

HyperSQL Utilities Guide

Edited by The HSQLB Development Group, Blaine Simpson, and Fred Toussi

HyperSQL Utilities Guide

by The HSQLB Development Group, Blaine Simpson, and Fred Toussi

\$Revision: 3539 \$

Published \$Date: 2010-03-17 11:47:21 -0400 (Wed, 17 Mar 2010) \$

Copyright 2002-2009 The HSQL Development Group. Permission is granted to distribute this document without any alteration under the terms of the HSQLDB license.

Table of Contents

Preface	viii
Available formats for this document	viii
1. SqlTool	1
Purpose, Coverage, Changes in Behavior	1
Platforms and SqlTool versions covered	2
Recent Functional Changes	3
New Features	4
The Bare Minimum	5
Embedding	7
Non-displayable Types	7
Desktop shortcuts	7
Loading sample data	8
Satisfying SqlTool's CLASSPATH Requirements	8
Accessing older HSQLDB Databases with SqlTool	8
App-specific Classes, Embedding, and non-HyperSQL Databases	9
Distributing SqlTool with your Apps	10
SqlTool Client PCs	10
RC File Authentication Setup	10
Switching Data Sources	13
Using Inline RC Authentication	14
Logging	14
Interactive Usage	14
SqlTool Command-Line Editing	15
Command Types	15
Command Types	16
Special Commands	17
Edit Buffer / History Commands	20
PL Commands	22
? Variable	23
Storing and retrieving binary files	23
Command History	24
Shell scripting and command-line piping	24
Emulating Non-Interactive mode	24
Non-Interactive	24
Giving SQL on the Command Line	25
SQL Files	25
Piping and shell scripting	27
Optimally Compatible SQL Files	28
Comments	28
Special Commands and Edit Buffer Commands in SQL Files	28
Automation	30
Getting Interactive Functionality with SQL Files	30
Character Encoding	30
Generating Text or HTML Reports	30
SqlTool Procedural Language	31
Variables	32
Macros	32
PL Sample	33
Logical Expressions	33
Flow Control	34
Example	34

Chunking	36
Why?	36
How?	36
Raw Mode	36
SQL/PSM, SQL/JRT, and PL/SQL	37
Delimiter-Separated-Value Imports and Exports	39
Simple DSV exports and imports using default settings	41
Specifying queries and options	42
Unit Testing SqlTool	43
2. Hsqldb Test Utility	45
3. Database Manager	47
Brief Introduction	47
Auto tree-update	47
Automatic Connection	48
RC File	48
Using the current DatabaseManagers with an older HSQLDB distribution.	48
DatabaseManagerSwing as an Applet	49
4. Transfer Tool	51
Brief Introduction	51
A. HyperSQL File Links	52

List of Tables

1. Available formats of this document	viii
---	------

List of Examples

1.1. Sample RC File	10
1.2. Inserting binary data into database from a file	24
1.3. Downloading binary data from database to a file	24
1.4. Piping input into SqlTool	27
1.5. Simple SQL file using PL	33
1.6. SQL File showing use of most PL features	35
1.7. Interactive Raw Mode example	37
1.8. PL/SQL Example	38
1.9. SQL/JRT Example	39
1.10. SQL/PSM Example	39
1.11. DSV Export Example	41
1.12. DSV Import Example	41
1.13. DSV Export of an Arbitrary SELECT Statement	42
1.14. Sample DSV headerswitch settings	43
1.15. DSV targettable setting	43

Preface

If you notice any mistakes in this document, please email the author listed at the beginning of the chapter. If you have problems with the procedures themselves, please use the HSQLDB support facilities which are listed at <http://hsqldb.org/web/hsqSupport.html>.

Available formats for this document

This document is available in several formats.

You may be reading this document right now at <http://hsqldb.org/doc/2.0>, or in a distribution somewhere else. I hereby call the document distribution from which you are reading this, your *current distro*.

<http://hsqldb.org/doc/2.0> hosts the latest production versions of all available formats. If you want a different format of the same *version* of the document you are reading now, then you should try your current distro. If you want the latest production version, you should try <http://hsqldb.org/doc/2.0>.

Sometimes, distributions other than <http://hsqldb.org/doc/2.0> do not host all available formats. So, if you can't access the format that you want in your current distro, you have no choice but to use the newest production version at <http://hsqldb.org/doc/2.0>.

Table 1. Available formats of this document

format	your distro	at http://hsqldb.org/doc/2.0
Chunked HTML	index.html	http://hsqldb.org/doc/2.0/util-guide/
All-in-one HTML	util-guide.html	http://hsqldb.org/doc/2.0/util-guide/util-guide.html
PDF	util-guide.pdf	http://hsqldb.org/doc/2.0/util-guide/util-guide.pdf

If you are reading this document now with a standalone PDF reader, the your distro links may not work.

Chapter 1. SqlTool

SqlTool Manual

Blaine Simpson, HSQL Development Group

\$Revision: 3607 \$

Published \$Date: 2010-06-01 07:18:26 -0400 (Tue, 01 Jun 2010) \$

Purpose, Coverage, Changes in Behavior



Note

Due to many important improvements to SqlTool, both in terms of stability and features, all users of SqlTool are advised to use the latest version of SqlTool, even if your database instances run with an older HSQLDB version. How to do this is documented in the [Accessing older HSQLDB Databases with SqlTool](#) section below.

This document explains how to use SqlTool, the main purpose of which is to read your SQL text file or stdin, and execute the SQL commands therein against a JDBC database. There are also a great number of features to facilitate both interactive use and automation. The following paragraphs explain in a general way why SqlTool is better than any existing tool for text-mode interactive SQL work, and for automated SQL tasks. Two important benefits which SqlTool shares with other pure Java JDBC tools is that users can use a consistent interface and syntax to interact with a huge variety of databases-- any database which supports JDBC; plus the tool itself runs on any Java platform. Instead of using `isql` for Sybase, `psql` for Postgresql, `sql*plus` for Oracle, etc., you can use SqlTool for all of them. As far as I know, SqlTool is the only production-ready, pure Java, command-line, generic JDBC client. Several databases come with a command-line client with limited JDBC abilities (usually designed for use with just their specific database).



Important

The SqlTool commands and settings are intuitive once you are familiar with the usage idioms. This Guide does not attempt to list every SqlTool command and option available. When you want to know what SqlTool commands or options are available for a specific purpose, you need to list the commands of the appropriate type with the relevant "?" command. For example, as explained below, to see all Special commands, you would run `\?`; and to see all DSV export options, you run `\x?`.

SqlTool is purposefully not a Gui tool like Toad or DatabaseManager. There are many use cases where a Gui SQL tool would be better. Where automation is involved in any way, you really need a text client to at least test things properly and usually to prototype and try things out. A command-line tool is really better for executing SQL scripts, any form of automation, direct-to-file fetching, and remote client usage. To clarify this last, if you have to do your SQL client work on a work server on the other side of a VPN connection, you will quickly appreciate the speed difference between text data transmission and graphical data transmission, even if using VNC or Remote Console. Another case would be where you are doing some repetitive or very structured work where variables or language features would be useful. Gui proponents may disagree with me, but scripting (of any sort) is more efficient than repetitive copy & pasting with a Gui editor. SqlTool starts up very quickly, and it takes up a tiny fraction of the RAM required to run a comparably complex Gui like Toad.

SqlTool is superior for interactive use because over many years it has evolved lots of features proven to be efficient for day-to-day use. Four concise help commands (`\?`, `/?`, `*?`, and `/?`) list all available commands of the corresponding type. SqlTool doesn't support up-arrow or other OOB escapes (due to basic Java I/O limitations), but it more than makes up for this limitation with macros, user variables, command-line history and recall, and command-line editing with extended Perl/Java regular expressions. The `\d` commands deliver JDBC metadata information as consistently

as possible (in several cases, database-specific work-arounds are used to obtain the underlying data even though the database doesn't provide metadata according to the JDBC specs). Unlike server-side language features, the same feature set works for any database server. Database access details may be supplied on the command line, but day-to-day users will want to centralize JDBC connection details into a single, protected RC file. You can put connection details (username, password, URL, and other optional settings) for scores of target databases into your RC file, then connect to any of them whenever you want by just giving SqlTool the ID ("urlid") for that database. When you Execute SqlTool interactively, it behaves by default exactly as you would want it to. If errors occur, you are given specific error messages and you can decide whether to roll back your session. You can easily change this behavior to auto-commit, exit-upon-error, etc., for the current session or for all interactive invocations. You can import or export delimiter-separated-value files. If you need to run a specific statement repeatedly, perhaps changing the WHERE clause each time, it is very simple to define a macro.

When you Execute SqlTool with a SQL script, it also behaves by default exactly as you would want it to. If any error is encountered, the connection will be rolled back, then SqlTool will exit with an error exit value. If you wish, you can detect and handle error (or other) conditions yourself. For scripts expected to produce errors (like many scripts provided by database vendors), you can have SqlTool continue-upon-error. For SQL script-writers, you will have access to portable scripting features which you've had to live without until now. You can use variables set on the command line or in your script. You can handle specific errors based on the output of SQL commands or of your variables. You can chain SQL scripts, invoke external programs, dump data to files, use prepared statements, Finally, you have a procedural language with `if`, `foreach`, `while`, `continue`, and `break` statements.

Platforms and SqlTool versions covered

SqlTool runs on any Java 1.5 or later platform. I know that SqlTool works well with Sun and OpenJDK JVMs. I haven't run other vendors' JVMs in years (IBM, JRockit, etc.). As my use with OpenJDK proves that I don't depend on Sun-specific classes, I expect it to work well with other (1.5-compatible) Java implementations.

SqlTool no longer writes any files without being explicitly instructed to. Therefore, it should work fine on read-only systems, and you'll never have orphaned temp files left around.

The command-line examples in this chapter work as given on all platforms (if you substitute in a normalized path in place of `$HSQLDB_HOME`), except where noted otherwise. When doing any significant command-line work on Windows (especially shell scripting), you're better off to completely avoid paths with spaces or funny characters. If you can't avoid it, use double-quotes and expect problems. As with any Java program, file or directory paths on the command line after "java" can use forward slashes instead of back slashes (this goes for System properties and the `CLASSPATH` variable too). I use forward slashes because they can be used consistently, and I don't have to contort my fingers to type them :).

If you are using SqlTool from a HyperSQL distribution of version 1.8.x or earlier, you should use the documentation with that distribution, because this manual documents many new features, several significant changes to interactive-only commands, and a few changes effecting backwards-compatibility (see next section about that). This document is now updated for the current versions of SqlTool and SqlFile at the time I am writing this (versions 3406 and 3604 correspondingly-- SqlFile is the class which actually processes the SQL content for SqlTool). Therefore, if you are using a version of SqlTool or SqlFile that is more than a couple revisions greater, you should find a newer version of this document. (The imprecision is due to content-independent revision increments at build time, and the likelihood of one or two behavior-independent bug fixes after public releases). The startup banner will report both versions when you run SqlTool interactively. (Dotted version numbers of SqlTool and SqlFile definitely indicate ancient versions).

This guide covers SqlTool as bundled with HSQLDB after 1.8.x. ¹

¹ To reduce the time I will need to spend maintaining this document, in this chapter I am giving the path to the `sample` directory as it is in HyperSQL 2.0.x distributions, namely, `HSQLDB_HOME/sample`. Users of HSQLDB before 2.0.x should translate these sample directory paths to use `HSQLDB_HOME/src/org/hsqldb/sample/...`

Recent Functional Changes

This section lists changes to SqlTool since the last major release of HSQLDB which may effect the portability of SQL scripts. For this revision of this document, this list consists of script-impacting changes made to SqlTool *after* the final 1.8.0.0 HSQLDB release. I'm specifically not listing changes to interactive(:)-only commands (with one legacy exception which is listed below), since these commands can't be used in SQL scripts; and I'm specifically not listing backwards-compatible feature additions and enhancements. The reason for limiting the change list to only portability-impacting changes is that a list of all enhancements since just 1.8.1.1 would be pages long.

- SqlTool is now bundled in the supplied file `sqltool.jar` instead of `hsqldb.jar`. Therefore, the usage idiom `java -jar ../hsqldb.jar` has changed to `java -jar ../sqltool.jar`. (depending on where you get your HyperSQL resources from, these files may be named with a version label, like `sqltool-1.2.3.4.jar`).
- The package of SqlTool and support classes has been changed from `org.hsqldb.util` to `org.hsqldb.cmdline`. There is no change to `java -jar ...` command-lines, but you will need to change the package name in SqlTool command lines of the form `java ... org.hsqldb ...`. This later usage is necessary if you need to modify the classpath for any reason, such as for embedding, to use a logging config file, or to use a non-HSQLDB JDBC driver.
- SqlTool now consistently outputs `\r\n` line breaks when on `\r\n`-linebreak platforms, like Windows. This includes output written to stdout, `\w` files, and `\o` files.
- Time type values are always output with the date as well as the time. This was required in order to produce consistent output for the wildly varying formats provided by different database vendors.
- DSV input now takes JDBC Timestamp format with date and optionally time of day.
- The command `:"` is now strictly an interactive command. If you want to repeat a command in an SQL scripts, just repeat the exact text of the command. Non-interactive use now has no dependency on command history.
- The command `:w` has replace the command `\w`. Unlike writing "output" to a file with `\w`, `:w` is used to write SQL "commands", and this is an interactive feature.
- Shell scripts using raw mode (e.g. PL/SQL scripts) must terminate the raw code with a line containing `:"`, which will also send the code to the database for execution. (The old `."` command has been changed to `:"` to make it very clear that that command is now an interactive command).
- You must use raw mode to *chunk* SQL statements to your DB server. I.e., with previous versions of SqlTool, SQL statements were only sent to the DB server when a semi-colon is read at the end of a line. SqlTool now has a much more powerful parser, and individual SQL statements are sent to the DB server as soon as they are terminated with a semi-colon, regardless of line delimiters. Therefore, to send multiple SQL statements to the DB server in one transmittal, you must use raw mode.
- The `--sql` argument will never automatically append a semicolon to the text you provide. If you want to execute a command ending with a semi-- then type a semi.
- Default encoding used is now the platform's default encoding instead of `US-ASCII`.
- To minimize side-effects (especially for instance-based programmatic usage), the only System properties used are those predefined by the JVM (incl. `javax.net.ssl.*`). Properties of the form `sqlfile.*` and `sqltool.*` are specifically no longer supported. (Less invasive configuration systems are provided to serve the same purposes).
- SqlTool no longer displays the usage banner if none of inline-RC, urlid, SQL files are supplied, because that now starts up SqlTool with no JDBC Connection. To see the usage banner, use the `--help` command-line switch.

- Requires Java 1.5 in order to build or run.
- Update and row counts are not displayed in non-interactive mode. The count values are readily available in a format more suitable for scripting uses through PL variables (like with `*{?}` or `* VARNAME _`).

Although it doesn't effect scripts, I will mention a significant recent change to interactive commands which could confuse existing power users. Special and PL commands are now stored to the edit buffer and to command history, so they can be recalled and edited just like SQL commands. Now, all commands other than edit/history : commands are stored to the buffer and history.

New Features

- DSV column and row delimiters are now specified separately for input and output. This allows the input delimiters to be regular expressions. (Which in turn, allows for the next item).
- The new default DSM import row delimiter works for standard UNIX, Windows, Mac lines. This makes DSV files portable across these platforms. (If using a change control system like Subversion, CVS, or whatever, you can now change control your .dsv files as native ASCII files).
- Both `/* This kind */` and `-- This kind` of comments are now handled generally and intuitively, in SQL statements and in SqlTool-specific commands. (There were previously several limitations to where they could be used).
- At the cost of adding another *command type*, command aliases were replaced by *macros*. Usage (i.e., execution) is basically the same, but the new macros are much easier to define and list; and macros can be used for both PL and Special commands now (not just for SQL statements).
- Reports Transaction Isolation level and JDBC Connection Read/Write or Read-Only state connection or request (with `\j`).
- New `\t` command to report database execution duration times.
- New `\v` command to set or report the Connection's Transaction Isolation LeVel.
- `\d object filter` commands now use the filter as a regular expression, where possible, and filter may have optional prefix `/to` mean to match the filter against all output (not just the object name).
- `\dX filter` commands now use the filter as a regular expression, where possible.
- New `*DSV_TRIM_ALL` to automatically handle import of data which is both positional and character-delimited.
- New `\l` command to log user-specified messages with `java.util.logging` or Log4j logging facility. Nothing at all is written directly to stderr.
- All warnings and messages now use logging facility. This allows for declarative fine control of what gets logged and where the messages go to, as well as allowing for embedded apps to manage SqlTool apps in an integrated fashion with other app log entries.
- New `*DSV_RECORDS_PER_COMMIT` setting to support user-specified tuning of large DSV imports.
- (Optional) DSV log report can be customized with style sheets.
- You can interactively (or in SQL scripts) switch JDBC data sources (with `\j`). SqlTool can be started and used without any data source, though you'll obviously need to connect to a data source before issuing SQL commands.
- Array types are now supported, including in DSV imports and exports, with the exception that DSV imports do not support element values containing commas.

The Bare Minimum

The Bare Minimum You Need to Know to Run SqlTool



Warning

If you are using an Oracle database server, it will commit your current transaction if you cleanly disconnect, regardless of whether you have set auto-commit or not. This will occur if you exit SqlTool (or any other client) in the normal way (as opposed to killing the process or using Ctrl-C, etc.). This is mentioned in this section only for brevity, so I don't need to mention it in the main text in the many places where auto-commit is discussed. This behavior has nothing to do with SqlTool. It is a quirk of Oracle.

If you want to use SqlTool, then you either have an SQL text file, or you want to interactively type in SQL commands. If neither case applies to you, then you are probably looking at the wrong program.

Procedure 1.1. To run SqlTool...

1. Copy the file `sample/sqltool.rc` ¹ of your HyperSQL distribution to your home directory and secure access to it if your computer is accessible to anybody else (most likely from the network). This file will work as-is for a Memory Only database instance; or if your target is a HyperSQL Server running on your local computer with default settings and the password for the "SA" account is blank (the SA password is blank when new HyperSQL database instances are created). Edit the file if you need to change the target Server URL, username, password, character set, JDBC driver, or TLS trust store as documented in the `RC File Authentication Setup` section. You could, alternatively, use the `--inlineRc` command-line switch or the `\j` special command to connect up to a data source, as documented below.
2. Find out where your `sqltool.jar` file resides. It typically resides at `HSQLDB_HOME/lib/sqltool.jar` where `HSQLDB_HOME` is the "hsqldb" directory inside the root level of your HyperSQL software installation. (For example, if you extract `hsqldb-9.1.0.zip` into `c:\temp`, your `HSQLDB_HOME` would be `c:/temp/hsqldb-9.1.0/hsqldb`. Your file may also have a version label in the file name, like `sqltool-1.2.3.4.jar`. The forward slashes work just fine on Windows). For this reason, I'm going to use `"$HSQLDB_HOME/lib/sqltool.jar"` as the path to `sqltool.jar` for my examples, but understand that you need to use the actual path to your own `sqltool.jar` file. (Unix users may set a real env. variable if they wish, in which case the examples may be used verbatim; Window users may do the same, but will need to dereference the variables like `%THIS%` instead of like `$THIS`).



Warning

My examples assume there are no spaces or funky characters in your file paths. This avoids bugs with the Windows `cmd` shell and makes for simpler syntax all-around. If you insist on using directories with spaces or shell metacharacters (including standard Windows home directories like `C:\Documents and Settings\blaine`), you will need to double-quote arguments containing these paths. (On UNIX you can alternatively use single-quotes to avoid variable dereferencing at the same time).

3. If you are just starting with SqlTool, you are best off running your SqlTool command from a shell *command-line* (as opposed to by using icons or the Windows' Start/Run... or Start/Start Search). This way, you will be sure to see error messages if you type the command wrong or if SqlTool can't start up for some reason. On recent versions of Windows, you can get a shell by running `cmd` from Start/Run... or Start/Start Search). On UNIX or Linux, any real or virtual terminal will work.

On your shell command line, run

```
java -jar $HSQLDB_HOME/lib/sqltool.jar --help
```

to see what command-line arguments are available. Note that you don't need to worry about setting the CLASS-PATH when you use the `-jar` switch to `java`.

To run SqlTool without a JDBC connection, run

```
java -jar $HSQLDB_HOME/lib/sqltool.jar
```

You won't be able to execute any SQL, but you can play with the SqlTool interface (including using PL features).

To execute SQL, you'll need the classes for the target database's JDBC driver (and database engine classes for *in-process* databases). As this section is titled *The Bare Minimum*, I'll just say that if you are running SqlTool from a HyperSQL product installation, you are all set to connect to any kind of HyperSQL database. This is because SqlTool will look for the file `hsqldb.jar` in the same directory as `sqltool.jar`, and that file contains all of the needed classes. (SqlTool supports all JDBC databases and does not require a HyperSQL installation, but these cases would take us beyond *the bare minimum*). So, with `hsqldb.jar` in place, you can run commands like

```
java -jar $HSQLDB_HOME/lib/sqltool.jar mem
```

for interactive use, or

```
java -jar $HSQLDB_HOME/lib/sqltool.jar --sql="SQL statement;" mem
```

or

```
java -jar $HSQLDB_HOME/lib/sqltool.jar mem filepath1.sql...
```

where `mem` is an *urlid*, and the following arguments are paths to text SQL files. For the filepaths, you can use whatever wildcards your operating system shell supports.

The *urlid* `mem` in these commands is a key into your RC file, as explained in the RC File Authentication Setup section. Since this is a *mem*: type catalog, you can use SqlTool with this *urlid* immediately with no database setup whatsoever (however, you can't persist any changes that you make to this database). The sample `sqltool.rc` file also defines the *urlid* "localhost-sa" for a local HyperSQL Listener. At the end of this section, I explain how you can load some sample data to play with, if you want to.



Tip

If SqlTool fails to connect to the specified *urlid* and you don't know why, add the invocation parameter `--debug`. This will cause SqlTool to display a stack trace from where the connection attempt fails. (If a connection attempt fails with the interactive `\j` command, details will always be displayed).



Important

SqlTool does not *commit* SQL changes by default. (You can use the `--autoCommit` command-line switch to have it auto-commit). This leaves it to the user's discretion whether to commit or rollback their modifications. If you do want your changes committed, remember to run `\=` before quitting SqlTool. (Most databases also support the SQL command `commit` ;),

If you put a file named `auto.sql` into your home directory, this file will be executed automatically every time that you run SqlTool interactively (unless you invoke with the `--noAutoFile` switch).

To use a JDBC Driver other than the HyperSQL driver, you can't use the `-jar` switch because you need to modify the classpath. You must add the `sqltool.jar` file and your JDBC driver classes to your classpath, and you must tell SqlTool what the JDBC driver class name is. The latter can be accomplished by either using the `--driver` switch,

or setting "driver" in your config file. The [RC File Authentication Setup](#) section explains the second method. Here's an example of the first method (after you have set the classpath appropriately).

```
java org.hsqldb.cmdline.SqlTool --driver=oracle.jdbc.OracleDriver urlid
```



Tip

If the tables of query output on your screen are all messy because of lines wrapping, the best and easiest solution is usually to resize your terminal emulator window to make it wider. (With some terms you click & drag the frame edges to resize, with others you use a menu system where you can enter the number of columns).

Embedding

Using SqlTool to execute SQL files from your own Java code

To repeat what is stated in the JavaDoc for the `SqlTool` class itself: *Programmatic users will usually want to use the `objectMain(String[])` method if they want arguments and behavior exactly like command-line SqlTool. If you don't need invocation parameter parsing, `auto.sql` execution, etc., you will have more control and efficiency by using the `SqlFile` class directly. The file `src/org/hsqldb/sample/SqlFileEmbedder.java` in the HyperSQL distribution provides an example for this latter strategy.*

Non-displayable Types

There are some SQL types which SqlTool (being a text-based program) can't display properly. This includes the SQL types BLOB, JAVA_OBJECT, STRUCT, and OTHER. When you run a query that returns any of these, SqlTool will save the very first such value obtained to the binary buffer and will not display any output from this query. You can then save the binary value to a file, as explained in the [Storing and retrieving binary files](#) section.

There are other types, such as BINARY, which JDBC can make displayable (by using `ResultSet.getString()`), but which you may very well want to retrieve in raw binary format. You can use the `\b` command to retrieve any-column-type-at-all in raw binary format (so you can later store the value to a binary file).

Another restriction which all text-based database clients have is the practical inability for the user to type in binary data such as photos, audio streams, and serialized Java objects. You can use SqlTool to load any binary object into a database by telling SqlTool to get the insert/update datum from a file. This is also explained in the [Storing and retrieving binary files](#) section.

Desktop shortcuts

Desktop shortcuts and quick launch icons are useful, especially if you often run SqlTool with the same set of arguments. It's really easy to set up several of them-- one for each way that you invoke SqlTool (i.e., each one would start SqlTool with all the arguments for one of your typical startup needs). One typical setup is to have one shortcut for each database account which you normally use (use a different `urlid` argument in each shortcut's Target specification).

Desktop icon setup varies depending on your Desktop manager, of course. I'll explain how to set up a SqlTool startup icon in Windows XP. Linux and Mac users should be able to take it from there, since it's easier with the common Linux and Mac desktops.

Procedure 1.2. Creating a Desktop Shortcut for SqlTool

1. Right click in the main Windows background.
2. New
3. Shortcut

4. Browse
5. Navigate to where your good JRE lives. For recent Sun JRE's, it installs to `C:\Program Files\Java*\bin` by default (the * will be a JDK or JRE identifier and version number).
6. Select `java.exe`.
7. OK
8. Next
9. Enter any name
10. Finish
11. Right click the new icon.
12. Properties
13. Edit the Target field.
14. Leave the path to `java.exe` exactly as it is, including the quotes, but append to what is there. Beginning with a space, enter the command-line that you want run.
15. Change Icon... to a pretty icon.
16. If you want a quick-launch icon instead of (or in addition to) a desktop shortcut icon, click and drag it to your quick launch bar. (You may or may not need to edit the Windows Toolbar properties to let you add new items).

Loading sample data

If you want some sample database objects and data to play with, execute the `sample/sampledata.sql` SQL file ¹. To separate the sample data from your regular data, you can put it into its own schema by running this before you import:

```
CREATE SCHEMA sampledata AUTHORIZATION dba;  
SET SCHEMA sampledata;
```

Run it like this from an SqlTool session

```
\i HSQLDB_HOME/sample/sampledata.sql
```

where **HSQLDB_HOME** is the base directory of your HSQLDB software installation ¹.

For memory-only databases, you'll need to run this every time that you run SqlTool. For other (persistent) databases, the data will reside in your database until you drop the tables.

Satisfying SqlTool's CLASSPATH Requirements

As discussed earlier, only the single file `sqltool.jar` is required to run SqlTool (the file name may contain a version label like `sqltool-1.2.3.4.jar`). But it's useless as an SQL *Tool* unless you can connect to a JDBC data source, and for that you need the target database's JDBC driver in the classpath. For *in-process* catalogs, you'll also need the database engine classes in the CLASSPATH. The *The Bare Minimum* section explains that the easiest way to use SqlTool with any HyperSQL database is to just use `sqltool.jar` in-place where it resides in a HyperSQL installation. This section explains how to satisfy the CLASSPATH requirements for other setups and use cases.

Accessing older HSQLDB Databases with SqlTool

If you are using SqlTool to access non-HSQLDB database(s), then you should use the latest and greatest-- just grab the newest public release of SqlTool (like from the latest public HyperSQL release) and skip this subsection.

You are strongly encouraged to use the latest SqlTool release to access older HSQLDB databases, to enjoy greatly improved SqlTool robustness and features. It is very easy to do this.

1. Obtain the latest `sqltool.jar` file. One way to obtain the latest `sqltool.jar` file is to download the latest HyperSQL distribution and extract that single file
2. Place (or copy) your new `sqltool.jar` file right alongside the `hsqldb.jar` file for your target database version. If you don't have a local copy of the `hsqldb.jar` file for your target database, just copy it from your database server, or download the full distribution for that server version and extract it.
3. (If you have used older versions of SqlTool before, notice that you now invoke SqlTool by specifying the `sqltool.jar` file instead of the `hsqldb.jar`). If your target database is a previous 2.x version of HyperSQL, then you are finished and can use the new SqlTool for your older database. Users upgrading from a pre-2.x version please continue...

Run SqlTool like this.

```
java -jar path/to/sqltool.jar --driver=org.hsqldb.jdbcDriver...
```

where you specify the pre-2.x JDBC driver name `org.hsqldb.jdbcDriver`. Give any other SqlTool parameters as you usually would.

Once you have verified that you can access your database using the `--driver` parameter as explained above, edit your `sqltool.rc` file, and add a new line

```
driver org.hsqldb.jdbcDriver
```

after each urlid that is for a pre-2.x database. Once you do this, you can invoke SqlTool as usual (i.e. you no longer need the `--driver` argument for your invocations).

App-specific Classes, Embedding, and non-HyperSQL Databases

For these situations, you need to add your custom, third-party, or SQL driver classes to your Java CLASSPATH. Java doesn't support adding arbitrary elements to the classpath when you use the `-jar`, so you must set a classpath containing `sqltool.jar` plus whatever else you need, then invoke SqlTool without the `-jar` switch, as briefly described at the end of the `The Bare Minimum` section. For embedded apps, invoke your own main class instead of SqlTool, and you can invoke `SqlTool` or `SqlFile` from your code base.

To customize the classpath, you need to set up your classpath by using your operating system or shell variable `CLASSPATH` or by using the `java` switch `-cp` (or the equivalent `-classpath`). I'm not going to take up space here to explain how to set up a Java CLASSPATH. That is a platform-dependent task that is documented well in tons of Java introductions and tutorials. What I'm responsible for telling you is *what* you need to add to your classpath. For the non-embedded case where you have set up your CLASSPATH environmental variable, you would invoke SqlTool like this.

```
java org.hsqldb.cmdline.SqlTool ...
```

If you are using the `-cp` switch instead of a CLASSPATH variable, stick it after `java`. After "SqlTool", give any SqlTool parameters exactly as you would put after `java -jar .../sqltool.jar` if you didn't need to customize the CLASSPATH. You can specify a JDBC driver class to use either with the `--driver` switch to SqlTool, or in your RC file stanza (the last method is usually more convenient).

Note that without the `-jar` switch, SqlTool will still automatically pull in HyperSQL JDBC driver or engine classes from HyperSQL jar files in the same directory. It's often a good practice to minimize your runtime classpath. To prevent the possibility of pulling in classes from other HyperSQL jar files, just copy `sqltool.jar` to some other directory (which does not contain other HyperSQL jar files) and put the path to that one in your classpath.

Distributing SqlTool with your Apps

You can distribute SqlTool along with your application, for standalone or embedded invocation. For embedded use, you will need to customize the classpath as discussed in the previous item. Either way, you should minimize your application footprint by distributing only those HyperSQL jar files needed by your app. You will obviously need `sqltool.jar` if you will use the `SqlTool` or `SqlFile` class in any way. If your app will only connect to external HyperSQL listeners, then build and include `hsqldb.jar`. If your app will also *run* a HyperSQL Listener, you'll need to include `hsqldb.jar`. If your app will connect directly to a *in-process* catalog, then include `hsqldbmain.jar`. Note that you never need to include more than one of `hsqldb.jar`, `hsqldbmain.jar`, `hsqldb.jar`, since the former jars include everything in the following jars.

SqlTool Client PCs

If you just want to be able to run SqlTool (interactively or non-interactively) on a PC, and have no need for documentation, then it's usually easiest to just copy `sqltool.jar` and `hsqldb.jar` to the PCs (plus JDBC driver jars for any other target databases). If you want to minimize what you distribute, then build and distribute `hsqldb.jar` or `hsqldbmain.jar` instead of `hsqldb.jar`, according to the criteria listed in the previous sub-section.

RC File Authentication Setup

RC file authentication setup is accomplished by creating a text RC configuration file. In this section, when I say *configuration* or *config* file, I mean an RC configuration file. RC files can be used by any JDBC client program that uses the `org.hsqldb.util.RCData` class-- this includes `SqlTool`, `DatabaseManager`, `DatabaseManagerSwing`.

You can use it for your own JDBC client programs too. There is example code showing how to do this at `src/org/hsqldb/sample/SqlFileEmbedder.java`.

The sample RC file shown here resides at `sample/sqltool.rc` in your HSQLDB distribution ¹.

Example 1.1. Sample RC File

```
# $Id: sqltool.rc 3353 2009-12-15 19:52:13Z unsaved $

# This is a sample RC configuration file used by SqlTool, DatabaseManager,
# and any other program that uses the org.hsqldb.lib.RCData class.
# See the documentation for SqlTool for various ways to use this file.

# If you have the least concerns about security, then secure access to
# your RC file.

# You can run SqlTool right now by copying this file to your home directory
# and running
#   java -jar /path/to/sqltool.jar mem
# This will access the first urlid definition below in order to use a
# personal Memory-Only database.
# "url" values may, of course, contain JDBC connection properties, delimited
# with semicolons.
# As of revision 3347 of SqlFile, you can also connect to datasources defined
# here from within an SqlTool session/file with the command "\j urlid".

# You can use Java system property values in this file like this: ${user.home}

# The only feature added recently is the optional "transiso" setting,
# which may be set to an all-caps transaction isolation level as listed
# in the Java API Spec for java.sql.Connection.
# Windows users are advised to use forward slashes instead of reverse slashes,
# and to avoid paths containing spaces or other funny characters. (This
# recommendation applies to any Java app, not just SqlTool).

# A personal Memory-Only (non-persistent) database.
```

```
urlid mem
url jdbc:hsqldb:mem:memdbid
username SA
password

# A personal, local, persistent database.
urlid personal
url jdbc:hsqldb:file:${user.home}/db/personal;shutdown=true
username SA
password
transiso TRANSACTION_READ_COMMITTED
# When connecting directly to a file database like this, you should
# use the shutdown connection property like this to shut down the DB
# properly when you exit the JVM.

# This is for a hsqldb Server running with default settings on your local
# computer (and for which you have not changed the password for "SA").
urlid localhost-sa
url jdbc:hsqldb:hsq://localhost
username SA
password

# Template for a urlid for an Oracle database.
# You will need to put the oracle.jdbc.OracleDriver class into your
# classpath.
# In the great majority of cases, you want to use the file classes12.zip
# (which you can get from the directory $ORACLE_HOME/jdbc/lib of any
# Oracle installation compatible with your server).
# Since you need to add to the classpath, you can't invoke SqlTool with
# the jar switch, like "java -jar ../hsqldb.jar..." or
# "java -jar ../hsqsqltool.jar...".
# Put both the HSQLDB jar and classes12.zip in your classpath (and export!)
# and run something like "java org.hsqldb.util.SqlTool...".

#urlid cardiff2
#url jdbc:oracle:thin:@aegir.admc.com:1522:TRAFFIC_SID
#username blaine
#password secretpassword
#driver oracle.jdbc.OracleDriver

# Template for a TLS-encrypted HSQLDB Server.
# Remember that the hostname in hsqldb (and https) JDBC URLs must match the
# CN of the server certificate (the port and instance alias that follows
# are not part of the certificate at all).
# You only need to set "truststore" if the server cert is not approved by
# your system default truststore (which a commercial certificate probably
# would be).

#urlid tls
#url jdbc:hsqldb:hsqldb://db.admc.com:9001/1m2
#username BLAINE
#password asecret
#truststore ${user.home}/ca/db/db-trust.store

# Template for a Postgresql database
#urlid blainedb
#url jdbc:postgresql://idun.africawork.org/blainedb
#username blaine
#password losung1
#driver org.postgresql.Driver

# Template for a MySQL database. MySQL has poor JDBC support.
```

```

#urlid mysql-testdb
#url jdbc:mysql://hostname:3306/dbname
#username root
#password hiddenpwd
#driver com.mysql.jdbc.Driver

# Note that "databases" in SQL Server and Sybase are traditionally used for
# the same purpose as "schemas" with more SQL-compliant databases.

# Template for a Microsoft SQL Server database
#urlid msprojsvr
#url jdbc:microsoft:sqlserver://hostname;DatabaseName=DbName;SelectMethod=Cursor
# The SelectMethod setting is required to do more than one thing on a JDBC
# session (I guess Microsoft thought nobody would really use Java for
# anything other than a "hello world" program).
# This is for Microsoft's SQL Server 2000 driver (requires mssqlserver.jar
# and msutil.jar).
#driver com.microsoft.jdbc.sqlserver.SQLServerDriver
#username myuser
#password hiddenpwd

# Template for a Sybase database
#urlid sybase
#url jdbc:sybase:Tds:hostname:4100/dbname
#username blaine
#password hiddenpwd
# This is for the jConnect driver (requires jconn3.jar).
#driver com.sybase.jdbc3.jdbc.SybDriver

# Template for Embedded Derby / Java DB.
#urlid derbyl
#url jdbc:derby:path/to/derby/directory;create=true
#username ${user.name}
#password any_noauthbydefault
#driver org.apache.derby.jdbc.EmbeddedDriver
# The embedded Derby driver requires derby.jar.
# There's also the org.apache.derby.jdbc.ClientDriver driver with URL
# like jdbc:derby://<server>[:<port>]/databaseName, which requires
# derbyclient.jar.
# You can use \= to commit, since the Derby team decided (why???)
# not to implement the SQL standard statement "commit"!!
# Note that SqlTool can not shut down an embedded Derby database properly,
# since that requires an additional SQL connection just for that purpose.
# However, I've never lost data by not shutting it down properly.
# Other than not supporting this quirk of Derby, SqlTool is miles ahead of ij.

```

As noted in the comment (and as used in a couple examples), you can use Java system properties like this: `${user.home}`. Windows users, please read the suggestion directed to you in the file.

You can put this file anywhere you want to, and specify the location to SqlTool/DatabaseManager/DatabaseManagerSwing by using the `--rcfile` argument. If there is no reason to not use the default location (and there are situations where you would not want to), then use the default location and you won't have to give `--rcfile` arguments to SqlTool/DatabaseManager/DatabaseManagerSwing. The default location is `sqltool.rc` or `dbmanager.rc` in your home directory (corresponding to the program using it). If you have any doubt about where your home directory is, just run SqlTool with a phony urlid and it will tell you where it expects the configuration file to be.

```
java -jar $HSQLDB_HOME/lib/sqltool.jar x
```

The config file consists of stanza(s) like this:

```

urlid web
url jdbc:hsqldb:hsq://localhost
username web
password webpassword

```

These four settings are required for every urlid. (There are optional settings also, which are described a couple paragraphs down). The URL may contain JDBC connection properties. You can have as many blank lines and comments like

```
# This comment
```

in the file as you like. The whole point is that the *urlid* that you give in your SqlTool/DatabaseManager command must match a *urlid* in your configuration file.



Important

Use whatever facilities are at your disposal to protect your configuration file.

It should be readable, both locally and remotely, only to users who run programs that need it. On UNIX, this is easily accomplished by using `chmod/chown` commands and making sure that it is protected from anonymous remote access (like via NFS, FTP or Samba).

You can also put the following optional settings into a urlid stanza. The setting will, of course, only apply to that urlid.

charset	This is used by the SqlTool program, but not by the DatabaseManager programs. See the Character Encoding section of the Non-Interactive section. This is used for input and output files, not for stdin or stdout, which are controlled by environmental variables and Java system properties. If you set no encoding for an urlid, input and outfiles will use the same encoding as for stdin/stdout. (As of right now, the charset setting here is not honored by the <code>\j</code> command, but only when SqlTool loads an urlid specified on the command-line).
driver	Sets the JDBC driver class name. You can, alternatively, set this for one SqlTool/DatabaseManager invocation by using the command line switch <code>--driver</code> . Defaults to <code>org.hsqldb.jdbc.JDBC.Driver</code> .
truststore	TLS trust keystore store file path as documented in the TLS section of the Listeners chapter of the HyperSQL User Guide [http://hsqldb.org/doc/2.0/guide/index.html] You usually only need to set this if the server is using a non-publicly-certified certificate (like a self-signed self-ca'd cert). Relative paths will be resolved relative to the <code>\${user.dir}</code> system property at JRE invocation time.
transiso	Specify the Transaction Isolation Level with an all-caps string, exactly as listed in the Field Summary of the Java API Spec for the class <code>java.sql.Connection</code> .

Property and SqlTool command-line switches override settings made in the configuration file.

Switching Data Sources

The `\j` command lets you switch JDBC Data Sources in your SQL files (or interactively). "`\?`" shows the syntax to make a connection by either RCData urlid or by name + password + JDBC Url. The urlid variant uses RC file of `$HOME/sqltool.rc`. We will add a way to specify an RC file if there is any demand for that.

You can start SqlTool without any JDBC Connection by specifying no Inline RC and urlid of `"-`" (just a hyphen). If you don't need to specify any SQL file paths, you can skip the hyphen, as in this example.

```
java -jar $HSQLDB_HOME/lib/sqltool.jar --setVar=v1=one
```

(The `"-`" is required when specifying one or more SQL files, in order to distinguish urlid-spec from file-spec). Consequently, if you invoke SqlTool with no parameters at all, you will get a SqlTool session with no JDBC Connection. You will obviously need to use `\j` before doing any database work.

Using Inline RC Authentication

Inline RC authentication setup is accomplished by using the `--inlineRc` command-line switch on SqlTool. The `--inlineRc` command-line switch takes a comma-separated list of key/value elements. The `url` and `user` elements are required. The rest are optional. The `--inlineRc` switch is the only case where you can give SQL file paths without a preceding urlid indicator (an urlid or `-`). The program knows not to look for an urlid if you give an inline.

<code>url</code>	The JDBC URL of the database you wish to connect to.
<code>user</code>	The username to connect to the database as.
<code>charset</code>	Sets the character encoding. Overrides the platform default, or what you have set by env variables or Java system properties. (Does not effect stdin or stdout).
<code>truststore</code>	The TLS trust keystore file path as documented in the TLS chapter. Relative paths will be resolved relative to the current directory.
<code>transiso</code>	<code>java.sql.Connection</code> transaction isolation level to connect with, as specified in the Java API spec.
<code>password</code>	You may only use this element to set empty password, like

```
password=
```

For any other password value, omit the `password` element and you will be prompted for the value.

(Use the `--driver` switch instead of `--inlineRc` to specify a JDBC driver class). Here is an example of invoking SqlTool to connect to a standalone database.

```
java -jar $HSQLDB_HOME/lib/sqltool.jar --inlineRc=url=jdbc:hsqldb:file:/home/dan/dandb,user=dan
```

For security reasons, you cannot specify a non-empty password as an argument. You will be prompted for a password as part of the login process.

Logging

Both the `\l` command and all warnings and error messages now use a logging facility. The logging facility hands off to Log4j if Log4j is found in the classpath, and otherwise will hand off to `java.util.logging`. The default behavior of `java.util.logging` should work fine for most users. If you are using log4j and are redirecting with pipes, you may want to configure a Console Appender with target of `"System.err"` so that error output will go to the error stream (all console output for `java.util.logging` goes to `stderr` by default). See the API specs for Log4j and for J2SE for how to configure either product. If you are embedding SqlTool in a product to process SQL files, I suggest that you use log4j. `java.util.logging` is neither scalable nor well-designed.

Run the command `\l?` to see how to use the logging command `\l` in your SQL files (or interactively), including what logging levels you may specify.

Interactive Usage

Do read the `The Bare Minimum` section before you read this section.

You run SqlTool interactively by specifying no SQL filepaths on the SqlTool command line. Like this.

```
java -jar $HSQLDB_HOME/lib/sqltool.jar urlid
```

Procedure 1.3. What happens when SqlTool is run interactively (using all default settings)

1. SqlTool starts up and connects to the specified database, using your SqlTool configuration file (as explained in the RC File Authentication Setup section).
2. SQL file `auto.sql` in your home directory is executed (if there is one),
3. SqlTool displays a banner showing the SqlTool and SqlFile version numbers and describes the different command types that you can give, as well as commands to list all of the specific commands available to you.

You exit your session by using the `"\q"` special command or ending input (like with Ctrl-D or Ctrl-Z).



Important

Any command may be preceded by space characters. Special Commands, Edit Buffer Commands, PL Commands, Macros always consist of just one line.

These rules do not apply at all to Raw Mode. Raw mode is for use by advanced users when they want to completely bypass SqlTool processing in order to enter a chunk of text for direct transmission to the database engine.

SqlTool Command-Line Editing

If you are really comfortable with `grep`, `perl`, or `vim`, you will instantly be an expert with SqlTool command-line editing. Due to limitations of Java I/O, we can't use up-arrow recall, which many people are used to from `DosKey` and `Bash shell`. If you don't know how to use regular expressions, and don't want to learn how to use them, then just forget `command-recall`.

Basic command entry (i.e., without regexps)

- Just type in your command, and use the backspace-key to fix mistakes on the same line.
- If you goof up a multi-line command, just hit the ENTER key twice to start over. (The command will be moved to the buffer where it will do no harm).
- Use the `":h"` command to view your command history. You can use your terminal emulator scroll bar and copy and paste facility to repeat commands.
- As long as you don't need to change text that is already in a command, you can easily repeat commands from the history like `":14;"` to re-run command number 14 from history.
- Expanding just a bit from the previous item, you can add on to a previous command by running a command like `":14a"` (where the "a" means *append*).
- See the `Macros` section about how to set and use macros.

If you use regular expressions to search through your command history, or to modify commands, be aware that the command type of commands in history are fixed. You can search and modify the text after a `\` or `*` prefix (if any), but you can't search on or change a prefix (or add or remove one).

Command Types

When you are typing into SqlTool, you are always typing part of the *immediate command*. If the immediate command is an SQL statement, it is executed as soon as SqlTool reads in the trailing (unquoted) semi-colon. Commands of the

other command types are executed as soon as you hit ENTER. The interactive `:` commands can perform actions with or on the edit buffer. The *edit buffer* usually contains a copy of the last command executed, and you can always view it with the `:b` command. If you never use any `:` commands, you can entirely ignore the edit buffer. If you want to repeat commands or edit previous commands, you will need to work with the edit buffer. The immediate command contains whatever (and exactly what) you type. The command history and edit buffer may contain any type of command other than comments and `:` commands (i.e., `:` commands and comments are just not copied to the history or to the edit buffer).

Hopefully an example will clarify the difference between the immediate command and the edit buffer. If you type in the edit buffer Substitution command `:s/tbl/table/`, the `:s` command that you typed is the immediate command (and it will never be stored to the edit buffer or history, since it is a `:` command), but the purpose of the substitution command is to modify the contents of the edit buffer (perform a substitution on it)-- the goal being that after your substitutions you would execute the buffer with the `:";` command. The `:"a` command is special in that when you hit ENTER to execute it, it copies the contents of the edit buffer to a new immediate command and leaves you in a state where you are *appending* to that *immediate* command (nearly) exactly as if you had just typed it in.

Command Types

Command types

SQL Statement

Any command that you enter which does not begin with `"\`", `:"`", `* "` or `/"` is an SQL Statement. The command is not terminated when you hit ENTER, like most OS shells. You terminate SQL Statements with either `;"` or with a blank line. In the former case, the SQL Statement will be executed against the SQL database and the command will go into the edit buffer and SQL command history for editing or viewing later on. In the latter case, *execute against the SQL database* means to transmit the SQL text to the database engine for execution. In the latter case (you end an SQL Statement with a blank line), the command will go to the edit buffer and SQL history, but will not be executed (but you can execute it later from the edit buffer).

(Blank lines are only interpreted this way when SqlTool is run interactively. In SQL files, blank lines inside of SQL statements remain part of the SQL statement).

As a result of these termination rules, whenever you are entering text that is not a Special Command, Edit Buffer / History Command, or PL Command, you are always *appending* lines to an SQL Statement or comment. (In the case of the first line, you will be appending to an empty SQL statement. I.e. you will be starting a new SQL Statement or comment).

Special Command

Run the command `"\?"` to list the Special Commands. All of the Special Commands begin with `"\`". I'll describe some of the most useful Special Commands below.

Edit Buffer / History Command

Run the command `:"?"` to list the Edit-Buffer/History Commands. All of these commands begin with `:"`". These commands use commands from the command history, or operate upon the edit "buffer", so that you can edit and/or (re-)execute previously entered commands.

PL Command

Procedural Language commands. Run the command `*?"` to list the PL Commands. All of the PL Commands begin with `* "`". PL commands are for setting and using scripting variables and conditional and flow control statements like `* if` and `* while`. A few PL features (such as macros and updating and selecting data directly from/to files) can be a real convenience for nearly all users,

so these features will be discussed briefly in this section. More detailed explanation of PL variables and the other PL features, with examples, are covered in the `SqlTool Procedural Language` section.

Macro Command	Macro definition and usage commands. Run the command <code>"/?"</code> to show the define, list, or use macros.
Raw Mode	The descriptions of command-types above do not apply to <code>Raw Mode</code> . In raw mode, SqlTool doesn't interpret what you type at all. It all just goes into the edit buffer which you can send to the database engine. Beginners can safely ignore raw mode. You will never encounter it unless you run the <code>\".</code> special command, or define a stored procedure or function. See the <code>Raw Mode</code> section for the details.

Special Commands

Essential Special Commands

<code>\?</code>	help
<code>\q</code>	quit
<code>\j...</code>	View JDBC Data Source details or connect up to a JDBC Data Source (replacing the current connection, if any). Run <code>\?</code> to see the syntax for the different usages.
<code>\i path/to/script.sql</code>	execute the specified SQL script, then continue again interactively. Since SqlTool is a Java program, you can safely use forward slashes in your file paths, regardless of your operating system.
<code>\=</code>	commit the current SQL transaction. Most users are used to typing the SQL statement <code>commit ;</code> , but this command is crucial for those databases which don't support the statement. It's obviously unnecessary if you have auto-commit mode on.
<code>\x?</code>	List a summary of DSV eXporting, and all available DSV options.
<code>\m?</code>	List a summary of DSV iMporting, and all available DSV options.
<code>\d?</code>	List a summary of the <code>\d</code> commands below.
<code>\dt [filter_substring]</code>	
<code>\dv [filter_substring]</code>	
<code>\ds [filter_substring]</code>	
<code>\di [table_name]</code>	
<code>\dS [filter_substring]</code>	
<code>\da [filter_substring]</code>	
<code>\dn [filter_substring]</code>	

\du [filter_substring]

\dr [filter_substring]

\d* [filter_substring]

Lists available objects of the given type.

- t: non-system Tables
- v: Views
- s: Sequences
- i: Indexes
- S: System tables
- a: Aliases
- n: schema Names
- u: database Users
- r: Roles
- *: all table-like objects

If your database supports schemas, then the schema name will also be listed.

If you supply an optional *filter substring*, then only items which match the specified substring. will be listed. In most cases, the specified filter will be treated as a regular expression matched against the candidate object names. In order to take advantage of extreme server-side performance benefits, however, in some cases the substring is passed to the database server and the filter will be processed by the server.



Important

The regexp test is case-sensitive! Even though in SQL queries and for the "\d objectname" command object names are usually case-insensitive, for the \dX commands, you must capitalize the filter substring exactly as it will appear in the special command output. This is an inconvenience, since the database engine will change names in SQL to default case unless you double-quote the name, but that is server-side functionality which cannot (portably) be reproduced by SqlTool. You can use spaces and other special characters in the string.



Tip

Filter substrings ending with "." are special. If a substring ends with ".", then this means to narrow the search by the exact, case-sensitive schema name given. For example, if I run "\d* BLAINE.", this will list all table-like database objects in the "BLAINE" schema. The capitalization of the schema must be exactly the same as how the schema name is listed by the "\dn" command. You can use spaces and other special characters in the

string. (I.e., enter the name exactly how you would enter it inside of double-quotes in an SQL command). This is an inconvenience, since the database engine will change names in SQL to default case unless you double-quote the name, but that is server-side functionality which cannot (portably) be reproduced by SqlTool.



Tip

The filter string "." (just a plain dot) means the current session schema, for databases which support the concept according to the SQL standard (HyperSQL database does).



Important

Indexes may not be searched for by *substring*, only by exact target table name. So if I1 is an index on table T1, then you list this index by running "\di T1". In addition, many database vendors will report on indexes only if a target table is identified. Therefore, "\di" with no argument will fail if your database vendor does not support it.

\d objectname [[/]regexp]

Lists names of columns in the specified table or view. `objectname` may be a base table name or a schema.object name.

If you supply a filter string, then only columns with a name matching the given regular expression will be listed. (If no special characters are used, this just means that names containing the specified substring will match). You'll find this filter is a great convenience compared to other database utilities, where you have to list all columns of large tables when you are only interested in one of them.

To narrow the displayed information based on all column outputs, instead of just the column names, just prefix the expression with /. For example, to list all INTERGER columns, you could run `\d mytable /INTEGER`.



Tip

When working with real data (as opposed to learning or playing), I often find it useful to run two SqlTool sessions in two side-by-side terminal emulator windows. I do all of my real work in one window, and use the other mostly for \d commands. This way I can refer to the data dictionary while writing SQL commands, without having to scroll.

This list here includes only the *essential* Special Commands, but n.b. that there are other useful Special Commands which you can list by running \?. (You can, for example, execute SQL from external SQL files, and save your interactive SQL commands to files). Some specifics of these other commands are specified immediately below, and the Generating Text or HTML Reports section explains how to use the "\o" and "\H" special commands to generate reports.

Be aware that the \! Special Command does not work for external programs that read from standard input. You can invoke non-interactive and graphical interactive programs, but not command-line interactive programs.

SqlTool executes \! programs directly, it does not run an operating system shell (this is to avoid OS-specific code in SqlTool). Because of this, you can give as many command-line arguments as you wish, but you can't use shell wildcards or redirection.

The `\w` command can be used to store any command in your SQL history to a file. Just restore the command to the edit buffer with a command like `"\-4"` before you give the `\w` command.

Edit Buffer / History Commands

Edit Buffer / History Commands

<code>?:</code>	help
<code>:b</code>	List the current contents of the edit buffer.
<code>:h</code>	Shows the Command History. For each command which has been executed (up to the max history length), the SQL command history will show the command; its command number (#); and also how many commands <i>back</i> it is (as a negative number). <code>:</code> commands are never added to the history list. You can then use either form of the command identifier to recall a command to the edit buffer (the command described next) or as the target of any of the following <code>:</code> commands. This last is accomplished in a manner very similar to the vi editor. You specify the target command number between the colon and the command. As an example, if you gave the command <code>:s/X/Y/</code> , that would perform the substitution on the contents of the edit buffer; but if you gave the command <code>:-3 s/X/Y/</code> , that would perform the substitution on the command 3 back in the command history (and copy the output to the edit buffer). Also, just like vi, you can identify the command to recall by using a regular expression inside of slashes, like <code>:/blue/ s/X/Y/</code> to operate on the last command you ran which contains "blue".
<code>:13 OR :-2 OR :/blue/</code>	Recalls a command from Command history to the edit buffer. Enter ":" followed by the positive command number from Command history, like ":13"... or ":" followed by a negative number like ":-2" for two commands back in the Command history... or ":" followed by a regular expression inside slashes, like <code>:/blue/</code> to recall the last command which contains "blue". The specified command will be written to the edit buffer so that you can execute it or edit it using the commands below. As described under the <code>:h</code> command immediately above, you can follow the command number here with any of the commands below to perform the given operation on the specified command from history instead of on the edit buffer contents. So, for example, <code>:4;</code> would load command 4 from history then execute it (see the <code>:"</code> command below).
<code>::</code>	Executes the SQL, Special or PL statement in the edit buffer (by default). This is an extremely useful command. It's easy to remember because it consists of ":", meaning <i>Edit Buffer Command</i> , plus a line-terminating ";", (which generally means to execute an SQL statement, though in this case it will also execute a special or PL command).
<code>:a</code>	Enter append mode with the contents of the edit buffer (by default) as the current command. When you hit ENTER, things will be nearly exactly the same as if you physically re-typed the command that is in the edit buffer. Whatever lines you type next will be appended to the immediate command. As always, you then have the choice of hitting ENTER to execute a Special or PL command, entering a blank line to store back to the edit buffer, or end a SQL statement with semi-colon and ENTER to execute it.

You can, optionally, put a string after the :a, in which case things will be exactly as just described except the additional text will also be appended to the new immediate command. If you put a string after the :a which ends with ;, then the resultant new immediate command will just be executed right away, as if you typed in and entered the entire thing.

If your edit buffer contains `SELECT x FROM mytab` and you run `a:le`, the resultant command will be `SELECT x FROM mytable`. If your edit buffer contains `SELECT x FROM mytab` and you run `a: ORDER BY y`, the resultant command will be `SELECT x FROM mytab ORDER BY y`. Notice that in the latter case the append text begins with a space character.

You may notice that you can't use the left-arrow or backspace key to back up over the original text. This is due to Java and portability constraints. If you want to edit existing text, then you shouldn't use the Append command.

:s/from regex/to string/switches

The Substitution Command is the primary method for SqlTool command editing-- it operates upon the current edit buffer by default. The "to string" and the "switches" are both optional (though the final "/" is not). To start with, I'll discuss the use and behavior if you don't supply any substitution mode switches.

Don't use "/" if it occurs in either "from string" or "to string". You can use any character that you want in place of "/", but it must not occur in the *from* or *to* strings. Example

```
:s@from string@to string@
```

The *to string* is substituted for the first occurrence of the (case-specific) *from string*. The replacement will consider the entire SQL statement, even if it is a multi-line statement.

In the example above, the from regex was a plain string, but it is interpreted as a regular expression so you can do all kinds of powerful substitutions. See the `perlre` man page, or the `java.util.regex.Pattern` [<http://http://java.sun.com/javase/6/docs/api/java/util/regex/Pattern.html>] API Spec for everything you need to know about extended regular expressions.

Don't end a *to* string with ";" in attempt to make a command execute. There is a substitution mode switch to use for that purpose.

You can use any combination of the substitution mode switches.

- Use "i" to make the searches for *from regex* case insensitive.
- Use "g" to substitute Globally, i.e., to substitute *all* occurrences of the *from regex* instead of only the first occurrence found.
- Use ";" to execute the command immediately after the substitution is performed.
- Use "m" for ^ and \$ to match each line-break in a multi-line edit buffer, instead of just at the very beginning and every end of the entire buffer.

If you specify a command number (from the command history), you end up with a feature very reminiscent of vi, but even more powerful, since the Perl/Java regular expression are a superset of the vi regular expressions. As an example,

```
:24 s/pin/needle/g;
```

would start with command number 24 from command history, substitute "needle" for all occurrences of "pin", then execute the result of that substitution (and this final statement will of course be copied to the edit buffer and to command history).

```
:w /path/to/file.sql
```

This appends the contents of the current buffer (by default) to the specified file. Since what is being written are Special, PL, or SQL commands, you are effectively creating an SQL script.

I find the `:/regex/` and `:/regex/;` constructs particularly handy for every-day usage.

```
:/\d/;
```

re-executes the last `\d` command that you gave (The extra `\` is needed to escape the special meaning of `\` in regular expressions). It's great to be able to recall and execute the last "insert" command, for example, without needing to check the history or keep track of how many commands back it was. To re-execute the last insert command, just run `:/insert/;`. If you want to be safe about it, do it in two steps to verify that you didn't accidentally recall some other command which happened to contain the string "insert", like

```
:/insert/  
:;
```

(Executing the last only if you are satisfied when SqlTool reports what command it restored). Often, of course, you will want to change the command before re-executing, and that's when you combine the `:s` and `:a` commands.

We'll finish up with a couple fine points about Edit/Buffer commands. You generally can't use PL variables in Edit/Buffer commands, to eliminate possible ambiguities and complexities when modifying commands. The `:w` command is an exception to this rule, since it can be useful to use variables to determine the output file, and this command does not do any "editing".

The `?:` help explains how you can change the default regular expression matching behavior (case sensitivity, etc.), but you can always use syntax like `(?i)` inside of your regular expression, as described in the Java API spec for class `java.util.regex.Pattern` [<http://http://java.sun.com/javase/6/docs/api/java/util/regex/Pattern.html>]. History-command-matching with the `/regex/` construct is purposefully liberal, matching any portion of the command, case sensitive, etc., but you can still use the method just described to modify this behavior. In this case, you could use `(?-i)` at the beginning of your regular expression to be case-sensitive.

PL Commands

Essential PL Command

* VARNAME = value

Set the value of a variable. If the variable doesn't exist yet, it will be created. The most common use for this is so that you can later use it in SQL statements, print statements, and PL conditionals, by using the `*{VARNAME}` or `*{:VARNAME}` construct. The only difference between `*{literal}` and `*{:VARNAME}` is that the former produces an error if VARNAME is not set, whereas the latter will expand to a zero-length string if VARNAME is not set.

If you put variable definitions into the SQL file `auto.sql` in your home directory, those variables will always be available for interactive use.

PL variables can be expanded within all commands other than `: edit/history` commands.

- * `load VARNAME /file/path.txt` Sets VARNAME to the content of the specified ASCII file.
- * `prepare VARNAME` Indicate that next command should be a SQL INSERT or UPDATE command containing one question mark. The value of VARNAME will be substituted for the ? variable. This does work for CLOB columns.
- * `VARNAME _` When next SQL command is run, instead of displaying the rows, just store the very first column value to variable VARNAME. This works for CLOB columns too. It also works with Oracle XML type columns if you use column labels and the `getclobval` function.
- * `VARNAME ~` Exactly the same as

* VARNAME ~

 except that the fetched results will be displayed in addition to setting the variable.
- * `dump VARNAME /file/path.txt` Store the value of VARNAME to the specified ASCII file.

? Variable

You don't set the ? variable. It is just like the Bourne shell variable \$? in that it is always automatically set to the first value of a result set (or the return value of other SQL commands). It works just like the `* VARNAME ~` command described above, but it all happens automatically. You can, of course, dereference ? like any PL variable, but it does not list with the `list` and `listvalues` commands. You can see the value whenever you want by running

```
\p *{?}
```

Note that PL commands are used to upload and download column values to/from local ASCII files, but the corresponding actions for binary files use the special `\b` commands. This is because PL variables are used for ASCII values and you can store any number of column values in PL variables. This is not true for binary column values. The `\b` commands work with a single binary byte buffer.

See the [SqlTool Procedural Language](#) section below for information on using variables in other ways, and information on the other PL commands and features.

Storing and retrieving binary files

You can upload binary files such as photographs, audio files, or serialized Java objects into database columns. SqlTool keeps one binary buffer which you can load from files with the `\bl` command, or from a database query by doing a one-row query for any non-displayable type (including BLOB, OBJECT, and OTHER). In the latter case, the data returned for the first non-displayable column of the first result row will be stored into the binary buffer.

Once you have data in the binary buffer, you can upload it to a database column (including BLOB, OBJECT, and OTHER type columns), or save it to a file. The former is accomplished by the special command `\bp` followed by a *prepared* SQL query containing one question mark place-holder to indicate where the data gets inserted. The latter is accomplished with the `\bd` command.

You can also store the output from normal, displayable column into the binary buffer by using the special command `\b`. The very first column value from the first result row of the next SQL command will be stored to the binary byte buffer.

Example 1.2. Inserting binary data into database from a file

```
\bl /tmp/favoritesong.mp3
\bpb
INSERT INTO musictbl (id, stream) VALUES(3112, ?);
```

Example 1.3. Downloading binary data from database to a file

```
SELECT stream FROM musictbl WHERE id = 3112;
\bdb /tmp/favoritesong.mp3
```

You can also store and retrieve text column values to/from ASCII files, as documented in the [Essential PL Command](#) section.

Command History

The SQL history shown by the `\h` command, and used by other commands, is truncated to 100 entries, since its utility comes from being able to quickly view the history list. You can change the history length by setting the system property `sqltool.historyLength` to the desire integer value (using any of the System Property mechanisms provided by Java). If there is any demand, I'll make the setting of this value more convenient.

The SQL history list contains all executed commands other than Edit Buffer commands and comments, even if the command has a syntax error or fails upon execution. The reason for including bad commands is so that you can recall and fix them if you wish to. The same applies to the edit buffer. If you copy a command to the edit buffer by entering blank line, or if you edit the edit buffer, that edit buffer value will never make it into the command history until and if you execute it.

Shell scripting and command-line piping

You normally use non-interactive mode for input piping. You specify `"-"` as the SQL file name. See the [Piping and shell scripting](#) subsection of the [Non-Interactive](#) chapter.

Emulating Non-Interactive mode

You can run SqlTool *interactively*, but have SqlTool behave exactly as if it were processing an SQL file (i.e., no command-line prompts, error-handling that defaults to fail-upon-error, etc.). Just specify `"-"` as the SQL file name in the command line. This is a good way to test what SqlTool will do when it encounters any specific command in an SQL file. See the [Piping and shell scripting](#) subsection of the [Non-Interactive](#) chapter for an example.

Non-Interactive

Read the [Interactive Usage](#) section if you have not already, because much of what is in this section builds upon that. You can skip all discussion about Command History and the edit buffer if you will not use those interactive features.



Important

If you're doing data updates, remember to issue a commit command or use the `--autoCommit` switch.

As you'll see, SqlTool has many features that are very convenient for scripting. But what really makes it superior for automation tasks (as compared to SQL tools from other vendors) is the ability to reliably detect errors and to control JDBC transactions. SqlTool is designed so that you can reliably determine if errors occurred within SQL scripts

themselves, and from the invoking environment (for example, from a perl, Bash, or Python script, or a simple cron tab invocation).

Giving SQL on the Command Line

If you just have a couple Commands to run, you can run them directly from the comand-line or from a shell script without an SQL file, like this.

```
java -jar $HSQLDB_HOME/lib/sqltool.jar --sql="SQL statement;" urlid
```



Note

The `--sql` switch automatically implies `--noinput`, so if you want to execute the specified SQL before *and in addition to* an interactive session (or stdin piping), then you must also give the `--stdininput` switch.

Since SqlTool transmits SQL statements to the database engine only when a line is terminated with ";", if you want feedback from multiple SQL statements in an `--sql` expression, you will need to use functionality of your OS shell to include linebreaks after the semicolons in the expression. With any Bourne-compatible shell, you can include linebreaks in the SQL statements like this.

```
java -jar $HSQLDB_HOME/lib/sqltool.jar --sql='
SQL statement number one;
SQL statement
number two;
SQL statement three;
' urlid
```



Note

The multi-line examples in this section will only work as-is with a Bourne-compatible shell. With some ugliness they can be converted to C shell. For Windows, you are better off to stick with SQL files for multi-line input.

If you don't need feedback, just separate the SQL commands with semicolons and the entire expression will be chunked .

The `--sql` switch is very useful for setting shell variables to the output of SQL Statements, like this.

```
# A shell script
USERCOUNT=`java -jar $HSQLDB_HOME/lib/sqltool.jar --sql='
select count(*) from usertbl;
' urlid` || {
# Handle the SqlTool error
}
echo "There are $USERCOUNT users registered in the database."
[ "$USERCOUNT" -gt 3 ] && { # If there are more than 3 users registered
# Some conditional shell scripting
```

SQL Files

Just give paths to sql text file(s) on the command line after the *urlid*.

Often, you will want to redirect output to a file, like

```
java -jar $HSQLDB_HOME/lib/sqltool.jar urlid file.sql... > /tmp/file.log 2>&1
```

You can also execute SQL files from an interactive session with the "\i" Special Command, but be aware that the default behavior in an interactive session is to continue upon errors. If the SQL file was written without any concern for error handling, then the file will continue to execute after errors occur. You could run `\c false` before `\i filename`, but then your SqlTool session will exit if an error is encountered in the SQL file. If you have an SQL file without error handling, and you want to abort that file when an error occurs, but not exit SqlTool, the easiest way to accomplish this is usually to add `\c false` to the top of the script.

If you specify multiple SQL files on the command-line, the default behavior is to exit SqlTool immediately if any of the SQL files encounters an error.

SQL files themselves have ultimate control over error handling. Regardless of what command-line options are set, or what commands you give interactively, if a SQL file gives error handling statements, they will take precedence.

You can also use `\i` in SQL files. This results in nested SQL files.

You can use the following SQL file, `sample/sample.sql`, from your HyperSQL distribution ¹. It contains SQL as well as Special Commands making good use of most of the Special Commands documented below.

```
/*
  $Id: sample.sql 3637 2010-06-07 00:59:13Z unsaved $
  Exemplifies use of SqlTool.
  PCTASK Table creation
*/

/* Ignore error for these two statements */
\c true
DROP TABLE pctasklist;
DROP TABLE pctask;
\c false

\p Creating table pctask
CREATE TABLE pctask (
  id integer identity,
  name varchar(40),
  description varchar(256),
  url varchar(80),
  UNIQUE (name)
);

\p Creating table pctasklist
CREATE TABLE pctasklist (
  id integer identity,
  host varchar(20) not null,
  tasksequence int not null,
  pctask integer,
  assigndate timestamp default current_timestamp,
  completedate timestamp,
  show boolean default true,
  FOREIGN KEY (pctask) REFERENCES pctask,
  UNIQUE (host, tasksequence)
);

\p Granting privileges
GRANT select ON pctask TO public;
GRANT all ON pctask TO tomcat;
GRANT select ON pctasklist TO public;
GRANT all ON pctasklist TO tomcat;

\p Inserting test records
INSERT INTO pctask (name, description, url) VALUES (
  'task one', 'Description for task 1', 'http://cnn.com');
INSERT INTO pctasklist (host, tasksequence, pctask) VALUES (
  'admc-masq', 101, (SELECT id FROM pctask WHERE name = 'task one'));
```

```
commit;
```

You can execute this SQL file with a Memory Only database with a command like

```
java -jar $HSQLDB_HOME/lib/sqltool.jar --sql='
create user tomcat password "x";
' mem path/to/hsqlldb/sample/sample.sql
```

This shows how you can mix SQL on the command line, and SQL inside an SQL file.



Note

The example above uses Bourne shell syntax. C shell syntax would be similar. You would need to use an SQL file to accomplish this on Windows.

(The `--sql="create...;"` arguments create an account which the script uses). You should see error messages between the `Continue-on-error...true` and `Continue-on-error...false`. The script purposefully runs commands that might fail there. The reason the script does this is to perform database-independent conditional table removals. (The SQL clause `IF EXISTS` is more graceful and succinct, so you may want to use that if you don't need to support databases which don't support `IF EXISTS`). If an error occurs when `continue-on-error` is false, the script would abort immediately.

Piping and shell scripting

You can of course, redirect output *from* SqlTool to a file or another program.

```
java -jar $HSQLDB_HOME/lib/sqltool.jar urlid file.sql > file.txt 2>&1
java -jar $HSQLDB_HOME/lib/sqltool.jar urlid file.sql 2>&1 | someprogram...
```

You can type commands in to SqlTool while being in non-interactive mode by supplying "-" as the file name. This is a good way to test how SqlTool will behave when processing your SQL files.

```
java -jar $HSQLDB_HOME/lib/sqltool.jar urlid -
```

This is how you have SqlTool read its input from another program:

Example 1.4. Piping input into SqlTool

```
echo "Some SQL commands with '$VARIABLES';" |
java -jar $HSQLDB_HOME/lib/sqltool.jar urlid -
```

For a shell not as graceful as the Bourne-compatible shells, you would need to type this all on the same line (or use a line-continuation trick).



Warning

Beware of null stdin to SqlTool (or SqlFile). At least with Java 6 on UNIX, `System.in` returns megabytes of garbage for reads if stdin is closed. I consider this an obvious bug. Therefore, unlike any other program you would invoke from scripts, check stdin before running any Java program that will read from it. I consider this a big ugly bug in Java. This is not just theoretical, because many remote execution environments will have stdin closed off.

Make sure that you also read the [Giving SQL on the Command Line](#) section. The `--sql` switch is a great facility to use with shell scripts.

Optimally Compatible SQL Files

If you want your SQL scripts optimally compatible among other SQL tools, then don't use any Special or PL Commands. SqlTool has default behavior which I think is far superior to the other SQL tools, but you will have to disable these defaults in order to have optimally compatible behavior.

These switches provide compatibility at the cost of poor control and error detection.

- `--continueOnErr=true`

The output will still contain error messages about everything that SqlTool doesn't like (malformatted commands, SQL command failures, empty SQL commands), but SqlTool will continue to run. Errors will not cause rollbacks (but that won't matter because of the following setting).

- `--autoCommit`

You don't have to worry about accidental expansion of PL variables, since SqlTool will never expand PL variables if you don't set any variables on the command line, or give any `"* "` PL commands. (And you could not have `"* "` commands in a compatible SQL file).

Comments

Comments of the form `/* . . . */` or `--` behave as a SQL programmer would expect, in all contexts other than in interactive edit/history commands.

If a comment occurs outside of an SQL statement, SqlTool will not send the comment to the database (to improve performance). Raw mode can be used to send just comments to the database. In order to proactively catch accidents, SqlTool will complain if you attempt to send an empty SQL statement (i.e., just whitespace) to the database, even in raw mode.

Special Commands and Edit Buffer Commands in SQL Files

Don't use Edit Buffer / History Commands in your sql files, because they won't work. Edit Buffer / History Commands are for interactive use only. (But, see the [Raw Mode](#) section for an exception). You can, of course, use any SqlTool command at all interactively. I just wanted to group together the commands most useful to script-writers.

`\q` [abort message]

Be aware that the `\q` command will cause SqlTool to completely exit. If a script `x.sql` has a `\q` command in it, then it doesn't matter if the script is executed like

```
java -jar ../sqltool.jar urlid a.sql x.sql z.sql
```

or if you use `\i` to read it in interactively, or if another SQL file uses `\i` to nest it. If `\q` is encountered, SqlTool will quit. See the [SqlTool Procedural Language](#) section for commands to abort an SQL file (or even parts of an SQL file) without causing SqlTool to exit.

`\q` takes an optional argument, which is an abort message. If you give an abort message, the message is displayed to the user and SqlTool will exit with a failure status. If you give no abort message, then SqlTool will exit quietly with successful status. As a result,

```
\q
```

means to make an immediate but graceful exit, whereas

```
\q Message
```

	means to abort immediately.
<code>\p [text to print]</code>	Print the given string to stdout. Just give " <code>\p</code> " alone to print a blank line.
<code>\i /path/to/file.sql</code>	Include another SQL file at this location. You can use this to nest SQL files. For database installation scripts I often have a master SQL file which includes all of the other SQL files in the correct sequence. Be aware that the current continue-upon-error behavior will apply to included files until such point as the SQL file runs its own error handling commands.
<code>\o [file/path.txt]</code>	Tee output to the specified file (or stop doing so). See the Generating Text or HTML Reports section.
<code>\=</code>	A database-independent way to commit your SQL session. Useful for database which have no <code>COMMIT SQL</code> statement.
<code>\a [true false]</code>	This turns on and off SQL transaction autocommits. Auto-commit defaults to false, but you can change that behavior by using the <code>--autoCommit</code> command-line switch.
<code>\c [true false]</code>	A "true" setting tells SqlTool to Continue when errors are encountered. The current transaction will not be rolled back upon SQL errors, so if <code>\c</code> is true, then run the <code>ROLLBACK;</code> command yourself if that's what you want to happen. The default for interactive use is to continue upon error, but the default for non-interactive use is to abort upon error. You can override this behavior by using the <code>--continueOnError</code> command-line switch.

With database setup scripts, I usually find it convenient to set "true" before dropping tables (so that things will continue if the tables aren't there), then set it back to false so that real errors are caught. `DROP TABLE tablename IF EXISTS;` is a more elegant, but less portable, way to accomplish the same thing.



Tip

It depends on what you want your SQL files to do, of course, but I usually want my SQL files to abort when an error is encountered, without necessarily killing the SqlTool session. If this is the behavior that you want, then put an explicit `\c false` at the top of your SQL file and turn on continue-upon-error only for sections where you really want to permit errors, or where you are using PL commands to handle errors manually. This will give the desired behavior whether your script is called by somebody interactively, from the SqlTool command-line, or included in another SQL file (i.e. nested).



Important

The default settings are usually best for people who don't want to put in any explicit `\c` or error handling code at all. If you run SQL files from the SqlTool command line, then any errors will cause SqlTool to roll back and abort immediately. If you run SqlTool interactively and invoke SQL files with `\i` commands, the scripts will continue to run upon errors (and will not roll back). This behavior was chosen because there are lots of SQL files out there that produce errors which can be ignored; but we don't want to ignore errors that a user won't see. I reiterate that any

and all of this behavior can (and often should) be changed by Special Commands run in your interactive shell or in the SQL files. Only you know whether errors in your SQL files can safely be ignored.

Automation

SqlTool is ideal for mission-critical automation because, unlike other SQL tools, SqlTool returns a dependable exit status and gives you control over error handling and SQL transactions. Autocommit is off by default, so you can build a completely dependable solution by intelligently using `\c` commands (Continue upon Errors) and commit statements, and by verifying exit statuses.

Using the SqlTool Procedural Language, you have ultimate control over program flow, and you can use variables for database input and output as well as for many other purposes. See the [SqlTool Procedural Language](#) section.

Getting Interactive Functionality with SQL Files

Some script developers may run into cases where they want to run with sql files but they also want SqlTool's interactive behavior. For example, they may want to do command recall in the sql file, or they may want to log SqlTool's command-line prompts (which are not printed in non-interactive mode). In this case, do not give the sql file(s) as an argument to SqlTool, but pipe them in instead, like

```
java -jar $HSQLDB_HOME/lib/sqltool.jar urlid < filepath1.sql > /tmp/log.html 2
```

or

```
cat filepath1.sql... |  
java -jar $HSQLDB_HOME/lib/sqltool.jar urlid > /tmp/log.html 2>&1
```

For a shell not as graceful as the Bourne-compatible shells, you would need to type this all on the same line (or use a line-continuation trick).

Character Encoding

There are several levels of encoding settings. First there are your platform defaults. These can be changed, temporarily or permanently, with system settings or environmental variables. Java system properties may be used to change the encodings for the JVM run. Finally, can specify a different encoding in your RC file, as documented in the [RC File Authentication Setup](#) section, though these will not effect stdin or stdout (as explained there). Programmatic users of `SqlFile` have complete control over encoding by setting up `Readers` and `PrintWriters`, or by using constructors with an encoding parameter. Developers should understand that where a `SqlFile` constructor takes a `Reader` or a `PrintWriter` parameter, we will not apply encoding settings to them, leaving that up to you.

Generating Text or HTML Reports

This section is about making a file containing the output of database queries. You can generate reports by using operating system facilities such as redirection, tee, and cutting and pasting. But it is much easier to use the `\o` and `\H` special commands.



Note

I've neglected the `\H` feature, because I don't know of anybody using it. Be aware that it writes very old-fashioned HTML, with no attention to encoding, metadata, style sheets, etc.

Procedure 1.4. Writing query output to an external file

1. By default, everything will be done in plain text. If you want your report to be in HTML format, then give the special command `\H`. If you do so, you will probably want to use filenames with an suffix of ".html" or ".htm" instead of ".txt" in the next step.
2. Run the command `\o path/to/reportfile.txt`. From this point on, output from your queries will be appended to the specified file. (I.e. another *copy* of the output is generated.) This way you can continue to monitor or use output as usual as the report is generated.
3. When you want SqlTool to stop writing to the file, run `\o` (or just quit SqlTool if you have no other work to do).
4. If you turned on HTML mode with `\H` before, you can run `\H` again to turn it back off, if you wish.

It is not just the output of "SELECT" statements that will make it into the report file, but

Kinds of output that get teed to `\o` files

- Output of SELECT statements.
- Output of all "\d" Special Commands. (I.e., "\dt", "\dv", etc., and "\d OBJECTNAME").
- Output of "\p" Special Commands. You will want to use this to add titles, and perhaps spacing, for the output of individual queries.

Other output will go to your screen or stdout, but will not make it into the report file. Be aware that no error messages will go into the report file. If SqlTool is run non-interactively (including if you give any SQL file(s) on the command line), SqlTool will abort with an error status if errors are encountered. The right way to handle errors is to check the SqlTool exit status. (The described error-handling behavior can be modified with SqlTool command-line switches and Special Commands).



Warning

Remember that `\o` *appends* to the named file. If you want a new file, then use a new file name or remove the pre-existing target file ahead of time.



Tip

So that I don't end up with a bunch of junk in my report file, I usually leave `\o` off while I perfect my SQL. With `\o` off, I perfect the SQL query until it produces on my screen exactly what I want saved to file. At this point I turn on `\o` and run `"/select/;"` to repeat the last SQL command containing the given string ("select" in this example). If I have several complex queries to run, I turn `\o` off and repeat until I'm finished. (Every time you turn `\o` on, it will append to the file, just like we need).

Usually it doesn't come to mind that I need a wider screen until a query produces lines that are too long. In this case, stretch your window and repeat the last command with the `"/;"` Edit Buffer Command.

SqlTool Procedural Language

Aka PL

Most importantly, run `SqlTool` interactively and give the `"*?"` command to see what PL commands are available to you. I've tried to design the language features to be intuitive. Readers experience with significant shell scripting in any language can probably learn everything they need to know by looking at (and running!) the sample script `sample/pl.sql` in your HyperSQL distribution¹ and using the `*?` command from within an interactive SqlTool session as a reference. (By *significant* shell scripting, I mean to the extent of using variables, for loops, etc.).

PL variables will only be expanded after you run a PL command (or set variable(s) from the command-line). We only want to turn on variable expansion if the user wants variable expansion. People who don't use PL don't have to worry about strings getting accidentally expanded.

All other PL commands imply the "*" command, so you only need to use the "*" statement if your script uses PL variables and it is possible that no variables may be set before-hand (and no PL commands have been run previously). In this case, without "*", your script would silently use a literal value like "{x}" instead of trying to expand it. With a preceding "*" command, PL will notice that the variable x has not been set and will generate an error. (If x had been set here will be no issue because setting a variable automatically turns on PL variable expansion).

PL is also used to upload and download column values to/from local ASCII files, analogously to the special \b commands for binary files. This is explained above in the Interactive Essential PL Command section above.

Variables

- Use the * list command to list some or all variables; or * listvalues to also see the values.
- You can set variables using the * VARNAME = value command. This document explains elsewhere how you can set a values to the contents of files, and to the return value of SQL statements and fetches.
- You can also set variables using the --setvar command-line switch. I give a very brief but useful example of this below.
- SqlTool does not allow for setting system variables. As explained below, they are expanded in the same way as PL variables.
- Variables are always expanded in SQL, Special, and PL commands if they are written like * {VARNAME} (assuming that a PL command has been run previously). Your SQL scripts can give good feedback by echoing the value of variables with the "\p" special command. Use the construct * { :VARNAME } to expand the variable, but to expand to a zero-length string instead of fail if VARNAME is not set.
- Java system variables work the same exact way, except you use \$ instead of * to dereference, like so: \${user.name}.
- Variables are normally written like *VARNAME in logical expressions to prevent them from being evaluated too early. See below about logical expressions.
- You can't do math with expression variables, but you can get functionality like the traditional for (i = 0; i < x; i++) by appending to a variable and testing the string length, like

```
* while (*i < ${x})
  * i = *{i}.
```

i will be a growing line of dots.

- Variable names must not contain white space, or the characters "}" or "=".

Macros

Macros are just shortcut commands that you can run in place of the full commands which they stand for. Macros stand for SQL, Special or PL commands, whereas PL variables can only be used for elements within a command. It is very easy to define, list, and use macros. Run the command "/?" to see how. If you often run a particular query, then for the effort of about 5 extra keystrokes, you can define a macro for it so that you can enter just "/q;" to run it, whether the original query is 1 line or 40 lines. (You can use any name in place of "q", and the target command can be any kind of SQL, special, or PL command).

When you run/use a macro, you can append to the macro value. *appendage* in the "/?" listing shows where you can append additional text to the original command. So, if you define


```
sql> /= myworkers SELECT name FROM employees
```

, you could narrow the query variously during different macro invocations, like

```
sql> /myworkers WHERE dept = 20;
sql> /myworkers WHERE name like 'Karen%';
```

Just like when recalling a command from history, you use ";" to execute even Special and PL macro commands.

```
sql> /= notate Work completed by
sql> /notate Blaine;
```

If you don't type the ;, you will just recall the command to the buffer (from which you can execute or edit it, if you wish to).

To make a macro for a multi-line SQL statement, you use the "/= name :" construct. First, get the target command into the command buffer. If you have already run the command, then run ":h" to see the command number and load it to the buffer like ":13". If you haven't run the command yet, then just enter the command, but end it with a blank line (and no semi-colon). You can check the buffer with ":b" to make sure it is what you want. Then just run "/= name :" to define a macro with name "name".

PL Sample

Here is a short SQL file that gives the specified user write permissions on some application tables.

Example 1.5. Simple SQL file using PL

```
/*
 grantwrite.sql

 Run SqlTool like this:
   java -jar path/to/sqltool.jar -setvar=USER=debbie grantwrite.sql
*/

/* Explicitly turn on PL variable expansion, in case no variables have
   been set yet. (Only the case if user did not set USER).
*/
*

GRANT all ON book TO *{USER};
GRANT all ON category TO *{USER};
```

Note that this script will work for any (existing) user just by supplying a different user name on the command-line. I.e., no need to modify the tested and proven script. There is no need for a `commit` statement in this SQL file since no DML is done. If the script is accidentally run without setting the `USER` variable, SqlTool will give a very clear notification of that.

The purpose of the plain "*" command is just so that the `*{USER}` variables will be expanded. (This would not be necessary if the `USER` variable, or any other variable, were set, but we don't want to depend upon that).

Logical Expressions

Logical expressions occur only inside of logical expression parentheses in PL statements. For example, `if (*var1 > astring)` and `while (*checkvar)`. (The parentheses after "foreach" do not enclose a logical expression, they just enclose a list).

There is a critical difference between `*{VARNAME}` and `*VARNAME` inside logical expressions. `*{VARNAME}` is expanded one time when the parser first encounters the logical expression. `*VARNAME` is re-expanded every time that

the expression is evaluated. So, you would never want to code `* while (*{X} < 5)` because the statement will always be true or always be false. (I.e. the following block will loop infinitely or will never run).

Don't use quotes or whitespace of any kind in `*{VARNAME}` variables in expressions. (They would expand and then the expression would most likely no longer be a valid expression as listed in the table below). Quotes and whitespace are fine in `*VARNAME` variables, but it is the entire value that will be used in evaluations, regardless of whether quotes match up, etc. I.e. quotes and whitespace are not *special* to the token evaluator.

Logical Operators

TOKEN	The token may be a literal, a <code>*{VARNAME}</code> which is expanded early, or a <code>*VARNAME</code> which is expanded late. (You usually do not want to use <code>*{VARNAME}</code> in logical expressions). False if the token is not set, empty, or "0". True otherwise.
TOKEN1 == TOKEN2	True if the two tokens are equivalent "strings".
TOKEN1 <> TOKEN2	Ditto.
TOKEN1 >< TOKEN2	Ditto.
TOKEN1 > TOKEN2	True if the TOKEN1 string is longer than TOKEN2 or is the same length but is greater according to a string sort.
TOKEN1 < TOKEN2	Similarly to TOKEN1 > TOKEN2.
! LOGICAL_EXPRESSION	Logical negation of any of the expressions listed above.

`*VARNAME`s in logical expressions, where the `VARNAME` variable is not set, evaluate to an empty string. Therefore `(*UNSETVAR = 0)` would be false, even though `(*UNSETVAR)` by itself is false and `(0)` by itself is false. Another way of saying this is that `*VARNAME` in a logical expression is equivalent to `*{:VARNAME}` out of a logical expression.

When developing scripts, you definitely use SqlTool interactively to verify that SqlTool evaluates logical expressions as you expect. Just run `* if` commands that print something (i.e. \p) if the test expression is true.

Flow Control

Flow control works by conditionally executing blocks of Commands according to conditions specified by logical expressions.

The conditionally executed blocks are called *PL Blocks*. These PL Blocks always occur between a PL flow control statement (like `* foreach`, `*while`, `* if`) and a corresponding `* end` PL Command (like `* end foreach`).

The values of control variables for `foreach` and `while` PL blocks will change as expected.

There are `* break` and `* continue`, which work as any shell scripter would expect them to. The `* break` command can also be used to quit the current SQL file without triggering any error processing. (I.e. processing will continue with the next line in the *including* SQL file or interactive session, or with the next SQL file if you supplied multiple on the command-line).

Example

Below is the example SQL file `sample/pl.sql`, which shows how to use most PL features ¹. If you have a question about how to use a particular PL feature, check this file in your distribution before asking for help. Definitely give it a run, like

```
java -jar $HSQLDB_HOME/lib/sqltool.jar mem $HSQLDB_HOME/pl.jar
```

Example 1.6. SQL File showing use of most PL features

```

/*
 $Id: pl.sql 3353 2009-12-15 19:52:13Z unsaved $
 SQL File to illustrate the use of SqlTool PL features.
 Invoke like
     java -jar ../hsqldb.jar ../pl.sql mem
                                     -- blaine
*/

* if (! *MYTABLE)
  \p MYTABLE variable not set!
  /* You could use \q to Quit SqlTool, but it's often better to just
     break out of the current SQL file.
     If people invoke your script from SqlTool interactively (with
     \i yourscripname.sql) any \q will kill their SqlTool session. */
  \p Use arguments "--setvar=MYTABLE=mytablename" for SqlTool
  * break
* end if

/* Turning on Continue-upon-errors so that we can check for errors ourselves.*/
\c true

\p
\p Loading up a table named '*{MYTABLE}'...

/* This sets the PL variable 'retval' to the return status of the following
   SQL command */
* retval ~
CREATE TABLE *{MYTABLE} (
  i int,
  s varchar(20);
);
\p CREATE status is *{retval}
\p

/* Validate our return status. In logical expressions, unset variables like
 *unsetvar are equivalent to empty string, which is not equal to 0
 (though both do evaluate to false on their own, i.e. (*retval) is false
 and (0) is false */
* if (*retval != 0)
  \p Our CREATE TABLE command failed.
  * break
* end if

/* Default Continue-on-error behavior is what you usually want */
\c false
\p

/* Insert data with a foreach loop.
   These values could be from a read of another table or from variables
   set on the command line like
*/
\p Inserting some data into our new table (you should see 3 row update messages)
* foreach VALUE (12 22 24 15)
  * if (*VALUE > 23)
    \p Skipping *{VALUE} because it is greater than 23
    * continue
    \p YOU WILL NEVER SEE THIS LINE, because we just 'continued'.
  * end if
  INSERT INTO *{MYTABLE} VALUES (*{VALUE}, 'String of *{VALUE}');
* end foreach
\p

* themax ~
/* Can put Special Commands and comments between "* VARNAME ~" and the target
   SQL statement. */

```

```

\p We're saving the max value for later.  You'll still see query output here:
SELECT MAX(i) FROM *{MYTABLE};

/* This is usually unnecessary because if the SELECT failed, retval would
   be undefined and the following print statement would make SqlTool exit with
   a failure status */
* if (! *themax)
  \p Failed to get the max value.
  /* It's possible that the query succeeded but themax is "0".
     You can check for that if you need to. */
  * break
  \p YOU WILL NEVER SEE THIS LINE, because we just 'broke'.
* end if

\p
\p #####
\p The results of our work:
SELECT * FROM *{MYTABLE};
\p MAX value is *{themax}

\p
\p Everything worked.

```

Chunking

We hereby call the ability to transmit multiple SQL commands to the database in one transmission *chunking*. Normally it's best to send SQL statements to the database one-at-a-time. That way, the database can give you or your program feedback about each statement. But there are situations where it is more important to transmit multiple-statements-at-a-time than to get feedback for each statement individually.

Why?

The first general reason to chunk SQL commands is performance. For standalone databases, the most common performance bottleneck is network latency. Chunking SQL commands can dramatically reduce network traffic.

The second reason is that there are a couple SQL commands which require the terminating ";" to be sent to the database engine. For simplicity and efficiency, it's usually better for general JDBC clients like SqlTool to strip off the final delimiter. Raw commands retains everything that the user types.

The third general reason to chunk SQL commands is if your database requires you to send multiple commands in one transmission. This is usually the case with the following types of commands:

- Nested SQL commands, like the nested CREATE SCHEMA variant, and most stored procedure, function, and trigger definitions.
- Commands containing non-quoted programming language to be interpreted by the database engine. Definitions of stored procedures, function, and triggers often contain code like this.

How?

Use raw mode. Go to the [Raw Mode](#) section to see how. You can enter any text at all, exactly how you want it to be sent to the database engine. Therefore, in addition to chunking SQL commands, you can give commands for non-SQL extensions to the database. For example, you could enter JavaScript code to be used in a stored procedure.

Raw Mode

You begin raw mode by issuing the Special Command "\.". You can then enter as much text in any format you want. When you are finished, enter a line consisting of only ";\." to store the input to the edit buffer and send it to the database server for execution.

This paragraph applies only to interactive usage. Interactive users may end the raw input with ":" instead of ";". This will just save the input to the edit buffer so that you can edit it and send it to the database manually. You can look at the edit buffer with the ":b" Buffer Command. You would normally use the command ";" to send the buffer to the database after you are satisfied with it. You'll notice that your prompt will be the "raw" prompt between entering "\." and terminating the raw input with ";" or ":".

Just by running commands beginning with `BEGIN`, `DECLARE`, `CREATE function`, or `CREATE procedure`, your SqlTool session will automatically be changed to Raw mode, exactly as if you had entered "\.". That's because these commands are universally used to define stored procedures or functions, and these commands require raw mode (as explained in the previous section).

Example 1.7. Interactive Raw Mode example

```
sql> \.
Enter RAW SQL. No \, :, * commands.
End with a line containing only ";" to send to database,
or ":" to store to edit buffer for editing or saving.
-----
raw> line one;
raw> line two;
raw> line three;
raw> :.
Raw SQL chunk moved into buffer. Run ";" to execute the chunk.
sql> ;;
Executing command from buffer:
line one;
line two;
line three;

SQL Error at 'stdin' line 13:
"line one;
line two;
line three;"
Unexpected token: LINE in statement [line]
sql>
```

The error message "Unexpected token: LINE in statement [line]" comes from the database engine, not SqlTool. All three lines were transmitted to the database engine.

Edit Buffer Commands are not available when running SqlTool non-interactively.

SQL/PSM, SQL/JRT, and PL/SQL

This section covers database-engine-embedded languages, which are often used in the definition of stored procedures, stored functions, and triggers. `SQL/PSM`, `SQL/JRT`, and `PL/SQL` are well known examples. We prefer `SQL/PSM` and `SQL/JRT` because unlike the alternatives, they are based on open SQL specifications.



Note

`PL/SQL` is **not** the same as `PL`. `PL` is the procedural language of `SqlFile` and is independent of your back-end database. `PL` commands always begin with `*`. `PL/SQL` is an Oracle-specific extension processed on the server side. You can not intermix `PL` and any server-embedded language (except for setting a `PL` variable to the output of execution), because when you enter server language to `SqlTool`, that input is not processed by `SqlFile`.

Use Raw Mode to send server-language code blocks to the database engine. You do not need to enter the "\." command to enter raw mode. Just begin a new `SqlTool` command line with `"DECLARE"`, `"BEGIN"`, `"CREATE FUNCTION"`,

or "CREATE PROCEDURE", and SqlTool will automatically put you into raw mode. See the [Raw Mode](#) section for details.

The following sample SQL file resides at `sample/plsql.sql` in your HyperSQL distribution ¹. This script will only work with Oracle, only if you have permission to create the table "T1" in the default schema, and if that object does not already exist.

Example 1.8. PL/SQL Example

```

/*
 * $Id: plsql.sql 826 2009-01-17 05:04:52Z unsaved $
 *
 * This example is copied from the "Simple Programs in PL/SQL"
 * example by Yu-May Chang, Jeff Ullman, Prof. Jennifer Widom at
 * the Standord University Database Group's page
 * http://www-db.stanford.edu/~ullman/fcdb/oracle/or-plsql.html .
 * I have only removed some blank lines (in case somebody wants to
 * copy this code interactively-- because you can't use blank
 * lines inside of SQL commands in non-raw mode SqlTool when running
 * it interactively); and, at the bottom I have replaced the
 * client-specific, non-standard command "run;" with SqlTool's
 * corresponding command "." and added a plain SQL SELECT command
 * to show whether the PL/SQL code worked. - Blaine
 */

CREATE TABLE T1(
    e INTEGER,
    f INTEGER
);

DELETE FROM T1;

INSERT INTO T1 VALUES(1, 3);

INSERT INTO T1 VALUES(2, 4);

/* Above is plain SQL; below is the PL/SQL program. */
DECLARE

    a NUMBER;

    b NUMBER;

BEGIN

    SELECT e,f INTO a,b FROM T1 WHERE e>1;

    INSERT INTO T1 VALUES(b,a);

END;

.;
/** The statement on the previous line, "." is SqlTool specific.
 * This command says to save the input up to this point to the
 * edit buffer and send it to the database server for execution.
 * I added the SELECT statement below to give imm
 */

/* This should show 3 rows, one containing values 4 and 2 (in this order)...*/
SELECT * FROM t1;

```

Note that, inside of raw mode, you can use any kind of formatting that your database engine needs or permits: Whatever you enter-- blank lines, comments, everything-- will be transmitted to the database engine.

This file resides at `testrun/sqltool/sqljrt.sql`

Example 1.9. SQL/JRT Example

```

/*
 * $Id: sqljrt.sql 3353 2009-12-15 19:52:13Z unsaved $
 *
 * Tests SQL/JRT
 */

create function dehexas(VARCHAR(80), INTEGER)
  returns INTEGER
  no sql
  language java
  external name 'CLASSPATH:java.lang.Integer.valueOf'
.;

CALL dehexas('12', 16);
*if (*? != 18)
  \q SQL/JRT function failed
*end if

```

This file resides at `testrun/sqltool/sqlpsm.sql`

Example 1.10. SQL/PSM Example

```

/*
 * $Id: sqlpsm.sql 826 2009-01-17 05:04:52Z unsaved $
 *
 * Tests SQL/JRT
 */

create table customers(
  id INTEGER default 0, firstname VARCHAR(50), lastname VARCHAR(50),
  entrytime TIMESTAMP);

create procedure new_customer(firstname varchar(50), lastname varchar(50))
  modifies sql data
  insert into customers values (
    default, firstname, lastname, current_timestamp)
.;

SELECT count(*) FROM customers;
*if (*? != 0)
  \q SQL/PSM preparation failed
*end if

CALL new_customer('blaine', 'simpson');

SELECT count(*) FROM customers;
*if (*? != 1)
  \q SQL/PSM procedure failed
*end if

```

Delimiter-Separated-Value Imports and Exports

SqlTool's DSV functionality encompasses what many users will recognize as CSV export, as well as portable backup or transfer of data. Those familiar with Oracle's SQL*Loader will recognize the extreme usefulness of the feature set. Besides database- and platform-independent data backups, exports can be used to deploy data sets with applications, to transfer data among multiple database instances (even drastically different database instances such as SQL Server and HyperSQL), and to properly change control data sets with a content management system such as a collaboration server or Subversion. To jump way ahead for a moment to whet your appetite, here is a sample *import reject report* which will can be generated automatically for you upon import just by setting the PL variable `*DSV_REJECT_REPORT` (to the desired destination HTML file name).

SqlTool DSV Reject Report - Mozilla Firefox

file:///tmp/sample-reject.html

Import performed **Mon Dec 14 20:40:40 EST 2009** with SqlTool.

Input DSV file: **sample.dsv**

Reject DSV file: **/tmp/sample-reject.dsv**

The corresponding records in '/tmp/sample-reject.dsv' are at line numbers of (reject # + 1), since the header record occupies the first line.

rej. #	input line #	bad column (if known)	<u>reason</u>
1	11	d	Bad date/time value 'not a date 0:00:00'. java.lang.IllegalArgumentException: Timestamp format must be yyyy-mm-dd hh:mm:ss[.fffffffff]
2	14		integrity constraint violation: NOT NULL check constraint: SYS_CT_10028 table: SAMPLETABLE

Import summary ('#'-skips / rejects / inserts): 3 / 2 / 3.

SqlFile revision 3347.

Done

A DSV Import reject report

If you wish to, you can review the reject report before deciding whether to commit or roll back the inserts.



Note

This feature is independent of HyperSQL Text Tables. (See the Text Tables chapter of the HyperSQL User Guide [http://hsqldb.org/doc/2.0/guide/index.html] for details about them). a server-side feature of HyperSQL. It makes no difference to SqlTool whether the source or target table of your export/import is a memory, cache, or text table. Indeed, like all features of SqlTool, it works fine with other JDBC databases. It works great, for example to migrate data from a table of one type to a table of another type, or to another schema, or to another database instance, or to another database system.

This feature is what most business people call "CSV", but these files are more accurately called *Delimiter Separated Value files* because the delimiter is usually not a comma, and, more importantly, we purposefully choose an effective delimiter instead of the CSV method of using a delimiter which works in some cases and then use quotes and backslashes to escape occurrence of the delimiter in the actual data. Just by choosing a delimiter which never needs escaping, we eliminate the whole mess, and the data in our files always looks just like the corresponding data in the database. To make this CSV / Delimiter-separated-value distinction clear, I use the suffix ".dsv" for my data files. This leads me to stipulate the abbreviation DSV for the *Delimiter Separated Value* feature of HyperSQL.

Use the `\x` command to eXport a table to a DSV file, and the `\m` command to iMport a DSV file into a pre-existing table.

The row and column delimiters may be any String (or even a regular expression for import), not just a single character. And just as the delimiter capability is more general than traditional CSV delimiters, the export function is also more general than just a table data exporter. Besides the trivial generalization that you may specify a view or other virtual table name in place of a table name, you can alternatively export the output of any query which produces normal text

output. A benefit to this approach is that it allows you to export only some columns of a table, and to specify a WHERE clause to narrow down the rows to be exported (or perform any other SQL transformation, mapping, join, etc.). One specific use for this would be to exclude columns of binary data (which can be exported by other means, such as a PL loop to store binary values to files with the \bd command).

Note that the import command will not create a new table. This is because of the impossibility of guessing appropriate types and constraints based only on column names and a data sampling (which is all that a DSV-importer has access to). Therefore, if you wish to populate a new table, create the table before running the import. The import file does not need to have data for all columns of a table. The only required columns are those required by database constraints (non-null, indexes, keys, etc.) One specific reason to omit columns is if you want values of some columns to be created automatically by column DEFAULT settings, triggers, HyperSQL identity sequences, etc. Another reason would be to skip binary columns.

Simple DSV exports and imports using default settings

Even if you need to change delimiters, table names, or file names from the defaults, I suggest that you run one export and import with default settings as a practice run. A memory-only HyperSQL instance is ideal for test runs like this.

This command exports the table `icf.projects` to the file `projects.dsv` in the current directory (where you invoked SqlTool from). By default, the output file name will be the specified source table name plus the extension `.dsv`.

Example 1.11. DSV Export Example

```
SET SCHEMA icf;  
\x projects
```

We could also have run `\x icf.projects` (which would have created a file named `icf.projects.dsv`) instead of changing the session schema. In this example we have chosen to make the export file name independent of the schema to facilitate importing it into a different schema.

Take a look at the output file. Notice that the first line consists of column names, not data. This line is present because it will be needed if the file is to be used for a DSV import. Notice the following characteristics about the export data. The column delimiter is the pipe character "|". The record delimiter is the default line delimiter character(s) for your operating system. The string used to represent database NULLs is `[null]`. See the next section for how to change these from their default values.



Warning

You can not DSV import Array values where any Array elements contain commas, for example an Array of VARCHARs which contain one or more commas. There is no such limitation on DSV exports, which you can use for purposes other than SqlTool importing, or you could use a script to change the commas to some other character.

This command imports the data from the file `projects.dsv` in the current directory (where you invoked SqlTool from) into the table `newschema.projects`. By default, the output table name will be the input filename after removing optional leading directory and trailing final extension.

Example 1.12. DSV Import Example

```
SET SCHEMA newschema;  
\m projects.dsv
```

If the DSV file was named with the target schema, you would have skipped the `SET SCHEMA` command, like `\m newschema.projects.dsv`. In order to allow for more flexibility, the default input delimiters are not ex-

actly the same as the output delimiters. The input delimiters are regular expressions. The input column delimiter happens to be the regular expression corresponding exactly to "|"; but the input record delimiter matches UNIX, Windows, Mac, and HTTP line breaks.

Specifying queries and options

For a hands on example of a DSM import which generates an import report and uses some other options, change to directory HSQLDB/sample and play with the working script `dsv-sample.sql`¹. You can execute it like

```
java -jar ../lib/sqltool.jar mem dsv-sample.sql
```

(assuming that you are using the supplied `sqltool.rc` file or have `urlid mem` set up).

The header line in the DSV file is required at this time. (If there is user demand, it can be made optional for exporting, but it will remain required for importing).

Your export will fail if the output column or record delimiter, or the null representation value occurs in the data being exported. You change these values by setting the PL variables `*DSV_COL_DELIM`, `*DSV_ROW_DELIM`, `*DSV_NULL_REP`. Notice that the asterisk is part of the variable names, to indicate that these variables are used by SqlTool internally. When specifying delimiters, you can use the escape sequences `\n`, `\r`, `\f`, `\t`, `\\`, and decimal, octal or hex specifications like `\20`, `\020`, `\0x20`. For example, to change the column delimiter to the tab character, you would give the command

```
* *DSV_COL_DELIM = \t
```

The input (`\m`) delimiter values, `*DSV_COL_SPLITTER` and `*DSV_ROW_SPLITTER`, are set using normal Perl/Java regexp syntax. There are escapes for specifying special characters, and anything else you would need. Input vs. output row and column delimiters are easily distinguished by containing "SPLITTER" for splitting input (`\m`) files; or "DELIM" for the delimiters that we will write (`\x`) among the data.

For imports, you must always specify the source DSV file path. If you want to *export* to a different file than one in the current directory named according to the source table, set the PL variable `*DSV_TARGET_FILE`, like

```
* *DSV_TARGET_FILE = /tmp/dtbl.dsv
```

For exports, you must always specify the source table name or query. If you want to *import* to a table other than that derived from the input DSV file name, set the PL variable `*DSV_TARGET_TABLE`. The table name may contain a schema name prefix.

You don't need to import all of the columns in a data file. To designate the fields to be skipped, iether set the PL variable `*DSV_SKIP_COLUMNS`, or replace the column names in the header line to "-" (hyphen). The value of `*DSV_SKIP_COLUMNS` is case-insensitive, and multiple column names are separated with white space and/or commas.

You can specify a query instead of a tablename with the `\x` command in order to filter or transform data from a table or view, or to export the output of a join, etc. You must set the PL variable `*DSV_TARGET_FILE`, as explained above (since there is no table name from which to automatically map a file name).

Example 1.13. DSV Export of an Arbitrary SELECT Statement

```
* *DSV_TARGET_FILE = outfile.txt
\x SELECT entrydate, 2 * aval "Double aval", modtime FROM bs.dtbl
```

Note that I specified the column label alias "Double aval" so that the label for that column in the DSV file header will not be blank.

By default, imports will abort as soon as a error is encountered during parsing the file or inserting data. If you invoke SqlTool with a SQL script on the command line, the failure will cause SqlTool to roll back and exit. If run interactively, you can decide whether to commit or roll back the rows that inserted before the failure. You can modify this behavior with the `\a` and `\c` settings.

If you set either a reject dsv file or a reject report file, then failures during imports will be reported but will not cause the import to abort. When run in this way, SqlTool will give you a report at the end about how many records were skipped, rejected, and successfully inserted. The reject dsv file is just a dsv file with exact copies of the dsv records that failed to insert. The reject report file is a HTML report which lists, for every rejected record, why that record was rejected. `\m?` will show you that the required PL variables for this functionality are `*DSV_REJECT_FILE` and `*DSV_REJECT_REPORT`. In both cases, you set the variable value to the path of the file which SqlTool will create.

To allow for user-friendly entry of headers, we require that tables for DSV import/exports use standard column names. I.e., no column names that require quoting. The DSV import and export parsers are very smart and user-friendly. The data types of columns are checked so that the parser can make safe assumptions about white space and blank entries in the data. If a column is a JDBC Boolean type, for example, then we know that a field value of " True " obviously means "True", and that a field value of "" obviously means null. Since we require vanilla style column names, we allow white space anywhere in the header column. We allow blank lines anywhere (where "lines" are delimited by `*DSV_ROW_DELIM`). By default, commented lines are ignored, and the comment character can be changed from its default value.

Run the command `"\x?"` or `"\m?"` to see the several system PL variables which you can set to adjust reject file behavior, commenting behavior, and other DSV features.

You can also define some settings right in the DSV file, and you can even specify multiple header lines in a single DSV file. I use this last feature to import data from one data set into multiple tables that are joined. Since I don't have any more time to dedicate to explaining all of these features, I'll give you some examples from working DSV files and let you take it from there.

Example 1.14. Sample DSV headerswitch settings

```
# RCS keyword was here.

headerswitch{
  itemdef:name|-|hardness|breakdc|-
  simpleitemdef:itemdef_name|maxvalue|weight|-|maxhp
}
```

I'll just note that the prefixes for the header rows must be of format `target-table-name + :`. You can use `*` for `target-table-name` here, for the obvious purpose.

Example 1.15. DSV targettable setting

```
targettable=t
```

This last example is from the SqlTool unit test file `dsv-trimming.dsv`. These special commands must be at the top of the file (before any normal data or header lines).

There is also the `*DSV_CONST_COLS` setting, which you can use to automatically write static, constant values to the specified columns of all inserted rows.

Unit Testing SqlTool

The SqlTool unit tests reside at `testrun/sqltool` in the HyperSQL source code repository. Just run the `runtests.bash` script from that directory to execute all of the tests. As you can see, the test runner, unfortunately, requires a Bash shell

at this time. Read the file `README.txt` to find out all about it, including everything you'd need to know to test your own scripts or to add more unit test scripts for SqlTool.

Chapter 2. Hsqldb Test Utility

The `org.hsqldb.test` package contains a number of tests for various functions of the database engine. Among these, the `TestUtil` class performs the tests that are based on scripts. To run the tests, you should compile the `hsqldbtest.jar` target with Ant and JUnit.

The `TestUtil` class should be run in the `/testrun/hsqldb` directory of the distributed files. It then runs the set of `TestSelf*.txt` files in the directory. To start the application in Windows, change to the directory and type:

```
java org.hsqldb.test.TestUtil
```

All files in the working directory with names matching `TestSelf*.txt` are processed in alphabetical order.

You can add your own scripts to test different series of SQL queries. The format of the `TestSelf*.txt` file is simple text, with some indentation and prefixes in the form of Java-style comments. The prefixes indicate what the expected result should be.

The class `org.hsqldb.test.TestScriptRunner` is a more general program which you can use to test any script files which you specify (with scripts of the same exact format as described below). For example,

```
java org.hsqldb.test.TestScriptRunner --urlid=mem script1.tsq1 script2.sql
```

You must have the HSQLDB classes, including the util and test classes, in your `CLASSPATH`. The `urlid` must be set up in an RC file as explained in the [RC File Authentication Setup](#) section. Use the `rcfile=` argument to specify an RC file other than the default of `testscriptrunner.rc` in the current directory. To see all invocation possibilities, just run `TestScriptRunner` with no arguments at all. `TestScriptRunner` can run tests sequentially (the default) or in simultaneous asynchronous threads.

- Comment lines must start with `--` and are ignored
- Lines starting with spaces are the continuation of the previous line (for long SQL statements)
- SQL statements with no prefix are simply executed.
- *The remaining items in this list exemplify use of the available command line-prefixes.*
- The `/*s*/` option stands for silent. It is used for executing queries regardless of results. Used for preparation of tests, not for actual tests.

```
/*s*/ Any SQL statement - errors are ignored
```

- The `/*c<rows>*/` option is for `SELECT` queries and asserts the number of rows in the result matches the given count.

```
/*c<rows>*/ SQL statement returning count of <rows>
```

- The `/*u*/` option is for queries that return an update count, such as `DELETE` and `UPDATE`. It asserts the update count matches.

```
/*u<count>*/ SQL statement returning an update count equal to <count>
```

- The `/*e*/` option asserts that the given query results in an error. It is mainly used for testing the error detection capabilities of the engine. The SQL State of the expected error can be defined, for example `/*e42578*/`, to verify the returned error. This option can be used with syntactically valid queries to assert a certain state in the database. For example a `CREATE TABLE` can be used to assert the table of the same name already exists.

```
/*e*/ SQL statement that should produce an error when executing
```

- The `/*r...*/` option asserts the `SELECT` query returns a single row containing the given set of field values.

```
/*r<string1>,<string2>*/ SQL statement returning a single row ResultSet equal to the specified value
```

- The extended `/*r...*/` option asserts the SELECT query returns the given rows containing the given set of field values.

```
/*r
  <string1>,<string2>
  <string1>,<string2>
  <string1>,<string2>
*/ SQL statement returning a multiple row ResultSet equal to the specified values
```

(note that the result set lines are indented).

- The `/*d*/` directive just displays the supplied text.

```
/*d*/ Some message
```

- The `/*w MILLIS*/` directive causes the test to Wait for a specified number of milliseconds.

```
/*w 1000*/ Optional message
```

- The `/*w ENFORCE_SEQUENCE WAITER_NAME*/` directive causes the test to Wait for the specified *Waiter*. A waiter is just name that is used to associate a `/*w*/` directive to its corresponding `/*p*/` directive. The ENFORCE_SEQUENCE argument must be set to `true` or `false` to specify whether to fail unless the `/*p*/` command runs after the `/*w*/` command is waiting.

```
/*w true script4*/ Optional message
```

- The `/*p ENFORCE_SEQUENC WAITER_NAME*/` directive is the peer directive to `/*w*/`, which causes a waiting thread to Proceed.

```
/*p true script4*/ Optional message
```

- All the options are lowercase letters. During development, an uppercase can be used for a given test to exclude a test from the test run. The utility will just report the test blocks that have been excluded without running them. Once the code has been developed, the option can be turned into lowercase to perform the actual test.

See the TestSelf*.txt files in the `/testrun/hsqlDb/` directory for actual examples.

The String `${timestamp}` may be used in script messages (like in `/*d*/`, `/*w*/`, `/*p*/`). It expands to the current time, down to the second. For example,

```
/*d*/ It is now ${timestamp}
```

Chapter 3. Database Manager

Fred Toussi, The HSQL Development Group
Blaine Simpson, The HSQL Development Group

Brief Introduction

The Database Manager tool is a simple GUI database query tool with a tree display of the tables. Both AWT and SWING versions of the tool are available and work almost identically. The AWT version class name is `org.hsqldb.util.DatabaseManager`; the SWING version, `org.hsqldb.util.DatabaseManagerSwing`. The SWING version has more refinements than the AWT version.

The AWT version of the database manager can be deployed as an applet in a browser. A demo HTML file with an embedded Database Manager is included in the `/demo` directory.

When the Database Manager is started, a dialogue allows you to enter the JDBC driver, URL, user and password for the new connection. A drop-down box, Type, offers preset values for JDBC driver and URL for most popular database engines, including HSQLDB. Once you have selected an item from this drop-down box, you should edit the URL to specify the details of the database or any additional properties to pass. You should also enter the username and password before clicking on the OK button.

The connection dialogue allows you to save the settings for the connection you are about to make. You can then access the connection in future sessions. To save a connection setting, enter a name in the Setting Name box before clicking on the OK button. Next time the connection dialogue is displayed, the drop-down box labeled Recent will include the name for all the saved connection settings. When you select a name, the individual settings are displayed in the appropriate boxes.

The small Clr button next to the drop-down box allows you to clear all the saved settings. If you want to modify an existing setting, first select it from the drop-down box then modify any of the text boxes before making the connection. The modified values will be saved.

Most SWING menu items have context-sensitive tool tip help text which will appear if you hold the mouse cursor still over the desired menu item. (Assuming that you don't turn Tooltips off under the Help menu.)

The database object tree in the SWING version allows you to right click on the name of a table or column and choose from common SQL statements for the object, for example `SELECT * FROM thistable ...`. If you click on one of the given choices, the sample statement is copied to the command window, where you can modify and complete it.

The DatabaseManagers do work with HSQLDB servers serving TLS-encrypted JDBC data. See the TLS section of the Listeners chapter of the [HyperSQL User Guide](http://distro_baseurl_DEFAULTVAL/guide/index.html) [distro_baseurl_DEFAULTVAL/guide/index.html]



Tip

If you are using `DatabaseManagerSwing` with Oracle, you will want to make sure that `Show row counts` and `Show row counts` are both off *before connecting to the database*. You may also want to turn off `Auto tree-update`, as described in the next section.

Auto tree-update

By default, the object tree in the left panel is refreshed when you execute DDL which may update those objects. If you are on a slow network or performance-challenged PC, use the `view / Auto-refresh tree` menu item to turn it off. You will then need to use the `viewRefresh tree` menu item every time that you want to refresh the tree.



Note

Auto-refresh tree does not automatically show all updates to database objects, it only refreshes when you submit DDL which may update database objects. (This behavior is a compromise between utility and performance).

Automatic Connection

You can use command-line switches to supply connection information. If you use these switch(es), then the connection dialog window will be skipped and a JDBC connection will be established immediately. Assuming that the `hsqldb.jar` (or an alternative jar) are in your `CLASSPATH`, this command will list the available command-line options.

```
java org.hsqldb.util.DatabaseManagerSwing --help
```

It's convenient to skip the connection dialog window if you always work with the same database account.



Warning

Use of the `--password` switch is not secure. Everything typed on command-lines is generally available to other users on the computer. The problem is compounded if you use a network connection to obtain your command line. The RC File section explains how you can set up automatic connections without supplying a password on the command line.

RC File

You can skip the connection dialog window securely by putting the connection information into an RC file and then using the `--urlid` switch to `DatabaseManager` or `DatabaseManagerSwing`. This strategy is great for adding launch menu items and/or launch icons to your desktop. You can set up one icon for each of the database accounts which you regularly use.

The default location for the RC file is `dbmanager.rc` in your home directory. The RC File Authentication Setup section explains how to put the connection information into this text file. If you also run `SqlTool`, then you can share the RC file with `SqlTool` by using a sym-link (if your operating system supports sym links), or by using the `--rcfile` switch for either `SqlTool` or `DatabaseManagerSwing`.



Warning

Use your operating system facilities to prevent others from reading your RC file, since it contains passwords.

To set up launch items/icons, first experiment on your command line to find exactly what command works. For example,

```
java -cp /path/to/hsqldb.jar org.hsqldb.util.DatabaseManagerSwing --urlid mem
```

Then, use your window manager to add an item that runs this command.

Using the current DatabaseManagers with an older HSQLDB distribution.

This procedure will allow users of a legacy version of HSQLDB to use all of the new features of the DatabaseManagers. You will also get the new version of the `SqlTool`! This procedure works for distros going back to 1.7.3.3 at least, probably much farther.

These instructions assume that you are capable of running an Ant build. See the Building Appendix of the HyperSQL User Guide [distro_baseurl_DEFAULTVAL/guide/index.html].

1. Download and extract a current HSQLDB distribution. If you don't want to use the source code, documentation, etc., you can use a temporary directory and remove it afterwards.
2. Cd to the build directory under the root directory where you extracted the distribution to.
3. Run `ant hsqldbutil`.
4. If you're going to wipe out the build directory, copy `hsqldbutil.jar` to a safe location first.
5. For now on, whenever you are going to run DatabaseManager*, make sure that you have this `hsqldbutil.jar` as the first item in your CLASSPATH.

Here's a UNIX example where somebody wants to use the new DatabaseManagerSwing with their older HSQLDB database, as well as with Postgresql and a local application.

```
CLASSPATH=/path/to/hsqldbutil.jar:/home/bob/myapp/classes:/usr/local/lib/pg.jdbc3.jar
export CLASSPATH
java org.hsqldb.util.DatabaseManagerSwing --urlid urlid
```

DatabaseManagerSwing as an Applet

DatabaseManagerSwing is also an applet. You can use it in HTML, JSPs, etc. Be aware that in Applet mode, actions to load or save local files will be disabled, and attempts to access any server other than the HTML-serving-host will fail.

Since the Applet can not store or load locally saved preferences, the only way to have persistent preference settings is by using Applet parameters.

DatabaseManagerSwing Applet Parameters

<code>jdbcUrl</code>	URL of a data source to auto-connect to. String value.
<code>jdbcDriver</code>	URL of a data source to auto-connect to. String value. Defaults to <code>org.hsqldb.driver.JDBCdriver</code> .
<code>jdbcUser</code>	User name for data source to auto-connect to. String value.
<code>jdbcPassword</code>	Password for data source to auto-connect to. String value. Defaults to zero-length string.
<code>schemaFilter</code>	Display only object from this schema in the object navigator. String value.
<code>laf</code>	Look-and-feel. String value.
<code>loadSampleData</code>	Auto-load sample data. Boolean value. Defaults to false.
<code>autoRefresh</code>	Auto-refresh the object navigator when DDL modifications detected in user SQL commands. Boolean value. Defaults to true.
<code>showRowCounts</code>	Show number of rows in each table in the object navigator. Boolean value. Defaults to false.
<code>showSysTables</code>	Show system tables in the object navigator. Boolean value. Defaults to false.
<code>showSchemas</code>	Show object names like <code>schema.name</code> in object navigator. Boolean value. Defaults to true.
<code>resultGrid</code>	Show query results in Gui grid (as opposed to in plain text). Boolean value. Defaults to true.

showToolTips Show help hover-text. Boolean value. Defaults to true.

Chapter 4. Transfer Tool

Fred Toussi, The HSQL Development Group

Brief Introduction

Transfer Tool is a GUI program for transferring SQL schema and data from one JDBC source to another. Source and destination can be different database engines or different databases on the same server.

Transfer Tool works in two different modes. Direct transfer maintains a connection to both source and destination and performs the transfer. Dump and Restore mode is invoked once to transfer the data from the source to a text file (Dump), then again to transfer the data from the text file to the destination (Restore). With Dump and Restore, it is possible to make any changes to database object definitions and data prior to restoring it to the target.

Dump and Restore modes can be set via the command line with `-d (--dump)` or `-r (--restore)` options. Alternatively the Transfer Tool can be started with any of the three modes from the Database Manager's Tools menu.

The connection dialogue allows you to save the settings for the connection you are about to make. You can then access the connection in future sessions. These settings are shared with those from the Database Manager tool. See the appendix on Database Manager for details of the connection dialogue box.

From version 1.8.0 Transfer Tool is no longer part of the `hsqldb.jar`. You can build the `hsqldbutil.jar` using the Ant command of the same name, to build a jar that includes Transfer Tool and the Database Manager.

When collecting meta-data, Transfer Tool performs `SELECT * FROM <table>` queries on all the tables in the source database. This may take a long time with some database engines. When the source database is HSQLDB, this means memory should be available for the result sets returned from the queries. Therefore, the memory allocation of the java process in which Transfer Tool is executed may have to be high.

The current version of Transfer is far from ideal, as it has not been actively developed for several years. The program also lacks the ability to create UNIQUE constraints and creates UNIQUE indexes instead. However, some bugs have been fixed in the latest version and the program can be used with most of the supported databases. The best way to use the program is the DUMP and RESTORE modes, which allow you to manually change the SQL statements in the dump file before restoring to a database. A useful idea is to dump and restore the database definition separately from the database data.

Appendix A. HyperSQL File Links

HyperSQL Files referred to in this Guide

HyperSQL files referred to in the text may be retrieved from the canonical HyperSQL documentation site, <http://hsqldb.org/doc/2.0>, or from the same location you are reading this page from.



Note

If you are reading this document with a standalone PDF reader, only the [http://hsqldb.org/doc/2.0/...](http://hsqldb.org/doc/2.0/) links will function.

Pairs of local + <http://hsqldb.org/doc/2.0> links for referenced files.

- Local: `../verbatim/sample/sqltool.rc`
<http://hsqldb.org/doc/2.0/verbatim/sample/sqltool.rc>
- Local: `../verbatim/sample/sampleddata.sql`
<http://hsqldb.org/doc/2.0/verbatim/sample/sampleddata.sql>
- Local: `../verbatim/sample/sample.sql`
<http://hsqldb.org/doc/2.0/verbatim/sample/sample.sql>
- Local: `../verbatim/sample/pl.sql`
<http://hsqldb.org/doc/2.0/verbatim/sample/pl.sql>
- Local: `../verbatim/sample/plsql.sql`
<http://hsqldb.org/doc/2.0/verbatim/sample/plsql.sql>
- Local: `../verbatim/sample/dsv-sample.sql`
<http://hsqldb.org/doc/2.0/verbatim/sample/dsv-sample.sql>
- Local: `../verbatim/testrun/sqltool/sqljrt.sql`
<http://hsqldb.org/doc/2.0/verbatim/testrun/sqltool/sqljrt.sql>
- Local: `../verbatim/testrun/sqltool/sqlpsm.sql`
<http://hsqldb.org/doc/2.0/verbatim/testrun/sqltool/sqlpsm.sql>
- Local: `../verbatim/src/org/hsqldb/sample/SqlFileEmbedder.java`
<http://hsqldb.org/doc/2.0/verbatim/src/org/hsqldb/sample/SqlFileEmbedder.java>
- Local: `../apidocs/org/hsqldb/cmdline/SqlFile.html`
<http://hsqldb.org/doc/2.0/apidocs/org/hsqldb/jcmdline/SqlFile.html>