite	37
Editing Favorites	38
Exporting Favorites	38
Importing Favorites	38
Favorites Menu Window Reference	38
Command Recall	39
Accessing the Command History Window	39
Recalling a Command	39
Copying Commands from the Command History	
Window	39
Saving Commands from the Command History	
Window	40
Removing Commands from the Command	
History Window	40
Logging Commands	40
Build Queries with the Query Editor	40
Creating a Query Using the Query Editor	41

Configuring the Query Editor	42
Query Editor Limitations	42
View Plans Using the Interactive SQL Plan Viewer	42
Starting the Plan Viewer	43
Viewing Graphical Plans in Interactive SQL	43
Configuring the Graphical Plan	44
Printing the Plan	44
Source Control Integration	44
Configuring Interactive SQL to Use Source Control	45
Opening Source Control Projects from Interactive SQL	46
Checking Out Files from Interactive SQL	46
Checking In Files from Interactive SQL	47
Additional Source Control Actions	47
SQL Statements for Interactive SQL	48
CLEAR Statement [Interactive SQL]	48
CONFIGURE Statement [Interactive SQL]	49
CONNECT Statement [ESQL] [Interactive SQL]	50
DISCONNECT Statement [Interactive SQL]	52
OUTPUT Statement [Interactive SQL]	54
PARAMETERS Statement [Interactive SQL]	57
READ Statement [Interactive SQL]	58
SET CONNECTION Statement [ESQL] [Interactive SQL]	61
SET OPTION Statement [Interactive SQL]	62
START DATABASE Statement [Interactive SQL]	63
START ENGINE Statement [Interactive SQL]	65
STOP DATABASE Statement [Interactive SQL]	66
STOP ENGINE Statement [Interactive SQL]	67
SQL Options for Interactive SQL	68
DEFAULT_ISQL_ENCODING Option [Interactive SQL]	68

ON_ERROR Option [Interactive SQL]	69
dblocate Database Administration Utility	71
dblog Database Administration Utility	73
dbping Database Administration Utility	77
dbstop Database Shutdown Utility	81
Stopping the Database Server with dbstop	81
dbstop Options	81
dbtran Database Administration Utility	83
dbvalid Database Administration Utility	89
iqdscp Configuration Utility	93
iqdsn Database Administration Utility	95
Encryption Connection Parameter [ENC]	98
LogFile Connection Parameter [LOG]	99
CommBufferSize Connection Parameter [CBSize]	100
DisableMultiRowFetch Connection Parameter	
[DMRF]	101
LivenessTimeout Connection Parameter [LTO]	101
iqheader Database Administration Utility	103
iqinit Database Administration Utility	107
iqocscfg Configuration Utility	117
isql Interactive SQL Utility	119
iqsqlpp SQL Preprocessor Utility	131
start_iq Database Server Startup Utility	135
Starting the start_iq Utility	135
Listing all start_iq Switches	135
start_iq Server Options	135
@data iqsrv16 database server option	136
@envvar iqsrv16 Server Option	137
@filename iqsrv16 Server Option	137
-c iqsrv16 Server Option	138
-al iqsrv16 Server Option	139
-ca iqsrv16 database server option	139
-cc iqsrv16 database server option	140
-ch iqsrv16 Server Option	140
-cl iqsrv16 database server option	141

-cp iqsrv16 database server option	142
-cr iqsrv16 database server option	143
-cs iqsrv16 database server option	143
-cv iqsrv16 database server option	143
-dt iqsrv16 database server option	144
-ec iqsrv16 database server option	145
-ep iqsrv16 database server option	147
-es iqsrv16 database server option	148
-fips iqsrv16 database server option	149
-ga iqsrv16 database server option	150
-gb iqsrv16 database server option	150
-gc iqsrv16 database server option	151
-gd iqsrv16 database server option	151
-ge iqsrv16 database server option	152
-gf iqsrv16 database server option	153
-gk iqsrv16 database server option	153
-gl iqsrv16 Server Option	153
-gm iqsrv16 database server option	154
-gn iqsrv16 Server Option	155
-gp iqsrv16 database server option	155
-gr iqsrv16 Server Option	156
-gss iqsrv16 Server Option	156
-gt iqsrv16 database server option	156
-gtc iqsrv16 database server option	157
-gu iqsrv16 database server option	158
-iqfrec iqsrv16 Server Option	160
-iqgovern iqsrv16 Server Option	160
-iqIm iqsrv16 Server Option	161
-iqmc iqsrv16 Server Option	161
-iqmpx_failover iqsrv16 Server Option	162
-iqmpx_ov iqsrv16 Server Option	162
-iqmpx_reclaimwriterfreelist iqsrv16 Server Option	162
-iqmpx_sn iqsrv16 Server Option	163
-iqmsgnum iqsrv16 Server Option	163

-iqmsgsz iqsrv16 Server Option	163
-iqmt iqsrv16 Server Option	164
-iqnotemp iqsrv16 Server Option	164
-iqnumbercpus iqsrv16 Server Option	164
-iqpartition iqsrv16 Server Option	165
-iqlvmem start_iq Server Option	165
-iqro iqsrv16 Server Option	166
-iqstart iqsrv16 Server Option	166
-iqtc iqsrv16 Server Option	167
-iqtss iqsrv16 Server Option	167
-iqwmem iqsrv16 Server Option	168
-k iqsrv16 database server option	168
-kl iqsrv16 database server option	168
-kr iqsrv16 database server option (deprecated)	169
-krb iqsrv16 database server option	170
-ks iqsrv16 database server option	171
-ksc iqsrv16 database server option	171
-ksd iqsrv16 database server option	171
-m iqsrv16 database server option	172
-n iqsrv16 database server option	172
-o iqsrv16 database server option	174
-oe iqsrv16 database server option	175
-on iqsrv16 database server option	175
-os iqsrv16 database server option	176
-ot iqsrv16 database server option	177
-p iqsrv16 database server option	177
-pc iqsrv16 database server option	177
-pt iqsrv16 database server option	178
-qi iqsrv16 database server option	179
-qp iqsrv16 database server option	179
-qs iqsrv16 database server option	179
-qw iqsrv16 database server option	180
-s iqsrv16 database server option	180
-sb iqsrv16 database server option	181

-sf iqsrv16 database server option	182
-sk iqsrv16 database server option	188
-su iqsrv16 database server option	189
-tdsl iqsrv16 database server option	190
-ti iqsrv16 database server option	191
-tl iqsrv16 database server option	191
-tmf iqsrv16 database server option	192
-tq iqsrv16 database server option	193
-u iqsrv16 database server option	193
-ud iqsrv16 Server Option	193
-uf iqsrv16 database server option	194
-ut iqsrv16 Server Option	194
-v iqsrv16 Server Option	194
-x iqsrv16 database server option	194
-xd iqsrv16 Server Option	195
-xs iqsrv16 database server option	196
-z iqsrv16 database server option	197
-ze iqsrv16 database server option	198
-zl iqsrv16 database server option	198
-zn iqsrv16 database server option	199
-zo iqsrv16 database server option	200
-zoc iqsrv16 database server option	200
-zr iqsrv16 database server option	201
-zs iqsrv16 database server option	202
start_iq Database File Parameters	203
start_iq Database Options	203
-a iqsrv16 database option	204
-ad iqsrv16 database option	204
-al iqsrv16 Database Option	205
-ar iqsrv16 database option	205
-as iqsrv16 database option	206
-dh iqsrv16 Database Option	207
-ds iqsrv16 database option	207
-ek iqsrv16 Database Option	208
-iqfreq iqsrv16 Database Option	208

-m iqsrv16 Database Option	208
-n iqsrv16 Database Option	209
-r iqsrv16 Database Option	209
-sm iqsrv16 Database Option	210
-sn iqsrv16 Database Option	210
-xp iqsrv16 Database Option	210
stop_iq Database Shutdown Utility	213
When to Stop and Restart the Server	213
Ways to Stop Database Servers	213
Permissions Required to Stop the Server	215
Stopping the Database Server with stop_iq	216
stop_iq Options	217
Stopping Servers in a cron or at Job	217
Servers with Long Paths	218
Appendix: dbisqlc Interactive SQL Classic Utility	
(Deprecated)	219
Interactive SQL Classic Syntax	219
Interactive SQL Classic Options	220
Function and Special Keys (UNIX)	220
Function and Special Keys (Windows)	221
Command Recall Keys (Windows)	222
Index	223

Audience

This guide is for SAP® Sybase® IQ utility program users who require reference material for the utility programs.

Utility programs are commands that you invoke directly from the operating system. Familiarity with relational database systems and introductory user-level experience with SAP Sybase IQ is assumed. Use this book to get information about available syntax, parameters, and options. Other books in the SAP Sybase IQ 16.0 documentation set provide more context on how to carry out particular tasks.

Audience

CP874toUTF8 Database Administration Utility

The **CP874toUTF8** utility converts data in the CP874 character set into a UTF8 collation that is supported by SAP Sybase IQ for the Thai language.

You can also use this utility to load data in the CP874 character set without converting it to UTF8.

Syntax

```
CP874toUTF8 [CP874InputFile]
```

Usage

You can run this utility only from the command prompt.

CP874toUTF8 reads the named file in the CP874 character set (or standard input if no files are named) and prints the UTF8 conversion to standard output.

Note: Files with embedded NULL characters ('\0') are not converted correctly. Remove such characters before running this utility.

Use caution with large data files; the UTF8 output can be up to three times larger than the input data. Input and output file size must both be within operating system limits.

CP874toUTF8 returns a 0 exit status upon successful completion. A nonzero exit status indicates an error occurred. The **CP874toUTF8** utility writes all error messages to stderr.

CP874toUTF8 calls the International Components for Unicode (ICU) library to perform the data conversion. If ICU cannot convert the data, **CP874toUTF8** reports a conversion error. If the conversion fails, divide the file into smaller sections for conversion, to isolate the point of failure. If you cannot resolve the problem, contact Technical Support.

dbbackup Database Administration Utility

The dbbackup utility makes a copy of the transaction log of a running SAP Sybase IQ database.

Note: To back up an entire SAP Sybase IQ database, always use **BACKUP DATABASE** instead of **dbbackup**. **BACKUP DATABASE** backs up all database files, and is the only way to back up the catalog store.

To truncate the transaction log on a simplex or multiplex server, use the **-m** option in your configuration file or **start_iq** command. The server may be running or stopped.

Syntax

```
dbbackup [options] target-directory
```

Parameters

This table lists the available options for the **dbbackup** utility.

Table 1. dbbackup Options

Option	Description
<i>@data</i>	<p>Reads options from the specified environment variable or configuration file. If both exist, and share the same name, SAP Sybase IQ uses the environment variable.</p> <p>To protect passwords or other information in the configuration file, use the File Hiding utility (dbfhide) to obfuscate configuration file contents.</p>
-c " <i>keyword=value; ...</i> "	Supplies database connection parameters. If you do not specify the connection parameters, connection parameters from the SQL CONNECT environment variable are used, if set. The user ID must have the BACKUP DATABASE system privilege.

Option	Description
<p>-l (lowercase L) file</p>	<p>Sends a live backup of the transaction log to a file. Enables a secondary system to be brought up rapidly in the event of server failure. A live backup does not terminate, but continues while the server runs. It runs until the primary server becomes unavailable. At that point, it shuts down, but the backed up log file is intact and can be used to quickly bring up a secondary system.</p> <p>The live backup of the transaction log is always the same length or shorter than the active transaction log. When a live backup is running and another backup restarts the transaction log (dbbackup -x), the live backup automatically truncates the live backup log and restarts the live backup at the beginning of the new transaction log.</p> <p>See <i>Administration: Backup, Restore, and Data Recovery > Data Backup, Recovery, and Archiving > How to Back Up Databases > Types of Backups > Making a Live Backup of a Transaction Log.</i></p>
<p>-o</p>	<p>Writes output messages to the named file.</p>
<p>-q</p>	<p>Quiet mode – does not display output messages. This option is available only when you run this utility from a command prompt.</p>
<p>-r</p>	<p>Copies the old transaction log to a new name and starts a new empty log. The following steps occur:</p> <ul style="list-style-type: none"> • SAP Sybase IQ copies and saves the current working transaction log to the directory specified in the command. • SAP Sybase IQ keeps the current transaction log in its current directory, but renames it using the format <code>yymmddxx.log</code>, where <code>xx</code> are sequential characters starting at <code>AA</code> and running to <code>ZZ</code>, and <code>yymmdd</code> represents the current year, month, and day. This file is then no longer the current transaction log. • SAP Sybase IQ generates a new transaction log file that contains no transactions. The new file has the name of the former current transaction log and becomes the current transaction log.
<p>-t</p>	<p>Backs up only the transaction log file. This can be used as an incremental backup since the transaction log can be applied to the most recently backed up copy of the database files.</p>

Option	Description
-xo filename	Truncate (delete and restart) the transaction log
target-directory	The directory to which the backup files are copied. If the directory does not exist, SAP Sybase IQ creates it. The parent directory must exist.

Usage

The **dbbackup** utility allows you to back up the transaction log while other applications or users are using the database. Backup file names are the same as the database file names.

The **dbbackup** utility truncates the database name to 70 characters and creates a target file with a truncated name. SAP Sybase IQ uses **dbbackup** when synchronizing secondary servers. Due to `dbbackup` restrictions, database names must be less than 70 characters long.

If you have adequate disk space, use **-r** to preserve the existing log file under a new name and start a new empty log. If disk space is limited, use **-xo** instead to truncate the existing log.

Exit codes are 0 (success) or nonzero (failure).

SERVER is the default logical server for **dbbackup**.

Examples

The following Windows command backs up the transaction log from the `iqdemo` database running on the `sample_server` server into the directory `iqbackup`, connecting as user ID `DBA` with password `sql`:

```
dbbackup -c "eng=sample_server;dbn=iqdemo;uid=DBA;pwd=sql" c:\sample\iqbackup
```

See also

- *dbfhide Database Administration Utility* on page 9
- *-m iqsrv16 Database Option* on page 208

dbfhide Database Administration Utility

The **dbfhide** utility is a file hiding utility you use to add simple encryption to configuration files and initialization files to hide the contents of each file.

Syntax

```
dbfhide original-configuration-file encrypted-configuration-file
```

Parameters

This table lists the available options for the **dbfhide** utility.

Option	Description
<i>original-configuration-file</i>	Name of the original file.
<i>encrypted-configuration-file</i>	Name for the new obfuscated file.

Usage

Configuration files are used by some utilities to hold command-line options. These options can contain a password. You can use the **dbfhide** utility to add simple encryption to configuration files and `.ini` files used by SAP Sybase IQ and its utilities, and thereby obfuscate the contents of the file. The original file is not modified. Once you add simple encryption to a file, there is no way to remove it. To change an obfuscated file, keep a copy of the original file that you can modify and obfuscate again.

Examples

Create a configuration file that starts the personal database server and the sample database. The file should set a cache of 10MB, and name this instance of the personal server “*Elora*”. The configuration file would be written as follows:

```
# Configuration file for server Elora -n Elora -c 10M path\iqdemo.db
```

Lines beginning with # are treated as comments.

Name the file `sample.txt`. To start the database using this configuration file, enter:

```
start_iq @sample.txt
```

Add simple encryption to the configuration.

```
dbfhide sample.txt encrypted_sample.txt
```

Use the `encrypted_sample.txt` file to start a database:

```
start_iq @encrypted_sample.txt
```

For more information about using configuration files, see the *Installation and Configuration Guide*.

dbfhide Database Administration Utility

The following command adds simple encryption to the `asaldap.ini` file:

```
dbfhide asaldap.ini encrypted_asaldap.ini
```

See also

- *dbbackup Database Administration Utility* on page 5
- *-m iqsrv16 Database Option* on page 208

Hiding the Contents of .ini Files Using dbfhide

Obfuscate an `.ini` file using the **dbfhide** utility.

1. Save the file with a different name.

```
rename saldap.ini saldap.ini.org
```

If you do not keep a copy of the original file, then you cannot modify the contents of the file once it has been obfuscated.

2. Obfuscate the file with the File Hiding utility, giving the obfuscated file the required file name:

```
dbfhide saldap.ini.org saldap.ini
```

3. Protect the `saldap.ini.org` file using file system or operating system protection, or store the file in a secure location.

To change the `saldap.ini` file, edit the `saldap.ini.org` file and repeat step 2.

Warning! You should not add simple encryption to the `.odbc.ini` system information file with the File Hiding utility (`dbfhide`) on UNIX unless you will only be using SAP Sybase IQ data sources. If you plan to use other data sources, then obfuscating the contents of the `.odbc.ini` file may prevent other drivers from functioning properly.

`dbfhide` does not accept the `@data` parameter to read in options from a configuration file.

dbinfo Database Administration Utility

The **dbinfo** utility displays information about a database catalog store. The information returned by **dbinfo** does not reflect the IQ main store.

Any valid user ID can run **dbinfo**, but to obtain page usage statistics you need the **MANAGE ANY DBSPACE** system privilege.

Syntax

```
dbinfo [ options ]
```

Parameters

This table lists the available options for the **dbinfo** utility.

Table 2. dbinfo Options

Option	Description
-c " <i>keyword=value; ...</i> "	Specify connection parameters. See <i>Connection and Communication Parameters Reference in Administration: Database</i> .
-o <i>filename</i>	Write output messages to the named file.
-q	Operate quietly; do not display output messages.
-u	Output page usage statistics. Display information about the usage and size of all catalog store tables, including system and user-defined tables. You can only request page usage statistics if no other users are connected to the database.

Usage

The **dbinfo** utility indicates when the database was created, the name of any transaction log file or log mirror that is maintained, the catalog store page size, the version of installed Java classes, and other information. Optionally, it can also provide catalog table usage statistics and details.

dbisql Interactive SQL Utility

Interactive SQL (**dbisql**) is a graphical utility included with SAP Sybase IQ that lets you execute SQL statements, build scripts, and display database data.

Note: For backward compatibility, SAP Sybase IQ includes the older Interactive SQL Classic (**dbisqlc**) utility. Use Interactive SQL rather than Interactive SQL Classic. Interactive SQL Classic is deprecated and will be removed in a future release of SAP Sybase IQ.

Starting Interactive SQL from the Windows Start Menu

Start Interactive SQL from the Windows Start menu.

1. Select **Start > Programs > Sybase > Sybase IQ 16.0 > Interactive SQL**.
2. In the **Connect** window, supply parameters.

See also

- *Interactive SQL Connection Parameters* on page 17
- *Interactive SQL Command Line Options* on page 14

Starting Interactive SQL from a Command Prompt

Start Interactive SQL from the command prompt on Windows or UNIX.

1. Do one of the following:
 - In a command shell, enter `dbisql`
 - Select **Start > Run** and enter `dbisql`
2. In the **Connect** window, supply parameters.

See also

- *Interactive SQL Connection Parameters* on page 17
- *Interactive SQL Command Line Options* on page 14

Interactive SQL Utility Syntax

Invoke Interactive SQL from a command prompt.

```
dbisql [ options ] [ dbisql-command | command-file ]
```

Note: Interactive SQL does not accept @filename parameters. Exit codes are 0 (success) or non-zero (failure).

- The following command, entered at a system prompt, runs the command file `mycom.sql` against the current default server, using the user ID `DBA` and the password `sql`. If there is an error in the command file, the process terminates.

```
dbisql -c "uid=DBA;pwd=sql" -onerror exit mycom.sql
```

- The following command, when entered on a single line at a command prompt, adds a user to the current default database:

```
dbisql -c "uid=DBA;pwd=sql" grant connect to joe identified by passwd
```

See also

- Interactive SQL Command Line Options* on page 14

Interactive SQL Command Line Options

Specify options when invoking Interactive SQL from a command prompt.

Table 3. Interactive SQL Options

Option	Description
-c <i>"keyword=value; ..."</i>	<p>Specifies connection parameters. If you do not specify this option, the environment variable <code>SQLCONNECT</code> is used. If Interactive SQL cannot connect, you see a dialog box where you can enter the connection parameters.</p> <hr/> <p>Note: Always specify connection parameters for Interactive SQL instead of relying on defaults. If you start more than one database on a server, for example, specify the database name, and in a network with subnets, specify the communications protocol parameter with host number.</p>
-d <i>delimiter</i>	<p>Specifies a command delimiter. Quotation marks around the delimiter are optional, except when the command shell itself interprets the delimiter in some special way.</p> <p>Command delimiters are used for all connections in that Interactive SQL session, regardless of the setting stored in the database (for the user, or the <code>PUBLIC</code> setting).</p>
-d1	<p>(The final character is the number 1, not a lower case L.) Echoes all statements that Interactive SQL executes to the Command window (STDOUT). This feedback is useful when debugging SQL scripts, or when Interactive SQL is processing a long SQL script</p>

Option	Description
-datasource <i>dsn-name</i>	Specifies an ODBC data source to connect to. You do not need to be using the SQL Anywhere® 16 JDBC driver to use this option. However, if the data source to which you are connecting is not configured to use TCP/IP, you must use the SQL Anywhere 16 JDBC driver to connect. By default, SAP Sybase IQ data sources are configured to use TCP/IP.
-f file name	Opens (but does not run) the file called <i>file name</i> . You must enclose the file name in quotes if the file name contains a blank; otherwise, quotes are optional. If the file does not exist, or if it is a directory instead of a file, Interactive SQL prints an error message to the console and quits. If the file name does not include a full drive and path specification, the file is assumed to be relative to the current directory.
-host <i>host name</i>	Specifies the host name or IP address of the computer on which the database server is running. You can use the name localhost to represent the current machine.
-nogui	<p>Runs Interactive SQL in a command-prompt mode, with no windowed user interface. This is useful for batch operations. If you specify either <i>dbisql-command</i> or <i>command-file</i>, then -nogui is assumed.</p> <p>In -nogui mode, Interactive SQL sets the program exit code to indicate success or failure. On Windows operating systems, the environment variable ERRORLEVEL is set to the program exit code. Exit codes are:</p> <ul style="list-style-type: none"> • 0 — Success. • 1 — General failure. At some point, a SQL or Interactive SQL statement did not execute successfully and the you chose to stop executing SQL statements. Alternatively, Interactive SQL noted an internal error. • 5 — User terminated Interactive SQL. When an error occurs during execution, you are prompted to ignore the error, stop, or exit Interactive SQL. If you opt to exit, the program returns code 5. The program also returns code 5 if an error occurs and the Interactive SQL option ON_ERROR is set to EXIT. • 9 — Unable to connect. • 255 — Bad command. The command line contained incomplete or invalid switches. <hr/> <p>Note: In -nogui mode, any SQL text you enter at the command prompt is executed when you press [Enter], even if you specify a command delimiter using the -d switch. Ensure the SQL command is complete before you press [Enter].</p>

Option	Description
-onerror (continue exit)	Controls what happens if an error is encountered while reading statements from a command file. This option overrides the on_error setting. This option is useful when using Interactive SQL in batch operations.
-port <i>portnumber</i>	Specifies the port number on which the database server is running. The default port number for SAP Sybase IQ is 2638.
-q	Runs in quiet mode—does not display output messages. This option is useful only if you start Interactive SQL with a command or command file. Specifying this option does not suppress error messages.
-version	Displays the version number of Interactive SQL.
-x	Scans commands but does not execute them. You may find this option useful for checking long command files for syntax errors.

See also

- *Interactive SQL Connection Parameters* on page 17
- *Starting Interactive SQL from a Command Prompt* on page 13
- *Starting Interactive SQL from the Windows Start Menu* on page 13
- *Interactive SQL Utility Syntax* on page 13

Setting Quiet Mode Prerequisites

If you use Interactive SQL with the **-q** option (quiet mode), and if the data extraction commands (primarily setting the option TEMP_EXTRACT_NAME1 to an output file) are in a command file, you must first set and make permanent the **Show all result sets** option.

If you do not set this option, the data extraction output file is not created.

1. Start Interactive SQL.
2. In the Interactive SQL window, select **Tools > Options**. The Options window appears.
3. Select **SAP Sybase IQ** in the left pane.
4. Click the **Results** tab.
5. In the Results Processing area, select **Show all results sets**.
6. Click **OK**.

Interactive SQL Connection Parameters

When no database is connected, SAP Sybase IQ displays the Connect window to request connection parameters. Information you enter in the Connect window is not preserved between sessions.

The connection parameters you specify in the Connect window are dependent on the number of databases running on the database server. To connect to a single database, you only need to complete the **User ID** and **Password** fields if the server was started on the local machine with *IQTMP16* environment variable set. If there are multiple databases running on the database server, you must specify additional connection parameters such as the server or database name.

Enter connection parameters in the Connect window using one of these methods:

- Using the Connect Assistant
- Manually

Note: You can bypass the Connect window by supplying the connection parameters on the command line with the **-c** option.

If the Connect window or an error message about missing information appears, you may need to enter the **-host** and **-port** or other missing information in the **Advanced** tab. If your database is on a remote server, enter the **-host** and **-port** parameters on separate lines, as in:

```
-host fiona
-port 1870
```

See also

- *Interactive SQL Command Line Options* on page 14
- *Starting Interactive SQL from a Command Prompt* on page 13
- *Starting Interactive SQL from the Windows Start Menu* on page 13

Supplying Connection Parameters Using the Connect Assistant

The Connect window has a Connect Assistant wizard to help you connect to a database. To display or hide the Connect Assistant, click the arrow in the top right corner of the window.

1. Click the **Next** button in the **Connect Assistant** area of the Connect window.
2. Follow the on-screen prompts.

Note: If you are connecting using an ODBC source, do not use the Connect Assistant. Close the Connect Assistant and click the **Identification** tab.

Supplying Connection Parameters Manually

Connect to a database using the Identification, Database, Network, and Advanced Options tabs on the Connect window.

1. From the **Authentication** list, specify either **Database** or **Windows Integrated Login** authentication.
2. Select an **Action**.

Table 4. Connect with an ODBC Data Source

Field/Button/Option	Details
ODBC data source name	Selects a data source, which is a stored set of connection parameters, for connecting to your database. This field is equivalent to the DSN connection parameter, which references a data source in the registry.
ODBC data source file	Selects a data source file for the connection.

Table 5. Connect to a Running Database on this Computer

Field/Button/Option	Details
Server name	The name of the database server you are connecting to. For remote servers, specify the server as <code>host name:port number</code> .
Database name	The name of the database you are connecting to. You need a database name only if there is more than one database running on the database server. If the database is not already running on the server, you should specify the database file instead. You can also select a recently used database, or click Browse to locate the database file. Note: If you specify both the database name and database file when trying to connect to a database that is already running, the database file is ignored.

Table 6. Connect to a Running Database on Another Computer

Field/Button/Option	Details
Host	The computer name on which the database server is running.

Field/Button/Option	Details
Port	The port the database server is using.
Server name	The name of the database server you are connecting to. For remote servers, specify the server as <code>host name:port number</code> .
Find	Looks for servers.

Table 7. Start and Connect to a Database on this Computer

Field/Button/Option	Details
Database file	Type the full path and name of the database file. Otherwise, the path of the file is relative to the working directory of the database server.
Encryption key	If the database file is encrypted, you must supply an encryption key to the database server every time the database server starts the database.
Database name	The name of the database you are connecting to. You need a database name only if there is more than one database running on the database server. If the database is not already running on the server, you should specify the database file instead. You can also select a recently used database, or click Browse to locate the database file. Note: If you specify both the database name and database file when trying to connect to a database that is already running, the database file is ignored.
Server name	The name of the database server you are connecting to. For remote servers, specify the server as <code>host name:port number</code> .
Start line	Enter a start line to connect to a local database server and to set your own start parameters. You must enter the full path of the database server.
Stop database after last disconnect	Automatically shut down the database after the last user disconnects.

Table 8. Start and Connect to a Database on Another Computer

Field/Button/Option	Details
Database file	Type the full path and name of the database file. Otherwise, the path of the file is relative to the working directory of the database server.
Encryption key	If the database file is encrypted, you must supply an encryption key to the database server every time the database server starts the database.
Database name	The name of the database you are connecting to. You need a database name only if there is more than one database running on the database server. If the database is not already running on the server, you should specify the database file instead. You can also select a recently used database, or click Browse to locate the database file. Note: If you specify both the database name and database file when trying to connect to a database that is already running, the database file is ignored.
Server name	The name of the database server you are connecting to. For remote servers, specify the server as <code>host name:port number</code> .
Host	The computer name on which the database server is running.
Port	The port the database server is using.
Stop database after last disconnect	Automatically shut down the database after the last user disconnects.

Table 9. Connect with a Connection String

Field/Button/Option	Details
Parameters	Connection string.

3. Click **Connect**.

Connect Window Tools

Click **Tools** at the bottom of the Connect window to access options.

Table 10. Connect Window Tools

Tool	Description
Test connection	This tool tests whether the information provided results in a proper connection.
Copy connection string to clipboard	This tool creates a connection string from the options you specified in the Connect window and copies the string into your clipboard.
Save as ODBC Data Source	Saves the new data source to your list of data sources.

Execute SQL Statements and Command Files

After you execute a **SELECT** statement, the result set appears on the **Results** tab in the Results pane. By default, row numbers appear to the left of the result set.

Using Text Completion

The text completion option supplies you with complete keywords and object names as you type.

For example, typing the letter S causes a text completion window to appear that includes a list of possible keywords or object names beginning with the letter S. You can configure the text completion settings from the Options window.

1. Type the first letter of a database object name in the **SQL Statements** pane.
2. By default, the text completion window opens automatically as you type. Alternatively, you can press **Ctrl+Space** or **Ctrl+Shift+Space** to open the window.
3. If you do not see the object name you want, press **Ctrl+A** to view a complete list of database objects (based on the filtering options you set—by default, all database objects appear in the list).
4. Select the object name from the list and then press **Enter**.

Text Completion Keyboard Shortcuts

Keyboard shortcuts are available when the text completion list is open.

Key	Description
Ctrl+A	Shows a context-free list of matches.

Key	Description
Ctrl+Double quote (") (or Ctrl+Shift+')	Completes the name, enclosing it in quotation marks, regardless of the setting of the quoted_identifier option.
Ctrl+Asterisk (*) (or Ctrl+Shift+8)	For tables, inserts a comma-separated list of columns including data types. For stored procedures, inserts the procedure name, followed by a comma-separated list of parameter names and their data types.
Ctrl+C	Changes the contents of the text completion list to show or hide columns.
Ctrl+F	Changes the contents of the text completion list to show or hide SQL functions.
Ctrl+P	Changes the contents of the text completion list to show or hide stored procedures and functions.
Ctrl+Plus Sign (+) (or Ctrl+Shift+=)	For tables, inserts a comma-separated list of columns. For stored procedures, inserts the procedure name, followed by a comma-separated list of parameter names.
Ctrl+S	Changes the contents of the text completion list to show or hide system objects.
Ctrl+Shift+Space or Ctrl+Space	Opens the text completion window. You can also use Ctrl+Space to open the text completion window.
Ctrl+T	Changes the contents of the text completion list to show or hide tables.
Ctrl+V	Changes the contents of the text completion list to show or hide views.
Esc	Closes the text completion window without adding any text.
Tab	Accepts the selection and closes the text completion window.

See also

- *Interactive SQL Keyboard Shortcuts* on page 30

Inserting Comments

Use comments to attach explanatory text to SQL statements or statement blocks.

You can turn text into a comment. The database server does not execute comments. Interactive SQL supports the following types of comments:

- -- (double hyphen)
- // (double slash)
- /* ... */ (slash-asterisk)

1. Select the text in the SQL Statements pane.
2. Press **Ctrl+Minus Sign (-)** to add double hyphen comment indicators or **Ctrl+Forward Slash (/)** to add double slash comment indicators.

If no text is selected, the comment indicator is added to the beginning of the current line.

Look Up Tables, Columns, and Procedures

While you are entering commands in Interactive SQL, you can look up the names of tables, columns, or procedures stored in the current database and insert them at your cursor position.

Looking Up the Names of Tables in the Database

Look up the names of tables stored in the database and insert them at your cursor position.

1. Select **Tools > Lookup Table Name** or press F7.
2. Find and select the table.
3. Click **OK** to insert the table name into the SQL Statements pane at the current cursor position.

Looking Up the Names of Columns in the Database

Look up the names of columns stored in the database and insert them at your cursor position.

1. Select **Tools > Lookup Table Name** or press F7.
2. Find and select the table containing the column.

Note: In the Lookup Table Name window, you can enter the first few characters of the table you are looking for. The list is narrowed to include only those items that start with the text you entered.

3. Click **Show Columns**.
4. Select the column and click **OK** to insert the column name into the SQL Statements pane at the current cursor position.

Looking Up the Names of Procedures in the Database

Look up the names of procedures stored in the database and insert them at your cursor position.

1. Select **Tools > Lookup Procedure Name** or press F8.
2. Find and select the procedure.

Note: In the Lookup Procedure Name window, you can enter the first few characters of the procedure you are looking for. The list is narrowed to include only those items that start with the text you entered.

3. Click **OK** to insert the procedure name into the SQL Statements pane at the current cursor position.

Indent SQL Statements

Increase or decrease indentation of SQL statements, or change the default number of spaces indented.

Adding or Increasing Indentation of SQL Statements

Add indentation to an SQL statement or increase the indentation of an already-indented SQL statement.

1. Select the text in the SQL Statements pane that you want to indent. If no text is selected, the indentation is applied to the current line.
2. Press **Ctrl+Shift+Period**.

Removing or Decreasing Indentation of SQL Statements

Remove indentation from an SQL statement or decrease the indentation of an already-indented SQL statement.

1. Select the text in the SQL Statements pane that you want to decrease the indentation. If no text is selected, the indentation is applied to the current line.
2. Press **Ctrl+Shift+Comma**.

Changing the Number of Spaces that are Indented

Change the default indent size.

1. Select **Tools > Options**.
2. Choose **Editor** and then click the **Tabs** tab.
3. Type a new number in the **Indent Size** field.

Clearing the SQL Statements Pane

Clear the SQL Statements pane when you no longer want to work with the SQL code displayed on the pane.

By default, the **Esc** key leaves the contents of the **SQL Statements** pane unchanged. Configure the **Esc** key to clear the SQL Statements pane and close any opened results sets.

1. Select **Tools > Options > Compatibility**.
2. Select **Pressing The Esc Key Clears SQL Statements And Closes Result Sets**.
3. Press **Esc**.

Executing All SQL Statements

If multiple SQL statements exist in the SQL statements pane, you can execute all statements at once.

1. Type your query in the SQL Statements pane.
2. Press **F5**, or choose **SQL Execute** to execute the statement.

See also

- *Executing Selected SQL Statements* on page 25
- *Executing SQL Statements One at a Time* on page 25
- *Configuring the Execute Statements Toolbar Button* on page 26

Executing Selected SQL Statements

If multiple SQL statements exist in the SQL statements pane, you can execute one or more selected statements.

1. Type your queries in the SQL Statements pane and select the query.
2. Press **F9**, or choose **SQL Execute Selection** to execute the statement.

See also

- *Executing All SQL Statements* on page 25
- *Executing SQL Statements One at a Time* on page 25
- *Configuring the Execute Statements Toolbar Button* on page 26

Executing SQL Statements One at a Time

To execute SQL statements individually, use the **Single Step** option. This is useful when debugging.

1. Type your query in the SQL Statements pane.

2. Place your cursor in the statement that you want to execute.
3. Select **SQL > Single Step** or press **Shift+F9**.
The selected SQL statement executes and the next SQL statement is selected.
4. Press **Shift+F9**.
The selected SQL statement executes.
5. Repeat the previous step until there are no more selected statements to execute.

See also

- *Executing All SQL Statements* on page 25
- *Executing Selected SQL Statements* on page 25
- *Configuring the Execute Statements Toolbar Button* on page 26

Configuring the Execute Statements Toolbar Button

Configure the **Execute Statements** button to either execute all SQL statements or only execute selected statements.

1. Select **Tools > Options > Toolbar**.
2. Perform one of these actions:
 - Select **Execute All Statement(s)**. This is the default setting.
 - Select **Execute Selected Statement(s)**.

See also

- *Executing All SQL Statements* on page 25
- *Executing Selected SQL Statements* on page 25
- *Executing SQL Statements One at a Time* on page 25

Cancelling a SQL Statement

A cancel operation stops the current processing and prompts for the next command. The **Interrupt The SQL Statement** button on the Interactive SQL toolbar cancels a command. Click **Interrupt The SQL Statement** on the Interactive SQL toolbar.

See also

- *ON_ERROR Option [Interactive SQL]* on page 69

Executing Command Files

Command files are text files that contain SQL statements and are useful if you want to run the same SQL statements repeatedly. You can use Interactive SQL to open, view, run, and save command files.

Execute the command file using one of these methods:

- Use the Interactive SQL **READ** statement to execute command files.
For example:

```
READ temp.sql;
```
- Load a command file into the SQL Statements pane and execute it directly from there.
- Load command files into the SQL Statements pane by choosing **File > Open**. Enter the file name when prompted.
- Run a command file without loading it by choosing **File > Run Script**.
- Supply a command file as a command line argument for Interactive SQL.

Executing COMMIT and ROLLBACK Statements

There are multiple ways to execute **COMMIT** and **ROLLBACK** statements.

You can execute a **COMMIT** statement by:

- Choosing **SQL > Commit**.
- Pressing Ctrl+Shift+C.
- Typing `Commit` into the **SQL Statements** pane, and then executing the statement.

You can execute a **ROLLBACK** statement by:

- Choosing **SQL > Rollback**.
- Pressing Ctrl+Shift+R.
- Typing `Rollback` into the **SQL Statements** pane, and then executing the statement.

Note: Executing a **COMMIT** or **ROLLBACK** via the **SQL** menu or a keyboard shortcut does not modify the contents of the **SQL Statements** pane; however, the **Results** tab in the **Results** pane is cleared.

Customizing the Interactive SQL Interface

You can configure settings for the tabs and panes in Interactive SQL using the Options window.

1. In Interactive SQL, choose **Tools > Options**.
2. In the left pane, click an option and specify the options that you want. You can change how results are displayed, specify whether or not to execute a commit after every statement (or only on exit or disconnect), and specify whether or not to check for updates.
3. Click **OK**.

Interactive SQL Window Reference

Use the Interactive SQL window to enter SQL statements and view results and messages.

Table 11. Interactive SQL Window Details

Pane	Column/Tab	Description
SQL Statements		Provides a place for you to type SQL statements to access and modify your data.
	Line Number Column	<p>A column on the left that shows line numbers. These line numbers allow you to do the following:</p> <ul style="list-style-type: none"> • Click a line number to select a line. Alternatively place your cursor in the line, and press Ctrl+comma(,). • Click and drag to select multiple lines. • Double-click a line to select the entire SQL statement that corresponds to the line. Alternatively place your cursor in the statement, and press Ctrl+period(.).
Results		The Results pane has two tabs: Results and Messages . The tabs appear at the bottom of the Results pane.
	Results Tab	The Results tab displays the results of commands that you execute. For example, if you use SQL statements to search for specific data in the database, the Results tab displays the columns and rows that match the search criteria in the pane above. You can edit the result set on the Results tab.
	Messages Tab	The Messages tab displays messages from the database server about the SQL statements that you execute in Interactive SQL.

Results of graphical plans for SAP Sybase IQ databases are displayed in separate Plan Viewer window(s).

See also

- *View Plans Using the Interactive SQL Plan Viewer* on page 42

Data Menu Window Reference

Use the windows available from the **Data** Menu to export a result set and import data.

Table 12. Data Menu Windows

Window	Description
Export	Opens the Export Wizard, which allows you to export a result set.
Import	Opens the Import Wizard, which allows you to import data from a file or database.

Tools Menu Window Reference

Use the windows available from the **Tools** menu to configure Interactive SQL settings, search for table and procedure names to insert into your queries, and edit your queries.

Table 13. Tools Menu Windows

Window	Description
Lookup Table Name	The Lookup Table Name window lets you browse table and column names and insert them into the SQL Statements pane.
Lookup Procedure Name	The Lookup Procedure Name window lets you browse procedure names and insert them into the SQL Statements pane.
Edit Query	The Query Editor provides a graphical way to create and edit SELECT statements in Interactive SQL.
Plan Viewer	The Plan Viewer is a graphical tool for viewing graphical plans for SAP Sybase IQ databases.
Options	The Options window sets options for commands, appearance, importing and exporting data, and messages in Interactive SQL.

See also

- *View Plans Using the Interactive SQL Plan Viewer* on page 42
- *Creating a Query Using the Query Editor* on page 41

Interactive SQL Keyboard Shortcuts

Learn about the keyboard shortcuts available in Interactive SQL.

Table 14. Interactive SQL Keyboard Shortcuts

Function key	Description
Tab	Accepts the selection and closes the text completion window.
Alt+F4	Exits Interactive SQL.
Ctrl+A	Shows a content-free list of matches.
Ctrl+C	Copies the selected rows and column headings to the clipboard in the Results pane. In the SQL Statements pane, copies the selected text to the clipboard.
Ctrl+End	Moves to the bottom of the current pane.
Ctrl+H	Displays the history of your executed SQL statements.
Ctrl+Home	Moves to the top of the current pane.
Ctrl+N	Clears the contents of the Interactive SQL window.
Ctrl+P	Prints the contents of the SQL Statements pane.
Ctrl+Q	Displays the Query Editor. The Query Editor helps you build SQL queries. When you have finished building your query, click OK to export it back into the SQL Statements pane.
Ctrl+S	Saves the contents of the SQL Statements pane.
Ctrl+Plus Sign (+)	Adds the item with its parameter list to the SQL Statements pane.
Ctrl+Asterisk	Adds the item with its parameter and type list to the SQL Statements pane.
Ctrl+Double quote ("")	Adds the item enclosed in quotation marks to the SQL Statements pane.
F2	Edits the selected value in the result set.
F5	Executes all text in the SQL Statements pane.
F7	Displays the Lookup Table Name dialog.
F8	Displays the Lookup Procedure Name dialog.
F9	Executes the text that is selected in the SQL Statements pane. If no text is selected, all of the statements are executed.

Function key	Description
Pgdn	Moves a page down in the current pane.
Pgup	Moves a page up in the current pane.

These keyboard shortcuts are available when the **SQL Statements** pane has the focus:

Table 15. Interactive SQL Keyboard Shortcuts for SQL Statements Pane

Function key	Description
Ctrl+]]	Moves the cursor to the matching parenthesis, braces, brackets, and angle brackets.
Ctrl+Backspace	Deletes the word to the left of the cursor.
Ctrl+Del	Deletes the word to the right of the cursor.
Ctrl+G	Opens the Go To dialog where you can specify the line you want to go to.
Ctrl+L	Deletes the current line from the SQL Statements pane and puts the line onto the clipboard.
Ctrl+Shift+]]	Extends the selection to the matching brace. Brace matching matches parentheses, braces, brackets, and angle brackets.
Ctrl+Shift+C	Executes a COMMIT statement.
Ctrl+Shift+L	Deletes the current line.
Ctrl+Shift+R	Executes a ROLLBACK statement.
Ctrl+Shift+U	Changes the selection to uppercase characters.
Ctrl+U	Changes the selection to lowercase characters.
F3	Finds the next occurrence of the selected text.
Home	Moves the cursor to the start of the current line or to the first word on the current line.
Shift+F3	Finds the previous occurrence of the selected text.
Shift+Home	Extends the selection to the start of the text on the current line.

Result Sets

Work with the results of commands that you execute.

Showing Multiple Results Sets (UNIX)

By default, Interactive SQL shows the first result set of the most-recently executed statement. Use this procedure to see all result sets on UNIX platforms.

1. If running Interactive SQL as a command line program (-nogui mode):
 - a) Navigate to your \$HOME directory and locate the file `.isqlPreferences11`.
 - b) Change the line`<entry>`
 `key="SybaseIQ.showMultipleResultSets">0</entry>` to `<entry`
 `key="SybaseIQ.showMultipleResultSets">1</entry>`
 - c) Save your changes to `.isqlPreferences11`.
2. If using GUI mode:
 - a) Click **Tools > Options** in the Interactive SQL window.
 The Interactive SQL Options window appears.
 - b) On the **SAP Sybase IQ** page, select the **Results** tab.
 - c) Choose **Show All Result Sets**.
 - d) Click **OK**.

Showing Multiple Result Sets (Windows)

By default, Interactive SQL shows the first result set of the most-recently executed statement. Use this procedure to see all result sets on Windows platforms.

1. Click **Tools > Options** in the Interactive SQL window.
 The Interactive SQL Options window appears.
2. On the **SAP Sybase IQ** page, select the **Results** tab.
3. Choose **Show all Result Sets**.
4. Click **OK**.

Edit Result Sets in Interactive SQL

Once you execute a query in Interactive SQL, you can sort and edit the result set to modify the database. You can also select rows from the result set and copy them for use in other applications.

Interactive SQL supports editing, inserting, and deleting rows. Editing the result set has the same effect as executing **UPDATE**, **INSERT**, and **DELETE** statements. After editing a result set, the equivalent **INSERT**, **UPDATE**, and **DELETE** statements are added to Interactive SQL's command history.

To edit a row or value in the result set, you must have the proper privileges on the table or column you want to modify values from. For example, if you want to delete a row, then you must have **DELETE** privilege for the table the row belongs to.

You cannot edit a result set if you:

- Select columns from a table with a primary key, but do not select all the primary key columns
- Attempt to edit the result set of a **JOIN** (for example, if there is data from more than one table in the result set).
- Attempt to edit a table that has its editing disabled

Editing the result set may fail if you:

- Attempt to edit a row or column you do not have privilege on.
- Enter an invalid value (for example, a string in a numeric column or a NULL in a column that does not allow NULLs).

When editing fails, an Interactive SQL error message appears explaining the error, and the database table values remain unchanged.

See also

- *Recalling a Command* on page 39

Editing Table Values from the Interactive SQL Result Set

You can change any or all of the values within existing rows in database tables, provided that you have UPDATE permission on the columns being modified. In addition table editing must not be disabled.

1. Execute a query in Interactive SQL.
2. On the **Results** tab, click the value you want to change.
3. Right-click the value and choose **Edit Row**, or press F2 to edit the result set.

A blinking cursor appears in the table cell containing the value.

4. Enter the new value. If you want to change other values in the row, press **Tab** or **Shift+Tab** to move to the other values.
5. Press **Enter** to update the database once you are done editing values in the row.

You can press the **Esc** key to cancel the change that was made to the selected value.

6. Execute a **COMMIT** statement to make your changes to the table permanent.

See also

- *Disabling Table Editing* on page 33

Disabling Table Editing

You can disable table editing via the Options window in Interactive SQL.

1. From the **Tools** menu, choose **Options**, and then choose **SAP Sybase IQ**.
2. Ensure that **Scrollable Table** is selected and select **Disable Editing**.
3. Click **OK**.

4. Execute a query. You must execute a new query for the changes to table editing to take effect.

See also

- *Editing Table Values from the Interactive SQL Result Set* on page 33

Insert Rows into the Database from the Interactive SQL Result Set

Interactive SQL allows you to add new rows to a table. You tab between columns in the result set to add values to the row. You must have INSERT permission on the table to add new rows.

Inserting a New Row into the Result Set

Add a new blank row to the table from the result set.

1. Right-click the result set and choose **Add Row**. A new blank row appears with a blinking cursor in the first value in the row.
2. Enter the new value and then press **Tab** to move to the next column.

You cannot enter invalid data types into a column. For example, you cannot enter a string into a column that accepts the INT data type. Repeat this step until all the column values are added.

3. Press **Enter** to update the database.

Inserting Values into Columns with Default Values

When adding a value in a column that has a default value, the cell editor contains a list with a (DEFAULT) item. Similarly, if a column accepts NULL values, (NULL) appears in the list. Select **(DEFAULT)** if you want to insert the default value.

If a column cannot be NULL and does not have a default value, you must enter a value.

When the Result Set Contains a Computed Column

If the result set contains a computed column and you do not specify a value for the computed column, the value is calculated when the database is updated. However, if you specify a value for the computed column, the database is updated with the specified value, and a value is not calculated for the computed column.

Inserting New Rows Using the INPUT Statement

An alternative to inserting new rows from the result set in Interactive SQL is to add rows using the **INPUT** statement with the **PROMPT** clause.

1. From Interactive SQL, add a row using the **INPUT** statement with the **PROMPT** clause. For example:

```
INPUT INTO Products PROMPT;
```

2. Specify the value for each column when prompted.

Deleting Rows from the Database Using Interactive SQL

You can also delete rows from a database table in Interactive SQL. You must have DELETE privilege on the table to delete rows.

1. Select the row(s) you want to delete using one of these methods:
 - a) Press and hold the **Shift** key while clicking the row(s).
 - b) Press and hold the **Shift** key while using the Up or Down Arrow.
2. Press **Delete**.
3. Execute a **COMMIT** to make the change permanent.

Copying Rows from an Interactive SQL Result Set

You can copy rows directly from the result set in Interactive SQL and then paste them into other applications.

1. Right-click any cell in the row and choose **Copy Data > Rows**. To copy multiple rows, hold the **Ctrl** key while clicking cells in the rows, and then right-click and choose **Copy Data > Rows**.
2. Paste the row(s) into other applications.

Copying Columns from the Interactive SQL Result Set

You can copy columns directly from the result set in Interactive SQL and then paste them into other applications.

1. Right-click any cell in the column and choose **Copy Data > Columns**. To copy multiple columns, hold the **Ctrl** key while clicking cells in the columns, and then right-click and choose **Copy Data > Columns**.
2. Paste the column(s) into other applications.

Copying Cells from the Interactive SQL Result Set

Copying an individual value copies only the data to the clipboard. No column headings are copied and no quoting is done.

1. Right-click the cell you want to copy and choose **Copy Data > Cells**. To copy multiple cells, hold the **Ctrl** key while clicking cells, and then right-click and choose **Copy Data > Cells**.
2. Paste the cell(s) into other applications.

Sorting Columns in an Interactive SQL Result Set

Sort the results by a specified column.

1. Click a column-header in the **Results** tab.
2. Fetch the remaining results if prompted.

Generating SQL Statements from Result Sets

You can create **INSERT**, **DELETE**, and **UPDATE** statements for selected rows in the result set.

1. Select the row(s) you want to generate a statement for.
2. Right-click the selection, choose **Generate**, and then choose **INSERT Statement**, **DELETE Statement**, or **UPDATE Statement**. The statement is copied to the clipboard.

Printing SQL Statements and Result Sets

You can print the contents of the SQL Statements pane or query results.

1. Type your query in the SQL Statements pane and execute the query, if desired.
2. Press **Ctrl+P** or select **File > Print**.
3. When prompted, choose to print the SQL statements or results.

Adding a Header

You can add a header to printed SQL statements or printed query results.

1. Select **Tools > Options**. The Interactive SQL Options window appears.
2. On the Editor page, click the **Print** tab.
3. In the **Header** field, specify the text that you want to appear in the header. You can also click the right arrow and choose items to include in the header.

Favorites List

You can add the current database connection and open SQL file to your Favorites list.

Adding a .sql File to Favorites

Store frequently-used SQL command files in a favorites list. The favorites list is specific to a single user and cannot be seen by other users.

1. Open the SQL command file that you want to add to your favorites.
2. Select **Favorites > Add to Favorites**
3. Select **Add the open file 'filename'**. In the **Name** field, type a name for the **.sql** file.
4. Click **OK**.

See also

- *Favorites Menu Window Reference* on page 38

Adding SQL Statements to Favorites

Store frequently-used SQL statements in a favorites list. The favorites list is specific to a single user and cannot be seen by other users.

1. Type the SQL commands that you want to add to your favorites in the **SQL Statements** pane. The contents in the SQL Statements pane cannot exceed 16384 characters.
2. From the **Favorites** menu, choose **Add To Favorites**.
3. Select **Add SQL Statements**. In the **Name** field, type a name for the favorite.
4. Click **OK**.

Adding a Connection to Favorites

Store frequently-used connections in a favorites list. The favorites list is specific to a single user and cannot be seen by other users.

1. Connect to a database.
2. Select **Favorites > Add To Favorites**
3. Select **Save The Connection Password**. In the Name field, type a name for the connection.
4. Click **OK**.

See also

- *Favorites Menu Window Reference* on page 38

Showing the Favorites

Show the list of frequently-used SQL command files and connections.

Select **Favorites > Show Favorites**.

The Favorites pane appears on the left side of the Interactive SQL window.

See also

- *Favorites Menu Window Reference* on page 38

Opening a Favorite

Open an SQL command file or connections from your list of favorites.

From the **Favorites** menu, choose the favorite you want to open.

See also

- *Favorites Menu Window Reference* on page 38

Editing Favorites

Edit the favorites list.

1. From the **Favorites** menu, choose **Show Favorites**.
2. Select a favorite, right-click, and then choose **Edit**.
3. Follow the instructions in the window.
4. Click **Save**.

Exporting Favorites

Export all of your favorites to a `.fav` file.

1. Choose **Favorites > Export Favorites**.
2. Specify a file name for the `.fav` file and then click **Export**.

Importing Favorites

Import favorites on to another computer or save the file as a backup.

1. Choose **Favorites > Import Favorites**.
2. Browse to the `filename.fav` file and then click **Import**.

Favorites Menu Window Reference

Use the windows available from the Favorites menu to manage SQL files and connection information as favorites.

Table 16. Favorite Menu Windows

Window	Description
Add to Favorites	Allows you to save SQL files and connection information as favorites.
Organize Favorites	Allows you to maintain and organize your favorites.
Show Favorites	Opens the Favorites window on the left side of the Interactive SQL window.

See also

- *Adding a .sql File to Favorites* on page 36
- *Adding a Connection to Favorites* on page 37
- *Showing the Favorites* on page 37
- *Opening a Favorite* on page 37

Command Recall

When you execute a command, Interactive SQL automatically saves it in a history list that persists between Interactive SQL sessions. You can also log your executed commands to a log file.

Interactive SQL maintains a record of up to 50 of the most recent commands.

You can view the entire list of commands in the Command History window. The most recent commands appear at the bottom of the list.

Accessing the Command History Window

Open the Command History window to view the record of up to 50 of the most recent commands.

Use one of the following methods:

- Press **Ctrl+H**.
- Click the **Open A List Of Past SQL Statements** button on the toolbar.

Recalling a Command

Recall commands you executed from the command history.

1. Open the Command History window and select the command.
2. Click **OK**.

Note: You can also recall commands without the Command History window. Use the **Recall Previous SQL Statement** and **Recall Next SQL Statement** icons in the toolbar to scroll back and forward through your commands, or press **Alt+Right Arrow** and **Alt+Left Arrow**, respectively.

See also

- *Edit Result Sets in Interactive SQL* on page 32

Copying Commands from the Command History Window

Copy commands from the Command History Window to the SQL Statements pane of Interactive SQL.

1. Open the Command History window.
2. Select the command or commands, and then press **Ctrl+C** or click **Copy**.
3. Click **OK**.

The selected statements copy to the SQL Statements pane.

Saving Commands from the Command History Window

Save the command history as a .sql file.

1. Open the Command History window.
2. Click the **Save History As .SQL File** button or press **Ctrl+S**.
3. In the Save As window, specify a location and name for the file. The command history file has a **.sql** extension.
4. Click **Save** when finished.

Removing Commands from the Command History Window

Clear a single command, multiple commands, or all commands from the command history.

1. Open the Command History window.
2. Remove the commands using one of these methods:
 - Select one or more commands and click the **Delete** button. This action cannot be undone.
 - Click **Clear History** to remove all the commands from the window. This action cannot be undone.

Logging Commands

You can record commands as you execute them. The recorded commands are stored in a log file so you can use the commands again.

1. Select **SQL > Start Logging**.
2. In the Save As window, specify a location and name for the log file. For example, name the file `mylogs.sql`.
3. Click **Save** when finished.
4. Select **SQL > Stop Logging**.

Build Queries with the Query Editor

The Query Editor is a tool in Interactive SQL that helps you build **SELECT** statements.

You can create SQL queries in the Query Editor, or you can import queries and edit them. When you have finished your query, click **OK** to export it back into Interactive SQL for processing.

You do not need to use SQL code to create queries with the Query Editor. However, you can use SQL with the Query Editor in the following ways:

- You can create a query in the SQL Statements pane in Interactive SQL and import it into the Query Editor by highlighting the code before you open the editor.
- At any time while using the Query Editor, you can click SQL at the bottom of the window to see the SQL code for the query you are building. You can directly edit the code, and the fields are automatically updated in the Query Editor.

Creating a Query Using the Query Editor

Use the Query Editor tools and tabs to write an SQL query.

1. Connect to a database from Interactive SQL.
2. Select **Tools > Edit Query**.

If you have SQL code selected in Interactive SQL, the selected code is automatically imported into the Query Editor

3. Create your query.
4. Click **OK** to write the query to the SQL Statements pane.
5. Use the tabs that guide you through the components of a SQL query:

Table 17. Query Editor Tabs

Tab	Description
Tables tab	Specifies the tables in your query.
Joins tab	Specifies a join strategy for combining the data in the tables. If you include more than one table in your query, you should specify a join strategy for combining the data in the tables. If you do not specify a join strategy for tables you added in the Tables tab, the Query Editor suggests one; if there is a foreign key relationship between the tables, it generates a join condition based on that relationship, or it suggests a cross product. When you open queries, the Query Editor accepts exactly the join strategy that you specified (and an unspecified JOIN is not defaulted to KEY JOIN).
Columns tab	Specifies the columns in your result set. If you do not specify columns, all columns appear..
INTO tab	Assign results to variables.
WHERE tab	Specifies conditions for restricting the rows in your result set.
GROUP BY tab	Group rows in the result set.

Tab	Description
HAVING tab	Restricts the rows in your result set based on group values.
ORDER BY tab	Sorts the rows.

6. Use the following tools:

Table 18. Query Editor Tools

Tool	Description
Expression Editor	Use the Expression Editor to build search conditions or define computed columns.
Derived Table	Use this window, which is nearly identical to the main Query Editor, to create derived tables and subqueries.

See also

- *View Plans Using the Interactive SQL Plan Viewer* on page 42
- *Tools Menu Window Reference* on page 29

Configuring the Query Editor

Configure the Query Editor to customize settings.

1. Select **Tools > Options > SAP Sybase IQ**
2. Click the **Query Editor** tab and change settings.

Query Editor Limitations

Learn about capabilities and syntax not supported by the Query Editor.

The Query Editor builds **SELECT** statements. It is not designed to create views, although you can create them in Interactive SQL and reference them in the Query Editor. Nor was it designed to create **UPDATE** statements or other non-**SELECT** SQL statements. It creates a single **SELECT** statement, so it does not build unions or intersects of **SELECT** statements. In addition, the Query Editor does not support Transact-SQL[®] syntax.

View Plans Using the Interactive SQL Plan Viewer

The Plan Viewer is a graphical tool for viewing graphical plans for databases.

The Plan Viewer window is divided into panes:

Table 19. Plan Viewer Panes

Pane	Description
SQL pane	Provides a place for you to type SQL statements that you want to generate plans for.
Results pane	Shows the graphical plan.
Details pane	Provides text details about the plan

See also

- *Interactive SQL Window Reference* on page 28
- *Tools Menu Window Reference* on page 29
- *Creating a Query Using the Query Editor* on page 41

Starting the Plan Viewer

Starting the Plan Viewer opens it in a separate window.

1. Open Interactive SQL.
2. Choose **Tools > Plan Viewer** (or press **Shift+F5**).

See also

- *Configuring the Graphical Plan* on page 44
- *Printing the Plan* on page 44
- *Viewing Graphical Plans in Interactive SQL* on page 43

Viewing Graphical Plans in Interactive SQL

You can view the query optimizer's execution plan for an SQL statement in the Plan Viewer window in Interactive SQL.

1. Type your query in the SQL Statements pane.
2. Select **Tools > Plan Viewer** or press **Shift+F5**.

The Plan Viewer appears in a separate window. Your specified query appears in the SQL pane.

3. Click **Get Plan** to generate a plan for the specified query.
4. Choose **Tools > Plan Viewer**.
5. Click **Open**.
6. Select a plan file (.saplan), and then click **Open**.

See also

- *Configuring the Graphical Plan* on page 44

- *Printing the Plan* on page 44
- *Starting the Plan Viewer* on page 43

Configuring the Graphical Plan

After executing the graphical plan, you can customize the appearance of items in the plan.

1. Right-click the plan in the lower left pane of the Plan Viewer and choose **Customize**.
2. Change the settings.
3. Click **OK** when finished.
4. Click **Get Plan** to generate the graphical plan with your changes.

See also

- *Printing the Plan* on page 44
- *Starting the Plan Viewer* on page 43
- *Viewing Graphical Plans in Interactive SQL* on page 43

Printing the Plan

You can print a plan in the Plan Viewer.

1. Select **Tools > Plan Viewer**.
2. Press the **Print** button or right-click the plan and select **Print**.

See also

- *Configuring the Graphical Plan* on page 44
- *Starting the Plan Viewer* on page 43
- *Viewing Graphical Plans in Interactive SQL* on page 43

Source Control Integration

Interactive SQL integrates with third-party source control systems, allowing you to perform many common source control operations on files from within Interactive SQL.

On Windows, Interactive SQL integrates with most source control products that support the Microsoft Common Source Code Control API (SCC), including Microsoft Visual SourceSafe. To use source control products that do not support the SCC API on Windows and other operating systems, specify a command line to run for each of the source control actions. Output from those commands appears in a log window.

Interactive SQL supports the following tasks (as long as the task is supported in the source control product):

- Open a source control project

- Get
- Check in
- Check out
- Undo check out
- Compare versions
- Show file history
- Show file properties
- Run the source control manager

If the underlying source control program does not support an action, its corresponding menu item is disabled. For example, Visual SourceSafe supports all of these actions, but using a custom (command line) source control system does not support opening a source control project, or running a source control manager.

You should be familiar with the operations of your source control program before attempting to use it from Interactive SQL.

Configuring Interactive SQL to Use Source Control

You must configure Interactive SQL to use source control before you can perform source control actions on files, such as checking files in and out, comparing different versions of a file, and viewing the history for a file.

If you are running Interactive SQL on a Windows computer that has a source control product that supports the Microsoft SCC API, you can use that product or use a custom (command line oriented) system.

Configuring Interactive SQL Source Control on Windows

Enable source control integration if you are running Interactive SQL on a Windows computer that has a source control product that supports the Microsoft SCC API.

1. Select **Tools > Options**.
2. In the left pane, click **Source Control**.
3. Click **Enable Source Control Integration**.
4. Click **OK**.

Configuring Interactive SQL Source Control Systems with a Command Line Interface

Enable source control integration if you are running Interactive SQL on a computer that has a custom command line oriented source control product.

1. Click **Tools > Options**.
2. In the left pane, click **Source Control**.
3. Click **Enable Source Control Integration**.

4. Click **Configure**.
5. In the Custom Source Control Options window, click **Reset**.
6. Select your source control system from the list, and then click **OK**.
7. Edit the commands in the list as necessary by selecting an action from the Source Control Actions list, and then typing the corresponding command in the Command Line pane.

When you are defining commands for your system in the Source Control Actions list, use the placeholder [FILENAME] to represent the name of the file that is used when you run the command.

If you do not specify a command line for an action, the item in the **File > Source Control** menu is disabled.

8. Click **OK**, and then click **OK** again.

Opening Source Control Projects from Interactive SQL

Some source control products require you to open a source control project before you can perform any other source control actions.

The exact definition of what a project is depends on the source control system you are using. Typically, it is a set of files that are under source control, along with a location on your local file system where working copies of the files are placed. You usually have to provide some credentials, such as a user ID and password, to the source control system to open a project.

If your source control system supports opening a source control project, the **File > Source Control > Open Source Control Project** menu item is enabled. Choosing this option from the **File** menu opens a source control-specific window for opening a project. Once you open a project, you do not have to open it again, even in subsequent Interactive SQL sessions. The project is opened automatically for you.

See also

- *Additional Source Control Actions* on page 47
- *Checking Out Files from Interactive SQL* on page 46
- *Checking In Files from Interactive SQL* on page 47

Checking Out Files from Interactive SQL

Once you open a file in Interactive SQL, use the command on the **File** menu to check the file out.

1. Choose **File > Open** and then browse to the file you want to open.

The file status (Checked In, Checked Out, or Not Controlled) appears on the status bar at the bottom of the Interactive SQL window.

2. Check out the file by choosing **File > Source Control > Check Out**.
3. Depending on which source control product you are using, you may be prompted for a comment or other options as part of the check out procedure.

Note: If you are using a SCC-compliant source control system, the status is always accurate. However, if you use the custom source control system, the status is based on whether the file is read-only or not. A read-only file is assumed to be checked in, but no assumptions are made about editable files because they could be either checked out or not controlled.

See also

- *Opening Source Control Projects from Interactive SQL* on page 46
- *Additional Source Control Actions* on page 47

Checking In Files from Interactive SQL

When you are finished making edits to your file, you can check it back in from Interactive SQL.

1. Choose **File > Source Control > Check In**.
2. Enter check in comments if you are prompted.

See also

- *Opening Source Control Projects from Interactive SQL* on page 46
- *Additional Source Control Actions* on page 47

Additional Source Control Actions

In addition to opening source control projects, and checking files in and out, Interactive SQL supports several other source control actions. The availability of these actions depends on the source control system you are using.

Access these actions from the **File > Source Control** menu in Interactive SQL.

Table 20. Additional Source Control Actions

Action	Description
Get	Gets the latest copy of the file you currently have open in the SQL Statements pane.
Undo Check Out	Discards your working copy of the file, and then downloads the copy of the file that is in the source control archive.
Compare Versions	This action compares the working copy of the file you have opened against the version in the source control archive.

Action	Description
History	Displays a list of source control actions (typically check-ins) that have been made to the file you have open.
Properties	Displays a list of source control properties that are associated with the file you have opened.
Run Source Control Manager	Launches the management program for your source control system. For example, if you are using Microsoft Visual SourceSafe, this launches Visual SourceSafe Explorer.

See also

- *Opening Source Control Projects from Interactive SQL* on page 46
- *Checking Out Files from Interactive SQL* on page 46
- *Checking In Files from Interactive SQL* on page 47

SQL Statements for Interactive SQL

Interactive SQL statements can only be used from within Interactive SQL.

CLEAR Statement [Interactive SQL]

Closes any open result sets in Interactive SQL (**dbisql**).

Quick Links:

Go to Usage on page 48

Go to Standards on page 49

Go to Permissions on page 49

Syntax

CLEAR

Usage

(back to top) on page 48

Closes any open result sets and leaves the contents of the SQL Statements pane unchanged

Side effects:

The **CLEAR** statement closes the cursor associated with the data being cleared.

Standards

(back to top) on page 48

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- SAP Sybase Database product—Not applicable.

Permissions

(back to top) on page 48

None

CONFIGURE Statement [Interactive SQL]

Activates the Interactive SQL (**dbisql**) configuration window.

Quick Links:

Go to Usage on page 49

Go to Standards on page 49

Go to Permissions on page 49

Syntax

CONFIGURE

Usage

(back to top) on page 49

The **dbisql** configuration window displays the current settings of all **dbisql** options. It does not display or let you modify database options.

If you select Permanent, the options are written to the `SYSOPTION` table in the database and the database server performs an automatic **COMMIT**. If you do not choose Permanent, and instead click OK, options are set temporarily and remain in effect for the current database connection only.

Standards

(back to top) on page 49

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- SAP Sybase Database product—Not supported by SAP Adaptive Server® Enterprise.

Permissions

(back to top) on page 49

None

CONNECT Statement [ESQL] [Interactive SQL]

Establishes a connection to the database identified by *database-name* running on the server identified by *engine-name*.

Quick Links:

Go to Parameters on page 50

Go to Examples on page 51

Go to Usage on page 51

Go to Standards on page 52

Go to Permissions on page 52

Syntax

Syntax 1

CONNECT

```
... [ TO engine-name ]  
...[ DATABASE database-name ]  
...[ AS connection-name ]  
...[ USER ] userid [ IDENTIFIED BY ]
```

Syntax 2

CONNECT USING *connect-string*

Parameters

(back to top) on page 50

- **engine-name** – identifier, string, or host-variable
- **database-name** – identifier, string, or host-variable
- **connection-name** – identifier, string, or host-variable
- **userid** – identifier, string, or host-variable
- **password** – identifier, string, or host-variable
- **connect-string** – a list of parameter settings of the form *keyword=value*, and must be enclosed in single quotes.
- **AS** – connection can optionally be named by specifying the clause. This allows multiple connections to the same database, or multiple connections to the same or different database servers, all simultaneously. Each connection has its own associated transaction.

You might even get locking conflicts between your transactions if, for example, you try to modify the same record in the same database from two different connections.

Examples

(back to top) on page 50

- **Example 1 – CONNECT usage within Embedded SQL:**

```
EXEC SQL CONNECT AS :conn_name
USER :userid IDENTIFIED BY :password;
EXEC SQL CONNECT USER "dba" IDENTIFIED BY "sql";
```

- **Example 2 – CONNECT usage from dbisql:**

Connect to a database from **dbisql**. Prompts display for user ID and password:

```
CONNECT
```

Connect to the default database as DBA, from **dbisql**. A password prompt displays:

```
CONNECT USER "DBA"
```

Connect to the demo database as the DBA, from **dbisql**, where *<machine>_iqdemo* is the engine name:

```
CONNECT
TO <machine>_iqdemo
USER "DBA"
IDENTIFIED BY sql
```

Connect to the demo database using a connect string, from **dbisql**:

```
CONNECT
USING 'UID=DBA;PWD=sql;DBN=iqdemo'
```

Usage

(back to top) on page 50

- **Embedded SQL behavior** – in Embedded SQL, if no *engine-name* is specified, the default local database server is assumed (the first database server started). If no *database-name* is specified, the first database on the given server is assumed.

The **WHENEVER** statement, **SET SQLCA**, and some **DECLARE** statements do not generate code and thus might appear before the **CONNECT** statement in the source file. Otherwise, no statements are allowed until a successful **CONNECT** statement has been executed.

The user ID and password are used for permission checks on all dynamic SQL statements. By default, the password is case-sensitive; the user ID is not. You can connect without a password by using a host variable for the password and setting the value of the host variable to be the null pointer.

- **dbisql behavior** – if no database or server is specified in the **CONNECT** statement, **dbisql** remains connected to the current database, rather than to the default server and database. If

dbisql Interactive SQL Utility

a database name is specified without a server name, **dbisql** attempts to connect to the specified database on the current server. You must specify the database name defined in the `-n` database switch, not the database file name. If a server name is specified without a database name, **dbisql** connects to the default database on the specified server. For example, if this batch is executed while connected to a database, the two tables are created in the same database.

```
CREATE TABLE t1( c1 int );  
CONNECT DBA IDENTIFIED BY sql;  
CREATE TABLE t2 (c1 int );
```

No other database statements are allowed until a successful **CONNECT** statement has been executed.

The user ID and password are used for checking the permissions on SQL statements. If the password or the user ID and password are not specified, the user is prompted to type the missing information. By default, the password is case-sensitive; the user ID is not.

Multiple connections are managed through the concept of a current connection. After a successful connect statement, the new connection becomes the current one. To switch to a different connection, use **SET CONNECTION**. Executing a **CONNECT** statement does not close the existing connection (if any). Use **DISCONNECT** to drop connections.

Static SQL statements use the user ID and password specified with the `-l` option on the **SQLPP** statement line. If no `-l` option is given, then the user ID and password of the **CONNECT** statement are used for static SQL statements also.

Standards

(back to top) on page 50

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- SAP Sybase Database product—Open Client Embedded SQL supports a different syntax for the **CONNECT** statement.

Permissions

(back to top) on page 50

None

DISCONNECT Statement [Interactive SQL]

Drops a connection with the database.

Quick Links:

Go to Parameters on page 53

Go to Examples on page 53

Go to Usage on page 53

Go to Standards on page 53

Go to Permissions on page 53

Syntax

```
DISCONNECT [ { connection-name | CURRENT | ALL } ]
```

Parameters

(back to top) on page 52

- **connection-name** – identifier, string, or host-variable
- **ALL** – drops all of the connections of the application to all database environments.
- **CURRENT** – (default) Drops the current connection.

Examples

(back to top) on page 52

- **Example 1** – how to use **DISCONNECT** in Embedded SQL:

```
EXEC SQL DISCONNECT :conn_name
```

- **Example 2** – how to use **DISCONNECT** from **dbisql** to disconnect all connections:

```
DISCONNECT ALL
```

Usage

(back to top) on page 52

The **DISCONNECT** statement drops a connection with the database server and releases all resources used by it. If the connection to be dropped was named on the **CONNECT** statement, then the name can be specified.

An implicit **ROLLBACK** is executed on connections that are dropped.

Standards

(back to top) on page 52

- SQL—ISO/ANSI SQL compliant.
- SAP Sybase Database product—Supported by Open Client/Open Server.

Permissions

(back to top) on page 52

None

OUTPUT Statement [Interactive SQL]

Writes the information retrieved by the current query to a file.

Quick Links:

Go to Parameters on page 54

Go to Examples on page 56

Go to Usage on page 57

Go to Standards on page 57

Go to Permissions on page 57

Syntax

```
OUTPUT TO filename
[ APPEND ] [ VERBOSE ]
[ FORMAT output-format ]
[ ESCAPE CHARACTER character ]
[ DELIMITED BY string ]
[ QUOTE string [ ALL ] ]
[ COLUMN WIDTHS ( integer, ... ) ]
[ HEXADECIMAL { ON | OFF | ASIS } ]
[ ENCODING encoding ]
[ WITH COLUMN NAMES ]
```

output-format

TEXT | FIXED | HTML | SQL | XML

Parameters

(back to top) on page 54

- **FORMAT** – the output format. If no FORMAT clause is specified, the Interactive SQL **OUTPUT_FORMAT** database option setting is used.
- **TEXT** – output is a TEXT format file with one row per line in the file. All values are separated by commas, and strings are enclosed in apostrophes (single quotes). The delimiter and quote strings can be changed using the DELIMITED BY and QUOTE clauses. If the ALL clause is specified in the QUOTE clause, all values (not just strings) are quoted. TEXT is the default output format.

Three other special sequences are also used. The two characters \n represent a newline character, \\ represents a single \, and the sequence \xDD represents the character with hexadecimal code DD.

If you are exporting Java methods that have string return values, you must use the HEXADECIMAL OFF clause.

- **FIXED** – output is fixed format with each column having a fixed width. The width for each column can be specified using the **COLUMN WIDTHS** clause. No column headings are output in this format.

If the **COLUMN WIDTHS** clause is omitted, the width for each column is computed from the data type for the column, and is large enough to hold any value of that data type. The exception is that **LONG VARCHAR** and **LONG BINARY** data defaults to 32KB.

- **HTML** – output is in the Hyper Text Markup Language format.
- **SQL** – output is an Interactive SQL **INPUT** statement required to recreate the information in the table.

Note: SAP Sybase IQ does not support the **INPUT** statement. You need to edit this statement to a valid **LOAD TABLE** (or **INSERT**) statement to use it to load data back in.

- **XML** – output is an XML file encoded in UTF-8 and containing an embedded DTD. Binary values are encoded in CDATA blocks with the binary data rendered as 2-hex-digit strings. The **LOAD TABLE** statement does not accept XML as a file format.
- **APPEND** – appends the results of the query to the end of an existing output file without overwriting the previous contents of the file. If the **APPEND** clause is not used, the **OUTPUT** statement overwrites the contents of the output file by default.

The **APPEND** clause is valid if the output format is **TEXT**, **FIXED**, or **SQL**.

- **VERBOSE** – error messages about the query, the SQL statement used to select the data, and the data itself are written to the output file. If the **VERBOSE** clause is omitted (the default), only the data is written to the file. The **VERBOSE** clause is valid if the output format is **TEXT**, **FIXED**, or **SQL**.
- **ESCAPE CHARACTER** – the default escape character for characters stored as hexadecimal codes and symbols is a backslash (\), so \x0A is the line feed character, for example.

This default can be changed using the **ESCAPE CHARACTER** clause. For example, to use the exclamation mark as the escape character, enter:

```
... ESCAPE CHARACTER '!'
```

- **DELIMITED BY** – for the **TEXT** output format only. The delimiter string is placed between columns (default comma).
- **QUOTE** – for the **TEXT** output format only. The quote string is placed around string values. The default is a single quote character. If **ALL** is specified in the **QUOTE** clause, the quote string is placed around all values, not just around strings.
- **COLUMN WIDTHS** – specifies the column widths for the **FIXED** format output.
- **HEXADECIMAL** – specifies how binary data is to be unloaded for the **TEXT** format only. When set to **ON**, binary data is unloaded in the format 0xabcd. When set to **OFF**, binary data is escaped when unloaded (\xab\xcd). When set to **ASIS**, values are written as

is, that is, without any escaping—even if the value contains control characters. ASIS is useful for text that contains formatting characters such as tabs or carriage returns.

- **ENCODING** – specifies the encoding that is used to write the file. The ENCODING clause can be used only with the TEXT format. Can be a string or identifier.

If the ENCODING clause is not specified, Interactive SQL determines the code page that is used to write the file as follows, where code page values occurring earlier in the list take precedence over those occurring later:

- The code page specified with the **DEFAULT_ISQL_ENCODING** option (if this option is set)
- The default code page for the computer Interactive SQL is running on

Examples

(back to top) on page 54

- **Example 1** – places the contents of the `Employees` table in a text file:

```
SELECT * FROM Employees;  
OUTPUT TO employees.txt FORMAT TEXT
```

- **Example 2** – places the contents of the `Employees` table at the end of an existing file, and include any messages about the query in this file as well:

```
SELECT * FROM Employees;  
OUTPUT TO employees.txt APPEND VERBOSE
```

- **Example 3** – exports a value that contains an embedded line feed character. A line feed character has the numeric value 10, which you can represent as the string `'\x0a'` in an SQL statement.

Execute this statement with **HEXADECIMAL ON**:

```
SELECT 'line1\x0aline2'; OUTPUT TO file.txt HEXADECIMAL ON
```

The result is a file with one line in it, containing this text:

```
line10x0aline2
```

Execute the same statement with **HEXADECIMAL OFF**:

```
line1\x0aline2
```

If you set **HEXADECIMAL** to **ASIS**, you get a file with two lines:

```
'line1  
line2'
```

Using **ASIS** generates two lines, because the embedded line feed character has been exported without being converted to a two-digit hex representation, and without a prefix.

Usage

(back to top) on page 54

The current query is the **SELECT** or **LOAD TABLE** statement that generated the information that appears on the Results tab in the Results pane. The **OUTPUT** statement reports an error if there is no current query.

Note: **OUTPUT** is especially useful in making the results of a query or report available to another application, but is not recommended for bulk operations. For high-volume data movement, use the ASCII and BINARY data extraction functionality with the **SELECT** statement. The extraction functionality provides much better performance for large-scale data movement, and creates an output file you can use for loads.

Side Effects

- In Interactive SQL, the Results tab displays only the results of the current query. All previous query results are replaced with the current query results.

Standards

(back to top) on page 54

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- SAP Sybase Database product—Not applicable.

Permissions

(back to top) on page 54

None

PARAMETERS Statement [Interactive SQL]

Specifies parameters to an Interactive SQL (**dbisql**) command file.

Quick Links:

Go to Examples on page 57

Go to Usage on page 58

Go to Standards on page 58

Go to Permissions on page 58

Syntax

```
PARAMETERS parameter1, parameter2, ...
```

Examples

(back to top) on page 57

dbisql Interactive SQL Utility

- **Example 1** – this **dbisql** command file takes two parameters:

```
PARAMETERS department_id, file ;
SELECT Surname
FROM Employees
WHERE DepartmentID = {department_id}
>#{file}.dat;
```

Usage

(back to top) on page 57

PARAMETERS specifies how many parameters there are to a command file and also names those parameters so that they can be referenced later in the command file.

Parameters are referenced by putting the named parameter into the command file where you want the parameter to be substituted:

```
{parameter1}
```

There must be no spaces between the braces and the parameter name.

If a command file is invoked with fewer than the required number of parameters, **dbisql** prompts for values of the missing parameters.

Standards

(back to top) on page 57

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- SAP Sybase Database product—Not applicable.

Permissions

(back to top) on page 57

None

READ Statement [Interactive SQL]

Reads Interactive SQL (**dbisql**) statements from a file.

Quick Links:

Go to Parameters on page 59

Go to Examples on page 60

Go to Standards on page 60

Go to Permissions on page 61

Syntax

```
READ [ ENCODING encoding ] filename [ parameter ] ...
```

Parameters

(*back to top*) on page 58

- **ENCODING** – an identifier or string, which lets you specify the encoding that is used to read the file.

The **READ** statement does not process escape characters when it reads a file. It assumes that the entire file is in the specified encoding. When running Interactive SQL, the encoding that is used to read the data is determined in the following order:

1. The encoding specified by the ENCODING clause (if this clause is specified).
2. The encoding specified by the byte order mark (BOM) in the file (if a BOM is specified).
3. The encoding specified with the `default_isql_encoding` option (if this option is set).
4. The default encoding for the platform you are running on. On English Windows computers, the default encoding is 1252.

- **filename** – If *filename* has no file extension, Interactive SQL searches for the same file name with the extension `.sql`.

If *filename* does not contain an absolute path, Interactive SQL searches for the file. The location of filename is determined based on the location of the READ statement, as follows:

- If the READ statement is executed directly in Interactive SQL, Interactive SQL first attempts to resolve the path to filename relative to the directory in which Interactive SQL is running. If unsuccessful, Interactive SQL looks for filename in the directories specified in the environment variable `SQLPATH`, and then the directories specified in the environment variable `PATH`.
- If the READ statements reside in an external file (for example, a `.sql` file), Interactive SQL first attempts to resolve the path to filename relative to the location of the external file. If unsuccessful, Interactive SQL looks for filename in a path relative to the directory in which Interactive SQL is running. If still unsuccessful, Interactive SQL looks in the directories specified in the environment variable `SQLPATH`, and then the directories specified in the environment variable `PATH`.
- **parameters** – can be listed after the name of the SQL script file. These parameters correspond to the parameters named in the `PARAMETERS` statement at the beginning of the statement file.

Parameter names must be enclosed in square brackets. Interactive SQL substitutes the corresponding parameter wherever the source file contains `{ parameter-name }`, where *parameter-name* is the name of the appropriate parameter.

The parameters passed to a script file can be identifiers, numbers, quoted identifiers, or strings. When quotes are used around a parameter, the quotes are put into the text during

dbisql Interactive SQL Utility

the substitution. Parameters that are not identifiers, numbers, or strings (contain spaces or tabs) must be enclosed in square brackets ([]). This allows for arbitrary textual substitution in the script file.

If not enough parameters are passed to the script file, Interactive SQL prompts for values for the missing parameters.

When executing a reload.sql file with Interactive SQL, you must specify the encryption key as a parameter. If you do not provide the key in the READ statement, Interactive SQL prompts for the key.

Examples

(back to top) on page 58

- **Example 1** – this example reads from the fictitious file `status.rpt` and `birthday.sql` and passes the parameter values to the variables within the file:

```
READ status.rpt '160'
```

```
READ birthday.sql [>= '1988-1-1'] [<= '1988-1-30']
```

- **Example 2** – this example uses the `PARAMETERS` clause to pass parameters to a script file:

```
[test1.sql]
PARAMETERS par1, par2;

BEGIN
DECLARE v_par1 int;
DECLARE v_par2 varchar(200)
```

```
SET v_par1 = {par1};
SET v_par2 = {par2};
```

```
MESSAGE STRING('PAR1 Value: ', v_par1 ) TO CLIENT;
MESSAGE STRING('PAR2 Value: ', v_par2 ) TO CLIENT;
```

```
END;
```

```
(USR1)> READ test1.sql 123 '041028'
PAR1 Value: 123
PAR2 Value: 041028
```

Note: The second parameter value 041028 must be enclosed in quotes, as `v_par2` is declared as a character data type.

Standards

(back to top) on page 58

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- SAP Sybase Database product—Not applicable.

Permissions

(*back to top*) on page 58

None

SET CONNECTION Statement [ESQL] [Interactive SQL]

Changes the active database connection.

Quick Links:

Go to Parameters on page 61

Go to Examples on page 61

Go to Usage on page 61

Go to Standards on page 62

Go to Permissions on page 62

Syntax

```
SET CONNECTION [connection-name]
```

Parameters

(*back to top*) on page 61

- **connection-name** – identifier, string, or host-variable

Examples

(*back to top*) on page 61

- **Example 1** – in Embedded SQL:

```
EXEC SQL SET CONNECTION :conn_name
```

- **Example 2** – from **dbisql**, set the current connection to the connection named “conn1”:

```
SET CONNECTION conn1
```

Usage

(*back to top*) on page 61

The current connection state is saved and is resumed when it again becomes the active connection. If *connection-name* is omitted and there is a connection that was not named, that connection becomes the active connection.

Note: When cursors are opened in Embedded SQL, they are associated with the current connection. When the connection is changed, the cursor names are not accessible. The cursors

dbisql Interactive SQL Utility

remain active and in position and become accessible when the associated connection becomes active again.

Standards

(back to top) on page 61

- SQL—**dbisql** use is a vendor extension to ISO/ANSI SQL grammar. Embedded SQL is a full-level feature.
- SAP Sybase Database product—Supported by Open Client/Open Server.

Permissions

(back to top) on page 61

None

SET OPTION Statement [Interactive SQL]

Changes Interactive SQL (**dbisql**) options.

Quick Links:

Go to Parameters on page 62

Go to Usage on page 63

Go to Permissions on page 63

Syntax

Syntax 1

```
SET [ TEMPORARY ] OPTION  
... [ userid. | PUBLIC.]option-name = [ option-value ]
```

Syntax 2

```
SET PERMANENT
```

Syntax 3

```
SET
```

Parameters

(back to top) on page 62

- **userid** – identifier, string, or host-variable
- **option-name** – identifier, string, or host-variable
- **option-value** – host-variable (indicator allowed), string, identifier, or number

Usage

(back to top) on page 62

SET PERMANENT (Syntax 2) stores all current **dbisql** options in the `SYSOPTION` system table. These settings are automatically established every time **dbisql** is started for the current user ID.

Syntax 3 is used to display all of the current option settings. If there are temporary options set for **dbisql** or the database server, these display; otherwise, permanent option settings are displayed.

If you incorrectly type the name of an option when you are setting the option, the incorrect name is saved in the `SYSOPTION` table. You can remove the incorrectly typed name from the `SYSOPTION` table by setting the option `PUBLIC` with an equality after the option name and no value:

```
SET OPTION PUBLIC.a_mistyped_name=;
```

Permissions

(back to top) on page 62

No specific system privileges are required to set your own options.

The `SET ANY PUBLIC OPTION` system privilege is required to set database options for another user.

The `SET ANY SYSTEM OPTION` system privilege is required to set a `SYSTEM` option for the `PUBLIC` user ID.

The `SET ANY SECURITY OPTION` system privilege is required to set a `SECURITY` option for the `PUBLIC` user ID.

START DATABASE Statement [Interactive SQL]

Starts a database on the specified database server.

Quick Links:

Go to Parameters on page 64

Go to Examples on page 64

Go to Usage on page 64

Go to Standards on page 64

Go to Permissions on page 65

Syntax

```
START DATABASE database-file  
... [ AS database-name ]
```

```
... [ ON engine-name ]  
... [ AUTOSTOP { YES | NO } ]  
... [ KEY key ]
```

Parameters

(back to top) on page 63

- **database-name** – if not specified, a default name is assigned to the database. This default name is the root of the database file. For example, a database in file `c:\sybase\16_0\demo\iqdemo.db` is given the default name `iqdemo`.
- **engine-name** – if not specified, the default database server is assumed. The default database server is the first started server among those currently running.
- **AUTOSTOP** – when set to YES (default), the database is unloaded when the last connection to it is dropped. When set to NO, the database is not unloaded.
- **KEY** – to enter the KEY value (password) for strongly encrypted databases.

Note: Start only one database on a given SAP Sybase IQ database server.

Examples

(back to top) on page 63

- **Example 1** – on a UNIX system, start the database file `/s1/sybase/sample_2.db` on the current server:

```
START DATABASE '/s1/sybase/sample_2.db'
```

- **Example 2** – on a Windows system, start the database file `c:\sybase\sample_2.db` as `sam2` on the server `eng1`:

```
START DATABASE 'c:\sybase\sample_2.db'  
AS sam2  
ON eng1
```

Usage

(back to top) on page 63

The database server must be running. The full path must be specified for the database file unless the file is located in the current directory.

The **START DATABASE** statement does not connect **dbisql** to the specified database: a **CONNECT** statement must be issued to make a connection.

Standards

(back to top) on page 63

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- SAP Sybase Database product—Not applicable.

Permissions

(back to top) on page 63

Requires the SERVER OPERATOR system privilege.

START ENGINE Statement [Interactive SQL]

Starts a database server.

Quick Links:

Go to Parameters on page 65

Go to Examples on page 65

Go to Usage on page 66

Go to Standards on page 66

Go to Permissions on page 66

Syntax

START ENGINE

```
AS engine-name [ STARTLINE command-string ]
```

Parameters

(back to top) on page 65

- **STARTLINE** – to specify a set of options for the server.
- **command-string** – valid command strings are those that conform to the database server command line description. See *start_iq Database Server Startup Utility* in the *Utility Guide*.

Examples

(back to top) on page 65

- **Example 1** – starts a database server named eng1 without starting any databases on it:

```
START ENGINE AS eng1
```

- **Example 2** – starts the same server with a cache of 8096KB:

```
START ENGINE AS eng1 STARTLINE 'start_iq -c 8096'
```

Usage

(back to top) on page 65

Several server options are required for SAP Sybase IQ to operate well. To ensure that you are using the right set of options, start your server by using either SAP Control Center or a configuration file with the **start_iq** command.

Standards

(back to top) on page 65

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- SAP Sybase Database product—Not applicable.

Permissions

(back to top) on page 65

None

STOP DATABASE Statement [Interactive SQL]

Stops a database on the specified database server.

Quick Links:

Go to Parameters on page 66

Go to Examples on page 67

Go to Standards on page 67

Go to Permissions on page 67

Syntax

```
STOP DATABASE database-name  
... [ ON engine-name ]  
... [ UNCONDITIONALLY ]
```

Parameters

(back to top) on page 66

- **database-name** – the name specified in the -n parameter when the database is started, or specified in the DBN (DatabaseName) connection parameter. This name is typically the file name of the database file that holds the catalog store, without the .db extension, but can be any user-defined name.
- **engine-name** – if not specified, all running engines are searched for a database of the specified name.

- **UNCONDITIONALLY** – if specified, the database is stopped, even if there are connections to the database. If not specified, the database is not stopped if there are connections to it.

Examples

(back to top) on page 66

- **Example 1** – stop the database named `sample` on the default server:

```
STOP DATABASE sample
```

Standards

(back to top) on page 66

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- SAP Sybase Database product—Not applicable.

Permissions

(back to top) on page 66

Requires the SERVER OPERATOR system privilege.

STOP ENGINE Statement [Interactive SQL]

Stops a database server.

Quick Links:

Go to Parameters on page 67

Go to Examples on page 68

Go to Standards on page 68

Go to Permissions on page 68

Syntax

```
STOP ENGINE engine-name [ UNCONDITIONALLY ]
```

Parameters

(back to top) on page 67

- **UNCONDITIONALLY** – if specified, the database server is stopped, even if there are connections to the server. If not specified, the database server is not stopped if there are connections to it.

Examples

(back to top) on page 67

- **Example 1** – stop the database server named `sample`:

```
STOP ENGINE sample
```

Standards

(back to top) on page 67

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- SAP Sybase Database product—Not applicable.

Permissions

(back to top) on page 67

None

SQL Options for Interactive SQL

Interactive SQL options are a subset of the SAP Sybase IQ SQL options. These change how Interactive SQL interacts with the database.

DEFAULT_ISQL_ENCODING Option [Interactive SQL]

Specifies the code page used by **READ** and **OUTPUT** statements.

Allowed Values

identifier or *string*

Default

Use system code page (empty string)

Scope

Can only be set as a temporary option, for the duration of the current connection.

Remarks

DEFAULT_ISQL_ENCODING is used to specify the code page to use when reading or writing files. It cannot be set permanently. The default code page is the default code page for the platform you are running on. On English Windows machines, the default code page is 1252.

Interactive SQL determines the code page that is used for a particular **OUTPUT** or **READ** statement as follows, where code page values occurring earlier in the list take precedence over those occurring later in the list:

- The code page specified in the **ENCODING** clause of the **OUTPUT** or **READ** statement
- The code page specified with the `DEFAULT_ISQL_ENCODING` option (if this option is set)
- The default code page for the computer on which Interactive SQL is running

Example

Set the encoding to UTF-16 (for reading Unicode files):

```
SET TEMPORARY OPTION DEFAULT_ISQL_ENCODING = 'UTF-16'
```

ON_ERROR Option [Interactive SQL]

Controls the action taken if an error is encountered while executing statements in Interactive SQL.

Allowed Values

Value	Description
STOP	Interactive SQL stops executing statements from the file and returns to the statement window for input.
PROMPT	Interactive SQL prompts the user to see if he or she wants to continue.
CONTINUE	Errors display in the Messages pane and Interactive SQL continues executing statements.
EXIT	Interactive SQL terminates.
NOTIFY_CONTINUE	The error is reported, and the user is prompted to press Enter or click OK to continue.
NOTIFY_STOP	The error is reported, and the user is prompted to press Enter or click OK to stop executing statements.
NOTIFY_EXIT	The error is reported, and the user is prompted to press Enter or click OK to terminate Interactive SQL.

Default

PROMPT

Remarks

Controls the action taken, if an error is encountered while executing statements. When you are executing a .SQL file, the values STOP and EXIT are equivalent.

See also

- *Canceling a SQL Statement* on page 26

dblocate Database Administration Utility

The **dblocate** utility is a server location utility that assists in diagnosing connection problems by locating databases on the immediate TCP/IP network.

Syntax

```
dblocate [ options ]
```

Parameters

This table lists the options available for the **dblocate** utility.

Table 21. dblocate Options

Option	Description
-d	Lists the server name and address, for each server found, followed by a comma-separated list of databases running on that server. If the list exceeds 160 characters, it is truncated and ends with an ellipsis (...).
-dn database-name	Lists the server name and address, for servers running a database with the specified name. If the list exceeds 160 characters, it is truncated and ends with an ellipsis (...).
-dv	Displays the server name and address, for each server found, listing each database running on that server on a separate line. The list is not truncated, so this option can be used to reveal lists that are truncated when the v option is used.
-n	Lists IP addresses in the output, rather than computer names. This may improve performance since looking up computer names may be slow.
-o filename	Writes output messages to the named file.
-p portnumber	Displays the server name and address only for servers using the specified TCP/IP port number. The TCP/IP port number must be between 1 and 65535.
-q	Runs in quiet mode—messages are not displayed.
-s name	Displays the server name and address only for servers with the specified server name. If this option is used, the -ss option should not be used (if both options are used, it is likely that no matching servers will be found).

dblocate Database Administration Utility

Option	Description
-ss <i>substr</i>	Displays the server name and address only for servers that contain the specified substring anywhere in the server name. If this option is used, the -s option should not be used (if both options are used, it is likely that no matching servers will be found).
-v	Displays the full server name. By default, dblocate truncates database server names that are longer than 40 bytes.

Usage

dblocate locates any SQL Anywhere or SAP Sybase IQ database servers running over TCP/IP on the immediate network. It prints a list of database servers and their addresses.

Depending on your network, the utility may take several seconds before printing its results.

See also

- *start_iq Database Options* on page 203

dblog Database Administration Utility

The **dblog** utility is a transaction log utility that displays or changes the name of the transaction log or transaction log mirror associated with your database.

You can also use **dblog** to stop a database from maintaining a transaction log mirror, or start maintaining a transaction log mirror.

SAP Sybase IQ automatically handles the creation and deletion of the transaction log for a database. The database must run with a transaction log. The server will not start without a transaction log. A transaction log mirror is a duplicate copy of a transaction log, maintained in tandem by the database.

Note: Although not required, use a transaction log mirror, especially if you do not frequently back up your SAP Sybase IQ database.

Syntax

```
dblog [options] database-file
```

Parameters

This table lists the options available for the **dblog** utility.

Table 22. dblog options

Option	Description
-ek <i>key</i>	Specify the encryption key for strongly encrypted databases directly in the command. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log in any way. For strongly encrypted databases, you must specify either -ek or -ep , but not both. The command will fail if you do not specify the correct key for a strongly encrypted database.
-ep	Specify that you want to be prompted for the encryption key. This option causes a window to appear, in which you enter the encryption key. It provides an extra measure of security by never allowing the encryption key to be seen in clear text. For strongly encrypted databases, you must specify either -ek or -ep , but not both. The command will fail if you do not specify the correct key for a strongly encrypted database.

Option	Description
-g <i>n</i>	Use this option if you are using the Log Transfer Manager to participate in a Replication Server installation. It can be used after a backup is restored, to set the generation number.
-il	Use this option if you have stopped using the Log Transfer Manager to participate in a Replication Server installation on this database, but continue to use SQL Remote or MobiLink synchronization. It resets the Log Transfer Manager log offset that is kept for the <code>delete_old_logs</code> option, allowing transaction logs to be deleted when they are no longer needed.
-is	Use this option if you have stopped using MobiLink synchronization on this database, but continue to use the Log Transfer Manager or SQL Remote. It resets the MobiLink log offset that is kept for the <code>delete_old_logs</code> option, allowing transaction logs to be deleted when they are no longer needed.
-m <i>mirror-name</i>	Specify the file name for a new transaction log mirror. If the database is not currently using a transaction log mirror, it starts using one. If the database is already using a transaction log mirror, it changes to using the new file as its transaction log mirror.
-n	Stop using a transaction log, and stop using a transaction log mirror. Without a transaction log, the database can no longer participate in data replication or use the transaction log in data recovery. If a SQL Remote, Log Transfer Manager, or dbmlsync truncation offset exists, the transaction log cannot be removed unless the corresponding ignore option (<code>-il</code> for the Log Transfer Manager, <code>-ir</code> for SQL Remote, or <code>-is</code> for dbmlsync) is also specified. You cannot stop using a transaction log if the database has auditing turned on (unless you first turn auditing off).
-o <i>filename</i>	Write output messages to the named file.
-q	Run in quiet mode—messages are not displayed.
-r	Maintain a single transaction log for databases that maintain a transaction log mirror.
-t <i>log-name</i>	Specify the file name for a new transaction log. If the database is not currently using a transaction log, it starts using one. If the database is already using a transaction log, it changes to using the new file as its transaction log.

Option	Description
-x n	Reset the transaction log current relative offset to n, so that the database can take part in replication. This option is used for reloading SQL Remote consolidated databases.
-z n	Reset the transaction log starting offset to n, so that the database can take part in replication. This option is used for reloading SQL Remote consolidated databases.

Usage

The **dblog** command line utility allows you to display or change the name of the transaction log or transaction log mirror. You can also stop or start maintaining a transaction log mirror.

The name of the transaction log is set when you create a database. The database cannot be running when you change its transaction log file name.

When you use the **RESTORE DATABASE** statement to move and/or rename a database, you can rename all of the files except the transaction log. Transactions continue to be written to the old log file name, in the location where the catalog store file (the `.db` file) is located after the database is restored.

When you rename or move all other files in the database, do the same for the log file. Use **dblog** to move or rename the log file. Run this utility after using **RESTORE DATABASE** with:

- A new database name
- The **RENAME** option

You can use **dblog** to rename the transaction log even if you have not restored the database, given these restrictions:

- The SAP Sybase IQ server must be stopped.
- After the log is renamed, retain the old log until the next database backup, in case the old log is needed for recovery from a media failure.

dblog displays additional information about the transaction log:

- Version number
- Starting offset, for use in replication
- Ending offset, for use in replication
- Page size
- Total number of pages
- Number of empty pages
- Percentage of the log file in use

dbping Database Administration Utility

The **dbping** utility is a ping utility that assists in diagnosing connection problems.

Syntax

```
dbping [options]
```

Parameters

This table lists the available options for dbping.

Table 23. dbping Options

Option	Description
-c " <i>keyword=value; ...</i> "	Supply database connection parameters. See <i>Appendix: Connection and Communication Parameters Reference in Administration: Database</i> . If no connection parameters are specified, connection parameters from the SQLCONNECT environment variable are used, if set.
-d	<p>Make a database connection if the server is found.</p> <p>If you do not supply the -d option, dbping reports success if it finds the server specified by the -c option. If you do supply the -d option, dbping reports success only if it connects to both server and database.</p> <p>For example, if you have a server named <code>blair</code> running the <code>iqdemo</code> database, this command succeeds:</p> <pre>dbping -c "eng=blair;dbn=iqdemo"</pre> <p>The following command fails, with the message <code>Ping database failed -- specified database not found</code></p> <pre>dbping -d -c "eng=blair;dbn=iqdemo"</pre>
-en	Specifies that you want dbping to exit with a failed return code when NULL is returned for any of the properties specified. By default, dbping prints NULL when the value for a property specified by -pc , -pd , or -ps is unknown, and exits with a success return code. This option can only be used with -pc , -pd , and -ps .

Option	Description
<p>-l <i>library</i></p>	<p>Use the specified ODBC driver or driver manager library. Specify the library to use (without its file extension). This option does not use the ODBC Driver Manager, and so is particularly useful on UNIX operating systems.</p> <p>For example, this command directly loads the ODBC driver:</p> <pre>dbping -m -c "dsn=IQ16iqdemo" -l dbodbc11</pre> <p>Use dbping to verify connectivity with the ODBC Driver Manager on UNIX systems. Unlike Interactive SQL or other tools, dbping allows you to explicitly test the components as you would with a third-party tool. For example:</p> <pre>dbping -m -c "dsn=dsname" -l /<full path>/libodbc.so</pre> <p>where <code>libodbc.so</code> is the third-party ODBC driver.</p>
<p>-m</p>	<p>Use the ODBC Driver Manager. Otherwise, connect using Embedded SQL™. Establish a connection using ODBC. By default, dbping attempts a connection using the embedded SQL interface.</p>
<p>-o <i>filename</i></p>	<p>Log output messages to a file.</p>
<p>-pc <i>property,...</i></p>	<p>Upon connection, display the specified connection properties. Supply the properties in a comma-separated list. You must specify enough connection information to establish a database connection.</p>
<p>-pd <i>property,...</i></p>	<p>Upon connection, display the specified database properties. Supply the properties in a comma-separated list. You must specify enough connection information to establish a database connection.</p>
<p>-ps <i>property,...</i></p>	<p>Upon connection, display the specified database server properties. Supply the properties in a comma-separated list.</p>
<p>-q</p>	<p>Operate quietly—do not print messages. If dbping fails, a message always appears.</p>

Option	Description
-s	Returns information about the performance of the network between the computer running dbping and the computer running the database server. Approximate connection speed, latency, and throughput are displayed. The -c option is usually required to specify the connection parameters to connect to a database on the server. You can only use dbping -s for embedded SQL connections. This option is ignored if -m or -l is also specified. By default, dbping -s loops through the requests for at least one second for each statistic it measures. A maximum of 200 connect and disconnect iterations are performed, regardless of the time they take, to avoid consuming too many resources. On slower networks, it can take several seconds to perform the minimum number of iterations for each statistic. The performance statistics are approximate, and are more accurate when both the client and server computers are fairly idle.
-st time	This option is the same as -s , except that it specifies the length of time, in seconds, that dbping loops through the requests for each statistic it measures. This option allows more accurate timing information to be obtained than -s .
-z	Display debugging information. This option is available only when an embedded SQL connection is being attempted. That is, it cannot be combined with -m or -l . It displays the network communication protocols used to attempt connection, and other diagnostic messages

Usage

Use **dbping** to help debug connection problems. You can enter a full or partial connection string; the utility returns a message indicating whether the attempt to locate a server or database, or to connect, was successful.

Use **dbping** for Embedded SQL or ODBC connections. You cannot use dbping for jConnect (TDS) connections.

Exit codes are 0 (success) or nonzero (failure).

This utility accepts @filename parameters.

dbstop Database Shutdown Utility

dbstop stops a database server. The dbstop utility is a command-line utility only, available on both UNIX and Windows platforms.

In UNIX, **dbstop** can shut down a server on any node on the network. You must specify a *server-name*, as well as any connection parameters you specified when you started the server. Without the proper connection parameters, **dbstop** does not know how to connect to the server to tell it to shutdown.

The **dbstop** command-line options control whether a server is stopped, even if there are active connections.

SERVER is the default logical server for **dbstop**.

Stopping the Database Server with dbstop

Run **dbstop** from the command line.

Issue a command in this format:

```
dbstop [ options ] server-name
```

For example, to stop a database named *iqdemo* on the server **myserver**, enter:

```
dbstop -c "uid=DBA;pwd=sql;eng=myserver;dbn=iqdemo"
```

The following example stops a server, **myserver**, regardless of the database running:

```
dbstop -c "uid=DBA;pwd=sql;eng=myserver;dbn=utility_db"
```

dbstop Options

This table lists the options available for the **dbstop** utility.

Table 24. dbstop Options

Switch	Description
@filename	Read in options from the specified environment variable or configuration file.
server-name	Server name of a running server to stop. If you supply a server name, do not supply connection parameters as well.

dbstop Database Shutdown Utility

Switch	Description
-c "keyword=value; ..."	<p>To stop a network server, supply a connection string with a user ID that has permissions to stop the server. By default, the SERVER OPERATOR system privilege is required on the network server, and all users can shut down a personal server; you can use the -gk server command-line option to change the default behavior.</p> <p>See <i>Appendix: Connection and Communication Parameters Reference</i> in <i>Administration: Database</i>.</p> <p>To stop a multiplex server, you need permissions to stop the server in the logical server context in which you connect to it. By default, dbstop logs in with SERVER context, so you would need permission to log into the SERVER logical server.</p> <p>If there are active connections, dbstop prompts whether you want to shut down the server. If you specify <code>unconditional=true</code> on the command line, the server shuts down without prompting, even if there are active connections.</p> <p>Avoid overriding the logical server in the connection string; if login redirection is enabled in the overriding user's login policy, the server may shut down as a connection is redirected.</p>
-d	Does not stop the database server. Instead, only stops the database specified in the connection string.
-o filename	Logs output messages to the named file.
-q	Quiet mode—does not print messages.
-x	Does not stop if there are active connections. Including this option prevents dbstop from prompting for confirmation if there are active connections.
-y	Stops without prompting even if there are active connections.

dbtran Database Administration Utility

Use the **dbtran** log translation utility, at the command prompt, to translate a transaction log into a `.sql` command file.

Syntax

Running against a database server:

```
dbtran [ options ] -c { connection-string } -n SQL-file
```

Running against a transaction log:

```
dbtran [ options ] [ transaction-log ] [ SQL-file ]
```

Parameters

Option	Description
<i>@data</i>	Reads in options from the specified environment variable or configuration file.
-a	Controls whether uncommitted transactions appear in the transaction log. The transaction log contains changes made only before the most recent COMMT by any transaction. If you do not specify -a, only committed transactions appear in the output file. If you specify -a, any uncommitted transactions found in the transaction log appear.
-c "keyword=value; ..."	Specifies the connection string when running the utility against a database server.
-d	Specifies that transactions are written in order from earliest to latest. This feature is intended for auditing database activity: do not apply dbtran output against a database.

Option	Description
-ek <i>key</i>	<p>Specifies the encryption key for strongly encrypted databases. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log. Specify either -ek or -ep, but not both. The command fails if you do not specify the correct encryption key. If you are running dbtran against a database server using the -c option, specify the key using a connection parameter instead of using the -ek option. For example, the following command gets the transaction log information about database <code>enc.db</code> from the database server <code>sample</code>, and saves its output in <code>log.sql</code>.</p> <pre data-bbox="548 557 1171 609">dbtran -n log.sql -c "ENG=sample;DBF=enc.db;UID=DBA;PWD=sql;DBKEY=mykey"</pre>
-ep	<p>Prompts for the encryption key. This option causes a window to appear, in which you enter the encryption key. It provides an extra measure of security by never allowing the encryption key to be seen in clear text. Specify either -ek or -ep, but not both. The command fails if you do not specify the correct encryption key. If you are running dbtran against a database server using the -c option, specify the key using a connection parameter, instead of using the -ep option. For example, the following command gets the transaction log information about database <code>enc.db</code> from the database server <code>sample</code>, and saves its output in <code>log.sql</code>.</p> <pre data-bbox="548 982 1171 1034">dbtran -n log.sql -c "ENG=sample;DBF=enc.db;UID=DBA;PWD=sql;DBKEY=mykey"</pre>
-f	Outputs only transactions completed since the last checkpoint.
-g	Adds auditing information to the transaction log if the auditing database option is turned on.
-ir <i>offset1,offset2</i>	Outputs a portion of the transaction log between two specified offsets.

Option	Description
-is <i>source</i> ,...	<p>Outputs operations on rows that have been modified by operations from one or more of the following sources, specified as a comma-separated list:</p> <ul style="list-style-type: none"> • All – all rows. This is the default setting. • SQLRemote – include only rows that were modified using SQL Remote. You can also use the short form “SR”. • RepServer – include only rows that were modified using the Replication Agent (LTM) and Replication Server. You can also use the short form “RS”. • Local – include only rows that are not replicated.
-it <i>owner.table</i> ,...	Outputs operations on the specified, comma-separated list of tables. Specify each table as owner.table.
-j date/time	Translates only transactions from the most recent checkpoint prior to the given date or time. The user-provided argument can be a date, time, or date and time, enclosed in quotes. you omit a time, the default is 00:00. If you omit a date, the current day is the default. The acceptable format for the date and time is: "YYYY/MMM/DD HH:NN".
-k	Prevents partial .sql files from being erased if an error is detected. If an error is detected while dbtran is running, the .sql file generated until that point is normally erased to ensure that a partial file is not used. Specifying this option may be useful if you are attempting to salvage transactions from a damaged transaction log.
-m	Specifies a directory that contains transaction logs. Use this option with the -n option.
-n <i>filename</i>	Specifies the output file that holds the SQL statements when you run dbtran against a database server.
-o <i>filename</i>	Writes output messages to the named file.
-r	Removes any uncommitted transactions. This is the default behavior.
-rsu <i>username</i> ,...	Specifies a comma-separated list of user names to override the default Replication Server user names. By default, the -is option assumes the default Replication Server user names of dbmaint and sa.

Option	Description
-s	Controls how UPDATE statements are generated. If you do not use this option, and there is no primary key or unique index on a table, dbtran generates UPDATE statements with a nonstandard FIRST keyword in case of duplicate rows. If you do use this option, the FIRST keyword is omitted for compatibility with the SQL standard.
-sr	Includes in the output file generated comments describing how SQL Remote distributes operations to remote sites.
-t	Controls whether triggers are included in the command file. By default, actions performed by triggers are not included in the command file. If the matching trigger is in the database, when the command file is run against the database, the trigger performs the actions automatically. Trigger actions should be included if the matching trigger does not exist in the database against which the command file is to run.
-u <i>userid</i> ,...	Limits the output from the transaction log to include only specified users.
-x <i>userid</i> ,...	Limits the output from the transaction log to exclude specified users.
-y	Replaces existing command files without prompting for confirmation. If you specify -q, you must also specify -y or the operation fails.
<i>transaction-log</i>	Specifies the log file to be translated. Cannot be used with -c or -m options.
<i>SQL-file</i>	Names the output file containing the translated information. For use with <i>transaction-log</i> only.

Usage

You can run **dbtran**:

- Against a database server — connects to the database server using the connection string specified following the -c option, and places output in a file specified with the -n option. The BACKUP DATABASE system privilege is required to run in this way. For example, this command translates log information from the **iqdemo** server and places the output in a file named `iqdemo.sql`:

```
dbtran -c "eng=iqdemo;dbn=iqdemo;dbf=iqdemo.db;uid=DBA;pwd=sql" -n iqdemo.sql
```

- Against a transaction log file — acts directly against a transaction log file. Protect your transaction log file from general access to prevent users from running this statement.


```
dbtran iqdemo.log iqdemo.sql
```

dbtran shows the earliest log offset in the transaction log, which you can use to determine the order in which multiple log files were generated.

dbtran-c attempts to translate the online transaction log file, and all the offline transaction log files in the same directory as the online transaction log file. If the directory contains transaction log files for more than one database, you may see an error. To avoid this, ensure that each directory contains transaction log files for only one database.

A transaction can span multiple transaction logs. If transaction log files contain transactions that span logs, translating a single transaction log file (for example, `dbtran demo.log`) might lose the spanning transactions. For **dbtran** to generate complete transactions, use the `-c` or `-m` options with the transaction log files in the directory.

Exit codes are 0 (success) or nonzero (failure).

This utility accepts `@filename` parameters.

dbvalid Database Administration Utility

The dbvalid utility is a validation utility that validates the indexes and keys on some or all of the SQL Anywhere tables in the catalog store.

The Validation utility scans the entire table and looks up each record in every index and key defined on the table. By default, the Validation utility uses the express check option.

Note: The **dbvalid** utility lets you easily validate catalog store tables, but does not validate IQ main store tables. Use the **sp iqcheckdb** stored procedure to validate IQ main store tables.

You can access the **dbvalid** utility at the system command-line level, which is useful for incorporating **dbvalid** into batch or command files.

Syntax

```
dbvalid [ options ] [object-name, ... ]
```

Parameters

This table lists the options available for the **dbvalid** utility.

Table 25. dbvalid Options

Option	Description
<i>object-name</i>	The name of a table or (if -i is used) an index to validate
-c " <i>keyword=value; ...</i> "	Supply database connection parameters. The user ID must have the VALIDATE ANY OBJECT system privilege. For example, the following validates the iqdemo database, connecting as user DBA with password sql: <pre>dbvalid -c "uid=DBA;pwd=sql;dbf-c:\sybase\IQ-16_0\demo\iqdemo.db"</pre>
-o <i>filename</i>	Log output messages to a file
-f	Validate tables with full check. In addition to the default validation checks, carry out both data checks (-fd) and index checks (-fi). This option corresponds to the WITH FULL CHECK option of the SQL Anywhere VALIDATE TABLE statement. Depending on the contents of your catalog store, this option may significantly extend the time required to validate.

Option	Description
-fd	<p>Validate tables with data check. In addition to the default validation checks, check that all of each LONG BINARY, LONG VARCHAR, TEXT or IMAGE data type can be read. Entries with these data types may span more than one page. In the catalog store:</p> <ul style="list-style-type: none"> • Domain – user-defined data type. • IMAGE – a domain to LONG BINARY . • TEXT – a domain to LONG VARCHAR . <p>This option instructs the database server to check all pages used by each entry. This corresponds to the WITH DATA CHECK option on the SQL Anywhere VALIDATE TABLE statement. Depending on the contents of your catalog store, this option may significantly extend the time required to validate.</p>
-fi	<p>Validate tables with index check. In addition to the default validation checks, validate each index on the table. This corresponds to the WITH INDEX CHECK option of the SQL Anywhere VALIDATE TABLE statement. Depending on the contents of your catalog store, this option may significantly extend the time required to validate.</p>
-fx	<p>Validate tables with express check. In addition to the default and data checks, check that the number of rows in the table matches the number of entries in the index. This corresponds to the WITH EXPRESS CHECK of the SQL Anywhere VALIDATE TABLE statement. This option does not perform individual index lookups for each row.</p>
-i	<p>Each <i>object-name</i> is an index. Instead of validating tables, validate indexes. Ensure that every row referenced in the index actually exists in the table. For foreign-key indexes, -i also ensures that the corresponding row exists in the primary table. If you supply a <i>table-name</i> instead of an <i>index-name</i>, validates the primary key index. In this case, for dbvalid, each of the <i>object-name</i> values supplied represents an index instead of a table and has a name in the following format:</p> <pre data-bbox="498 1272 973 1298">[[owner.] table-name.] index-name</pre> <p>Must be the owner of the table on which the index is created or have the VALIDATE ANY OBJECT system privilege.</p>
-q	<p>Operate quietly—do not print output messages.</p>

Option	Description
-s	Validate database pages using checksums. Checksums are used to determine whether a database page has been modified on disk. If you created a database with checksums enabled, you can validate the catalog store using checksums. Checksum validation reads each page of the catalog store from disk and calculates its checksum. If the calculated checksum is different from the checksum stored on the page, the page has been modified on disk and an error is returned. The page numbers of any invalid catalog store pages appear in the server messages window. You cannot use -s with -i , -t , or any of the -f options.
-t	The list of <i>object-name</i> values is a list of tables. This is also the default behavior.

Usage

With the **dbvalid** command-line utility, you can validate the indexes and keys on some or all of the tables in the catalog store. **dbvalid** scans the entire table and confirms that each row exists in the appropriate indexes. It is the same as running the SQL Anywhere **VALIDATE TABLE** statement on each catalog store table.

Note: **VALIDATE TABLE** is not supported in SAP Sybase IQ. **sp_iqcheckdb** provides a similar function for IQ main store tables.

By default, the Validation utility uses the express check option. However, the express check option is not used if you specify **-f**, **-fd**, **-fi**, **-fn**, or **-i**.

If the catalog store table is inconsistent, **dbvalid** reports an error. If errors are reported, you can drop all of the indexes and keys on a table and re-create them. You must also re-create any foreign keys to the table.

Warning! Validate a table or entire catalog store only when no connections are making changes to the database; otherwise, spurious errors may be reported, indicating some form of database inconsistency even though no inconsistency actually exists.

Program exit code	Description
0	Database validated successfully
1	General failure in utility
2	Error validating database
7	Cannot find database to connect to (database name is wrong)
8	Cannot connect to database (user ID/password is wrong)
11	Cannot find server to connect to (server name is wrong)

dbvalid Database Administration Utility

Program exit code	Description
12	Incorrect encryption key for starting database

Example

The following command validates the catalog store of the iqdemo database, connecting as user DBA with password sql:

```
dbvalid -c "uid=DBA;pwd=sql;dbf-c:\sybase\IQ-16_0\demo\iqdemo.db"
```

iqdscp Configuration Utility

The **iqdscp** utility is a UNIX-only Open Client and Open Server configuration utility you can use to configure the interfaces file and to configure a directory service.

See *Open Server 15.5 > Open Client and Open Server Configuration Guide for UNIX > Using dscp*.

iqdsn Database Administration Utility

The **iqdsn** utility is a data source utility you use for creating, changing, deleting, describing, and listing SAP Sybase IQ ODBC data sources.

The **iqdsn** utility is a cross-platform alternative to the ODBC Administrator.

On Windows operating systems, data sources are held in the Registry. On UNIX operating systems, data sources are held in the `.odbc.ini` file. Use the **iqdsn** utility for batch operations.

Syntax

```
iqdsn [ modifier-options ]
{
  -l
  | -d dsn
  | -g dsn
  | -w dsn [ details-options;... ]
  | -cl }
```

Parameters

These tables list the available options for the **iqdsn** utility.

Table 26. iqdsn Major Options

Major Option	Description
-l	Lists the available SAP Sybase IQ ODBC data sources. You can modify the list format using the -b or -v options.
-d dsn	Deletes the named data source. If you supply -y , any existing data source is overwritten without confirmation.
-g dsn	List the definition of the named data source. You can modify the format of the output using the -b or -v option.
-w dsn [details-options]	Creates a new data source, or overwrites one if one of the same name exists. If you supply -y , any existing data source is overwritten without confirmation.
-cl	This convenience option lists the connection parameters supported by the iqdsn utility.

Table 27. iqdsn Modifier-Options

Modifier option	Description
-b	Brief. Format the output of the list as a single line connection string.

Modifier option	Description
-or	Creates a data source for the SQL Anywhere 16 – Oracle ODBC driver when specified with the -c option. For example: <code>dbdsn -w MyOracleDSN -or -c Userid=DBA;Password=sql;SID=abcd;ArraySize=500;ProcResults=y</code> You can specify the -cl option with the -or option to obtain a list of the connection parameters for the SQL Anywhere 16 Oracle ODBC driver.
-q	Quiet. Do not print the informational banner.
-v	Verbose. Format the output of the list over several lines, as a table.
-va	Verbose All. Print connection parameters in same format as -v , but also include other hidden parameters. Use this option to display ODBC driver qualifier needed for remote data access on those UNIX platforms that support such access, or for some third-party driver managers.
-y	Automatically delete or overwrite each file without prompting you for confirmation.

Table 28. iqdsn Details-Options

Details Option	Description
-c "keyword= value;..."	Specify connection parameters as a connection string. The connection string consists of a set of parameters, separated by semicolons.
-ec encryption type	Encrypt packets sent between the client application and the server. See <i>Encryption Connection Parameter [ENC]</i> .
-o filename	Write output messages to the named file. By default, messages are written to the console. See <i>LogFile Connection Parameter [LOG]</i> .
-p size	Set the maximum packet size for network communications, in bytes. The value must be greater than 500, and less than 16000. The default setting is 1460. See <i>CommBufferSize Connection Parameter [CBSize]</i> .
-r	Disable multiple record fetching. By default, when the database server gets a simple fetch request, the application asks for extra rows. Use the -r option to disable this behavior. See <i>DisableMultiRowFetch Connection Parameter [DMRF]</i> .

Details Option	Description
-tl <i>seconds</i>	Client liveness timeout period. Terminates connections when they are no longer intact. The value is in seconds. The default is the server setting, which in turn has a default of 120 seconds. See <i>LivenessTimeout Connection Parameter [LTO]</i> .
-z	Provide diagnostic information on communications links on startup.
<i>server-name</i>	Connect to the named database server. Only the first 40 characters are used.

Usage

The **iqdsn** modifier options can occur before or after the major option specification. The order makes a difference only if you specify a connection parameter value more than once. In such a case, the last value specified is used.

Examples

This example writes a definition of the data source `newdsn` and does not prompt for confirmation if the data source already exists.

```
iqdsn -y tcpip -w newdsn -c "uid=DBA;pwd=sql" -v
```

Lists all known user data sources, one data source name per line:

```
iqdsn -l
```

Lists all data sources along with their associated connection string:

```
iqdsn -l -b
```

Reports the connection string for user data source `MyDSN`:

```
iqdsn -g MyDSN
```

Deletes the data source `BadDSN`, but first list the connection parameters for `BadDSN` and prompt for confirmation:

```
iqdsn -d BadDSN -v
```

Deletes the data source `BadDSN` without prompting for confirmation.

```
iqdsn -d BadDSN -y
```

Creates a data source named `NewDSN` for the database server `MyServer`:

```
iqdsn -w NewDSN -c "uid=DBA;pwd=sql;eng=bar"
```

If a `NewDSN` already exists, the utility asks you if you want to overwrite it.

The following example connects to the `sample` database server. The server name **sample** overrides the previous specified value of `MyServer`:

```
iqdsn -w NewDSN -c "uid=DBA;pwd=sql;eng=MyServer" sample
```

Lists all connection parameter names and their aliases:

```
iqdsn -cl
```

Encryption Connection Parameter [ENC]

Encrypts packets sent between the client application and the database server using transport-layer security or simple encryption.

Usage

For **TLS**, TCP/IP only

For **NONE** or **SIMPLE**, anywhere

Values

String

Default

NONE

If an Encryption value is not set, encryption is controlled by the setting on the server, which defaults to no encryption.

Description

You can use this parameter if you are concerned about the security of network packets. Encryption marginally affects performance. The **Encryption (ENC)** connection parameter accepts these arguments:

- **None** accepts communication packets that are unencrypted.
- **Simple** accepts communication packets that are encrypted with simple encryption supported on all platforms and on pre-12.6 versions of SAP Sybase IQ. Simple encryption does not provide server authentication, RSA encryption, or other features of transport-layer security.
- **TLS** accepts communication packets that are encrypted using RSA encryption technology. For FIPS-certified RSA encryption, specify FIPS=Y. RSA FIPS uses a separate certified library, but is compatible with SAP Sybase IQ servers specifying RSA. TLS connections are supported between client and server on all supported platforms but FIPS=Y is only supported on LinuxAMD and Windows (32 and 64-bit). To authenticate the server, the software verifies that the server's certificate values match any values you supply about the client using the following arguments:
 - **trusted_certificate** specifies the certificate file the client uses to authenticate the server.

- **certificate_company** specifies the value for the organization field. The server's value and the client's value must match.
- **certificate_unit** specifies the value for the organization unit field. The server's value and the client's value must match.
- **certificate_name** specifies the certificate's common name. The server's value and the client's value must match.

Warning! Use the sample certificate only for testing purposes. It provides no security in deployed situations because it and the corresponding password are widely distributed with SAP Sybase IQ software. To protect your system, create your own certificate.

You can use the **CONNECTION_PROPERTY** system function to retrieve the encryption settings for the current connection.

Examples

- This connection string fragment connects to a database server myeng with a TCP/IP link, using RSA encryption and the sample trusted certificate:

```
"ENG=myeng; LINKS=tcPIP;
Encryption=(FIPS=N;TRUSTED_CERTIFICATE=iq-16_0/samples/
certificates/rsaroot.crt) "
```

See also

- *-ec iqsrv16 database server option* on page 145

LogFile Connection Parameter [LOG]

Sends client error messages and debugging messages to a file.

Usage

Anywhere

Values

String

Description

To save client error messages and debugging messages in a file, use the **LogFile (LOG)** parameter.

If the file name includes a path, it is relative to the current working directory of the client application.

The LogFile (LOG) connection parameter is connection-specific, so from a single application you can set different LogFile arguments for different connections.

Examples

This command line fragment specifies that client messages for this connection should be sent to the file `error.log` in the current working directory for the client:

```
...
LogFile=error.log
...
```

CommBufferSize Connection Parameter [CSize]

Sets the maximum size of communication packets, in bytes. Use `k` to specify units of kilobytes.

Syntax

{ CommBufferSize | CBSIZE }=size[k]

Usage

Anywhere

Allowed Values

- **size** – specifies the maximum size of communication packets. The default value is in bytes, but you can use `k` to specify units of kilobytes. The minimum value of `CommBufferSize` is 500 bytes, and the maximum is 65535 bytes.

Default

If no `CommBufferSize` value is set, the `CommBufferSize` is controlled by the setting on the server, which defaults to 7300 bytes.

Remarks

The protocol stack sets the maximum size of a packet on a network. If you set the `CommBufferSize` to be larger than that permitted by your network, the communication packets are broken up by the network software. The default size is a multiple of the standard ethernet TCP/IP maximum packet size (1460 bytes).

A larger packet size may improve performance for multi-row fetches and fetches of larger rows, but it also increases memory usage for both the client and the server.

If `CommBufferSize` is not specified on the client, the connection uses the server's buffer size. If `CommBufferSize` is specified on the client, the connection uses the `CommBufferSize` value.

Using the `-p` database server option to set the `CommBufferSize` causes all clients that do not specify their own `CommBufferSize` to use the size specified by the `-p` database server option.

Example

To set the buffer size to 1460 bytes:

```
...
CommBufferSize=1460
...
```

Alternatively, you can set this parameter by entering its value in the **Buffer size** text box of the **Network** tab of the connection window.

DisableMultiRowFetch Connection Parameter [DMRF]

Turns off multi record fetches across the network.

Usage

Anywhere

Default

No

By default, when the database server gets a simple fetch request, the application asks for extra rows. You can disable this behavior by setting this parameter to ON.

Setting the **DisableMultiRowFetch** parameter to ON is equivalent to setting the PREFETCH option to OFF.

Example

This connection string fragment prevents prefetching:

```
DMRF=Yes
```

LivenessTimeout Connection Parameter [LTO]

Controls the termination of connections when they are no longer intact.

Usage

Network server on TCP/IP communications protocols.

All platforms except non threaded UNIX applications.

Values

Integer (in seconds)

Default

120

If no **LivenessTimeout** value is set, the liveness timeout is controlled by the setting on the server, which defaults to 120 seconds.

Description

A liveness packet is periodically sent across a client/server TCP/IP communications protocol to confirm that a connection is intact. If the client runs for the liveness timeout period without detecting a liveness request or response packet, communication is severed.

Liveness packets are sent when a connection has not sent any packets for between one third and two thirds of the **LivenessTimeout** value.

When communication is severed, the client machine forgets the address of the server. It looks the address up next time there is a connection to the server from that machine, dropping all current connections to that server.

When there are more than 200 connections to a server, the server automatically calculates a higher **LivenessTimeout** value based on the stated **LivenessTimeout** value. This enables the server to more efficiently handle a large number of connections.

Alternatively, you can set this parameter by entering its value in the LivenessTimeout text box of the Network tab of the ODBC Configuration dialog.

Example

This sets a Liveness timeout value of 60 seconds:

```
LTO=60
```


iqheader Database Administration Utility

The **iqheader** utility is a dbspace header utility that determines which server, if any, is using a particular device, file, or LUN (Logical Unit Number) as a dbspace to analyze disk usage or to configure a multiplex query server.

The **iqheader** utility reports the configuration of an arbitrary device regardless of whether it is currently in use by an SAP Sybase IQ server.

The user interface is a standalone console application called **iqheader** (**header.exe** on Windows). The **iqheader** tool searches the device for an SAP Sybase IQ dbspace header and reports the header information in a user-readable format.

Note: LUN is a logical unit number and is used to identify SCSI devices so the host can address and access the data on each disk drive in an array.

Syntax

```
iqheader [ [ dbspace_path ]
```

Parameters

The **iqheader** application takes a single parameter, which is the device to be checked.

Usage

- **iqheader usage** – When invoked with no parameters, a usage summary is reported and a nonzero status is returned:

```
>iqheader
Usage: iqheader [dbspace_path]
```

- **iqheader error** – When the specified target is not a dbspace, an error message is reported and a nonzero status is returned:

```
>iqheader /dev/null
Not an IQ file: Error 0
```

- **Operating system errors** – When the specified target is unreadable, or any file operation fails due to an error returned from the operating system, the native operating system error appears and a nonzero status is returned:

```
>iqheader /dev/rdsk/c1t32d0s1 <
Open Failed: No such file or directory

>iqheader /dev/rdsk/c1t3d0s1 <
Open Failed: Permission denied
```

- **iqheader output** – When a valid dbspace is specified, **iqheader** prints the dbspace configuration to the console and returns a 0 exit status.

Table 29. iqheader output

Field	Description
File Name	Name of the file.
Full Path	Full path after symbolic link resolution
Version	Dbospace file format version
File ID	Unique number assigned to each dbospace
Create Time	Time of dbospace creation
RW Mode	Current read-write mode: RW, RO, RW, N/A (Upgraded)
Last RW Mode	Last dbospace mode
Size (MB)	DBSpace size, in megabytes
Reserve (MB)	DBSpace reserve size, in megabytes
Block Size	Size of block, in bytes
Page Size	Size of page, in bytes
First Block	First block number mapped to dbospace
Block Count	Number of blocks that map to actual disk blocks
Reserve Blocks	Number of blocks that may be added to this dbospace
Last Real Block	Last block number that maps to an actual disk block
Last Mapped Block	Last block number mapped to dbospace
OFlag	Online status (YES/NO)
Create ID	Commit ID in which dbospace was created
Alter ID	Last commit ID in which dbospace was altered
DBID1	Location of first database identity
DBID2	Location of second database identity
DBSpace ID	Unique identifier assigned to each dbospace
_NextFLAllocLowerBank	Lower bound of pre-allocate space for dbospace
_NextFLAllocUpperBank	Upper bound of pre-allocate space for dbospace
Pre-alter commit ID	Commit_id in the system tables ISYSDBFILE and ISYSIQDBSPACE
_ReqNumFreeListBlocks	Number of blocks of type 'F' (free list blocks)

Example

This example shows output for iqheader:

```
File Name: file1.iq
Full Path: /dev/dsk/file1.iq
DBFile Header Info
Version: 2
File ID: 16395
Create Time: 2008-06-02 21:57:00
RW Mode: RO
Last RW Mode: RW
Size (MB): 20
Reserve (MB): 20
Block Size: 8192
Page Size: 131072
First Block: 9408960
Block Count: 2560
Reserve Blocks: 2560
Last Real Block: 9411519
Last Mapped Block: 10454399
OFlags: 1
Create ID: 6905
Alter ID: 6964
DBID1: 0
DBID2: 0
DBSpace ID: 16395
  NextFLAllocLowerBank: 0
  NextFLAllocUpperBank: 0
Pre-alter commit ID: 6925
Dropped: NO
```


iqinit Database Administration Utility

The **iqinit** utility starts a database that gives the user executing the utility permission to create a database. The user must have access to the computer and file system.

iqinit is the SAP Sybase IQ version of the SQL Anywhere **dbinit** utility. **iqinit** lets you create either an an SAP Sybase IQ or SQL Anywhere database from the command line without starting a database:

- If no **iqinit** command line parameters are specified, **iqinit** creates a SQL Anywhere database.
- If the **-iqpath** command line parameter is specified, **iqinit** creates an SAP Sybase IQ database.

Syntax

```
iqinit [ options ] new-database-file
```

Parameters

This table lists the available options for the **iqinit** utility.

Table 30. iqinit Options

Option	Description
<i>@data</i>	Reads in options from the specified environment variable or configuration file. If you want to protect passwords or other information in the configuration file, you can use the File Hiding utility to obfuscate the contents of the configuration file.
-a	Causes string comparisons to respect accent differences between letters (for example, e is less than é if the Unicode Collation Algorithm (UCA) is used for either CHAR or NCHAR data types (see -z and -zn). With the exception of Japanese databases created with a UCA collation, by default, accents are ignored (meaning e is equal to é). If all base letters (letters with accents and case removed) are otherwise equal, then accents are compared from left to right. The default accent sensitivity of a UCA collation when creating a Japanese database is sensitive. That is, accents are respected.

Option	Description
-af	<p>Causes string comparisons to respect accent differences between letters (for example, e is less than é) if the UCA is used for either CHAR or NCHAR data types (see -z and -zn). By default, accents are ignored (meaning e is equal to é). If all base letters (letters with accents removed) are otherwise equal, then accents are compared from right to left, consistent with the rules of the French language.</p>
-b	<p>Blank pads the database.</p> <p>The database server compares all strings as if they are varying length and stored using the VARCHAR domain. This includes string comparisons involving fixed length CHAR or NCHAR columns. In addition, SQL Anywhere never trims or pads values with trailing blanks when the values are stored in the database.</p> <p>By default, the database server treats blanks as significant characters. So, the value 'a ' (the character 'a' followed by a blank) is not equivalent to the single-character string 'a'. Inequality comparisons also treat a blank as any other character in the collation.</p> <p>If blank padding is enabled (the iqinit -b option), the semantics of string comparisons more closely follow the ANSI/ISO SQL standard. With blank-padding enabled, SQL Anywhere ignores trailing blanks in any comparison.</p> <p>In the example above, an equality comparison of 'a ' to 'a' in a blank-padded database returns TRUE. With a blank-padded database, fixed-length string values are padded with blanks when they are fetched by an application. The ansi_blanks connection option controls whether the application receives a string truncation warning on such an assignment.</p>
-c	<p>Considers all values case sensitive in comparisons and string operations. Identifiers in the database are case insensitive, even in case sensitive databases.</p> <p>With the exception of Japanese databases created with a UCA collation, the default behavior is that all comparisons are case insensitive. The default case sensitivity of a UCA collation when creating a Japanese database is sensitive.</p> <p>This option is provided for compatibility with the ISO/ANSI SQL standard.</p>

Option	Description
<p>-dba [<i>DBA-user</i>] [<i>pwd</i>]</p>	<p>Specifies the DBA user ID and password. If you specify a new name for the DBA user for the database, you can no longer connect to the database as the user DBA. You can also specify a different password for the DBA database user. If you do not specify a password, the default password sql is used. If you do not specify this option, the default user ID DBA with password sql is created.</p> <p>The following command creates a database with a DBA user named testuser with the default password sql:</p> <pre>iqinit -dba testuser, mydb.db</pre> <p>The following command uses the default user ID DBA with password mypwd:</p> <pre>iqinit -dba ,mypwd mydb.db</pre> <p>The following command changes the DBA user to user1 with password mypwd:</p> <pre>iqinit -dba user1,mypwd mydb.db</pre>
<p>-dbs <i>size</i>[k m g p]</p>	<p>Pre-allocates space for the database. Pre-allocating space for the database helps reduce the risk of running out of space on the drive the database is located on. As well, it can help improve performance by increasing the amount of data that can be stored in the database before the database server needs to grow the database, which can be a time-consuming operation.</p> <p>By default, the size is in bytes. You can use k, m, g, or p to specify units of kilobytes, megabytes, or gigabytes, or pages, respectively.</p>

Option	Description
<p>-ea <i>algorithm</i></p>	<p>Specifies the encryption algorithm used for database or table encryption (-et). Specify -ea simple for simple encryption (do not specify -ek or -ep). Simple encryption is equivalent to obfuscation and is intended only to keep data hidden in the event of casual direct access of the database file, to make it more difficult for someone to decipher the data in your database using a disk utility to look at the file.</p> <p>For greater security, specify AES or AES256 for 128-bit or 256-bit strong encryption, respectively. Specify AES_FIPS or AES256_FIPS for 128-bit or 256-bit FIPS-certified encryption, respectively. For strong encryption, you must also specify the -ek or -ep option.</p> <p>To create a database that is not encrypted, specify -ea none, or do not include the -ea option (and do not specify -et, -ep, or -ek). If you do not specify the -ea option, the default behavior is as follows:</p> <ul style="list-style-type: none"> • -ea none, if -ek, -ep, or -et is not specified • -ea AES, if -ek or -ep is specified (with or without -et) • -ea simple, if -et is used without -ek or -ep <p>Algorithm names are case insensitive.</p> <p>File compression utilities cannot compress encrypted database files as much as unencrypted ones.</p>
<p>-ek <i>key</i></p>	<p>Specifies that you want to create a strongly encrypted database by specifying an encryption key directly in the command. The -ek option is used with an AES algorithm, optionally specified using the -ea option. If you specify the -ek option without specifying the -ea option, AES is used by default. When specified with -et, the database is not encrypted. Instead, table encryption is enabled.</p>
<p>-ep</p>	<p>Specifies that you want to create a strongly encrypted database by inputting the encryption key in a window. This provides an extra measure of security by never allowing the encryption key to be seen in clear text. You must input the encryption key twice to confirm that it was entered correctly. If the keys don't match, the initialization fails. When specified with -et, the database is not encrypted. Instead, table encryption is enabled.</p>

Option	Description
-et	<p>Enables table encryption using the encryption algorithm (and key) specified for the -ea option. Use this option when you want to create encrypted tables instead of encrypting the entire database. If you specify -et with -ek or -ep, but not -ea, the AES algorithm is used by default. When you specify only -et, simple encryption is used.</p> <p>Enabling table encryption does not mean your tables are encrypted. You must encrypt tables individually, after database creation.</p> <p>When table encryption is enabled, table pages for the encrypted table, associated index pages, and temporary file pages are encrypted, and the transaction log pages that contain transactions on encrypted tables.</p> <p>The following example creates the database new.db with strong encryption enabled for tables using the key abc, and the AES_FIPS encryption algorithm:</p> <pre>iqinit -et -ek abc -ea AES_FIPS new.db</pre>
-i	<p>Excludes jConnect™ for JDBC™ system objects from the database. To use the jConnect JDBC driver to access system catalog information, you need jConnect catalog support (installed by default). When you specify this option, you can still use JDBC, as long as you do not access system information. You can add jConnect support at a later time using the ALTER DATABASE statement.</p>
-iqblksize	<p>The I/O transfer block size in bytes.</p>
-iqmsgpath	<p>The path name of the segment containing the SAP Sybase IQ message trace file.</p>
-iqpath	<p>The path name of the main segment file containing the SAP Sybase IQ data.</p>
-iqpgsize	<p>The page size in bytes for the SAP Sybase IQ segment of the database.</p>
-iqreservesize	<p>The size in MB of the space to reserve for the IQ main store.</p>
-iqsize	<p>The size in MB of either the raw partition or OS file with the -iqpath.</p>
-iqtmppath	<p>The path name of the temporary segment file.</p>
-iqtmpreservesize	<p>The size in MB to reserve for the temporary IQ store.</p>
-iqtmpsize	<p>The size in MB of either the raw partition or OS file for the -iqtmppath.</p>
-k	<p>Does not create the SYSCOLUMNS and SYSINDEXES views. By default, database creation generates the views SYS.SYSCOLUMNS and SYS.SYSINDEXES for compatibility with system tables that were available in Watcom SQL (versions 4 and earlier of this software). These views conflict with the Sybase Adaptive Server Enterprise compatibility views dbo.syscolumns and dbo.sysindexes.</p>

Option	Description
-l	Lists the available collation sequences and then stops. No database is created. To specify a collation sequence, use the <code>-z</code> option.
-le	<p>Lists the available character set encodings and then stops. No database is created. Each character set encoding is identified by one or more labels. These are strings that can be used to identify the encoding. Each line of text that appears lists the encoding label and alternate labels by which the encoding can be identified. These labels fall into one of several common categories: SA (the SQL Anywhere label), IANA (Internet Assigned Numbers Authority), MIME (Multipurpose Internet Mail Extensions), ICU (International Components for Unicode), JAVA, or ASE (Adaptive Server Enterprise).</p> <p>If you want to view a list of character set encodings that includes the alternate labels, specify the <code>-le+</code> option.</p> <p>When the <code>iqinit</code> utility reports the character set encoding, it always reports the database server version of the label. For example, the following command reports the CHAR character set encoding <code>windows-1250</code>:</p> <pre>iqinit -ze cp1250 -z uca test.db</pre>
-m <i>filename</i>	Creates a transaction log mirror. A transaction log mirror is an identical copy of a transaction log, usually maintained on a separate device, for greater protection of your data. By default, the database server does not use a transaction log mirror.
-n	Creates a database without a transaction log. Creating a database without a transaction log saves disk space, but can result in poorer performance because each commit causes a checkpoint. Also, if your database becomes corrupted and you are not running with a transaction log, data is not recoverable. The transaction log is required for data replication and provides extra security for database information during a media or system failure.
-o <i>filename</i>	Writes output messages to the named file.

Option	Description
-p <i>page-size</i>	<p>Specifies the page size for the database. The page size for a database can be (in bytes) 2048, 4096, 8192, 16384, or 32768, with 4096 being the default. Use <i>k</i> to specify units of kilobytes (for example, -p 4k).</p> <p>Large databases can benefit from a larger page size. For example, the number of I/O operations required to scan a table is generally lower, as a whole page is read in at a time. However, there are additional memory requirements for large page sizes. It is strongly recommended that you do performance testing (and testing in general) when choosing a page size. Then choose the smallest page size that gives satisfactory results. For most applications, 16 KB or 32 KB page sizes are not recommended. You should not use page sizes of 16 KB or 32 KB in production systems unless you can be sure that a large database server cache is always available, and only after you have investigated the trade offs of memory and disk space with its performance characteristics. If a large number of databases are going to be started on the same server, pick a reasonable page size.</p>
-pd	<p>If this option is not specified (default), all privileged system procedures execute with the privileges of the invoker. If specified, pre-16.0 privileged system procedures execute with the privileges of the definer, while 16.0 or later privileged system procedures execute with the privileges of the invoker.</p>
-q	<p>Runs in quiet mode—messages are not displayed.</p>
-s	<p>Adds global checksums (a checksum is added to each database page). By default, this option is on. Checksums are used to determine whether a database page has been modified on disk. When you create a database with global checksums enabled, a checksum is calculated for each page just before it is written to disk. The next time the page is read from disk, the page's checksum is recalculated and compared to the checksum stored on the page. If the checksums are different, then the page has been modified or corrupted on disk, and an error occurs. Critical database pages are always checksummed by the database server, regardless of the value of the -s option.</p> <p>Checksums are automatically enabled for databases running on storage devices such as removable drives to help provide early detection if the database becomes corrupt.</p> <p>If a database is created with global checksums disabled, you can still add checksums to pages when they are written by using the -wc option or the START DATABASE statement.</p>

Option	Description
<p>-t <i>transaction-log-name</i></p>	<p>Specifies the name of the transaction log file. The transaction log is a file where the database server logs all changes, made by all users, no matter what application is being used. The transaction log plays a key role in backup and recovery, and in data replication. If the file name has no path, it is placed in the same directory as the database file. If you run iqinit without specifying -t or -n, a transaction log is created with the same file name as the database file, but with extension <code>.log</code>.</p>
<p>-z <i>coll [collation-tailoring-string]</i></p>	<p>Specifies the collation sequence for the database. The collation sequence is used for sorting and comparing character data types (CHAR, VARCHAR, and LONG VARCHAR). The collation provides character comparison and ordering information for the encoding (character set) being used. It is important to choose your collation carefully. It cannot be changed after the database has been created without unloading and reloading the database. If the collation is not specified, the database server chooses a collation based on the operating system language and character set. To view the available collation sequences, see the -l option.</p> <p>Optionally, you can specify collation tailoring options (<i>collation-tailoring-string</i>) for additional control over the sorting and comparing of characters. These options take the form of keyword=value pairs, assembled in parentheses, following the collation name. For example:</p> <p>Case and accent settings specified in the <i>collation-tailoring-string</i> override case and accent options for iqinit (-c, -a, and -af), if you specify both.</p>
<p>-ze <i>encoding</i></p>	<p>Specifies the encoding for the collation. Most collations specified by -z dictate both the encoding (character set) and ordering. For those collations, -ze should not be specified.</p> <p>If the collation specified by -z is Unicode Collation Algorithm (UCA), then -ze can specify UTF-8 or any single-byte encoding for CHAR data types.</p> <p>By default, the database server uses UTF-8. Use -ze to specify a locale-specific encoding and get the benefits of the UCA for comparison and ordering.</p>

Option	Description
<p>-zn <i>coll</i> [<i>collation-tailoring-string</i>]</p>	<p>Specifies the collation sequence used for sorting and comparing of national character data types (NCHAR, NVARCHAR, and LONG NVARCHAR). The collation provides character ordering information for the UTF-8 encoding (character set) being used. Values are UCA (the default), or UTF8BIN which provides a binary ordering of all characters whose encoding is greater than 0x7E. If the dbicu16 and dbicudt16 DLLs are not installed, then the default NCHAR collation is UTF8BIN</p> <p>Optionally, you can specify collation tailoring options (<i>collation-tailoring-string</i>) for additional control over the sorting and comparing of characters. These options take the form of keyword=value pairs, assembled in parentheses, following the collation name. For example:</p> <pre>iqinit -c -zn UCA(case=LowerFirst) sens.db</pre> <p>Case and accent settings specified in the <i>collation-tailoring-string</i> override case and accent options for iqinit (-c, -a, and -af), if you specify both.</p>

Example

This command creates a SAP Sybase IQ database called `bar.iq`.

```
$ iqinit -iqpath bar.iq -iqsize 20M
-iqpggsize 2048
-iqreservesize 10M bar.db
SQL Anywhere Initialization Utility Version 16.0.0.5530
Debug
CHAR collation sequence:
ISO BINENG(CaseSensitivity=Respect)
CHAR character set encoding: ISO_8859-1:1987
NCHAR collation sequence:
UCA(CaseSensitivity=UpperFirst;
AccentSensitivity=Respect;
PunctuationSensitivity=Primary)
NCHAR character set encoding: UTF-8
Creating system tables
Creating system views
Setting option values
Database "bar.db" created successfully
```


iqocscfg Configuration Utility

The **iqocscfg** utility is a Windows-only Open Client and Open Server configuration utility you can use to configure environment variables, directory drivers, and security drivers.

See *Open Server 15.5 > Software Developer's Kit 15.5 > Open Client and Open Server Configuration Guide for Microsoft Windows > Using ocscfg*.

isql Interactive SQL Utility

The **isql** utility is a command-line Interactive SQL utility that uses the Adaptive Server Enterprise Open Client API.

You cannot create user-defined database options in **isql**. If you need to add your own database options, use the **dbisql** Interactive SQL utility instead.

Syntax

```
isql
[-b] [-e] [-F] [-p] [-n] [-v] [-W] [-X] [-Y] [-Q]
  [-a display_charset]
  [-A packet_size]
  [-c cmdend]
  [-D database]
  [-E editor]
  [-h header]
  [-H hostname]
  [-i inputfile]
  [-I interfaces_file]
  [-J client_charset]
  [-K keytab_file]
  [-l login_timeout]
  [-m errorlevel]
  [-o outputfile]
  [-P password]
  [-R remote_server_principal]
  [-s colseparator]
  [-S server_name]
  [-t timeout]
  -U username
  [-V [security_options]]
  [-w columnwidth]
  [-z locale_name]
  [-Z security_mechanism]
  [--conceal]
  [--URP]
```

Parameters

This table lists the available options for the **isql** utility.

Table 31. isql Options

Option	Description
-b	Disables display of the table headers output.

Option	Description
-e	Echoes input.
-F	Enables the FIPS flagger. When you specify the -F parameter, the server returns a message when it encounters a non-standard SQL command. This option does not disable SQL extensions. Processing completes when you issue the non-ANSI SQL command.
-p	Prints performance statistics.
-n	Removes numbering and the prompt symbol (>) from the echoed input lines in the output file when used in conjunction with -e .
-v	Prints the version number and copyright message for isql and then exits.
-X	<p>Initiates the login connection to the server with client-side password encryption. -X enables both extended password encrypted connections and password encrypted connections without plain text password reconnection. isql (the client) specifies to the server that password encryption is desired. The sever sends back an encryption key, which isql uses to encrypt your password, and the server uses the key to authenticate your password when it arrives.</p> <p>If isql crashes, the system creates a core file that contains your password. If you did not use the encryption option, the password appears in plain text in the file. If you used the encryption option, your password is not readable</p>
-W	Disables both extended password and password encrypted negotiations.
-Y	Tells the Adaptive Server to use chained transactions.
-Q	Provides clients with failover property.

Option	Description
-a <i>display_charset</i>	<p>Runs isql from a terminal whose character set differs from that of the machine on isql is running. Use -a with -J to specify the character set translation file (.xlt file) required for the conversion. Use -a without -J only if the client character set is the same as the default character set.</p> <hr/> <p>Note: The <code>ascii_7</code> character set is compatible with all character sets. If either the Adaptive Server character set or the client character set is set to <code>ascii_7</code>, any 7-bit ASCII character can pass unaltered between client and server. Other characters produce conversion errors.</p>
-A <i>packet_size</i>	<p>Specifies the network packet size to use for this isql session. For example, the following sets the packet size to 4096 bytes for this isql session:</p> <pre>isql -A 4096</pre> <ul style="list-style-type: none"> • To check your network packet size, enter: <code>SELECT * FROM sysprocesses</code> • The value is displayed under the <code>network_pktsz</code> heading. • <i>size</i> must be between the values of the default network packet size and maximum network packet size configuration parameters, and must be a multiple of 512. • Use larger-than-default packet sizes to perform I/O-intensive operations, such as readtext or writetext operations. • Setting or changing Adaptive Server's packet size does not affect the packet size of remote procedure calls.
-c <i>cmdend</i>	<p>Changes the command terminator. By default, you terminate commands and send them to by typing “go” on a line by itself. When you change the command terminator, do not use SQL reserved words or control characters.</p>
-D <i>database</i>	<p>Selects the database in which the isql session begins.</p>

Option	Description
-E <i>editor</i>	Specifies an editor other than the default editor, vi . To invoke the editor, enter its name as the first word of a line in isql .
-h <i>headers</i>	Specifies the number of rows to print between column headings. The default prints headings only once for each set of query results.
-H <i>hostname</i>	Specifies the number of rows to print between column headings. The default prints headings only once for each set of query results.
-i <i>interfaces_file</i>	Specifies the name and location of the interfaces file to search when connecting to Adaptive Server. If you do not specify -I , isql looks for a file named <code>interfaces</code> in the directory specified by your SYBASE environment variable.
-J <i>client_charset</i>	<p>Specifies the character set to use on the client. -J client_charset requests that Adaptive Server convert to and from <i>client_charset</i>, the character set used on the client. A filter converts input between <i>client_charset</i> and the Adaptive Server character set.</p> <p>-J with no argument sets character set conversion to NULL. No conversion takes place. Use this if the client and server use the same character set.</p> <p>Omitting -J sets the character set to a default for the platform. The default may not necessarily be the character set that the client is using. For more information about character sets and the associated flags, see “Configuring Client/Server Character Set Conversions,” in the <i>Adaptive Server Enterprise System Administration Guide, Volume One</i>.</p>
-K <i>keytab_file</i>	Specifies the path to the keytab file used for authentication in DCE.

Option	Description
-l <i>login_timeout</i>	Specifies the maximum timeout value allowed when connecting to Adaptive Server. The default is 60 seconds. This value affects only the time that isql waits for the server to respond to a login attempt. To specify a timeout period for command processing, use the -t <i>timeout</i> parameter.
-m <i>errorlevel</i>	Customizes the error message display. For errors of the severity level specified or higher, only the message number, state, and error level are displayed; no error text appears. For error levels lower than the specified level, nothing appears.
-o <i>outputfile</i>	Specifies the name of an operating system file to store the output from isql . Specifying the parameter as -o <i>outputfile</i> is similar to > <i>outputfile</i> .
-P <i>password</i>	Specifies your Adaptive Server password. If you do not specify the -P flag, isql prompts for a password. If your password is NULL, use the -P flag without any password.
-R <i>remote_server_principal</i>	Specifies the principal name for the server as defined to the security mechanism. By default, a server's principal name matches the server's network name (which is specified with the -S parameter or the DSQUERY environment variable). Use the -R parameter when the server's principal name and network name are not the same.
-s <i>colseparator</i>	Resets the column separator character, which is blank by default. To use characters that have special meaning to the operating system (for example, "[", ":", "&", "<", ">"), enclose them in quotes or precede them with a backslash.
-S <i>server_name</i>	Specifies the name of the Adaptive Server to which to connect. isql looks this name up in the interfaces file. If you specify -S with no argument, isql looks for a server named SYBASE. If you do not specify -S , isql looks for the server specified by your DSQUERY environment variable.

Option	Description
<p>-t <i>timeout</i></p>	<p>Specifies the number of seconds before a SQL command times out. If you do not specify a timeout, the command runs indefinitely. This affects commands issued from within isql, not the connection time. The default timeout for logging into isql is 60 seconds.</p>
<p>-U <i>username</i></p>	<p>Specifies a login name. Login names are case sensitive.</p>
<p>--URP <i>string</i></p>	<p>Allows you to supply a login redirection string for an SAP Sybase IQ server. This example specifies the logical server, database, node type and redirection setting:</p> <pre data-bbox="716 652 1137 730">isql -Usa -P --URP "LS=sal-ogsrv;dbname=dbone;node=writer;redirect=no"</pre> <p>or</p> <p>Enables setting the universal remote password <i>remotepassword</i> for clients accessing Adaptive Server. The application sets the universal remote password. For example, ctlib uses <code>ct_remote_pwd()</code> and jConnect uses the <code>setRemotePassword</code> method.</p>

Option	Description
<p>-V <i>security_options</i></p>	<p>Specifies network-based user authentication. With this option, the user must log in to the network's security system before running the utility. In this case, users must supply their network user name with the -U option; any password supplied with the -P option is ignored.</p> <p>You can follow -V with a <i>security_options</i> string of key-letter options to enable additional security services. These key letters are:</p> <ul style="list-style-type: none"> • -c – Enable data confidentiality service • -i – Enable data integrity service • -m – Enable mutual authentication for connection establishment • -o – Enable data origin stamping service • -q – Enable out-of-sequence detection • -r – Enable data replay detection
<p>-w <i>columnwidth</i></p>	<p>Sets the screen width for output. The default is 80 characters. When an output line reaches its maximum screen width, it breaks into multiple lines.</p>
<p>-z <i>locale_name</i></p>	<p>Specifies the official name of an alternate language to display isql prompts and messages. Without -z, isql uses the server's default language. You can add languages to an Adaptive Server during installation or afterward, using the langinstall utility (langinst in Windows) or the sp_addlanguage stored procedure.</p>
<p>-Z <i>security_mechanism</i></p>	<p>Specifies the name of a security mechanism to use on the connection.</p> <p>Security mechanism names are defined in the <code>libtcl.cfg</code> configuration file located in the <code>ini</code> subdirectory below the Sybase installation directory. If no <i>security_mechanism name</i> is supplied, the default mechanism is used. For more information on security mechanism names, see the description of the <code>libtcl.cfg</code> file in the <i>Open Client and Open Server Configuration Guide</i>.</p>

Option	Description
--conceal	<p>Hides your input during an isql session. The --conceal option is useful when entering sensitive information, such as passwords. The syntax for the --conceal option is:</p> <pre>isql --conceal [':?' 'wild-card']</pre> <p>The maximum length of <i>wildcard</i>, a 32-byte variable, is 80 characters.</p> <hr/> <p>Note: --conceal is silently ignored in batch mode.</p>

Usage

Use this syntax for **isql_r** if you are using threaded drivers.

Use this syntax for **isql** if you are using threaded drivers in the IBM platform.

You must set the SYBASE environment variable to the location of the current version of Adaptive Server before you can use isql.

This table lists the commands you can execute within interactive **isql**:

Table 32. Commands Available Within isql

Command	Description
:r filename	<p>Reads an operating system file into the command buffer.</p> <p>Do not include the command terminator in the file; once you have finished editing, enter the terminator interactively on a line by itself.</p>
:R filename	<p>Reads an operating system file into the command buffer and then displays it.</p> <p>Do not include the command terminator in the file; once you have finished editing, enter the terminator interactively on a line by itself.</p>
use database_name	Changes the current database.
!! os_command	Executes an operating system command. Place at the start of a line.

Command	Description
> <i>file_name</i>	Redirects the output of the T-SQL command to <i>file_name</i> . The following example inserts the server version into <i>file_name</i> : <pre>select @@version go > file_name</pre>
>> <i>file_name</i>	Appends the output of the T-SQL command to <i>file_name</i> . The following example appends the server version into <i>file_name</i> : <pre>select @@version go >> file_name</pre>
<i>command</i>	Pipes the output of the T-SQL command to an external command. The following example finds all instances of “sa” in the listing produced by sp_who : <pre>sp_who go grep sa</pre>
vi (UNIX) edit (Windows)	Calls the default editor.
reset	Clears the query buffer.
quit or exit	Exits isql .

The 5701 (“changed database”) server message is no longer displayed after login or issuing a **use database** command.

Error message format differs from earlier versions of **isql**. If you have scripts that perform routines based on the values of these messages you may need to rewrite them.

When you include the **-X** parameter, the password-enabled connection proceeds according to server capabilities:

- If the server can handle both extended password and password encryption, extended password encryption negotiations are used.
- If the server can handle password encryption only, password encryption negotiations are used.
- If the server cannot handle password encryption or extended password encryption, the first connection attempt fails and the client attempts to reconnect using a plain text password.

To use **isql** interactively, give the command **isql** (and any of the optional parameters) at your operating system prompt. The **isql** program accepts SQL commands and sends them to

isql Interactive SQL Utility

Adaptive Server. The results are formatted and printed on standard output. Exit **isql** with **quit** or **exit**.

Terminate a command by typing a line beginning with the default command terminator `go` or another command terminator, if the `-c` parameter is used. You can follow the command terminator with an integer to specify how many times to run the command. For example, to execute this command 100 times, type:

```
select x = 1
go 100
```

The results display once at the end of execution.

If you enter an option more than once on the command line, `isql` uses the last value. For example, if you enter the following command, “send”, the second value for `-c`, overrides “.”, the first value:

```
isql -c"." -csend
```

This enables you to override any aliases you set up.

To call an editor on the current query buffer, enter its name as the first word on a line. Define your preferred callable editor by specifying it with the `EDITOR` environment variable. If `EDITOR` is not defined, the default is `vi` on UNIX and `edit` on Windows.

For example, if your `EDITOR` environment variable is set to “`emacs`,” then you must invoke it from within `isql` with “`emacs`” as the first word on the line.

To clear the existing query buffer, type `reset` on a line by itself. `isql` discards any pending input. You can also press `Ctrl-c` anywhere on a line to cancel the current query and return to the `isql` prompt.

Read in an operating system file containing a query for execution by `isql` as follows:

```
isql -U alma -P password < input_file
```

The file must include a command terminator. The results appear on your terminal. Read in an operating system file containing a query and direct the results to another file as follows

```
isql -U alma -P password < input_file > output_file
```

isql displays only six digits of float or real data after the decimal point, rounding off the remainder.

You can include comments in a Transact-SQL statement submitted to Adaptive Server by **isql**. Open a comment with “`/*`”. Close it with “`*/`”, as shown in the following example:

```
select au_lname, au_fname
/*retrieve authors' last and first names*/
from authors, titles, titleauthor
where authors.au_id = titleauthor.au_id
and titles.title_id = titleauthor.title_id
/*this is a three-way join that links authors
**to the books they have written.**/
```

If you want to comment out a **go** command, it should not be at the beginning of a line. For example, use the following to comment out the **go** command:

```
/*
**go
*/
```

Do not use the following:

```
/* go */
```

isql defines the order of the date format as month, date, and year (mm dd yyyy hh:mmAM (or PM)) regardless of the locale environment. To change this default order, use the **convert** function.

In an **isql** session, the default prompt label is either the default wildcard `?:` or the value of *wildcard*. You can customize the prompt label by providing a one-word character string, with a maximum length of 80 characters after a wildcard. If you specify a prompt label that is more than one word, the characters after the first word are ignored.

Note: In an **isql** session, **isql** recognizes `?:`, or the value of *wildcard*, as wildcards only when these characters are placed at the beginning of an **isql** line.

Examples

Puts you in a text file where you can edit the query. When you write and save the file, you are returned to **isql**. The query appears; type "go" on a line by itself to execute it:

```
isql -Ujoe -Pabracadabra
1> select *
2> from authors
3> where city = "Oakland"
4> vi
```

reset clears the query buffer. **quit** returns you to the operating system:

```
isql -Ualma
Password:
1> select *
2> from authors
3> where city = "Oakland"
4> reset
1> quit
```

Specifies that you are running **isql** from a Macintosh against a server that is using the roman8 character set:

```
isql -a mac -J roman8
```

Creates column separators using the “#” character in the output in the pubs2 database for store ID 7896:

```
isql -Uasa -P -s#
1> use pubs2
2> go
1> select * from sales where stor_id = "7896"
#stor_id#ord_num          #date          #
```

isql Interactive SQL Utility

```
#-----#-----#-----#
#7896 #124152 # Aug 14 1986 12:00AM#
#7896 #234518 # Feb 14 1991 12:00AM#

(2 rows affected)
```

Changes password without displaying the password entered. This example uses “old” and “new” as prompt labels:

```
$ isql -Uguest -Pguest -Smyase --conceal
sp_password
:~old
/
:~:~ new
-----
old
new
Confirm new
Password correctly set.
(Return status 0)
```

Activates a role for the current user. This example uses a custom wildcard and the prompt labels “role” and “password:”

```
$ isql -UmyAccount --conceal '*'Password:
set role
* role
with passwd
** password
on
go

role
password
Confirm password
```

iqsqlpp SQL Preprocessor Utility

The SAP Sybase IQ SQL preprocessor utility **iqsqlpp** translates the SQL statements in an input file (.sql) into C language source that is put into an output file (.c).

Embedded SQL is a database programming interface for the C and C++ programming languages. Embedded SQL consists of SQL statements intermixed with (embedded in) C or C++ source code. These SQL statements are translated by an SQL preprocessor into C or C++ source code, which you then compile.

Syntax

```
iqsqlpp [ options ] <in filename > [<out filename> ]
```

Parameters

This table lists the options available for the **iqsqlpp** utility.

Table 33. iqsqlpp Options

Option	Description
-d	Favor data size.
-e <level>	<p>Flag nonconforming SQL syntax as an error.</p> <p>The allowed values of <level> are:</p> <ul style="list-style-type: none"> • c03 – Flag syntax that is not core SQL/2003 syntax • p03 – Flag syntax that is not full SQL/2003 syntax • c99 – Flag syntax that is not core SQL/1999 syntax • p99 – Flag syntax that is not full SQL/1999 syntax • e92 – Flag syntax that is not entry-level SQL/1992 syntax • i92 – Flag syntax that is not intermediate-level SQL/1992 syntax • f92 – Flag syntax that is not full-SQL/1992 syntax • t – Flag syntax that is not full-SQL/1992 syntax • u – Flag non-standard host variable types <p>The following are also supported for compatibility with previous versions: e (for entry-level SQL92), i, (for intermediate-level SQL92), f (for full SQL92), and w (to allow all supported syntax).</p>
-h <width>	Limit the maximum line length of output.
-k	Include user declaration of SQLCODE.
-n	Line numbers.

Option	Description
-o <O/S spec>	Target operating system specification (WINDOWS, WINNT or UNIX).
-q	Quiet mode—do not print banner.
-r-	Generate reentrant code.
-s <len>	Maximum string constant length for the compiler.
-w <level>	Flag nonconforming SQL syntax as a warning. The allowed values of <level> are: <ul style="list-style-type: none"> • c03— Flag syntax that is not core SQL/2003 syntax • p03 – Flag syntax that is not full SQL/2003 syntax • c99 – Flag syntax that is not core SQL/1999 syntax • p99 – Flag syntax that is not full SQL/1999 syntax • e92 – Flag syntax that is not entry-level SQL/1992 syntax • i92 – Flag syntax that is not intermediate-level SQL/1992 syntax • f92 – Flag syntax that is not full-SQL/1992 syntax • t – Flag syntax that is not full-SQL/1992 syntax • u – Flag non-standard host variable types
-x	Change multibyte SQL strings to escape sequences.
-z <cs>	Specify the collation sequence. For a list of recommended collation sequences, enter iqinit -l at a command prompt.

Usage

The SQL preprocessor processes a C or C++ program containing Embedded SQL before the compiler is run. **iqsqlpp** translates the SQL statements in the input file *sql-filename* into C language source that is put into the *output-filename*. The normal extension for source programs with Embedded SQL is *.sqc*. The default output file name is the *sql-filename* with an extension of *.c*. If the *sql-filename* has a *.c* extension, the default output file name extension is *.CC*.

Options

- **-d** – Favor data size. Generate code that reduces data space size. Data structures are reused and initialized at execution time before use. This increases code size.
- **e <flag>** – This option flags, as an error, any Embedded SQL that is not part of a specified set of SQL92.

The allowed values of <flag> and their meanings are as follows:

- **e** – Flag syntax that is not entry-level SQL92 syntax.

- **i** – Flag syntax that is not intermediate-level SQL92 syntax.
 - **f** – Flag syntax that is not full-SQL92 syntax.
 - **t** – Flag non-standard host variable types.
 - **w** – Allow all supported syntax.
 - **u** – Flag syntax that is not supported by UltraLite.
 - **c99** – Flag syntax that is not core SQL/1999 syntax.
 - **c03** – Flag syntax that is not core SQL/2003 syntax.
 - **p99** – Flag syntax that is not full SQL/1999 syntax.
 - **p03** – Flag syntax that is not full SQL/2003 syntax.
- **-h width** – Limits the maximum length of lines output by **iqlpp** to *width*. The continuation character is a backslash (\), and the minimum value of *width* is ten.
 - **-k** – Notifies the preprocessor that the program to be compiled includes a user declaration of SQLCODE.
 - **-n** – Generate line number information in the C file. This consists of *#line* directives in the appropriate places in the generated C code. If the compiler you are using supports the *#line* directive, this option makes the compiler report errors on line numbers in the SQC file (the file with the Embedded SQL) as opposed to reporting errors on line numbers in the C file generated by the SQL preprocessor. Also, the *#line* directives are used indirectly by the source level debugger so that you can debug while viewing the SQC source file.
 - **-o <O/S spec>** Specify the target operating system. This option must match the operating system where you run the program. A reference to a special symbol is generated in your program. This symbol is defined in the interface library. If you use the wrong operating system specification or the wrong library, an error is detected by the linker. The supported operating systems are:
 - **WINDOWS** – Microsoft Windows
 - **UNIX** – Use this option if you are creating a 32-bit UNIX application.
 - **UNIX64** – Use this option if you are creating a 64-bit UNIX application.
 - **-q** – Operate quietly. Do not print the banner.
 - **-r** – Generate reentrant code.
 - **-s <len>** – Set the maximum size string that the preprocessor puts into the C file. Strings longer than this value are initialized using a list of characters (“a,” “b,” “c,” and so on). Most C compilers have a limit on the size of string literal they can handle. This option is used to set that upper limit. The default value is 500.
 - **-w <flag>** – This option flags any Embedded SQL that is not part of a specified set of SQL92 as a warning.

The allowed values of *<flag>* and their meanings are as follows:

iqsqlpp SQL Preprocessor Utility

- **e** – Flag syntax that is not entry-level SQL92 syntax.
- **i** – Flag syntax that is not intermediate-level SQL92 syntax.
- **f** – Flag syntax that is not full-SQL92 syntax.
- **t** – Flag non-standard host variable types.
- **u** – Flag syntax that is not supported by UltraLite.
- **w** – Allow all supported syntax.
- **c03** – Flag syntax that is not core SQL/2003 syntax.
- **c99** – Flag syntax that is not core SQL/1999 syntax.
- **p99** – Flag syntax that is not full SQL/1999 syntax.
- **p03** – Flag syntax that is not full SQL/2003 syntax.
- **-x** – Change multibyte strings to escape sequences so that they can pass through compilers.
- **-z <cs>** – This option specifies the collation sequence. For a list of recommended collation sequences, enter **iqinit -l** at a command prompt.

The collation sequence helps the preprocessor understand the characters used in the source code of the program, for example, in identifying alphabetic characters suitable for use in identifiers. If you do not specify **-z**, the preprocessor attempts to determine a reasonable collation to use, based on the operating system and the IQLANG and IQCHARSET environment variables.

start_iq Database Server Startup Utility

The database startup utility **start_iq** starts a SAP Sybase IQ network database server.

The UNIX versions of SAP Sybase IQ provide the script **start_iq**, and the Windows version provides **start_iq.exe**. These scripts verify that your environment is set correctly and start the server with all required switches preset to recommended defaults (along with any switches you add). The **start_iq** utility also includes some parameters and calculates others. For switches that are specific to your operating system, see the *Installation and Configuration Guide*.

Note: SAP Control Center provides an easy graphical interface for starting servers and is the recommended method for starting multiplex servers.

Starting the start_iq Utility

Start the **start_iq** database startup utility from the command line on Windows and UNIX.

1. Issue a command in this format:

```
start_iq [ server-options ] [ database-file  
[ database-options ], ...]
```

2. Specify the **-n** switch in [*server-options*] to prevent you from unintentionally connecting to the wrong server.

Listing all start_iq Switches

You can display a list of all available switches for the *server-options*, *database-file*, and *database-options* parameters.

Issue the following command:

```
start_iq -?
```

start_iq Server Options

Available switches for the **start_iq** *server-options* parameters.

See also

- *Database Server Naming Restrictions* on page 174

@data iqsrv16 database server option

Reads in options from the specified environment variable or configuration file.

Syntax

```
iqsrv16 @data ...
```

Applies to

All operating systems and database servers. It is supported for all database utilities except the Language Selection utility (dblang), the Certificate Creation utility (createcert), the Certificate Viewer utility (viewcert), the Microsoft ActiveSync provider install utility (mlasinst), and the File Hiding utility (dbfhide).

- the Language Selection utility (dblang)
- the Certificate Creation utility (createcert)
- the Certificate Viewer utility (viewcert)
- the Microsoft ActiveSync provider install utility (mlasinst)
- the File Hiding utility (dbfhide)

Remarks

Use this option to read in command-line options from the specified environment variable or configuration file. If both exist with the same name that is specified, the environment variable is used.

Configuration files can contain line breaks, and can contain any set of options.

To protect the information in a configuration file (for example, because it contains passwords), use the File Hiding (dbfhide) utility to obfuscate the contents of configuration files.

The @data parameter can occur at any point in the command, and parameters contained in the file are inserted at that point. Multiple files can be specified, and the file specifier can be used with command line options.

Example

The following configuration file holds a set of options for a server named `myserver` that starts with a cache size of 4 MB and loads the sample database:

```
-c 4096  
-n myserver  
"c:\mydatabase.db"
```

If this configuration file is saved as `c:\config.txt`, it can be used in a command as follows:

```
iqsrv16 @c:\config.txt
```

The following configuration file contains comments:

```
#This is the server name:  
-n MyServer
```

```
#These are the protocols:
-x tcpip
#This is the database file
my.db
```

The following statement sets an environment variable that holds options for a database server that starts with a cache size of 4 MB and loads the sample database.

```
SET envvar=-c 4096 "c:\mydatabase.db";
```

The following command starts the database server using an environment variable named envvar.

```
iqsrv16 @envvar
```

@envvar iqsrv16 Server Option

Reads in command-line switches from the supplied environment variable.

Syntax

@envvar

Remarks

The environment variable may contain any set of command-line switches. For example, the first of the following pair of Windows statements sets an environment variable holding a set of command-line switches and loads the sample database. The second statement starts the database server:

```
set envvar= -gp 4096 -gm 15
c:\sybase\IQ-16_0\demo\start_iq -n myserver @envvar iqdemo.db
```

Note: If you have both a file and an environment variable with the value of your @ command-line switch, the result is unpredictable. Use only one of these methods to set a given @ command-line switch.

@filename iqsrv16 Server Option

Read in switches from configuration file.

Syntax

@filename

Remarks

The file may contain line breaks, and may contain any set of command-line switches. For example, the following Windows command file holds a set of command-line switches for a server named iqdemo that allows 10 connections, and sets the maximum catalog page size to 4096 bytes, and starts the iqdemo database:

```
# iqdemo.cfg
# -----
# Default startup parameters for the IQ demo database
```

start_iq Database Server Startup Utility

```
# -----  
-n iqdemo  
-x tcpip{port=2638}  
# The following parameters are also found in the configuration file  
# %IQDIR%\scripts\default.cfg. Any parameters not specified below  
# and not in the start up parameter list, will be added by start_iq  
# using default.cfg as a guide.  
  
-c 48m  
-gc 20  
-gd all  
-gl all  
-gm 10  
-gp 4096  
-iqmc 32  
-iqtc 24
```

If this configuration file is saved as `c:\config.txt`, the file can be used in a command line as follows:

```
start_iq @c:\config.txt
```

-c iqsrv16 Server Option

Sets initial memory reserved for caching catalog store pages and other server information.

Syntax

-c *cache-size* [**k** | **m** | **g** | **p**]

Remarks

The database server uses extra memory for caching database pages if the memory is set aside in the cache. Any cache size less than 10000 is assumed to be KB (1K = 1024 bytes). Any cache size 10000 or greater is assumed to be in bytes. You can also specify the cache size nK, nM or nP (1M = 1024 KB), where P is a percentage of the physical system memory.

In the `default.cfg` file, the default value of `-c` and `start_iq` is 32MB (`-c 32M`) for Windows platforms, and 48MB (`-c 48M`) for UNIX platforms. Use this default or set `-c` to a higher value.

You can use % as an alternative to P, but as most non-UNIX operating systems use % as an environment variable escape character, you must escape the % character. For example, to use 20 percent of the physical system memory, specify:

```
start_iq -c 20%% ...
```

Do not use `-c` in the same configuration file or command line with `-ch` or `-cl`. For related information, see the `-ch cache-size` option and the `-ca 0` option.

-al iqsrv16 Server Option

Extends LOGIN_MODE for LDAPUA only to a select number of users using Standard authentication

Syntax

```
-al "user1;user2;user3" server_name.cfg database-name.db
```

Remarks

- Up to five user IDs can be specified, separated by semi-colons, and enclosed in double quotation marks.
- When run at the server level, the **-al** switch remains in effect until the next time the server is restarted.

-ca iqsrv16 database server option

Enforces a static cache size.

Syntax

```
iqsrv16 -ca 0 ...
```

Applies to

Windows, Unix

Remarks

You can disable automatic cache size tuning by specifying **-ca 0** option. If you do not include the **-ca 0** option, the database server automatically increases the cache size. If you specify this option, the cache size is still adjusted if the database server would otherwise run into an error indicating that the dynamic memory is exhausted.

This server option must only be used in the form **-ca 0**.

Example

The following example starts a database server named myserver that has a static cache that is 40% of the total physical memory and loads the sample database, and the database server does not automatically adjust the cache size:

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

The following example is for Windows:

```
start_iq @"%
  ALLUSERSPROFILE%\SybaseIQ\demo\iqdemo.cfg" -c 40P
  -ca
  0 -n myserver "%
  ALLUSERSPROFILE%\SybaseIQ\demo\iqdemo.db"
```

The following example is for Unix:

start_iq Database Server Startup Utility

```
start_iq
_@$IQDIR16/demo/iqdemo.cfg -c 40P -ca 0
$IQDIR16/demo/iqdemo.db
```

-cc iqsrv16 database server option

Collects information about database pages to be used for cache warming the next time the database is started.

Syntax

```
iqsrv16 -cc{ + | - } ...
```

Applies to

All operating systems and database servers.

Remarks

By default, page collection is turned on. When collection is turned on, the database server keeps track of each database page that is requested. Collection stops when the maximum number of pages has been collected, the database is shut down, or the collection rate falls below the minimum value. You cannot configure the maximum number of pages collected or specify the value for the collection rate (the value is based on cache size and database size). Once collection stops, information about the requested pages is recorded in the database so those pages can be used to warm the cache the next time the database is started with the `-cr` option. Collection of referenced pages is turned on by default.

-ch iqsrv16 Server Option

Set catalog store cache size upper limit in bytes.

Syntax

```
-ch size[ k| m| g| p]
```

Remarks

By default, the upper limit is approximately the lower of 256MB and 90% of the physical memory of the machine.

You specify the cache-size using the K, M, and P characters as in the `-c` option. For the meaning and usage of the cache size argument and the K, M, and P characters, see `-c cache-size`.

In some cases the standard catalog cache size may be too small, for example, to accommodate certain queries that need a lot of parsing. In these cases, you may find it helpful to set `-cl` and `-ch`. For example, on 32-bit platforms, try these settings:

```
-cl 128M
-ch 512M
```

Warning! To control catalog store cache size explicitly, you must do either of the following, but not both, in your configuration file (.cfg) or on the UNIX command line for server startup:

- Set the **-c** parameter.
- Set specific upper and lower limits for the catalog store cache size using the **-ch** and **-cl** parameters.

Specifying different combinations of the parameters above can produce unexpected results.

See also

- *-c iqsrv16 Server Option* on page 138
- *-cl iqsrv16 database server option* on page 141

-cl iqsrv16 database server option

Sets a minimum cache size as a lower limit to dynamic cache resizing.

Syntax

```
iqsrv16 -cl { size[ k | m | g | p ] } ...
```

Default

2 MB on Windows

8 MB on Unix

Applies to

Windows, Unix

Remarks

This option sets a lower limit to the cache. If **-c** is specified, and **-cl** is not specified, then the minimum cache size is set to the initial cache size (the **-c** setting). If neither **-c** nor **-cl** is set, the minimum cache is set to a low, constant value, so that the cache can shrink if necessary. On Windows platforms, this value is 2 MB.

The *size* is the amount of memory, in bytes. Use *k*, *m*, or *g* to specify units of kilobytes, megabytes, or gigabytes, respectively.

The unit *p* is a percentage either of the physical system memory, or of the maximum supported cache size, whichever is lower. The maximum supported cache size depends on the operating system. For example:

- 2.5 GB for Windows 32-bit Advanced Server, Enterprise Server, and Datacenter Server
- 3.5 GB for the 32-bit database server running on Windows x64 Edition
- 1.5 GB on all other 32-bit systems
- On 64-bit database servers, the cache size can be considered unlimited

start_iq Database Server Startup Utility

If you use `p`, the argument is a percentage. You can use `%` as an alternative to `p`, but on Windows, which uses `%` as an environment variable escape character, you must escape the `%` character.

For example, to set the minimum cache size to 50 percent of the physical system memory, run the following command:

```
iqsrv16 -cl 50%% ...
```

Note: If you attempt to set your initial or minimum cache sizes to a value that is less than one eighth of the maximum cache size, the initial and minimum cache sizes are automatically increased relative to the supported cache size.

Example

The following example starts a database server named `silver` that has a minimum cache size of 5 MB and loads the database file `example.db`:

```
iqsrv16 -cl 5m -n silver "c:\example.db"
```

-cp iqsrv16 database server option

Specifies an additional set of directories or JAR files in which to search for classes.

Syntax

```
iqsrv16 -cp location[ ;location ... ] ...
```

Applies to

All operating systems and database servers.

Remarks

It is recommended that all classes and JAR files that are being used with Java in the database be installed in the database. When you store the classes and JAR files within the database, the database can be easily moved to a different computer or operating system. Another benefit of installing classes and JAR files into the database is that the database server's class loader can fetch the classes and resources from the database, allowing each connection that is using Java in the database to have its own instance of these classes and its own copy of static variables within these classes.

However, when a class or JAR file must be loaded by the system class loader, it can be specified with the `java_class_path` database option or the `-cp` database server option. Both options add classes and JAR files to the classpath that the database server builds for launching the Java VM. The `java_class_path` database option is useful when the server is running multiple databases and each database has a different set of JARs and directories that need to be loaded by the system class loader. The `-cp` database server option is useful when all databases on the server require the same classes or JAR files.

-cr iqsrv16 database server option

Reloads (warms) the cache with database pages using information collected the last time the database was run.

Syntax

```
iqsrv16 -cr{ + | - } ...
```

Applies to

All operating systems and database servers.

Remarks

You can instruct the database server to warm the cache using pages that were referenced the last time the database was started (page collection is turned on using the -cc option). Cache warming is turned on by default. When a database is started, the server checks the database to see if it contains a collection of pages requested the last time the database was started. If the database contains this information, the previously-referenced pages are then loaded into the cache.

Warming the cache with pages that were referenced the last time the database was started can improve performance when the same query or similar queries are executed against a database each time it is started.

-cs iqsrv16 database server option

Displays statistics related to dynamic cache sizing in the database server messages window.

Syntax

```
iqsrv16 -cs ...
```

Applies to

Windows, Unix

Remarks

For troubleshooting purposes, this option displays statistics in the database server messages window that database server is using to determine how to tune size of the cache.

-cv iqsrv16 database server option

Controls the appearance of messages about cache warming in the database server messages window.

Syntax

```
iqsrv16 -cv{ + | - } ...
```

start_iq Database Server Startup Utility

Default

Cache warming messages are suppressed.

Applies to

All operating systems and database servers.

Remarks

When `-cv+` is specified, a message appears in the database server messages window when any of the following cache warming activities occur:

- collection of requested pages starts or stops (controlled by the `-cc` server option)
- page reloading starts or stops (controlled by the `-cr` server option)

Example

The following command starts the database `mydatabase.db` with database page collection and page loading turned on, and logs messages about these activities to the database server messages window:

```
iqsrv16 -cc+ -cr+ -cv+ mydatabase.db
```

-dt iqsrv16 database server option

Specifies the directory where temporary files are stored.

Syntax

```
iqsrv16 -dt temp-file-dir ...
```

Applies to

All servers and operating systems, except shared memory connections on Unix.

Remarks

The database server creates two types of temporary files:

- **Database server-related temporary files (created on all platforms)** – You can use the `-dt` option to specify a directory for database server-related temporary files. If you do not specify this option when starting the database server, the database server examines the following environment variables, in the order shown, to determine the directory in which to place the temporary file.

Windows/Unix
<ol style="list-style-type: none">1. IQTMP162. TMP3. TMPDIR4. TEMP

If none of the environment variables are defined, the database server places its temporary file in the current directory on Windows, and in the /tmp directory on Unix.

- **Communications-related temporary files (created only on Unix for both the client and the database server)** – Temporary files for communications on Unix are not placed in the directory specified by -dt. Instead, the database server examines the following environment variables, in the order shown, to determine the directory in which to place the temporary file:

Unix
<ol style="list-style-type: none"> 1. IQTMP16 2. TMP 3. TMPDIR 4. TEMP

If none of the environment variables are defined, the database server places its temporary file in the /tmp directory on Unix.

On Unix, both the client and the database server must set IQTMP16 to the same value when connecting via shared memory.

Example

To locate the database server-related temporary files, use the DB_PROPERTY system function with the TempFileName property:

```
SELECT DB_PROPERTY ( 'TempFileName' );
```

-ec iqsrv16 database server option

Uses transport-layer security or simple encryption to encrypt all command sequence communication protocol packets (such as DBLib and ODBC) transmitted to and from all clients. TDS packets aren't encrypted.

Syntax

```
iqsrv16 -ec encryption-options ...
```

```
encryption-options :
```

```
{ NONE |
  SIMPLE |
  TLS ( [ FIPS={ Y | N }; ]
  IDENTITY=server-identity-filename;
  IDENTITY_PASSWORD=password ) }, ...
```

Allowed values

- **NONE** – accepts connections that aren't encrypted.

start_iq Database Server Startup Utility

- **SIMPLE** – accepts connections that are encrypted with simple encryption. This type of encryption is supported on all platforms, and on previous versions of the database server and clients. Simple encryption doesn't provide server authentication, RSA encryption, or other features of transport-layer security.
- **TLS** – accepts connections that are encrypted with RSA encryption. The TLS parameter accepts the following arguments:
 - **FIPS** – For FIPS-certified RSA encryption, specify `FIPS=Y`. RSA FIPS-certified encryption uses a separate certified library, but is compatible with version 9.0.2 or later clients specifying RSA.

For a list of FIPS-certified components, see <http://www.sybase.com/detail?id=1061806>.

The algorithm must match the encryption used to create your certificates.

- **server-identity-filename** – is the path and file name of the server identity certificate. If you are using FIPS-certified RSA encryption, you must generate your certificates using the RSA algorithm.
- **password** – is the password for the server private key. You specify this password when you create the server certificate.

Applies to

NONE and SIMPLE apply to all servers and operating systems.

TLS applies to all servers and operating systems.

For information about supporting FIPS-certified encryption, see <http://www.sybase.com/detail?id=1061806>.

Remarks

You can use this option to secure communication packets between client applications and the database server using transport-layer security.

The `-ec` option instructs the database server to accept only connections that are encrypted using one of the specified types. You must specify at least one of the supported parameters in a comma-separated list. Connections over the TDS protocol, which include Java applications using jConnect, are always accepted and are never encrypted, regardless of the usage of the `-ec` option. Setting the TDS protocol option to NO disallows these unencrypted TDS connections.

By default, communication packets aren't encrypted, which poses a potential security risk. If you are concerned about the security of network packets, use the `-ec` option. Encryption affects performance only marginally.

If the database server accepts simple encryption, but does not accept unencrypted connections, then any non-TDS connection attempts using no encryption automatically use simple encryption.

Starting the database server with `-ec SIMPLE` tells the database server to only accept connections using simple encryption. TLS connections (RSA and RSA FIPS-certified encryption) fail, and connections requesting no encryption use simple encryption.

If you want the database server to accept encrypted connections over TCP/IP, but also want to be able to connect to the database from the local computer over shared memory, you can specify the `-es` option with the `-ec` option when starting the database server.

The `dbrsa16.dll` file contains the RSA code used for encryption and decryption. The file `dbfips16.dll` contains the code for the FIPS-certified RSA algorithm. When you connect to the database server, if the appropriate file cannot be found, or if an error occurs, a message appears in the database server messages window. The server doesn't start if the specified types of encryption cannot be initiated.

The client's and the server's encryption settings must match or the connection fails except in the following cases:

- If `-ec SIMPLE` is specified on the database server, but `-ec NONE` is not, then connections that do not request encryption can connect and automatically use simple encryption.
- If the database server specifies RSA and the client specifies FIPS-certified encryption, or vice versa, the connection succeeds. In these cases, the Encryption connection property returns the value specified by the database server.

Note: All strong encryption technologies are subject to export regulations.

Example

The following example specifies that connections with no encryption and simple encryption are allowed.

```
iqsrv16 -ec NONE,SIMPLE -x tcpip c:\myiqdemo.db
```

The following example starts a database server that uses the RSA server certificate `rsaserver.id`.

```
iqsrv16 -ec TLS(IDENTITY=rsaserver.id;IDENTITY_PASSWORD=test) -x  
tcpip c:\myiqdemo.db
```

The following example starts a database server that uses the FIPS-approved RSA server certificate `rsaserver.id`.

```
iqsrv16 -ec TLS(FIPS=Y;IDENTITY=rsaserver.id;IDENTITY_PASSWORD=test)  
-x tcpip c:\myiqdemo.db
```

-ep iqsrv16 database server option

Prompts the user for the encryption key upon starting a strongly encrypted database.

Syntax

```
iqsrv16 -ep ...
```

start_iq Database Server Startup Utility

Applies to

All operating systems and database servers.

Remarks

The `-ep` option instructs the database server to display a window where the user enters the encryption key for databases started on the command line that require an encryption key. This server option provides an extra measure of security by never allowing the encryption key to be seen in clear text.

When used with the database server, the user is prompted for the encryption key when the following are all true:

- the `-ep` option is specified
- the server is just starting up
- a key is required to start a database
- the database server is either not a Windows service, or it is a Windows service with the `interact with desktop` option turned ON
- the server isn't a daemon (Unix)

To secure communication packets between client applications and the database server, use the `-ec` server option and transport-layer security.

Example

The user is prompted for the encryption key when the `myencrypted.db` database is started:

```
iqsrv16 -ep -x tcpip myencrypted.db
```

-es iqsrv16 database server option

Allows unencrypted connections over shared memory.

Syntax

```
iqsrv16 -ec encryption-options -es ...
```

Applies to

All servers and operating systems.

Remarks

This option is only effective when specified with the `-ec` option. The `-es` option instructs the database server to allow unencrypted connections over shared memory. Connections over TCP/IP must use an encryption type specified by the `-ec` option. This option is useful in situations where you want remote clients to use encrypted connections, but for performance reasons you also want to access the database from the local computer with an unencrypted connection.

Example

The following example specifies that connections with simple encryption and unencrypted connections over shared memory are allowed.

```
iqsrv16 -ec SIMPLE -es -x tcpip c:\myiqdemo.db
```

-fips iqsrv16 database server option

Requires that only FIPS-certified algorithms can be used for strong database and communication encryption.

Syntax

```
iqsrv16 -fips ...
```

Applies to

Windows, Unix, and Linux.

Remarks

Specifying this option forces all database server encryption to use FIPS-certified algorithms. This option applies to strong database encryption, client/server transport-layer security, and web services transport-layer security. You can still use unencrypted connections and databases when the `-fips` option is specified, but you cannot use simple encryption.

Note: All strong encryption technologies are subject to export regulations.

For strong database encryption, the `-fips` option causes new databases to use the FIPS-certified equivalent of AES and AES256 if they are specified in the `ALGORITHM` clause of the `CREATE DATABASE` statement.

When the database server is started with `-fips`, you can run databases encrypted with AES, AES256, AES_FIPS, or AES256_FIPS encryption, but not databases encrypted with simple encryption. Unencrypted databases can also be started on the server when `-fips` is specified.

For transport-layer security, the `-fips` option causes the server to use the FIPS-certified RSA encryption algorithm, even if RSA is specified.

For transport-layer security for web services, the `-fips` option causes the server to use FIPS-certified HTTPS, even if HTTPS is specified.

When you specify `-fips`, the `ENCRYPT` and `HASH` functions use the FIPS-certified RSA encryption algorithm, and password hashing uses the SHA-256 FIPS algorithm rather than the SHA-256 algorithm.

-ga iqsrv16 database server option

Unloads the database after the last non-HTTP client connection disconnects.

Syntax

```
iqsrv16 -ga ...
```

Applies to

All operating systems.

Remarks

Specifying this option on the network server causes each database to be unloaded after the last non-HTTP client connection disconnects. In addition to unloading each database after the last non-HTTP connection disconnects, the database server shuts down when the last database is stopped.

If the only connection to a database is an HTTP connection, and the database is configured to stop automatically, when the HTTP connection disconnects, the database is not unloaded. As well, if you specify the -ga option, and the database has an HTTP connection and a command sequence or TDS connection, when the last command sequence or TDS connection disconnects, the database stops automatically, and any HTTP connections are dropped.

-gb iqsrv16 database server option

Sets the server process priority class.

Windows syntax

```
iqsrv16 -gb { idle | normal | high | maximum } ...
```

Unix syntax

```
iqsrv16 -gb level ...
```

Allowed values

- **Unix** – On Unix, the *level* is an integer from -20 to 19. The default value on Unix is the same as the nice value of the parent process. Lower *level* values represent a more favorable scheduling priority. All restrictions placed on setting a nice value apply to the -gb option. For example, on most Unix platforms, only the root user can lower the priority level of a process (for example, changing it from 0 to -1).
- **Windows** – On Windows, normal and high are the commonly-used settings. The value idle is provided for completeness. The value maximum may interfere with the running of your computer.

Applies to

Windows, Unix.

-gc iqsrv16 database server option

Sets the maximum interval between checkpoints.

Syntax

```
iqsrv16 -gc minutes ...
```

Default

60 minutes

Allowed values

- **minutes** – The default value is the setting of the checkpoint_time database option, which defaults to 60 minutes. If a value of 0 is entered, the default value of 60 minutes is used.

Applies to

All operating systems and database servers.

Remarks

Use this option to set the maximum length of time, in minutes, that the database server runs without doing a checkpoint on each database.

Checkpoints generally occur more frequently than the specified time.

-gd iqsrv16 database server option

Sets the privileges required to start or stop a database on a running database server.

Syntax

```
iqsrv16 -gd { DBA | all | none } ...
```

Allowed values

- **DBA** – Only users with the SERVER OPERATOR system privilege can start or stop databases.
- **all** – All users can start or stop databases. Not recommended for network servers that can be accessed by remote clients.
- **none** – Starting and stopping databases isn't allowed except when the database server itself is started and stopped.

Default

The default setting is DBA for the network database server. Both uppercase and lowercase syntax are allowed.

Applies to

All operating systems and database servers.

start_iq Database Server Startup Utility

Remarks

This option specifies the level of privilege required for a user to cause a new database file to be loaded by the database server, or to stop a database on a running database server.

When the option is set to DBA, the client application must use an existing connection to another database running on the same server to start and stop databases. You cannot start a database that is not already running by using the DatabaseFile connection parameter.

You can obtain the setting of the -gd option using the StartDBPermission server property:

```
SELECT PROPERTY ( 'StartDBPermission' );
```

The privileges for stopping a database server are specified by the -gk option.

Example

The following steps illustrate how to use the -gd option for the network database server.

1. Start the network database server:

```
iqsrv16 -su mypwd -gd DBA -n my_server
```

2. Connect to the utility database from Interactive SQL:

```
dbisql -c "UID=DBA;PWD=mypwd;DBN=utility_db"
```

3. Start a database:

```
START DATABASE 'iqdemo.db';
```

4. Connect to the database:

```
CONNECT USING 'DBN=demo;UID=DBA;PWD=sql';
```

-ge iqsrv16 database server option

Sets the stack size for external functions.

Syntax

```
iqsrv16 -ge integer ...
```

Default

32 KB

Applies to

Windows.

Remarks

Sets the stack size for threads running external functions, in bytes.

-gf iqsrv16 database server option

Disables firing of triggers by the server.

Syntax

```
iqsrv16 -gf ...
```

Applies to

All operating systems and database servers.

Remarks

The -gf server option instructs the server to disable the firing of triggers, including referential integrity triggers (such as cascading updates and deletes).

-gk iqsrv16 database server option

Sets the privileges required to stop the database server.

Syntax

```
iqsrv16 -gk { DBA | all | none } ...
```

Allowed values

- **DBA** – Only users with the SERVER OPERATOR system privilege can stop the database server. This is the default for the network server.
- **all** – No privileges are required to shut down the database server.
- **none** – The database server cannot be stopped.

Applies to

All operating systems and database servers.

Remarks

The -gd database server option applies to the dbstop utility as well as to the following statements:

- ALTER DATABASE *dbname* FORCE START statement.
- STOP DATABASE statement

-gl iqsrv16 Server Option

Set the permission required to load data using **LOAD TABLE**.

Syntax

-gl level

start_iq Database Server Startup Utility

Remarks

The **LOAD TABLE** statement reads files from the database server machine. To control access to the file system using these statements, the **-gl** command-line switch allows you to control the level of database permission that is required to use these statements. *level* is either:

- **DBA** – only users with the **LOAD ANY TABLE**, **ALTER ANY TABLE** or **ALTER ANY OBJECT** system privilege can load data.
- **ALL** – all users can load data.
- **NONE** – data cannot be loaded.

You can use either uppercase and lowercase syntax for the options.

The default settings are **all** for servers started with **start_iq** and **dba** for other servers. For consistency with earlier versions, use the **all** value on all systems. The **all** setting is used in the `iqdemo.cfg` and `default.cfg` configuration files.

-gm iqsrv16 database server option

Limits the number of concurrent connections to the database server.

Syntax

```
iqsrv16 -gm integer ...
```

Default

The default value is 50, but **start_iq** sets this value to 10. This number will be reduced by internal temporary connections utilized by the server during operation.

Applies to

All operating systems and database servers.

Remarks

Defines the connection limit for the server. If this number is greater than the number that is allowed under licensing and memory constraints, it has no effect. Computer resources typically limit the number of connections to a network server to a lower value than the default.

The database server allows one extra **DBA** connection above the connection limit to allow a user with the **DROP CONNECTION** system privilege to connect to the database server and drop other connections.

Maximum values are 1001 for UNIX/Linux and 200 for Windows.

-gn iqsrv16 Server Option

Sets the number of execution threads that are used for the catalog store and connectivity while running with multiple users.

Syntax

-gn *integer*

Remarks

This parameter applies to all operating systems and servers. Each connection uses a thread for each request, and when the request is completed, the thread is returned to the pool for use by other connections. As no connection can have more than one request in progress at one time, no connection uses more than one thread at a time.

An exception to this rule is if a Java application uses threads. Each thread in the Java application is a database server execution thread.

On Windows, specify this parameter in **start_iq**. To calculate its value, use:

```
gn_value >= gm_value * 1.5
```

Set the **-gn** value to at least 1.5 times the value of **-gm**. Specify a minimum of 25. The total number of threads cannot exceed a platform-specific maximum; see **-iqmt num** for details.

See also

- *-iqmt iqsrv16 Server Option* on page 164

-gp iqsrv16 database server option

Sets the maximum allowed database page size.

Syntax

```
iqsrv16 -gp { 2048 | 4096 | 8192 | 16384 | 32768 } ...
```

Default

4096 (if a database server is started with no databases loaded)

Applies to

All operating systems and database servers.

Remarks

Database files with a page size larger than the page size of the server cannot be loaded. This option explicitly sets the page size of the server, in bytes.

By default, the server page size is the same as the largest page size of the databases on the command line.

start_iq Database Server Startup Utility

On all platforms, if you do not use this option and start a server with no databases loaded, the default value is 4096.

-gr iqsrv16 Server Option

Sets the maximum length of time, in minutes, that the database server takes to recover from system failure.

Syntax

-gr *num*

-gss iqsrv16 Server Option

Sets, in part, the stack size for server execution threads that execute requests for server connections.

Syntax

-gss { *integer* | *integerK* | *integerM* }

Remarks

SAP Sybase IQ calculates the stack size of these server threads using the formula: (**-gss** + **-iqtss**). See **-iqtss**.

On Windows, the default value for **-gss** is 4MB. The maximum stack size is 256MB.

See also

- **-iqtss iqsrv16 Server Option** on page 167

-gt iqsrv16 database server option

Sets the maximum number of physical processors that can be used (up to the licensed maximum). This option is only useful on multiprocessor systems.

Syntax

```
iqsrv16 -gt num-processors ...
```

Allowed values

- **num-processors** – This integer can be a value between 1 and the minimum of:
 - the number of physical processors on the computer
 - the maximum number of CPUs that the server is licensed for if CPU-licensing is in effect

If the **-gt** value specified lies outside this range, the lower or upper limit is imposed.

Remarks

With per-seat licensing, the network database server uses all CPUs available on the computer. With CPU-based licensing, the network database server uses only the number of processors

you are licensed for. The number of CPUs that the network database server can use may also be restricted by your SAP Sybase IQ edition.

When you specify a value for the `-gt` option, the database server adjusts its affinity mask (if supported on that hardware platform) to restrict the database server to run on only that number of physical processors. If the database server is licensed for n processors, the server, by default, runs on all logical processors (hyperthreads and cores) of n physical processors. This behavior can be further restricted with the `-gtc` option.

The `-gt` option cannot be used with the `-gta` option.

-gtc iqsrv16 database server option

Controls the maximum processor concurrency that the database server allows.

Syntax

```
iqsrv16 -gtc logical-processors-to-use ...
```

Applies to

Linux, Solaris, and Windows operating systems executing on Intel-compatible x86 and x64 platforms.

Remarks

When you start the database server, the number of physical and logical processors detected by the database server appears in the database server messages window.

Physical processors are sometimes referred to as **packages** or **dies**, and are the CPUs of the computer. Additional logical processors exist when the physical processors support hyperthreading or are themselves configured as **multiprocessors** (usually referred to as multi-core processors). The operating system schedules threads on logical processors.

The `-gtc` option allows you to specify the number of logical processors that can be used by the database server. Its effect is to limit the number of database server threads that are created at server startup. This limits the number of active database server tasks that can execute concurrently at any one time. By default, the number of threads created is 1 + the number of logical processors on all licensed physical processors.

By default, the database server allows concurrent use of all logical processors (cores or hyperthreads) on each licensed physical processor.

For example, when using a one-CPU license on a two-CPU system where each CPU contains four cores with two threads per core, the network database server permits eight threads to run concurrently on one CPU and zero threads on the other. If the `-gtc` option is specified, and the number of logical processors to be used is less than the total available for the number of physical processors that are licensed, then the database server allocates logical processors based on round-robin assignment. Specifying 1 for the `-gtc` option implicitly disables intra-

start_iq Database Server Startup Utility

query parallelism (parallel processing of individual queries). Intra-query parallelism can also be explicitly limited or disabled using the `max_query_tasks` option.

The `-gtc` option cannot be used with the `-gta` option.

Example

Consider the following examples for a Windows-based SMP computer. In each case, assume a 4-processor system with two cores on each physical processor for a total of eight logical processors. The physical processors are identified with letters and the logical processors (cores in this case) are identified with numbers. This 4-processor system therefore has processing units A0, A1, B0, B1, C0, C1, D0, and D1.

Scenario	Network database server settings
A single CPU license or <code>-gt 1</code> specified	<code>-gt 1</code> <code>-gtc 2</code> <code>-gnh 20</code> Threads can execute on A0 and A1.
No licensing restrictions on the CPU with <code>-gtc 5</code> specified	<code>-gt 4</code> <code>-gtc 5</code> <code>-gnh 20</code> Threads can execute on A0, A1, B0, C0, and D0.
A database server with a three CPU license and <code>-gtc 5</code> specified	<code>-gt 3</code> <code>-gtc 5</code> <code>-gnh 20</code> Threads can execute on A0, A1, B0, B1, and C0.
No licensing restrictions on the CPU with <code>-gtc 1</code> specified	<code>-gt 4</code> <code>-gtc 1</code> <code>-gnh 20</code> Threads can execute only on A0.

-gu iqsrv16 database server option

Sets the privilege required for executing database file administration statements such as for creating or dropping databases.

Syntax

```
iqsrv16 -gu { all | none | DBA | utility_db } ...
```


Allowed values

-gu option	Effect	Applies to
all	<i>This option is deprecated.</i> Anyone can execute file administration statements.	Any database including utility database
none	Executing file administration statements is not allowed.	Any database including utility database
DBA	Only users with the SERVER OPERATOR system privilege can execute file administration statements	Any database including utility database
utility_db	Only the users who can connect to the utility database can execute file administration statements	Only the utility database

Default

DBA

Applies to

All operating systems and database servers.

Remarks

Restricts the users who can execute the following database file administration statements:

- ALTER DATABASE dbfile ALTER TRANSACTION LOG
- CREATE DATABASE statement
- CREATE DECRYPTED DATABASE statement
- CREATE DECRYPTED FILE statement
- CREATE ENCRYPTED DATABASE statement
- CREATE ENCRYPTED FILE statement
- DROP DATABASE statement
- RESTORE DATABASE statement.

When utility_db is specified, these statements can only be run from the utility database. When DBA is specified, these statements can only be run by a user with the SERVER OPERATOR system privilege. When none is specified, no user can execute these statements.

Examples

To prevent the use of the file administration statements, start the database server using the none privilege level of the -gu option. The following command starts a database server and names it

start_iq Database Server Startup Utility

TestSrv. It loads the mytestdb . db database, but prevents anyone from using that server to create or delete a database, or execute any other file administration statement regardless of their resource creation rights, or whether they can load and connect to the utility database.

```
iqsrv16 -n TestSrv -gu none c:\mytestdb.db
```

To permit only the users knowing the utility database password to execute file administration statements, start the server by running the following command.

```
iqsrv16 -n TestSrv -su secret -gu utility_db
```

The following command starts Interactive SQL as a client application, connects to the server named TestSrv, loads the utility database, and connects the user.

```
dbisql -c  
"UID=DBA;PWD=secret;DBN=utility_db;Host=host1;Server=TestSrv"
```

Having executed the above command successfully, the user connects to the utility database, and can execute file administration statements.

-iqfrec iqsrv16 Server Option

Open database in forced recovery mode.

Syntax

-iqfrec *dbname*

-iqgovern iqsrv16 Server Option

Sets the number of concurrent queries allowed by the server.

Syntax

-iqgovern *num*

Default

Setting of the RECOVERY_TIME database option, which defaults to 2 minutes.

Remarks

The number of concurrent queries is not the same as the number of connections. A single connection can have multiple open cursors. **-iqgovern** can help SAP Sybase IQ optimize paging of buffer data out to disk and avoid overcommitting memory. The default value of this switch is equal to 2 times the number of CPUs on your machine, plus 10. You may find that another value, such as 2 times the number of CPUs plus 4, provides better throughput, especially when large numbers of users are connected.

-iqlm iqsrv16 Server Option

Specifies the maximum amount of memory in MB that SAP Sybase IQ can dynamically request from the operating system for temporary use.

Syntax

-iqlm *size*

Default

2048MB

Remarks

Some load operations may require more large memory than the 2GB default provides. If memory requirements exceed the default, use the **-iqlm** startup option to increase the memory that SAP Sybase IQ can dynamically request from the OS. Set **-iqlm** as a switch as part of the command or configuration file that starts the server.

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

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

Note: Always specify the size value, without including the units of measurement. For example, specify **-iqlm 2048** rather than **-iqlm 2048MB**. If you specify the unit of measurement, **start_iq** ignores this switch.

-iqmc iqsrv16 Server Option

Specifies the main IQ store buffer cache size, in MB.

Syntax

-iqmc *size*

Default

64MB

Remarks

The switch overrides the default value of 64MB, and applies to all databases started from the time the server is started until the server is shut down. In other words, if you start one database at server startup and another later, you need 2 * **-iqmc** available for the main buffer cache.

Large memory requirements represent one third of all available physical memory. To ensure adequate memory for the main store, set the **-iqmc** startup parameter to one third of all available physical memory.

Always specify the size value, without including the units of measurement; for example, specify **-iqmc 32** rather than **-iqmc 32MB**. If you specify the unit of measurement, **start_iq** ignores this switch, unlike SQL Anywhere, which requires a unit of measurement.

Note: Do not run multiple databases with an SAP Sybase IQ server.

-iqmpx_failover iqsrv16 Server Option

Initiates multiplex coordinator failover to establish the designated failover Secondary node as the new coordinator.

Remarks

Starting the coordinator with this option has no effect. Users must be licensed for the Multiplex Grid Option to run secondary nodes. For `iqmpx_failover` values, see *Administration: Multiplex*.

-iqmpx_ov iqsrv16 Server Option

Performs multiplex configuration override for the current node.

Remarks

Used to change node properties during startup in the event that a node's location or other property has changed. Users must be licensed for the Multiplex Grid Option to run secondary nodes. For `iqmpx_ov` values, see *Administration: Multiplex*.

-iqmpx_reclaimwriterfreelist iqsrv16 Server Option

This option applies only while restarting a coordinator node.

Remarks

The coordinator will forcefully reclaim the free-list of the writer node identified by server-name. This switch is only needed in the event that a writer fails and cannot be restarted. Users must be licensed for the Multiplex Grid Option to run secondary nodes. For `iqmpx_reclaimwriterfreelist` values, see *Administration: Multiplex*.

-iqmpx_sn iqsrv16 Server Option

Runs the current node in multiplex in single-node mode.

Remarks

This mode is used exclusively for fixing problems with the multiplex configuration and should be used with extreme caution. Requires all other nodes in the multiplex to be shut down.

Recommended only for use on the coordinator node. Users must be licensed for the Multiplex Grid Option to run secondary nodes. For `iqmpx_sn` values, see *Administration: Multiplex*.

Multiplex.

-iqmsgnum iqsrv16 Server Option

Specifies the number of archives of the old message log maintained by the server.

Syntax

-iqmsgnum *num*

Remarks

Allowed values are integers 0 – 64 (inclusive). Default value is 0, which means that messages are wrapped in the main message log file. Takes effect only if **-iqmsgsz** or the **IQMsgMaxSize** server property is a value other than zero. The **IQMsgNumFiles** server property corresponds to **-iqmsgnum** and takes precedence over the value of **-iqmsgnum**.

A **-iqmsgnum** value *n* greater than 0 means that the server maintains *n* message log archives. For example, this command specifies that the server maintain 3 archives of the message log:

```
start_iq -n iqdemo iqdemo.db ... <other options> ... -iqmsgsz 100 -iqmsgnum 3
```

-iqmsgsz iqsrv16 Server Option

Limits the maximum size of the message log.

Syntax

-iqmsgsz *size*

Remarks

-iqmsgsz is an integer 0-2047 (inclusive), in MB. The default value is 0, which specifies that there is no limit on message log size. The **IQMsgMaxSize** server property corresponds to the **-iqmsgsz** server switch and takes precedence over the value of **-iqmsgsz**.

A **-iqmsgsz** value *n* greater than 0 means that the message log can grow up to *n* megabytes in size. For example, the following command limits the size of the message log to 100MB:

```
start_iq -n iqdemo iqdemo.db ... <other options> ... -iqmsgsz 100
```

-iqmt iqsrv16 Server Option

Specifies the number of SAP Sybase IQ threads to create.

Syntax

-iqmt *num*

Remarks

The default is 60 threads per CPU for the first 4 CPUs and 50 threads per CPU for the remainder, with 3 more for system use, plus threads needed for database connections and background tasks. For example, on a system with 12 CPUs and 10 connections: $60 * 4 + 50 * (\text{numCPUs} - 4) + \text{numConnections} + 6 = 656$.

The minimum value of *num* is `num_conn + 3`.

The total number of server threads cannot exceed 4096.

-iqnotemp iqsrv16 Server Option

Creates a temporary file in place of the defined temporary dbspace.

Syntax

-iqnotemp *size*

Remarks

size is file size, in MB. This parameter causes the server to ignore all SAP Sybase IQ temporary dbfile definitions when starting a database. You can use **-iqnotemp** to solve temporary dbfile problems by dropping damaged files and replacing them later.

If you start the utility database server with **-iqnotemp** when restoring, SAP Sybase IQ ignores all temporary file definitions in the backed up database during the restore operation, including intermediate starts. You can thus restore a database to a different temporary file topology without recreating and using old temporary file definitions.

The only temporary file operation allowed on a database started with **-iqnotemp** is **ALTER DBSPACE IQ_SYSTEM_TEMP DROP FILE**.

-iqnumbercpus iqsrv16 Server Option

Specifies the number of CPUs available to SAP Sybase IQ, overriding the physical number of CPUs, for resource planning purposes.

Syntax

-iqnumbercpus *num*

Remarks

The value of **-iqnumbercpus** defaults to the total number of CPUs, but the range of available values is 1 – 128.

Use **-iqnumbercpus** only on:

- Machines with Intel CPUs and hyperthreading enabled, setting **-iqnumbercpus** to the number of CPU cores available
- Machines where an operating system utility has been used to restrict SAP Sybase IQ to a subset of the CPUs within the machine

Setting **-iqnumbercpus** higher than the number of available CPUs may affect performance.

-iqpartition iqsrv16 Server Option

Sets the number of IQ main and temp buffer cache partitions.

Syntax

-iqpartition *num*

Remarks

Specifies the number of partitions in the IQ main and temp buffer caches. Must be a power of 2. By default, allowed values are: 0 (default), 1, 2, 4, 8, 16, 32, 64, 128, 256. By default, SAP Sybase IQ computes the number of partitions automatically as *number_of_cpus/8*, rounded to the nearest power of 2, up to a maximum of 64. You may be able to improve performance by adjusting the number of cache partitions. The **-iqpartition** switch sets this value for an SAP Sybase IQ server, and overrides the value set by the `Cache_Partitions` database option.

- Excludes jConnect™ for JDBC™ system objects from the database. To use the jConnect JDBC driver to access system catalog information, you need jConnect catalog support (installed by default). When you specify this option, you can still use JDBC, as long as you do not access system information. You can add jConnect support at a later time using the **ALTER DATABASE** statement.

-iqrlvmem start_iq Server Option

Specifies the amount of memory, in megabytes, available to the RLV store.

Syntax

-iqrlvmem *size*

Default

2048 (megabytes)

start_iq Database Server Startup Utility

Remarks

If you specify 0 or an invalid value, then the default (2048 MB) is used. If the value exceeds 2/3rd of the system virtual memory limit, and error message appears, and the server shuts down.

Usage

-iqrlvmem is used at server startup to tell the server how much memory to reserve for row-level versioning.

-iqro iqsrv16 Server Option

Runs the current node read-only.

Remarks

Used to verify simplex and multiplex migration and in legacy database unloads. Always specify the Boolean *I* (the only valid parameter) as an argument to this option.

-iqstart iqsrv16 Server Option

Provides startup diagnostics for dbspaces.

Syntax

-iqstart *N*

Remarks

The input parameter *N* is a number value that represents an integer bit mask. You may combine values to provide more than one feature. Output generated before the IQ message file is generated goes to the console. The **-z** startup switch provides additional startup and connection information.

The allowed values are as follows:

- **N=1**—returns basic information about the file names from `SYSIQFILES` that are used when opening the dbspace. It then displays the fully qualified names used. You can use this option to create a record of the files in use by the database in the IQ message file.
- **N=2**—stops after the transaction log replay before executing `RecoveryComplete` allowing you to examine the database without opening it all the way. You can combine **N=2** with other options. In certain modes using **N=2** may rewrite the `commit_identity`, but does not otherwise modify the database in a permanent manner—the checkpoint that would commit the recovery actions is not allowed to complete. All recovery actions reexecute the next time the database is opened.
- **N=4**—returns full diagnostic information, including all rows of `SYSIQFILE`, the subset of file names selected if the database is a multiplex database, the fully resolved file names, each individual dbspace file header block, the `database_identity`, the `commit_identity`, each checkpoint log entry, and each transaction log entry.

- N=8—allows the file paths in `SYSIQFILE` to be overridden. Instead of the `SYSIQFILE` values, the file names `iqmsg.iqmsg`, `iqmain_1`, `iqmain_2`, ..., `iqtemp_1`, `iqtemp_2`, and so on. will be used. These may be links and must be in the same directory as the `.db` file. You may use a link to the actual `.db` file but if the server is given a link to a `.db` file that uses a transaction log relative to the database, the server looks for the transaction log relative to the link, rather than to the database. In this case, create a link for the transaction log also

-iqtc iqsrv16 Server Option

Specifies IQ temporary store cache size, in MB.

Syntax

-iqtc *size*

Default

64MB

Remarks

The switch overrides the default value of 64MB, and applies to all databases started from the time the SAP Sybase IQ server is started until the server is shut down. In other words, if you start one database at server startup and another later, you need 2 * **-iqtc** available for the temp cache. In general, avoid running multiple databases with an SAP Sybase IQ server.

Large memory requirements represent one third of all available physical memory. To ensure adequate memory for the IQ temporary store cache, set the **-iqtc** startup parameter to one third of all available physical memory.

Always specify the size value, without including the units of measurement; for example, specify **-iqtc 32** instead of **-iqtc 32MB**. If you specify the unit of measurement, **start_iq** ignores this switch, unlike SQL Anywhere, which requires a unit of measurement.

Ensure that the `IQ_SYSTEM_TEMP` dbspace is at least as large as **-iqtc**.

-iqtss iqsrv16 Server Option

Specifies the stack size, in KB, for server execution threads running either in the background or as part of a thread team assisting the main server connection thread.

Syntax

-iqtss *size*

Remarks

The default is 512KB on 64-bit platforms, and 200KB on 32-bit platforms.

See also

- *-gss iqsrv16 Server Option* on page 156

-iqwmem iqsrv16 Server Option

Creates a pool of “wired” memory on HP and Sun UNIX systems.

Syntax

-iqwmem *size*

Remarks

This memory is locked down so it cannot be paged by the operating system. Specify the memory size, in MB. Use this switch only if you have enough memory to dedicate for this purpose. Otherwise, you may cause serious performance degradation.

-k iqsrv16 database server option

Controls the collection of Performance Monitor statistics.

Syntax

```
iqsrv16 -k ...
```

Default

Performance Monitor statistics are collected

Applies to

All operating systems and database servers.

Remarks

If you specify **-k** when you start the database server, then the database server does not collect Performance Monitor statistics. The **-k** option does not affect the collection of column statistics used by the query optimizer.

This option should only be used in situations where the database server is running on a multi-processor computer where it can be shown by testing to improve performance. For most workloads, the benefit will be negligible, so use of this option is not recommended. When you disable the performance counters, this information is not available for analyzing performance problems.

You can also change the setting for the collection of Performance Monitor statistics using the `sa_server_option` system procedure.

-kl iqsrv16 database server option

Specifies the file name of the Kerberos GSS-API library (or shared object on Unix) and enables Kerberos authenticated connections to the database server.

Syntax

```
iqsrv16 -kl GSS-API-library-file ...
```

Applies to

All operating systems.

Remarks

This option specifies the location and name of the Kerberos GSS-API. This option is only required if the Kerberos client uses a different file name for the Kerberos GSS-API library than the default, or if there are multiple GSS-API libraries installed on the computer running the database server. A Kerberos client must already be installed and configured, and SSPI cannot be used by the database server.

Specifying this option enables Kerberos authentication to the database server.

Example

The following command starts a database server that uses the `libgssapi_krb5.so` shared object for Kerberos authentication.

```
iqsrv16 -kl libgssapi_krb5.so -n my_server_princ /opt/myapp/
kerberos.db
```

-kr iqsrv16 database server option (deprecated)

Specifies the realm of the Kerberos server principal and enables Kerberos authenticated connections to the database server.

Note: The use of the `-kr` option is deprecated. It is recommended that you use the `-kp` option to specify the Kerberos server principal. When specifying `-kp` the server principal must have been extracted to the Kerberos keytab file on the computer running the database server.

Syntax

```
iqsrv16 -kr server-realm ...
```

Applies to

All operating systems.

Remarks

This option specifies the realm of the Kerberos server principal. Normally, the principal used by the database server for Kerberos authentication is `server-name@default-realm`, where `default-realm` is the default realm configured for the Kerberos client. Use this option if you want the server principal to use a different realm than the default realm, in which case the server principal used is `server-name@server-realm`.

Specifying this option enables Kerberos authentication to the database server.

The `-kr` option cannot be specified if the `-kp` option is specified.

Example

The following command starts a database server that accepts Kerberos logins and uses the principal `my_server_princ@MYREALM` for authentication.

```
iqsrv16 -kr MYREALM -n my_server_princ C:\kerberos.db
```

-krb iqsrv16 database server option

Enables Kerberos-authenticated connections to the database server.

Syntax

```
iqsrv16 -krb ...
```

Applies to

All operating systems.

Remarks

This option enables Kerberos authentication to the database server. You must specify one or more of the `-krb`, `-kl`, and `-kr` options for the database server to be able to authenticate clients using Kerberos.

Before you can use Kerberos authentication, a Kerberos client must already be installed and configured on both the client and database server computers. Additionally, the principal `server-name@REALM` must already exist in the Kerberos KDC, and the keytab for the principal `server-name@REALM` must already have been securely extracted to the keytab file on the database server computer. The database server will not start if the `-krb` option is specified, but this setup has not been performed.

Note: The database server name cannot contain any of the following characters: `/`, `\`, or `@`, and database server names with multibyte characters cannot be used with Kerberos.

The `login_mode` database option must be set to allow Kerberos logins, and Kerberos client principals must be mapped to database user IDs using the `GRANT KERBEROS LOGIN` statement.

Example

For a Kerberos principal for the database server named `my_server_princ@MYREALM`, the following command starts a database server named `my_server_princ`.

```
iqsrv16 -krb -n my_server_princ C:\kerberos.db
```

-ks iqsrv16 database server option

Disables the creation of shared memory that the Performance Monitor uses to collect counter values from the database server.

Syntax

```
iqsrv16 -ks 0 ...
```

Applies to

Windows

Remarks

When you specify this option, the Performance Monitor does not show any database server, database, or connection statistics for the current database server.

-ksc iqsrv16 database server option

Specifies the maximum number of connections that the Performance Monitor can monitor.

Syntax

```
iqsrv16 -ksc integer ...
```

Default

10

Applies to

Windows

-ksd iqsrv16 database server option

Specifies the maximum number of databases that the Performance Monitor can monitor.

Syntax

```
iqsrv16 -ksd integer ...
```

Default

2

Applies to

Windows

-m iqsrv16 database server option

Truncates the transaction log when a checkpoint is done.

Syntax

```
iqsrv16 -m ...
```

Applies to

All operating systems and database servers.

Remarks

This option truncates the transaction log when a checkpoint is done, either at shutdown or as a result of a checkpoint scheduled by the server.

Warning! When this option is selected, there is no protection against media failure on the device that contains the database files.

This option provides a way to automatically limit the growth of the transaction log. Checkpoint frequency is still controlled by the `checkpoint_time` and `recovery_time` options (which you can also set on the command line).

The `-m` option is useful for limiting the size of the transaction log in situations where high volume transactions requiring fast response times are being processed, and the contents of the transaction log aren't being relied upon for recovery or replication. The `-m` option provides an alternative to operating without a transaction log at all, in which case a checkpoint would be required following each COMMIT and performance would suffer as a result. When the `-m` option is specified, there is no protection against media failure on the device that contains the database files. Other alternatives for managing the transaction log (for example, using the BACKUP statement and events) should be considered before using the `-m` option.

To avoid database file fragmentation, it is recommended that where this option is used, the transaction log be placed on a separate device or partition from the database itself.

When this option is used, no operations can proceed while a checkpoint is in progress.

-n iqsrv16 database server option

Sets the name of the database server.

Syntax

```
iqsrv16 -n server-name database-filename ...
```

Default

The name of the first database file (with the path and extension removed) that is started on the database server.

Applies to

All operating systems and database servers.

Remarks

When a database server starts, it attempts to become the default database server on that computer. The first database server to start when there is no default server becomes the default database server. Shared memory connection attempts on that computer that do not explicitly specify a database server name connect to the default server.

Note: It is recommended that you use the `-xd` option for database servers being used by deployed applications, and that all clients explicitly specify the name of the database server to which they should connect by using the `ServerName (Server)` connection parameter. This ensures that the database connects to the correct database server when a computer is running multiple database servers.

There is no character set conversion performed on the server name. If the client character set and the database server character set are different, using extended characters in the server name can cause the server to not be found. If your clients and servers are running on different operating systems or locales, you should use 7-bit ASCII characters in the server name.

Database server names must be valid identifiers. Long database server names are truncated to different lengths depending on the protocol. Database server names cannot:

- begin with white space, single quotes, or double quotes
- end with white space
- contain semicolons, forward slashes (/), or backslashes (\)
- be longer than 250 bytes
- contain spaces when they are running on Unix

Note: On Windows and Unix, version 9.0.2 and earlier clients cannot connect to version 10.0.0 and later database servers with names longer than the following lengths:

- 40 bytes for Windows shared memory
- 31 bytes for Unix shared memory
- 40 bytes for TCP/IP

The server name specifies the name to be used in the `ServerName (Server)` connection parameter of client application connection strings or profiles. With shared memory, unless `-xd` is specified, a default database server is used if a server name is not specified and there is at least one database server running on the computer.

Running multiple database servers with the same name is not recommended.

Note: There are two `-n` options. The `-n` option is positional. If it appears before any database file names, it is a server option and names the server. If it appears after a database file name, it is a database option and names the database.

start_iq Database Server Startup Utility

For example, the following command names the database server `SERV` and the database `DATA`:

```
iqsrv16 -n SERV sales.db -n DATA
```

Database Server Naming Restrictions

If you use the `-n` switch in `start_iq [server-options]`, certain naming restrictions apply.

No character set conversion is performed on the server name. If the client character set and the database server character set differ, using extended characters in the server name can cause the server to not be found. If clients and servers run on different operating systems or locales, use 7-bit ASCII characters in the server name.

Database server names must be valid identifiers. Long database server names are truncated to different lengths depending on the protocol. Database server names cannot:

- Begin with white space, single quotes, or double quotes
- End with white space
- Contain semicolons
- Exceed 128 bytes

The server name specifies the name to be used on client application connection strings or profiles. Running multiple database servers with the same name is not recommended.

See also

- *start_iq Server Options* on page 135

-o iqsrv16 database server option

Prints all database server messages to the database server message log file.

Syntax

```
iqsrv16 -o filename ...
```

Applies to

All operating systems and database servers.

Remarks

Print all database server messages, including informational messages, errors, warnings, and MESSAGE statement output, to the specified file, and to the database server messages window. If you specify the `-qi` option with `-o`, all messages appear only in the database server message log file.

It is recommended that you do not end the file name with `.log` because this can create problems for utilities that perform operations using the transaction log.

To find the name of the database server message log file, execute the following statement:


```
SELECT PROPERTY ( 'ConsoleLogFile' );
```

start_iq Error Reporting

If there is a problem starting the server, **start_iq** returns a non zero value.

If you did not specify a log file after the **-o** switch on startup, the error is written to the first one of the following that is defined:

- (Windows) %ALLUSERSPROFILE%\SybaseIQ\16_0\logfiles\
- (UNIX / Linux) \$IQDIR16/logfiles/

-oe iqsrv16 database server option

Specifies a file name to log startup errors, fatal errors, and assertions.

Syntax

```
iqsrv16 -oe filename ...
```

Applies to

All operating systems and database servers.

Remarks

Each line in the output log file is prefixed with the date and time. Startup errors include such errors as:

- Couldn't open/read database file: *database file*.
- A database server with that name has already started.

Fatal errors and assertions are logged to the Windows Application Event Log or the Unix system log regardless of whether **-oe** is specified.

It is recommended that you do not end the file name with `.log` because this can create problems for utilities that perform operations using the transaction log.

-on iqsrv16 database server option

Specifies a maximum size for the database server message log, after which the file is renamed with the extension `.old` and a new file is started.

Syntax

```
iqsrv16 -on { size[ k | m | g ] } ...
```

Applies to

All operating systems and database servers.

Remarks

The *size* is the maximum file size for the database server message log, in bytes. Use *k*, *m*, or *g* to specify units of kilobytes, megabytes, or gigabytes respectively. The minimum size limit is 10 KB. By default, there is no maximum size limit.

When the database server message log reaches the specified size, the database server renames the file with the extension `.old`, and starts a new file with the original name.

Note: If the `.old` database server message log file already exists, it is overwritten. To avoid losing old database server message log files, use the `-os` option instead.

This option cannot be used with the `-os` option.

It is recommended that you do not end the database server message log file name with `.log` because this can create problems for utilities that perform operations using the transaction log.

-os iqsrv16 database server option

Specifies a maximum size for the database server message log file, at which point the file is renamed.

Syntax

```
iqsrv16 -os { size[ k | m | g ] } ...
```

Applies to

All operating systems and database servers.

Remarks

The *size* is the maximum file size for logging database server messages, in bytes. Use *k*, *m*, or *g* to specify units of kilobytes, megabytes, or gigabytes respectively. The minimum size limit is 10 KB. By default, there is no maximum size limit.

Before the database server logs output messages to the database server message log file, it checks the current file size. If the log message will make the file size exceed the specified size, the database server renames the database server message log file to `yyymmddxx.slg`, where *yyymmdd* represents the year, month, and day the file was created, and *xx* is a number that starts at 00 and continues incrementing.

This option allows you to identify old database server message log files that can be deleted to free up disk space.

This option cannot be used with the `-on` option.

It is recommended that you do not end the database server message log file name with `.log` because this can create problems for utilities that perform operations using the transaction log.

-ot iqsrv16 database server option

Truncates the database server message log file and appends output messages to it.

Syntax

```
iqsrv16 -ot logfile ...
```

Applies to

All operating systems and database servers.

Remarks

The functionality is the same as the -o option except the database server message log file is truncated before any messages are written to it. To find the name of the database server message log file, execute the following statement:

```
SELECT PROPERTY ( 'ConsoleLogFile' );
```

It is recommended that you do not end the database server message log file name with .log because this can create problems for utilities that perform operations using the transaction log.

-p iqsrv16 database server option

Sets the maximum size of communication packets.

Syntax

```
iqsrv16 -p integer ...
```

Default

7300 bytes

Applies to

All operating systems and database servers.

Remarks

The minimum value is 500 bytes and the maximum value is 65535.

You can change the communication buffer size for a connection by setting the CommBufferSize (CBSIZE) connection parameter.

-pc iqsrv16 database server option

Compresses all connections except for same-computer connections.

Syntax

```
iqsrv16 -pc ...
```

start_iq Database Server Startup Utility

Applies to

All operating systems and network servers, except web servers.

Remarks

The packets sent between a client and database server can be compressed using the `-pc` option. Compressing a connection may improve performance under some circumstances. Large data transfers with highly compressible data tend to get the best compression rates. This option can be overridden for a particular client by specifying `COMPRESS=NO` in the client's connection parameters.

By default, connections are not compressed. Specifying the `-pc` option compresses all connections except same-computer connections, web services connections, and TDS connections. TDS connections (including jConnect) do not support communication compression.

Same-computer connections over any communication link are not compressed, even if the `-pc` option or `COMPRESS=YES` connection parameter is used.

-pt iqsrv16 database server option

Increases or decreases the size limit at which packets are compressed.

Syntax

```
iqsrv16 -pt size ...
```

Default

120 bytes

Applies to

All operating systems and network servers.

Remarks

This parameter takes an integer value representing the minimum byte-size of packets to be compressed. Values less than 80 are not recommended.

Under some circumstances, changing the compression threshold can help performance of a compressed connection by allowing you to compress packets only when compression will increase the speed at which the packets are transferred. The default setting should be appropriate for most cases.

If both client and server specify different compression threshold settings, the client setting applies.

-qi iqsrv16 database server option

Controls whether database server system tray icon and database server messages window appear.

Syntax

```
iqsrv16 -qi ...
```

Applies to

Windows

Remarks

This option leaves no visual indication that the server is running, other than possible startup error windows. You can use either (or both) the -o or -oe log files to diagnose errors.

-qp iqsrv16 database server option

Specifies that messages about performance do not appear in the database server messages window.

Syntax

```
iqsrv16 -qp ...
```

Applies to

All operating systems and database servers.

Remarks

Do not display messages about performance in the database server messages window.

Messages that are suppressed include the following:

- No unique index or primary key for table '*table-name*'
- Database file "*mydatabase.db*" consists of *nnn* fragments

-qs iqsrv16 database server option

Suppresses startup error windows.

Syntax

```
iqsrv16 -qs ...
```

Applies to

Windows

start_iq Database Server Startup Utility

Remarks

This option suppresses startup error windows. Examples of startup errors include the database server not being able to open or read a database file or a database server not starting because another database server with the specified name is already running.

On Windows platforms, if the server isn't being started automatically, these errors appear in a window and must be cleared before the server stops. These windows do not appear if the -qs option is used.

If there is an error loading the language DLL, no window appears if -qs was specified on the command line and not in @data. This error isn't logged to the -o or -oe logs, but rather to the Windows Application Event Log

Usage errors are suppressed if -qs is on the command line, but not in @data expansion.

-qw iqsrv16 database server option

Specifies that the database server messages window does not appear.

Syntax

```
iqsrv16 -qw ...
```

Applies to

All operating systems and database servers.

Remarks

This option suppresses the database server messages window. On Windows platforms, the database server system tray icon is still visible. You can use either (or both) the -o or -oe log files to diagnose errors.

-s iqsrv16 database server option

Sets the user ID for Syslog messages.

Syntax

```
iqsrv16 -s { none | user | daemon | localn } ...
```

Applies to

Unix

Remarks

Sets the system user ID used in messages to the Syslog facility. The default is user for database servers that are started in the foreground, and daemon for those that are run in the background (for example, started by dbspawn, started automatically by a client, or started with the -ud database server option).

A value of none prevents any Syslog messages from being logged. The `localn` argument allows you to use a facility identifier to redirect messages to a file. You can specify a number between 0 and 7, inclusive, for *n*. For more information, refer to the Unix Syslog(3) man page.

-sb iqsrv16 database server option

Specifies how the database server reacts to broadcasts.

Syntax

```
iqsrv16 -sb { 0 | 1 } ...
```

Applies to

TCP/IP.

Remarks

The database server starts one or more UDP listeners so that it can respond to three types of broadcasts:

1. client connection broadcasts looking for this server.
2. database server enumeration broadcasts (such as those from the `dblocate` utility or the **Find Servers Wizard** of the administration tools **Connect** window).
3. broadcasts sent by another database server that is looking for any database servers with the same name.

For embedded SQL connections, the `dblocate` utility and `db_locate_servers` function cause broadcast packets to be sent out on local networks in an attempt to find all database servers. The UDP listener within the database server then responds back to the sender with information about how to connect to the database server.

The `-sb` option controls the behavior of the UDP listeners within the database server.

When the `-sb` option is not specified, the database server responds to all three types of requests.

If you specify `-sb 0`, then the database server responds to (3) only. This forces clients to use a Host connection parameter or `HOST` protocol option when connecting to the database server. In addition if the server is using a port other than the default port (2638), then the client must also specify the server's port. Since the database server does not respond to database server enumeration broadcasts, the database server is not included in the output of `dblocate`.

If you specify `-sb 1`, then the database server responds to (1) and (3). Since the database server does not respond to database server enumeration broadcasts, the database server is not included in the output of `dblocate`.

-sf iqsrv16 database server option

Controls whether users have access to features for databases running on the current database server. A secured feature can only be accessed by a user with appropriate privileges, while an unsecured feature can be accessed by all users.

Syntax

```
iqsrv16 -sf feature-list ...
```

feature-list :

```
feature-name | feature-set [ , feature-name | feature-set ] ...
```

Feature set	Included features (feature sets in bold)
none	All features are unsecured except manage_features , manage_keys , and disk_sandbox .
manage_server	processor_affinity
manage_security	manage_features manage_keys manage_disk_sandbox
server_security	disk_sandbox trace_system_event

Feature set	Included features (feature sets in bold)
all	<p>client –</p> <p>read_client_file write_client_file</p> <p>remote –</p> <p>remote_data_access send_udp send_email web_service_client</p> <p>local –</p> <ul style="list-style-type: none"> • local_call – cmdshell external_procedure java • local_db – backup restore database dbspace • local_env – getenv • local_io – create_trace_file read_file write_file directory sp_list_directory sp_create_directory sp_copy_directory sp_move_directory sp_delete_directory sp_copy_file

Feature set	Included features (feature sets in bold)
	sp_move_file sp_delete_file • local_log – request_log console_log webclient_log

Parameters

- **none** – Specifies that no features are secured.
- **manage_server** – Prevents users from accessing all database server-related features. This set consists of the following features:
 - **processor_affinity** – Prevents users from changing the processor affinity (the number of logical processors being used) of the database server.
 - **manage_security** – Prevents users from accessing features that allow the management of database server security. By default, these features are secured.
 - **manage_features** – Prevents users from modifying the list of features that can be secured on the database server.
 - **manage_keys** – Prevents the creation, modification, deletion, or listing of secure feature keys.

A user that has access to the manage_keys feature but not the manage_features feature cannot define a key with more secure features than those assigned to the user.
 - **manage_disk_sandbox** – Prevents users from temporarily changing disk sandbox settings by using the sa_server_option system procedure or the sa_db_option system procedure. The manage_disk_sandbox secure feature cannot be turned off for all databases or users—it can only be turned off for individual connections by using the sp_use_secure_feature_key system procedure.
- **server_security** – Prevents users from accessing features that can temporarily bypass security settings. By default, these features are secured.
 - **disk_sandbox** – Prevents users from performing read-write file operations on the database outside the directory where the main database file is located.
 - **trace_system_event** – Prevents users from creating user-defined trace events.
- **all** – Prevents users from accessing the following groups:
 - **client** – Prevents users from accessing all features that allow access to client-related input and output. This feature controls access to the client computing environment. This set consists of the following features:

- **read_client_file** – Prevents the use of statements that can cause a client file to be read. For example, the READ_CLIENT_FILE function and the LOAD TABLE statement.
- **write_client_file** – Prevents the use of all statements that can cause a client file to be written to. For example, the UNLOAD statement and the WRITE_CLIENT_FILE function.
- **remote** – Prevents users from accessing all features that allow remote access or communication with remote processes. This set consists of the following features:
 - **remote_data_access** – Prevents the use of any remote data access services, such as proxy tables.
 - **send_udp** – Prevents the ability to send UDP packets to a specified address by using the sa_send_udp system procedure.
 - **send_email** – Prevents the use of email system procedures, such as xp_sendmail.
 - **web_service_client** – Prevents the use of web service client stored procedure calls (stored procedures that issue HTTP requests).
- **local** – Prevents users from accessing all local-related features. This feature controls access to the server computing environment. This set consists of the local_call, local_db, local_io, and local_log feature subsets.
 - **local_call** – Prevents users from accessing all features that provide the ability to execute code that is not directly part of the database server and is not controlled by the database server. This set consists of the following features:
 - **cmdshell** – Prevents the use of the xp_cmdshell procedure.
 - **external_procedure** – Prevents the use of external stored procedures. This setting does not disable the use of the xp_* system procedures (such as xp_cmdshell, xp_readfile, and so on) that are built into the database server. Separate feature control options are provided for these system procedures.
 - **external_procedure_v3 - External C or C++ procedure. See the *User-Defined Functions* guide for information on C and C++ UDFs.** – See the User-Defined Functions guide.
 - **java** – Prevents the use of Java-related features, such as Java procedures.
- **local_db** – Prevents users from accessing all features related to database files. This set consists of the following features:
 - **backup** – Prevents the use of the BACKUP statement, and with it, the ability to run server-side backups. You can still perform client-side backups by using the dbbackup utility.
 - **restore** – Prevents the use of the RESTORE DATABASE statement.
 - **database** – Prevents the use of the CREATE DATABASE, ALTER DATABASE, DROP DATABASE, CREATE ENCRYPTED FILE, CREATE

DECRYPTED FILE, CREATE ENCRYPTED DATABASE, and CREATE DECRYPTED DATABASE statements.

- **dbspace** – Prevents the use of the CREATE DBSPACE, ALTER DBSPACE, and DROP DBSPACE statements.
- **local_env** – Prevents users from accessing all features related to environment variables. This set consists of the following features:
 - **getenv** – Prevents users from reading the value of any environment variable.
- **local_io** – Prevents users from accessing all features that allow direct access to files and their contents. This set consists of the following features:
 - **create_trace_file** – Prevents the use of statements that create an event tracing target.
 - **read_file** – Prevents the use of statements that can cause a local file to be read. For example, the xp_read_file system procedure, the LOAD TABLE statement, and the use of OPENSTRING(FILE...). The alternate names load_table and xp_read_file are deprecated.
 - **write_file** – Prevents the use of all statements that can cause a local file to be written to. For example, the UNLOAD statement and the xp_write_file system procedure. The alternate names unload_table and xp_write_file are deprecated.
 - **delete_file** – Prevents the use of all statements that can cause a local file to be deleted. For example, securing this feature causes the dbbackup utility to fail if the -x or -xo options are specified.
 - **directory** – Prevents the use of directory class proxy tables. This feature is disabled when remote_data_access is disabled.
 - **sp_list_directory** – Prevents the use of the sp_list_directory system procedure.
 - **sp_create_directory** – Prevents the use of the sp_create_directory system procedure.
 - **sp_copy_directory** – Prevents the use of the sp_copy_directory system procedure.
 - **sp_move_directory** – Prevents the use of the sp_move_directory system procedure.
 - **sp_delete_directory** – Prevents the use of the sp_delete_directory system procedure.
 - **sp_copy_file** – Prevents the use of the sp_copy_file system procedure.
 - **sp_move_file** – Prevents the use of the sp_move_file system procedure.
 - **sp_delete_file** – Prevents the use of the sp_delete_file system procedure.
- **local_log** – Prevents users from accessing all logging features that result in creating or writing data directly to a file on disk. This set consists of the following features:
 - **request_log** – Prevents the ability to change the request log file name and also prevents the ability to increase the limits of the request log file size or number of files. You can specify the request log file and limits on this file in the command

to start the database server; however, they cannot be changed once the database server is started. When request log features are disabled, you can still turn request logging on and off and reduce the maximum file size and number of request logging files.

- **console_log** – Prevents the ability to change the database server message log file name using the ConsoleLogFile option of the sa_server_option system procedure. Securing this feature also prevents the ability to increase the maximum size of the log file using the ConsoleLogMaxSize option of the sa_server_option system procedure. You can specify a server log file and its size when starting the database server.
- **webclient_log** – Prevents the ability to change the web service client log file name using the WebClientLogFile option of the sa_server_option system procedure. You can specify a web service client log file when starting the database server.

Applies to

All operating systems and database servers.

Remarks

This option allows the owner of the database server to control whether users have access to features for databases running on the database server. The -sk option allows the owner of the database server to create a system secure feature key that prevents users from accessing features specified by the -sf option.

If you start a database without specifying a system secure feature key, the default secure features are secured, and you cannot change the secure feature settings for the database server or any databases running on it. You cannot create the system secure feature key later—you must shut down the database server and specify a system secure feature key when you restart it.

The *feature-list* is a comma-separated list of feature names or feature sets to secure for the database server. Securing a feature makes it inaccessible to all database users other than administrators. Specifying a feature set secures all the features included in the set. To secure one or more, but not all, of the features in the feature set, specify the individual feature name.

Note: Sub-features of feature sets that are secured by default, cannot be unsecured from the command line. In other words the following command will not work:

```
-sf manage_security, -manage_keys
```

Use *feature-name* to indicate that the feature should be secured (made inaccessible), and -*feature-name* or *feature-name-* to indicate that the feature should be unsecured (accessible to all database users). For example, the following command indicates that only dbspace features are accessible to all users:

```
iqsrv16 -n secure_server -sf all,-dbspace
```

start_iq Database Server Startup Utility

Example

The following command starts a database server named `secure_server` with access to the request log and with all remote data access features secured. The key specified by the `-sk` option can be used later with the `sp_use_secure_feature_key` system procedure to make these features accessible to all users on the current connection.

```
iqsrv16 -n secure_server -sf remote,-request_log -sk j978kls12
```

If a user connected to a database running on the `secure_server` database server uses the `sp_use_secure_feature_key` system procedure with the `authorization_key` parameter set to the same value as that specified by `-sk`, that connection has access to the remote data access features:

```
CALL sp_use_secure_feature_key ( 'MyKey' , 'j978kls12' );
```

The following command secures all features, with the exception of local database features:

```
iqsrv16 -n secure_server -sf all,-local_db
```

-sk iqsrv16 database server option

Specifies a system secure feature key that can be used to allow access to features that are secured for the database server.

Syntax

```
iqsrv16 -sk key ...
```

Applies to

All operating systems and database servers.

Remarks

When you secure features for a database server by using the `-sf` option, you can also include the `-sk` option, which specifies a key that can be used with the `sp_use_secure_feature_key` system procedure to allow access to secured features for a connection. That connection can also use the `sa_server_option` system procedure to modify the features or feature sets that are secured for all databases running on the database server.

The key must be a non-empty string of at least six characters, and it cannot contain double quotes, control characters (any character less than 0x20), or backslashes. There is a limit of 1000 secure feature keys per database.

If the value for the `authorization_key` parameter of the `sp_use_secure_feature_key` system procedure is set to any value other than the one specified by `-sk`, no error is given and the features specified by `-sf` remain secured for the connection.

If you specify `-sk` without `-sf`, only the default secure features are enabled, but you can use the system secure feature key while the database server is running to change the secure feature settings.

Example

The following command starts a database server named `secure_server` with the backup feature secured. The key specified by the `-sk` option can be used later to allow access to these features for a specific connection.

```
iqsrv16 -n secure_server -sf backup -sk j978kls12
```

Setting the `authorization_key` parameter to the value specified by `-sk` for a connection to a database running on the `secure_server` database server allows that connection to perform backups or change the features that are secured on the `secure_server` database server:

```
CALL sp_use_secure_feature_key ( 'MyKey' , 'j978kls12' );
```

The user can then secure all features for databases running on `secure_server` by executing the following statement:

```
CALL sa_server_option( 'SecureFeatures', 'all' );
```

-su iqsrv16 database server option

Sets the password for the DBA user of the utility database (`utility_db`), or disable connections to the utility database.

Syntax

```
iqsrv16 -su password ...
```

Applies to

All operating systems and database servers.

Remarks

This option specifies the initial password for the DBA user of the utility database. The password is case sensitive. You can specify `none` for the password to disable all connections to the utility database.

- If you are using the network database server and do not specify the `-su` option, connections to the utility database are not allowed unless the `util_db.ini` file exists and the user ID is DBA with a password that matches the password in the `util_db.ini` file. On a network server, if both `-su` and `util_db.ini` are used, `util_db.ini` is ignored. The `util_db.ini` file is deprecated.

Specifying a password for the utility database for a network database server is useful for the cases when the database server is running, but it is not possible to connect to the database.

You can execute a `CREATE USER DBA IDENTIFIED BY new-password` statement while connected to `utility_db` to change the password for the DBA user of the utility database. The `REVOKE CONNECT FROM DBA` statement can be used to disable connections to the `utility_db` database. Not all SQL statements are supported for the utility database.

start_iq Database Server Startup Utility

To avoid having the utility database password in clear text on the command line, you can use the `dbfhide` utility to obfuscate a file containing the password, and then reference the obfuscated file on the command line.

Example

The following command disables all connections to the utility database:

```
iqsrv16 -su none c:\inventory.db
```

In the following example, the file named `util_db_pwd.cfg` that contains the utility database password is obfuscated using `dbfhide` and renamed `util_db_pwd_hide.cfg`:

```
dbfhide util_db_pwd.cfg util_db_pwd_hide.cfg
```

The `util_db_pwd_hide.cfg` file can then be used to specify the utility database password:

```
iqsrv16 -su @util_db_pwd_hide.cfg -n my_server c:\inventory.db
```

-tdsl iqsrv16 database server option

Sets the TDS login mode.

Syntax

```
iqsrv16 -tdsl { all | RSA | RSANonce }...
```

Default

all

Applies to

All operating systems and database servers.

Remarks

This option restricts the type of TDS login requests that the database server supports.

Mode type	Description
all	Encrypted and unencrypted passwords are allowed (including those with a nonce) in TDS login requests:
RSA	Only RSA encrypted passwords (including those with a nonce) are allowed in TDS login requests:
RSANonce	Only RSA encrypted passwords with a nonce are allowed in TDS login requests:

When you make login requests from a TDS application that supports RSA without a nonce, the database server generates a new set of encryption key for the login requests. Generating new encryption keys can be time-consuming.

When you make login requests from a TDS application that supports RSA with a nonce, the database server reuses a set of RSA encryption keys. These encryption keys are re-generated every 24 hours. By reusing the RSA encryption keys, performance can improve while protecting the database server from replay attacks.

Both jConnect and Open Client support RSA login requests with and without a nonce.

-ti iqsrv16 database server option

Disconnects inactive connections.

Syntax

```
iqsrv16 -ti minutes ...
```

Default

240 (4 hours)

Applies to

All operating systems and database servers.

Remarks

Disconnects connections that haven't submitted a request for the specified number of *minutes*. The maximum value is 32767. A client computer in the middle of a database transaction holds locks until the transaction is ended or the connection is disconnected. The -ti option is provided to disconnect inactive connections, freeing their locks.

The -ti option is very useful when used in conjunction with iqsrv16 since most connections will be over network links (TCP).

Setting the value to zero disables checking of inactive connections, so that no connections are disconnected. If the Idle connection parameter is not used, then the idle timeout value for TCP/IP connections is controlled by the -ti database server option. If both the -ti database server option and the Idle connection parameter are specified, then the idle timeout value is controlled by the connection parameter.

-tl iqsrv16 database server option

Sets the period at which to send liveness packets.

Syntax

```
iqsrv16 -tl seconds ...
```

Applies to

All database servers using TCP/IP.

Remarks

A liveness packet is sent periodically across a client/server TCP/IP communications protocol to confirm that a connection is intact. If the server runs for a LivenessTimeout period (default 2 minutes) without detecting a liveness packet on a connection, the communication is severed, and the server drops the connection associated with that client. Unix non-threaded clients and TDS connections do not do liveness checking.

The -tl option on the server sets the LivenessTimeout value for all clients that do not specify a liveness period.

Liveness packets are sent when a connection hasn't sent any packets for between one third and two thirds of the LivenessTimeout value.

When there are more than 200 connections, the server automatically calculates a higher LivenessTimeout value based on the stated LivenessTimeout value, so the server can handle a large number of connections more efficiently. Liveness packets are sent between one third and two thirds of the LivenessTimeout on each idle connection. Large numbers of liveness packets aren't sent at the same time. Liveness packets that take a long time to send could be sent after two thirds of the LivenessTimeout. A warning appears in the database server message log if the liveness sends take a long time. If this warning occurs, consider increasing the LivenessTimeout value.

Although it isn't generally recommended, you can disable liveness by specifying the following:

```
iqsrv16 -tl 0 -n my_server
```

Rather than disabling the LivenessTimeout option, consider increasing the value to one hour as follows:

```
iqsrv16 -tl 3600 -n my_server
```

-tmf iqsrv16 database server option

Forces transaction manager recovery for distributed transactions.

Syntax

```
iqsrv16 -tmf ...
```

Applies to
Windows.

Remarks

Used during recovery of distributed transactions when the distributed transaction coordinator isn't available. If DTC cannot be located, the outstanding operations are rolled back and recovery continues. It can also be used when starting a database with distributed transactions in the transaction log on a platform where the DTC isn't available.

Warning! If you use this option, distributed transactions are not recovered properly. It is not intended for routine use.

-tq iqsrv16 database server option

Shuts down the server at a specified time.

Syntax

```
iqsrv16 -tq { datetime | time } ...
```

Applies to

All operating systems and database servers.

Remarks

This option is useful for setting up automatic off-line backup procedures.

The format for the time is in *hh.mm* (24 hour clock), and can be preceded by an optional date. If a date is specified, the date and time must be enclosed in double quotes and be in the format *YYYY/MM/DD HH.MM*.

-u iqsrv16 database server option

Opens files using the operating system disk cache.

Syntax

```
iqsrv16 -u ...
```

Applies to

Windows, Unix.

Remarks

Files are opened using the operating system disk cache in addition to the database cache.

While the operating system disk cache may improve performance, better performance is often obtained by using only the database cache.

If the server is running on a dedicated computer, you shouldn't use the -u option, as the database cache itself is generally more efficient. You may want to use the -u option if the server is running on a computer with several other applications (so that a large database cache may interfere with other applications) and yet IO-intensive tasks are run intermittently on the server (so that a large cache will improve performance).

-ud iqsrv16 Server Option

(UNIX servers) Causes the process to run as a daemon in the root directory. SAP recommends that you do not use this switch in IQ servers.

-uf iqsrv16 database server option

Specifies the action that the database server takes when a fatal error or assertion failure occurs on a database server.

Syntax

```
iqsrv16 -uf action ...
```

Default

default

Allowed values

- **abort** – The database server is shut down and a core file is generated.
- **default** – The database server behaves in the same manner as abort, except when a device-full fatal error occurs. In this case, it behaves in the same manner as defunct. This action prevents the system from trying to write a core file on a full device.
- **defunct** – The database server continues running and does not call abort. Any new connection attempts that are made to the database server receive the original error.

Applies to

All operating systems.

-ut iqsrv16 Server Option

(UNIX servers) Causes the server to touch catalog store temporary files at intervals specified by *min*.

Syntax

-ut *min*

-v iqsrv16 Server Option

Displays the database server version in a message box (Windows) or in a version string (UNIX / Linux).

Syntax

-v or **-v2**

-x iqsrv16 database server option

Specifies server-side network communications protocols.

Syntax

```
iqsrv16 -x { all | none | tcpip [ ( parm=value;... ) ] ...
```

Allowed values

Regardless of which settings you choose for the `-x` option, the database server always accepts shared memory connections.

By default, the network database server starts the shared memory and TCP/IP protocols.

You can specify the following values for the `-x` option:

- **TCPIP (TCP)** – Listen for connection attempts by the client using the shared memory and TCP/IP protocols.
- **NONE** – Listen for connection attempts by the client using only the shared memory protocol.
- **ALL** – Listen for connection attempts by the client using shared memory and TCP/IP protocols.

Applies to

All operating systems and database servers.

Remarks

Use the `-x` option to specify which communications protocols you want to use to listen for client connections.

For the TCP/IP protocol, additional parameters may be provided, in the following format:

```
-x tcpip(PARM1=value1;PARM2=value2;...)
```

For Unix, quotation marks are required if more than one parameter is supplied or certain punctuation characters are used, such as round brackets. For example:

```
-x "tcpip(PARM1=value1;PARM2=value2;...)"
```

When the database server listens for TCP/IP connections, by default it listens to all network cards on port 2638.

Examples

Allow shared memory and TCP/IP connections on the network server:

```
iqsrv16 -n server_name "%IQDIRSAMP16%\iqdemo.db;"
```

Use the `-x` option to specify protocol options to tune the behavior of TCP/IP. Allow the network server to use two specific network cards:

```
iqsrv16 -x "tcpip(MyIP=192.75.209.12,192.75.209.32)" "%IQDIRSAMP16%\iqdemo.db;"
```

-xd iqsrv16 Server Option

Prevents the database server from becoming the default database server.

When a database server starts, it attempts to become the default database server on that computer. The first database server to start when there is no default server becomes the default

start_iq Database Server Startup Utility

database server. Shared memory connection attempts on that computer that do not explicitly specify a database server name connect to the default server.

Specifying this option prevents the database server from becoming the default database server. If this option is specified, clients that do not specify a database server name cannot find the database server over shared memory. The **-xd** option also prevents the database server from using the default TCP port. If a TCP port is not specified, the database server uses a port other than port 2638.

-xs iqsrv16 database server option

Specifies server-side web services communications protocols.

Syntax

```
iqsrv16 -xs [,...] { protocol[,...] }
```

```
protocol : {  
NONE  
| HTTP [ ( option=value;... ) ]  
| HTTPS [ ( option=value;... ) ]
```

Allowed values

You can specify any of the following:

- **option** – A network protocol option
- **HTTP** – Listen for web requests by the client using the HTTP protocol. The default port on which to listen is 80.
- **HTTPS** – Listen for web requests by the client using the HTTPS protocol. The default port on which to listen is 443. You must specify the server's certificate and password to use HTTPS. The password must be an RSA certificate because HTTPS uses RSA encryption.

The HTTP server supports HTTPS connections using SSL version 3.0 and TLS versions 1.0 and 1.1.

You can specify **HTTPS**, or **HTTPS** with **FIPS=Y** for FIPS-certified RSA encryption. FIPS-certified HTTPS uses a separate certified library, but is compatible with HTTPS.

- **NONE** – Do not listen for web requests. This is the default.

Applies to

All operating systems and database servers.

Remarks

Use the **-xs** option to specify which web protocols you want to use to listen for requests.

If you do not specify the **-xs** option, the database server doesn't attempt to listen for web requests.

To specify multiple protocols, you can specify the `-xs` option for each protocol or you can specify multiple protocols with one `-xs` option. The server listens for web requests using all the specified protocols.

Note: To start multiple web servers at the same time, change the port for one of them since they both have the same default port.

You can use the HTTPS or FIPS-certified HTTPS protocols for transport-layer security.

On Unix, quotation marks are required if more than one parameter is supplied:

```
-xs "HTTP (OPTION1=value1; OPTION2=value2; ...)"
```

Note: All strong encryption technologies are subject to export regulations.

Example

Listen for HTTP web requests on port 80:

```
iqsrv16 web.db -xs HTTP (PORT=80)
```

Listen for web requests using HTTPS:

```
iqsrv16 web.db -xs  
HTTPS (FIPS=N; PORT=82; IDENTITY=ecserver.id; IDENTITY_PASSWORD=test)
```

Listen on ports 80 and 8080:

```
iqsrv16 -xs HTTP (port=80) , HTTP (PORT=8080)
```

or

```
iqsrv16 -xs HTTP (port=80) -xs HTTP (PORT=8080)
```

-z iqsrv16 database server option

Displays diagnostic communication messages, and other messages, for troubleshooting purposes.

Syntax

```
iqsrv16 -z ...
```

Applies to

All operating systems and database servers.

Remarks

This should only be used when tracking problems. The information appears in the database server messages window.

-ze iqsrv16 database server option

Displays database server environment variables in the database server messages window.

Syntax

```
iqsrv16 -ze ...
```

Applies to

All operating systems and database servers.

Remarks

When you specify the `-ze` option, environment variables are listed in the database server messages window on startup. You can log the contents of the database server messages window to a file by specifying the `-o` option when starting the database server.

Example

The following command starts a database server named `myserver`, and outputs the environment variables set for the server to the database server messages window and the file `server-log.txt`.

```
iqsrv16 -n myserver -ze -o server-log.txt
```

-zl iqsrv16 database server option

Turns on capturing of the most recently-prepared SQL statement for each connection to databases on the server.

Syntax

```
iqsrv16 -zl ...
```

Applies to

All operating systems and database servers.

Remarks

This feature can also be turned on using the `RememberLastStatement` server setting. You can obtain the most recently-prepared SQL statement for a connection using the `LastStatement` value of the `CONNECTION_PROPERTY` function. The `sa_conn_activity` stored procedure allows you to obtain the most recently-prepared SQL statement for all current connections to databases on the server.

The `LastStatement` value is set when a statement is prepared, and is cleared when a statement is dropped. Only one statement string is remembered for each connection.

If `sa_conn_activity` reports a non-empty value for a connection, it is most likely the statement that the connection is currently executing. If the statement had completed, it would likely have been dropped and the property value would have been cleared. If an application prepares

multiple statements and retains their statement handles, the LastStatement value does not reflect what a connection is currently doing.

For stored procedure calls, only the outermost procedure call appears, not the statements within the procedure.

Warning! When `-zl` is specified or when the RememberLastStatement server setting is turned on, any user can call the `sa_conn_activity` system procedure or obtain the value of the LastStatement connection property to find out the most recently-prepared SQL statement for any other user. This option should be used with caution and turned off when it isn't required.

-zn iqsrv16 database server option

Specifies the number of request log file copies to retain.

Syntax

```
iqsrv16 -zn integer
```

Applies to

All operating systems and database servers.

Remarks

If request logging is enabled over a long period of time, the request log file can become large. The `-zn` option allows you to specify the number of request log file copies to retain. It only takes effect if `-zs` is also specified. The `-zs` option allows you to create a new log file and rename the original log file when the original log file reaches a specified size.

For example, if you redirect request logging information to the file `req.out`, and specify five request log file copies using the `-zn` option, the server creates files in the following order: `req.out.1`, `req.out.2`, `req.out.3`, `req.out.4`, and `req.out.5`. When these files exist and the active request log fills again, the following happens:

- `req.out.1` is deleted
- the files `req.out.2` to `req.out.5` are renamed `req.out.1` to `req.out.4`
- the copy of the active log is renamed `req.out.5`

Request logging is turned on using the `-zr` option and redirected to a separate file using the `-zo` option. You can also set the number of request logs using the `sa_server_option` system procedure where *nn* specifies the number of request log file copies:

```
CALL sa_server_option('RequestLogNumFiles', nn);
```

Example

In the following example, entered all on one line, request logging information is output to a request log file named `mydatabase.log`, which has a maximum size of 10 KB, and three copies of the request log are kept:

```
iqsrv16 "c:\my data\mydatabase.db" -zr all -zn 3
      -zs 10 -zo mydatabase.log
```

-zo iqsrv16 database server option

Redirects request logging information to a file separate from the regular log file.

Syntax

```
iqsrv16 -zo filename ...
```

Applies to

All operating systems and database servers.

Remarks

Request logging is turned on using the -zr option. You can direct the output from this file to a different file that is not the regular log file by specifying the -zo option.

This option also prevents request logging from appearing in the database server messages window.

-zoc iqsrv16 database server option

Redirects web service client information to a file.

Syntax

```
iqsrv16 -zoc filename ...
```

Applies to

All operating systems and database servers.

Remarks

The web service client log file contains HTTP requests and transport data recorded for outbound web service client calls. The web service client log file can also be specified using the sa_server_option system procedure:

```
CALL sa_server_option( 'WebClientLogFile', 'clientinfo.txt' );
```

Logging is enabled automatically when you specify the -zoc server option. You can enable and disable logging to this file using the sa_server_option system procedure:

```
CALL sa_server_option( 'WebClientLogging', 'ON' );
```

Example

The following command starts the database server so that it listens for HTTP web requests on port 80, and logs outbound web service client information to the file clientinfo.txt:

```
iqsrv16 web.db -xs HTTP(PORT=80) -zoc clientinfo.txt
```

-zr iqsrv16 database server option

Enables request logging of operations.

Syntax

```
iqsrv16 -zr { SQL | HOSTVARS | PLAN | PROCEDURES | TRIGGERS |  
OTHER | BLOCKS | REPLACE | ALL | NONE } ...
```

Allowed values

- **SQL** – enables logging of the following:
 - START DATABASE statements
 - STOP DATABASE statements
 - STOP SERVER statements
 - Statement preparation and execution
 - EXECUTE IMMEDIATE statement
 - Option settings
 - COMMIT statements
 - ROLLBACK statements
 - PREPARE TO COMMIT operations
 - Connects and disconnects
 - Beginnings of transactions
 - DROP STATEMENT statements
 - Cursor explanations
 - Cursor open, close, and resume
 - Errors
- **PLAN** – enables logging of execution plans (short form). Execution plans for procedures are also recorded if logging of procedures (PROCEDURES) is enabled.
- **HOSTVARS** – enables logging of host variable values. If you specify HOSTVARS, the information listed for SQL is also logged.
- **PROCEDURES** – enables logging of statements executed from within procedures and user-defined functions.
- **TRIGGERS** – enables logging of statements executed from within triggers.
- **OTHER** – enables logging of additional request types not included by SQL, such as FETCH and PREFETCH. However, if you specify OTHER but do not specify SQL, it is the equivalent of specifying SQL+OTHER. Including OTHER can cause the log file to grow rapidly and could negatively affect server performance.
- **BLOCKS** – enables logging of details showing when a connection is blocked and unblocked on another connection.

start_iq Database Server Startup Utility

- **REPLACE** – at the start of logging, the existing request log is replaced with a new (empty) one of the same name. Otherwise, the existing request log is opened and new entries are appended to the end of the file.
- **ALL** – logs all supported information. This setting is equivalent to specifying SQL+PLAN+HOSTVARS+PROCEDURES+TRIGGERS+OTHER+BLOCKS. This setting can cause the log file to grow rapidly and could negatively affect server performance.
- **NO or NONE** – turns off logging to the request log.

Applies to

All operating systems and database servers.

Remarks

This option should only be used when tracking problems. The information appears in the database server messages window or is sent to the request log. When you specify multiple values, they are separated with a , or a +.

Once the database server is started, you can change the request log settings to log more or less information using the sa_server_option system procedure.

You can find the current value of the RequestLogging setting using the following query:

```
SELECT PROPERTY( 'RequestLogging' );
```

-zs iqsrv16 database server option

Limits the size of the request log.

Syntax

```
iqsrv16 -zs { size[ k | m | g ] } ...
```

Applies to

All operating systems and database servers.

Remarks

Request logging is turned on using the -zr option, and redirected to a separate file using the -zo option. You can limit the size of the file using the -zs option.

The *size* is the maximum file size for the request log, in bytes. Use k, m, or g to specify units of kilobytes, megabytes, or gigabytes respectively.

If you specify -zs 0, then there is no maximum size for the request logging file, and the file is never renamed. This is the default value.

When the request log file reaches the size specified by either the -zs option or the sa_server_option system procedure, the file is renamed with the extension .old appended (replacing an existing file with the same name if one exists). The request log file is then restarted.

Example

The following example shows how the `-zs` option is used to control log file size. Suppose you start a database server with the following command line:

```
iqsrv16 -zr all -zs 10k -zo mydatabase.log
```

A new log file `mydatabase.log` is created. When this file reaches 10 KB in size, any existing `mydatabase.old` files are deleted, `mydatabase.log` is renamed to `mydatabase.old`, and a new `mydatabase.log` file is started. This process is repeated each time the `mydatabase.log` file reaches the specified size (in this case 10 KB).

start_iq Database File Parameters

This topic lists the parameters of the database server/database file. You specify the database file after the server options in the command syntax.

Table 34. start_iq Database File Parameters

Parameter	Description
<code>-n server-name</code>	Specifies the name of the database server.
database-file	<p>Specifies the database file name. If <i>database-file</i> is specified without a file extension, SAP Sybase IQ looks for <i>database-file</i> with extension <code>.db</code>.</p> <p>If you use a relative path, the path is read relative to the current working directory of the server. You can supply a full path.</p> <p>On Windows you can supply a path that conforms to the Universal Naming Convention (UNC) format:</p> <pre>\\server\volume\path\file.ext</pre>

Warning! The database file must be on the same machine as the database server. Managing a database file that is located on a network drive can lead to file corruption.

start_iq Database Options

Available switches for the **start_iq database-options** parameters.

Specify these options after the database file. These options apply only to that database. These options apply only to the preceding database in the command syntax.

There are two forms of syntax for **start_iq** database options:

- When specifying options in a configuration file, do not enclose the option value in quotation marks. For example:

start_iq Database Server Startup Utility

```
iqdemo.db -ek xxx
```

- When specifying options on the command line, enclose the option value in quotation marks. For example:

```
start_iq @iqdemo.cfg iqdemo.db -ek 'xxx'
```

See also

- *dblocate Database Administration Utility* on page 71

-a iqsrv16 database option

Applies the named transaction log.

Syntax

```
iqsrv16 [ server-options ] database-file -a log-filename ...
```

Applies to

All operating systems and database servers.

Remarks

This option is used to recover from media failure on the database file. When this option is specified, the database server applies the log and then shuts down—it doesn't continue to run. If you need to apply multiple transaction logs, you must know the correct order in which to apply them when using -a. The database server automatically applies multiple transaction logs in the correct order if you use the -ad or -ar option.

The -a database option must be specified after the *database-file*, and applies only to that database.

Specifying a cache size when starting the server can reduce recovery time.

Example

The following example, entered all on one line, applies the log file to a backup copy of the mydemo database.

```
iqsrv16 "c:\backup\myiqdemo.db" -a "c:\backup\mydemo.log"
```

-ad iqsrv16 database option

Specifies the directory containing transaction log files to be applied to the database.

Syntax

```
iqsrv16 [ server-options ] database-file -ad log-directory ...
```

Applies to

All operating systems and database servers.

Remarks

When you include the `-ad` option, the specified directory is scanned for transaction log files associated with the database. Transaction log files with starting log offsets greater than or equal to the start log offset stored in the database file are applied, in log offset order. Once all the transaction log files have been applied, the database is stopped. You must also specify the `-as` option if you want the database to continue running once the transaction log files have been applied.

If the specified log directory uses a relative path, then this directory is read relative to the *database-file* directory.

The `-ad` database option must be specified after the *database-file*, and applies only to that database.

Example

The database server applies the log files in the `backup` directory to the `mysample.db` database and then stops the database once the log files have been applied.

```
iqsrv16 "c:\mysample.db" -ad "c:\backup"
```

The database server applies the log files in the `backup` directory to the `mysample.db` database and the database continues running once the log files have been applied.

```
iqsrv16 "c:\mysample.db" -ad "c:\backup" -as
```

-al iqsrv16 Database Option

Extends LOGIN_MODE for LDAPUA only to a select number of users using Standard authentication.

Syntax

```
-al "user1;user2;user3" server_name.cfg database_name.db
```

Remarks

- Up to five user IDs can be specified, separated by semi-colons, and enclosed in double quotation marks.
- When run at the database level, it remains in effect until the next time the database is stopped/started.

-ar iqsrv16 database option

Specifies that any transaction log files located in the same directory as the current transaction log should be applied to the database.

Syntax

```
iqsrv16 [ server-options ] database-file -ar ...
```

start_iq Database Server Startup Utility

Applies to

All operating systems and database servers.

Remarks

When you include the `-ar` option, the database server looks for transaction log files associated with the database that are located in the same directory as the current transaction log. The transaction log location is obtained from the database. Transaction log files with starting log offsets greater than or equal to the start log offset stored in the database are applied, in log offset order. Once all the transaction log files have been applied, the database is stopped. You must also specify the `-as` option if you want the database to continue running once the transaction log files have been applied.

The `-ar` database option must be specified after the *database-file*, and applies only to that database.

Example

The database server applies the transaction log files (whose location is obtained from the database) to the `mysample.db` database. The database continues running after the transaction log files have been applied.

```
iqsrv16 "c:\mysample.db" -ar -as
```

-as iqsrv16 database option

Specifies that the database should continue to run after transaction logs have been applied (used in conjunction with `-ad` or `-ar`).

Syntax

```
iqsrv16 [ server-options ] database-file { -ad log-dir | -ar } -  
as ...
```

Applies to

All operating systems and database servers.

Remarks

The `-as` option must be specified in conjunction with either the `-ad` or `-ar` option. When you include `-as`, the database continues running after the transaction logs are applied to it.

The `-as` database option must be specified after the *database-file*, and applies only to that database.

Example

The database server applies the transaction log files to the `mysample.db` database. In this case, because `-ar` is specified, the database server obtains the location of the transaction logs from the database. The database continues running after the log files have been applied.


```
iqsrv16 "c:\mysample.db" -ar -as
```

The database server applies the log files in the backup directory to the mysample.db database. The database continues running after the log files have been applied.

```
iqsrv16 "c:\mysample.db" -ad "c:\backup" -as
```

-dh iqsrv16 Database Option

Makes a database undetectable when the Server Location utility **dblocate -d** is run against the server.

-ds iqsrv16 database option

Specifies the directory where the dbspaces for the database and the transaction log are located.

Syntax

```
iqsrv16 database-file -ds dspace-directory ...
```

Applies to

All operating systems and database servers.

This option only applies to Catalog dbspaces, not IQ dbspaces.

Remarks

When a dspace directory is specified, the database server only searches this directory for dbspaces. The location of the dspace appears in the database server messages window.

If your backup includes dbspaces with full path names, you can use this option to start the backed up copy of the database on the same computer as the original database while the original database is still running.

The -ds database option must be specified after the *database-file*, and applies only to that database.

If a transaction log file is not found in the directory specified by this option, then one is created in this location.

Warning! The -ds option should only be used for recovery. If you specify this option and the database has a current, live transaction log that is not located in the directory specified by the -ds option, then a new transaction log is created in the specified location.

Example

The following example starts a database server that looks for dbspaces in the directory c:\backup\Nov15:

```
iqsrv16 c:\backup\Nov15\my.db -ds c:\backup\Nov15\
```

The following example starts a database server that looks for dbspaces in the current directory:

```
iqsrv16 my.db -ds .
```

-ek iqsrv16 Database Option

Specifies the database encryption key.

Syntax

-ek key

Remarks

Provided after the file name of a strongly encrypted database. Requires the key value as an argument to start an encrypted database. The key value is a string, including mixed cases, numbers, letters, and special characters. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log; if you do not, the command fails. For a strongly encrypted database, you must specify either **-ek** or **-ep**, but not both.

-iqfreq iqsrv16 Database Option

Marks the specified database as in use and restores the IQ main store portion of the database to its last known consistent state.

Syntax

-iqfreq dbname

Remarks

Do not use **-iqfreq** during normal operations; use it only while force-recovering a database after seeing s_buf or free list errors during recovery after SAP Sybase IQ server failure. The *dbname* must be the physical database name, not a logical name or nickname.

Note: The option **-iqfreq** applies only to the IQ main store portion of the database, not to the IQ catalog store. **-iqfreq** does not enable a forced recovery on the catalog store.

Follow correct system recovery and database repair procedures when using **-iqfreq**.

-m iqsrv16 Database Option

Truncates (deletes) the transaction log when a checkpoint is done, either at shutdown or as a result of a checkpoint scheduled by the server.

Truncation provides a way to automatically limit the growth of the transaction log. Checkpoint frequency is still controlled by the CHECKPOINT_TIME and RECOVERY_TIME options (also definable on the command line).

The **-m** option is useful where high volume transactions requiring fast response times are being processed, and the contents of the transaction log are not being relied upon for recovery or replication. When this option is selected, there is no protection provided against media failure on the device containing the database files.

To avoid database file fragmentation, place the transaction log on a separate device or partition than the database itself.

This option is the same as the **-m** server option, but applies only to the current database or the database identified by the *database-file* command-line variable.

Note: Do not use the **-m** option with databases that are being replicated, as replication inherently relies on transaction log information. For this reason, never use the **-m** option on a multiplex database.

See also

- *dbbackup Database Administration Utility* on page 5
- *dbfhide Database Administration Utility* on page 9

-n iqsrv16 Database Option

Provides an alternate name, or nickname, for the database.

Syntax

-n *name*

Remarks

Using a nickname simplifies connections. For Open Client, the **-n** nickname must be the same as the entry in the *interfaces* file.

Note: Since a database server can load several databases, the database name is used to distinguish the different databases. However, run only one database on an SAP Sybase IQ server. If you must run two databases, start two SAP Sybase IQ database servers on different ports.

By default, the database receives as a name the file name with the path and extension removed. For example, you start a server on `c:\sybase\IQ-16_0\demo\iqdemo.db` and do not specify the **-n** option, then the name of the database is **iqdemo**. To avoid using the default name, always specify a server name.

For naming conventions, see the **-n** server option.

Note: There are two **-n** switches. If **-n** does not follow a database file name, the option names the server. If **-n** appears after a database file name, the switch is a database switch.

See also

- *-n iqsrv16 database server option* on page 172

-r iqsrv16 Database Option

Forces all databases that start on the database server to be read-only.

No changes to the database are allowed: the database server doesn't modify the database file.

-sm iqsrv16 Database Option

Provides an alternate database server name that can be used to access the read-only mirror database.

Syntax

-sm *name*

Remarks

The alternate-server-name is only active when the database server is acting as mirror for the database. By using the **-sm** and **-sn** command-line options, an application can always connect to the database on the primary or the mirror server, without knowing which physical server is acting as primary or mirror.

-sn iqsrv16 Database Option

Provides an alternate server name for a single database running on a database server.

Syntax

-sn *name*

Remarks

The database server can be configured to listen for more than one server name for a particular database server. Server names other than the real server name are called alternate server names, and are specific to a particular database running on the database server. Clients using the alternate server name to connect can only connect to the database that specified the alternate server name. Alternate server names must be unique on the network; otherwise, the database fails to start. If the database is started in the server command and the alternate server name is not unique, the server fails to start.

-xp iqsrv16 Database Option

Provides information to an operational server that allows it to connect to its partner and to the arbiter when database mirroring is being used.

Syntax

```
iqsrv16 [ server-options ] database-file -xp on
```

Remarks

The **-xp** database option must be specified after the *database-file*, and applies only to that database.

If the connection parameters specified in the **-xp** option are invalid, and there are multiple databases running on the server, then the mirrored database fails to start and does not attempt

to reconnect. If the mirrored database is the only database running on the database server, then the database server does not start.

- **on** – You can only use database mirroring if you specify the **-xp** option when the database server is started, even if mirroring information is stored in the database. When you specify **-xp on**, you cannot specify other mirroring options with the **-xp** option. The value **off** is not supported. Database mirroring settings are defined in the database using the following statements:

- **CREATE MIRROR SERVER**
- **SET MIRROR OPTION**

When you specify **-xp on**, you should also specify the name of the database server in the mirroring system using the **-n** option. It is also recommended that you also include the **-su** option to specify the password for the utility database. Then, you can use the utility database to shut down the database server, or force the mirror server to become the primary server, if necessary.

start_iq Database Server Startup Utility

stop_iq Database Shutdown Utility

On UNIX and Linux platforms, you can stop the database server using the **stop_iq** utility.

Use **stop_iq** to stop your server and close all user connections to your server.

stop_iq shuts down a server without regard for user connections or load process status.

Normally, do not shut down a server while it is still connected to one or more clients.

Note: The **dbstop** utility offers a finer level of control, providing options to control whether a server is stopped based on user connections.

When to Stop and Restart the Server

In some situations, the server may need to be stopped and restarted.

For example:

- To install a new version of SAP Sybase IQ
- To reset some server command line options
- To cause a small number of server-wide database options to take effect
- Before closing the operating system session

Ways to Stop Database Servers

The preferred methods of stopping the database server in SAP Sybase IQ are as follows:

- In UNIX, use the **stop_iq** utility at the operating system command line.
When you run **stop_iq**, the following message displays:


```
"Please note that 'stop_iq' will shut down a server completely without regard for users, connections, or load process status. For more control, use the 'dbstop' utility, which has options that control stopping servers based on active connections."
```
- In Windows, click **Shutdown** on the database server display or right-click the IQ icon in the system tray and select **Exit**.
- In Windows, if the server is run as a service, open the Service Manager in Control Panel. Select the service and click **Stop**.

Avoid shutting down a server that is still connected to one or more clients. Doing so returns a warning that any uncommitted transactions will be lost. Disconnect or close all the clients and try again.

You can also stop the database server in the following ways:

stop_iq Database Shutdown Utility

- At the operating system command line, issue the **DBSTOP** command with appropriate parameters. Use the same parameters as when you started the server. Without the proper connection parameters **DBSTOP** does not know how to connect to the server to tell it to shut down.
- In an Interactive SQL window or command file, issue the **STOP ENGINE** command to stop a named database server.
- In UNIX, in the window where the database server was started, type:

```
q
```

This command does not work if you have redirected input to a different device or if you started the server with **start_iq**. It only works with **iqsrv16**.

- In a UNIX **cron** or **at** job, use **stop_iq** with the appropriate **-stop** option. The utility stops one or all servers associated with the user who starts the **cron** or **at** job depending on the parameter specified. The user must be the same one who started the server. No operator prompting occurs, and no operator action is required.

To use **stop_iq** in such jobs, specify the utility with the appropriate **-stop** option:

```
stop_iq -stop one
```

Setting **-stop one** shuts down a single server, when exactly one running server was started by the user ID that starts the **cron** or **at** job.

```
stop_iq -stop all
```

Setting **-stop all** shuts down all servers that were started by the user ID that starts the **cron** or **at** job.

Note: You must specify the full path name to the **stop_iq** executable in the **cron** statement.

Example — Stop a Server With stop_iq

The following example uses the **stop_iq** utility in a UNIX operating system command line to shut down a SAP Sybase IQ server and close all user connections to it.

When you issue the **stop_iq** command, SAP Sybase IQ lists all the servers owned by other users, followed by the server(s) you own. It then asks if you want to stop your server. For example:

```
% stop_iq
```

```
Checking system for IQ 16 Servers ...
The following 2 server(s) are owned by other users.
##      Owner      PID      Started      CPU_Time      Additional Information
--      -
          handari  19895   15:43:44    183:38
start_iq @iqdemo.cfg iqdemo.db -gn 105 -o /server1/users/surya/
IQ-16_0/logfiles/surya_ibm2.001.srvlog -hn 8          pamela      409802
18:05:02          0:05   SVR:ibm1_iqdemo2 DB:iqdemo
PORT:2678/ibm1/users/sybase/iql60/IQ-16_0/bin64/iqsrv16 @iqdemo.cfg
iqdemo.db -ti 4400 -gn 25 -o /ibm1/users/sybase/iql60/IQ
16_0/logfiles/ibm64qa_iq
The following 1 server(s) are owned by 'kermit'
##      Owner      PID      Started      CPU_Time      Additional Information
```



```

-----
l:      kermit  422034  15:11:37    0:07  SVR:myserver_iqdemo
DB:iqdemo PORT:2638 /myserver/users/sybase/iql60/IQ-16_0/bin64/
iqsrv16
@iqdemo.cfg iqdemo.db -ti 4400 -gn 25 -o /myserver/users/sybase/
iql60/IQ-
16_0/logfiles/myserver_iq
start_iq -c 32m -gd all -gm 10 -gn 25 -gp 4096 -ti 4400 -tl 300
@iqdemo.cfg

```

```

--
    Please note that 'stop_iq' will shut down a server completely
    without regard for users connections or load processes status.
    For more control, use the 'dbstop' utility, which has options
    that control stopping servers based on active connections.

```

Do you want to stop the server displayed above <Y/N>?

To shut down the server, type **Y** (yes). Messages like the following display:

```

Shutting down server (422034) ...
Checkpointing server (422034) ...
Server shutdown.

```

To leave the server running, type **N** (no). You return to the system prompt and IQ does not shut down the server.

If no running servers were started by your user ID, SAP Sybase IQ displays information about servers run by other users, then a message like the following:

```

There are no servers owned by 'kermit'

```

Example — Stop a Server From Interactive SQL

The following example stops a server from Interactive SQL:

```

STOP ENGINE Ottawa UNCONDITIONALLY

```

The optional keyword **UNCONDITIONALLY** specifies that the database server will be stopped even if there are connections to it.

Note: You can stop a server from Interactive SQL if you are connected as DBA to one of the databases running on that server (including the `utility_db` database), or if the server was started with the **-gk ALL** option.

Permissions Required to Stop the Server

When you start a server, you can use the **-gk** option to set the level of permissions required for users to stop the server with **DBSTOP** or **STOP ENGINE**.

The default level of permissions required is **DBA**, which requires the **SERVER OPERATOR** system privilege, but you can also set the value to **ALL** or **NONE**. If you set it to **NONE**, even the

stop_iq Database Shutdown Utility

a user with **SERVER OPERATOR** system privileges cannot execute **STOP ENGINE**. In a production environment, only the DBA should be allowed to stop the database server.

Running **stop_iq** at the UNIX command line, or Shutdown on Windows platforms, still allows you to stop the server and databases on the machine where the server was started.

Stopping the Database Server with stop_iq

Run **stop_iq** from the command line.

1. Issue a command in this format:

```
stop_iq [ -cleanup ] [ -stop [ one | all ] ] [ -user <user_name> ] [
-version [ 12 | 15 | all ]> ] [ -wait <seconds> ]
```

2. When prompted to stop your server, respond Y (yes).
3. If the server is connected to one or more clients, respond to the warning that uncommitted transactions will be lost.

The following example illustrates how to use **stop_iq** interactively to list all the servers that are running, and then shut down one server.

```
% stop_iq

Checking system ...
The following 1 server(s) are owned by 'TEST'

##      Owner      PID      Started  CPU_Time  Additional Information
--      -
1:     TEST       22399  08:56:39    1:43  SVR:QA_sun7qa DB:iqdemo PORT:
8888
/sun7qa1/users/QA/090513/IQ-16_0/bin64/iqsrv16 @iqdemo.cfg iqdemo.db
-ti 4400
--

Please note that 'stop_iq' will shutdown a server completely
without regard for users connections or load processes status.
For a finer level of detail the utility 'dbstop' has the options
to control whether a server is stopped based on active
connections.

Do you want to stop the server displayed above <Y/N>? Y

Shutting down server (22399) ...
Checkpointing server (22399) ...
Server shutdown.
```

stop_iq Options

This table lists the options available for the **stop_iq** utility.

Table 35. stop_iq Options

Switch	Description
-cleanup	Removes the orphan SAP Sybase IQ process on Linux.
-stop [one all]	Removes user interaction with stop_iq . Assumes a yes response to all questions.
-user	Performs two functions: <ul style="list-style-type: none"> • If, due to truncation or substitution, stop_iq cannot find the server or agent owned by the current user, the system manager can specify the name and id found in the process table in the -user argument to shut down the server/agent with the stop_iq utility. • A user with root privileges can shut down another user's server or agent without having to log in as that user. The stop_iq utility has no superuser (su) or root powers, so a non-privileged user cannot shut down a server owned by another user.
-version	Specifies the major version of SAP Sybase IQ that is being used. The default is the current major version (16.0).
-wait	Specify the time to wait for the server to shut down before timeout expires.

Stopping Servers in a cron or at Job

If using **stop_iq** in a **cron** or **at** job, use the appropriate **-stop** option:

Table 36. stop_iq -stop Options

Option	Description
-stop one	Shuts down a single server, when exactly one running server was started by the user ID that starts the cron or at job. This prevents the wrong server from being shut down if several are running.

stop_iq Database Shutdown Utility

Option	Description
-stop all	Shuts down all servers that were started by the user ID that starts the server.

For example:

```
stop_iq -stop one
```

```
stop_iq -stop all
```

Note: In a **cron** statement, you must specify the complete path name to the **stop_iq** executable.

Servers with Long Paths

Depending on the operating system, the **stop_iq** utility may fail to report running servers when the path to the server exceeds 74 characters.

As a workaround, use a **ps -ef** command to display servers with long paths. For example:

```
ps -ef|grep myserver  
  
rsmithson 1133      1      0 07:04:32 ?                223:35  
/sunsys1234/users/rsmithson/mybigtest1234_withdeletion  
_allcol/IQ-16_0/bin64/myserver  
  
rsmithson 2046     862    0 10:02:30 pts/3          0:00 grep  
myserver
```

Appendix: **dbisqlc** Interactive SQL Classic Utility (Deprecated)

The Interactive SQL Classic (**dbisqlc**) utility executes SQL statements against a database. This utility is deprecated.

The utility is similar to the Interactive SQL utility, except that it is not implemented in Java. The absence of Java can be an advantage if you are deploying the utility to a computer with limited resources.

Note: Interactive SQL Classic is deprecated; however, there are currently no plans to remove it. Interactive SQL Classic is provided for backwards compatibility for running SQL scripts and as a lightweight tool for deployment. It does not support all the features that Interactive SQL supports and may not support all the features available in the current version of SAP Sybase IQ. Use the Interactive SQL utility instead.

The Interactive SQL Classic utility is supported on Microsoft Windows and UNIX.

Interactive SQL Classic Syntax

Invoke Interactive SQL Classic from a command prompt.

```
dbisqlc [ options ] [ dbisqlc-command | command-file ]
```

If you specify *dbisqlc-command*, Interactive SQL Classic executes the command. You can also specify a command file name. If you do not specify *dbisqlc-command* or *command-file* argument, Interactive SQL Classic enters interactive mode, where you can type a command into a command window.

Interactive SQL Classic Options

Specify these options when invoking Interactive SQL Classic from a command prompt.

Table 37. Interactive SQL Classic Options

Option	Description
-c <i>"keyword=value; ..."</i>	Specifies connection parameters. If you do not specify any connection parameters, the environment variable SQLCONNECT is used. If Interactive SQL cannot connect, enter the appropriate parameters in the dialog box that appears. Note: Always specify connection parameters for Interactive SQL Classic instead of relying on defaults, whether you specify them in a command line or an initialization file such as <code>.odbc.ini</code> on UNIX, or <code>odbc.ini</code> on Windows. If you start more than one database on a server, for example, specify the database name, and in a network with subnets, specify the communications protocol parameter with host number.
-d <i>delimiter</i>	Specifies a command delimiter. By default, the delimiter is the semicolon.
-q	Quiet mode—does not display output messages. This option is useful only if you start Interactive SQL Classic with a command or command file.
-r	Returns the error "Not enough fields allocated in sqllda" if the defined result set of the stored procedure does not match the actual result set. This option may be useful when you are querying stored procedures.
-x	Checks syntax only. Scans commands but does not execute them. You may find this option useful for checking long command files for syntax errors.

Function and Special Keys (UNIX)

Use function keys and special keys to move data and list database tables.

Table 38. Interactive SQL Classic Function and Special Keys on UNIX

Function key	Description
F5	Move data to the left by one column in the data window.
Shift+F5	Move data to the left by one character.

Function key	Description
F6	Move data to the right by one column.
Shift+F6	Move data to the right by one character.
F7	Display a list of the tables in the database. Use the up and down arrow keys to scroll through the table names changing the highlighted table name. While the list appears, press Enter to insert the current table name into the command window at the cursor position. While the list appears, press F7 to see a list of columns for the highlighted table. Again, use Enter to select the highlighted column name and put it into the command window at the cursor position.
Ctrl+PgUp	Move to the top of data.
Ctrl+PgDn	Move to bottom of data.

Function and Special Keys (Windows)

Use function keys and special keys to move data and list database tables.

Table 39. Interactive SQL Classic Function and Special Keys on Windows

Function key	Description
F5	Move data to the left by one column in the data window.
Shift+F5	Move data to the left by one character.
F6	Move data to the right by one column.
Shift+F6	Move data to the right by one character.
F7	Display a list of the tables in the database. Use the up and down arrow keys to scroll through the table names changing the highlighted table name. While the list appears, press Enter to insert the current table name into the command window at the cursor position. While the list appears, press F7 to see a list of columns for the highlighted table. Again, use Enter to select the highlighted column name and put it into the command window at the cursor position.
F9	Execute the command that is in the command window. This operation can also be performed with the mouse by clicking Execute.
F10	Activate the menus at the top of the window.
Page Up	Move data up a page.
Page Down	Move data down a page.
Ctrl+PageUp	Move to top of data.

Function key	Description
Ctrl+PageDown	Move to bottom of data.

Command Recall Keys (Windows)

Use key sequences to recall previous commands

Table 40. Interactive SQL Classic Recall Keys

Key sequence	Description
Ctrl+r	Brings up the command recall window
Ctrl+p	Cycles backwards through previously executed commands. Retrieved commands are placed into the command window
Ctrl+n	Cycles forward through previously executed commands

Index

B

- backup
 - live 5
 - on running server 5
- backup utility 5
- buffer cache
 - partitioning 165

C

- cache size
 - setting for catalog store 138
- catalog store
 - setting cache size 138
 - validating 89
- CBSize connection parameter 100
- character set translation
 - enabling 140
- character strings 133
- CLEAR statement
 - syntax 48
- client side
 - Encryption [ENC] connection parameter 98
- code pages
 - DEFAULT_ISQL_ENCODING option 68
- collation
 - CP874toUTF8 utility 3
- command files
 - parameters 57
- command limiter 14
- command line
 - database server 135
 - in configuration file 137
- command line utilities
 - dbfhide syntax 9
 - file hiding syntax 9
- CommBufferSize connection parameter 100
- communications
 - Encryption [ENC] connection parameter 98
- configuration files
 - adding simple encryption with dbfhide 9
 - hiding 9
- CONFIGURE statement
 - syntax 49
- conformance 132

- CONNECT statement
 - syntax 50
- connection parameters
 - dbisql 14
 - Encryption [ENC] 98
- connections
 - dbisql 52
- CP874toUTF8 utility 3

D

- data
 - exporting from tables into files 54
- data sources
 - creating ODBC with iqdsn 95
- database information 11
- database pages
 - displaying size of 11
- database server
 - command line 135
 - starting 135
 - stopping 81, 213
- database servers
 - starting 65
 - stopping 67
- database utilities
 - data source 95
 - dbfhide 9
 - file hiding 9
 - iqdsn 95
- database utility
 - log translation 83
- databases
 - creating 107
 - information 11
 - page usage 11
 - starting 63
 - stopping 66
- dbbackup utility 5
- dbfhide
 - about 9
- dbfhide utility
 - syntax 9
- dbinfo utility 11
- dbinit utility 107

Index

- dbisql
 - connecting to a database 50
 - options 62
- dbisql utility 13
- dbisqlc utility 219
- dblocate utility 71
- dblog utility 73
 - command line 73
- dbping utility 77
- dbspace header 103
- dbstop utility 81
 - usage 81
- dbtran utility 83
 - exit codes 83
 - syntax 83
- dbvalid utility 89
 - exit codes 91
- default server 195
- DEFAULT_ISQL_ENCODING option
 - description 68
- DisableMultiRowFetch connection parameter 101
- DISCONNECT statement
 - syntax 52
- DMRF connection parameter 101

E

- embedded SQL 133
 - introduction 131
- ENC connection parameter
 - description 98
- encryption
 - dbfhide utility 9
 - Encryption [ENC] connection parameter 98
 - file hiding utility 9
 - INI files 9
 - strong 98
- Encryption connection parameter
 - description 98
- environment variable
 - switches 137
- escape character
 - OUTPUT SQL statement 54
- exit codes
 - validation utility (dbvalid) 91
- exporting data
 - from tables into files 54

F

- file hiding utility
 - about 9

- syntax 9
- files
 - exporting data from tables into 54
- FIPS 132

H

- hyperthreading
 - server switch 164

I

- information utility 11
- INI files
 - adding simple encryption with dbfhide 9
- initialization files
 - adding simple encryption with dbfhide 9
- Interactive SQL
 - OUTPUT statement syntax 54
 - specifying code page for reading and writing to files 68
- Interactive SQL Classic utility 219
 - command recall keys 222
 - function keys 220
 - syntax 219
- Interactive SQL options
 - DEFAULT_ISQL_ENCODING 68
- Interactive SQL utility 13
 - comments 23
 - connect window 18
 - connection parameters 17
 - executing statements 21
 - indenting 24
 - keyboard shortcuts 30
 - main window 28
 - multiple result sets 32
 - options 14, 68
 - plan viewer 42
 - query editor 40
 - quiet mode 16
 - source control integration 44
 - syntax 13
- iqdscp utility 93
- iqdsn utility
 - about 95
- iqinit utility 107
- iqmsgnum startup switch 163
- iqmsgsz startup switch 163

- iqnumbercpus
 - server switch 164
- iqocscfg utility 117
- iqpartition startup switch 165
- iqrlvmem startup switch 165
- iqsqlpp length limit 133

L

- line numbers 133
- live backup 5
- LivenessTimeout connection parameter 101
- load performance
 - iqpartition server option 165
- loading data 3
- lock contention
 - iqpartition server option 165
- log files 73
- LogFile
 - connection parameter 99
- LTM
 - transaction log options 73
- LTO connection parameter 101

M

- message log
 - iqmsgnum server option 163
 - iqmsgsz server option 163
 - specifying number of archives 163
 - specifying size 163

O

- ODBC data sources
 - creating with iqdsn 95
- ON_ERROR option
 - description 69
- options
 - DEFAULT_ISQL_ENCODING 68
 - ON_ERROR 69
 - setting dbisql options 49
 - setting DBISQL options 49
 - setting temporary 62
- OUTPUT statement
 - SQL syntax 54

P

- pages
 - displaying usage in database files 11

- PARAMETERS statement
 - syntax 57

- ping
 - utility 77

R

- READ statement
 - syntax 58
- recycling the server 213
- replication
 - replication server 73
- resource planning
 - iqnumbercpus switch 164
- return codes
 - validation utility (dbvalid) 91

S

- security
 - adding simple encryption to configuration files
 - 9
 - dbfhide utility 9
 - Encryption [ENC] connection parameter 98
 - file hiding utility 9
- server
 - default 195
- servers
 - live backup 5
 - recycling 213
- SET CONNECTION statement
 - syntax 61
- SET OPTION statement
 - syntax 62
- SET TEMPORARY OPTION statement
 - syntax 62
- shared memory
 - finding default server 195
- software
 - dbinfo 11
 - dblog 73
- SQL preprocessor
 - command line 131
 - command line options 132
 - introduction 131
 - syntax 131
- SQL statements
 - OUTPUT syntax 54
- SQL92 132

Index

- sqlpp
 - command line 131
 - command line options 131
 - introduction 131
 - options 131
 - parameters 131
 - syntax 131
- stack size
 - SQL Anywhere threads 156
- START DATABASE statement
 - syntax 63
- START ENGINE statement
 - syntax 65
- start_iq
 - database option 203
 - server options 135
 - syntax 135
- starting
 - database servers 65
 - databases 63
- statements
 - OUTPUT syntax 54
- STOP DATABASE statement
 - syntax 66
- STOP ENGINE statement
 - syntax 67
- stop_iq
 - syntax 216
 - usage 213
 - utility 213
- stopping databases 66
- strong encryption
 - Encryption [ENC] connection parameter 98

T

- tables
 - exporting data into files from 54
- Thai language 3
- threads
 - limiting stack size 156
- transaction log
 - backup 5

- dblog 73
- utilities 73

U

- UTF8 collation 3
- utilities
 - backup 5
 - CP874toUTF8 3
 - data source 95
 - dbfhide 9
 - dbfhide syntax 9
 - dbinfo 11
 - dbinit 107
 - dblocate 71
 - dblog 73
 - dbping 77
 - dbspace header 103
 - dbstop 81
 - dbvalid 89
 - file hiding 9
 - file hiding syntax 9
 - information 11
 - Interactive SQL 13
 - Interactive SQL Classic 219
 - iqdscp 93
 - iqinit 107
 - iqocscfg 117
 - isql 119
 - Server Location utility 71
 - SQL preprocessor 131
 - stop_iq 213
 - Transaction Log utility 73
 - Validation utility 89

V

- validating
 - catalog store 89
- validation utility (dbvalid)
 - exit codes 91