



cubeSQL

Admin Manual

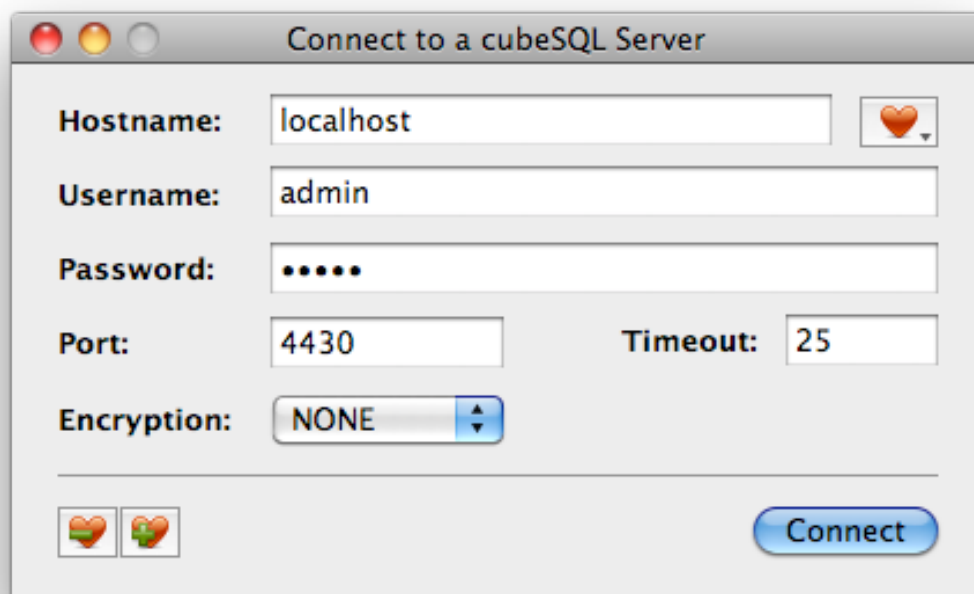
INTRODUCTION	3
Connecting to the Server	4
The Admin Window	5
Quitting the Server	5
Registering the Server	5
SERVER SECTION	7
Status Panel	8
Databases Panel	9
Uploading a Database	9
Downloading a Database	11
Stopping a Database	12
Tables & Indexes Panel	13
Console Panel	19
Clients Panel	20
SECURITY SECTION	21
Users & Groups Panel	22
Privileges Panel	25
ADVANCED SECTION	28
Schedules Panel	29
Restore Panel	31
Backup Panel	32
MVCC Panel	33
ADMINISTRATOR SECTION	34
Plugins Panel	35
Commands Panel	36
Settings Panel	37
Bugs Panel	39
Log Panel	40

INTRODUCTION

You can use cubeSQL Admin application to administer many aspects of the core cubeSQL server engine. It can be used to administer both locally running servers and servers running on remote machines.

Connecting to the Server

When you first launch cubeSQL Admin, you will be presented with a Connect dialog box. This dialog enables you to connect to either the server on the same computer or one that is visible on your network.



The presented values for Host and Port will connect to the server engine on the local computer running on its default 4430 port (officially registered at [IANA](https://iana.org) since 2008).

A newly installed server will have a default user named "admin" with the password set to "admin". The "admin" username is provided as the default in the connection dialog. If you are connecting to a newly installed server, then you can enter the default password and click the Connect button. If you are connecting to a server that has already been configured, be sure to provide the correct admin password. If you are connecting to a remote server, be sure that the firewalls are configured so that you can communicate on the selected port.

The admin user has full privileges for administering the server. That is, no operations are prohibited. Although it is possible to perform some administrative tasks as a user other than an admin user, it is strongly recommended that you connect to the server either as the admin user or as someone in the admin group. For information on adding users and groups, see the sections "Users & Groups Panel". Please refer to the section "Privileges Panel" for information on granting privileges to groups.

Because the admin user has full privileges, it is strongly recommended that you change admin's password as soon as it is convenient. For information on changing a user's password, see the section "Modifying Users".

The Admin Window

The Admin window's panels are organized into the following topics: Server, Security, Advanced, and Administrator. The Admin application refreshes the current panel periodically to keep the information in that panel up-to-date. For example, if you add a database by dragging it into the databases folder, it will appear on the Databases panel automatically.

All the operations that can be done via the cubeSQL Admin's GUI can also be performed programmatically, via the special custom commands that are described in the Command Reference. When you execute any such command, the admin's GUI is updated accordingly.

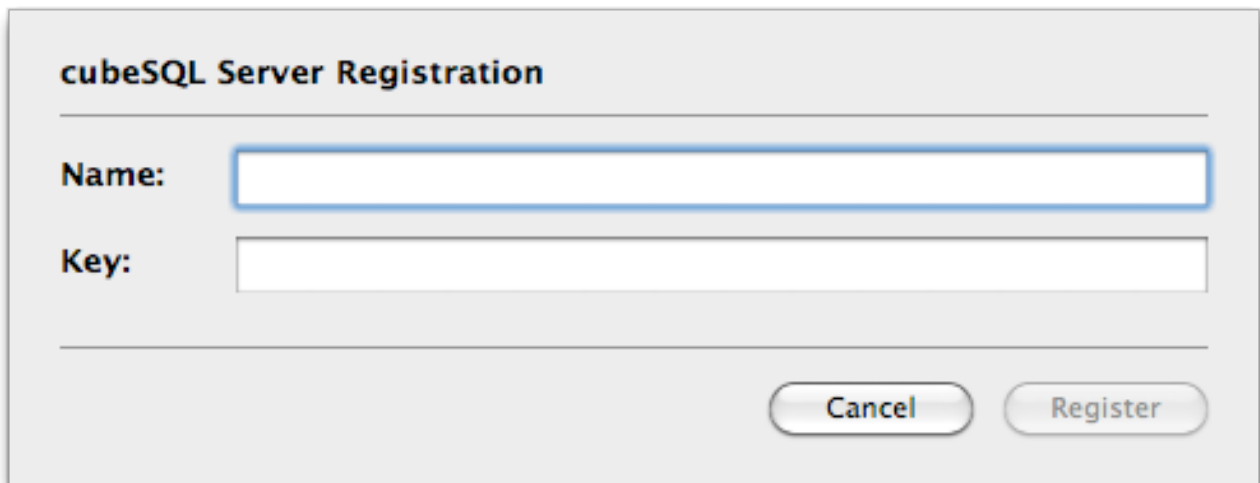
You can connect to more than one server at the same time. Simply choose File → Connect from an open Admin application window or from the menu bar. A new Connect dialog will appear, enabling you to establish a connection to another server. You will then get a new Admin window for that server.

Quitting the Server

To quit a running server, just select Quit Server from the Server menu (once connected). Sometimes the server needs to wait until every connected user has finished with any currently executing operations. Connected users will be disconnected automatically as soon as their current operation is finished.

Registering the Server

Without a proper registration key server runs in restricted mode where only custom commands are allowed and just two concurrent connections are accepted. In order to register the server select "Register Server" from the Server menu and enter the registration information you have obtained from SQLabs:



The image shows a 'cubeSQL Server Registration' dialog box. It has a title bar with the text 'cubeSQL Server Registration'. Below the title bar, there are two input fields. The first is labeled 'Name:' and the second is labeled 'Key:'. Both fields are empty. At the bottom right of the dialog box, there are two buttons: 'Cancel' and 'Register'.

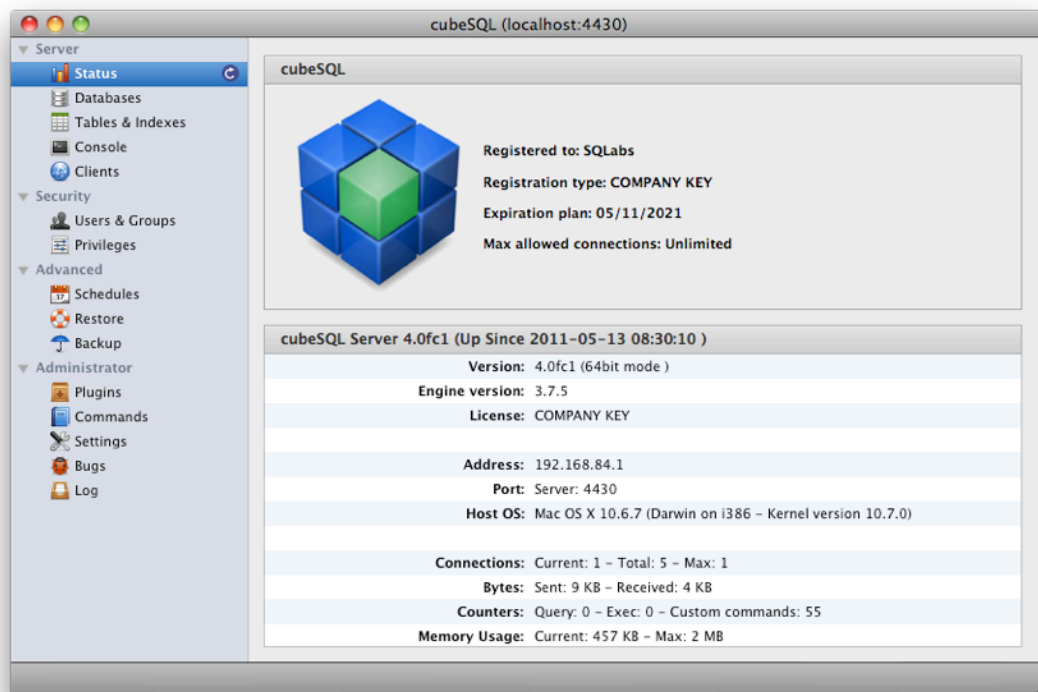
If you want to evaluate the server before your purchase select "Get a Developer Key" from the Server menu and you'll be redirected to a form page on the SQLabs web site where you can request an evaluation key.

Please note that with a Developer Key, server is fully registered but enters in restricted mode every two hours and then it needs to be restarted.

SERVER SECTION

Status Panel


The Status panel shows you overall information about a running server and displays useful details about the installed registration key.



Information like the number of connections or the IP address and port used by the server machine, the version of the server, the OS on which the server is running, plus some basic statistics about connections, memory usage and executed commands are displayed into this area.

Databases Panel

You use the Databases panel to manage the databases that the server serves. You can create new databases, start, stop, and drop databases here.



Creating a Database

In the Name field enter the name of the new database. If you want to encrypt it, enter the key for the database or leave the Key field blank if you want to create the database unencrypted. If you decide to encrypt the database at a later point, you can do so by issuing the SET KEY FOR DATABASE command or just selecting it and then choosing "Set Encryption Key" from the Server menu. If you want to change the encoding for the database, choose an encoding value from the pop-up menu. Click Create to begin database creation process on the server.

The new database will then be listed in the Databases list. It will have no tables until you add them.

To delete an existing database, select the database and click the Drop button.

Uploading a Database


If you have an existing database that is not currently being served by cubeSQL, you can install it in the databases folder by uploading it. You can upload from either the Admin application or with the UPLOAD DATABASE command in the language.

If you want to replace a database from the Admin application, first delete the database to be replaced and then upload the replacement.

1. Display the Databases panel in the Admin application.

A standard open-file dialog box appears.

A dialog appears, enabling you to rename the database and enter the encryption key for this database if it is encrypted.



When you click Upload, cubeSQL copies the database into the databases folder and immediately begins serving it.

You can also upload a database to the server by dragging it into the “databases” folder from the desktop.

If the database that you uploaded is encrypted and the server does not know its encryption key, then it cannot serve it. For that reason, it appears in gray in the Databases panel. You need to set its encryption key (either using the Admin application or via the proper custom command).

In the Admin application, select the database in the list and choose Server → Set Database Key. A dialog will appear in which you can enter the database’s encryption key.

In the language, call the SET KEY FOR DATABASE command, passing the database’s key in the command.

Downloading a Database

Downloading a database reverses the process. It takes a database that is installed in the “databases” folder and copies it to another location on your hard disk. When you download a database, it does not remove it from the “databases” folder and it continues to be served.

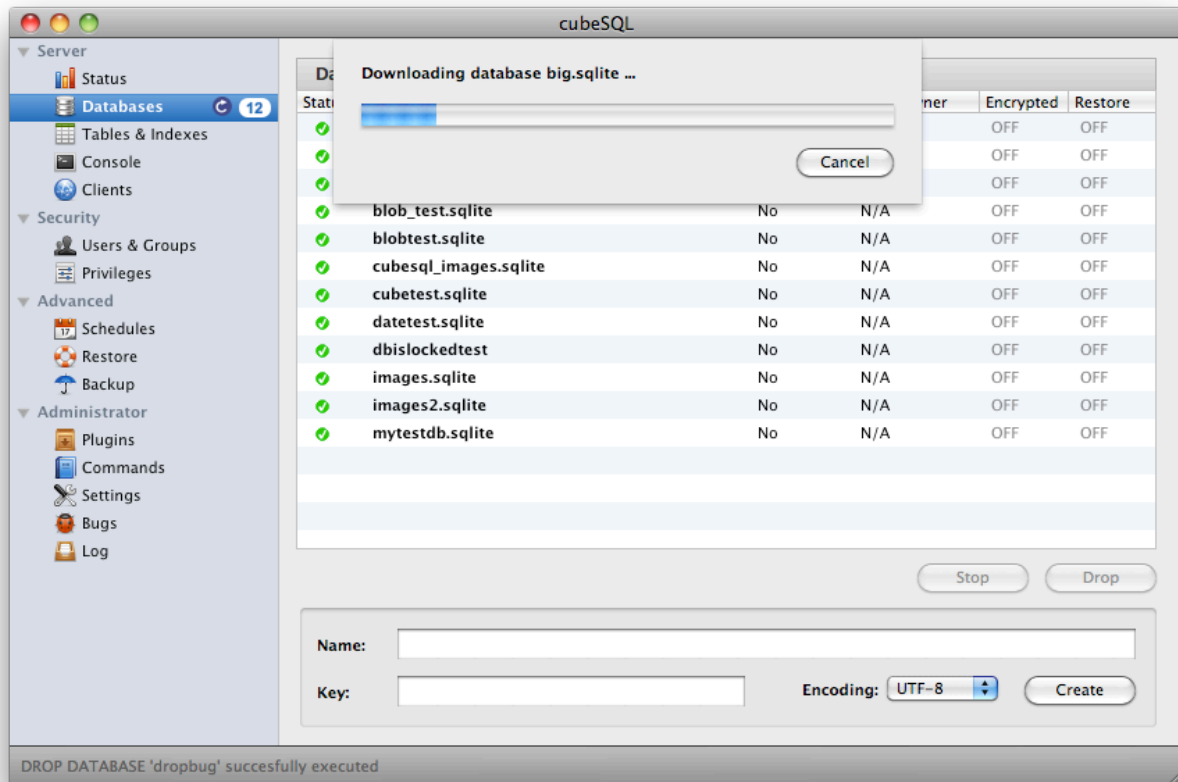
You can also download a database using the language, with the DOWNLOAD DATABASE command.

To download a database from the Admin application, do this:

- 1. Switch to the Databases panel in the Admin application.**
- 2. Click on the database you wish to download in the Databases list.**
- 3. Choose Server → Download Database.**

A standard save-file dialog box appears.

Navigate to the folder in which you want to save the database and click Choose. cubeSQL Admin displays a progress dialog.



4. Click Start to begin the download.

cubeSQL copies the database to the location you chose.

Stopping a Database

To stop a database, select it in the databases list and press the Stop button. To start a stopped database, select it and press the Start button.


The running or stopped status of a database is indicated by a bubble located to the left of the database's name. Green indicates a running database and red indicates a stopped database. By default, a newly created database is running. Stopping a database is a way to safely temporary disable database access.

Tables & Indexes Panel

You use the Tables & Indexes panel to create tables, fields and indexes for your databases.

Creating a Table

To create a new table for a database, switch to the Tables & Indexes panel.



The pop-up menu at the top of the panel lists all the databases on the server. If a database contains tables or indexes, then they are listed in the body of the panel.

To add a table to a database, do this:

1. Choose the database from the pop-up menu and click Add Table.

The Table Editor appears.

Table Editor

Table Name:

Name	Type	Default
col1	TEXT	

Field Name: ▼

Default Value: ▼

2. Enter the name of the table in the Table Name field.

You must now add at least one field to the table.

3. Enter a field name in the Field Name field and choose its data type from the pop-up menu to its right.

Your choices are as follows (in **BOLD** native sqlite type, others are REALbasic types):

Data Type	Description
TEXT	The value is a text string, stored using the database encoding (UTF-8, UTF-16BE or UTF-16LE).
INTEGER	The value is a signed integer, stored in 1, 2, 3, 4, 6, or 8 bytes depending on the magnitude of the value.
REAL	The value is a floating point value, stored as an 8-byte IEEE floating point number.
BLOB	The value is a blob of data, stored exactly as it was input.
CURRENCY	This is a 64-bit fixed-point number format that holds 15 digits to the left of the decimal point and 4 digits to the right.

Data Type	Description
BOOLEAN	SQLite does not have a separate Boolean storage class. Instead, Boolean values are stored as integers 0 (false) and 1 (true).
DATE	Stores year, month, and day values of a date in the format YYYY-MM-DD. The year value is four digits; the month and day values are two digits.
TIME	Stores hour, minute, and second values of a time in the format HH:MM:SS. The hours and minutes are two digits. The seconds values is also two digits, may include a optional fractional part, e.g., 09:55:25.248. The default length of the fractional part is zero.
TIMESTAMP	Stores both date and time information in the format YYYY-MM-DD HH:MM:SS. The lengths of the components of a TimeStamp are the same as for Time and Date, except that the default length of the fractional part of the time component is six digits rather than zero. If a TimeStamp values has no fractional component, then its length is 19 digits. If it has a fractional component, its length is 20 digits, plus the length of the fractional component.
DOUBLE	See REAL.
FLOAT	See REAL.
VARCHAR	See TEXT.
BINARY	See BLOB.

4. (Optional) Specify a default value for the field.

5. (Optional) Specify any field attributes from the Attributes pop-up menu, located to the right of the Default Value field.

Your choices are:

- Primary Key
- Primary Key auto increment
- Unique (single)
- Unique (multi)
- Not NULL

If you want to clear the selection, select the text and press Delete.

6. When all the attributes of the field are specified, click the Add Field button to add the field to the field list.

The new field is added to the list.

Continue adding fields as needed by clearing the entries and replacing them with new values. When you are finished, click Add Table to save the table to the database.

Creating an Index

After adding your tables, you may want to index certain fields. Indexes improve the performance of the database. You should index fields on which you plan to do the majority of your searches and sorts. Information retrieval and relational operations among tables are faster when indexed fields are used. Primary key fields are indexed automatically.

Fields that are not good candidates for indexing are those that only take on a few values (i.e., gender or race), fields that are rarely searched on, and any fields in small tables where search and sort times are unlikely to be long.

You should not index too many fields because each index adds to the size of the database and, as records are added and deleted, it takes time for the system to update each index.


Indexes are added via the Tables & Indexes panel. Select a database that has at least one table. Select the table in the list and click the Add Index button.

To create the indexes for a table, do this:

1. Choose a database from the pop-up menu and the name of a table in the Tables list.

2. Click the Add Index button below the list.

The Index Editor appears.



The Index Editor shows the fields in the selected table. To create an index, click all the fields that will comprise the index. If you choose more than one field, it will create a composite index rather than more than one distinct index.

Enter the name of the index in the field at the top of the dialog.
Usually you will use the names of the field or fields that comprise the index.

Click the checkbox for the field or fields that you want to include in the index.
If you click two or more fields, a composite index will be created, not separate indexes for each field.

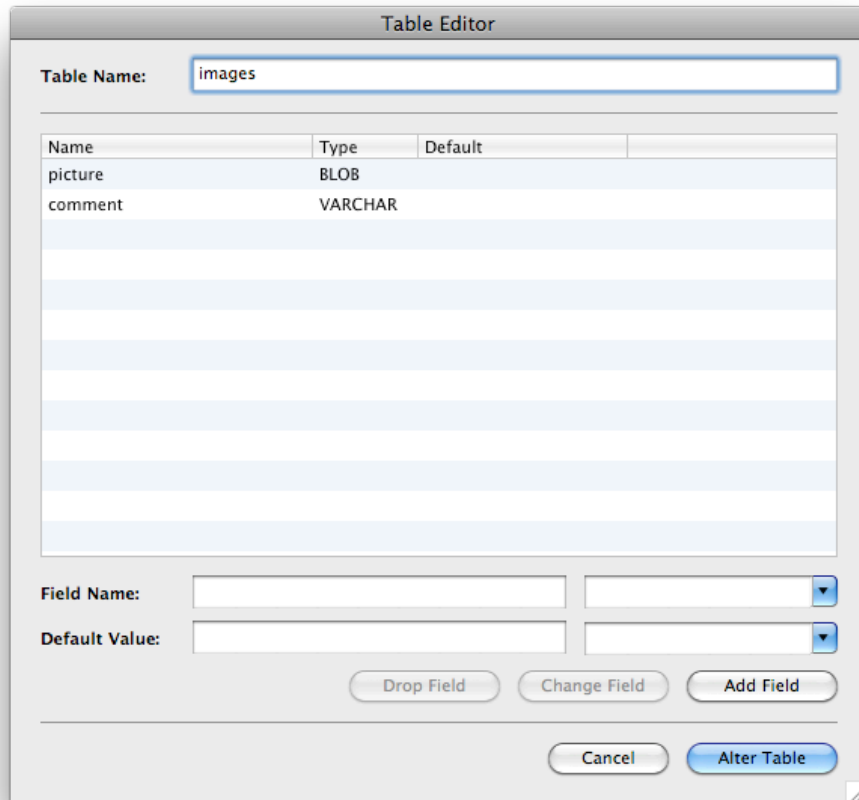
For example, you may want to create a composite index on LastName and FirstName fields so that the database can quickly work with different people with the same last name.

Click Save to save the index.
If desired, repeat this process to create additional indexes. To create indexes in another database, select another database in the Databases list and repeat this process.

Editing a Table Schema

You can edit the table schema for any existing table. To do so, display the Structures panel, select the database you want to modify. Click the Alter Table button. The Table Editor for that database appears.

When you are in the process of creating fields, you can alter the schema by deleting new fields and adding other fields.



The Table Editor dialog box is shown with the title 'Table Editor'. It features a 'Table Name:' field containing the text 'images'. Below this is a table with three columns: 'Name', 'Type', and 'Default'. The table contains two rows: 'picture' with type 'BLOB' and 'comment' with type 'VARCHAR'. Below the table are two rows of input fields for 'Field Name:' and 'Default Value:', each with a dropdown arrow. At the bottom are four buttons: 'Drop Field', 'Change Field', 'Add Field', and 'Alter Table' (highlighted in blue). There is also a 'Cancel' button.

Name	Type	Default
picture	BLOB	
comment	VARCHAR	

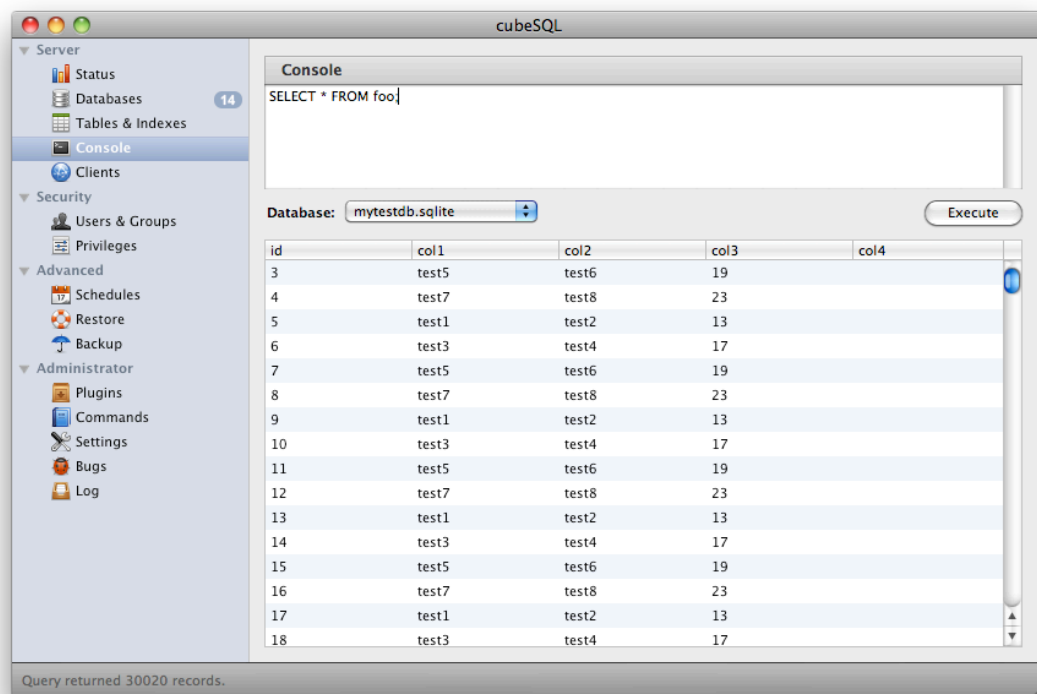
To delete a field, highlight the field in the list and click the Delete button.

To add a field, deselect any existing field and then repeat the process of adding fields described in the previous section.

When you click Alter Table to display the Table Editor, you cannot drop or change fields, only add fields. When you are finished, click Alter Table to save your changes and put away the dialog box.

Console Panel


The Console panel enables you to execute cubeSQL commands interactively. In addition to standard sql statements any command in the Command reference can be executed.



Since cubeSQL is sqlite based ,you can get a list of supported sql commands into the official [sqlite](https://www.sqlite.org/) website.

Clients Panel

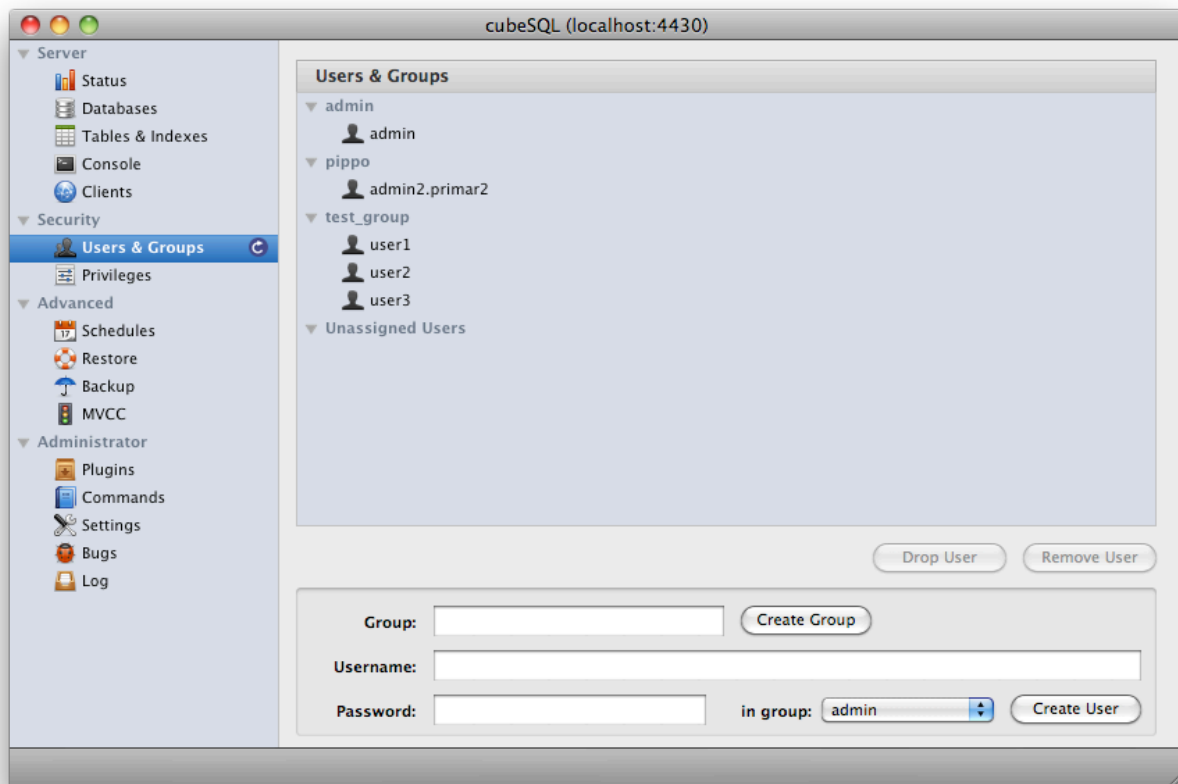
The Clients panel displays a list of connected clients that are connected to the server that the instance of the Admin window is connected to. It lists the client ID, IP address, the username, date of connection, last activity date, database on which the client is connected, version, and type. You have the ability to disconnect a client from this panel.



SECURITY SECTION

Users & Groups Panel

Use the Users and Groups panel to create and edit the list of users who have access to the server and the databases. This includes both server administrators and end users.



By default there is only one group with one user. (the admin group with the admin user). You can add other users to the admin group and create new groups that will be assigned fewer privileges.

You use the Users & Groups panel to first create and edit the list of groups. You can then assign users to groups.

To create the privileges system, do this:

- Create groups that will have different access privileges.
- Create a username and password for each user and place users into groups.
- Use the Privileges panel, to grant privileges to each group.

The users in a group automatically acquire the privileges of their group. A user can be in more than one group. Such a user gets the privileges of those groups.

The admin user is the default user that ships with the server and it is a member of the admin group.

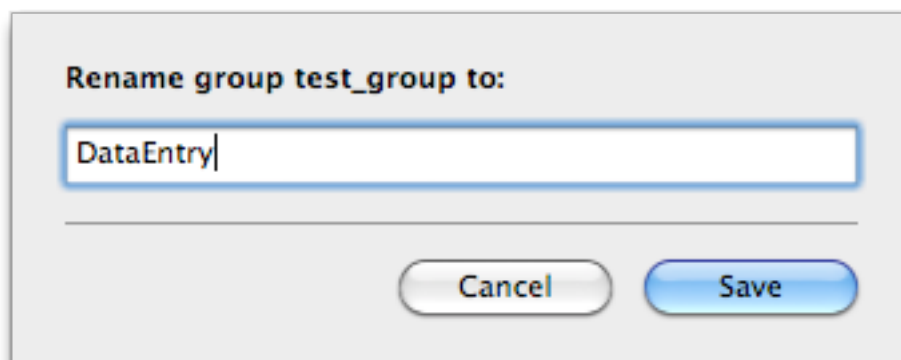
The admin group has been granted all possible privileges and this cannot be changed. Privileges are granted to groups, not users. A user automatically gets all the privileges that have been granted to the groups that user is in. For example, you might have a Data Entry group that has permissions to write data only to the tables that concern them.

In addition to administrative functions, privileges control whether a user can select, add, modify, or delete individual records. Privileges can be granted at the level of individual tables, for the database as a whole, and for the server itself.

To create a new group, enter a new group name click the Create Group button located under the Users & Groups list. The new group is added to the list of groups.

To delete an existing group, select the group and click the Drop Group button. You cannot delete the Admin group.

To rename a group, double-click on the group name . The Rename Group dialog appears.



Enter the new name for the group and click Save.
You cannot rename or delete the admin group.

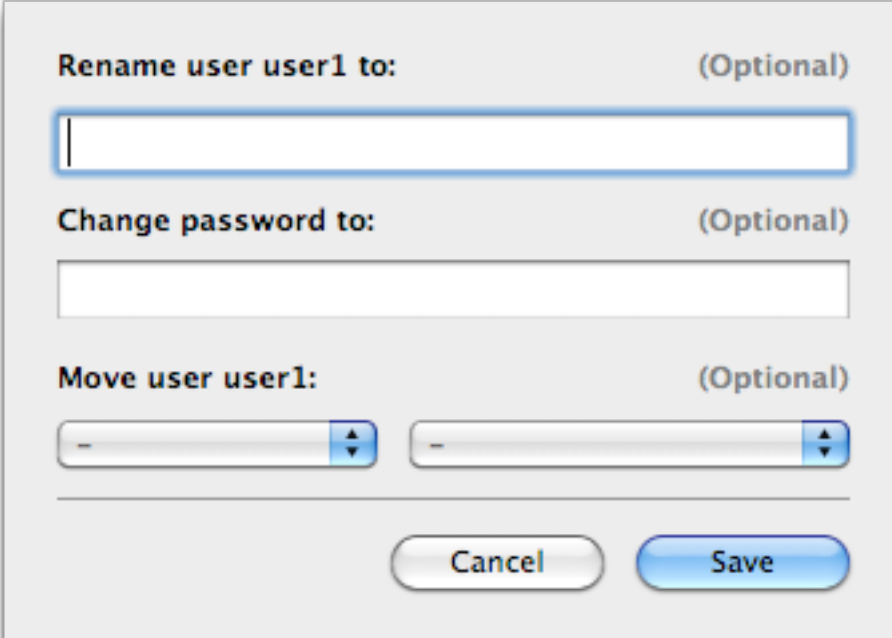
Creating Users

To create a new user, enter the user name and password into the username and password fields. Choose a group that the user will be in from the In Group pop-up menu and click Create User. (You can assign a user to more than one group later.)

Using the CREATE USER command, it is possible to create a user without assigning him or her to a group. In this case, the user will appear in the "Unassigned Users" category. You can assign them to groups as explained in the following section.

Modifying Users

To rename a user, change the password, or add a user to another group, double-click the user in the Users list. The following dialog shown will appear.



The dialog box is titled "Modifying Users" and contains three sections for user modification:

- Rename user user1 to:** (Optional) - A text input field.
- Change password to:** (Optional) - A text input field.
- Move user user1:** (Optional) - Two dropdown menus for selecting a new group.

At the bottom of the dialog are two buttons: "Cancel" and "Save".

If you want to rename the user or change his password, enter the new information in the appropriate fields. You cannot change the "admin" user's name but you can change the admin's password. To add the user to another group, use the Move User area.

The pop-up menu has items for moving the user and adding the user. Select the desired option and choose the group from the second pop-up menu. select the new group from the Add User pop-up menu.

If you want to assign the user to another group, repeat this process. When you are finished making changes, click Save.

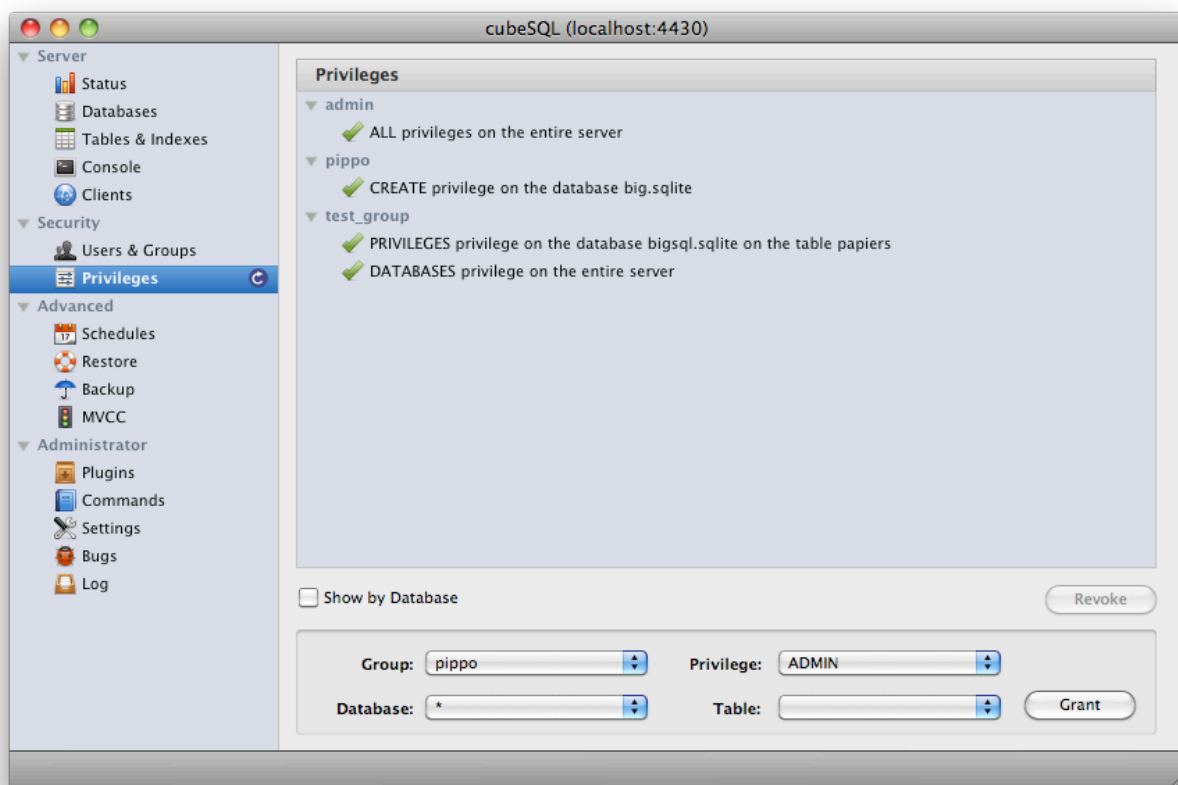
Deleting Users

You can delete a user from the list of users by and clicking Drop User. The Drop User command drops the user from all groups and the server. If you want to remove a user from only one group, select the user in that group and click Remove User.

Privileges Panel

cubeSQL uses a system of access privileges that you use to allow or deny actions on server resources. Privileges are assigned to server groups. Users in a group have all of the privileges of that group. A user can belong to more than one group. If a user is in more than one group, then the user has all the privileges of every group in which the user is a member.

By default, cubeSQL has one group named admin. Every user in the admin group is considered an admin user, and as such, has every possible privilege. The admin user is a permanent member of the admin group and cannot be renamed or removed. Neither the admin user nor the admin group may be removed from a server. Also, privileges cannot be removed from the admin group. However, users other than admin can be removed from the admin group, stripping them of admin privileges.



In order to give a user a particular set of privileges, a group must exist with those privileges. By adding the user to the group, the user will automatically gain all the privileges of the group. Deleting a group removes the group's privileges from the members of that group.

There are three types of privileges:

- **Server privileges:** the privilege is granted for the entire server

- **Database privileges:** the privilege is granted only for a specified database
- **Table privileges:** the privilege is granted only for a specified table in a database.

You set the scope of a privilege when you grant a privilege to a group. The bottom portion of the Privileges panel contains pop-up menus for Group, Database, and Table. If a privilege is to be granted for the entire server, you specify only the Group and set Database to "*". To grant a privilege at the database level, then specify the database from this pop-up and leave Table blank. To grant a privilege for only a table in a database, specify the database and the table from those two pop-up menus.

Privileges include:

- **Admin:** has all the privileges of the Admin group
- **Privileges:** can edit group privileges
- **Preferences:** can see and edit server settings
- **Databases:** can create and delete databases
- **Upload:** can upload databases to the server
- **Download:** can download all databases from the server
- **Create:** can create tables in the database
- **Drop:** can drop tables in the database
- **Select:** can query any table in the database
- **Insert:** can insert records into any table in the database
- **Update:** can update records in any table in the database
- **Delete:** can delete records in any table in the database
- **Upload:** can upload the database to the server, replacing it
- **Download:** can download the database
- **Pragma:** can execute PRAGMA commands to change database settings
- **Plugin:** can execute the plug-in commands, DISABLE PLUGIN and ENABLE PLUGIN
- **Restore:** can restore a database
- **Backup:** can execute backups
- **Select:** can query the table
- **Insert:** can insert records into the table
- **Update:** can update records in the table
- **Delete:** can delete records in the table

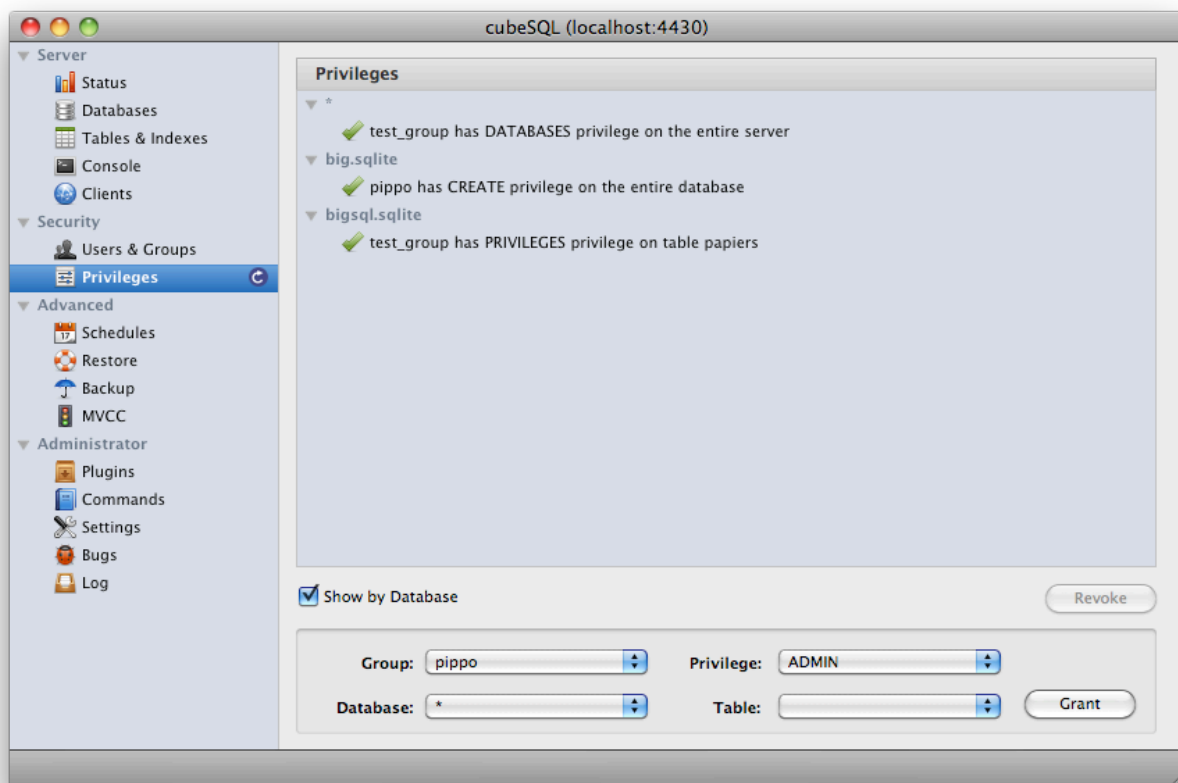
In the Command Reference, the privileges that are required to execute each command are noted.

Using the Privileges Panel

You use the Privileges panel to assign privileges to groups. To give a user privileges, create a group with the desired privileges and then add the user to that group. You may also change the privileges for a group after users have been added to the group, of course.

The Privileges panel is divided into two sections:

The top section lists all the groups and the privileges assigned to them. The bottom section is used to assign privileges to groups. Privileges are nested in groups. The default view organizes the information by group. Click the Show by Database checkbox to organize the information by database. In this display, the privileges are divided into a section for the entire server and separate sections for each database on the server.



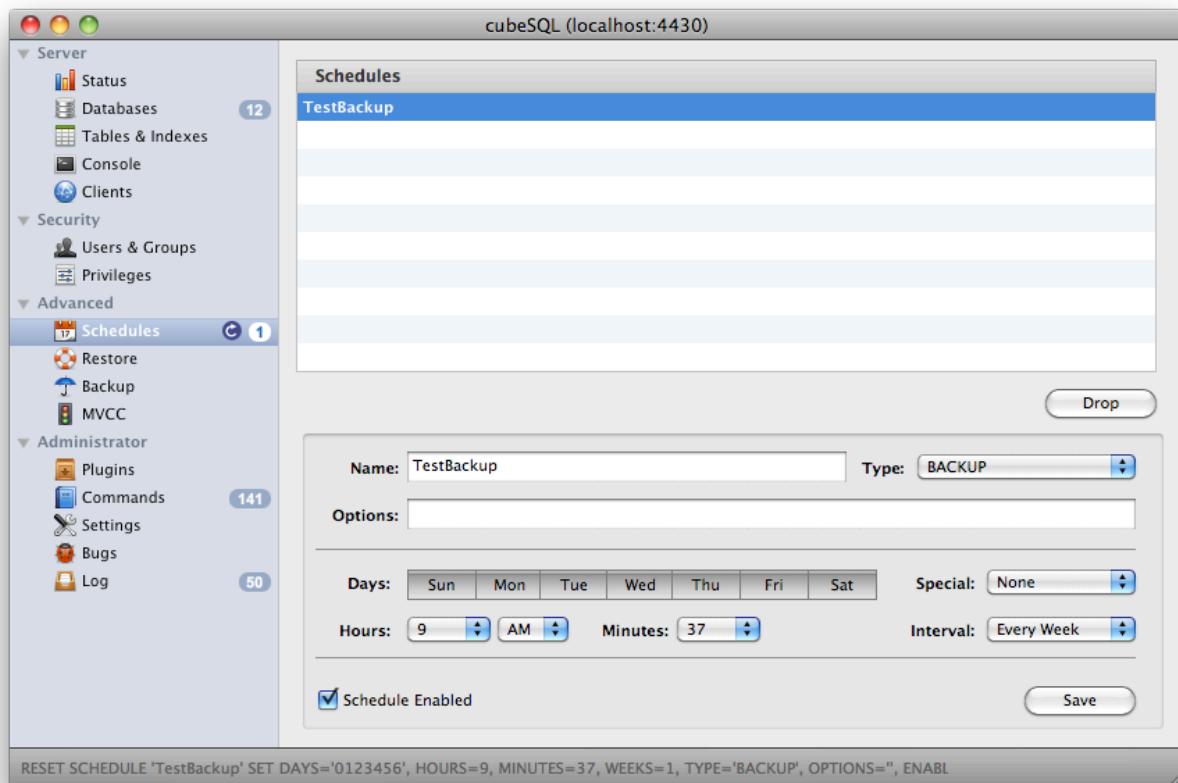
The bottom section is used to assign privileges. You use the Database and Table pop-up menus to specify the scope of the privilege to be granted. To specify Server privileges, specify only the Group to be granted the privilege and the privilege to be granted. To specify Database privileges, specify the Group, the database, and the privilege. To specify Table privileges, specify the Group, the database, the Table, and the privilege.

There are no limitations or constraints: you can check as many or as few privileges as you like.

ADVANCED SECTION

Schedules Panel

cubeSQL includes a powerful integrated scheduling subsystem. Built-in schedules enables you to schedule operations to be performed at a given date/time.



To create or modify a schedule click the Schedules item in the Advanced group.


Set schedule characteristics like its name, its type, the Days/Time in which it should run and make sure that the Schedule Enabled checkbox is checked.

Three type of schedule can be created, the most important is a BACKUP schedule that enables you to perform a backup at a predefined date and time. cubeSQL supports also a SHELL schedule that enables you to perform a shell command (specified into the Options field) and a SQL schedule used to perform an SQL command at a give date (SQL command to be executed specified into the Options field). Scheduling an SQL command is disabled in this version of the server.

A schedule can be delete using the Drop button or temporary disable just unchecking the Schedule Enable checkbox and then Save the changes.

Once a schedule has been created, it can be attached to one or more databases.


Just select the Databases panel, click on a database and from the Server menu select "Manage Schedules":



In the "Manage Schedules" window you'll be able to Attach and Detach schedules to a given database.

Restore Panel

cubeSQL supports an advance feature called Restore Points. When RESTORE is activated for a given backup (selected from the popup menu) then every WRITE operation is logged into a special restore file (you can have a list of all logged write operations using the "Show Log" button).




While backups are important in order to preserve your data from corruption due to power-loss or due to operating system bugs, usually there is nothing you can do from user's mistakes. For example, what happens if a valid user send a DELETE FROM table command? or a DROP TABLE command?

Well... you'll lose all the information contained in that table and most of the time this is not a very good event but thanks to the RESTORE feature you'll be safe even in case of user's mistakes. RESTORE is like an unlimited UNDO queue linked to the entire database. So in case of a DROP TABLE command, you can just RESTORE your database one line before that wrong command and all your data will be automatically restored.

Note that enabling RESTORE on a database will slow down all WRITE operations because data and sql commands needs to be written twice.

Backup Panel

In the Admin application, you can perform backups and/or manage backups using the powerful Backups panel. A popup menu will list all installed databases and you can use the "Backup Now" button to create a backup copy of a running database file.



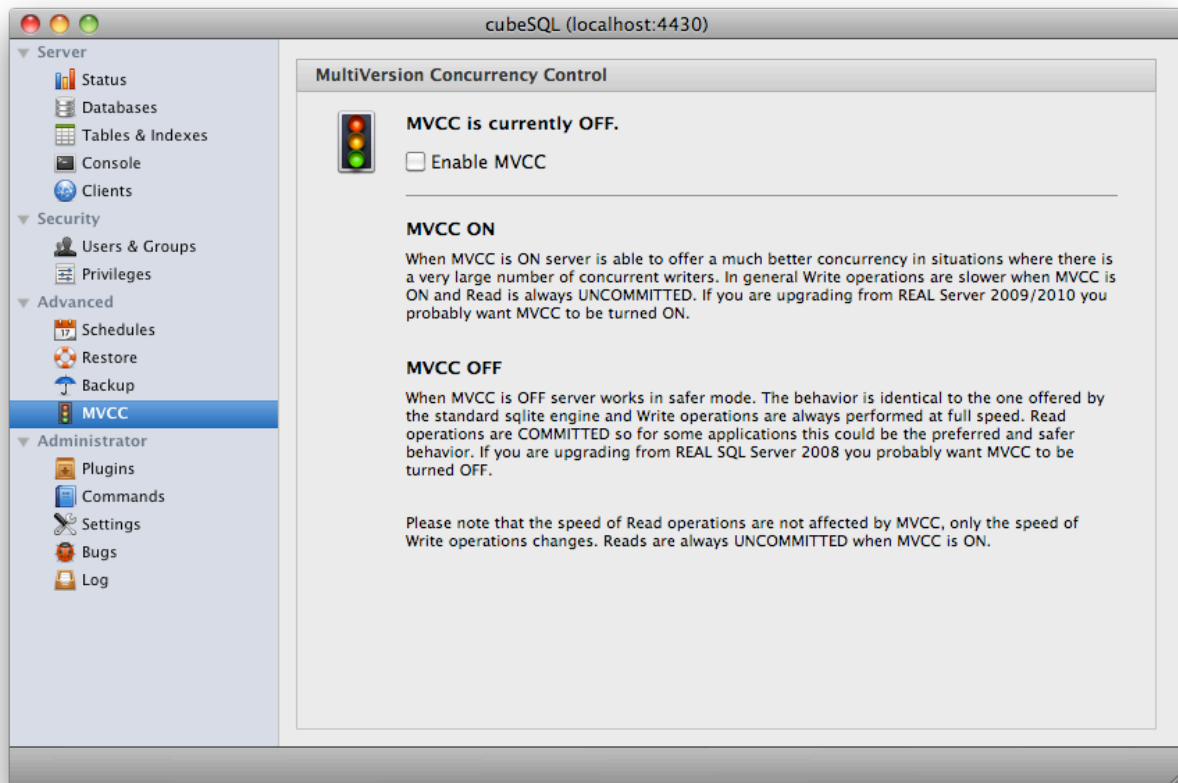
Note that while a backup copy is performed, that database will be fully accessible for both readers and writers.

Once a backup copy has been done you'll be able to delete it (using the Delete Backup button) or in case of database problems or database corruption you'll be able to fully restore a backup copy just selecting it from the ListBox and pressing the Restore Backup button.

At any time you can obtain a list of available backups for a given database using the "Show Backups" button. Note that a backup can also be scheduled (see "Schedule Panel" section).

If you want to manage backups with the language instead of the Admin application, take a look at the example custom application "Backup" in the Examples folder in your cubeSQL Server folder.

MVCC Panel



In order to increase concurrency cubeSQL supports MultiVersion Concurrency Control (MVCC), a standard technique for avoiding conflicts between reads and writes of the same object. MVCC guarantees that each transaction sees a consistent view of the database.


When MVCC is ON server is able to offer a much better concurrency in situations where there is a very large number of concurrent writers. In general Write operations are slower when MVCC is ON and Read is always UNCOMMITTED. If you are upgrading from REAL Server 2009/2010 you probably want MVCC to be turned ON.

When MVCC is OFF server works in safer mode. The behavior is identical to the one offered by the standard sqlite engine and Write operations are always performed at full speed. Read operations are COMMITTED so for some applications this could be the preferred and safer behavior. If you are upgrading from REAL SQL Server 2008 you probably want MVCC to be turned OFF.

ADMINISTRATOR SECTION

Plugins Panel

cubeSQL has a modular plug-in architecture and you can use it to extend the SQL language, change how your results are sorted or even create your own custom server commands.

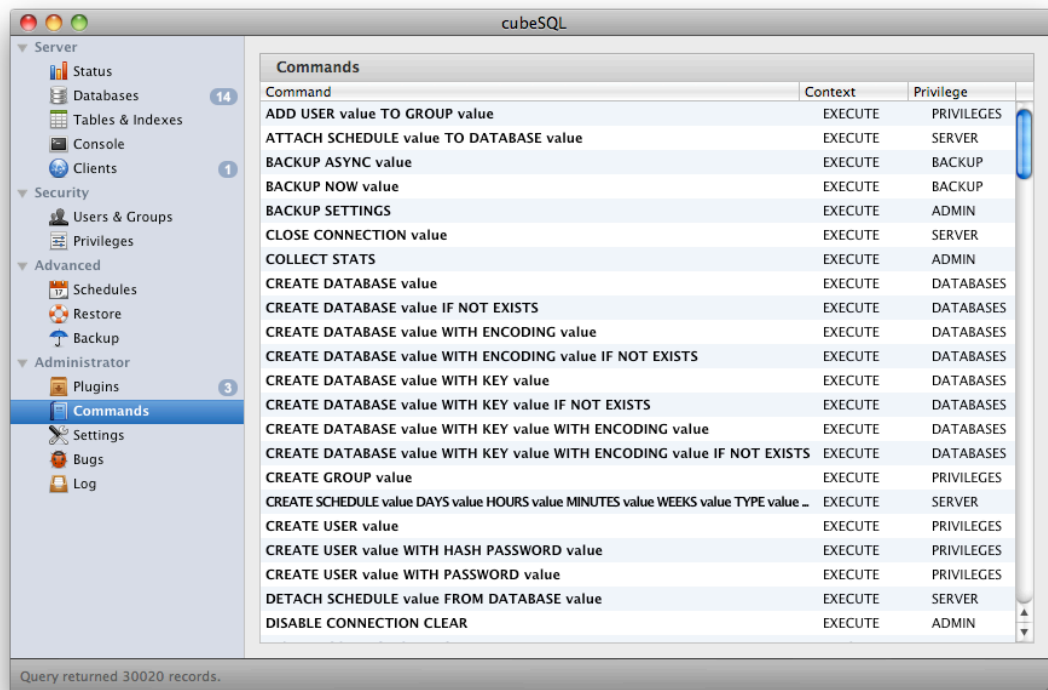


The plug-in architecture gives you the ability to enhance the power of cubeSQL in ways limited only by your imagination. A native plug-in SDK is included within each server's installation.

This panel enables you to list and enable/disable all the installed plugins.

Commands Panel

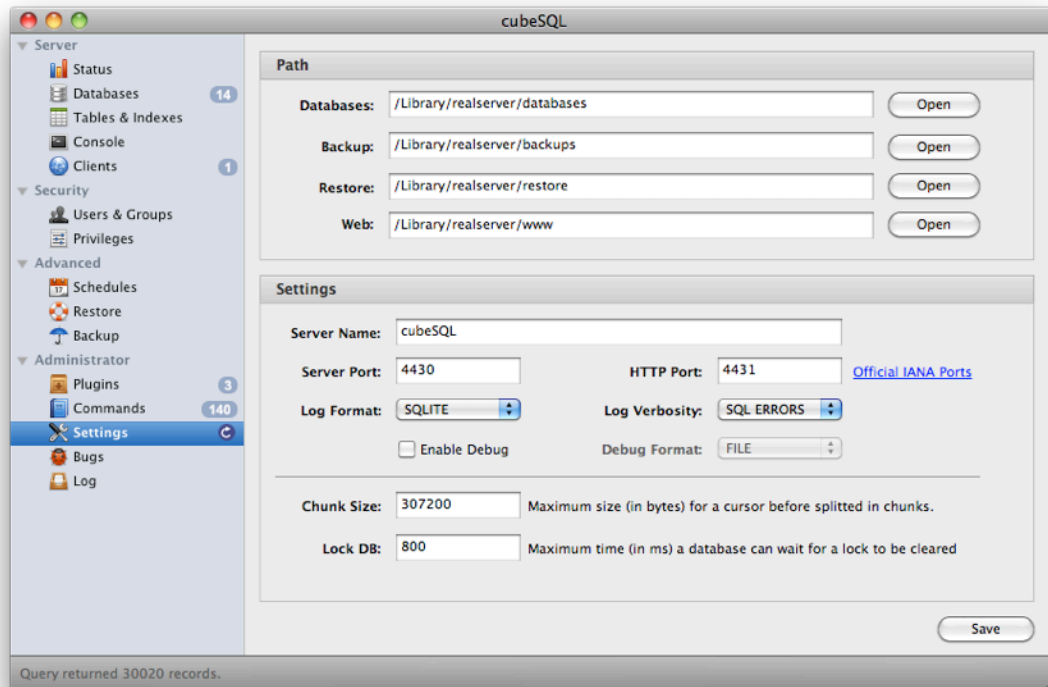
The Commands panel lists the syntaxes of all the special cubeSQL commands available on the server. It also lists the privileges that are required to execute the commands and the Context they needs (can be EXECUTE or SELECT and SELECT means that a RecordSet is returned).



Note that the Commands Panel lists also plug-ins defined custom commands in addition to the one built-into the server.

Settings Panel

To make changes to the server's settings, click the Settings panel in the Administrator group in the browser.



Path Settings

The Path area enables you to set the paths to the Databases, Backups, and Restore folders. You can modify these paths entering a different path (Web path is currently unused in this version).

Server Settings

If desired, enter a new name for the server. You don't lose functionality if the server has no name.

The Server Port property is the port that the Admin application uses to communicate with the server engine. Under most circumstances, the default value of 4430 will work fine. If you need to change the port for some reason, enter a new port number. If you change the port, the new port setting won't take affect until the server is restarted. You may also need to update the computers' firewall settings to permit communication on the selected port. If you encounter problems connecting, check your firewall settings. (HTTP Port is currently unused in this version).

The Verbose Logging option controls the type of information that is written to the log. There are four settings: None, SQL Errors, and SQL Commands, and Debug. Debug is

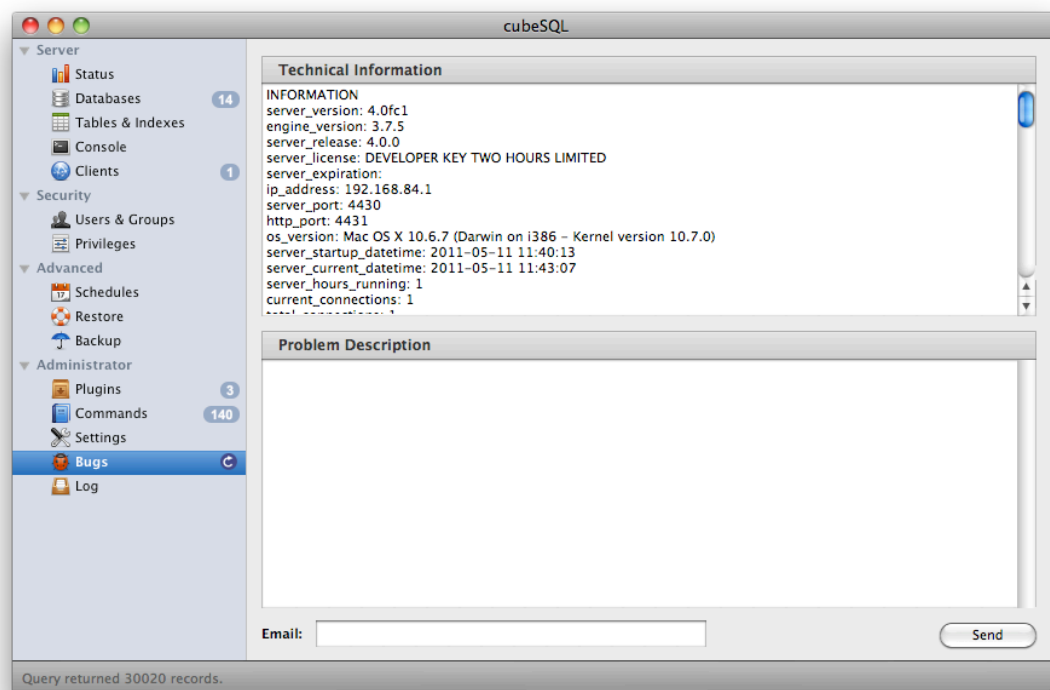
selected by checking the Enable Debug checkbox. You can write out the log in either of two formats, SQLite and Text. The two formats are shown in the section "Log Panel" .

The Chuck Size value specify the maximum size (in bytes) for a cursor. When the result of a query reaches the Chuck Size limit than the result is split and each chuck is sent to the client. Increase Chuck Size to speed-up very big queries, decrease that value in order to save memory.

The Locks DB settings specify the maximum time (in milliseconds) that a database can wait for a write lock to be cleared in a given table. Increase this value is you have many clients that tries to write to the same table concurrently.

Bugs Panel

Use the bugs panel to submit any bugs that you find in cubeSQL.




The top section contains information on the state of the server and its operating environment. The bottom section is for you to describe the bug in enough detail so the SQLabs engineers can replicate it.

Enter your email in the email field and click Send to report your bug.

Log Panel

You use the Log panel to review the server log. The Log panel displays the last XX entries in the log. Enter the number of entries that you want to see into "Entries to show" field in the Log panel. If you want to see entries other than the last XX rows, use the SHOW LOG command. It allows you to specify a date range instead of the last rows.

If you want to change the format, change value of the Log Format in the Settings panel (see "Changing Server Settings"). You can also set the verbosity (SQL errors, SQL commands, or none), enable or disable debug mode, and choose the debug format.



The log includes columns for the datetime stamp, description, operation, address of the client machine from which the command originated, the username, and the database, if relevant. Only the timestamp and the description are shown within the Admin application. You can access all the fields by submitting the SHOW LOG command.