
pgAdmin 4 Documentation

Release 3.2

The pgAdmin Development Team

Jul 25, 2019

CONTENTS



Welcome to pgAdmin 4. pgAdmin is the leading Open Source management tool for Postgres, the world's most advanced Open Source database. pgAdmin 4 is designed to meet the needs of both novice and experienced Postgres users alike, providing a powerful graphical interface that simplifies the creation, maintenance and use of database objects.

Contents:

GETTING STARTED

Pre-compiled and configured installation packages for pgAdmin 4 are available for a number of desktop environments; we recommend using an installer whenever possible.

In a Server Deployment, the pgAdmin application is deployed behind a webserver or with the WSGI interface. If you install pgAdmin in server mode, you will be prompted to provide a role name and pgAdmin password when you initially connect to pgAdmin. The first role registered with pgAdmin will be an administrative user; the administrative role can use the pgAdmin *User Management* dialog to create and manage additional pgAdmin user accounts. When a user authenticates with pgAdmin, the pgAdmin tree control displays the server definitions associated with that login role.

Contents:

1.1 Server Deployment

pgAdmin may be deployed as a web application by configuring the app to run in server mode and then deploying it either behind a webserver running as a reverse proxy, or using the WSGI interface.

The following instructions demonstrate how pgAdmin may be run as a WSGI application under Apache HTTP, using `mod_wsgi`.

See also:

For detailed instructions on building and configuring pgAdmin from scratch, please see the README file in the top level directory of the source code. For convenience, you can find the latest version of the file [here](#), but be aware that this may differ from the version included with the source code for a specific version of pgAdmin.

1.1.1 Requirements

Important: Some components of pgAdmin require the ability to maintain affinity between client sessions and a specific database connection (for example, the Query Tool in which the user might run a BEGIN command followed by a number of DML SQL statements, and then a COMMIT). pgAdmin has been designed with built-in connection management to handle this, however it requires that only a single Python process is used because it is not easily possible to maintain affinity between a client session and one of multiple WSGI worker processes.

On Windows systems, the Apache HTTP server uses a single process, multi-threaded architecture. WSGI applications run in `embedded` mode, which means that only a single process will be present on this platform in all cases.

On Unix systems, the Apache HTTP server typically uses a multi-process, single threaded architecture (this is dependent on the MPM that is chosen at compile time). If `embedded` mode is chosen for the WSGI application, then there will be one Python environment for each Apache process, each with its own connection manager which will lead to loss of connection affinity. Therefore one should use `mod_wsgi`'s `daemon` mode, configured to use a single process. This will launch a single instance of the WSGI application which is utilised by all the Apache worker processes.

Whilst it is true that this is a potential performance bottleneck, in reality pgAdmin is not a web application that's ever likely to see heavy traffic unlike a busy website, so in practice should not be an issue.

Future versions of pgAdmin may introduce a shared connection manager process to overcome this limitation, however that is a significant amount of work for little practical gain.

1.1.2 Configuration

In order to configure pgAdmin to run in server mode, it may be necessary to configure the Python code to run in multi-user mode, and then to configure the web server to find and execute the code.

Note that there are multiple configuration files that are read at startup by pgAdmin. These are as follows:

- `config.py`: This is the main configuration file, and should not be modified. It can be used as a reference for configuration settings, that may be overridden in one of the following files.
- `config_distro.py`: This file is read after `config.py` and is intended for packagers to change any settings that are required for their pgAdmin distribution. This may typically include certain paths and file locations.
- `config_local.py`: This file is read after `config_distro.py` and is intended for end users to change any default or packaging specific settings that they may wish to adjust to meet local preferences or standards.

Python

From pgAdmin 4 v2 onwards, server mode is the default configuration. If running under the desktop runtime, this is overridden automatically. There should typically be no need to modify the configuration simply to enable server mode to work, however it may be desirable to adjust some of the paths used.

In order to configure the Python code, follow these steps:

1. Create a `config_local.py` file alongside the existing `config.py` file.
2. Edit `config_local.py` and add the following settings. In most cases, the default file locations should be appropriate:

NOTE: You must ensure the directories specified are writeable by the user that the web server processes will be running as, e.g. `apache` or `www-data`.

```
LOG_FILE = '/var/log/pgadmin4/pgadmin4.log'  
SQLITE_PATH = '/var/lib/pgadmin4/pgadmin4.db'  
SESSION_DB_PATH = '/var/lib/pgadmin4/sessions'  
STORAGE_DIR = '/var/lib/pgadmin4/storage'
```

4. Run the following command to create the configuration database:

```
# python setup.py
```

5. Change the ownership of the configuration database to the user that the web server processes will run as, for example, assuming that the web server runs as user `www-data` in group `www-data`, and that the SQLite path is `/var/lib/pgadmin4/pgadmin4.db`:

```
# chown www-data:www-data /var/lib/pgadmin4/pgadmin4.db
```

Apache HTTPD Configuration (Windows)

Once Apache HTTP has been configured to support `mod_wsgi`, the pgAdmin application may be configured similarly to the example below:


```
<VirtualHost *>
  ServerName pgadmin.example.com
  WSGIScriptAlias / "C:\Program Files\pgAdmin4\web\pgAdmin4.wsgi"
  <Directory "C:\Program Files\pgAdmin4\web">
    Order deny,allow
    Allow from all
  </Directory>
</VirtualHost>
```

Now open the file `C:\Program Files\pgAdmin4\web\pgAdmin4.wsgi` with your favorite editor and add the code below which will activate Python virtual environment when Apache server runs.

```
activate_this = 'C:\Program Files\pgAdmin4\venv\Scripts\activate_this.py'
execfile(activate_this, dict(__file__=activate_this))
```

Note: The changes made in `pgAdmin4.wsgi` file will revert when `pgAdmin4` is either upgraded or downgraded.

Apache HTTPD Configuration (Linux/Unix)

Once Apache HTTP has been configured to support `mod_wsgi`, the `pgAdmin` application may be configured similarly to the example below:

```
<VirtualHost *>
  ServerName pgadmin.example.com

  WSGIDaemonProcess pgadmin processes=1 threads=25 python-home=/path/to/python/
  ↪virtualenv
  WSGIScriptAlias / /opt/pgAdmin4/web/pgAdmin4.wsgi

  <Directory /opt/pgAdmin4/web>
    WSGIProcessGroup pgadmin
    WSGIApplicationGroup %{GLOBAL}
    Order deny,allow
    Allow from all
  </Directory>
</VirtualHost>
```

Note: If you're using Apache HTTPD 2.4 or later, replace the lines:

```
Order deny,allow
Allow from all
```

with:

```
Require all granted
```

Adjust as needed to suit your access control requirements.

1.2 pgAdmin Login Dialog

Use the *pgAdmin Login* dialog to log in to `pgAdmin`:

pgAdmin 4 Login

Login

Language English

[Forgotten your password?](#)

Use the fields in the *pgAdmin Login* dialog to authenticate your connection:

1. Provide the email address associated with your account in the *Email Address* field.
2. Provide your password in the *Password* field.
3. Click the *Login* button to securely log into pgAdmin.

Recovering a Lost Password

If you cannot supply your password, click the *Forgotten your password?* button to launch a password recovery utility.

Recover pgAdmin 4 Password

Enter the email address for the user account you wish to recover the password for:

1. Provide the email address associated with your account in the *Email Address* field.
2. Click the *Recover Password* button to initiate recovery. An email, with directions on how to reset a password, will be sent to the address entered in the *Email Address* field.

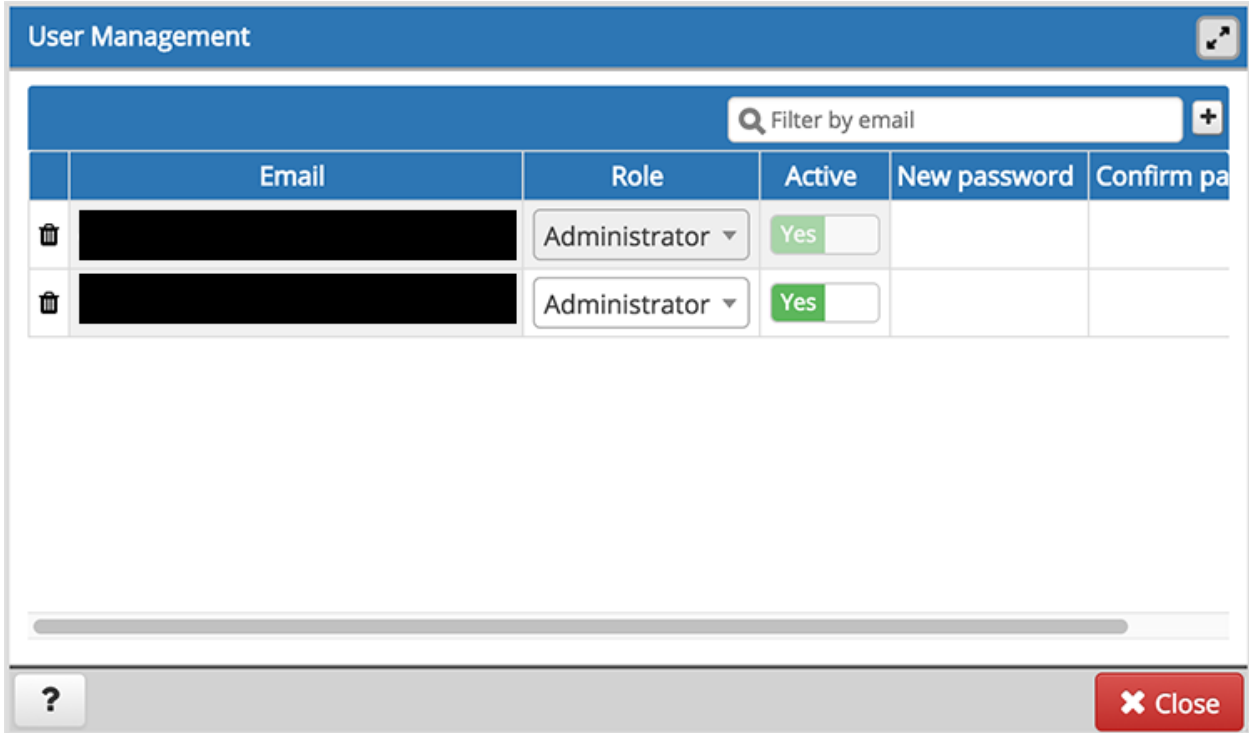
If you have forgotten the email associated with your account, please contact your administrator.

1.3 User Management Dialog

When invoking pgAdmin in desktop mode, a password is randomly generated, and then ignored. If you install pgAdmin in server mode, you will be prompted for an administrator email and password for the pgAdmin client.

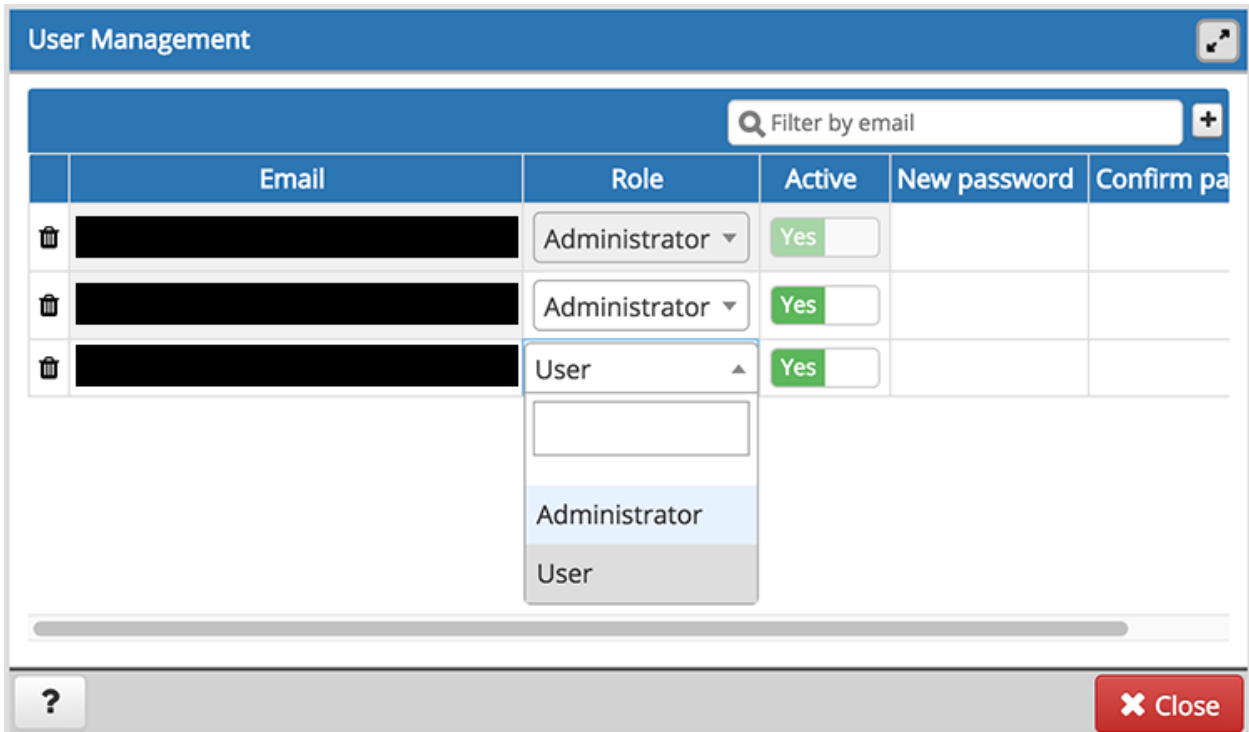
When you authenticate with pgAdmin, the server definitions associated with that login role are made available in the tree control. An administrative user can use the *User Management* dialog to

- add or delete pgAdmin roles
- assign privileges
- manage the password associated with a role



Use the *Filter by email* search field to find a user; enter a user's email address to find a user. If the user exists, the *User Management* table will display the user's current information.

To add a user, click *Add* to add new role.



Provide information about the new pgAdmin role in the row:

- Click in the *Email* field, and provide an email address for the user; this address will be used to recover the

password associated with the role should the password be lost.

- Use the drop-down listbox next to *Role* to select whether a user is an *Administrator* or a *User*.
 - Select *Administrator* if the user will have administrative privileges within the pgAdmin client.
 - Select *User* to create a non-administrative user account.
- Move the *Active* switch to the *No* position if the account is not currently active; the default is *Yes*. Use this switch to disable account activity without deleting an account.
- Use the *New password* field to provide the password associated with the user specified in the *Email* field.
- Re-enter the password in the *Confirm password* field.

To discard a user, and revoke access to pgAdmin, click the trash icon to the left of the row and confirm deletion in the *Delete user?* dialog.

Users with the *Administrator* role are able to add, edit and remove pgAdmin users, but otherwise have the same capabilities as those with the *User* role.

- Click the *Help* button (?) to access online help.
- Click the *Close* button to save work. You will be prompted to return to the dialog if your selections cannot be saved.

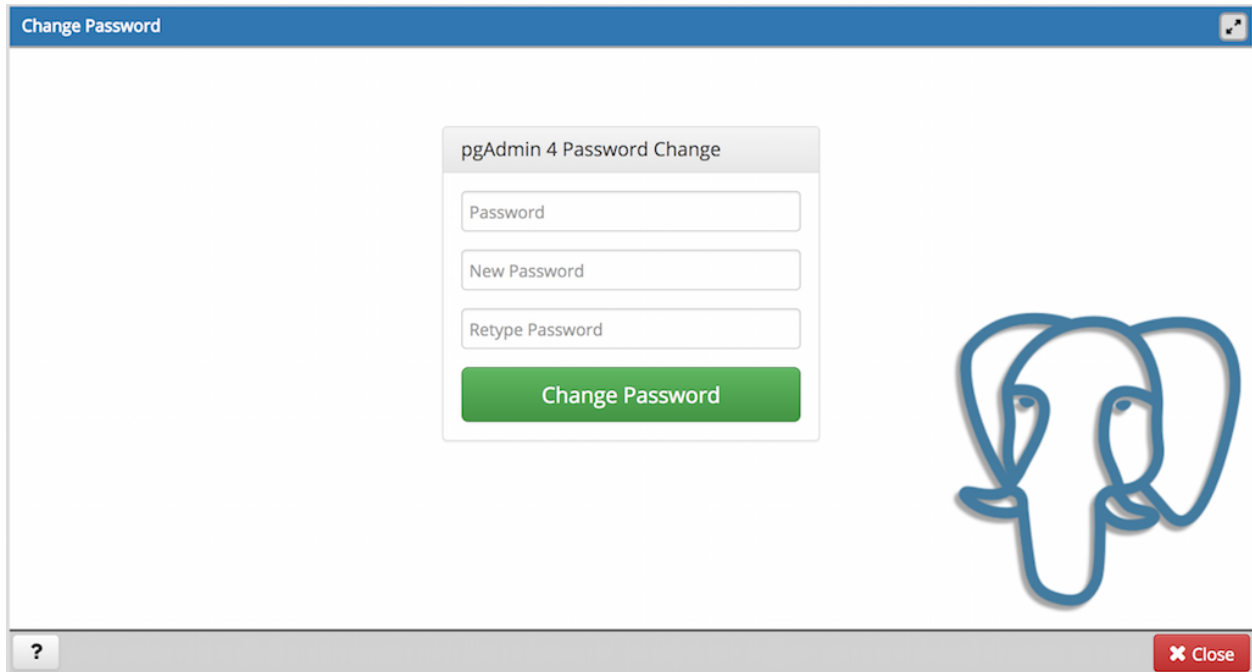
1.4 Change User Password Dialog

It is a good policy to routinely change your password to protect data, even in what you may consider a ‘safe’ environment. In the workplace, failure to apply an appropriate password policy could leave you in breach of Data Protection laws.

Please consider the following guidelines when selecting a password:

- Ensure that your password is an adequate length; 6 characters should be the absolute minimum number of characters in the password.
- Ensure that your password is not open to dictionary attacks. Use a mixture of upper and lower case letters and numerics, and avoid words or names. Consider using the first letter from each word in a phrase that you will remember easily but is an unfamiliar acronym.
- Ensure that your password is changed regularly; at minimum, change it every ninety days.

The above should be considered a starting point: It is not a comprehensive list and it **will not guarantee security**.



Use the *Change Password* dialog to change your password:

- Enter your existing password in the *Current Password* field.
- Enter the desired password for in the *New Password* field.
- Re-enter the new password in the *Confirm Password* field.

Click the *Change Password* button to change your password; click *Close* to exit the dialog.

In a Desktop Deployment, the pgAdmin application is configured to use the desktop runtime environment to host the program on a supported platform. Typically, users will install a pre-built package to run pgAdmin in desktop mode, but a manual desktop deployment can be installed and though it is more difficult to setup, it may be useful for developers interested in understanding how pgAdmin works.

Contents:

1.5 Desktop Deployment

pgAdmin may be deployed as a desktop application by configuring the application to run in desktop mode and then utilising the desktop runtime to host the program on a supported Windows, Mac OS X or Linux installation.

The desktop runtime is a system-tray application that when launched, runs the pgAdmin server and launches a web browser to render the user interface. If additional instances of pgAdmin are launched, a new browser tab will be opened and be served by the existing instance of the server in order to minimise system resource utilisation. Clicking the icon in the system tray will present a menu offering options to open a new pgAdmin window, configure the runtime, view the server log and shut down the server.

Note: Pre-compiled and configured installation packages are available for a number of platforms. These packages should be used by end-users wherever possible - the following information is useful for the maintainers of those packages and users interested in understanding how pgAdmin works.

See also:

For detailed instructions on building and configuring pgAdmin from scratch, please see the README file in the top level directory of the source code. For convenience, you can find the latest version of the file [here](#), but be aware that this may differ from the version included with the source code for a specific version of pgAdmin.

1.5.1 Configuration

From pgAdmin 4 v2 onwards, the default configuration mode is server, however, this is overridden by the desktop runtime at startup. In most environments, no Python configuration is required unless you wish to override other default settings.

There are multiple configuration files that are read at startup by pgAdmin. These are as follows:

- `config.py`: This is the main configuration file, and should not be modified. It can be used as a reference for configuration settings, that may be overridden in one of the following files.
- `config_distro.py`: This file is read after `config.py` and is intended for packagers to change any settings that are required for their pgAdmin distribution. This may typically include certain paths and file locations.
- `config_local.py`: This file is read after `config_distro.py` and is intended for end users to change any default or packaging specific settings that they may wish to adjust to meet local preferences or standards.

Note: If the `SERVER_MODE` setting is changed in `config_distro.py` or `config_local.py`, you will most likely need to re-set the `LOG_FILE`, `SQLITE_PATH`, `SESSION_DB_PATH` and `STORAGE_DIR` values as well as they will have been set based on the default configuration or overridden by the runtime.

Runtime

When executed, the runtime will automatically try to execute the pgAdmin Python application. If execution fails, it will prompt you to adjust the Python Path to include the directories containing the pgAdmin code as well as any additional Python dependencies. You can enter a list of paths by separating them with a semi-colon character, for example:

```
/Users/dpage/.virtualenvs/pgadmin4/lib/python2.7/site-packages//Users/dpage/python-
↳libs/
```

The configuration settings are stored using the `QSettings` class in Qt, which will use an INI file on Unix systems (`~/config/pgadmin/pgadmin4.conf`), a plist file on Mac OS X (`~/Library/Preferences/org.pgadmin.pgadmin4.plist`), and the registry on Windows (`HKEY_CURRENT_USER\Software\pgadmin\pgadmin4`).

The configuration settings:

Key	Type	Purpose
<code>ApplicationPath</code>	String	The directory containing <code>pgAdmin4.py</code>
<code>BrowserCommand</code>	String	An alternate command to run instead of the default browser.
<code>ConnectionTimeout</code>	Integer	The number of seconds to wait for application server startup.
<code>PythonPath</code>	String	The Python module search path

Note: Pre-compiled and configured installation packages are available for a number of platforms. These packages should be used by end-users wherever possible - the following information is useful for the maintainers of those packages and users interested in understanding how pgAdmin works.

The pgAdmin 4 client features a highly-customizable display that features drag-and-drop panels that you can arrange to make the best use of your desktop environment.

The tree control provides an elegant overview of the managed servers, and the objects that reside on each server. Right-click on a node within the tree control to access context-sensitive menus that provide quick access to management tasks for the selected object.

The tabbed browser provide quick access to statistical information about each object in the tree control, and pgAdmin tools and utilities (such as the Query tool and the debugger). pgAdmin opens additional feature tabs each time you access the extended functionality offered by pgAdmin tools; you can open, close, and re-arrange feature tabs as needed.

Use the *Preferences* dialog to customize the content and colors of the pgAdmin display. To open the *Preferences* dialog, select *Preferences* from the *File* menu.

Help buttons in the lower-left corner of each dialog will open the online help for the dialog. You can access additional Postgres help by navigating through the *Help* menu, and selecting the name of the resource that you wish to open.

Contents:

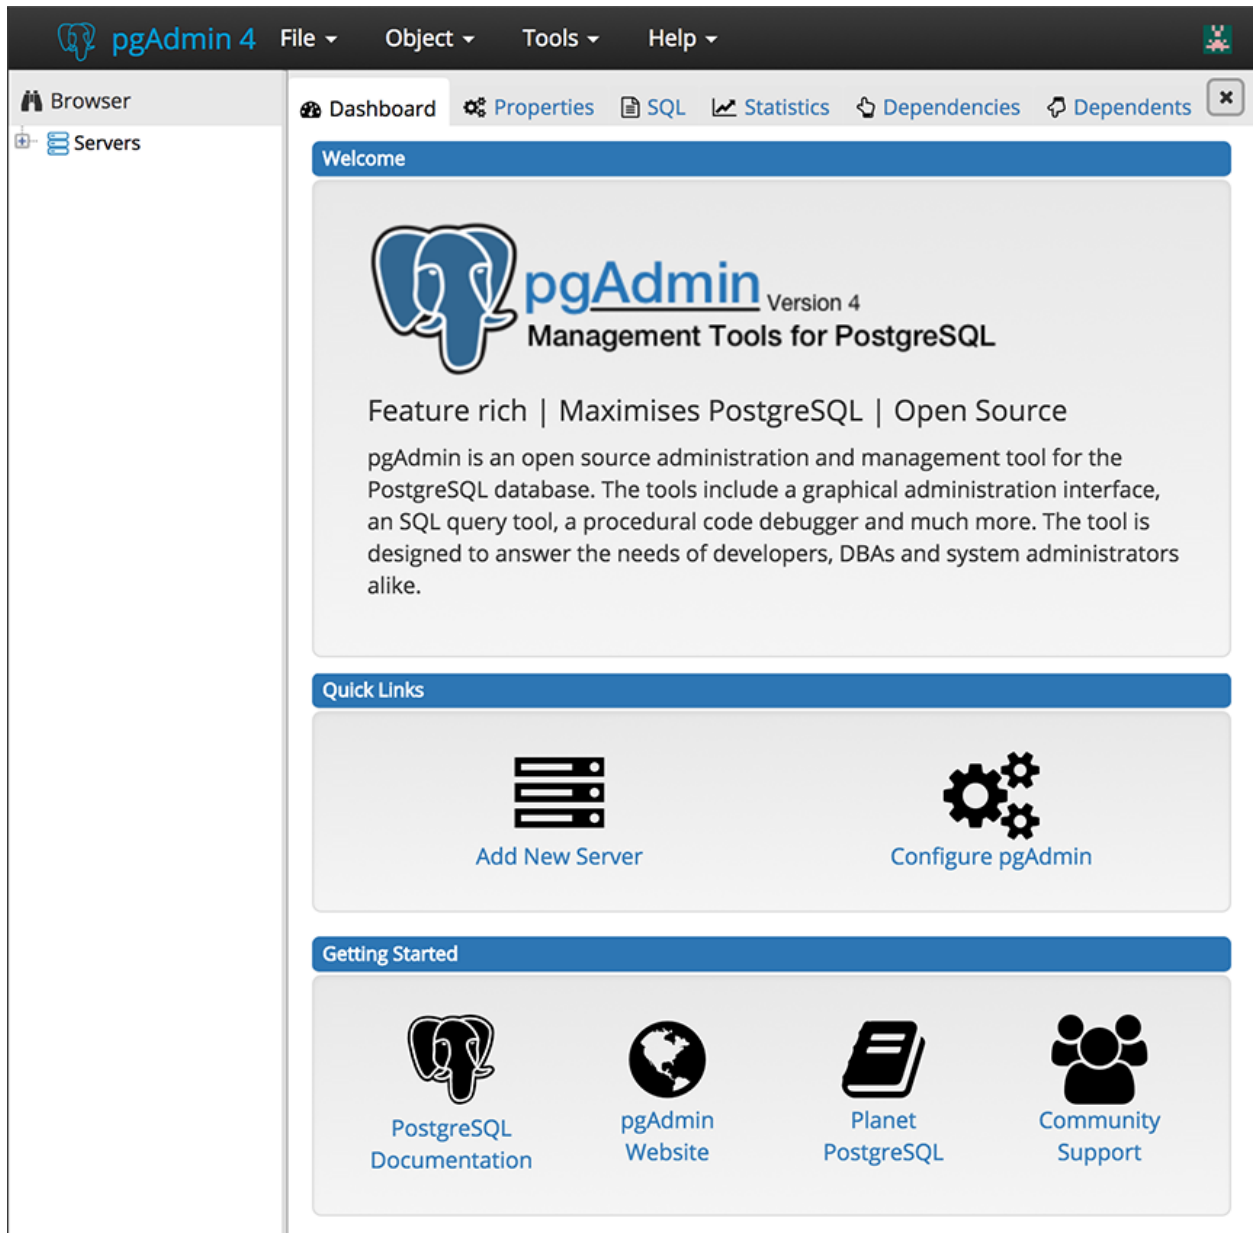
1.6 The pgAdmin 4 Client

pgAdmin 4 supports all PostgreSQL features, from writing simple SQL queries to developing complex databases. It is designed to query an active database (in real-time), allowing you to stay current with modifications and implementations.

Features of pgAdmin 4 include:

- auto-detection and support for objects discovered at run-time
- a live SQL query tool with direct data editing
- support for administrative queries
- a syntax-highlighting SQL editor
- redesigned graphical interfaces
- powerful management dialogs and tools for common tasks
- responsive, context-sensitive behavior
- supportive error messages
- helpful hints
- online help and information about using pgAdmin dialogs and tools.

When pgAdmin opens, the interface features a menu bar and a window divided into two panes: the *Browser* tree control in the left pane, and a tabbed browser in the right pane.



Select an icon from the *Quick Links* panel on the *Dashboard* tab to:

- Click the *Add New Server* button to open the *Create - Server dialog* to add a new server definition.
- Click the *Configure pgAdmin* button to open the *Preferences dialog* to customize your pgAdmin client.

Links in the *Getting Started* panel open a new browser tab that provide useful information for Postgres users:

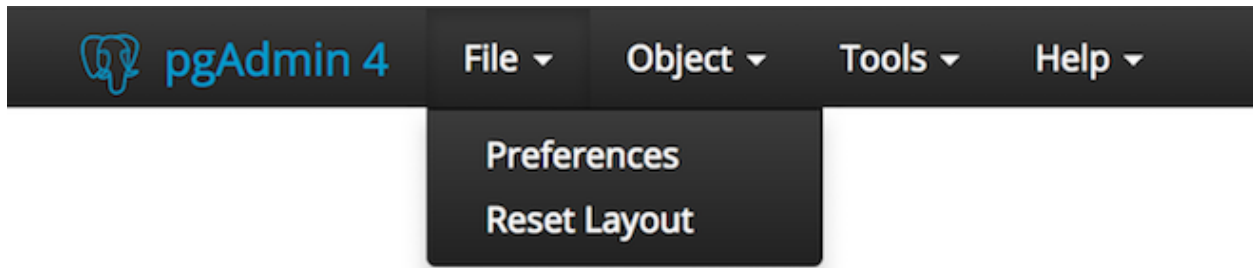
- Click the *PostgreSQL Documentation* link to navigate to the *Documentation* page for the PostgreSQL open-source project; once at the project site, you can review the manuals for the currently supported versions of the PostgreSQL server.
- Click the *pgAdmin Website* link to navigate to the pgAdmin project website. The pgAdmin site features news about recent pgAdmin releases and other project information.
- Click the *Planet PostgreSQL* link to navigate to the blog aggregator for Postgres related blogs.
- Click the *Community Support* link to navigate to the *Community* page at the PostgreSQL open-source project

site; this page provides information about obtaining support for PostgreSQL features.

1.7 pgAdmin Menu Bar

The pgAdmin menu bar provides drop-down menus for access to options, commands, and utilities. The menu bar displays the following selections: *File*, *Object*, *Tools**, and *Help*. Selections may be grayed out which indicates they are disabled for the object currently selected in the *pgAdmin* tree control.

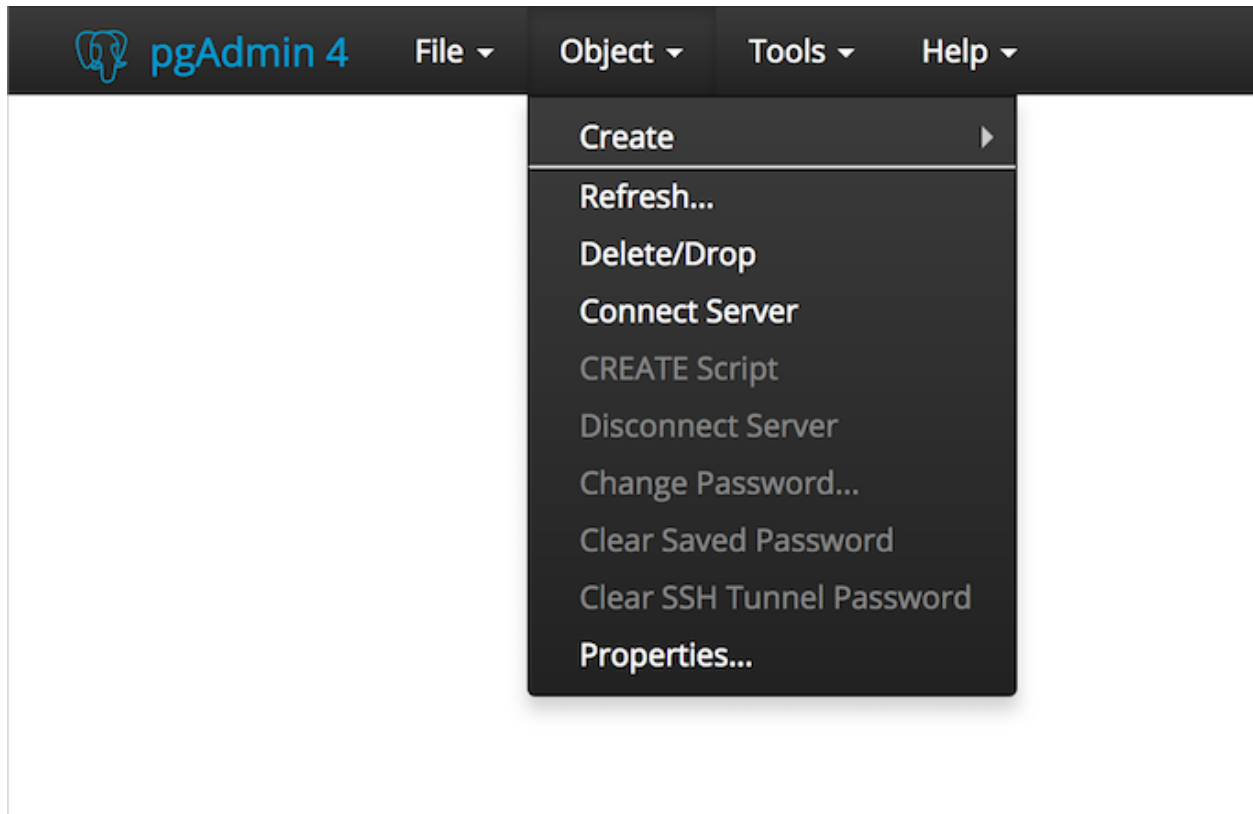
The File Menu



Use the *File* menu to access the following options:

Option	Action
<i>Preferences</i>	Click to open the <i>Preferences</i> dialog to to customize your pgAdmin settings.
<i>Reset Layout</i>	If you have modified the workspace, click to restore the default layout.

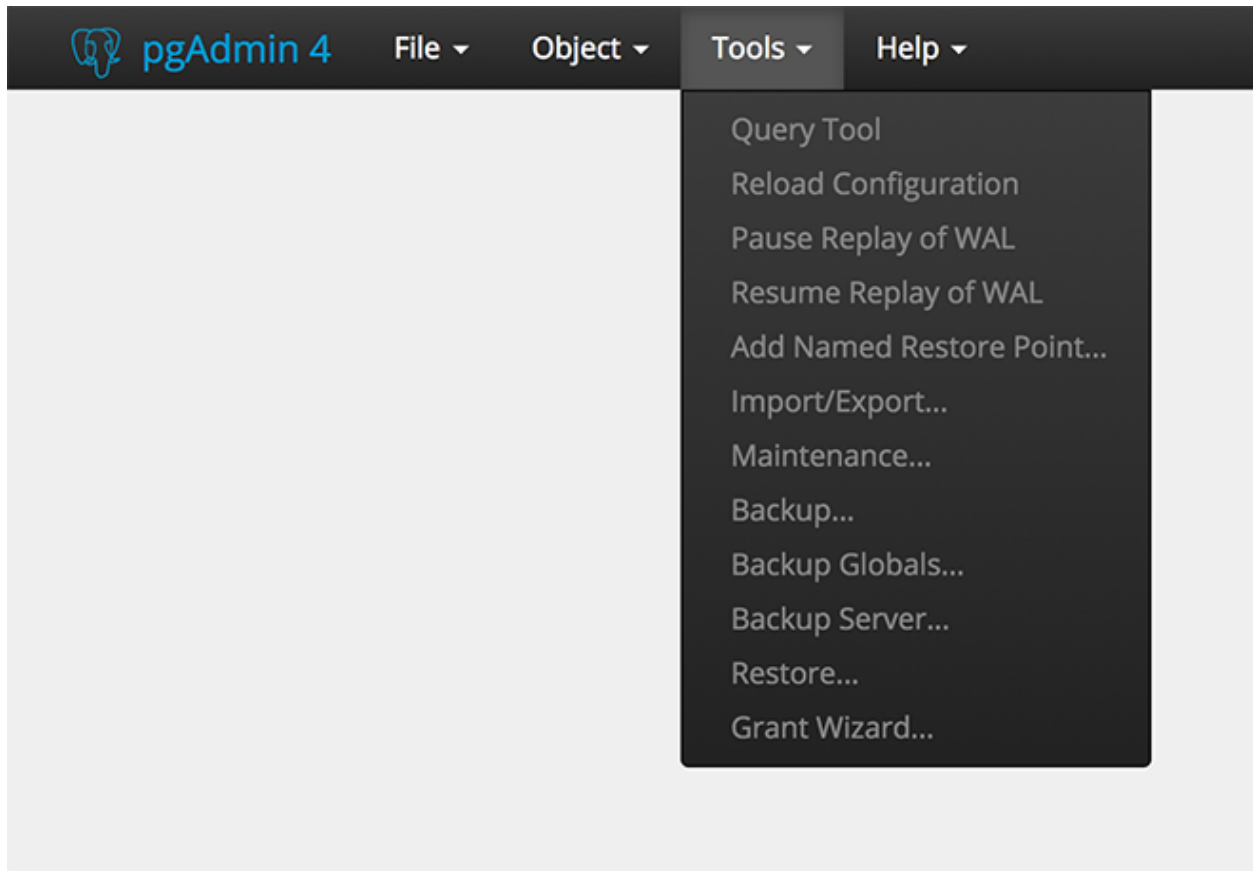
The Object Menu



The *Object* menu is context-sensitive. Use the *Object* menu to access the following options (in alphabetical order):

Option	Action
<i>Change Password...</i>	Click to open the <i>Change Password...</i> dialog to change your password.
<i>Clear Saved Password</i>	If you have saved the database server password, click to clear the saved password. Enable only when password is already saved.
<i>Clear SSH Tunnel Password</i>	If you have saved the ssh tunnel password, click to clear the saved password. Enable only when password is already saved.
<i>Connect Server...</i>	Click to open the <i>Connect to Server</i> dialog to establish a connection with a server.
<i>Create</i>	Click <i>Create</i> to access a context menu that provides context-sensitive selections. Your selection opens a <i>Create</i> dialog for creating a new object.
<i>Delete/Drop</i>	Click to delete the currently selected object from the server.
<i>Disconnect Server...</i>	Click to refresh the currently selected object.
<i>Drop Cascade</i>	Click to delete the currently selected object and all dependent objects from the server.
<i>Properties...</i>	Click to review or modify the currently selected object's properties.
<i>Refresh...</i>	Click to refresh the currently selected object.
<i>Scripts</i>	Click to open the <i>Query tool</i> to edit or view the selected script from the flyout menu.
<i>Trigger(s)</i>	Click to <i>Disable</i> or <i>Enable</i> trigger(s) for the currently selected table. Options are displayed on the flyout menu.
<i>Truncate</i>	Click to remove all rows from a table (<i>Truncate</i>) or to remove all rows from a table and its child tables (<i>Truncate Cascade</i>). Options are displayed on the flyout menu.
<i>View Data</i>	Click to access a context menu that provides several options for viewing data (see below).

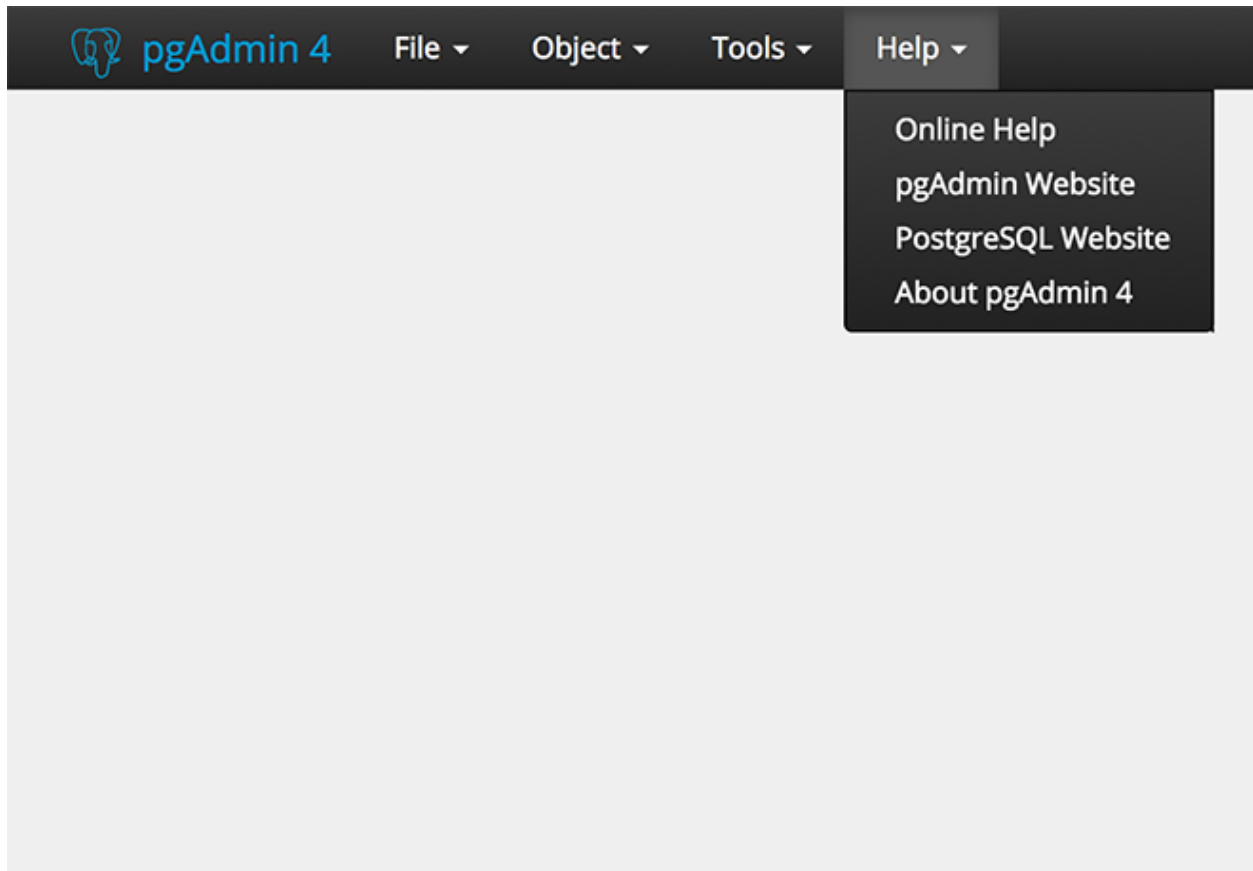
The Tools Menu



Use the *Tools* menu to access the following options (in alphabetical order):

Option	Action
<i>Add named restore point</i>	Click to open the <i>Add named restore point...</i> dialog to take a point-in-time snapshot of the current server state.
<i>Backup...</i>	Click to open the <i>Backup...</i> dialog to backup database objects.
<i>Backup Globals...</i>	Click to open the <i>Backup Globals...</i> dialog to backup cluster objects.
<i>Backup Server...</i>	Click to open the <i>Backup Server...</i> dialog to backup a server.
<i>Grant Wizard...</i>	Click to access the <i>Grant Wizard</i> tool.
<i>Import/Export...</i>	Click to open the <i>Import/Export data...</i> dialog to import or export data from a table.
<i>Maintenance...</i>	Click to open the <i>Maintenance...</i> dialog to VACUUM, ANALYZE, REINDEX, or CLUSTER.
<i>Pause replay of WAL</i>	Click to pause the replay of the WAL log.
<i>Query tool</i>	Click to open the <i>Query tool</i> for the currently selected object.
<i>Reload Configuration...</i>	Click to update configuration files without restarting the server.
<i>Restore...</i>	Click to access the <i>Restore</i> dialog to restore database files from a backup.
<i>Resume replay of WAL</i>	Click to resume the replay of the WAL log.

The Help Menu

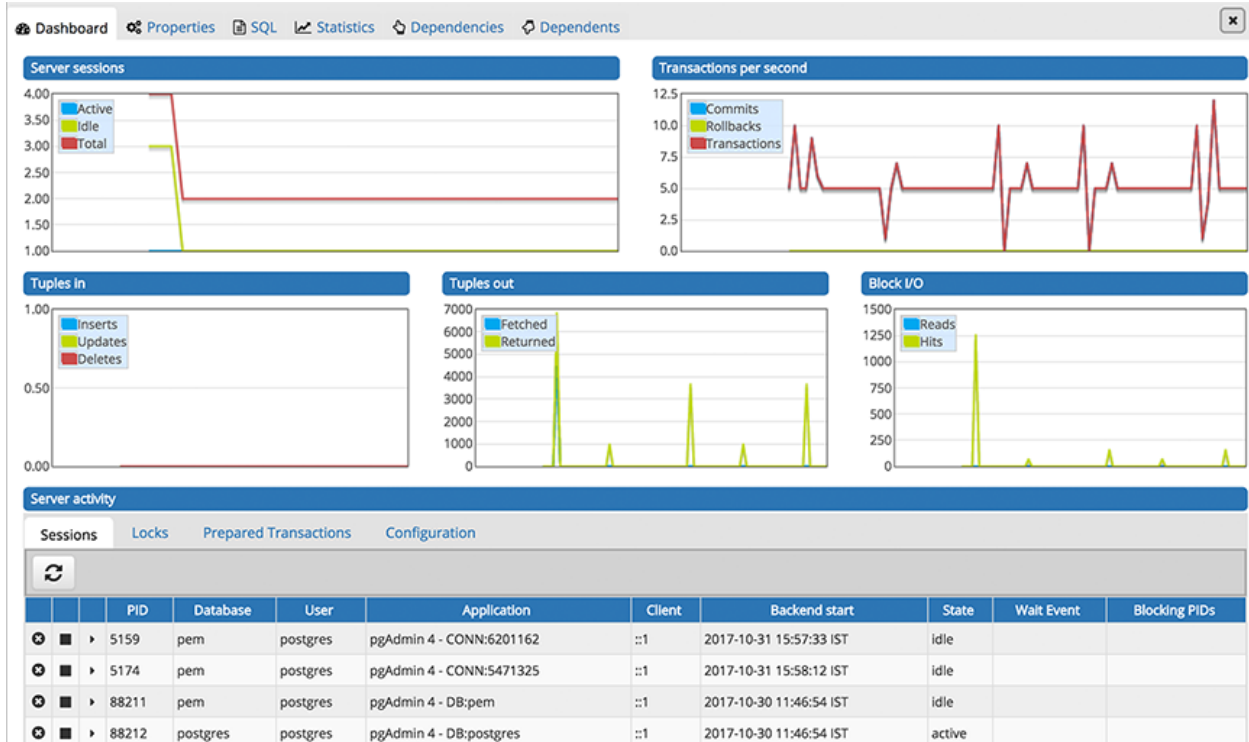


Use the options on the *Help* menu to access online help documents, or to review information about the pgAdmin installation (in alphabetical order):

Option	Action
<i>About pgAdmin 4</i>	Click to open a window where you will find information about pgAdmin; this includes the current version and the current user.
<i>Online Help</i>	Click to open documentation support for using pgAdmin utilities, tools and dialogs. Navigate (in the newly opened tab?) help documents in the left browser pane or use the search bar to specify a topic.
<i>pgAdmin Website</i>	Click to open the <i>pgAdmin.org</i> website in a browser window.
<i>PostgreSQL Website</i>	Click to access the PostgreSQL core documentation hosted at the PostgreSQL site. The site also offers guides, tutorials, and resources.

1.8 pgAdmin Tabbed Browser

The right pane of the *pgAdmin* window features a collection of tabs that display information about the object currently selected in the *pgAdmin* tree control in the left window. Select a tab to access information about the highlighted object in the tree control.



The graphs on the *Dashboard* tab provides an active analysis of the usage statistics for the selected server or database:


- The *Server sessions* or *Database sessions* graph displays the interactions with the server or database.
- The *Transactions per second* graph displays the commits, rollbacks, and total transactions per second that are taking place on the server or database.
- The *Tuples in* graph displays the number of tuples inserted, updated, and deleted on the server or database.
- The *Tuples out* graph displays the number of tuples fetched and returned from the server or database.
- The *Block I/O* graph displays the number of blocks read from the filesystem or fetched from the buffer cache (but not the operating system's file system cache) for the server or database.

The *Server activity* panel displays information about sessions, locks, prepared transactions, and server configuration (if applicable). The information is presented in context-sensitive tables. Use controls located above the table to:

- Click the *Refresh* button to update the information displayed in each table.
- Enter a value in the *Search* box to restrict the table content to one or more sessions that satisfy the search criteria. For example, you can enter a process ID to locate a specific session, or a session state (such as *idle*) to locate all of the sessions that are in an idle state.


You can use icons in the *Sessions* table to review or control the state of a session:

- Use the *Terminate* icon (located in the first column) to stop a session and remove the session from the table. Before the server terminates the session, you will be prompted to confirm your selection.
- Use the *Cancel* icon (located in the second column) to terminate an active query without closing the session. Before canceling the query, the server will prompt you to confirm your selection. When you cancel a query, the value displayed in the *State* column of the table will be updated from *Active* to *Idle*. The session will remain in the table until the session is terminated.
- Use the *Details* icon (located in the third column) to open the *Details* tab; the tab displays information about the selected session.




Database	Owner	Comment
pem	postgres	
postgres	postgres	default administrative connection database
template0	postgres	unmodifiable empty database
template1	postgres	default template for new databases

The *Properties* tab displays information about the object selected. Click the *Edit* icon in the toolbar under the browser tabs to launch the *Properties* dialog for the selected object.



To preserve any changes to the *Properties* dialog, click the *Save* icon; your modifications will be displayed in the updated *Properties* tab.



The Properties dialog shows the following details:

- General**
 - Name: FTS Config
 - OID: 16407
 - Owner: postgres
 - Comment: (empty text area)
- Definition**
 - Parser: default

Details about the object highlighted in the tree control are displayed in one or more collapsible panels. You can use the arrow to the left of each panel label to open or close a panel.

```

Dashboard Properties SQL Statistics Dependencies Dependents Query - pem on ...
1  -- Column: pg_catalog.pg_shdescription.description
2
3  -- ALTER TABLE pg_catalog.pg_shdescription DROP COLUMN description;
4
5  ALTER TABLE pg_catalog.pg_shdescription
6      ADD COLUMN description text COLLATE pg_catalog."default" NOT NULL;
    
```

The *SQL* tab displays the SQL script that created the highlighted object, and when applicable, a (commented out) SQL statement that will *DROP* the selected object. You can copy the SQL statements to the editor of your choice using cut and paste shortcuts.

Table name	Tuples inserted	Tuples updated	Tuples deleted	Tuples HOT updated	Live tuples	Dead tuples	Last
audit_configuration	0	0	0	0	0	0	0 2017
audit_logs	0	0	0	0	0	0	0 2017
auto_discover_servers	0	0	0	0	0	0	0 2017
background_writer_statistics	0	0	0	0	0	0	0 2017
cpu_usage	0	0	0	0	0	0	0 2017
data_log_file_analysis	0	0	0	0	0	0	0 2017
database_frozenxid	0	0	0	0	0	0	0 2017
database_size	0	0	0	0	0	0	0 2017
database_statistics	0	0	0	0	0	0	0 2017
disk_busy_info	0	0	0	0	0	0	0 2017
disk_space	0	0	0	0	0	0	0 2017
efm_cluster_info	0	0	0	0	0	0	0 2017
efm_cluster_node_status	0	0	0	0	0	0	0 2017
function_statistics	0	0	0	0	0	0	0 2017
index_size	0	0	0	0	0	0	0 2017
index_statistics	0	0	0	0	0	0	0 2017
installed_packages	0	0	0	0	0	0	0 2017
io_analysis	0	0	0	0	0	0	0 2017
load_average	0	0	0	0	0	0	0 2017
lock_info	0	0	0	0	0	0	0 2017
log_configuration	0	0	0	0	0	0	0 2017
memory_usage	0	0	0	0	0	0	0 2017
mview_bloat	0	0	0	0	0	0	0 2017

The *Statistics* tab displays the statistics gathered for each object on the tree control; the statistics displayed in the table vary by the type of object that is selected. Click a column heading to sort the table by the data displayed in the column; click again to reverse the sort order. The following table lists some of the statistics that are available:

Panel	Description
<i>PID</i>	The process ID associated with the row.
<i>User</i>	The name of the user that owns the object.
<i>Database</i>	displays the database name.
<i>Backends</i>	displays the number of current connections to the database.
<i>Backend start</i>	The start time of the backend process.
<i>Xact Committed</i>	displays the number of transactions committed to the database within the last week.
<i>Xact Rolled Back</i>	displays the number of transactions rolled back within the last week.
<i>Blocks Read</i>	displays the number of blocks read from memory (in megabytes) within the last week.
<i>Blocks Hit</i>	displays the number of blocks hit in the cache (in megabytes) within the last week.
<i>Tuples Returned</i>	displays the number of tuples returned within the last week.
<i>Tuples Fetched</i>	displays the number of tuples fetched within the last week.
<i>Tuples Inserted</i>	displays the number of tuples inserted into the database within the last week.
<i>Tuples Updated</i>	displays the number of tuples updated in the database within the last week.
<i>Tuples Deleted</i>	displays the number of tuples deleted from the database within the last week.
<i>Last statistics reset</i>	displays the time of the last statistics reset for the database.
<i>Tablespace conflicts</i>	displays the number of queries canceled because of recovery conflict with dropped tablespaces in database.
<i>Lock conflicts</i>	displays the number of queries canceled because of recovery conflict with locks in database.
<i>Snapshot conflicts</i>	displays the number of queries canceled because of recovery conflict with old snapshots in database.
<i>Bufferpin conflicts</i>	displays the number of queries canceled because of recovery conflict with pinned buffers in database.
<i>Temporary files</i>	displays the total number of temporary files, including those used by the statistics collector.
<i>Size of temporary files</i>	displays the size of the temporary files.
<i>Deadlocks</i>	displays the number of queries canceled because of a recovery conflict with deadlocks in database.
<i>Block read time</i>	displays the number of milliseconds required to read the blocks read.
<i>Block write time</i>	displays the number of milliseconds required to write the blocks read.
<i>Size</i>	displays the size (in megabytes) of the selected database.

Dashboard Properties SQL Statistics Dependencies Dependents		
Type	Name	Restriction
Domain	information_schema.character_data	normal
Domain	information_schema.sql_identifier	normal
Schema	information_schema	normal

The *Dependencies* tab displays the objects on which the currently selected object depends. If a dependency is dropped, the object currently selected in the pgAdmin tree control will be affected. To ensure the integrity of the entire database structure, the database server makes sure that you do not accidentally drop objects that other objects depend on; you must use the DROP CASCADE command to remove an object with a dependency.

The *Dependencies* table displays the following information:

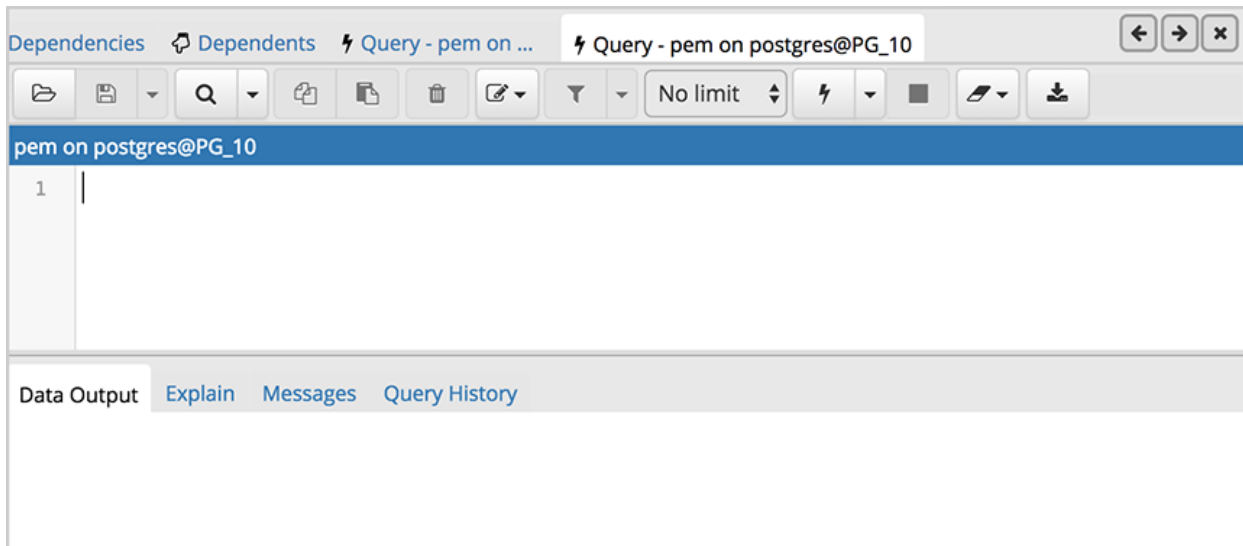
- The *Type* field specifies the parent object type.
- The *Name* field specifies the identifying name of the parent object.
- **The *Restriction* field describes the dependency relationship between the currently selected object and the parent.**

- If the field is *auto*, the selected object can be dropped separately from the parent object, and will be dropped if the parent object is dropped.
- If the field is *internal*, the selected object was created during the creation of the parent object, and will be dropped if the parent object is dropped.
- If the field is *normal*, the selected object can be dropped without dropping the parent object.
- If the field is *blank*, the selected object is required by the system, and cannot be dropped.

Type	Name	Restriction
Type	information_schema.pg_foreign_data_wrappers	internal
Rule	_RETURN ON information_schema.pg_foreign_data_wrappers	internal
Rule	_RETURN ON information_schema.pg_foreign_data_wrappers	normal
Rule	_RETURN ON information_schema.foreign_data_wrapper_options	normal
Rule	_RETURN ON information_schema.foreign_data_wrappers	normal

The *Dependents* tab displays a table of objects that depend on the object currently selected in the *pgAdmin* browser. A dependent object can be dropped without affecting the object currently selected in the *pgAdmin* tree control.

- The *Type* field specifies the dependent object type.
- The *Name* field specifies the identifying name for the dependent object.
- The *Database* field specifies the database in which the object resides.

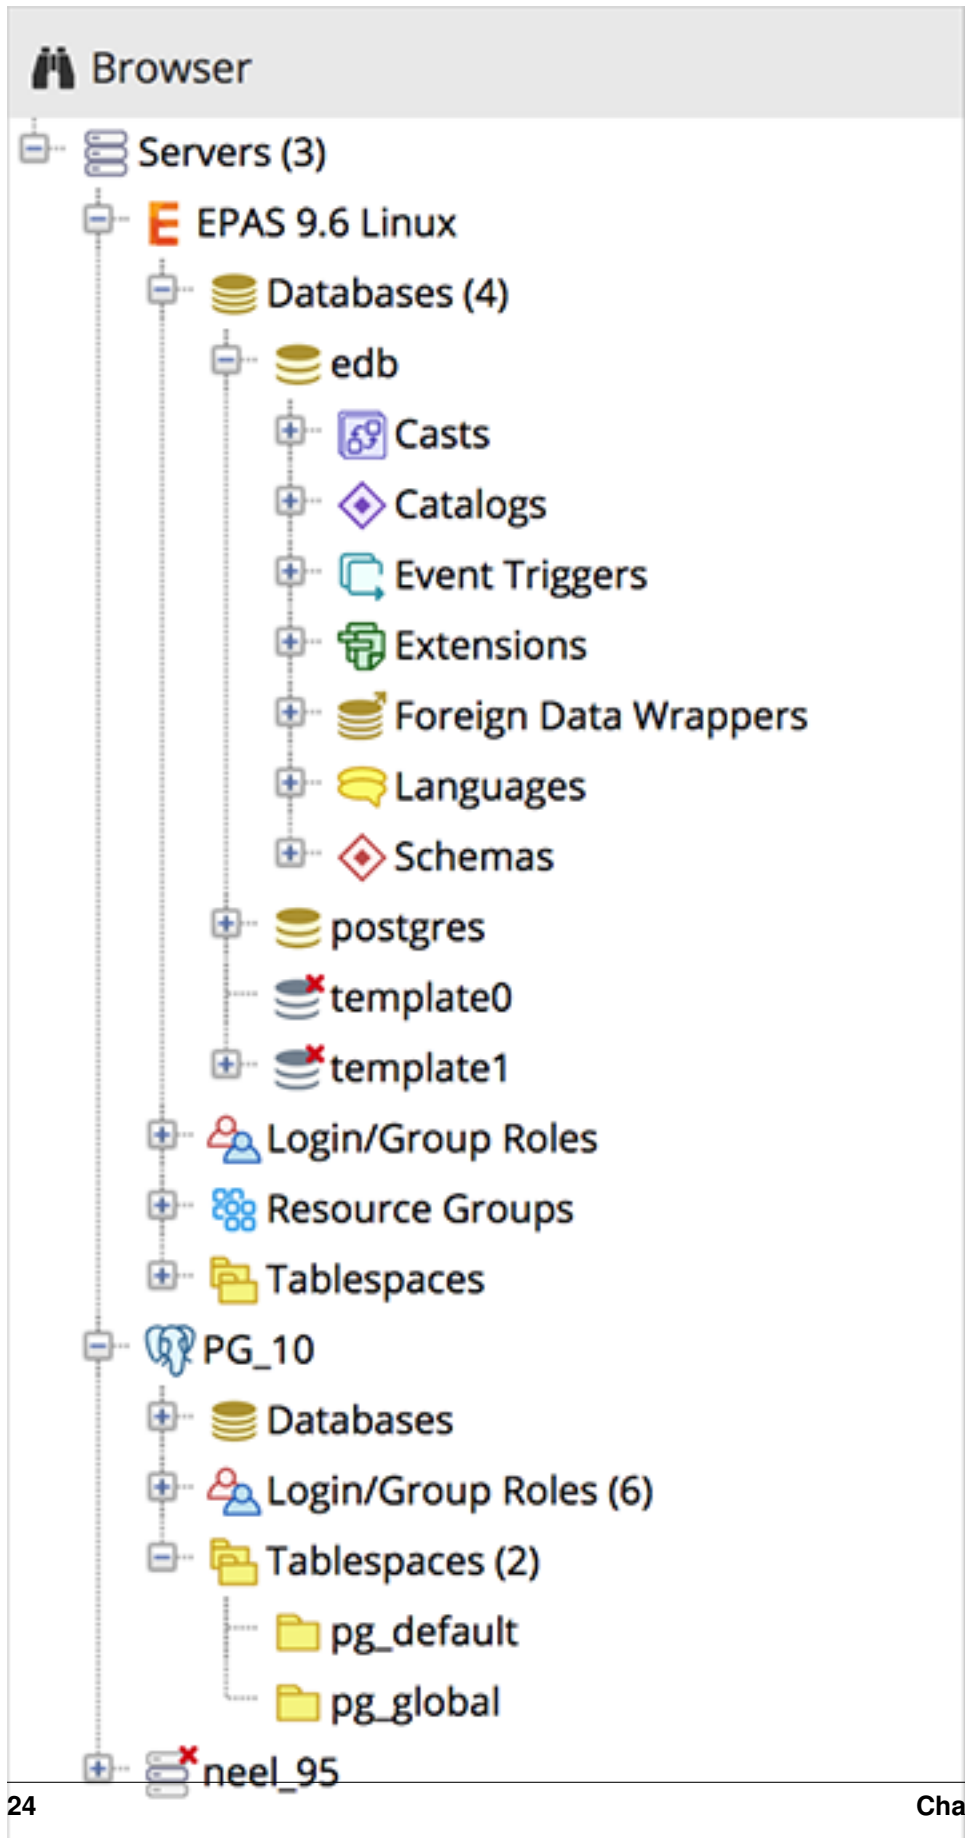


Additional tabs open when you access the extended functionality offered by *pgAdmin* tools (such as the Query tool, Debugger, or SQL editor). Use the close icon (X) located in the upper-right corner of each tab to close the tab when you are finished using the tool. Like permanent tabs, these tabs may be repositioned in the *pgAdmin* client window.

By default, each time you open a tool, *pgAdmin* will open a new browser tab. You can control this behavior by modifying the *Display* node of the *Preferences* dialog for each tool. To open the *Preferences* dialog, select *Preferences* from the *File* menu.

1.9 pgAdmin Tree Control

The left pane of the main window displays a tree control (the *pgAdmin* tree control) that provides access to the objects that reside on a server.



You can expand nodes in the tree control to view the database objects that reside on a selected server. The tree control expands to display a hierarchical view:

- Use the plus sign (+) to the left of a node to expand a segment of the tree control.
- Click the minus sign (-) to the left of a node to close that node.

Access context-sensitive menus by right-clicking on a node of the tree control to perform common tasks. Menus display options that include one or more of the following selections (options appear in alphabetical order):

Option	Action
<i>Add named restore point</i>	Click to create and enter the name of a restore point.
<i>Backup...</i>	Click to open the <i>Backup...</i> dialog to backup database objects.
<i>Backup Globals...</i>	Click to open the <i>Backup Globals...</i> dialog to backup cluster objects.
<i>Backup Server...</i>	Click to open the <i>Backup Server...</i> dialog to backup a server.
<i>Connect Server...</i>	Click to open the <i>Connect to Server</i> dialog to establish a connection with a server.
<i>Create</i>	Click to access a context menu that provides context-sensitive selections. Your selection opens a <i>Create</i> dialog for creating a new object.
<i>CREATE Script</i>	Click to open the <i>Query tool</i> to edit or view the CREATE script.
<i>Debugging</i>	Click through to open the <i>Debug</i> tool or to select <i>Set breakpoint</i> to stop or pause a script execution.
<i>Delete/Drop</i>	Click to delete the currently selected object from the server.
<i>Disconnect Database...</i>	Click to terminate a database connection.
<i>Disconnect Server...</i>	Click to refresh the currently selected object.
<i>Drop Cascade</i>	Click to delete the currently selected object and all dependent objects from the server.
<i>Debugging</i>	Click to access the <i>Debugger</i> tool.
<i>Grant Wizard</i>	Click to access the <i>Grant Wizard</i> tool.
<i>Maintenance...</i>	Click to open the <i>Maintenance...</i> dialog to VACUUM, ANALYZE, REINDEX, or CLUSTER.
<i>Properties...</i>	Click to review or modify the currently selected object's properties.
<i>Refresh...</i>	Click to refresh the currently selected object.
<i>Reload Configuration...</i>	Click to update configuration files without restarting the server.
<i>Restore...</i>	Click to access the <i>Restore</i> dialog to restore database files from a backup.
<i>View Data</i>	Use the <i>View Data</i> option to access the data stored in a selected table with the <i>Data Output</i> tab of the <i>Query Tool</i> .

The context-sensitive menus associated with *Tables* and nested *Table* nodes provides additional display options (options appear in alphabetical order):

Option	Action
<i>Import/Export...</i>	Click open the <i>Import/Export...</i> dialog to import data to or export data from the selected table.
<i>Reset Statistics</i>	Click to reset statistics for the selected table.
<i>Scripts</i>	Click to open the <i>Query tool</i> to edit or view the selected script from the flyout menu.
<i>Truncate</i>	Click to remove all rows from a table.
<i>Truncate Cascade</i>	Click to remove all rows from a table and its child tables.
<i>View First 100 Rows</i>	Click to access a data grid that displays the first 100 rows of the selected table.
<i>View Last 100 Rows</i>	Click to access a data grid that displays the last 100 rows of the selected table.
<i>View All Rows</i>	Click to access a a data grid that displays all rows of the selected table.
<i>View Filtered Rows...</i>	Click to access the <i>Data Filter</i> popup to apply a filter to a set of data.

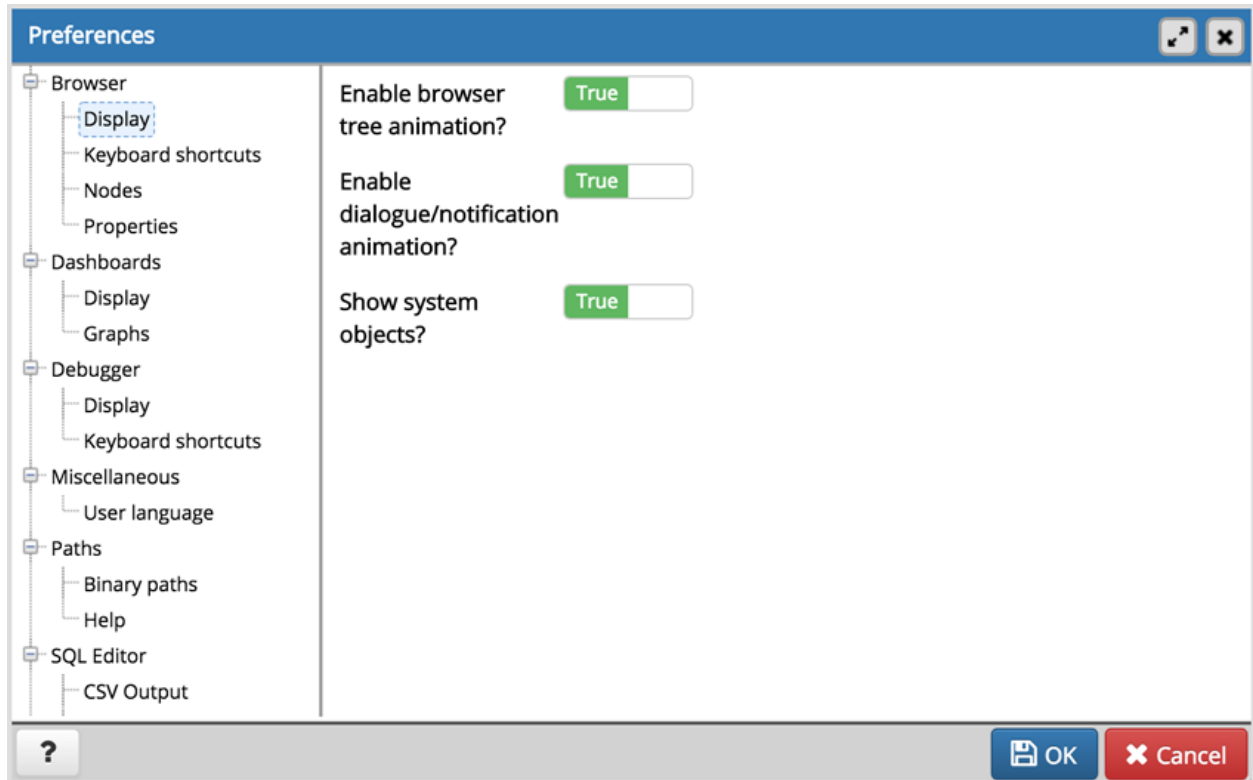
1.10 Preferences Dialog

Use options on the *Preferences* dialog to customize the behavior of the client. To open the *Preferences* dialog, select *Preferences* from the *File* menu. The left pane of the *Preferences* dialog displays a tree control; each node of the tree control provides access to options that are related to the node under which they are displayed.

- Use the plus sign (+) to the left of a node name to expand a segment of the tree control.
- Use the minus sign (-) to the left of a node name to close that node.

The Browser Node

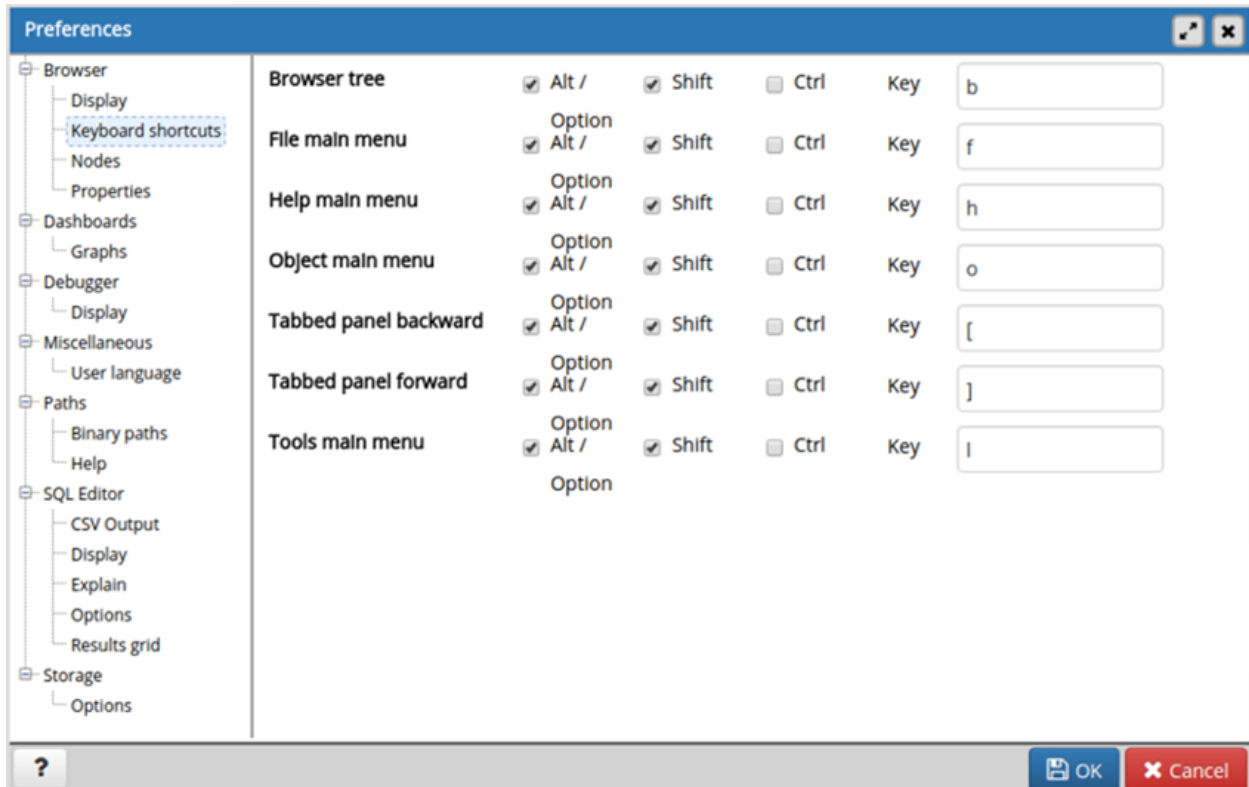
Use preferences found in the *Browser* node of the tree control to personalize your workspace.



Use the fields on the *Display* panel to specify general display preferences:

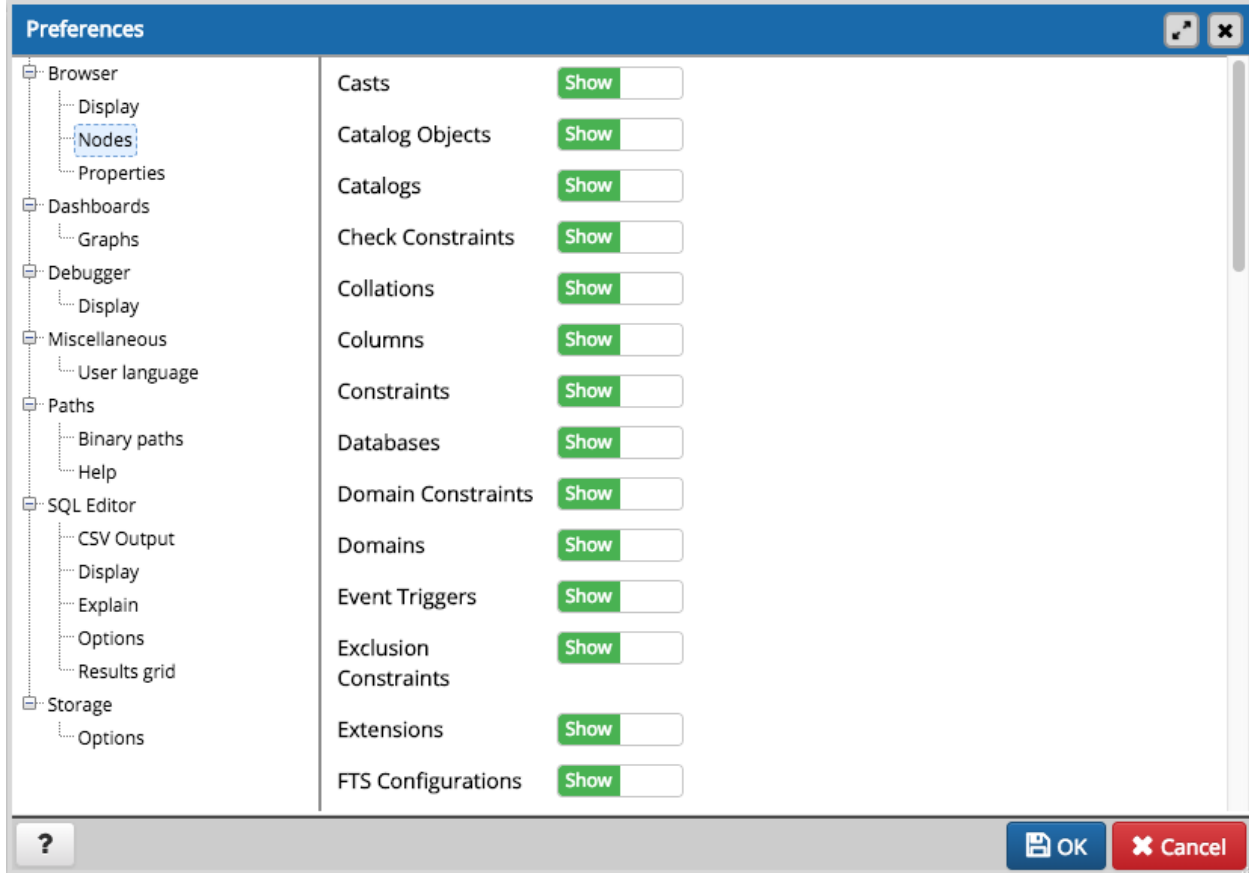
- When the *Show system objects?* switch is set to *True*, the client will display system objects such as system schemas (for example, *pg_temp*) or system columns (for example, *xmin* or *ctid*) in the tree control.
- When the *Enable browser tree animation?* switch is set to *True*, the client will display the animated tree control otherwise it will be unanimated.
- When the *Enable dialogue/notification animation?* switch is set to *True*, the client will display the animated dialogues/notifications otherwise it will be unanimated.

Use the fields on the *Keyboard shortcuts* panel to configure shortcuts for the main window navigation:



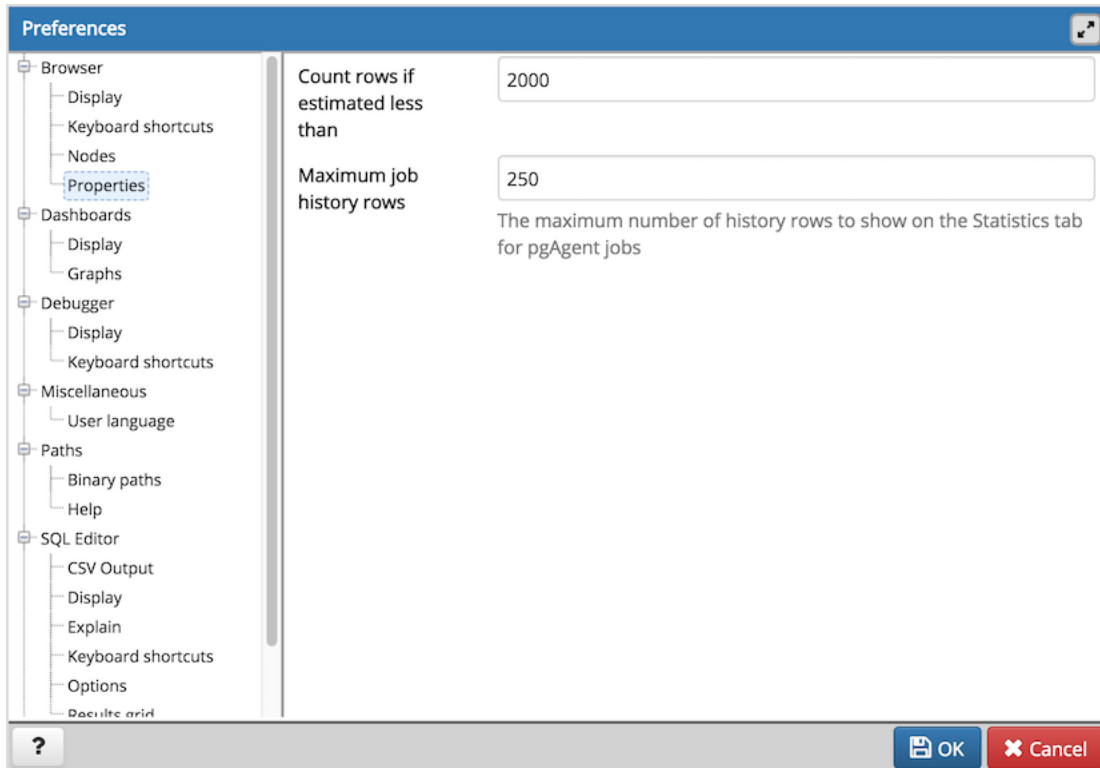
- The panel displays a list of keyboard shortcuts available for the main window; select the combination of the modifier keys along with the key to configure each shortcut.

Use the fields on the *Nodes* panel to select the object types that will be displayed in the *Browser* tree control:



- The panel displays a list of database objects; slide the switch located next to each object to *Show* or *Hide* the database object. When querying system catalogs, you can reduce the number of object types displayed to increase speed.

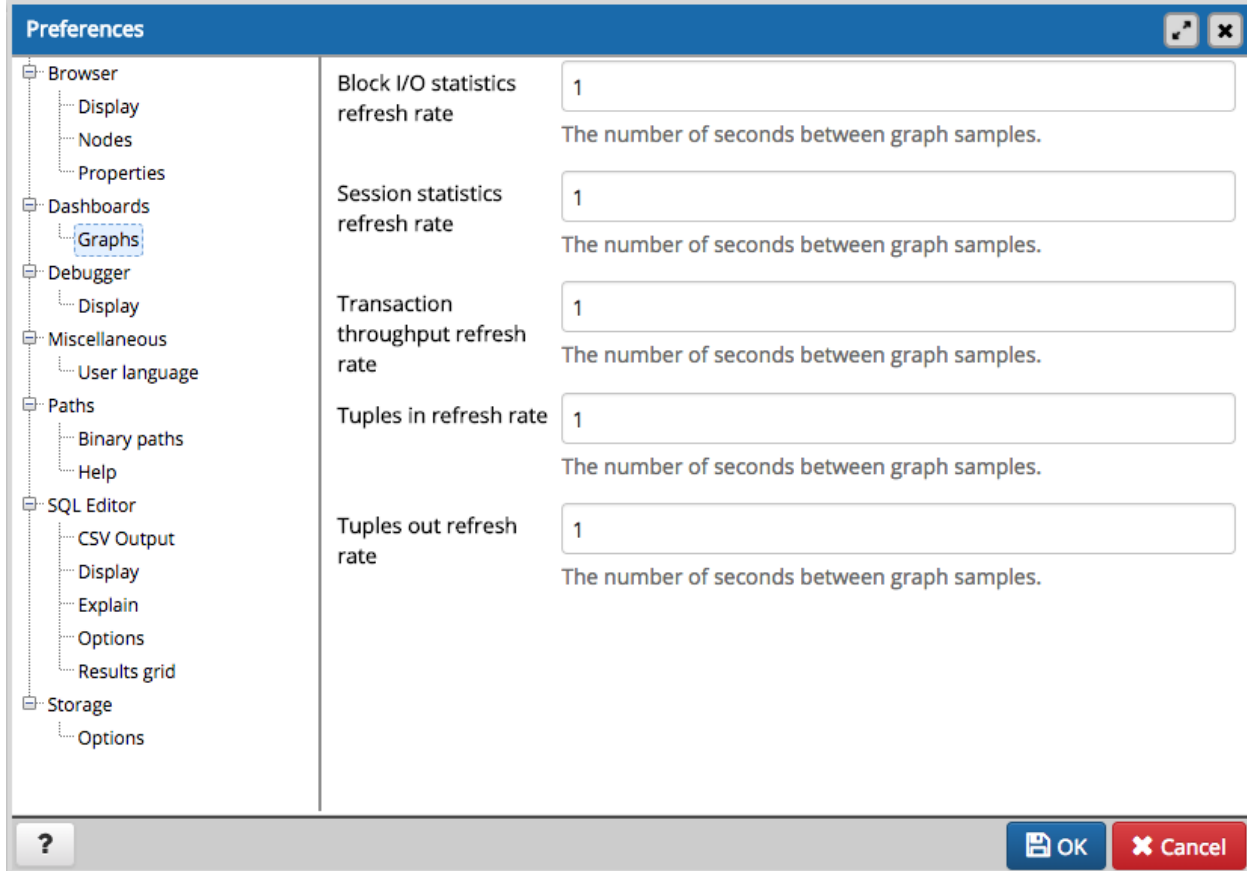
Use fields on the *Properties* panel to specify browser properties:



- Include a value in the *Count rows if estimated less than* field to perform a `SELECT count(*)` if the estimated number of rows in a table (as read from the table statistics) is below the specified limit. After performing the `SELECT count(*)`, pgAdmin will display the row count. The default is 2000.
- Provide a value in the *Maximum job history rows* field to limit the number of rows to show on the statistics tab for pgAgent jobs. The default is 250.

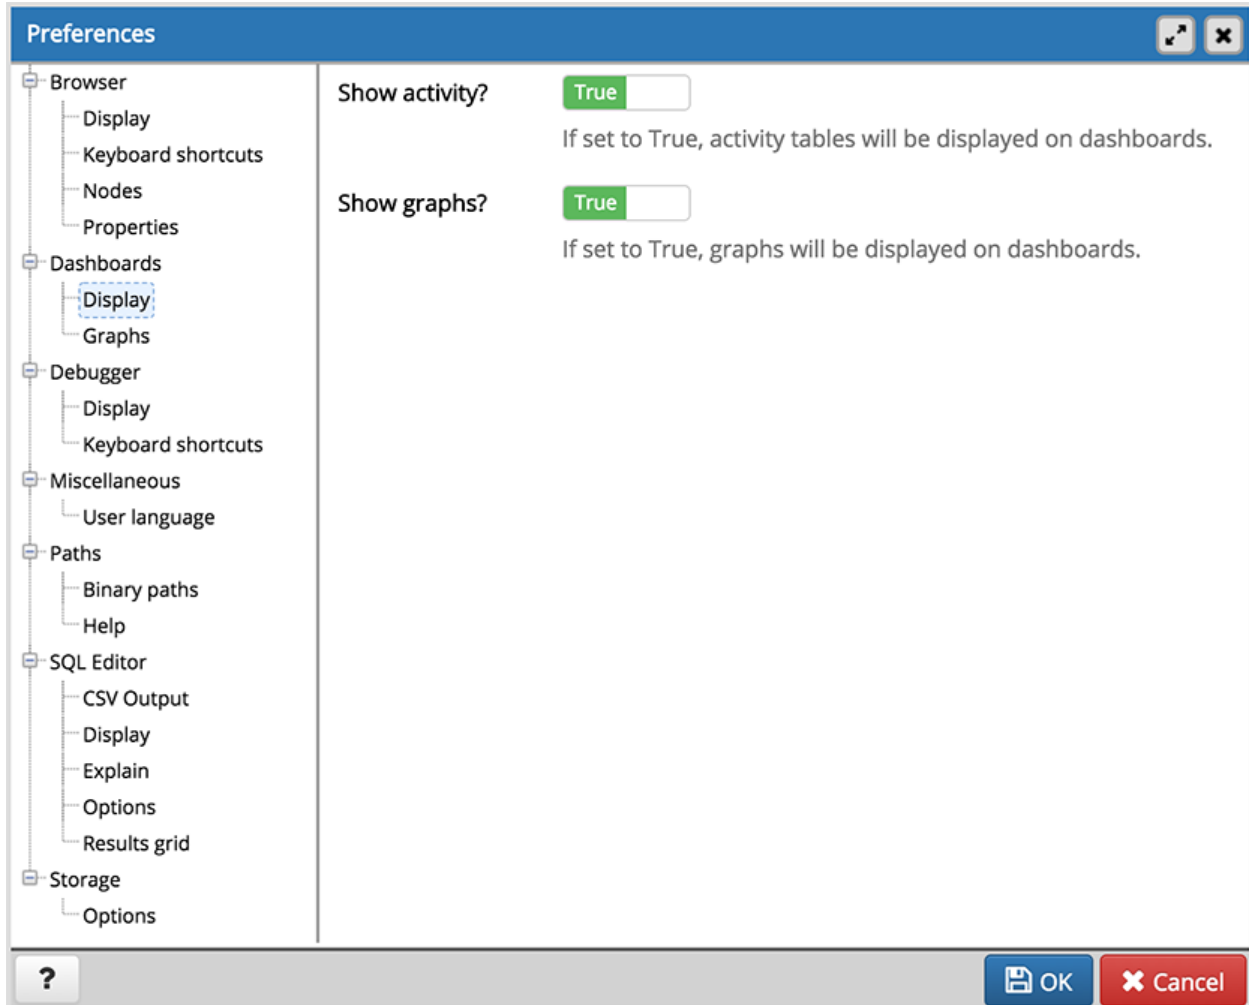
The Dashboards Node

Expand the *Dashboards* node to specify your dashboard display preferences.



Use the fields on the *Graphs* panel to specify your display preferences for the graphs on the *Dashboard* tab:

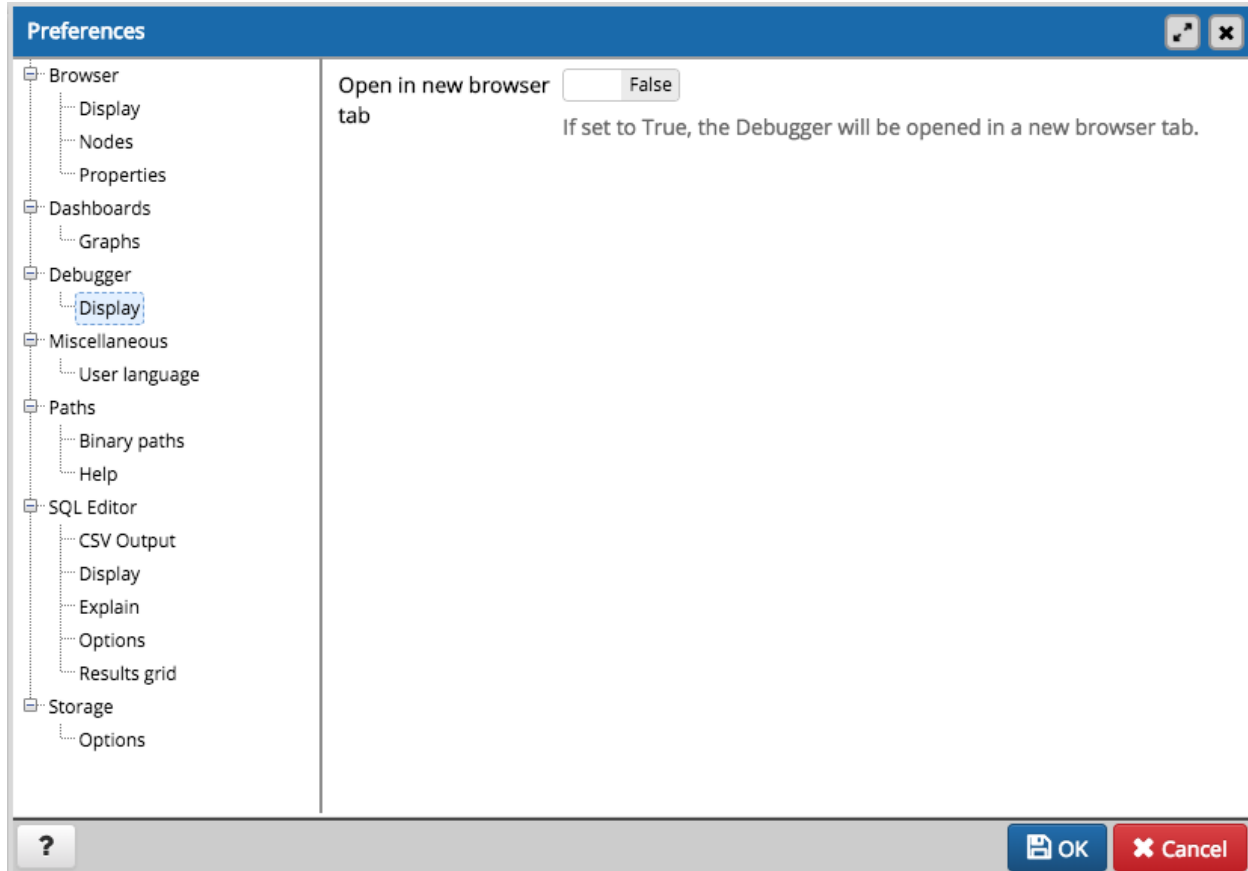
- Use the *Block I/O statistics refresh rate* field to specify the number of seconds between block I/O statistic samples displayed in graphs.
- Use the *Session statistics refresh rate* field to specify the number of seconds between session statistic samples displayed in graphs.
- Use the *Transaction throughput refresh rate* field to specify the number of seconds between transaction throughput samples displayed in graphs.
- Use the *Tuples in refresh rate* field to specify the number of seconds between tuples-in samples displayed in graphs.
- Use the *Tuples out refresh rate* field to specify the number of seconds between tuples-out samples displayed in graphs.



- When the *Show graphs?* switch is set to *True*, graphs will be displayed on dashboards.
- When the *Show activity?* switch is set to *True*, activity tables will be displayed on dashboards.

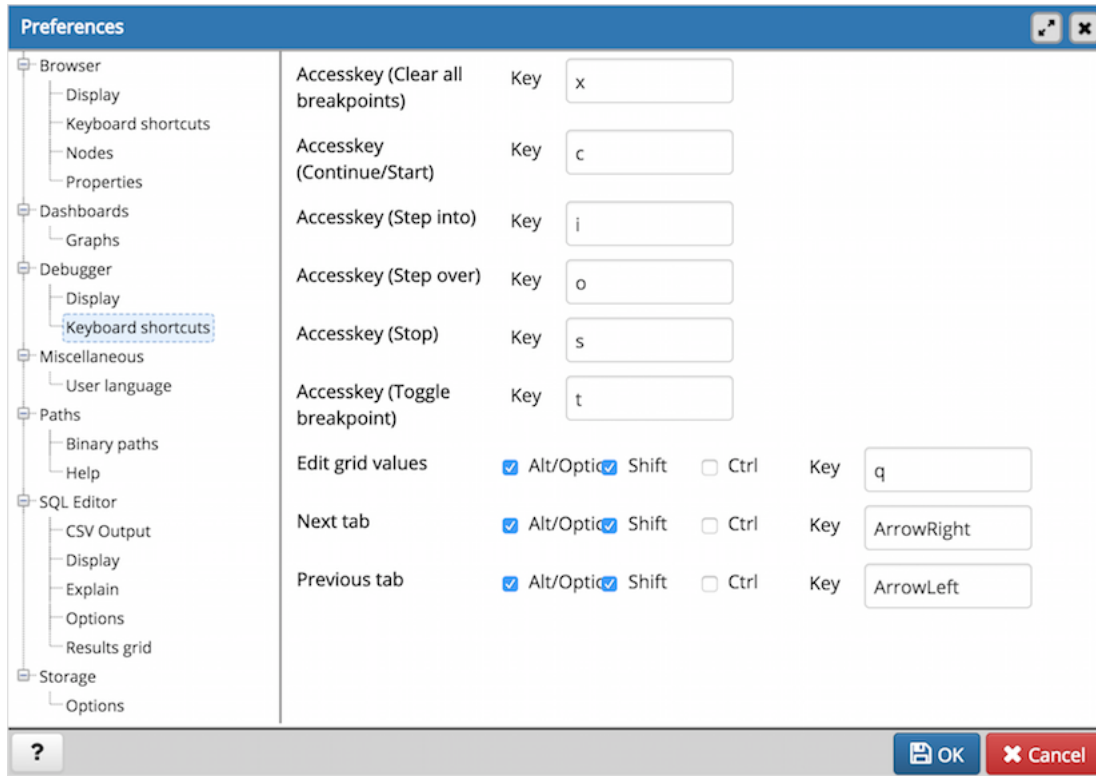
The Debugger Node

Expand the *Debugger* node to specify your debugger display preferences.



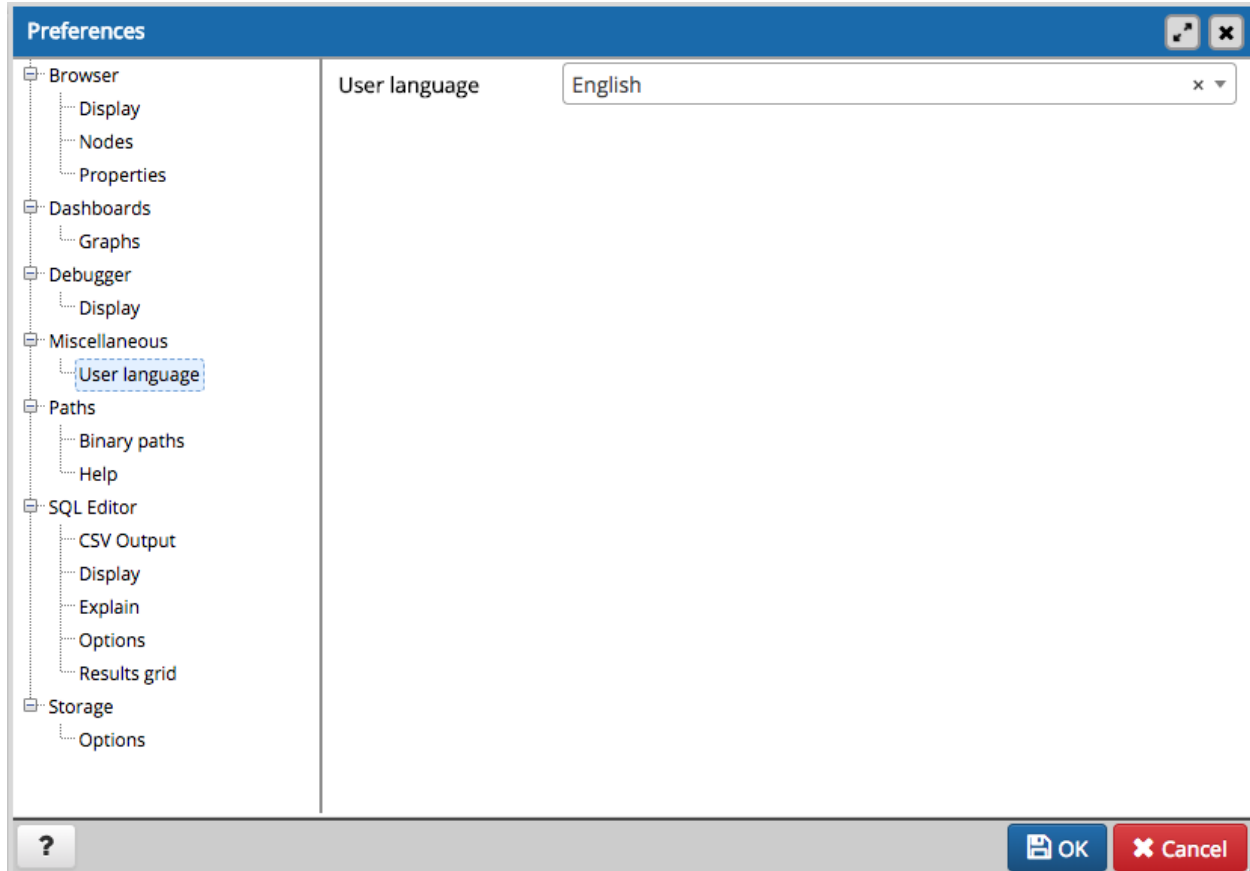
- When the *Open in new browser tab* switch is set to *True*, the Debugger will open in a new browser tab when invoked.

Use the fields on the *Keyboard shortcuts* panel to configure shortcuts for the debugger window navigation:



The Miscellaneous Node

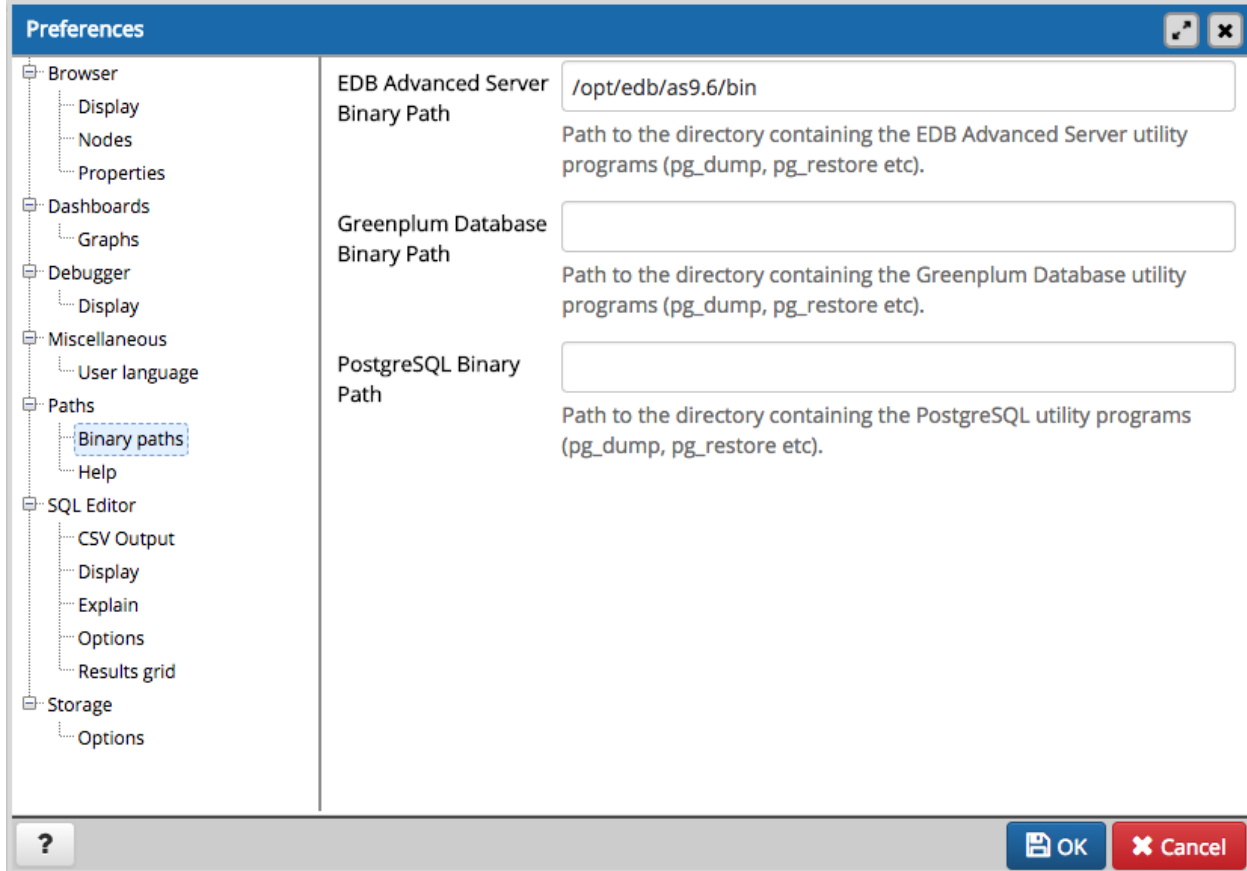
Expand the *Miscellaneous* node to specify miscellaneous display preferences.



- Use the *User language* drop-down listbox to select the display language for the client.

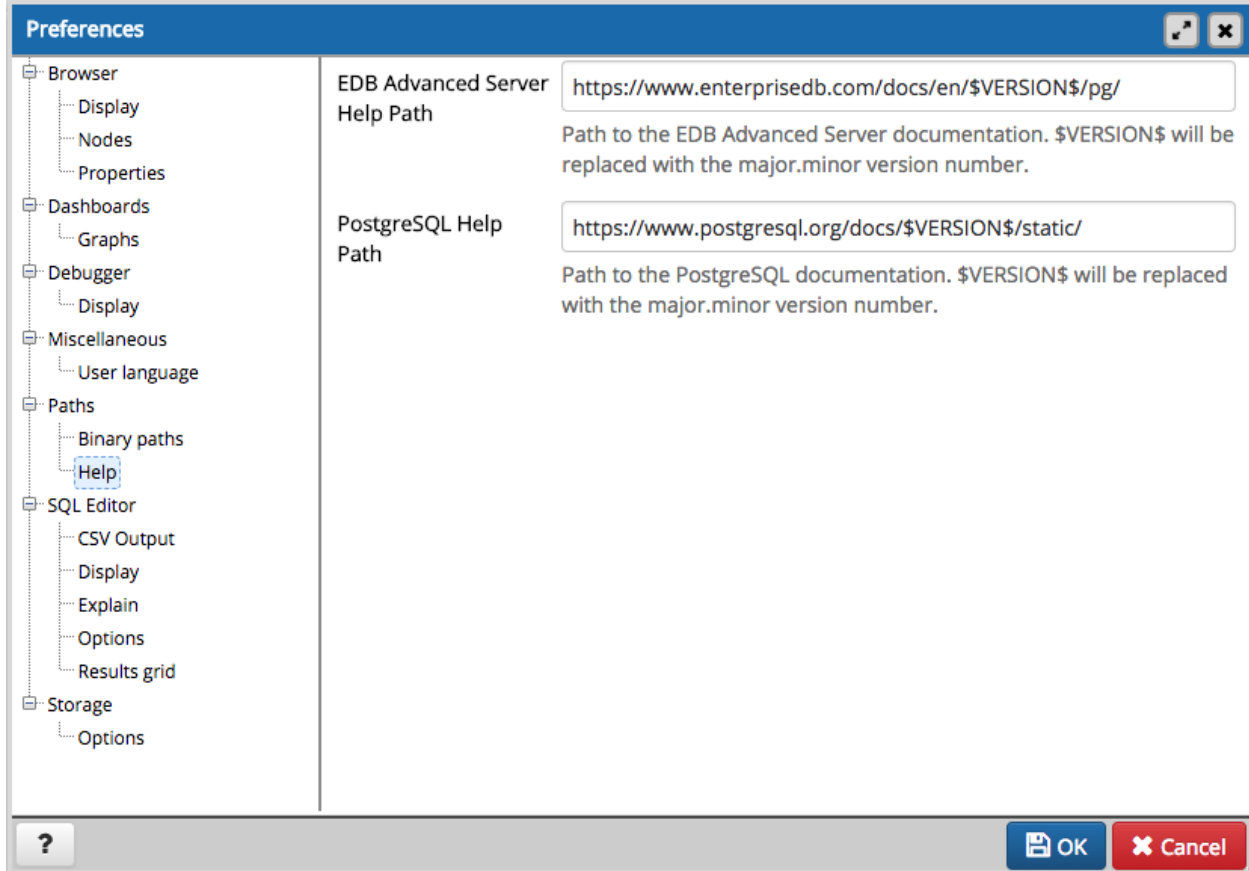
The Paths Node

Expand the *Paths* node to specify the locations of supporting utility and help files.



Use the fields on the *Binary paths* panel to specify the path to the directory that contains the utility programs (pg_dump, pg_restore, and pg_dumpall) for monitored databases:

- Use the *PostgreSQL Binary Path* field to specify the location of the PostgreSQL utility programs. If this path is not set, pgAdmin will attempt to find the utilities in standard locations used by PostgreSQL.



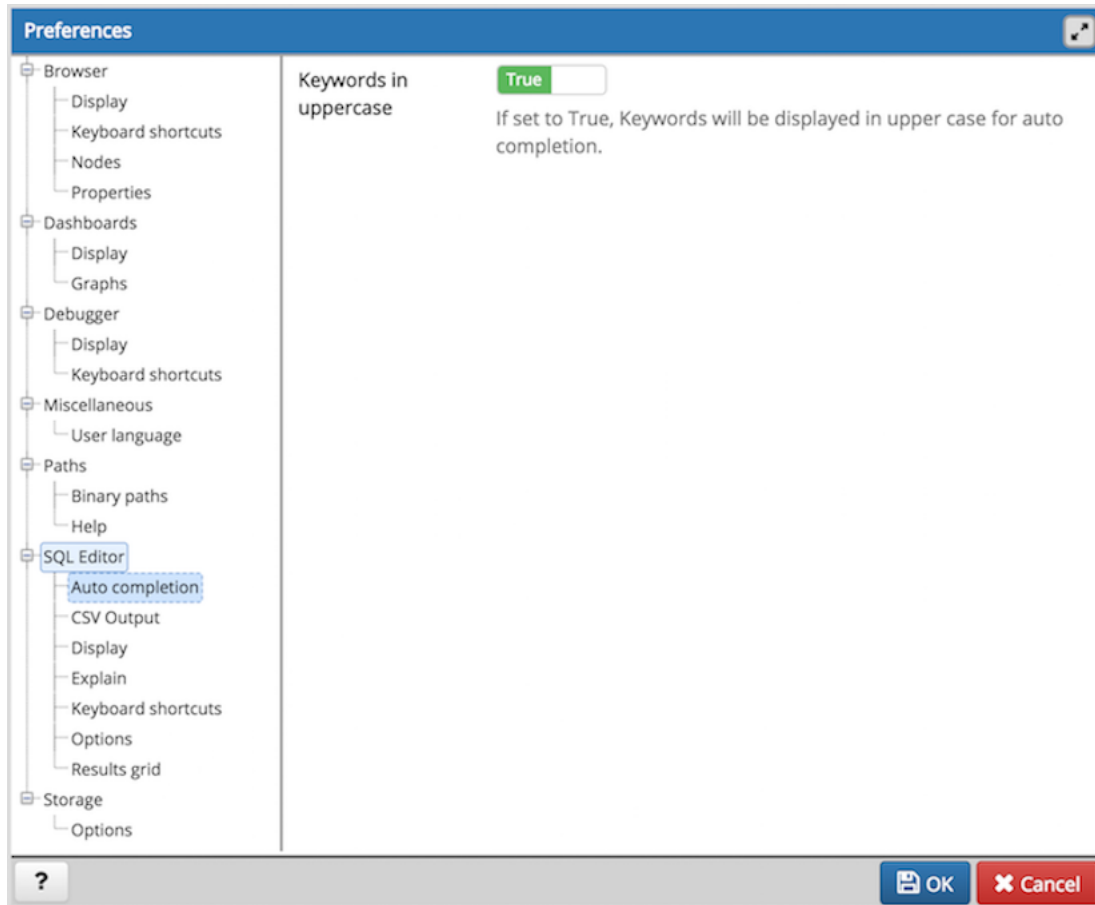
Use the fields on the *Help* panel to specify the location of help files.

- Use the *PostgreSQL Help Path* to specify the path to PostgreSQL documentation.

Please note: the default help paths include the *VERSION* placeholder; the *\$VERSION\$* placeholder will be replaced by the current database version.

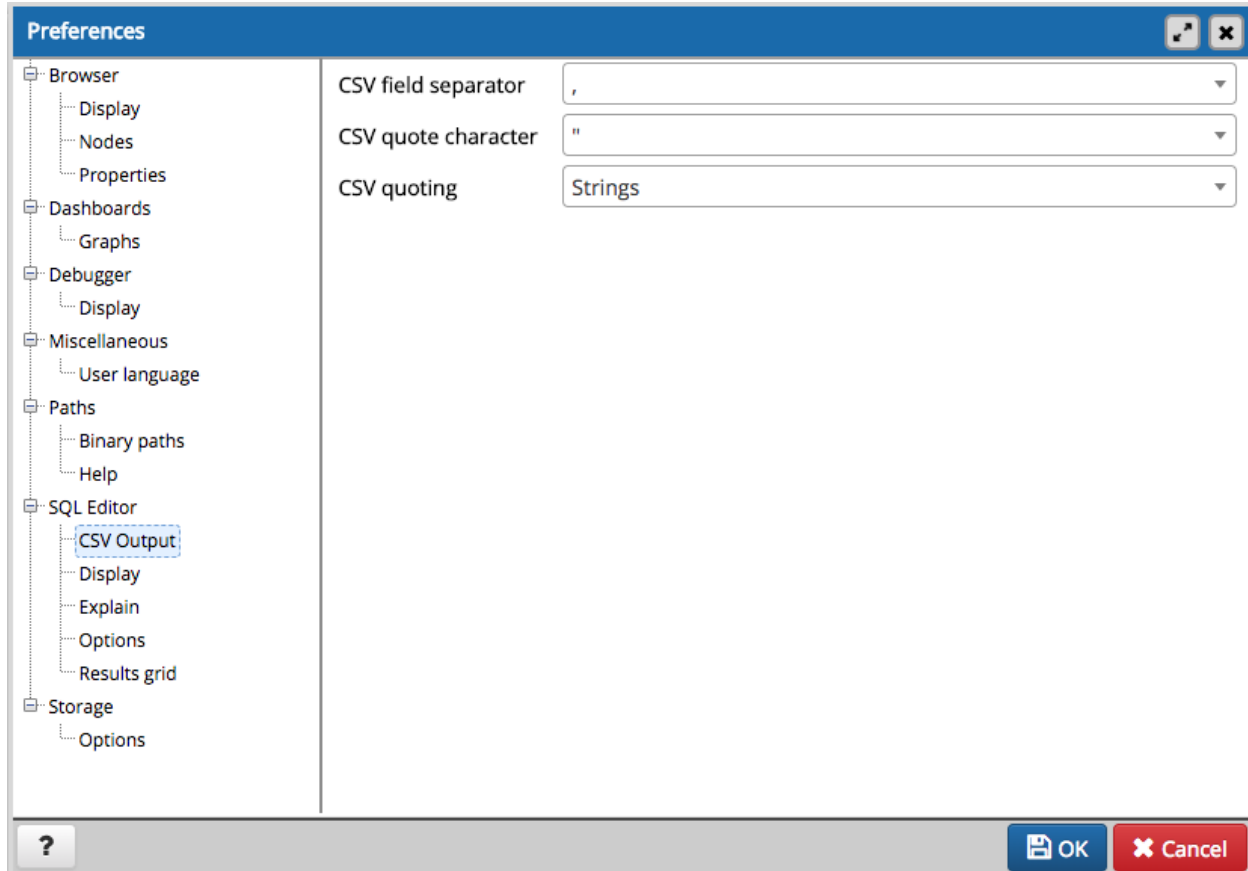
The SQL Editor Node

Expand the *SQL Editor* node to access panels that allow you to specify your preferences for the SQL Editor tool.



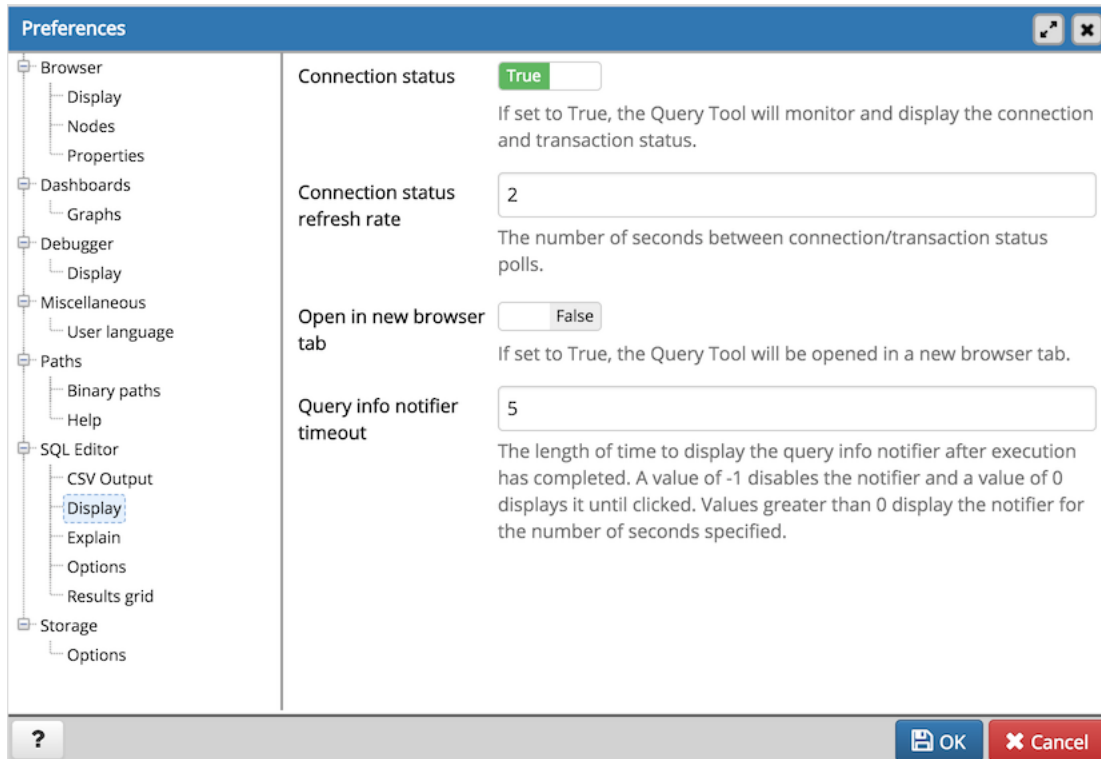
Use the fields on the *Auto Completion* panel to set the auto completion options.

- When the *Keywords in uppercase* switch is set to *True* then keywords are shown in upper case.




Use the fields on the *CSV Output* panel to control the CSV output.

- Use the *CSV field separator* drop-down listbox to specify the separator character that will be used in CSV output.
- Use the *CSV quote character* drop-down listbox to specify the quote character that will be used in CSV output.
- Use the *CSV quoting* drop-down listbox to select the fields that will be quoted in the CSV output; select *Strings*, *All*, or *None*.



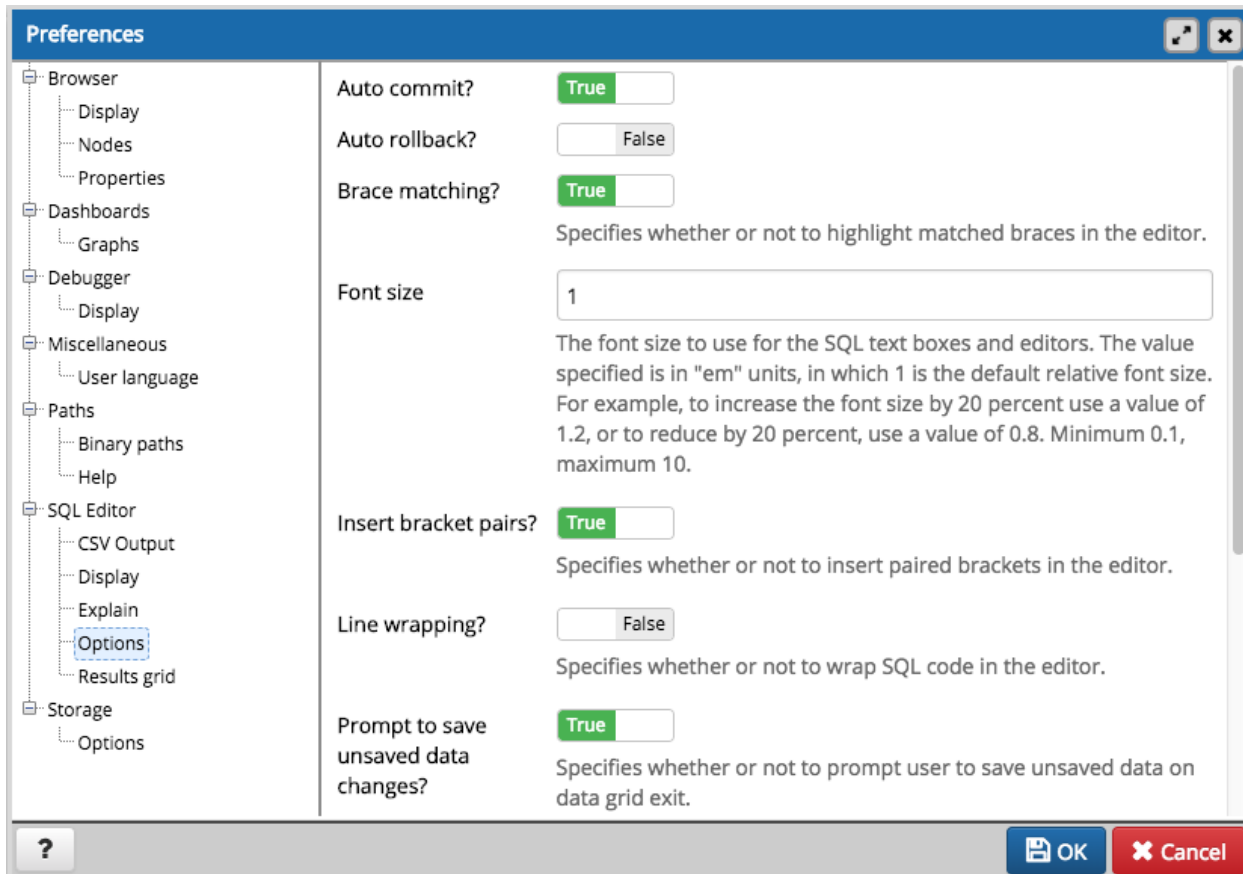
Use the fields on the *Display* panel to specify your preferences for the SQL Editor display.

- When the *Connection status* switch is set to *True*, each new instance of the SQL Editor will display connection and transaction status.
- Use the *Connection status refresh rate* field to specify the number of seconds between connection/transaction status updates.
- When the *Open in new browser tab* switch is set to *True*, each new instance of the SQL Editor will open in a new browser tab.
- Use the *Query info notifier timeout* field to control the behaviour of the notifier that is displayed when query execution completes. A value of *-1* will disable the notifier, and a value of *0* will display it until clicked. If a positive value above zero is specified, the notifier will be displayed for the specified number of seconds. The default is *5*.




Use the fields on the *Explain* panel to specify the level of detail included in a graphical EXPLAIN.

- When the *Show Buffers?* switch is set to *True*, graphical explain details will include information about buffer usage.
- When the *Show Costs?* switch is set to *True*, graphical explain details will include information about the estimated startup and total cost of each plan, as well as the estimated number of rows and the estimated width of each row.
- When the *Show Timing?* switch is set to *True*, graphical explain details will include the startup time and time spent in each node in the output.
- When the *Verbose output?* switch is set to *True*, graphical explain details will include extended information about the query execution plan.



Use the fields on the *Options* panel to manage editor preferences.

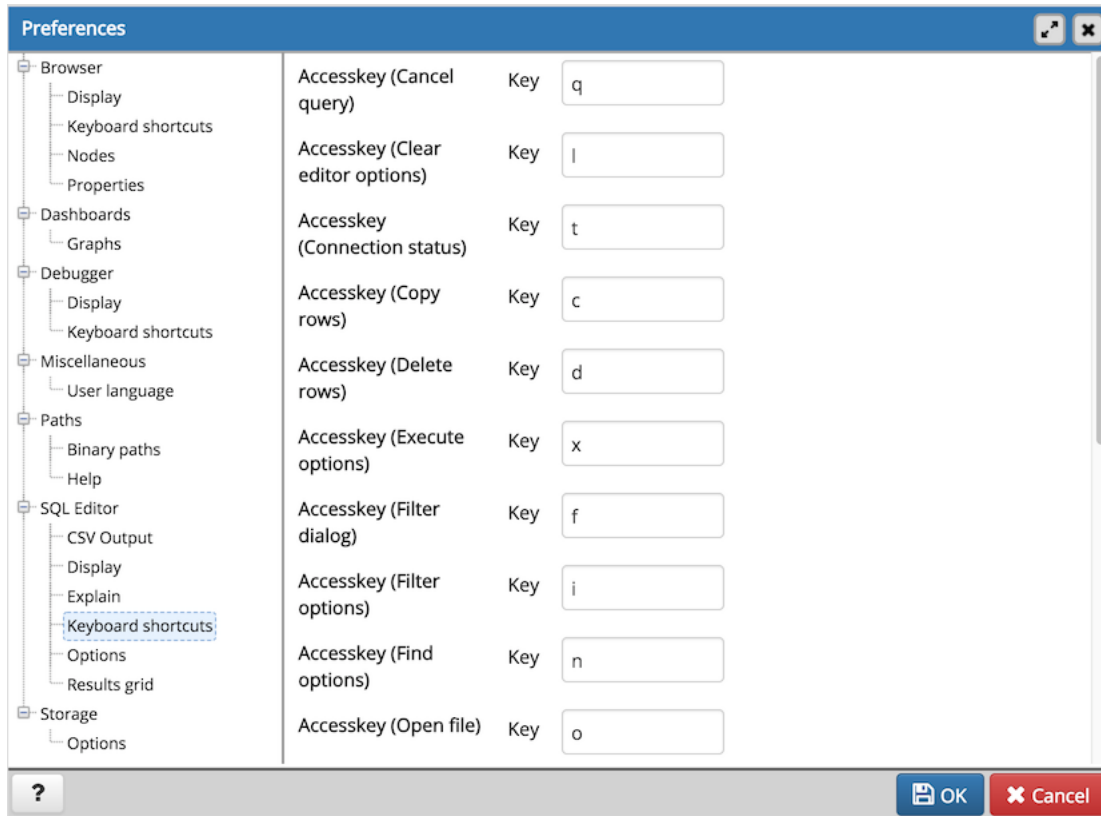
- When the *Auto-Commit?* switch is set to *True*, each successful query is committed after execution.
- When the *Auto-Rollback?* switch is set to *True*, failed queries are rolled back.
- When the *Brace matching?* switch is set to *True*, the editor will highlight pairs of matched braces.
- Use the *Font size* field to specify the font size that will be used in text boxes and editors.
- When the *Insert bracket pairs?* switch is set to *True*, the editor will automatically insert paired brackets.
- When the *Line wrapping* switch is set to *True*, the editor will implement line-wrapping behavior.
- When the *Prompt to save unsaved data changes?* switch is set to *True*, the editor will prompt the user to saved unsaved data when exiting the data editor.
- When the *Prompt to save unsaved query changes?* switch is set to *True*, the editor will prompt the user to saved unsaved query modifications when exiting the query tool.
- Use the *Tab size* field to specify the number of spaces per tab character in the editor.
- When the *Use spaces* switch is set to *True*, the editor will insert spaces (instead of tab characters) when the tab key or auto-indent are used.



Use the fields on the *Results grid* panel to specify your formatting preferences for copied data.

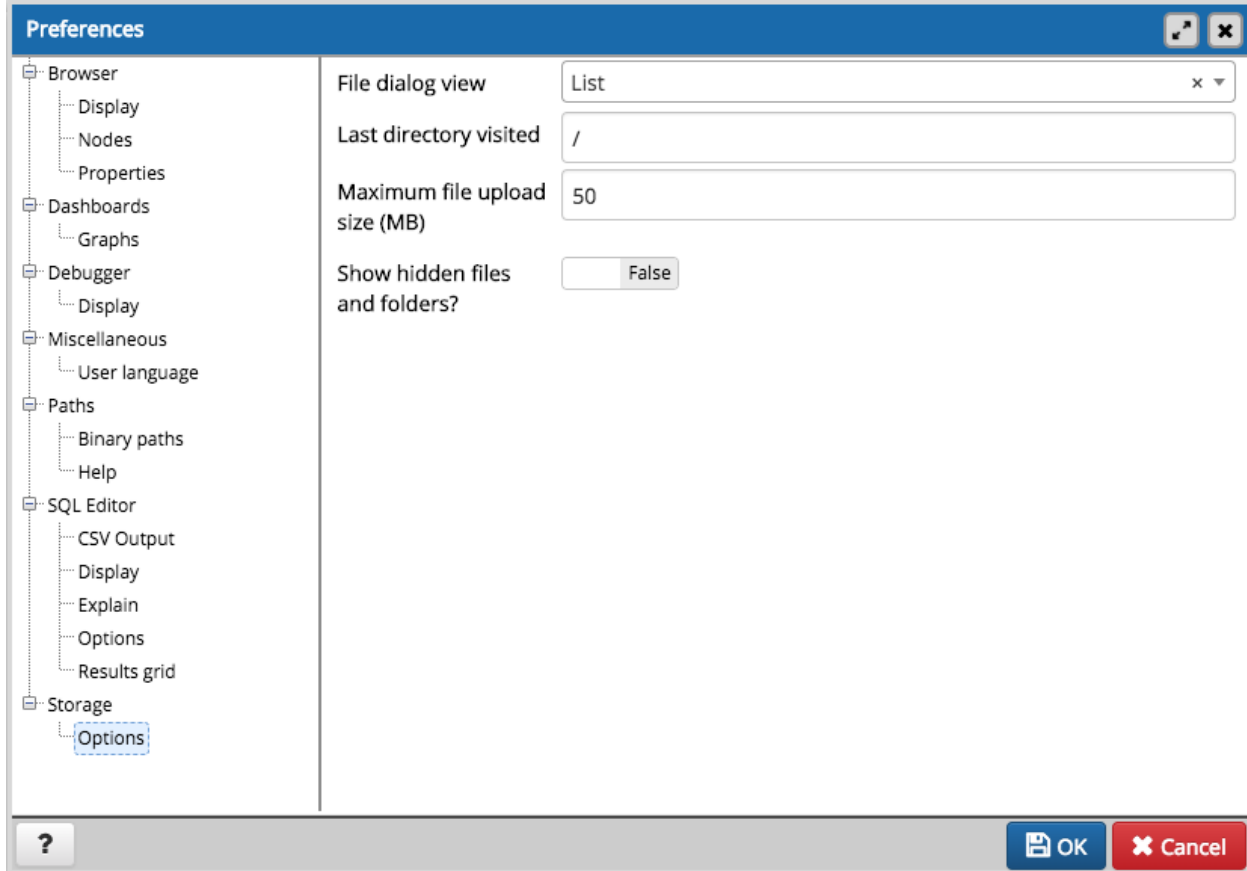
- Use the *Result copy field separator* drop-down listbox to select the field separator for copied data.
- Use the *Result copy quote character* drop-down listbox to select the quote character for copied data.
- Use the *Result copy quoting* drop-down listbox to select which type of fields require quoting; select *All*, *None*, or *Strings*.

Use the fields on the *Keyboard shortcuts* panel to configure shortcuts for the sql editor window navigation:



The Storage Node

Expand the *Storage* node to specify your storage preferences.



Use the fields on the *Options* panel to specify storage preferences.

- Use the *File dialog view* drop-down listbox to select the style of icons and display format that will be displayed when you open the file manager; select *List* to display a list view, or *Grid* to display folder icons.
- Use the *Last directory visited* field to specify the name of the folder in which the file manager will open.
- Use the *Maximum file upload size(MB)* field on the *Options* panel of the **Storage** node to specify the maximum file size for an upload.
- When the *Show hidden files and folders?* switch is set to *True*, the file manager will display hidden files and folders.

1.11 Keyboard Shortcuts

Keyboard shortcuts are provided in pgAdmin to allow easy access to specific functions. Alternate shortcuts can be configured through File > Preferences if desired.”

Main Browser Window

When using main browser window, the following keyboard shortcuts are available:

Shortcut for all platforms	Function
Alt+Shift+F	Open the File menu
Alt+Shift+O	Open the Object menu
Alt+Shift+L	Open the Tools menu
Alt+Shift+H	Open the Help menu
Alt+Shift+B	Focus the browser tree
Alt+Shift+[Move tabbed panel backward
Alt+Shift+]	Move tabbed panel forward
Alt+Shift+Q	Open the Query Tool in the current database
Alt+Shift+V	View Data in the selected table/view
Alt+Shift+C	Open the context menu
Alt+Shift+N	Create an object
Alt+Shift+E	Edit object properties
Alt+Shift+D	Delete the object
Alt+Shift+G	Direct debugging

Dialog tab shortcuts

Use the shortcuts below to navigate the tabsets on dialogs:

Shortcut for all platforms	Function
Control+Shift+[Dialog tab backward
Control+Shift+]	Dialog tab forward

SQL Editors

When using the syntax-highlighting SQL editors, the following shortcuts are available:

Shortcut (Windows/Linux)	Shortcut (Mac)	Function
Alt+Left	Option+Left	Move to the beginning of the line
Alt+Right	Option+Right	Move to the end of the line
Ctrl+Alt+Left	Cmd+Option+Left	Move left one word
Ctrl+Alt+Right	Cmd+Option+Right	Move right one word
Ctrl+/	Cmd+/	Comment selected code (Inline)
Ctrl+.	Cmd+.	Uncomment selected code (Inline)
Ctrl+Shift+/	Cmd+Shift+/	Comment/Uncomment code (Block)
Ctrl+A	Cmd+A	Select all
Ctrl+C	Cmd+C	Copy selected text to the clipboard
Ctrl+R	Cmd+R	Redo last edit un-done
Ctrl+V	Cmd+V	Paste text from the clipboard
Ctrl+Z	Cmd+Z	Undo last edit
Tab	Tab	Indent selected text
Shift+Tab	Shift+Tab	Un-indent selected text
Alt+G	Alt+G	Jump (to line:column)
Ctrl+Space	Ctrl+Space	Auto-complete
Ctrl+F	Cmd+F	Find
Ctrl+G	Cmd+G	Find next
Ctrl+Shift+G	Cmd+Shift+G	Find previous
Ctrl+Shift+F	Cmd+Shift+F	Replace

Query Tool

When using the Query Tool, the following shortcuts are available:

Shortcut (Windows/Linux)	Shortcut (Mac)	Function
F5	F5	Execute query
F7	F7	EXPLAIN query
Shift+F7	Shift+F7	EXPLAIN ANALYZE query
F8	F8	Execute query to CSV file
<accesskey> + o	<accesskey> + o	Open file
<accesskey> + s	<accesskey> + s	Save file
<accesskey> + n	<accesskey> + n	Find option drop down
<accesskey> + c	<accesskey> + c	Copy row(s)
<accesskey> + p	<accesskey> + p	Paste row(s)
<accesskey> + d	<accesskey> + d	Delete row(s)
<accesskey> + f	<accesskey> + f	Filter dialog
<accesskey> + i	<accesskey> + i	Filter options drop down
<accesskey> + r	<accesskey> + r	Row limit
<accesskey> + q	<accesskey> + q	Cancel query
<accesskey> + l	<accesskey> + l	Clear option drop down
<accesskey> + x	<accesskey> + x	Execute option drop down
<accesskey> + t	<accesskey> + t	Display connection status
<accesskey> + y	<accesskey> + y	Copy SQL on history panel

Debugger

When using the Debugger, the following shortcuts are available:

Shortcut (Windows/Linux)	Shortcut (Mac)	Function
<accesskey> + i	<accesskey> + i	Step in
<accesskey> + o	<accesskey> + o	Step over
<accesskey> + c	<accesskey> + c	Continue/Restart
<accesskey> + t	<accesskey> + t	Toggle breakpoint
<accesskey> + x	<accesskey> + x	Clear all breakpoints
<accesskey> + s	<accesskey> + s	Stop
Alt + Shift + q	Alt + Shift + q	Enter or Edit values in Grid

Inner panel navigation

When using the Query Tool and Debugger, the following shortcuts are available for inner panel navigation:

Shortcut (Windows/Linux)	Shortcut (Mac)	Function
Alt + Shift + Right Arrow	Alt + Shift + Right Arrow	Move to next inner panel
Alt + Shift + Left Arrow	Alt + Shift + Left Arrow	Move to previous inner panel

Note: <accesskey> is browser and platform dependant. The following table lists the default access keys for supported browsers.

	Windows	Linux	Mac
Internet Explorer	Alt	Alt	
Chrome	Alt	Alt	Ctrl+Alt
Firefox	Alt+Shift	Alt+Shift	Ctrl+Alt
Safari	Alt		Ctrl+Alt

Before using pgAdmin to manage objects that reside on a server, you must define a connection to the server; for more information please see *Connecting to a Server*:

Contents:

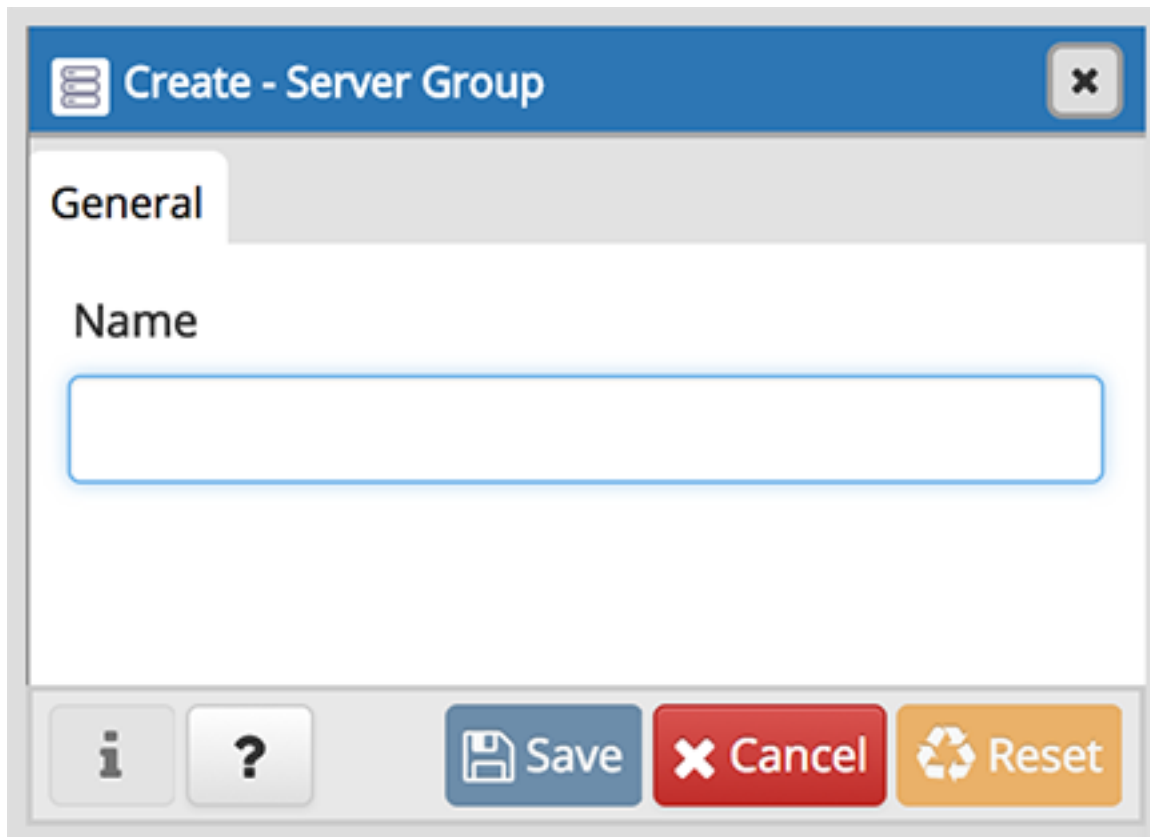
1.12 Connecting to a Server

Before you can use the pgAdmin client to manage the objects that reside on your Postgres server, you must define a connection to the server. You can (optionally) use the *Server Group* dialog to create server groups to organize the server connections within the tree control for easier management. To open the *Server Group* dialog, right-click on the *Servers* node of the tree control, and select *Server Group* from the *Create* menu.

Contents:

1.12.1 Server Group Dialog

Use the *Server Group* dialog to add a new server group. Assign servers to server groups to simplify management of multiple servers. Server groups are displayed as part of the *pgAdmin* tree control.



Use the *Name* field on the *Server Group* dialog to specify a name that will identify the server group in the *pgAdmin* tree control.

- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

To create server connections in a server group, right click on the named server group and select the *Create* option to open the *Create - Server* dialog.

Use the fields on the *Server* dialog to define the connection properties for each new server that you wish to manage with pgAdmin. To open the *Server* dialog, right-click on the *Servers* node of the tree control, and select *Server* from the *Create* menu.

Contents:

1.12.2 Server Dialog

Use the *Server* dialog to describe a connection to a server. Note: you must ensure that the `pg_hba.conf` file of the server from which you are connecting allows connections from the host of the client.

The screenshot shows the 'Create - Server' dialog box in pgAdmin 4. The dialog has a blue title bar with a close button. Below the title bar are tabs for 'General', 'Connection', 'SSL', 'SSH Tunnel', and 'Advanced'. The 'General' tab is active. It contains the following fields and controls:

- Name:** A text input field containing 'Server1'.
- Server group:** A drop-down list box showing 'Servers'.
- Background:** A color picker control.
- Foreground:** A color picker control.
- Connect now?:** A checked checkbox.
- Comments:** A large text area.

At the bottom of the dialog are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.


Use the fields in the *General* tab to identify the server:

- Use the *Name* field to add a descriptive name for the server; the name specified will be displayed in the *Browser* tree control.
- Use the drop-down list box in the *Server group* field to select the parent node for the server; the server will be displayed in the *Browser* tree control within the specified group.
- Use the color-picker in the *Background* field to specify the background color for the server.
- Use the color-picker in the *Foreground* field to specify the foreground color for the server.
- If the *Connect now?* checkbox is checked, the client will attempt a connection to the server upon completion of

the dialog; this is the default

- Provide a comment about the server in the *Comments* field.

Click the *Connection* tab to continue.



Create - Server

General **Connection** SSL SSH Tunnel Advanced

Host name/address: 172.10.11.12

Port: 5432

Maintenance database: postgres

Username: postgres

Password:

Save password?

Role:

Service:


Save Cancel Reset

Use the fields in the *Connection* tab to configure a connection:

- Specify the IP address of the server host, or the fully qualified domain name in the *Host name/address* field. If you provide a unix domain socket, the directory name must begin with a “/”.
- Enter the listener port number of the server host in the *Port* field. The default is 5432.

- Use the *Maintenance database* field to specify the name of the initial database to which the client will connect. If you will be using pgAgent or adminpack objects, the pgAgent schema and adminpack objects should be installed on that database.
- Use the *Username* field to specify the name of a role that will be used when authenticating with the server.
- Use the *Password* field to provide a password that will be supplied when authenticating with the server.
- Check the box next to *Save password?* to instruct pgAdmin to save the password for future use. Use *Clear Saved Password* to remove the saved password.
- Use the *Role* field to specify the name of a role that has privileges that will be conveyed to the client after authentication with the server. This selection allows you to connect as one role, and then assume the permissions of this specified role after the connection is established. Note that the connecting role must be a member of the role specified.
- Use the *Service* field to specify the service name. For more information, see [Section 33.16 of the Postgres documentation](#).

Click the *SSL* tab to continue.



Use the fields in the *SSL* tab to configure SSL:

- Use the drop-down list box in the *SSL* field to select the type of SSL connection the server should use. For more information about using SSL encryption, see [Section 33.18 of the Postgres documentation](#).

If pgAdmin is installed in Server mode (the default mode), you can use the platform-specific File manager dialog to upload files that support SSL encryption to the server. To access the File manager dialog, click the icon that is located to the right of each of the following fields.


- Use the *Client certificate* field to specify the file containing the client SSL certificate. This file will replace the default `~/.postgresql/postgresql.crt` if pgAdmin is installed in Desktop mode, and `<STORAGE_DIR>/<USERNAME>/.postgresql/postgresql.crt` if pgAdmin is installed in Web mode. This parameter

is ignored if an SSL connection is not made.

- Use the *Client certificate key* field to specify the file containing the secret key used for the client certificate. This file will replace the default `~/.postgresql/postgresql.key` if pgAdmin is installed in Desktop mode, and `<STORAGE_DIR>/<USERNAME>/.postgresql/postgresql.key` if pgAdmin is installed in Web mode. This parameter is ignored if an SSL connection is not made.
- Use the *Root certificate* field to specify the file containing the SSL certificate authority. This file will replace the default `~/.postgresql/root.crt`. This parameter is ignored if an SSL connection is not made.
- Use the *Certificate revocation list* field to specify the file containing the SSL certificate revocation list. This list will replace the default list, found in `~/.postgresql/root.crl`. This parameter is ignored if an SSL connection is not made.
- When *SSL compression?* is set to *True*, data sent over SSL connections will be compressed. The default value is *False* (compression is disabled). This parameter is ignored if an SSL connection is not made.

WARNING: In Server mode, certificates, private keys, and the revocation list are stored in the per-user file storage area on the server, which is owned by the user account under which the pgAdmin server process is run. This means that administrators of the server may be able to access those files; appropriate caution should be taken before choosing to use this feature.

Click the *SSH Tunnel* tab to continue.



The screenshot shows the 'Create - Server' dialog box in pgAdmin 4, with the 'SSH Tunnel' tab selected. The dialog has a blue title bar with a close button. Below the title bar are five tabs: 'General', 'Connection', 'SSL', 'SSH Tunnel', and 'Advanced'. The 'SSH Tunnel' tab contains the following fields and controls:

- Use SSH tunneling:** A toggle switch currently set to 'No'.
- Tunnel host:** An empty text input field.
- Tunnel port:** A text input field containing the value '22'.
- Username:** An empty text input field.
- Authentication:** A toggle switch currently set to 'Password'.
- Identity file:** A text input field with a file selection icon (three dots) on the right.
- Password:** An empty text input field.
- Save password?:** An unchecked checkbox.

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two smaller buttons: an information icon ('i') and a help icon ('?').


Use the fields in the *SSH Tunnel* tab to configure SSH Tunneling:

You can use the “SSH Tunnel” tab to connect pgAdmin (through an intermediary proxy host) to a server that resides on a network to which the client may not be able to connect directly.

- Set “Use SSH tunneling” to *Yes* to specify that pgAdmin should use an SSH tunnel when connecting to the specified server.
- Specify the name or IP address of the SSH host (through which client connections will be forwarded) in the *Tunnel host* field.
- Specify the port of the SSH host (through which client connections will be forwarded) in the *Tunnel port* field.

- Specify the name of a user with login privileges for the SSH host in the *Username* field.
- Specify the type of authentication that will be used when connecting to the SSH host in the *Authentication* field.
 - Select the *Password* option to specify that pgAdmin will use a password for authentication to the SSH host. This is the default.
 - Select the *Identity file* to specify that pgAdmin will use a private key file when connecting.
- If the SSH host is expecting a private key file for authentication, use the *Identity file* field to specify the location of the key file.
- If the SSH host is expecting a password of the user name or an identity file if being used, use the *Password* field to specify the password.
- Check the box next to *Save password?* to instruct pgAdmin to save the password for future use. Use *Clear SSH Tunnel Password* to remove the saved password.

Click the *Advanced* tab to continue.



The screenshot shows the 'Create - Server' dialog box in pgAdmin 4, with the 'Advanced' tab selected. The dialog has a blue title bar with a close button. Below the title bar are tabs for 'General', 'Connection', 'SSL', 'SSH Tunnel', and 'Advanced'. The 'Advanced' tab contains four input fields: 'Host address' (empty), 'DB restriction' (containing 'Select from the list'), 'Password file' (empty with a file selection button), and 'Connection timeout (seconds)' (containing '0'). At the bottom of the dialog are buttons for 'i', '?', 'Save', 'Cancel', and 'Reset'.

Use the fields in the *Advanced* tab to configure a connection:

- Specify the IP address of the server host in the *Host address* field. Using this field to specify the host IP address may save time by avoiding a DNS lookup on connection, but it may be useful to specify both a host name and address when using Kerberos, GSSAPI, or SSPI authentication methods, as well as for verify-full SSL certificate verification.
- Use the *DB restriction* field to provide a SQL restriction that will be used against the `pg_database` table to limit the databases that you see. For example, you might enter: `live_db test_db` so that only `live_db` and `test_db` are shown in the pgAdmin browser. Separate entries with a comma or tab as you type.
- Use the *Password File* field to specify the location of a password file (`.pgpass`). A `.pgpass` file allows a user to

login without providing a password when they connect. For more information, see [Section 33.15 of the Postgres documentation](#).

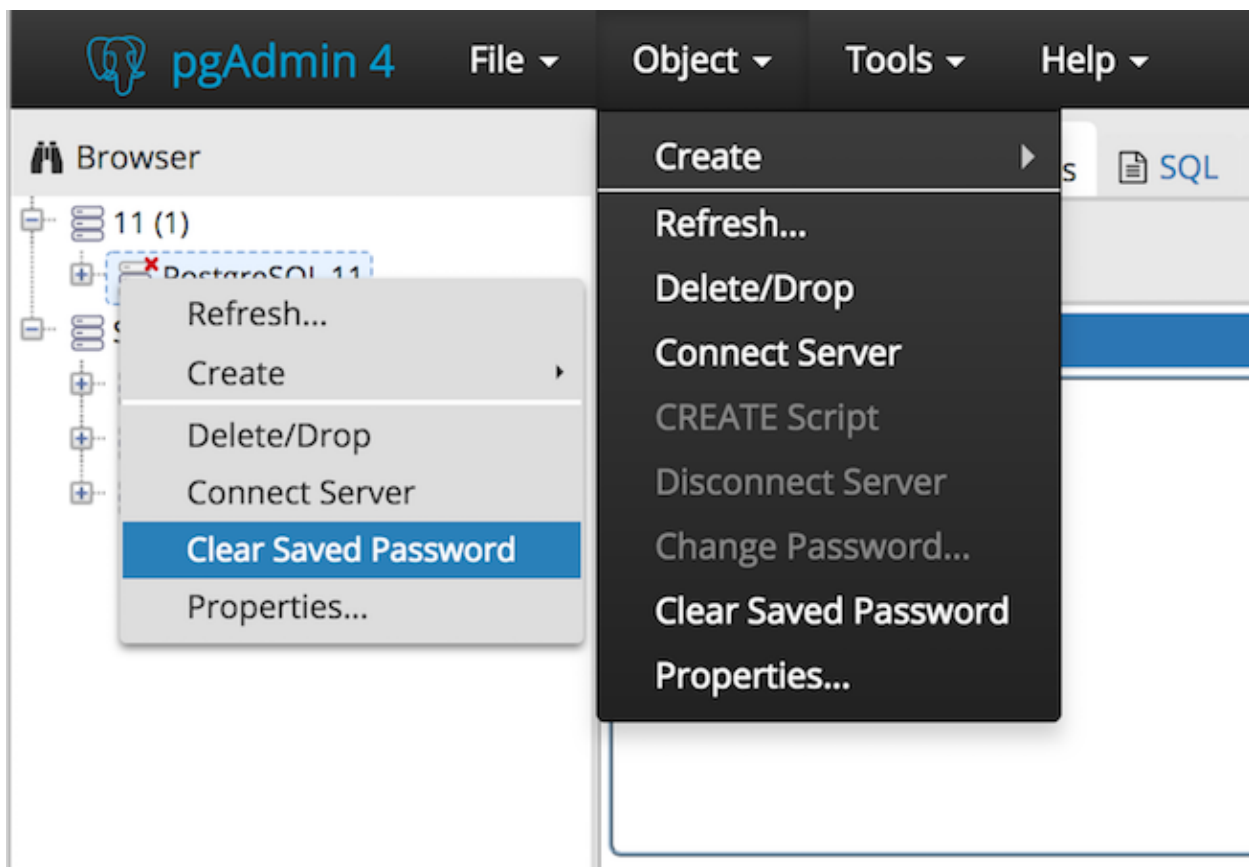
- Use the *Connection timeout* field to specify the maximum wait for connection, in seconds. Zero or not specified means wait indefinitely. It is not recommended to use a timeout of less than 2 seconds.

NOTE: The password file option is only supported when pgAdmin is using libpq v10.0 or later to connect to the server.

- Click the *Save* button to save your work.
- Click the *Cancel* button to exit without saving your work.
- Click the *Reset* button to return the values specified on the Server dialog to their original condition.

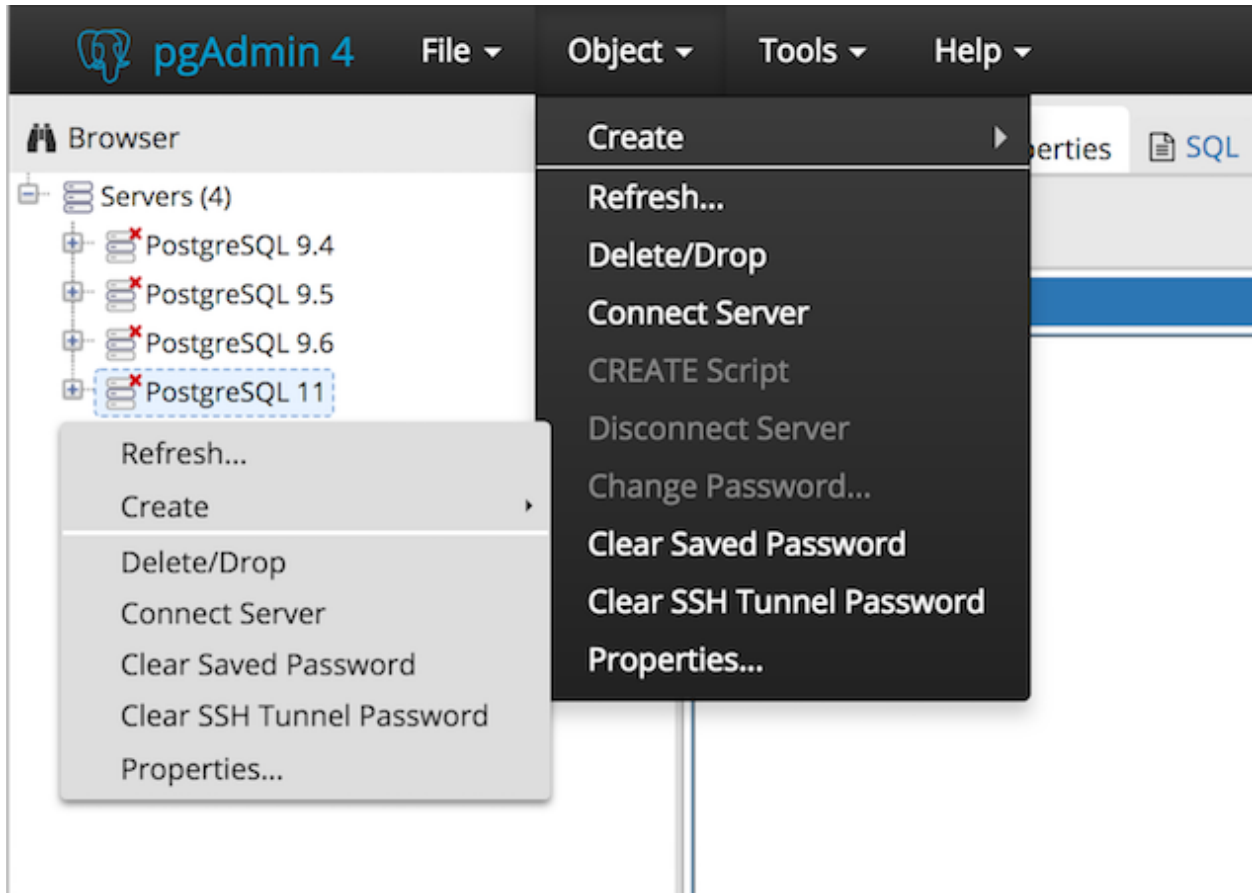
Clear Saved Passwords

Use *Clear Saved Password* functionality to clear the saved password for the database server.



Clear Saved Password shows in the context menu for the selected server as well as under the *Object* menu on the top menu bar.

Use *Clear SSH Tunnel Password* functionality to clear the saved password of SSH Tunnel to connect to the database server.



Clear SSH Tunnel Password shows in the context menu for the selected server as well as under the *Object* menu on the top menu bar.

Note: It will be enabled/visible when the password for the selected database server is already saved.

After defining a server connection, right-click on the server name, and select *Connect to server* to authenticate with the server, and start using pgAdmin to manage objects that reside on the server.

Contents:

1.12.3 Connect to server

Use the *Connect to Server* dialog to authenticate with a defined server and access the objects stored on the server through the pgAdmin tree control. To access the dialog, right click on the server name in the *pgAdmin* tree control, and select *Connect Server...* from the context menu.

Connect to Server [X]

Please enter the password for the user 'enterprisedb' to connect the server - "EPAS 9.6 Linux"

Password

Save Password

OK Cancel

Provide authentication information for the selected server:

- Use the *Password* field to provide the password of the user that is associated with the defined server.
- Check the box next to *Save Password* to instruct the server to save the password for future connections; if you save the password, you will not be prompted when reconnecting to the database server with this server definition.

In case of SSH Tunneling, *Connect to Server* dialog will prompt for SSH Tunnel and Database server passwords if not already saved.

Connect to Server

SSH Tunnel password for the user 'akshay' to connect the server "127.0.0.1"

Save Password

Database server password for the user 'postgres' to connect the server "PostgreSQL 11"

Save Password

OK Cancel

Provide authentication information for the selected server:

- Use the *Password* field to provide the password of the user that is associated with the defined server.
- Check the box next to respective *Save Password* to instruct the server to save the password for future connections; if you save the password, you will not be prompted when reconnecting to the database server with this server definition.

The pgAdmin client displays a message in a green status bar in the lower right corner when the server connects successfully.

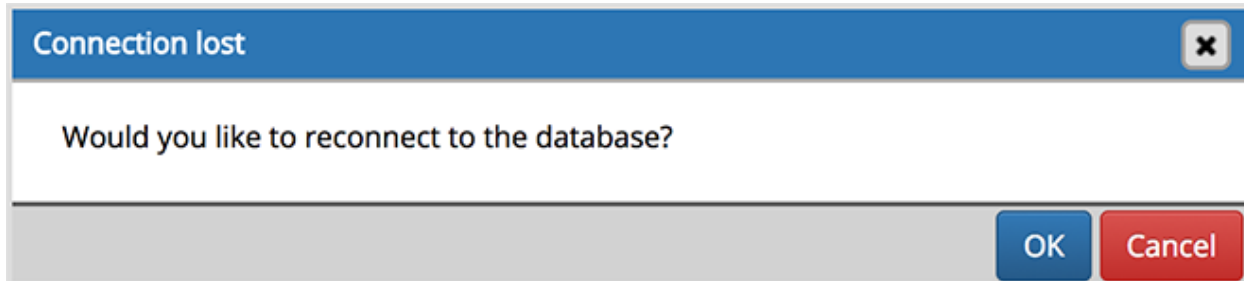
If you receive an error message while attempting a connection, verify that your network is allowing the pgAdmin host and the host of the database server to communicate. For detailed information about a specific error message, please see the *Connection Error* help page.

To review or modify connection details, right-click on the name of the server, and select *Properties...* from the context menu.

1.12.4 Connection error

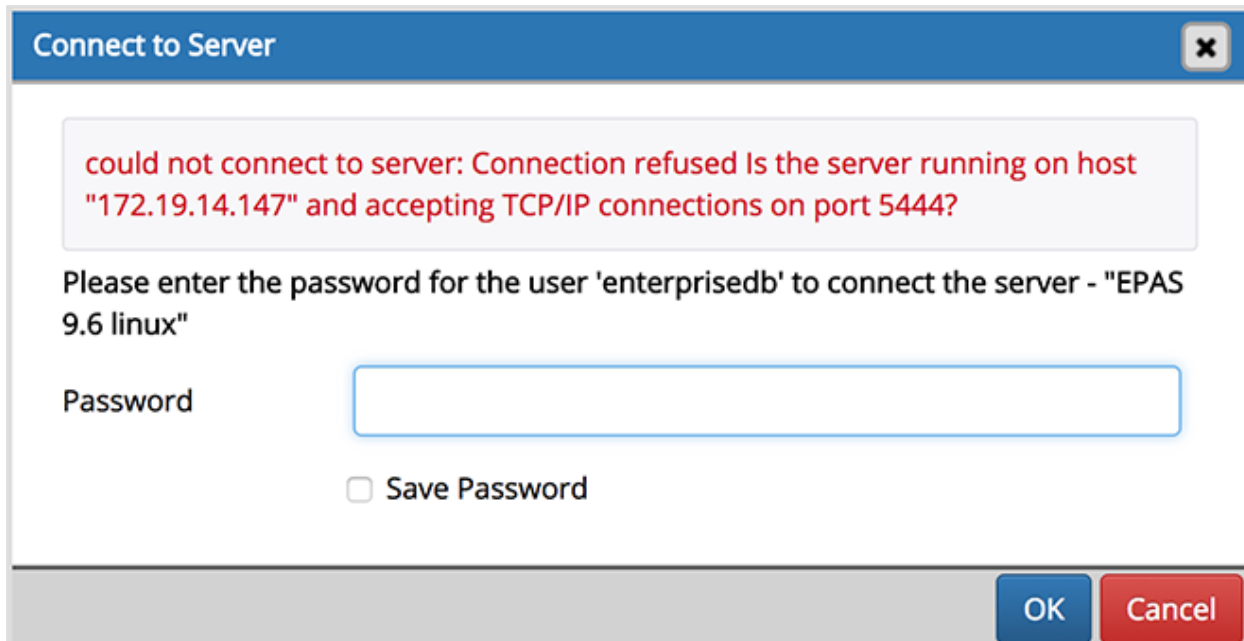
When connecting to a PostgreSQL server, you may get an error message. If you encounter an error message, please review the message carefully; each error message attempts to incorporate the information you'll need to resolve the problem. For more details about specific errors, please locate the error message in the list below:

Connection to the server has been lost



This error message indicates that the connection attempt has taken longer than the specified threshold; there may be a problem with the connection properties provided on the *Server* dialog, network connectivity issues, or the server may not be running.

could not connect to Server: Connection refused




If pgAdmin displays this message, there are two possible reasons for this:

- the database server isn't running - simply start it.
- the server isn't configured to accept TCP/IP requests on the address shown.

For security reasons, a PostgreSQL server “out of the box” doesn't listen on TCP/IP ports. Instead, it must be enabled to listen for TCP/IP requests. This can be done by adding **tcpip = true** to the `postgresql.conf` file for Versions 7.3.x and 7.4.x, or **listen_addresses='*'** for Version 8.0.x and above; this will make the server accept connections on any IP interface.

For further information, please refer to the PostgreSQL documentation about [runtime configuration](#).

FATAL: no pg_hba.conf entry



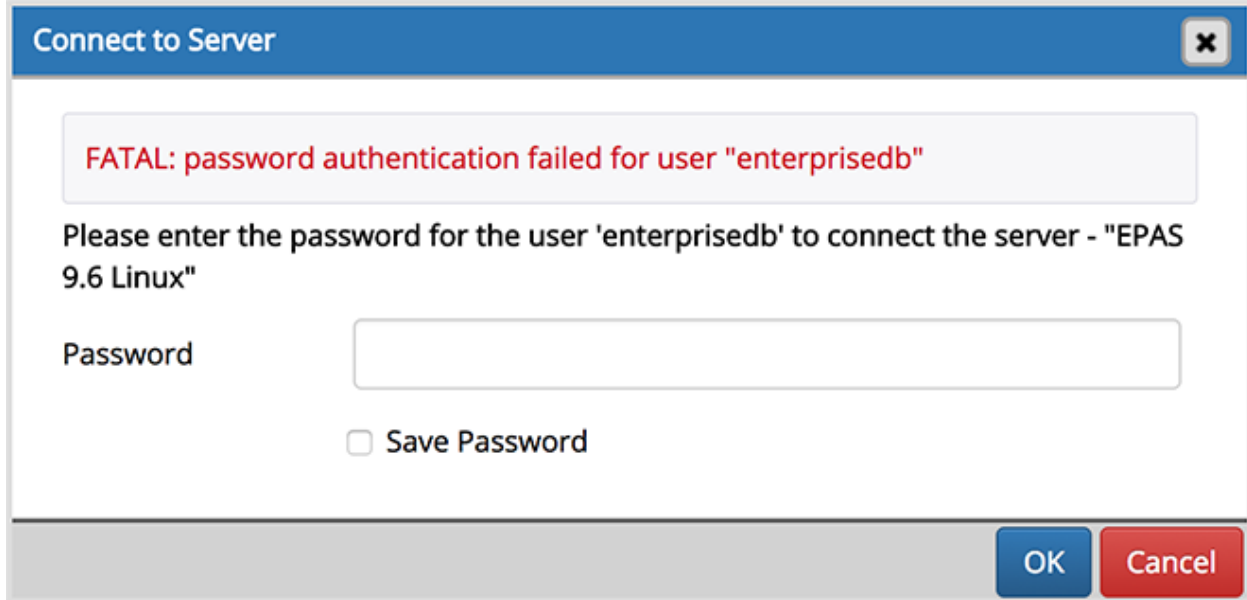
If pgAdmin displays this message when connecting, your server can be contacted correctly over the network, but is not configured to accept your connection. Your client has not been detected as a legal user for the database.

To connect to a server, the `pg_hba.conf` file on the database server must be configured to accept connections from the host of the pgAdmin client. Modify the `pg_hba.conf` file on the database server host, and add an entry in the form:

- **host template1 postgres 192.168.0.0/24 md5** for an IPV4 network
- **host template1 postgres ::ffff:192.168.0.0/120 md5** for an IPV6 network

For more information, please refer to the PostgreSQL documentation about [client authentication](#).

FATAL: password authentication failed



- The *password authentication failed for user* error message indicates there may be a problem with the password you entered. Retry the password to confirm you entered it correctly. If the error message returns, make sure that you have the correct password, that you are authorized to access the server, and that the access has been correctly configured in the server's `postgresql.conf` configuration file.

MANAGING CLUSTER LEVEL OBJECTS


Some object definitions reside at the cluster level; pgAdmin 4 provides dialogs that allow you to create these objects, manage them, and control their relationships to each other. To access a dialog that allows you to create a database object, right-click on the object type in the pgAdmin tree control, and select the *Create* option for that object. For example, to create a new database, right-click on the *Databases* node, and select *Create Database...*

Contents:

2.1 Database Dialog

Use the *Database* dialog to define or modify a database. To create a database, you must be a database superuser or have the CREATE privilege.

The *Database* dialog organizes the development of a database through the following dialog tabs: *General*, *Definition*, *Security*, and *Parameters*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Database' dialog box in pgAdmin 4. The window title is 'Create - Database' and it has a close button (X) in the top right corner. The dialog has five tabs: 'General', 'Definition', 'Security', 'Parameters', and 'SQL'. The 'General' tab is selected and active. It contains three main fields: 'Database' (a text input field), 'Owner' (a drop-down listbox with 'postgres' selected), and 'Comment' (a large text area). At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two smaller buttons: 'i' (information) and '?' (help).

Use the fields in the *General* tab to identify the database:

- Use the *Database* field to add a descriptive name for the database. The name will be displayed in the *pgAdmin* tree control.
- Select the owner of the database from the drop-down listbox in the *Owner* field.
- Store notes about the database in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Database' dialog box with the 'Definition' tab selected. The fields and their values are as follows:

Field	Value
Encoding	UTF8
Template	pem
Tablespace	pg_default
Collation	C
Character type	C
Connection limit	-1

Use the *Definition* tab to set properties for the database:

- Select a character set from the drop-down listbox in the *Encoding* field. The default is *UTF8*.
- Select a template from the drop-down listbox in the *Template* field. If you do not specify a template, the database will use template1.
- Select a tablespace from the drop-down listbox in the *Tablespace* field. The selected tablespace will be the default tablespace used to contain database objects.
- Select the collation order from the drop-down listbox in the *Collation* field.
- Select the character classification from the drop-down listbox in the *Character Type* field. This affects the categorization of characters, e.g. lower, upper and digit. The default, or a blank field, uses the character classification of the template database.

- Specify a connection limit in the *Connection Limit* field to configure the maximum number of connection requests. The default value (-1) allows unlimited connections to the database.

Click the *Security* tab to continue.

The screenshot shows the 'Create - Database' dialog box in pgAdmin 4, with the 'Security' tab selected. The dialog has four tabs: 'General', 'Definition', 'Security', 'Parameters', and 'SQL'. The 'Security' tab contains two main sections: 'Privileges' and 'Security Labels'.

The 'Privileges' section is a table with three columns: 'Grantee', 'Privileges', and 'Grantor'. The 'Grantee' column has a dropdown menu with 'pem_admin' selected. The 'Privileges' column has a grid of checkboxes for various privileges: ALL, CREATE, TEMPORARY, and CONNECT, each with a 'WITH GRANT OPTION' checkbox. All these checkboxes are checked. The 'Grantor' column has a dropdown menu with 'postgres' selected.

The 'Security Labels' section is a table with two columns: 'Provider' and 'Security Label'. The 'Provider' column has a dropdown menu with 'myProvider' selected. The 'Security Label' column has a text field with 'mySecurity' entered.

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left.

Use the *Security* tab to assign privileges and define security labels.

Use the *Privileges* panel to assign privileges to a role. Click the *Add* icon (+) to set privileges for database objects:

- Select the name of the role from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.
- Select the name of the role from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click add to set additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the database. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

To discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Parameters* tab to continue.

The screenshot shows the 'Create - Database' dialog box in pgAdmin 4, with the 'Parameters' tab selected. The dialog has a title bar with a close button (X) and a hamburger menu icon. Below the title bar are tabs for 'General', 'Definition', 'Security', 'Parameters', and 'SQL'. The 'Parameters' tab contains a table with the following structure:

	Name	Value	Role
	role	2	pem_admin

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two small icons: an information icon (i) and a help icon (?).

Use the *Parameters* tab to set parameters for the database. Click the *Add* icon (+) to add each parameter:

- Use the drop-down listbox in the *Name* field to select a parameter.
- Use the *Value* field to set a value for the parameter.
- Use the drop-down listbox next to *Role* to select a role to which the parameter setting specified will apply.

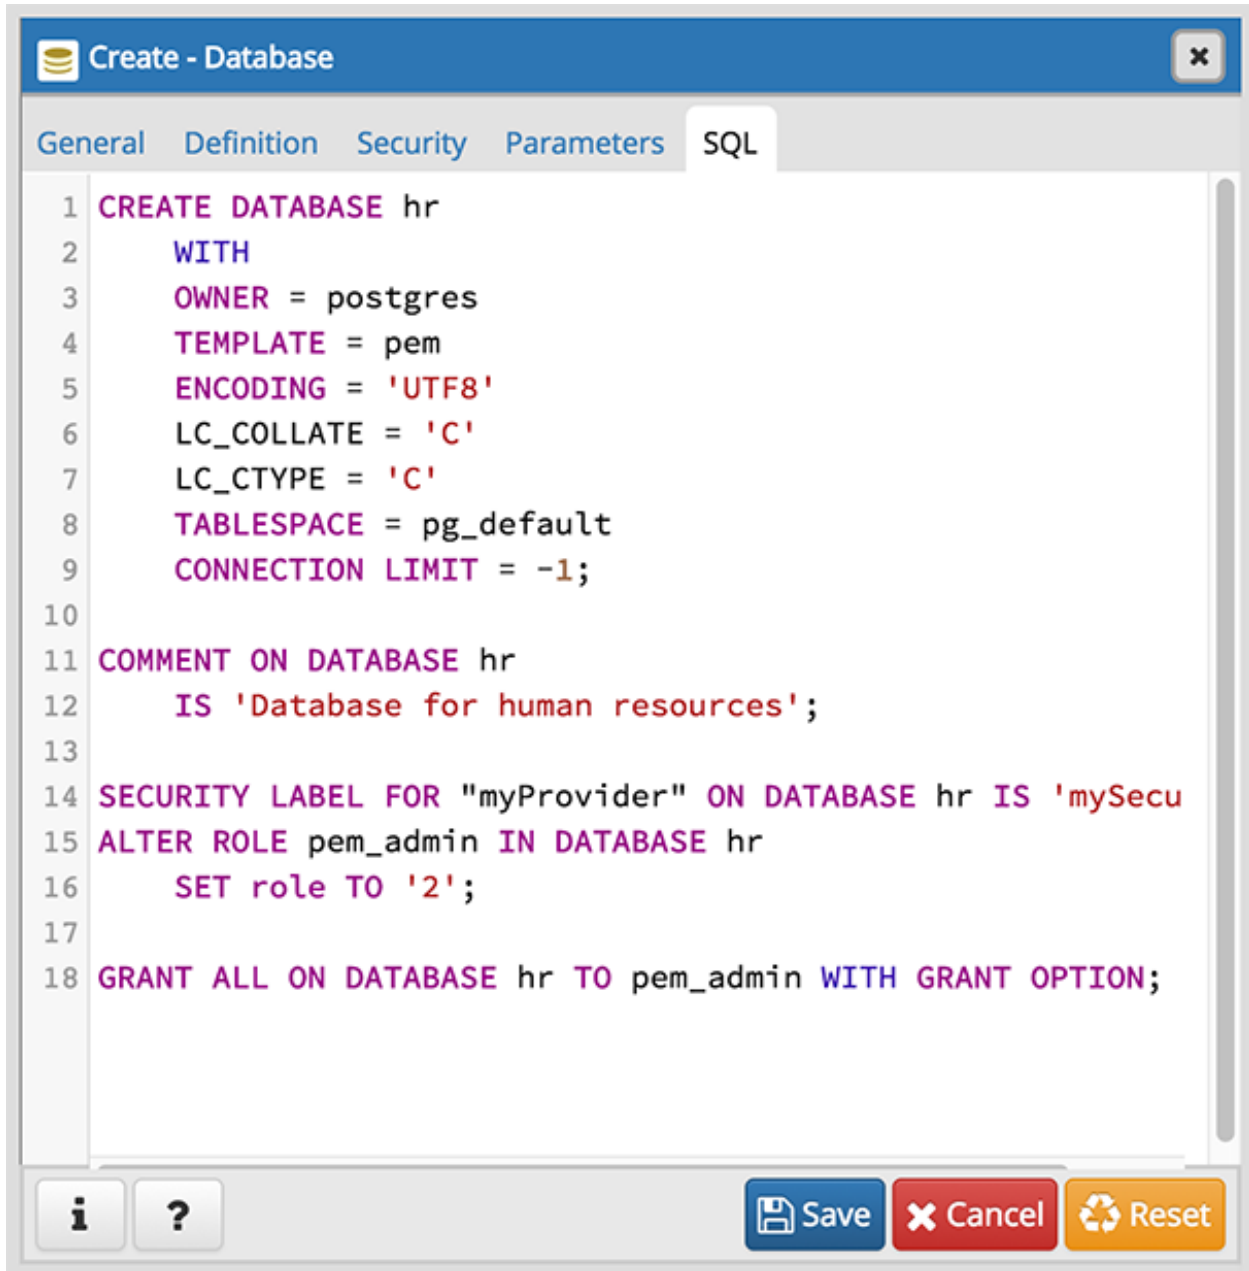
Follow these steps to add additional parameter value definitions; to discard a parameter, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Database* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Database* dialog:



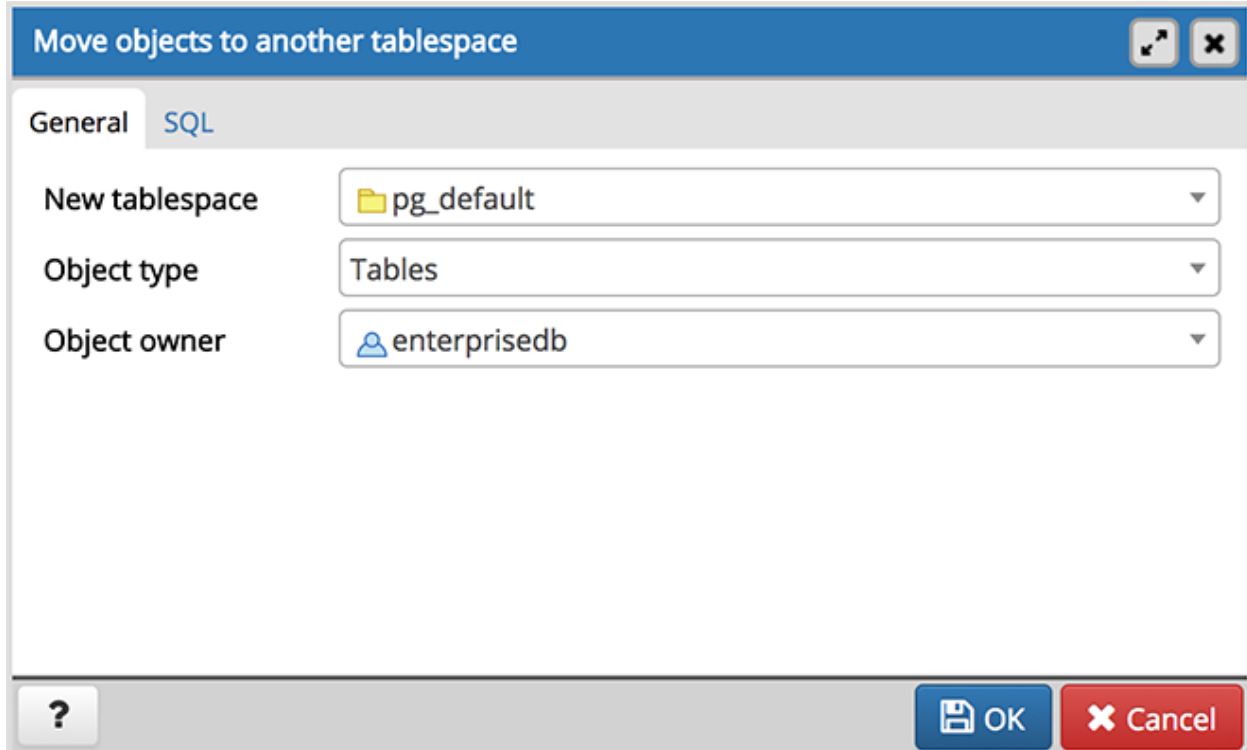
The example creates a database named *hr* that is owned by *postgres*. It allows unlimited connections, and is available to all authenticated users.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

2.2 Move Objects Dialog

Use the *Move Objects* dialog to to move database objects from one tablespace to another tablespace.

The *Move Objects* dialog organizes the movement of database objects with the *General* tab; the *SQL* tab displays the SQL code generated by dialog selections.



Use the fields in the *General* tab to identify the items that will be moved and the tablespace to which they will be moved:

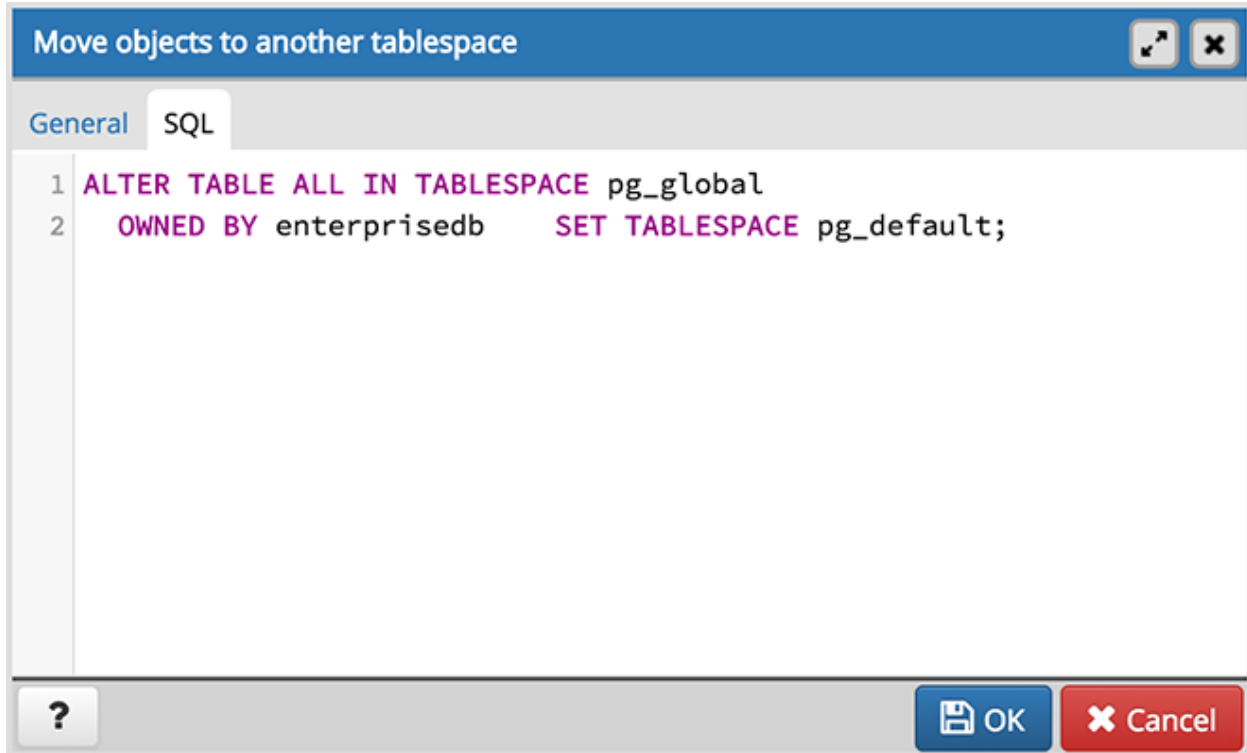
- Use the *New tablespace* drop-down listbox to select a pre-existing tablespace to which the object will be moved. (To create a tablespace, use the *Tablespace* dialog; access the dialog by right clicking *Tablespaces* in the *pgAdmin* tree control and selecting *Create Tablespace...* from the context-menu.)
- Use the *Object type* drop-down listbox to select from the following:
 - Select *All* to move all tables, indexes, and materialized views from the current tablespace (currently selected in the *pgAdmin* tree control) to the new tablespace.
 - Select *Tables* to move tables from the current tablespace to the new tablespace.
 - Select *Indexes* to move indexes from the current tablespace to the new tablespace.
 - Select *Materialized views* to move materialized views from the current tablespace to the new tablespace.
- Use the *Object owner* drop-down listbox to select the role that owns the objects selected in the *Object type* field. This field is optional.

Click the *SQL* tab to continue.

Your entries in the *Move Objects* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit the *General* tab to modify the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Move Objects* dialog:




The example shown demonstrates moving materialized views owned by Alice from tablespace *tblspace_01* to *tblspace_02*.

- Click the *Help* button (?) to access online help.
- Click the *OK* button to save work.
- Click the *Cancel* button to exit without saving work.

2.3 Login/Group Role Dialog

Use the *Login/Group Role* dialog to define a role. A role may be an individual user (with or without login privileges) or a group of users. Note that roles defined at the cluster level are shared by all databases in the cluster.

The *Login/Group Role* dialog organizes the creation and management of roles through the following dialog tabs: *General*, *Definition*, *Privileges*, *Parameters*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Login/Group Role' dialog box in pgAdmin 4. The dialog has a blue title bar with a close button. Below the title bar are tabs for 'General', 'Definition', 'Privileges', 'Membership', 'Parameters', and 'Security'. The 'General' tab is selected, showing a 'Name' text input field and a 'Comments' text area. Below the form fields are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left.

Use the fields on the *General* tab to identify the role.

- Use the *Name* field to provide the name of the role. The name will be displayed in the tree control.
- Provide a note about the role in the *Comments* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Login/Group Role' dialog box with the 'Definition' tab selected. The 'Password' field is highlighted in yellow and contains six dots. The 'Account expires' field contains the text 'YYYY-MM-DD HH:mm:ss Z' and a calendar icon. The 'Connection limit' field contains the value '-1'. At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left side of the bottom bar.

Use the *Definition* tab to set a password and configure connection rules:

- Provide a password that will be associated with the role in the *Password* field.
- Provide an expiration date for the password in the *Account Expires* field (the role does not expire). The expiration date is not enforced when a user logs in with a non-password-based authentication method.
- If the role is a login role, specify how many concurrent connections the role can make in the *Connection Limit* field. The default value (-1) allows unlimited connections.

Click the *Privileges* tab to continue.

The screenshot shows the 'Create - Login/Group Role' dialog box in pgAdmin 4, with the 'Privileges' tab selected. The dialog contains the following settings:

Privilege	Value
Can login?	No
Superuser	No
Create roles?	No
Create databases?	No
Update catalog?	No
Inherit rights from the parent roles?	Yes
Can initiate streaming replication and backups?	No

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left.

Use the *Privileges* tab to grant privileges to the role.

- Move the *Can login?* switch to the *Yes* position if the role has login privileges. The default value is *No*.
- Move the *Superuser* switch to the *Yes* position if the role is a superuser within the database. The default value is *No*.
- Move the *Create roles?* switch to the *Yes* position to specify whether a role is permitted to create roles. A role with this privilege can alter and drop roles. The default value is *No*.
- Move the *Create databases* switch to the *Yes* position to control whether a role can create databases. The default value is *No*.
- The *Update catalog?* switch is disabled until the role is given superuser privileges. Move the *Update catalogs?* switch to the *No* position to control whether a role can update catalogs. The default value is *Yes* when the *Superuser* switch is in the *Yes* position.
- Move the *Inherit rights from the parent roles?* switch to the *No* position if a role does not inherit privileges. The default value is *Yes*.


- Move the *Can initiate streaming replication and backups?* switch to the *Yes* position to control whether a role can initiate streaming replication or put the system in and out of backup mode. The default value is *No*.



The screenshot shows the 'Create - Login/Group Role' dialog box in pgAdmin 4, with the 'Membership' tab selected. The dialog has a blue title bar and a tabbed interface with tabs for 'General', 'Definition', 'Privileges', 'Membership', 'Parameters', 'Security', and 'SQL'. The 'Roles' field is a multi-select list containing one entry: 'pem_user' with a checked checkbox to its right and a delete 'x' icon to its left. Below the list, there is a text instruction: 'Select the checkbox for roles to include WITH ADMIN OPTION.' At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange), along with information and help icons on the left.

- Specify members of the role in the *Role Membership* field. Click inside the *Roles* field to select role names from a drop down list. Confirm each selection by checking the checkbox to the right of the role name; delete a selection by clicking the *x* to the left of the role name. Membership conveys the privileges granted to the specified role to each of its members.

Click the *Parameters* tab to continue.

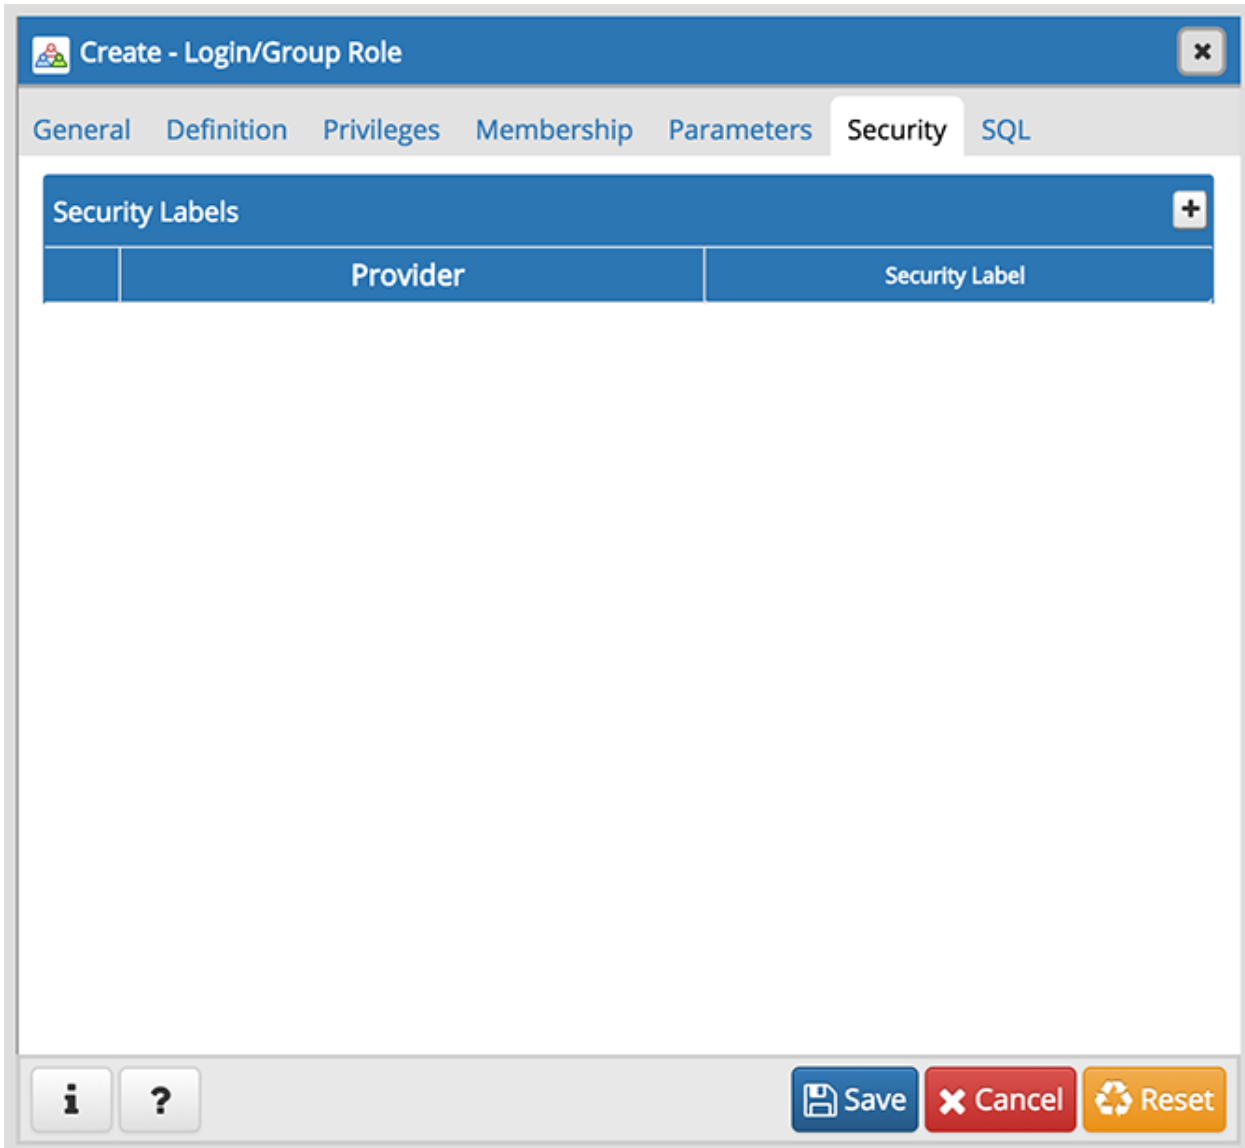


Use the fields on the *Parameters* tab to set session defaults for a selected configuration parameter when the role is connected to a specified database. This tab invokes the `ALTER ROLE... SET configuration_parameter` syntax. Click the *Add* icon (+) to assign a value for a parameter.

- Use the drop-down listbox in the *Name* field to select a parameter.
- Use the *Value* field to specify a value for the parameter.
- Use the drop-down listbox in the *Database* field to select a database.

Click the *Add* icon (+) to specify each additional parameter; to discard a parameter, click the trash icon to the left of the row and confirm the deletion in the *Delete Row* popup.

Click the *Security* tab to continue.



Use the *Security* tab to define security labels applied to the role. Click the *Add* icon (+) to add each security label selection.

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

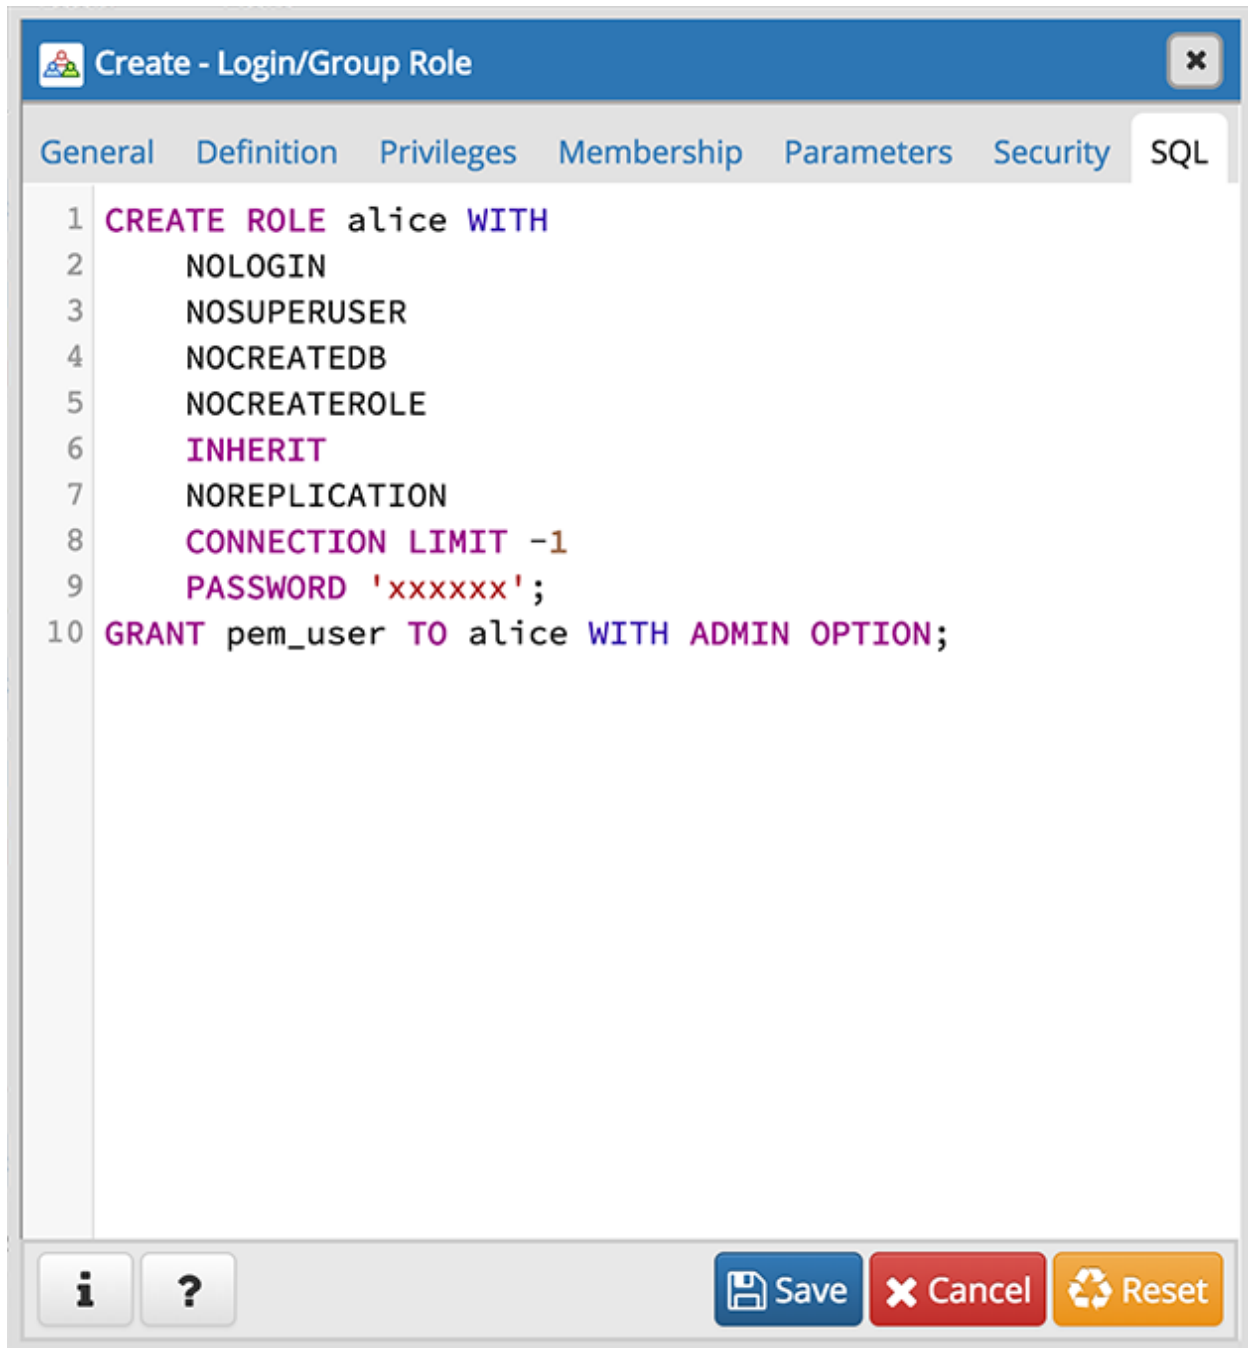
To discard a security label, click the trash icon to the left of the row and confirm the deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Login/Group Role* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Login/Group Role* dialog:



The example creates a login role named *alice* with *pem_user* privileges; the role can make unlimited connections to the server at any given time.

- Click the Info button (*i*) to access online SQL help.
- Click the Help button (?) to access the documentation for the dialog.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

MANAGING DATABASE OBJECTS

pgAdmin 4 provides simple but powerful dialogs that you can use to design and create database objects. Each dialog contains a series of tabs that you use to describe the object that will be created by the dialog; the SQL tab displays the SQL command that the server will execute when creating the object.

To access a dialog that allows you to create a database object, right-click on the object type in the pgAdmin tree control, and select the *Create* option for that object. For example, to create a new cast, right-click on the *Casts* node, and select *Create Cast...*

Contents:

3.1 Cast Dialog

Use the *Cast* dialog to define a cast. A cast specifies how to convert a value from one data type to another.

The *Cast* dialog organizes the development of a cast through the following dialog tabs: *General* and *Definition*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Cast' dialog box in pgAdmin 4. The dialog has a blue title bar with the text 'Create - Cast' and a close button. Below the title bar are three tabs: 'General', 'Definition', and 'SQL'. The 'General' tab is selected and contains two fields: 'Name' and 'Comment'. The 'Name' field is a text box that is disabled (grayed out). The 'Comment' field is a larger text area. At the bottom of the dialog, there are four buttons: an information icon (i), a help icon (?), a 'Save' button (blue with a floppy disk icon), a 'Cancel' button (red with an 'X' icon), and a 'Reset' button (orange with a refresh icon).

Use the fields in the *General* tab to identify the cast:

- The *Name* field is disabled. The name that will be displayed in the *pgAdmin* tree control is the *Source* type concatenated with the *Target* type, and is generated automatically when you make selections on the *Cast* dialog *Definition* tab.
- Store notes about the cast in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Cast' dialog box with the 'Definition' tab selected. The 'Source type', 'Target type', and 'Function' fields are drop-down menus, all currently displaying 'Select from the list'. The 'Context' field is a toggle switch with 'EXPLICIT' selected. The bottom of the dialog features an information icon, a help icon, and three action buttons: 'Save', 'Cancel', and 'Reset'.

Use the fields in the *Definition* tab to define parameters:

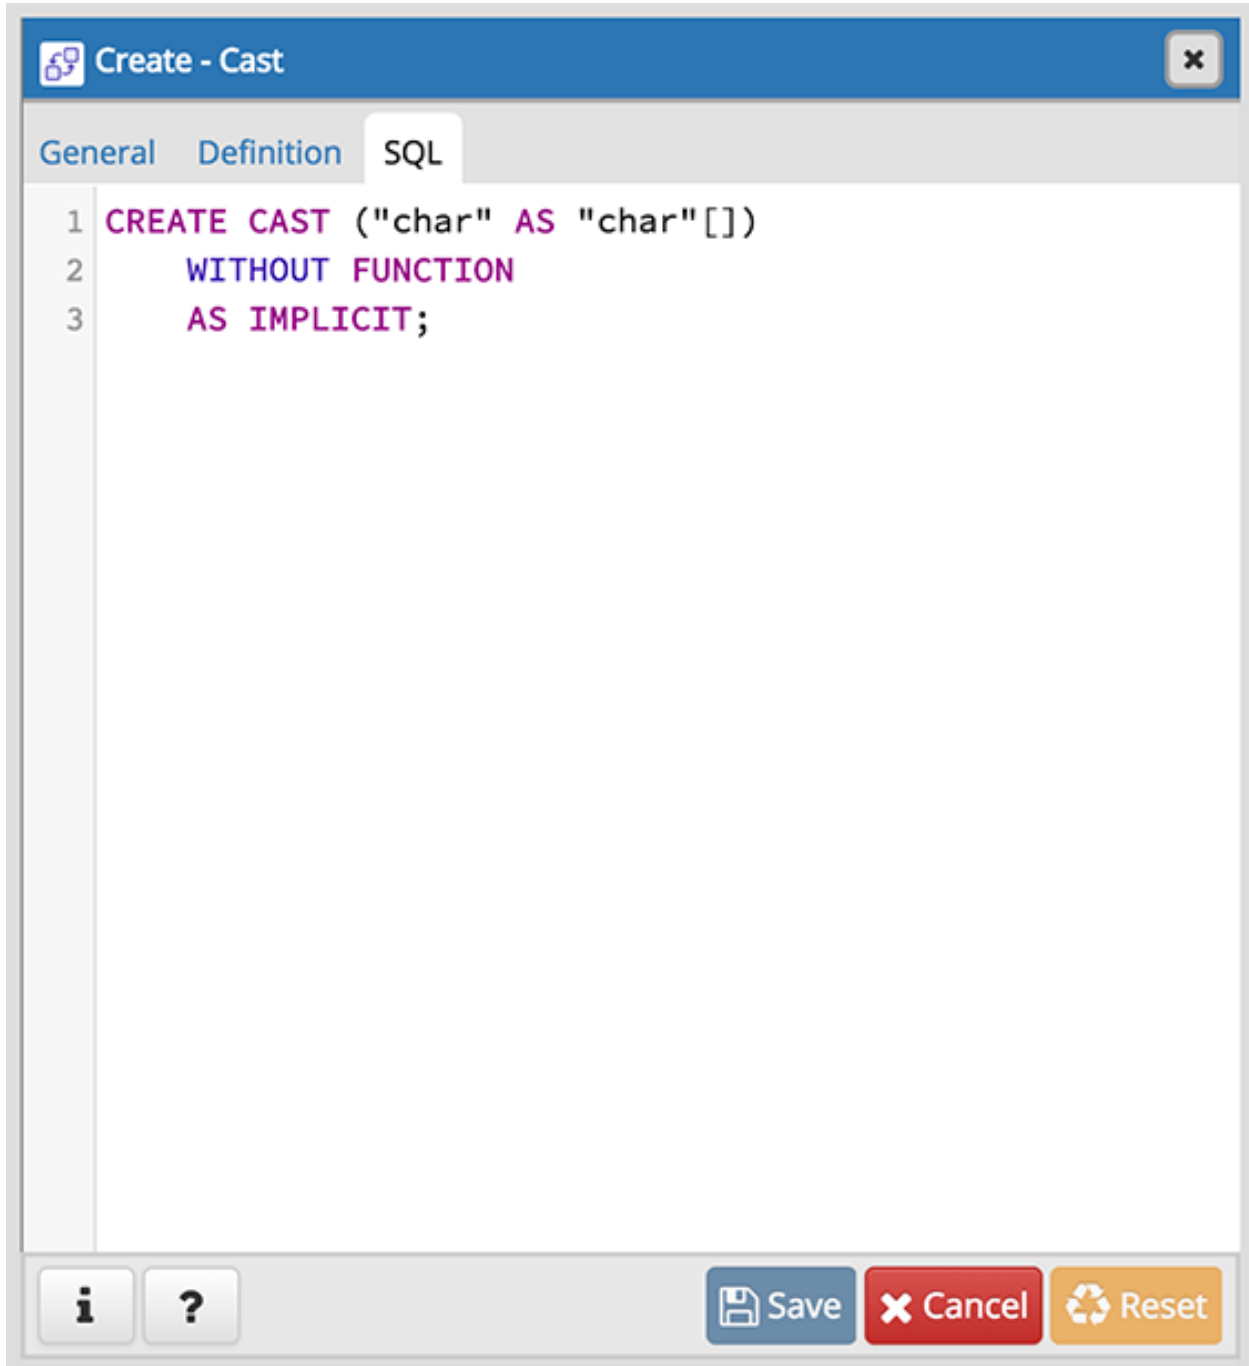
- Use the drop-down listbox next to *Source type* to select the name of the source data type of the cast.
- Use the drop-down listbox next to *Target type* to select the name of the target data type of the cast.
- Use the drop-down listbox next to *Function* to select the function used to perform the cast. The function's result data type must match the target type of the cast.
- Move the *Context* switch to the *Implicit* position if the cast is implicit. By default, a cast can be invoked only by an explicit cast request. If the cast is marked *Implicit* then it can be invoked implicitly in any context, whether by assignment or internally in an expression.

Click the *SQL* tab to continue.

Your entries in the *Cast* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Cast* dialog:



The cast uses a function named *int4(bigint)* to convert a bigint data type to an integer.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.

- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.2 Collation Dialog

Use the *Collation* dialog to define a collation. A collation is an SQL schema object that maps a SQL name to operating system locales. To create a collation, you must have a `CREATE` privilege on the destination schema.

The *Collation* dialog organizes the development of a collation through the following dialog tabs: *General* and *Definition*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Collation' dialog box in pgAdmin 4. The dialog has a blue title bar with the text 'Create - Collation' and a close button. Below the title bar are three tabs: 'General', 'Definition', and 'SQL'. The 'General' tab is selected. The form contains the following fields:

- Name:** A text input field with a vertical cursor.
- Owner:** A dropdown menu showing 'postgres' with a user icon and a close button.
- Schema:** A dropdown menu showing 'public' with a schema icon and a close button.
- Comment:** A large text area for entering notes.

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two small buttons with an 'i' and a '?' icon.

Use the fields in the *General* tab to identify the collation:

- Use the *Name* field to provide a name for the collation. The collation name must be unique within a schema. The name will be displayed in the *pgAdmin* tree control.
- Select the name of the owner from the drop-down listbox in the *Owner* field.
- Select the name of the schema in which the collation will reside from the drop-down listbox in the *Schema* field.
- Store notes about the collation in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Collation' dialog box with the 'Definition' tab selected. The dialog has a blue title bar with the text 'Create - Collation' and a close button. Below the title bar are three tabs: 'General', 'Definition', and 'SQL'. The 'Definition' tab contains four input fields: 'Copy collation' (a dropdown menu with the text 'Select from the list'), 'Locale' (a text input field), 'LC_COLLATE' (a text input field), and 'LC_TYPE' (a text input field). At the bottom of the dialog are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also two small buttons with an 'i' and a '?' icon on the left side of the bottom bar.

Use the fields in the *Definition* tab to specify the operating system locale settings:

- Use the drop-down listbox next to *Copy collation* to select the name of an existing collation to copy. The new collation will have the same properties as the existing one, but will be an independent object. If you choose to copy an existing collation, you cannot modify the collation properties displayed on this tab.
- Use the *Locale* field to specify a locale; a locale specifies language and language formatting characteristics. If you specify this, you cannot specify either of the following parameters. To view a list of locales supported by your Linux system use the command `locale -a`.
- Use the *LC_COLLATE* field to specify a locale with specified string sort order. The locale must be applicable to the current database encoding. (See `CREATE DATABASE` for details.)

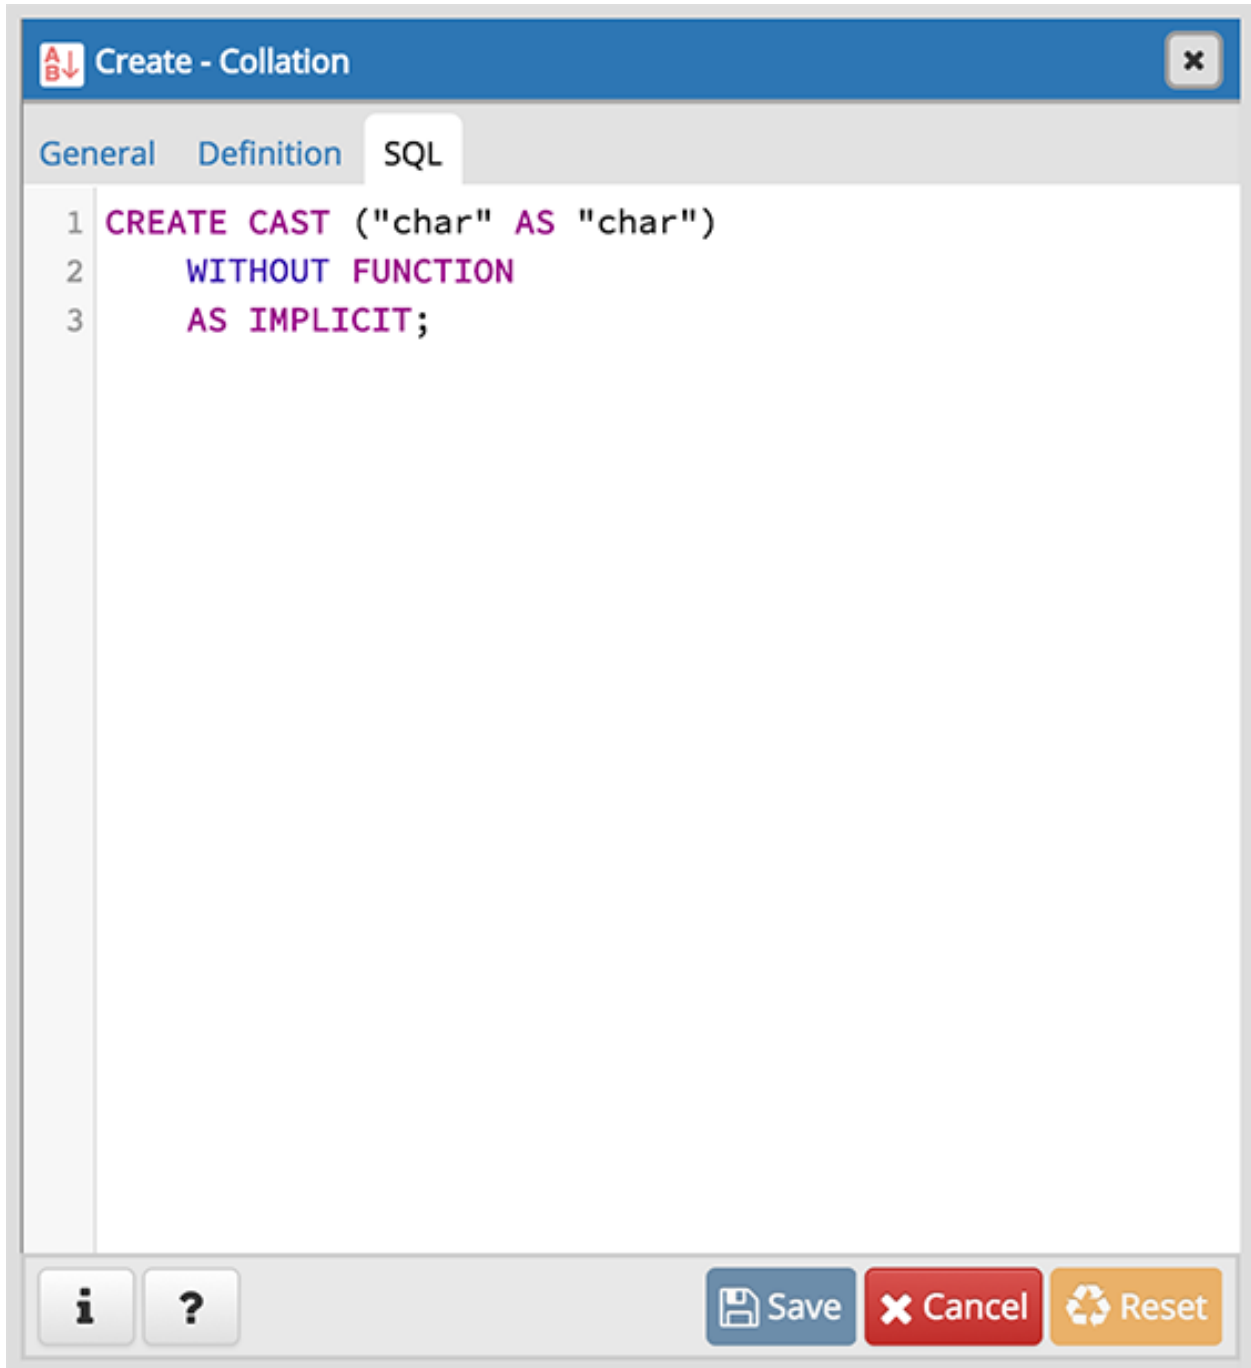
- Use the `LC_CTYPE` field to specify a locale with specified character classification. The locale must be applicable to the current database encoding. (See `CREATE DATABASE` for details.)

Click the *SQL* tab to continue.

Your entries in the *Collation* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Collation* dialog:



The example shown demonstrates creating a collation named *french* that uses the rules specified for the locale,

fr_FR.utf8. The collation is owned by **postgres*.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation. For more information about setting a locale, see Chapter 22.1 Locale Support of the PostgreSQL core documentation:


<http://www.postgresql.org/docs/9.5/static/locale.html>

- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.3 Domain Dialog

Use the *Domain* dialog to define a domain. A domain is a data type definition that may constrain permissible values. Domains are useful when you are creating multiple tables that contain comparable columns; you can create a domain that defines constraints that are common to the columns and re-use the domain definition when creating the columns, rather than individually defining each set of constraints.

The *Domain* dialog organizes the development of a domain through the following tabs: *General*, *Definition*, *Constraints*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Domain' dialog box in pgAdmin 4. The dialog has a blue title bar with a home icon and a close button. Below the title bar are five tabs: 'General', 'Definition', 'Constraints', 'Security', and 'SQL'. The 'General' tab is selected. It contains four fields: 'Name' (a text input field), 'Owner' (a dropdown menu with 'postgres' selected), 'Schema' (a dropdown menu with 'public' selected), and 'Comment' (a large text area). At the bottom of the dialog are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also two small buttons with an 'i' and a '?' icon on the left side of the bottom bar.

Use the fields on the *General* tab to identify a domain:

- Use the *Name* field to add a descriptive name for the domain. The name will be displayed in the *pgAdmin* tree control.
- Use the drop-down listbox next to *Owner* to select a role that will own the domain.
- Select the name of the schema in which the domain will reside from the drop-down listbox in the *Schema* field.
- Store notes about the domain in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Domain' dialog box with the 'Definition' tab selected. The fields are as follows:

- Base type:** A dropdown menu with the text 'Select from the list'.
- Length:** A text input field.
- Precision:** A text input field.
- Default:** A text input field with the placeholder text 'Enter an expression or a value.'
- Not Null?:** A toggle switch currently set to 'No'.
- Collation:** A dropdown menu with the text 'Select from the list'.

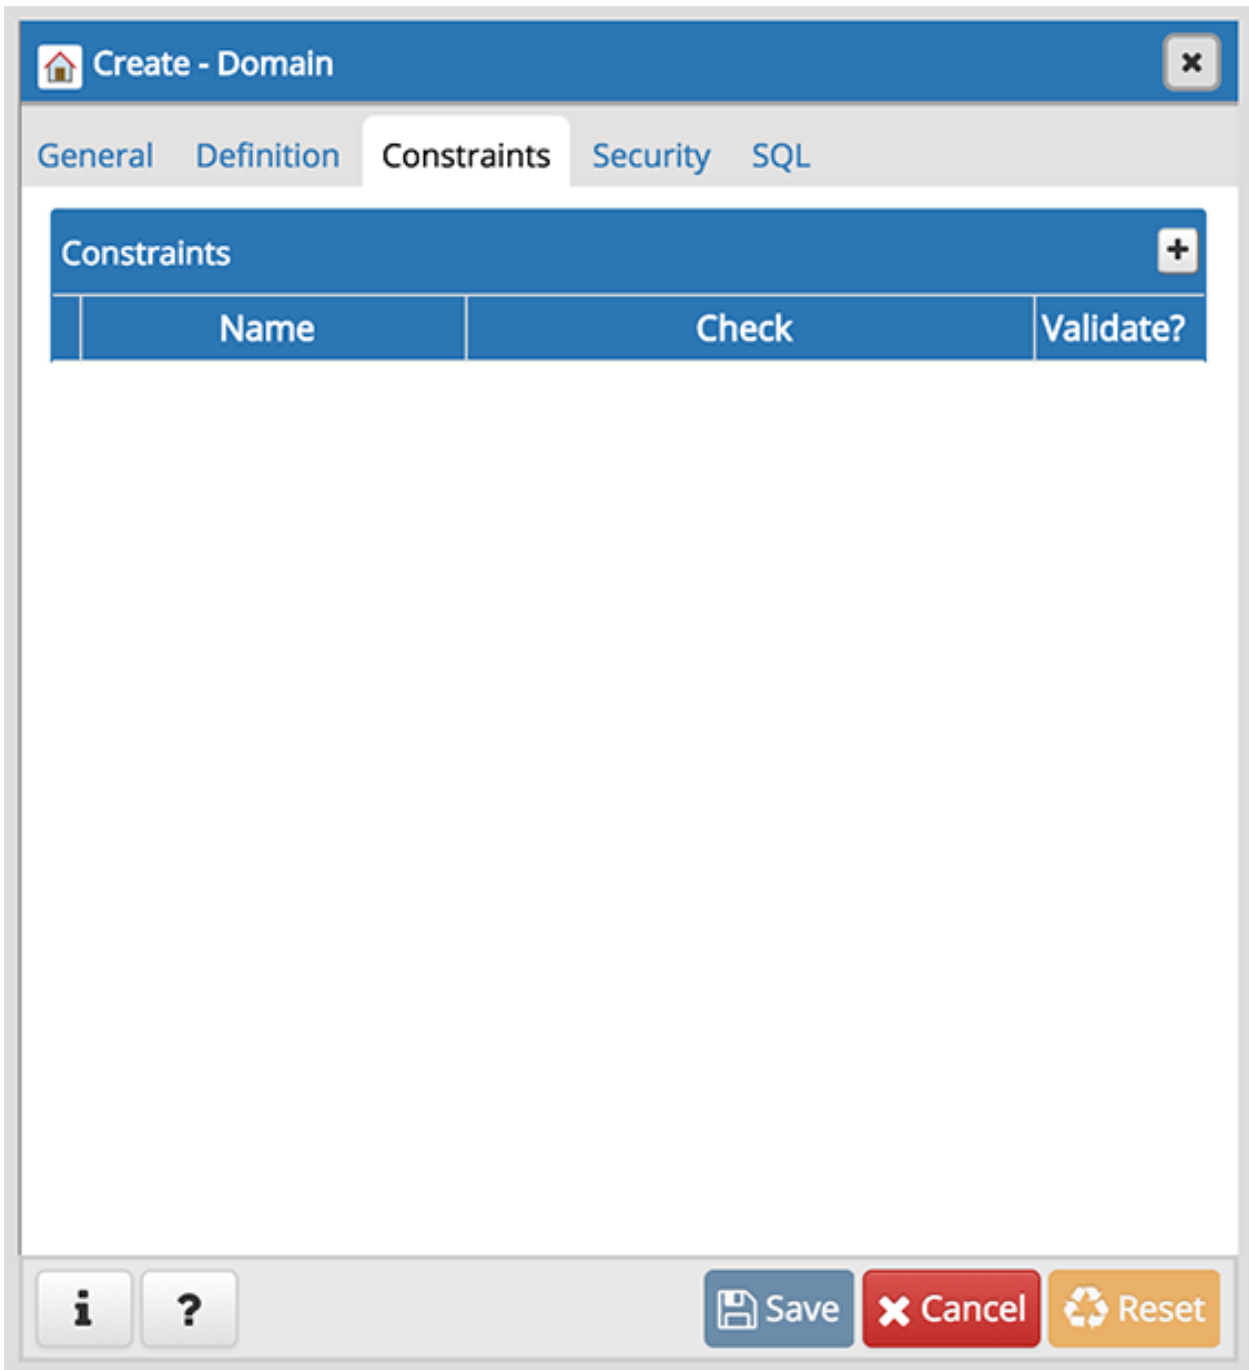
At the bottom of the dialog, there are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields in the *Definition* tab to describe the domain:

- Use the drop-down listbox next to *Base type* to specify a data type.
- Use the context-sensitive *Length* field to specify a numeric length for a numeric type.
- Use the context-sensitive *Precision* field to specify the total count of significant digits for a numeric type.
- Specify a default value for the domain data type in the *Default* field. The data type of the default expression must match the data type of the domain. If no default value is specified, then the default value is the null value.
- Move the *Not Null* switch to specify the values of this domain are prevented from being null.

- Use the drop-down listbox next to *Collation* to apply a collation cast. If no collation is specified, the underlying data type's default collation is used. The underlying type must be collatable if COLLATE is specified.

Click the *Constraints* tab to continue.



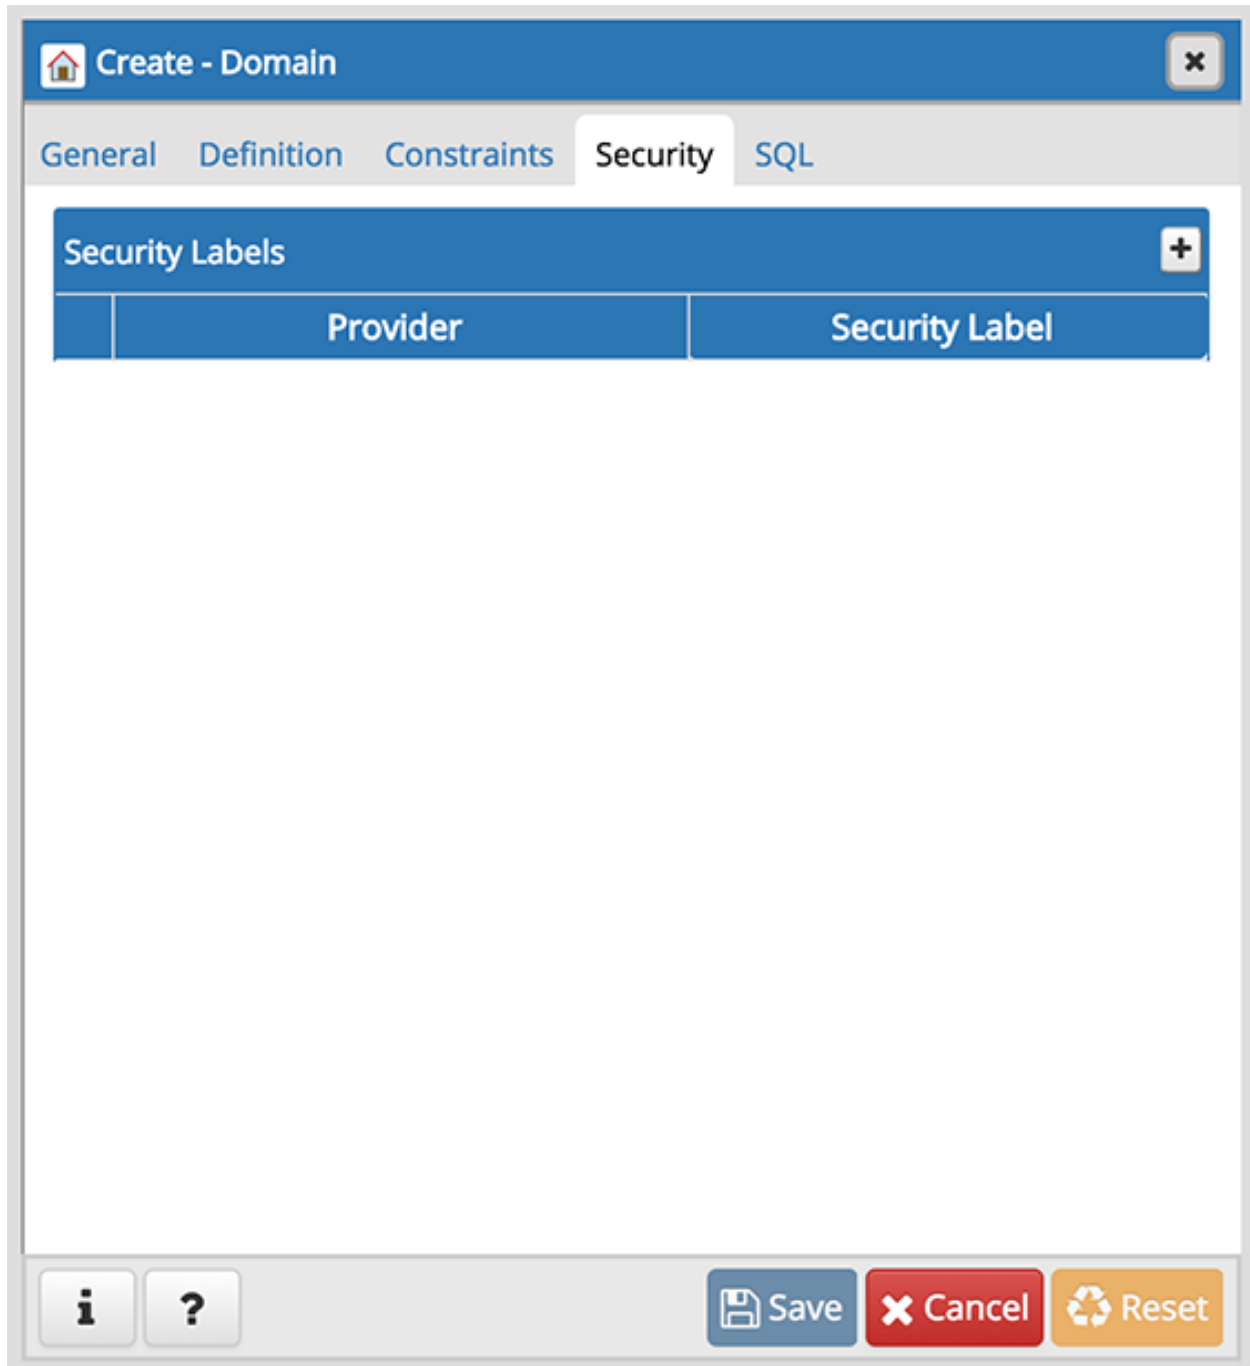
Use the fields in the *Constraints* tab to specify rules for the domain. Click the *Add* icon (+) to set constraints:

- Use the *Name* field to specify a name for the constraint.
- Use the *Check* field to provide an expression for the constraint.
- Use the *Validate* checkbox to determine whether the constraint will be validated. The default checkbox is checked and sets a validation requirement.

A CHECK clause specifies an integrity test which values of the domain must satisfy. Each constraint must be an expression that produces a Boolean result. Use the key word VALUE to refer to the value being tested. Expressions evaluating to TRUE or UNKNOWN succeed. If the expression produces a FALSE result, an error is reported and the value is not allowed to be converted to the domain type. A CHECK expression cannot contain subqueries nor refer to variables other than VALUE. If a domain has multiple CHECK constraints, they will be tested in alphabetical order by name.

Click the *Add* icon (+) to set additional constraints; to discard a constraint, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Security* tab to continue.



Use the *Security Labels* panel to assign security labels. Click the *Add* icon (+) to add a label:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

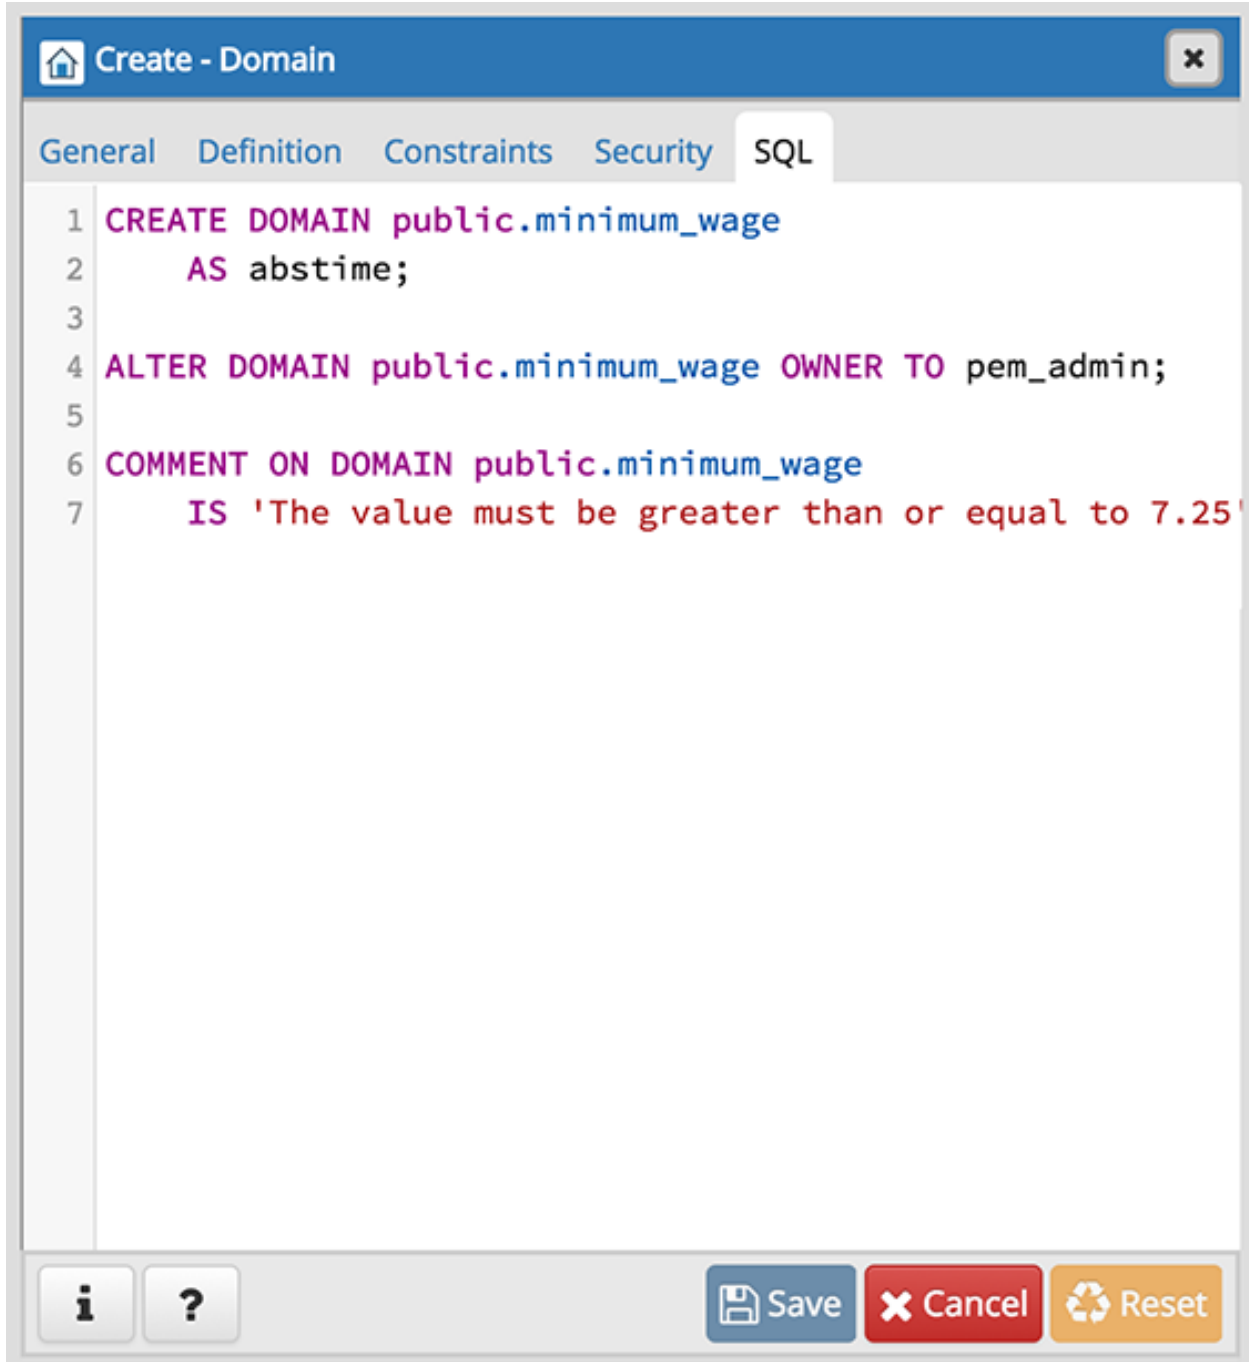
Click the *Add* icon (+) to specify each additional label; to discard a label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Domain* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by selections made in the *Domain* dialog:



The example shown demonstrates creating a domain named *minimum-wage* that confirms that the value entered is greater than or equal to 7.25.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.4 Domain Constraints Dialog

Use the *Domain Constraints* dialog to create or modify a domain constraint. A domain constraint confirms that the values provided for a domain meet a defined criteria. The *Domain Constraints* dialog implements options of the ALTER DOMAIN command.

The *Domain Constraints* dialog organizes the development of a domain constraint through the following dialog tabs: *General* and *Definition*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows a dialog window titled "Create - Domain Constraints". It has three tabs: "General", "Definition", and "SQL". The "General" tab is selected. In the "General" tab, there is a "Name" field containing the text "timesheets" and a "Comment" text area below it. At the bottom of the dialog, there are three buttons: "Save" (blue), "Cancel" (red), and "Reset" (orange). To the left of these buttons are two small icons: an information icon (i) and a help icon (?).

Use the fields in the *General* tab to identify the domain constraint:

- Use the *Name* field to add a descriptive name for the constraint. The name will be displayed in the *pgAdmin* tree control.
- Store notes about the constraint in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Domain Constraints' dialog box with the 'Definition' tab selected. The 'Check' field contains the SQL expression 'VALUE = 'Friday''. The 'Validate?' checkbox is checked. The dialog has tabs for 'General', 'Definition', and 'SQL'. At the bottom are buttons for 'Save', 'Cancel', and 'Reset'.

Use the fields in the *Definition* tab to define the domain constraint:

- Use the *Check* field to provide a CHECK expression. A CHECK expression specifies a constraint that the domain must satisfy. A constraint must produce a Boolean result; include the key word VALUE to refer to the value being tested. Only those expressions that evaluate to TRUE or UNKNOWN will succeed. A CHECK

expression cannot contain subqueries or refer to variables other than VALUE. If a domain has multiple CHECK constraints, they will be tested in alphabetical order.

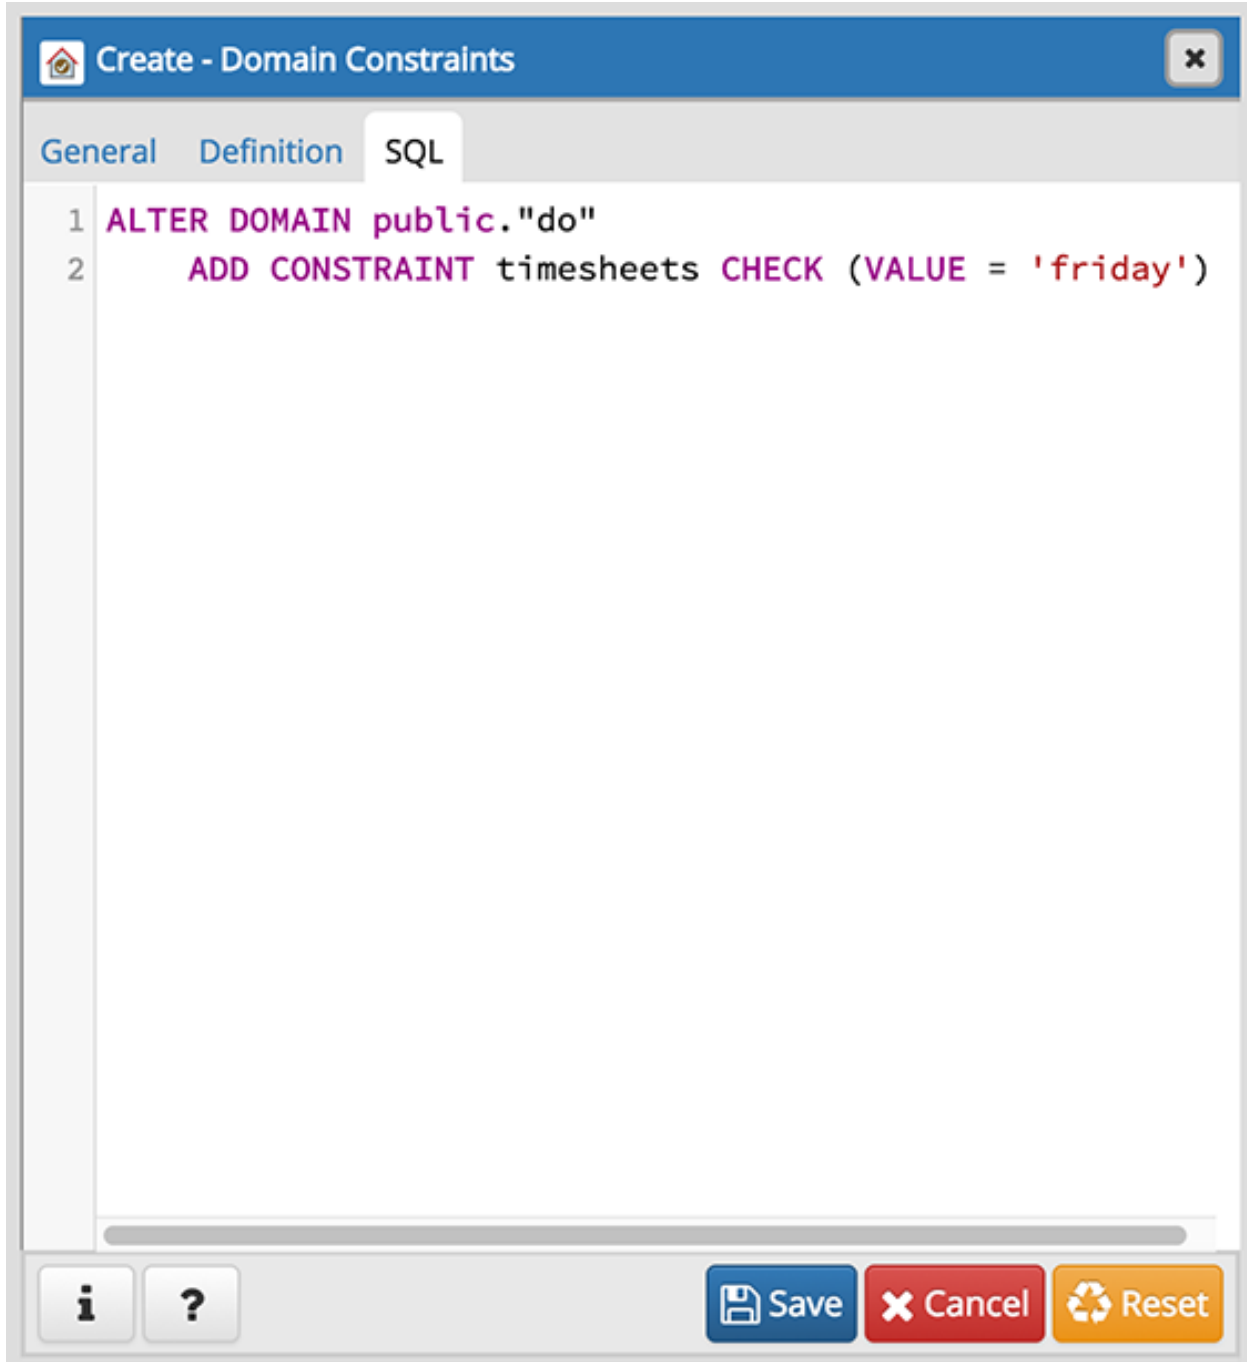
- Move the *Validate?* switch to the *No* position to mark the constraint NOT VALID. If the constraint is marked NOT VALID, the constraint will not be applied to existing column data. The default value is *Yes*.

Click the *SQL* tab to continue.

Your entries in the *Domain Constraints* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Domain Constraints* dialog:



The example shown demonstrates creating a domain constraint on the domain *timesheets* named *weekday*. It constrains a value to equal *Friday*.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.5 Event Trigger Dialog

Use the *Domain Trigger* dialog to define an event trigger. Unlike regular triggers, which are attached to a single table and capture only DML events, event triggers are global to a particular database and are capable of capturing DDL events. Like regular triggers, event triggers can be written in any procedural language that includes event trigger support, or in C, but not in SQL.

The *Domain Trigger* dialog organizes the development of a event trigger through the following dialog tabs: *General*, *Definition*, and *Security Labels*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the "Create - Event Trigger" dialog box. The "General" tab is active, showing the following fields:


- Name:** accounts
- Owner:** postgres
- Comment:** (empty text area)

At the bottom of the dialog, there are three buttons: "Save" (blue), "Cancel" (red), and "Reset" (orange). There are also information and help icons on the left.

Use the fields in the *General* tab to identify the event trigger:

- Use the *Name* field to add a descriptive name for the event trigger. The name will be displayed in the *pgAdmin* tree control.
- Use the drop-down listbox next to *Owner* to specify the owner of the event trigger.
- Store notes about the event trigger in the *Comment* field.

Click the *Definition* tab to continue.



The screenshot shows the 'Create - Event Trigger' dialog box with the 'Definition' tab selected. The dialog has a title bar with a close button. Below the title bar are four tabs: 'General', 'Definition', 'Security', and 'SQL'. The 'Definition' tab is active and contains the following fields:

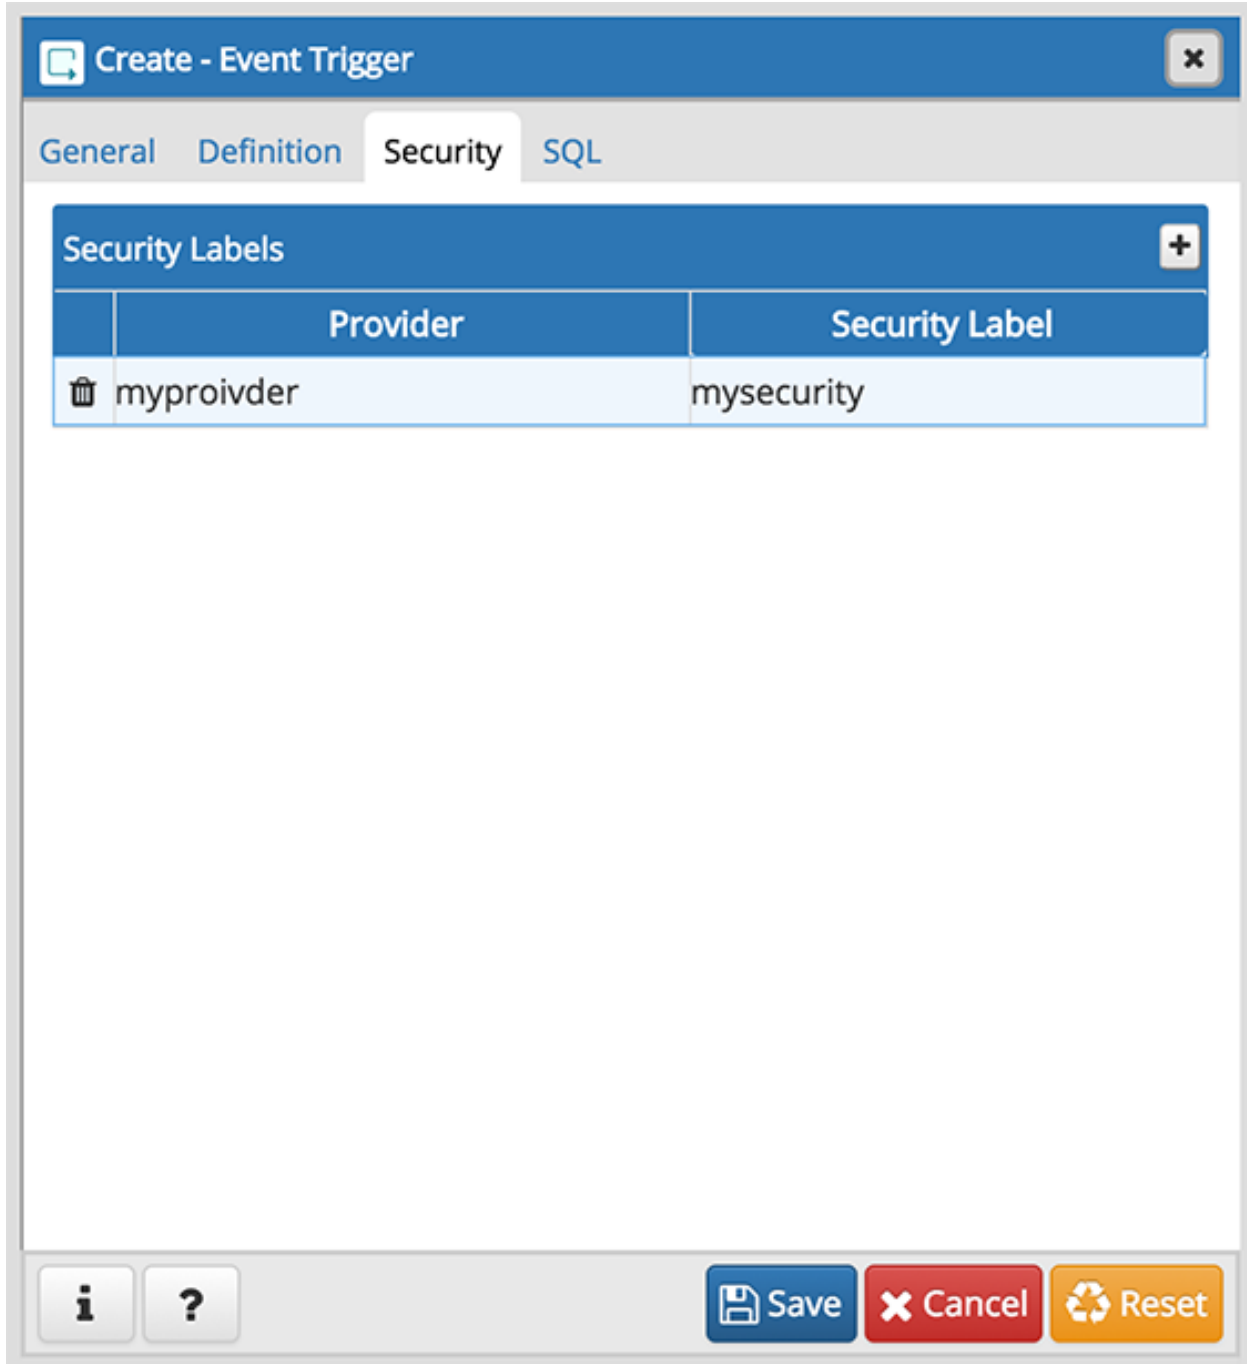
- Enabled status:** Radio buttons for 'Enable' (selected), 'Disable', 'Replica', and 'Always'.
- Trigger function:** A text box containing 'public.fnnn' and a dropdown arrow.
- Events:** Radio buttons for 'DDL COMMAND START' (selected), 'DDL COMMAND END', and 'SQL DROP'.
- When:** A large empty text area for defining the trigger condition.

At the bottom of the dialog are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left side of the bottom bar.

Use the fields in the *Definition* tab to define the event trigger:

- Select a radio button in the *Enabled Status* field to specify a status for the trigger: *Enable Disable, Replica Always*.
- Use the drop-down listbox next to *Trigger function* to specify an existing function. A trigger function takes an empty argument list, and returns a value of type `event_trigger`.
- Select a radio button in the *Events* field to specify when the event trigger will fire: *DDL COMMAND START, DDL COMMAND END, or SQL DROP*.
- Use the *When* field to write a condition for the event trigger that must be satisfied before the event trigger can execute.

Click the *Security Labels* tab to continue.



Use the *Security* tab to define security labels applied to the trigger. Click the *Add* icon (+) to add each security label.

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

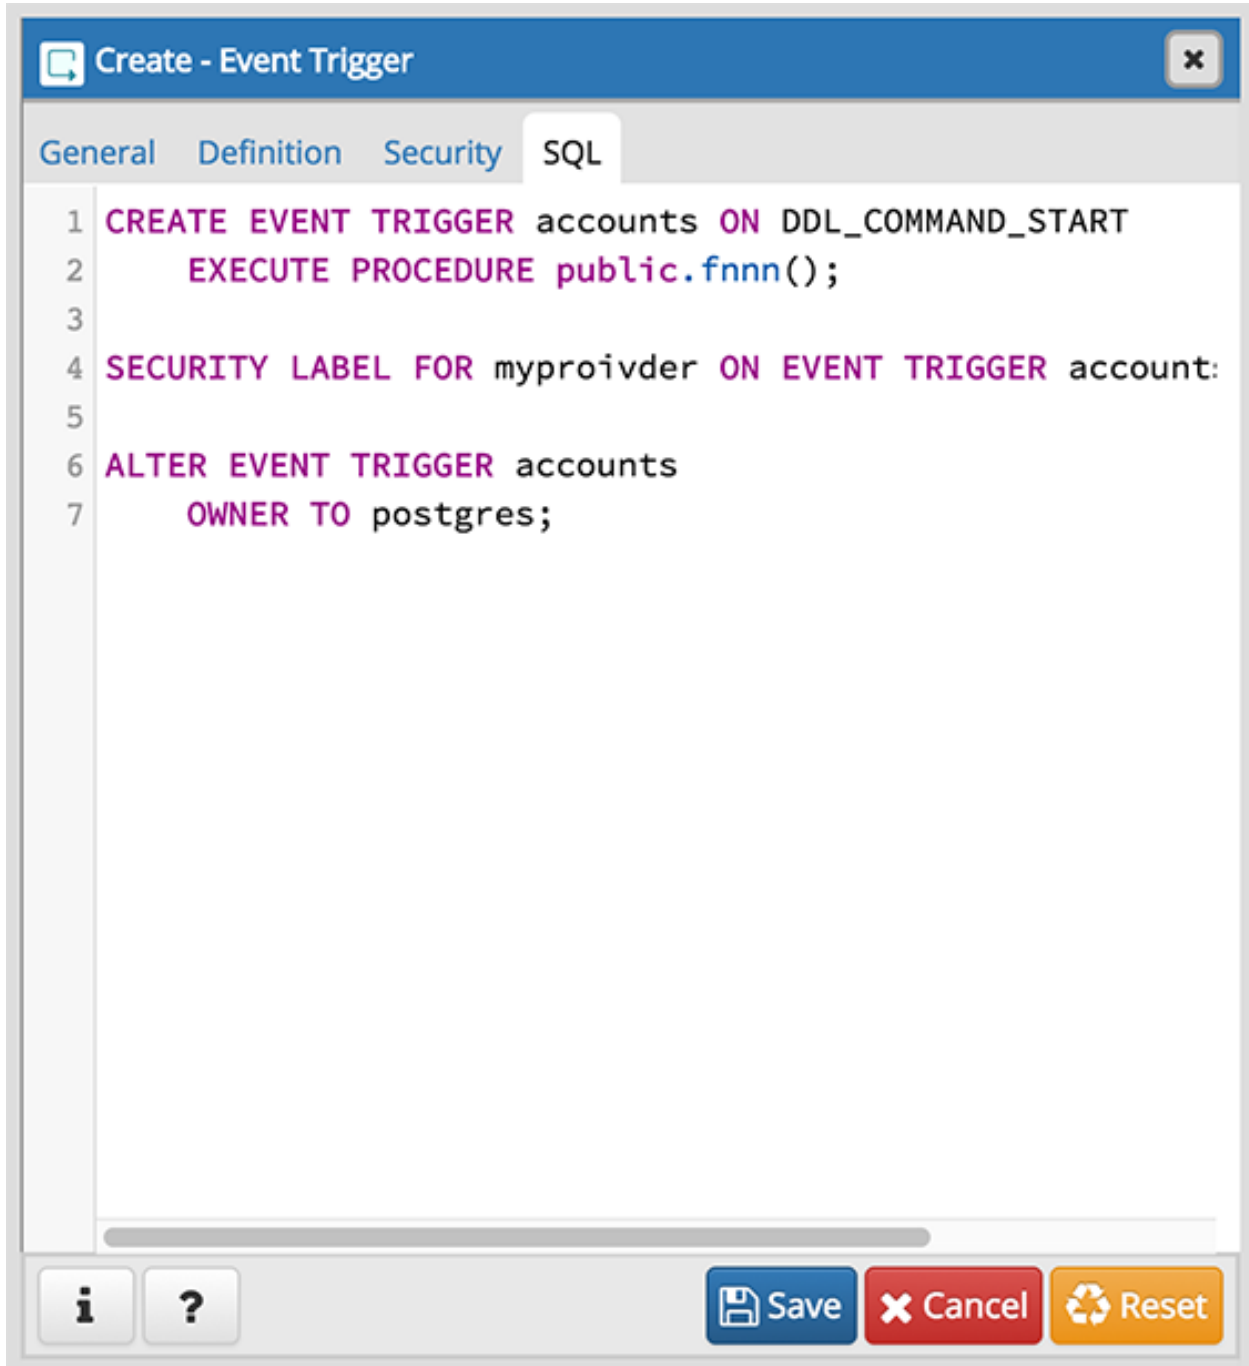
Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Domain Trigger* dialog generate a generate a SQL command. Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Domain Trigger* dialog:



The command creates an event trigger named *accounts* that invokes the procedure named *acct_due*.


- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.

- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.6 Extension Dialog

Use the *Extension* dialog to install a new extension into the current database. An extension is a collection of SQL objects that add targeted functionality to your Postgres installation. The *Extension* dialog adds the functionality of an extension to the current database only; you must register the extension in each database that use the extension. Before you load an extension into a database, you should confirm that any pre-requisite files are installed.

The *Extension* dialog allows you to implement options of the CREATE EXTENSION command through the following dialog tabs: *General* and *Definition*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Extension' dialog box. The title bar reads 'Create - Extension' with a close button. Below the title bar are three tabs: 'General', 'Definition', and 'SQL'. The 'General' tab is selected. In the 'General' tab, there is a 'Name' field with a drop-down listbox containing the text 'Select from the list'. Below the 'Name' field is a 'Comment' field, which is a large, empty text area. At the bottom of the dialog, there are five buttons: an information button (i), a help button (?), a 'Save' button, a 'Cancel' button, and a 'Reset' button.

Use the fields in the *General* tab to identify an extension:

- Use the drop-down listbox in the *Name* field to select the extension. Each extension must have a unique name.
- Store notes about the extension in the *Comment* field.

Click the *Definition* tab to continue.



The screenshot shows the 'Create - Extension' dialog box. The title bar reads 'Create - Extension' with a close button. Below the title bar are three tabs: 'General', 'Definition', and 'SQL'. The 'Definition' tab is selected. Under the 'Definition' tab, there are two labels: 'Schema' and 'Version'. Each label is followed by a drop-down menu containing the text 'Select from the list'. At the bottom of the dialog, there are four buttons: an information icon (i), a help icon (?), a 'Save' button (blue with a floppy disk icon), a 'Cancel' button (red with an 'X' icon), and a 'Reset' button (orange with a refresh icon).

Use the *Definition* tab to select the *Schema* and *Version*:

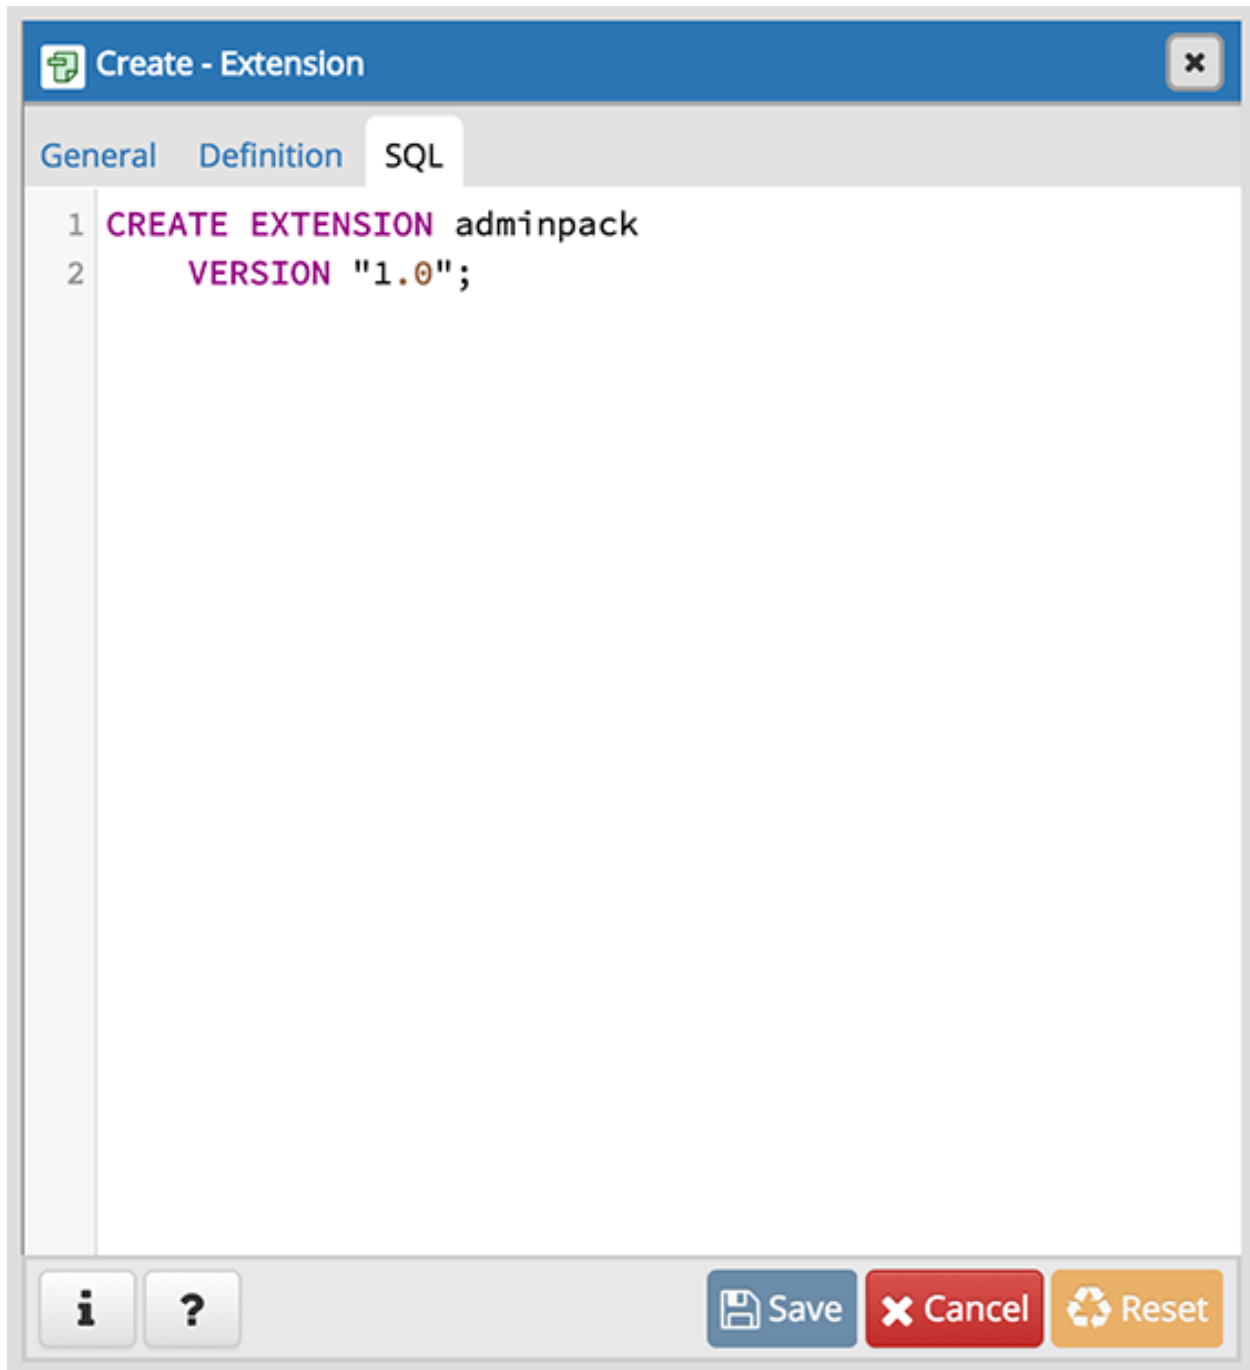
- Use the drop-down listbox next to *Schema* to select the name of the schema in which to install the extension's objects.
- Use the drop-down listbox next to *Version* to select the version of the extension to install.

Click the *SQL* tab to continue.

Your entries in the *Extension* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Extension* dialog:



The command creates the *chkpass* extension in the *public* schema. It is version *1.0* of *chkpass*.


- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.7 Foreign Data Wrapper Dialog

Use the *Foreign Data Wrapper* dialog to create or modify a foreign data wrapper. A foreign data wrapper is an adapter between a Postgres database and data stored on another data source.

You must be a superuser to create a foreign data wrapper.

The *Foreign Data Wrapper* dialog organizes the development of a foreign data wrapper through the following dialog tabs: *General*, *Definition*, *Options*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Foreign Data Wrapper' dialog box. The title bar is blue with a close button (X) on the right. Below the title bar are five tabs: 'General', 'Definition', 'Options', 'Security', and 'SQL'. The 'General' tab is selected and contains the following fields:

- Name:** A text input field with a vertical cursor.
- Owner:** A dropdown menu showing 'postgres' with a user icon on the left and a downward arrow on the right.
- Comment:** A large, empty text area with a small double-slash icon in the bottom right corner.

At the bottom of the dialog, there are three main buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two smaller buttons with an 'i' and a '?' icon.

Use the fields in the *General* tab to identify the foreign data wrapper:

- Use the *Name* field to add a descriptive name for the foreign data wrapper. A foreign data wrapper name must be unique within the database. The name will be displayed in the *pgAdmin* tree control.
- Use the drop-down listbox next to *Owner* to select the name of the role that will own the foreign data wrapper.
- Store notes about the foreign data wrapper in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Foreign Data Wrapper' dialog box in pgAdmin 4. The dialog has a blue title bar with a close button. Below the title bar are tabs for 'General', 'Definition', 'Options', 'Security', and 'SQL'. The 'Definition' tab is active, showing two drop-down menus: 'Handler' and 'Validator', both currently displaying 'Select from the list'. At the bottom of the dialog are buttons for 'i', '?', 'Save', 'Cancel', and 'Reset'.

Use the fields in the *Definition* tab to set parameters:

- Select the name of the handler function from the drop-down listbox in the *Handler* field. This is the name of an existing function that will be called to retrieve the execution functions for foreign tables.

- Select the name of the validator function from the drop-down listbox in the *Validator* field. This is the name of an existing function that will be called to check the generic options given to the foreign data wrapper, as well as options for foreign servers, user mappings and foreign tables using the foreign data wrapper.

Click the *Options* tab to continue.


Options		+
Option	Value	

Use the fields in the *Options* tab to specify options:

- Click the the *Add* icon (+) button to add an option/value pair for the foreign data wrapper. Supported option/value pairs will be specific to the selected foreign data wrapper.
- Specify the option name in the *Option* field and provide a corresponding value in the *Value* field.

Click the *Add* icon (+) to specify each additional pair; to discard an option, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Security* tab to continue.



The screenshot shows the 'Create - Foreign Data Wrapper' dialog box in pgAdmin 4, with the 'Security' tab selected. The dialog has a title bar with a close button (X) and a menu icon. Below the title bar are tabs for 'General', 'Definition', 'Options', 'Security', and 'SQL'. The 'Security' tab is active, showing a 'Privileges' section with a blue header and a '+' icon. Below the header is a table with three columns: 'Grantee', 'Privileges', and 'Grantor'. The table is currently empty. At the bottom of the dialog are buttons for 'Save', 'Cancel', and 'Reset', along with information (i) and help (?) icons.

Privileges		
Grantee	Privileges	Grantor

Use the *Security* tab to assign security privileges. Click the *Add* icon (+) to assign a set of privileges.

- Select the name of the role from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privileges to the specified user.
- Select the name of the role granting the privileges from the drop-down listbox in the *Grantor* field. The default

grantor is the owner of the foreign data wrapper.

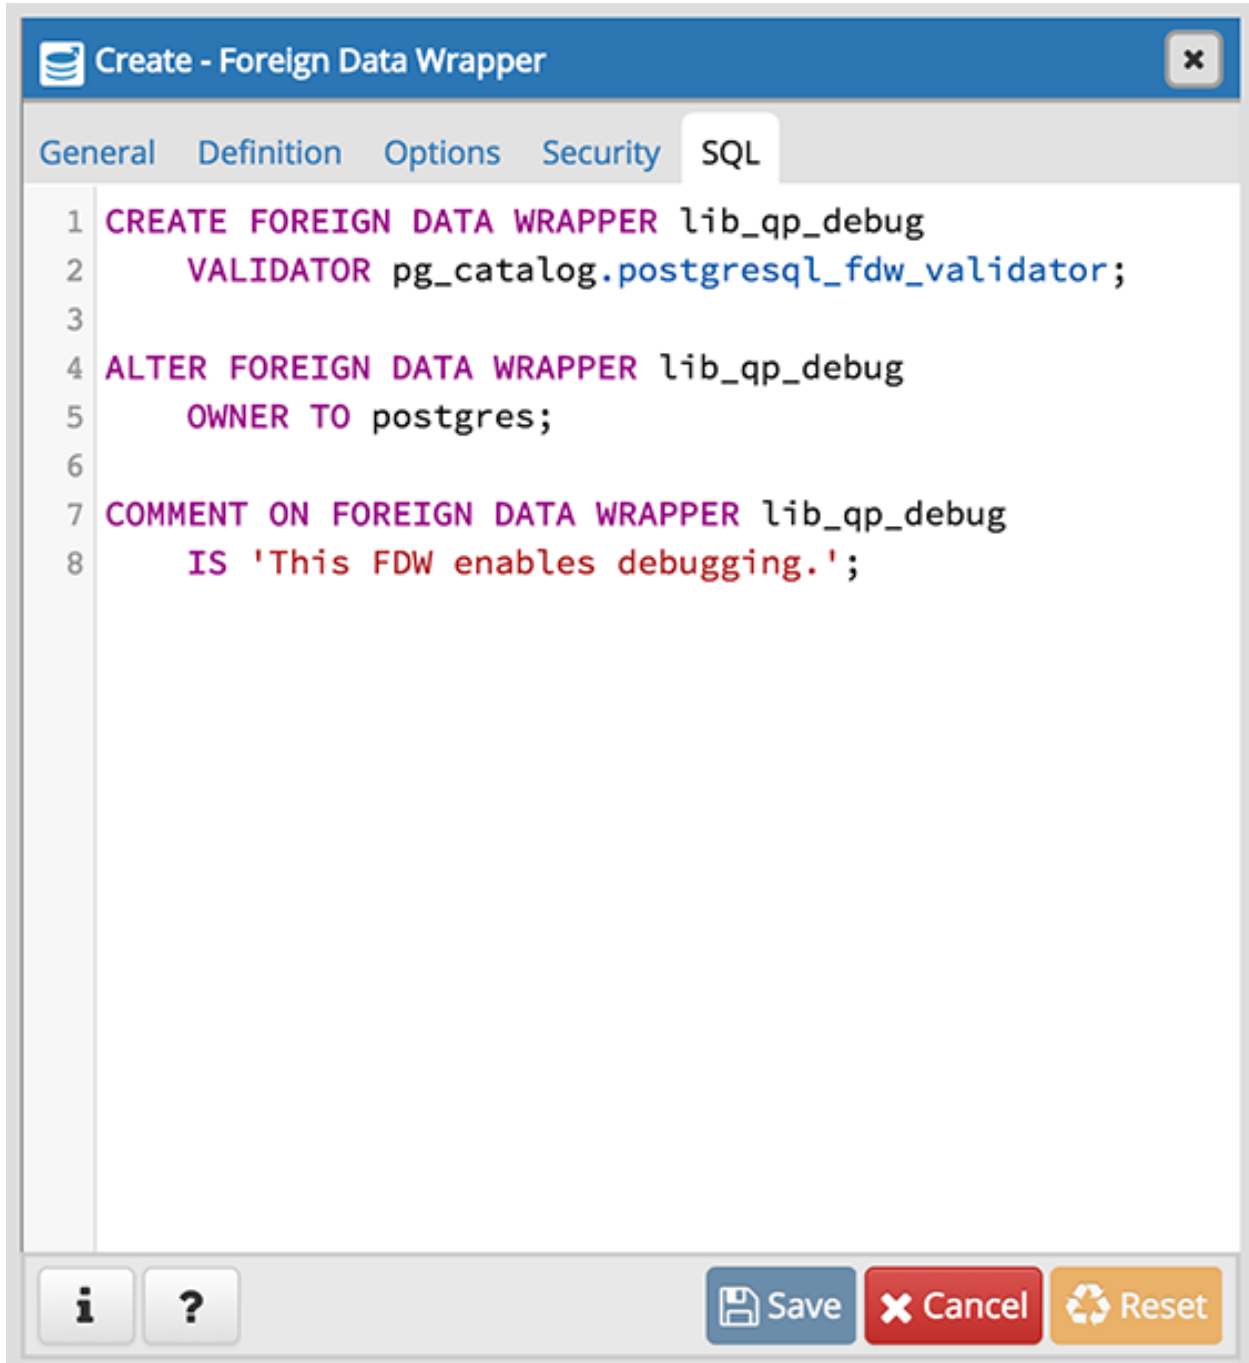
Click add to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Foreign Data Wrapper* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Foreign Data Wrapper* dialog:



The example creates a foreign data wrapper named *libpq_debug* that uses pre-existing validator and handler functions, *dblink_fdw_validator* and *libpq_fdw_handler*. Selections on the *Options* tab set *debug* equal to *true*. The foreign data wrapper is owned by *postgres*.

- Click the *Help* button (?) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.8 Foreign Server Dialog

Use the *Foreign Server* dialog to create a foreign server. A foreign server typically encapsulates connection information that a foreign-data wrapper uses to access an external data resource. Each foreign data wrapper may connect to a different foreign server; in the *pgAdmin* tree control, expand the node of the applicable foreign data wrapper to launch the *Foreign Server* dialog.


The *Foreign Server* dialog organizes the development of a foreign server through the following dialog tabs: *General*, *Definition*, *Options*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.

The screenshot shows the 'Create - Foreign Server' dialog box. The title bar is blue with a close button. Below the title bar are five tabs: 'General', 'Definition', 'Options', 'Security', and 'SQL'. The 'General' tab is selected and contains three fields: 'Name' (a text input field), 'Owner' (a dropdown menu showing 'postgres'), and 'Comment' (a large text area). At the bottom of the dialog are four buttons: an information icon, a question mark icon, a 'Save' button, a 'Cancel' button, and a 'Reset' button.

Use the fields in the *General* tab to identify the foreign server:

- Use the *Name* field to add a descriptive name for the foreign server. The name will be displayed in the *pgAdmin* tree control. It must be unique within the database.
- Use the drop-down listbox next to *Owner* to select a role.
- Store notes about the foreign server in the *Comment* field.

Click the *Definition* tab to continue.

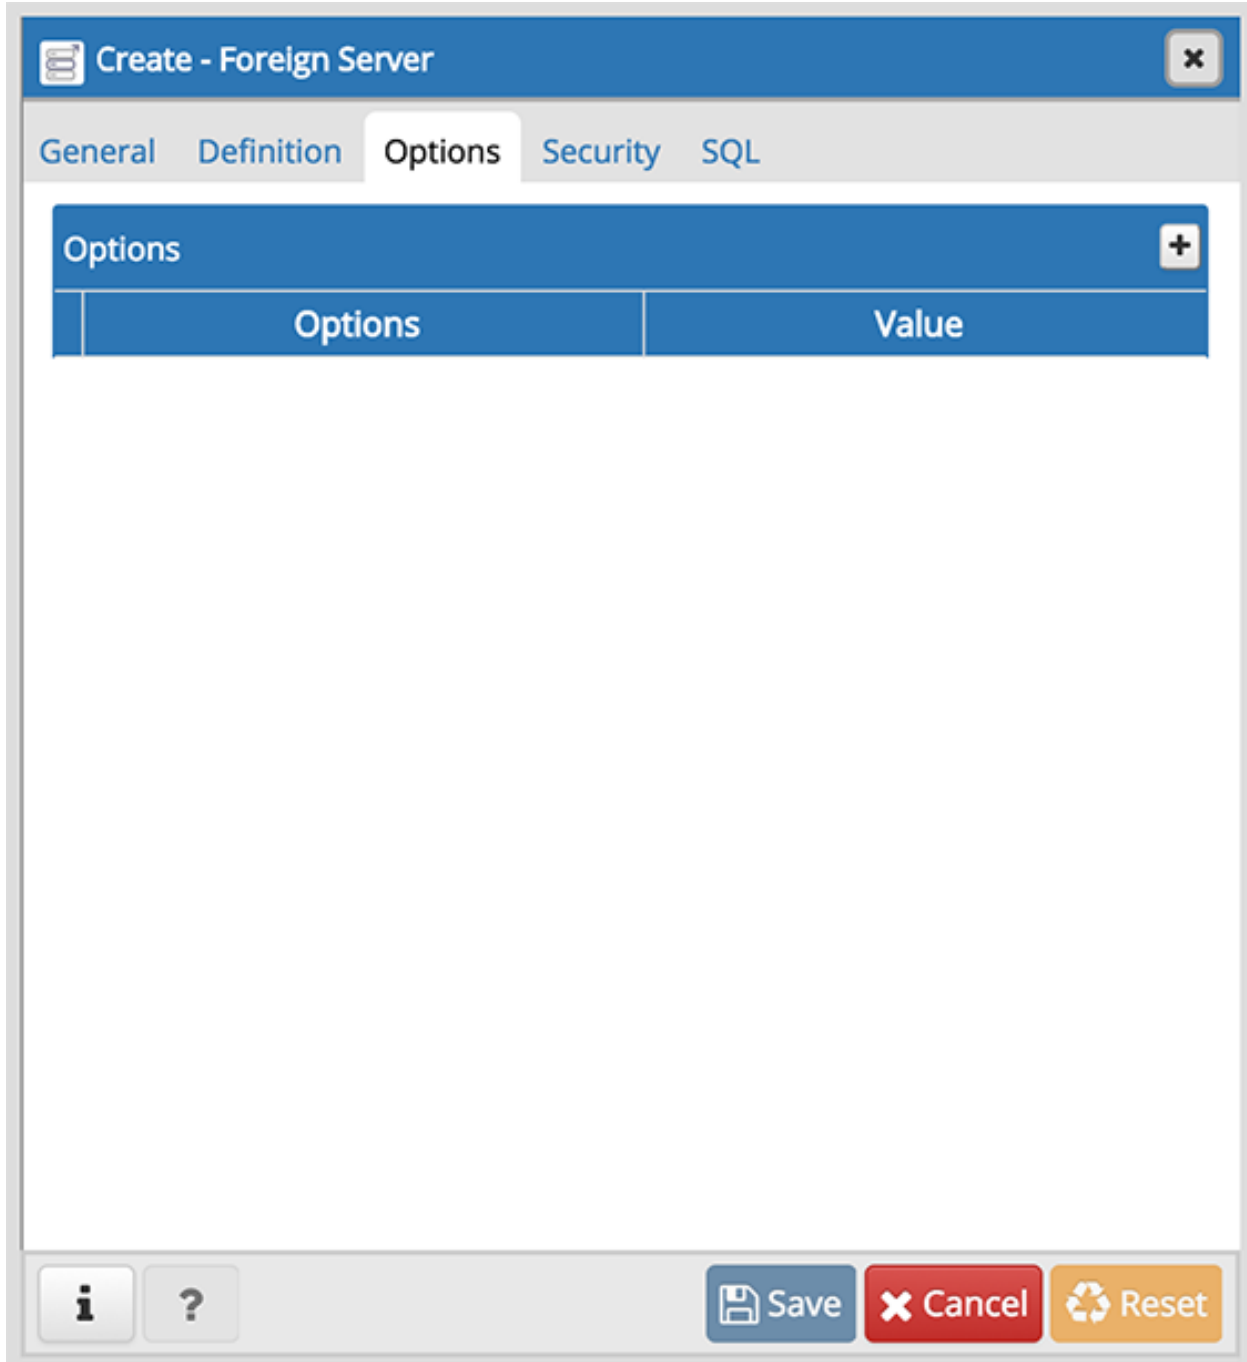


The screenshot shows a dialog box titled "Create - Foreign Server" with a close button (X) in the top right corner. The dialog has four tabs: "General", "Definition", "Options", "Security", and "SQL". The "Definition" tab is currently selected. Inside this tab, there are two text input fields: "Type" and "Version". At the bottom of the dialog, there are three buttons: "Save" (blue), "Cancel" (red), and "Reset" (orange). To the left of these buttons are two smaller buttons: an information icon (i) and a help icon (?).

Use the fields in the *Definition* tab to set parameters:

- Use the *Type* field to specify a server type.
- Use the *Version* field to specify a server version.

Click the *Options* tab to continue.

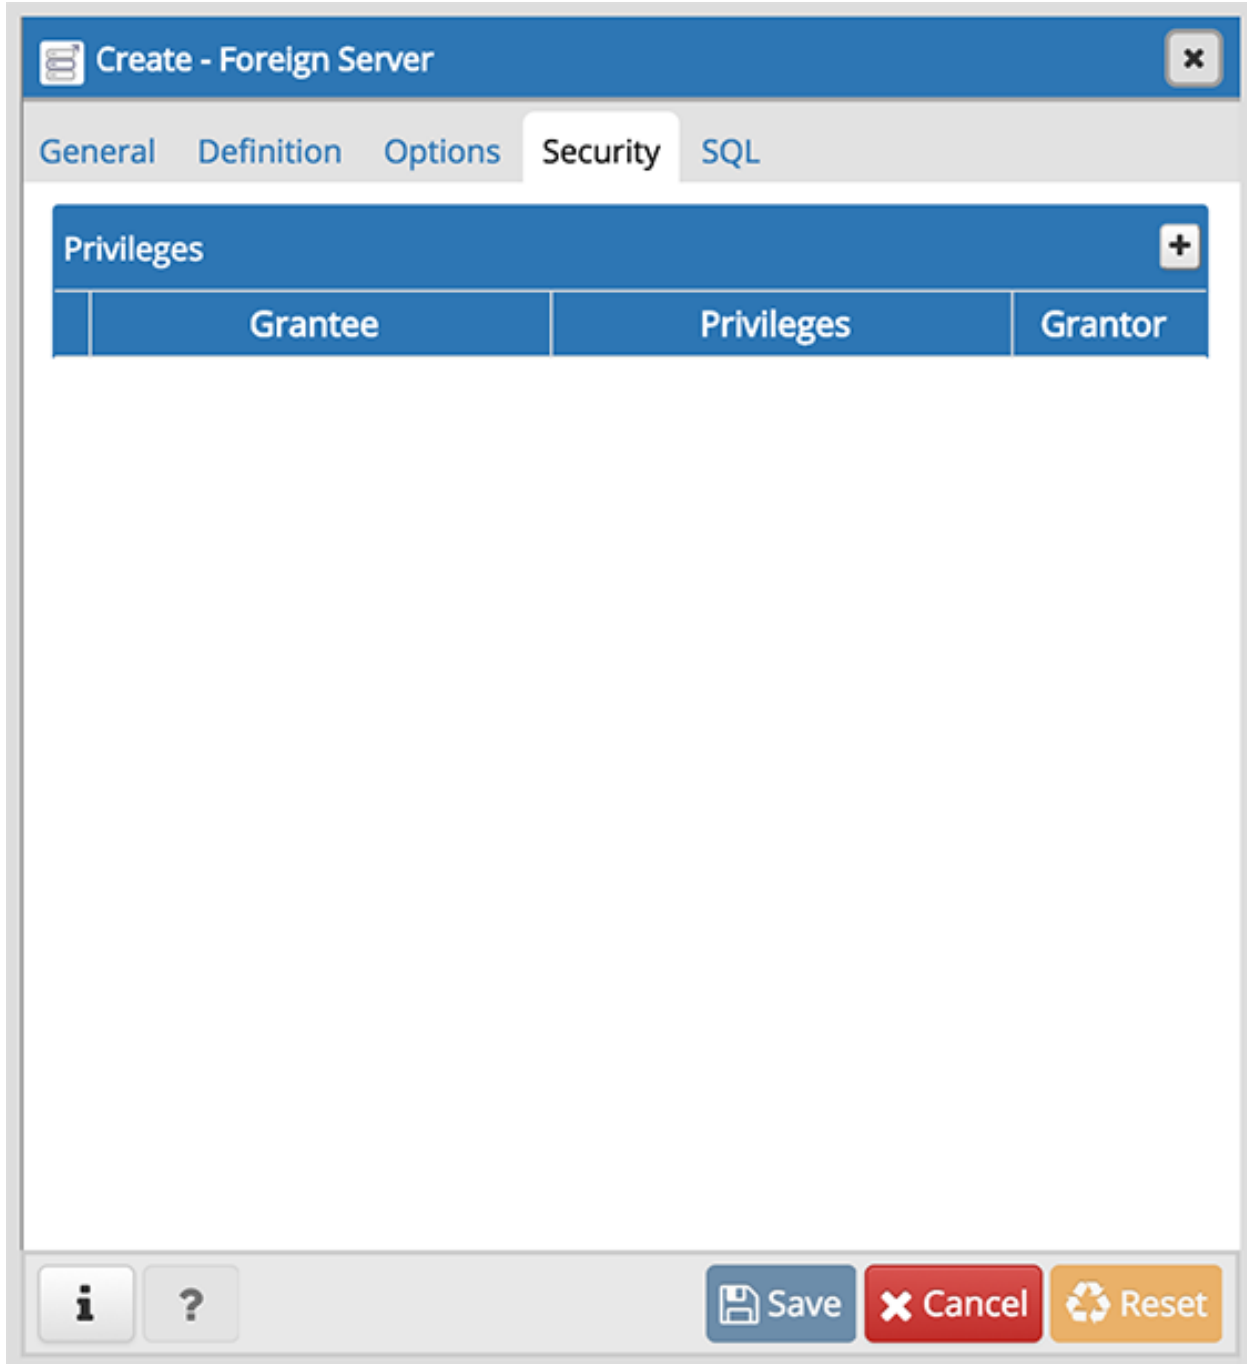


Use the fields in the *Options* tab to specify options. Click the *Add* button to create an option clause for the foreign server.

- Specify the option name in the *Option* field.
- Provide a corresponding value in the *Value* field.

Click *Add* to create each additional clause; to discard an option, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Security* tab to continue.



Use the *Security* tab to assign security privileges to the foreign server. Click *Add* before you assign a set of privileges.

- Select the name of the role from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privileges to the specified user.
- Select the name of the role from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the foreign server. This is a required field.

Click *Add* to assign a new set of privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* dialog.

Click the *SQL* tab to continue.

Your entries in the *Foreign Server* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Foreign Server* dialog:



The example shown demonstrates creating a foreign server for the foreign data wrapper *hdfs_fdw*. It has the name *hdfs_server*; its type is *hiveserver2*. Options for the foreign server include a host and a port.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new


tab displays the PostgreSQL core documentation.

- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.9 Foreign Table Dialog

Use the *Foreign Table* dialog to define a foreign table in the current database. Foreign tables define the structure of an external data source that resides on a foreign server.

The *Foreign Table* dialog organizes the development of a foreign table through the following dialog tabs: *General*, *Definition*, *Columns*, *Constraints*, *Options*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



Create - Foreign Table

General Definition Columns Constraints Options Security SQL

Name: weblogs

Owner: postgres

Schema: public

Comment:

Save Cancel Reset

Use the fields in the *General* tab to identify the foreign table:

- Use the *Name* field to add a descriptive name for the foreign table. The name of the foreign table must be distinct from the name of any other foreign table, table, sequence, index, view, existing data type, or materialized view in the same schema. The name will be displayed in the *pgAdmin* tree control.
- Use the drop-down listbox next to *Owner* to select the name of the role that will own the foreign table.
- Select the name of the schema in which the foreign table will reside from the drop-down listbox in the *Schema* field.
- Store notes about the foreign table in the *Comment* field.

Click the *Definition* tab to continue.

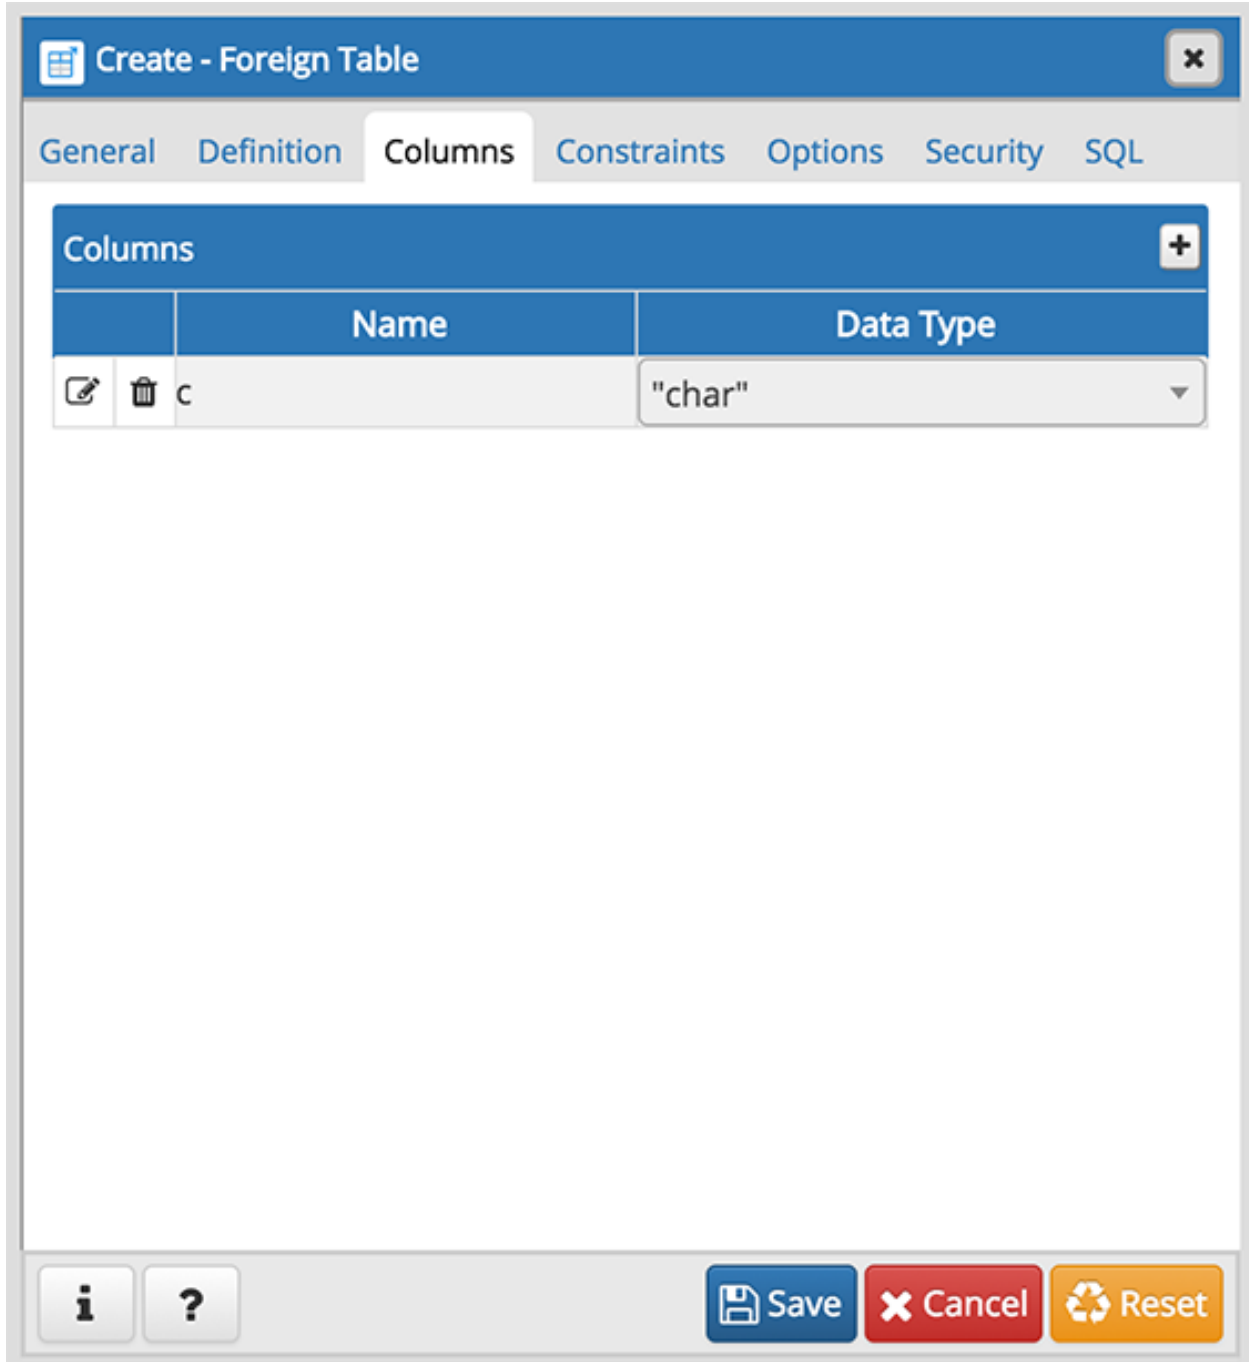


The screenshot shows the 'Create - Foreign Table' dialog box in pgAdmin 4. The 'Definition' tab is selected, showing the 'Foreign server' field set to 'hdfs_server' and the 'Inherits' field set to 'public.t'. The dialog has tabs for 'General', 'Definition', 'Columns', 'Constraints', 'Options', 'Security', and 'SQL'. At the bottom are buttons for 'Save', 'Cancel', and 'Reset'.

Use the fields in the *Definition* tab to define the external data source:

- Use the drop-down listbox next to *Foreign server* to select a foreign server. This list is populated with servers defined through the *Foreign Server* dialog.
- Use the drop-down listbox next to *Inherits* to specify a parent table. The foreign table will inherit all of its columns. This field is optional.

Click the *Columns* tab to continue.



Use the fields in the *Columns* tab to add columns and their attributes to the table. Click the *Add* icon (+) to define a column:

- Use the *Name* field to add a descriptive name for the column.
- Use the drop-down listbox in the *Data Type* field to select a data type for the column. This can include array specifiers. For more information on which data types are supported by PostgreSQL, refer to Chapter 8 of the core documentation.

Click the *Add* icon (+) to specify each additional column; to discard a column, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Constraints* tab to continue.

The screenshot shows the 'Create - Foreign Table' dialog box in pgAdmin 4, with the 'Constraints' tab selected. The dialog has a title bar with a close button and a tabbed interface with 'General', 'Definition', 'Columns', 'Constraints', 'Options', 'Security', and 'SQL'. The 'Constraints' tab contains a table with the following data:

	Name	Check	No Inherit	Validate?
	myconstraint	mycheck	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Below the table is a 'General' form with the following fields:

- Name:
- Check:
- No Inherit:
- Validate?:

At the bottom of the dialog are buttons for 'i', '?', 'Save', 'Cancel', and 'Reset'.


Use the fields in the *Constraints* tab to apply a table constraint to the foreign table. Click the *Add* icon (+) to define a constraint:

- Use the *Name* field to add a descriptive name for the constraint. If the constraint is violated, the constraint name is present in error messages, so constraint names like *col must be positive* can be used to communicate helpful information.
- Use the *Check* field to write a check expression producing a Boolean result. Each row in the foreign table is expected to satisfy the check expression.

- Check the *No Inherit* checkbox to specify that the constraint will not propagate to child tables.
- Uncheck the *Validate* checkbox to disable validation. The database will not assume that the constraint holds for all rows in the table.

Click the *Add* icon (+) to specify each additional constraint; to discard a constraint, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Options* tab to continue.



The screenshot shows the 'Create - Foreign Table' dialog box in pgAdmin 4, with the 'Options' tab selected. The dialog has a title bar with a close button (X) and a menu icon. Below the title bar are tabs for 'General', 'Definition', 'Columns', 'Constraints', 'Options', 'Security', and 'SQL'. The 'Options' tab is active, showing a table with two columns: 'Option' and 'Value'. The table has a header row and one empty data row. A trash icon is visible to the left of the data row. At the top right of the table area is an 'Add' (+) icon. At the bottom of the dialog are buttons for 'Save', 'Cancel', and 'Reset', along with information (i) and help (?) icons.

Options		+
	Option	Value
🗑️		

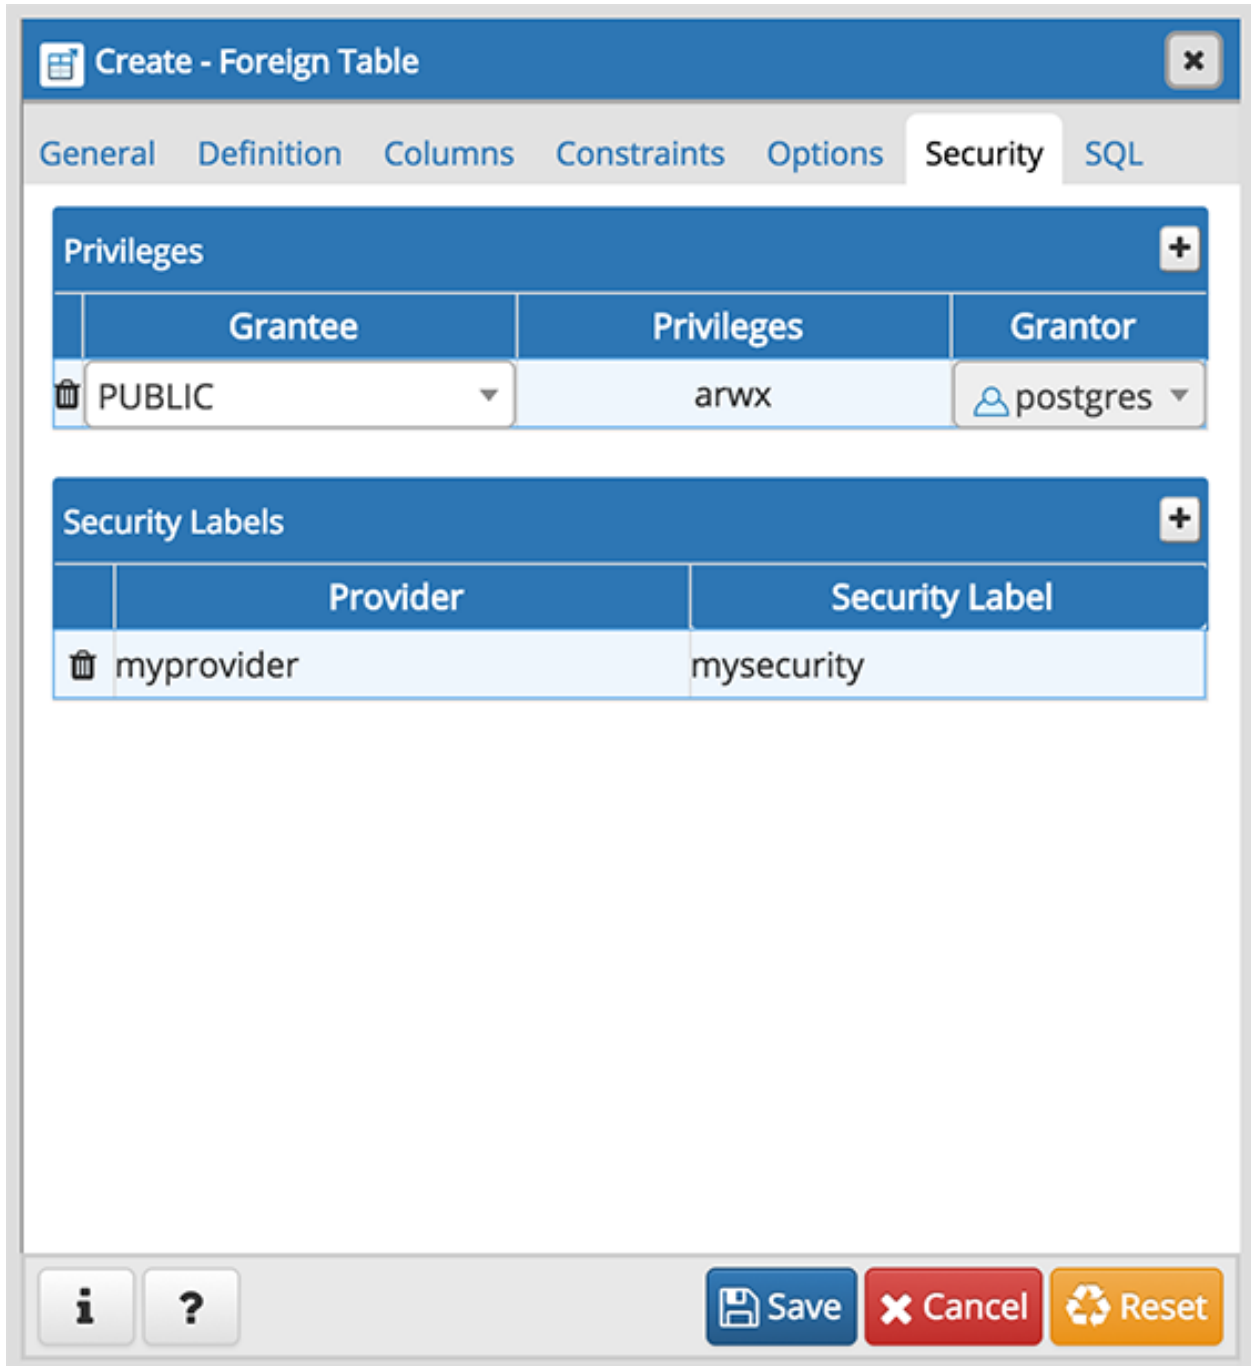
Use the fields in the *Options* tab to specify options to be associated with the new foreign table or one of its columns; the accepted option names and values are specific to the foreign data wrapper associated with the foreign server. Click

the *Add* icon (+) to add an option/value pair.

- Specify the option name in the *Option* field. Duplicate option names are not allowed.
- Provide a corresponding value in the *Value* field.

Click the *Add* icon (+) to specify each additional option/value pair; to discard an option, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Security* tab to continue.



Use the *Security* tab to assign privileges and define security labels.

Use the *Privileges* panel to assign privileges to a role. Click the *Add* icon (+) to set privileges for database objects:

- Select the name of the role to which privileges will be assigned from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.
- Select the name of the role that owns the foreign table from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click the *Add* icon (+) to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the function. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

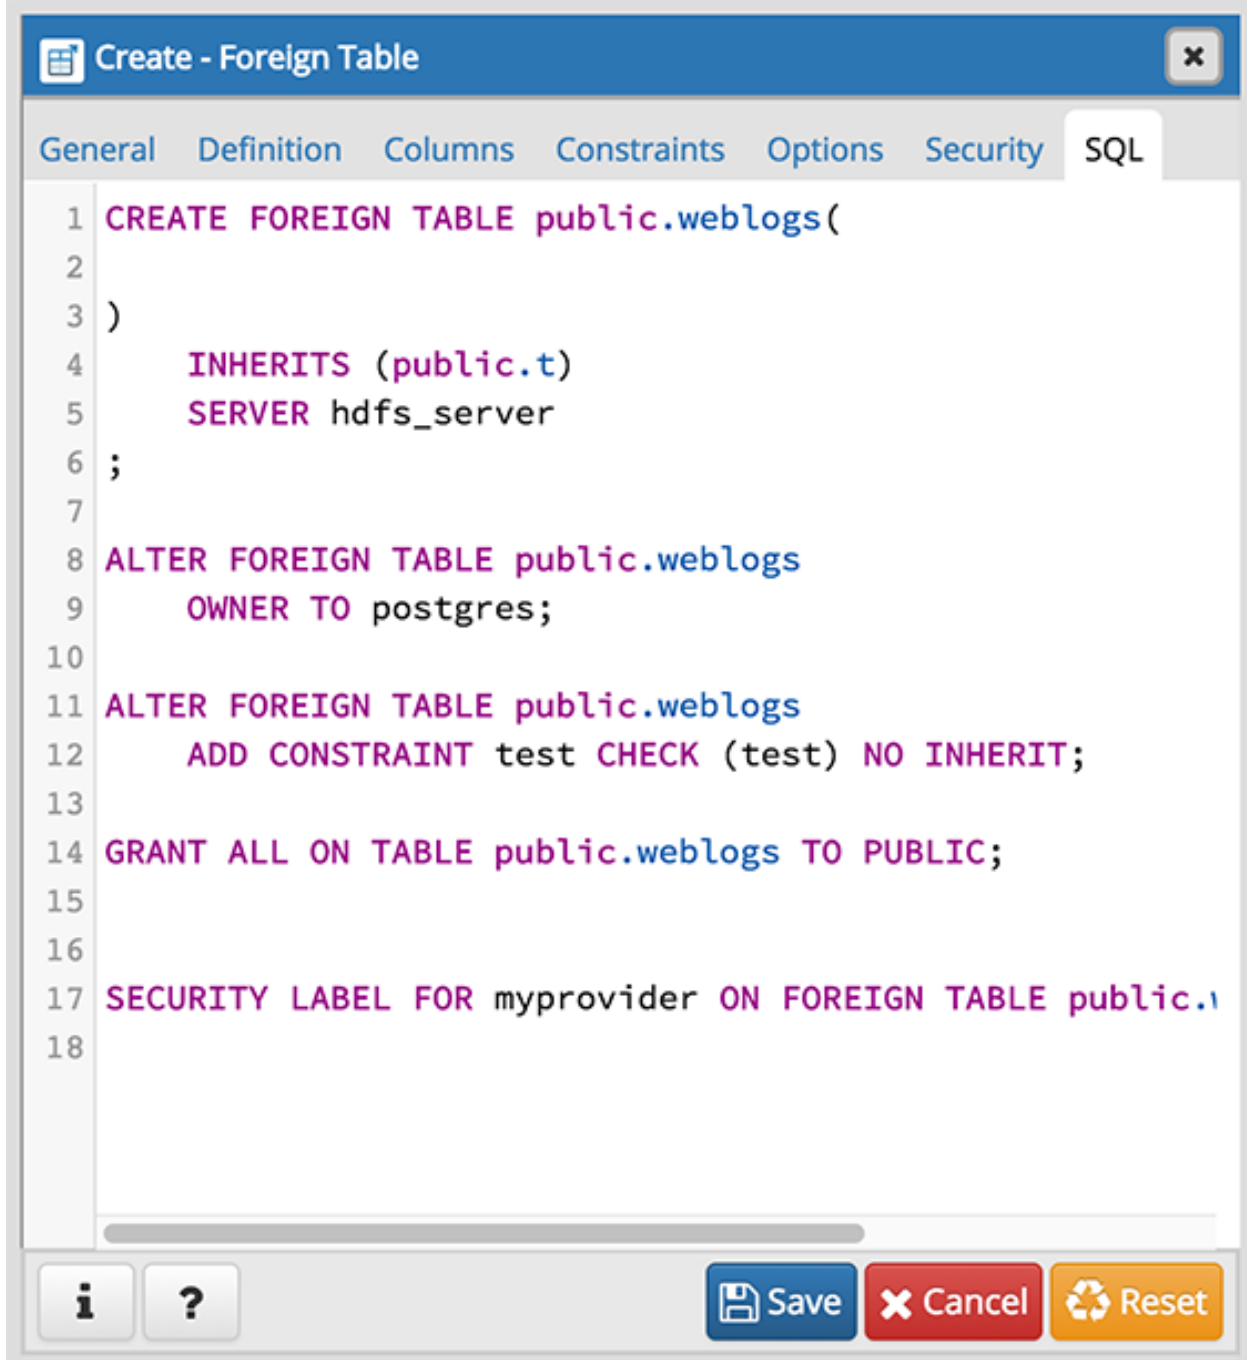
Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Foreign Table* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Foreign Table* dialog:



The example shown demonstrates creating a foreign table *weblogs* with multiple columns and two options.


- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.10 FTS Configuration dialog

Use the *FTS Configuration* dialog to configure a full text search. A text search configuration specifies a text search parser that can divide a string into tokens, along with dictionaries that can identify searchable tokens.

The *FTS Configuration* dialog organizes the development of a FTS configuration through the following dialog tabs: “*General*, *Definition*, and *Tokens*. The *SQL* tab displays the SQL code generated by dialog selections.

Click the *General* tab to begin.



The screenshot shows the 'Create - FTS Configuration' dialog box. The 'General' tab is active, showing the following fields:

- Name:** meme_phrases
- Owner:** postgres
- Schema:** public
- Comment:** (empty text area)

At the bottom of the dialog, there are three buttons: **Save**, **Cancel**, and **Reset**. There are also information and help icons on the left.

Use the fields in the *General* tab to identify a FTS configuration:

- Use the *Name* field to add a descriptive name for the FTS configuration. The name will be displayed in the *pgAdmin* tree control.
- Use the drop-down listbox next to *Owner* to specify the role that will own the configuration.
- Select the name of the schema in which the FTS configuration will reside from the drop-down listbox in the *Schema* field.
- Store notes about the FTS configuration in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows a dialog box titled "Create - FTS Configuration" with a close button (X) in the top right corner. Below the title bar are four tabs: "General", "Definition", "Tokens", and "SQL". The "Definition" tab is selected. In this tab, there are two fields: "Parser" and "Copy Config". The "Parser" field is a text input with "default" entered and a small "x" and dropdown arrow on the right. The "Copy Config" field is a dropdown menu that is currently empty. At the bottom of the dialog, there are four buttons: an information icon (i), a help icon (?), a blue "Save" button with a floppy disk icon, a red "Cancel" button with an "x" icon, and an orange "Reset" button with a refresh icon.

Use the fields in the *Definition* tab to define parameters:

- Select the name of the text search parser from the drop-down listbox in the *Parser* field.
- Select a language from the drop-down listbox in the *Copy Config* field.

Click the *Tokens* tab to continue.

The screenshot shows a dialog box titled "Create - FTS Configuration" with a close button (X) in the top right corner. The dialog has four tabs: "General", "Definition", "Tokens" (which is selected), and "SQL". Below the tabs is a blue header bar. Underneath, there is a "Tokens" label followed by a text input field and a plus sign (+) icon. Below this is a table with two columns: "Token" and "Dictionaries". The table is currently empty. At the bottom of the dialog, there are three buttons: "Save" (blue), "Cancel" (red), and "Reset" (orange). To the left of these buttons are two small icons: an information icon (i) and a help icon (?).

Use the fields in the *Tokens* tab to add a token:

- Use the *Tokens* field to specify the name of a token.
- Click the *Add* icon (+) to create a token.
- Use the *Dictionaries* field to specify a dictionary.

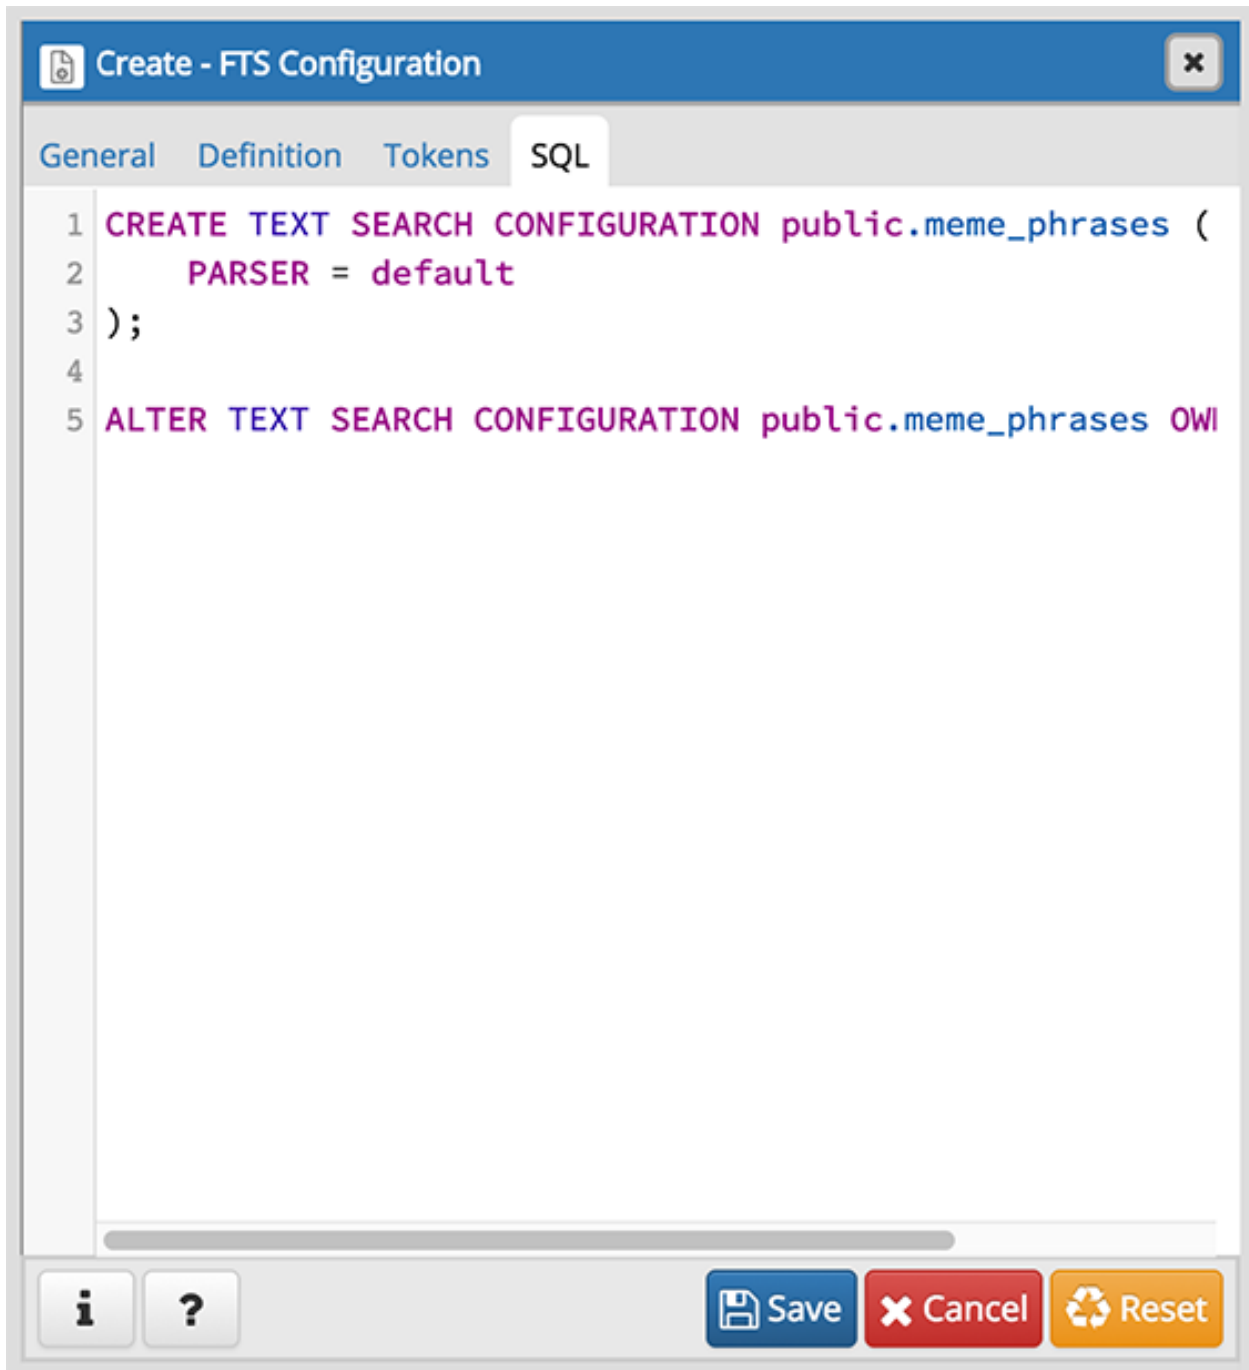
Repeat these steps to add additional tokens; to discard a token, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *FTS Configuration* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *FTS Configuration* dialog:




The example shown demonstrates creating a FTS configuration named *meme_phrases*. It uses the *default* parser.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.11 FTS Dictionary Dialog

Use the *FTS Dictionary* dialog to create a full text search dictionary. You can use a predefined templates or create a new dictionary with custom parameters.

The *FTS Dictionary* dialog organizes the development of a FTS dictionary through the following dialog tabs: *General*, *Definition*, and *Options*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - FTS Dictionary' dialog box in pgAdmin 4. The 'General' tab is active, displaying the following fields:

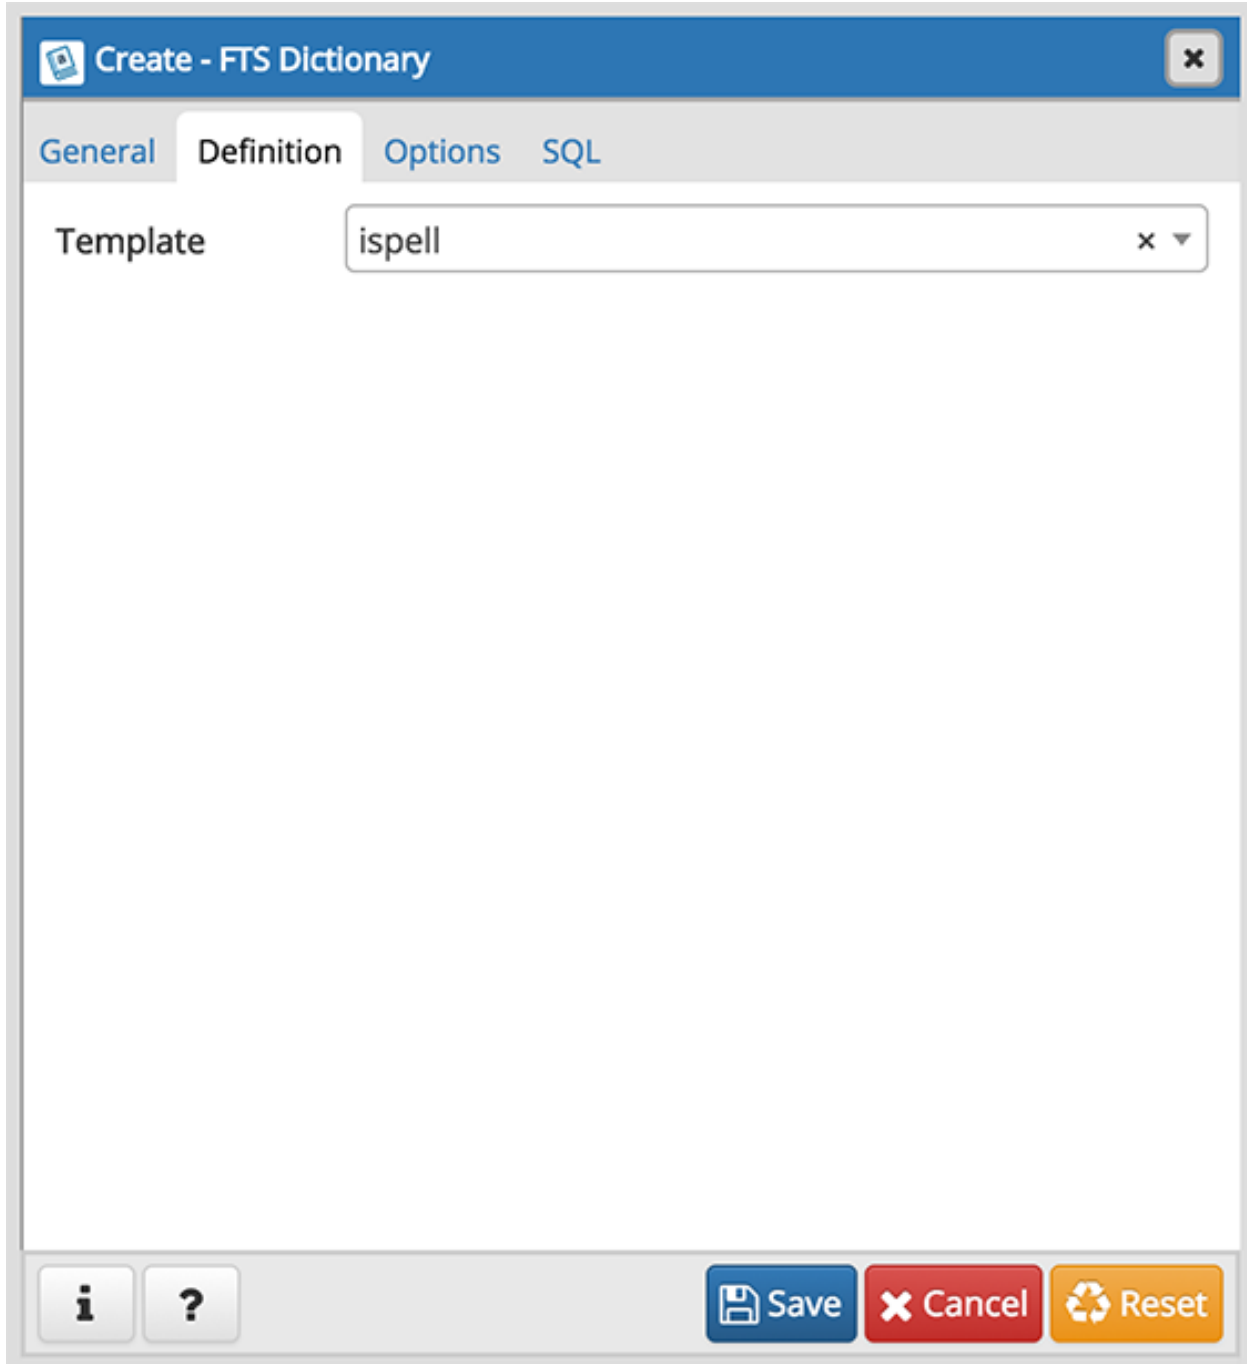
- Name:** more_stopwords
- Owner:** postgres
- Schema:** public
- Comment:** (empty text area)

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left.

Use the fields in the *General* tab to identify the dictionary:

- Use the *Name* field to add a descriptive name for the dictionary. The name will be displayed in the *pgAdmin* tree control.
- Use the drop-down listbox next to *Owner* to select the role that will own the FTS Dictionary.
- Select the name of the schema in which the dictionary will reside from the drop-down listbox in the *Schema* field.
- Store notes about the dictionary in the *Comment* field.

Click the *Definition* tab to continue.

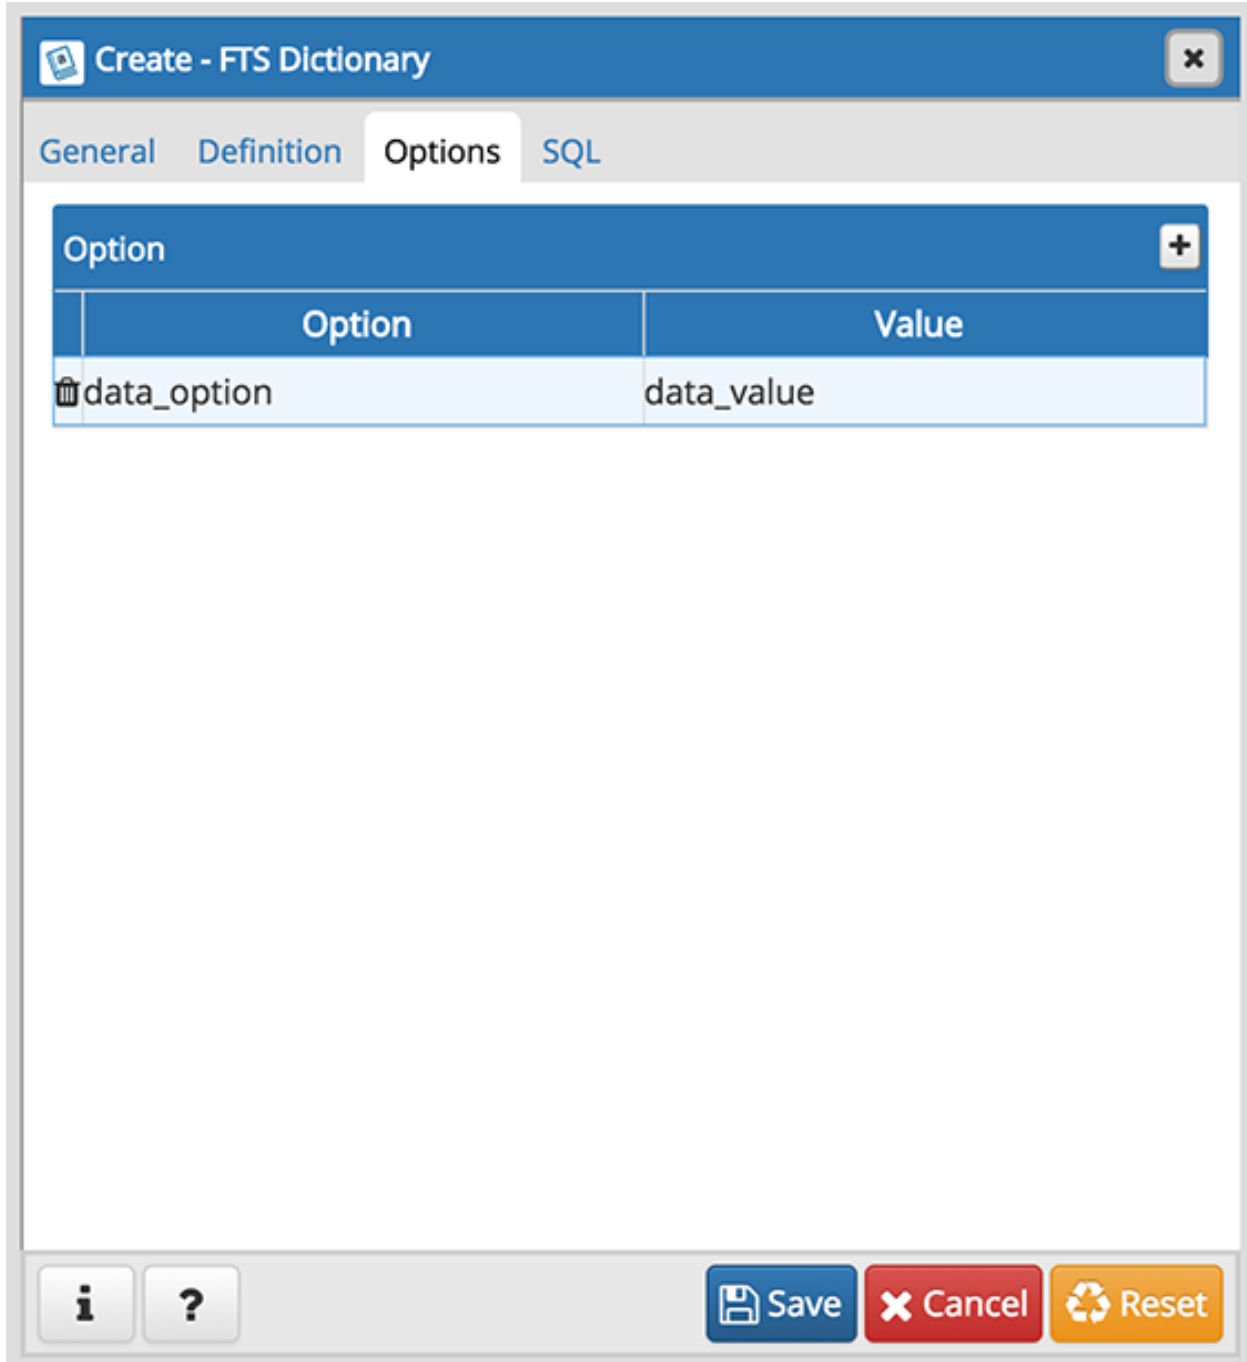


Use the field in the *Definition* tab to choose a template from the drop-down listbox:

*Select *ispell* to select the Ispell template. The Ispell dictionary template supports morphological dictionaries, which can normalize many different linguistic forms of a word into the same lexeme. For example, an English Ispell dictionary can match all declensions and conjugations of the search term bank, e.g., banking, banked, banks, banks', and bank's. Ispell dictionaries usually recognize a limited set of words, so they should be followed by another broader dictionary; for example, a Snowball dictionary, which recognizes everything. *Select *simple* to select the simple template. The simple dictionary template operates by converting the input token to lower case and checking it against a file of stop words. If it is found in the file then an empty array is returned, causing the token to be discarded. If not, the lower-cased form of the word is returned as the normalized lexeme. Alternatively, the dictionary can be configured to report non-stop-words as unrecognized, allowing them to be passed on to the next dictionary in the list. *Select *snowball* to

select the Snowball template. The Snowball dictionary template is based on a project by Martin Porter, inventor of the popular Porter's stemming algorithm for the English language. Snowball now provides stemming algorithms for many languages (see the Snowball site for more information). Each algorithm understands how to reduce common variant forms of words to a base, or stem, spelling within its language. A Snowball dictionary recognizes everything, whether or not it is able to simplify the word, so it should be placed at the end of the dictionary list. It is useless to have it before any other dictionary because a token will never pass through it to the next dictionary. *Select *synonym* to select the synonym template. This dictionary template is used to create dictionaries that replace a word with a synonym. Phrases are not supported (use the thesaurus template (Section 12.6.4) for that). A synonym dictionary can be used to overcome linguistic problems, for example, to prevent an English stemmer dictionary from reducing the word Paris to *pari*. *Select *thesaurus* to select the thesaurus template. A thesaurus dictionary replaces all non-preferred terms by one preferred term and, optionally, preserves the original terms for indexing as well. PostgreSQL's current implementation of the thesaurus dictionary is an extension of the synonym dictionary with added phrase support.

Click the *Options* tab to continue.



Use the fields in the *Options* tab to provide template-specific options. Click the *Add* icon (+) to add an option clause:

- Specify the name of an option in the *Option* field
- Provide a value for the option in the *Value* field.

Click the *Add* icon (+) to specify each additional option/value pair; to discard an option, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *FTS Dictionary* dialog generate a generate a SQL command. Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *FTS Dictionary* dialog:



The example shown demonstrates creating a custom dictionary named *more_stopwords* which is based on the simple template and is configured to use standard English.


- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.

- Click the *Reset* button to restore configuration parameters.

3.12 FTS Parser Dialog

Use the *FTS Parser* dialog to create a new text search parser. A text search parser defines a method for splitting a text string into tokens and assigning types (categories) to the tokens.

The *FTS Parser* dialog organizes the development of a text search parser through the following dialog tabs: *General*, and *Definition*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - FTS Parser' dialog box. The title bar reads 'Aa Create - FTS Parser'. The dialog has three tabs: 'General', 'Definition', and 'SQL'. The 'General' tab is selected. The 'Name' field contains 'data_paraser'. The 'Schema' field is a dropdown menu showing 'public'. The 'Comment' field is a large empty text area. At the bottom, there are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields in the *General* tab to identify a text search parser:

- Use the *Name* field to add a descriptive name for the parser. The name will be displayed in the *pgAdmin* tree control.
- Select the name of the schema in which the parser will reside from the drop-down listbox in the *Schema* field.
- Store notes about the domain in the *Comment* field.

Click the *Definition* tab to continue.

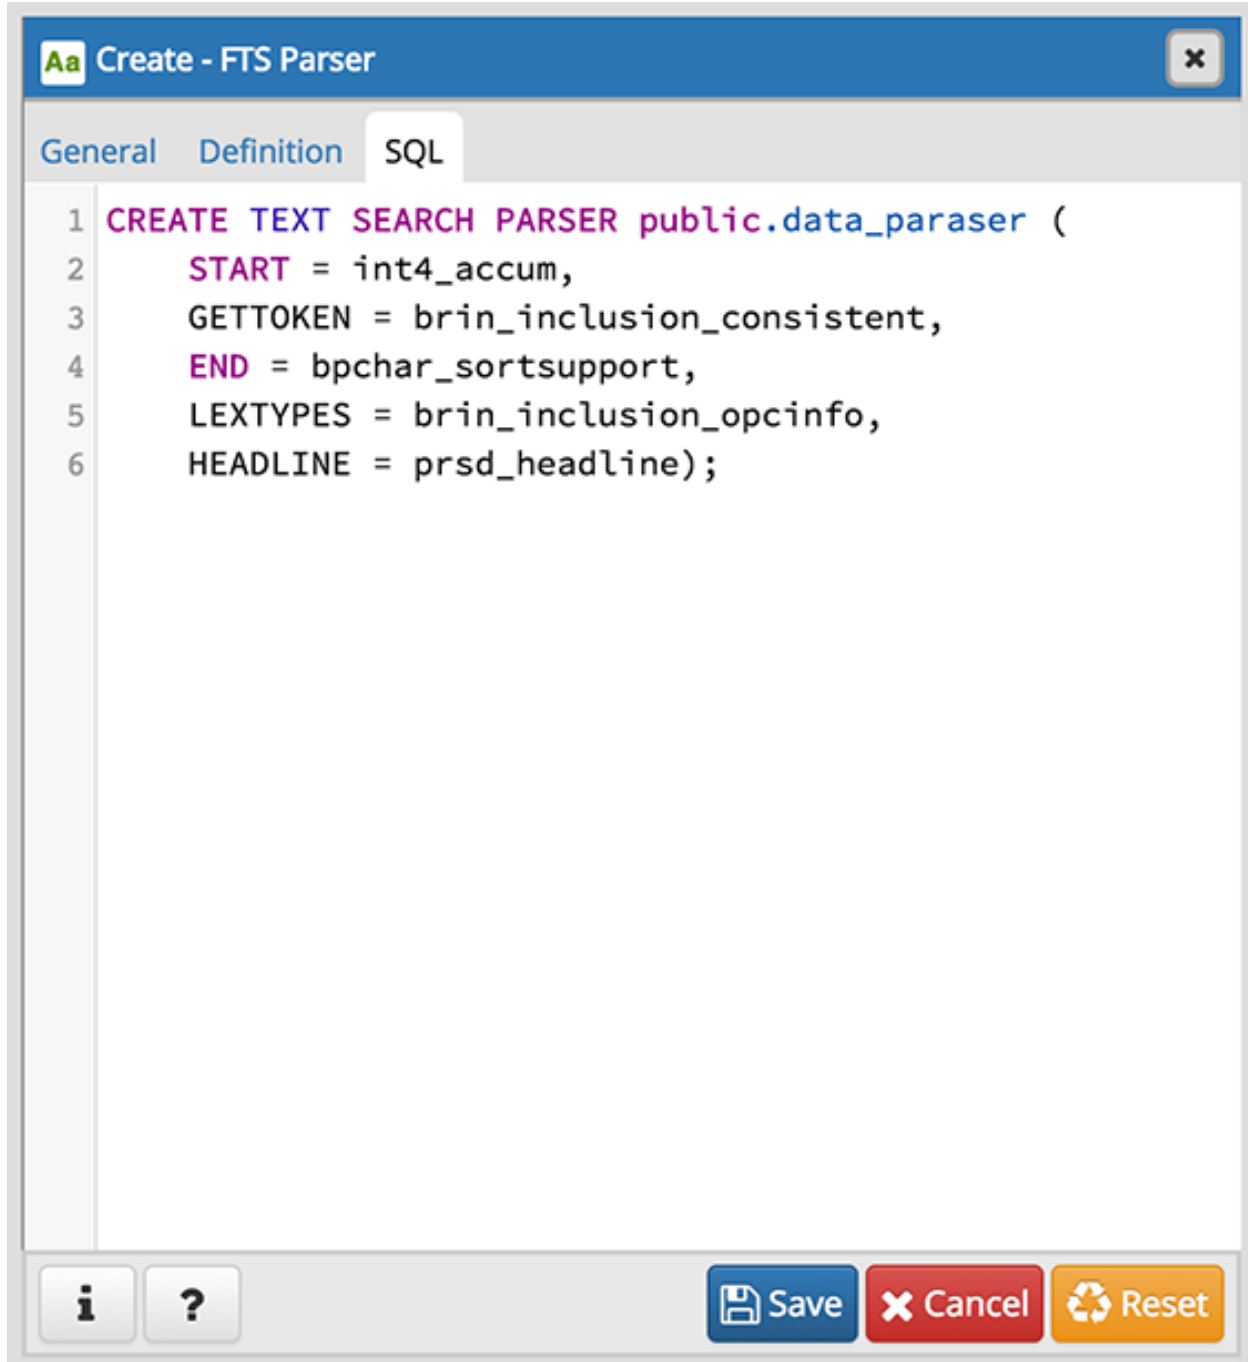
The screenshot shows the 'Create - FTS Parser' dialog box in pgAdmin 4. The 'Definition' tab is selected, and the following functions are configured:

Function Name	Value
Start function	int4_accum
Get next token function	brin_inclusion_consistent
End function	bpchar_sortsupport
Letypes function	brin_inclusion_opcinfo
Headline function	prsd_headline

Use the fields in the *Definition* tab to define parameters:

- Use the drop-down listbox next to *Start function* to select the name of the function that will initialize the parser.
- Use the drop-down listbox next to *Get next token function* to select the name of the function that will return the next token.
- Use the drop-down listbox next to *End function* to select the name of the function that is called when the parser is finished.
- Use the drop-down listbox next to *Lextypes function* to select the name of the lextypes function for the parser. The lextypes function returns an array that contains the id, alias, and a description of the tokens used by the parser.
- Use the drop-down listbox next to *Headline function* to select the name of the headline function for the parser. The headline function returns an excerpt from the document in which the terms of the query are highlighted.

Click the *SQL* tab to continue.



Your entries in the *FTS Parser* dialog generate a generate a SQL command. Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.13 FTS Template Dialog

Use the *FTS Template* dialog to create a new text search template. A text search template defines the functions that implement text search dictionaries.

The *FTS Template* dialog organizes the development of a text search Template through the following dialog tabs: *General*, and *Definition*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - FTS Template' dialog box. The title bar is blue with a search icon and the text 'Create - FTS Template'. Below the title bar are three tabs: 'General', 'Definition', and 'SQL'. The 'General' tab is selected. The 'Name' field contains 'ru.template'. The 'Schema' field is a dropdown menu showing 'public'. The 'Comment' field is a large empty text area. At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). On the left side of the bottom bar, there are two small buttons: 'i' (info) and '?' (help).


Use the fields in the *General* tab to identify a template:

- Use the *Name* field to add a descriptive name for the template. The name will be displayed in the *pgAdmin* tree

control.

- Select the name of the schema in which the template will reside from the drop-down listbox in the *Schema* field.
- Store notes about the template in the *Comment* field.

Click the *Definition* tab to continue.



The screenshot shows a dialog box titled "Create - FTS Template" with three tabs: "General", "Definition", and "SQL". The "Definition" tab is active. It contains two fields: "Init function" with a dropdown menu showing "brin_inclusion_opcinfo" and "Lexize function" with a dropdown menu showing "dispell_lexize". At the bottom of the dialog, there are three buttons: "Save" (blue), "Cancel" (red), and "Reset" (orange). To the left of these buttons are two small icons: an information icon (i) and a help icon (?).

Use the fields in the *Definition* tab to define function parameters:

- Use the drop-down listbox next to *Init function* to select the name of the init function for the template. The init function is optional.

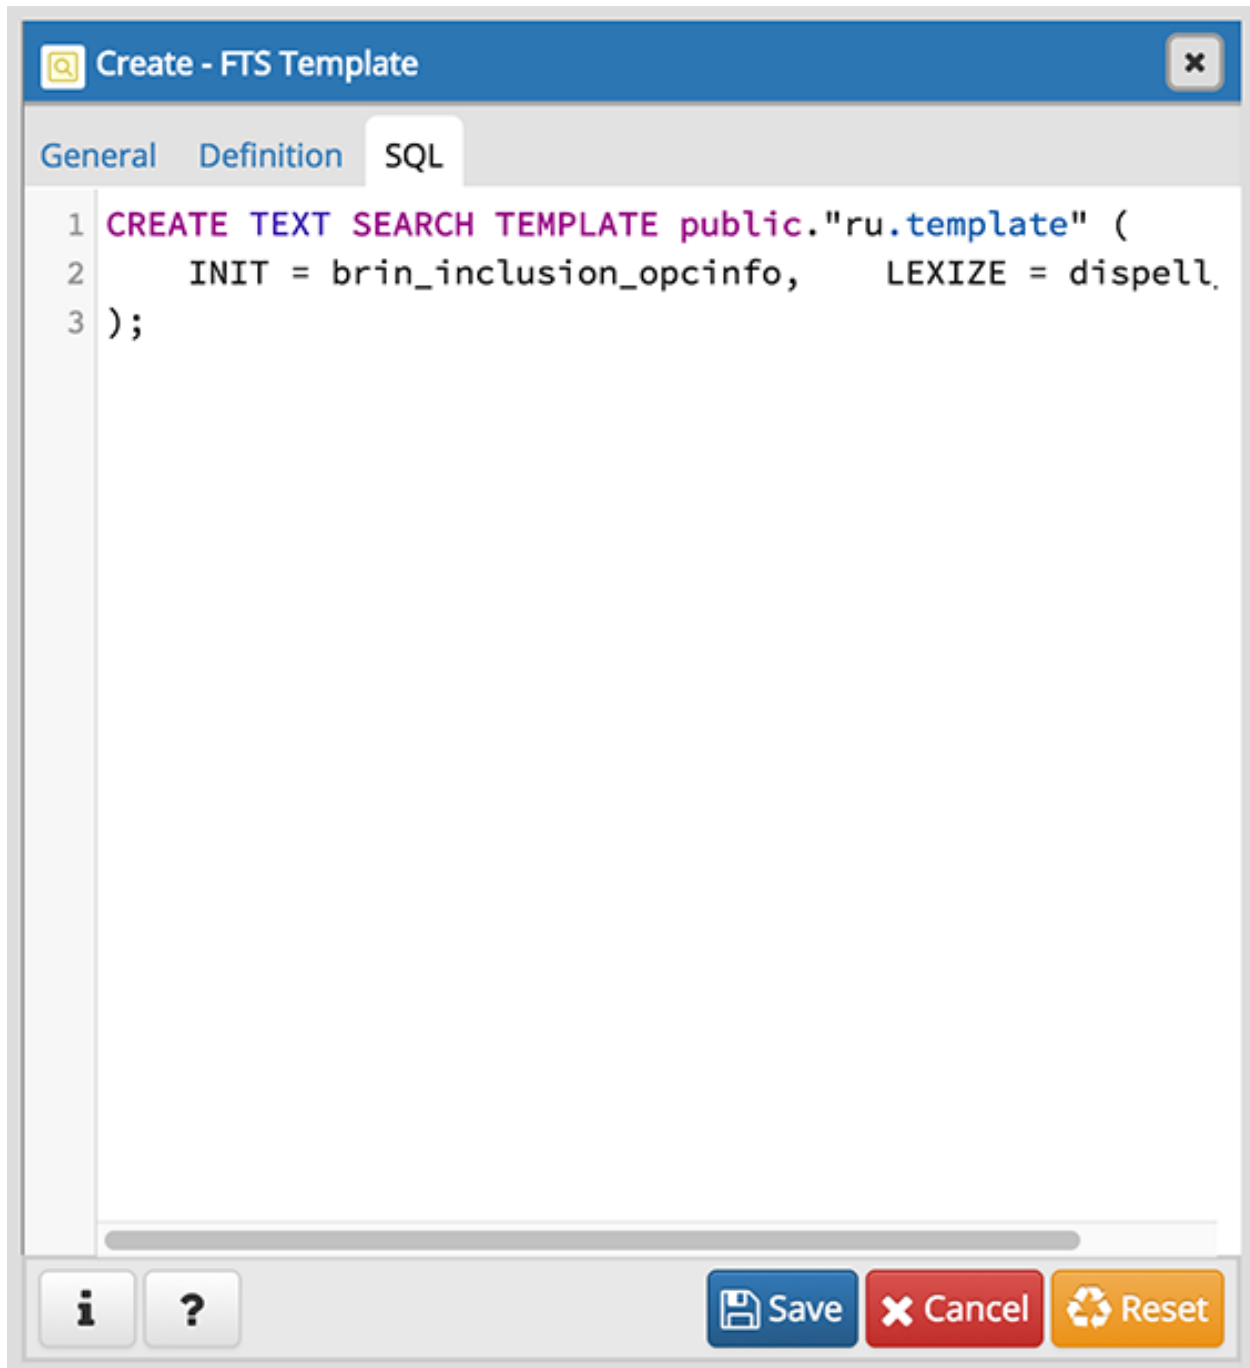
- Use the drop-down listbox next to *Lexize function* to select the name of the lexize function for the template. The lexize function is required.

Click the *SQL* tab to continue.

Your entries in the *FTS Template* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *FTS Template* dialog:




The example shown demonstrates creating a fts template named *ru_template* that uses the ispell dictionary.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.14 Function Dialog

Use the *Function* dialog to define a function. If you drop and then recreate a function, the new function is not the same entity as the old; you must drop existing rules, views, triggers, etc. that refer to the old function.

The *Function* dialog organizes the development of a function through the following dialog tabs: *General*, *Definition*, *Options*, *Arguments*, *Parameters*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



Create - Function

General Definition Options Arguments Parameters Security SQL

Name: list_emp

Owner: postgres

Schema: public


Comment: This procedure returns list of employees

i ? Save Cancel Reset

Use the fields in the *General* tab to identify a function:

- Use the *Name* field to add a descriptive name for the function. The name will be displayed in the *pgAdmin* tree control.
- Use the drop-down listbox next to *Owner* to select the name of the role that will own the function.
- Use the drop-down listbox next to *Schema* to select the schema in which the function will be created.
- Store notes about the function in the *Comment* field.

Click the *Definition* tab to continue.



Use the fields in the *Definition* tab to define the function:

- Use the drop-down listbox next to *Return type* to select the data type returned by the function, if any.
- Use the drop-down listbox next to *Language* to select the implementation language. The default is *sql*.
- Use the *Code* field to write the code that will execute when the function is called.

Click the *Options* tab to continue.

The screenshot shows the 'Create - Procedure' dialog box in pgAdmin 4, with the 'Options' tab selected. The dialog contains the following fields and controls:

- Volatility:** A drop-down menu set to 'VOLATILE'.
- Returns a set?:** A checkbox with 'Yes' selected.
- Strict?:** A checkbox with 'Yes' selected.
- Security of definer?:** A checkbox with 'Yes' selected.
- Window?:** A checkbox with 'Yes' selected.
- Parallel:** A drop-down menu set to 'UNSAFE'.
- Estimated cost:** An empty text input field.
- Estimated rows:** An empty text input field.
- Leak proof?:** A checkbox with 'Yes' selected.

At the bottom of the dialog, there are buttons for 'i' (Help), '?' (Info), 'Save', 'Cancel', and 'Reset'.

Use the fields in the *Options* tab to describe or modify the action of the function:

- Use the drop-down listbox next to *Volatility* to select one of the following. *VOLATILE* is the default value.
 - *VOLATILE* indicates that the function value can change even within a single table scan, so no optimizations can be made.
 - *STABLE* indicates that the function cannot modify the database, and that within a single table scan it will consistently return the same result for the same argument values.
 - *IMMUTABLE* indicates that the function cannot modify the database and always returns the same result when given the same argument values.

- Move the *Returns a Set?* switch to indicate if the function returns a set that includes multiple rows. The default is *No*.
- Move the *Strict?* switch to indicate if the function always returns NULL whenever any of its arguments are NULL. If *Yes*, the function is not executed when there are NULL arguments; instead a NULL result is assumed automatically. The default is *No*.
- Move the *Security of definer?* switch to specify that the function is to be executed with the privileges of the user that created it. The default is *No*.
- Move the *Window?* switch to indicate that the function is a window function rather than a plain function. The default is *No*. This is currently only useful for functions written in C. The WINDOW attribute cannot be changed when replacing an existing function definition. For more information about the CREATE FUNCTION command, see the PostgreSQL core documentation available at:

<http://www.postgresql.org/docs/9.5/static/functions-window.html>

- Use the *Estimated cost* field to specify a positive number representing the estimated execution cost for the function, in units of `cpu_operator_cost`. If the function returns a set, this is the cost per returned row.
- Use the *Estimated rows* field to specify a positive number giving the estimated number of rows that the query planner should expect the function to return. This is only allowed when the function is declared to return a set. The default assumption is 1000 rows.
- Move the *Leak proof?* switch to indicate whether the function has side effects. The default is *No*. This option can only be set by the superuser.

Click the *Arguments* tab to continue.

The screenshot shows the 'Create - Function' dialog box with the 'Arguments' tab selected. The dialog has a title bar 'Create - Function' and a close button. Below the title bar are tabs for 'General', 'Definition', 'Options', 'Arguments', 'Parameters', 'Security', and 'SQL'. The 'Arguments' tab contains a table with the following structure:

Arguments				
	Data Type	Mode	Argument Name	Default Value
	"char"			

At the bottom of the dialog are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields in the *Arguments* tab to define an argument. Click the *Add* icon (+) to set parameters and values for the argument:

- Use the drop-down listbox in the *Data type* field to select a data type.
- Use the drop-down listbox in the *Mode* field to select a mode. Select *IN* for an input parameter; select *OUT* for an output parameter; select *INOUT* for both an input and an output parameter; or, select *VARIADIC* to specify a VARIADIC parameter.
- Provide a name for the argument in the *Argument Name* field.
- Specify a default value for the argument in the *Default Value* field.

Click the *Add* icon (+) to define another argument; to discard an argument, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Parameters* tab to continue.



The screenshot shows the 'Create - Function' dialog box in pgAdmin 4, with the 'Parameters' tab selected. The dialog has a title bar with a close button (X) and a header with tabs for 'General', 'Definition', 'Options', 'Arguments', 'Parameters', 'Security', and 'SQL'. The 'Parameters' tab contains a table with the following structure:

Parameters		+
	Name	Value
	array_nulls	False

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). On the left side of the bottom bar, there are two icons: an information icon (i) and a help icon (?).

Use the fields in the *Parameters* tab to specify settings that will be applied when the function is invoked. Click the *Add* icon (+) to add a *Name/Value* field in the table.

- Use the drop-down listbox in the *Name* column in the *Parameters* panel to select a parameter.
- Use the *Value* field to specify the value that will be associated with the selected variable. This field is context-sensitive.

Click the *Security* tab to continue.

Create - Function

General Definition Options Arguments Parameters **Security** SQL

Privileges

	Grantee	Privileges	Grantor
	pem_admin	X*	postgres

Security Labels

	Provider	Security Label
	myProvider	mySecurity

i ? Save Cancel Reset

Use the *Security* tab to assign privileges and define security labels.

Use the *Privileges* panel to assign usage privileges for the function to a role.

- Select the name of the role from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.
- Select the name of the role from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click the *Add* icon (+) to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the function. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Function* dialog generate a generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

3.15 Language Dialog

Use the CREATE LANGUAGE dialog to register a new procedural language.

The *Language* dialog organizes the registration of a procedural language through the following dialog tabs: *General*, *Definition*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.

Create - Language

General Definition security SQL

Name plperl

Owner pem_agent

Comment

i ? Save Cancel Reset

Use the fields in the *General* tab to identify a language:

- Use the drop-down listbox next to *Name* to select a language script.
- Use the drop-down listbox next to *Owner* to select a role.
- Store notes about the language in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Language' dialog box in pgAdmin 4, with the 'Definition' tab selected. The dialog has a blue title bar with a close button. Below the title bar are four tabs: 'General', 'Definition', 'security', and 'SQL'. The 'Definition' tab is active, showing a 'Trusted?' checkbox with a green 'Yes' label and a white 'No' label. Below this are three drop-down menus: 'Handler Function', 'Inline Function', and 'Validator Function'. At the bottom of the dialog are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left side of the bottom bar.

Use the fields in the *Definition* tab to define parameters:

- Move the *Trusted?* switch to the *No* position to specify only users with PostgreSQL superuser privilege can use this language. The default is *Yes*.
- When enabled, use the drop-down listbox next to *Handler Function* to select the function that will be called to execute the language's functions.
- When enabled, use the drop-down listbox next to *Inline Function* to select the function that will be called to execute an anonymous code block (DO command) in this language.
- When enabled, use the drop-down listbox next to *Validator Function* to select the function that will be called

when a new function in the language is created, to validate the new function.

Click the *Security* tab to continue.

The screenshot shows the 'Create - Language' dialog box with the 'security' tab selected. The dialog is divided into four tabs: 'General', 'Definition', 'security', and 'SQL'. The 'security' tab is active and contains two main sections: 'Privileges' and 'Security Labels'.

The 'Privileges' section has a table with the following structure:

Grantee	Privileges	Grantor
PUBLIC		postgres

The 'Security Labels' section has a table with the following structure:

Provider	Security Label
myprovider	mysecurity

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left.

Use the *Security* tab to assign privileges and define security labels.

Use the *Privileges* panel to assign privileges to a role. Click the *Add* icon (+) to set privileges for database objects:

- Select the name of the role from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.

- Select the name of the role from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click the *Add* icon (+) to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the function. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

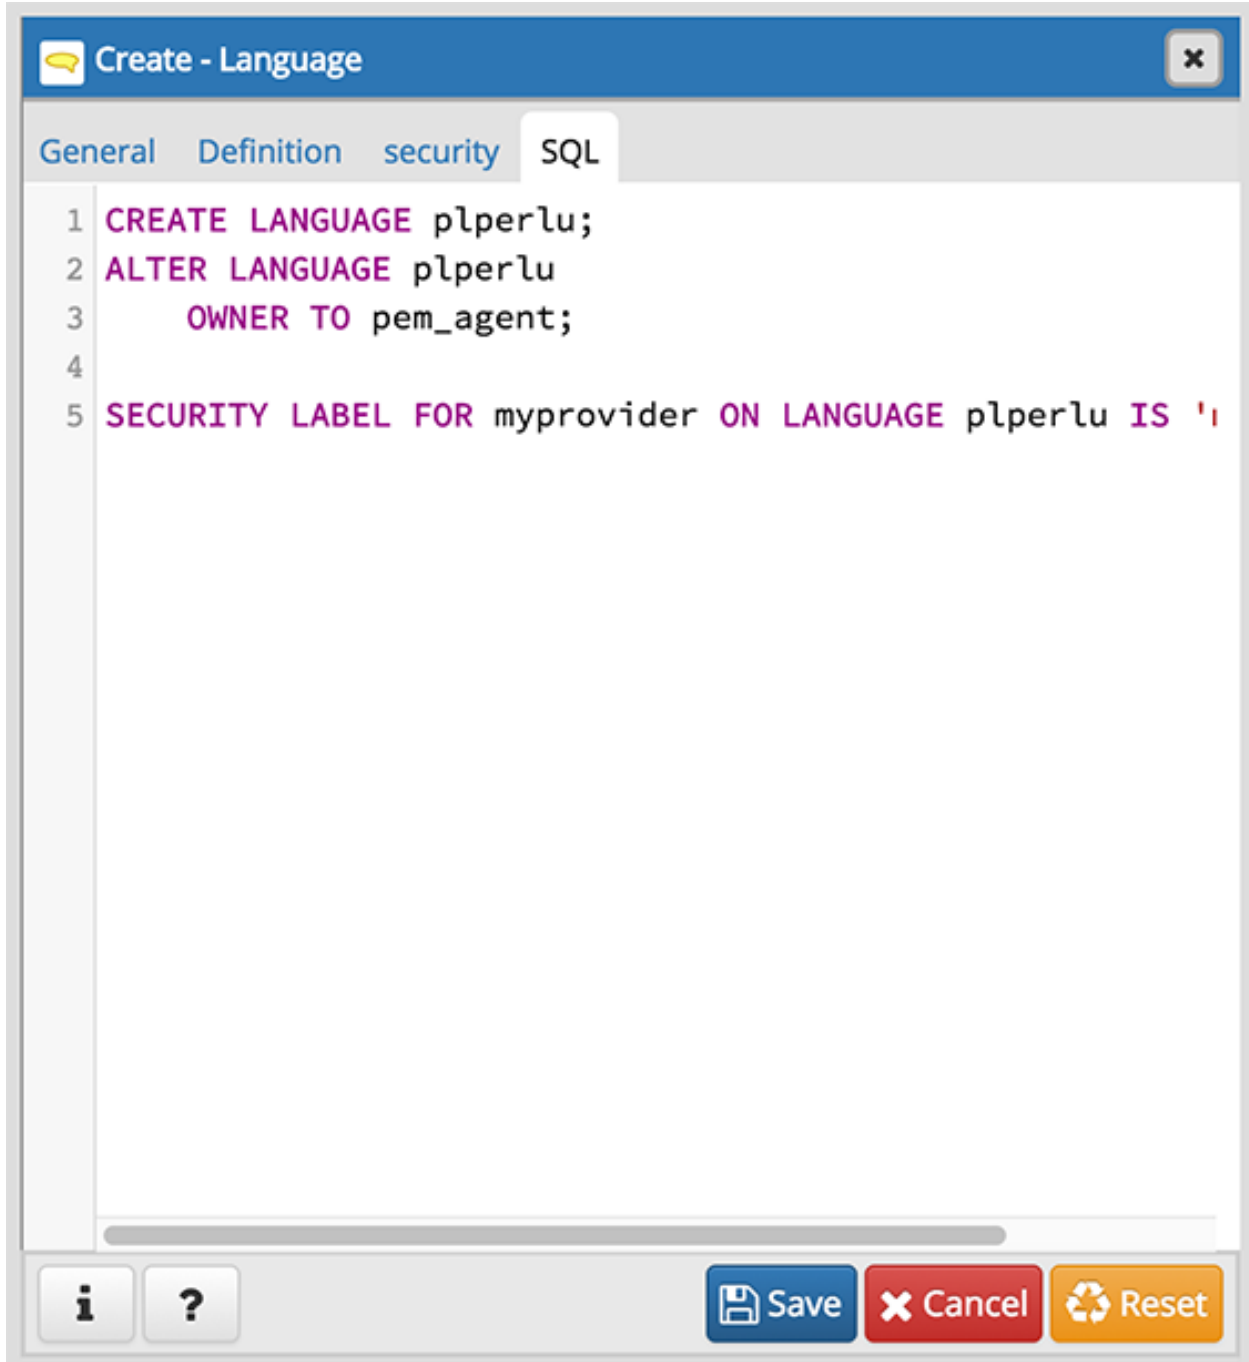
Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Language* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Language* dialog:




“The example shown demonstrates creating the procedural language named *plperl*.”

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.16 Materialized View Dialog

Use the *Materialized View* dialog to define a materialized view. A materialized view is a stored or cached view that contains the result set of a query. Use the `REFRESH MATERIALIZED VIEW` command to update the content of a materialized view.

The *Materialized View* dialog organizes the development of a materialized_view through the following dialog tabs: *General*, *Definition*, *Storage*, *Parameter*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the "Create - Materialized View" dialog box. The "General" tab is active, showing the following fields:

- Name:** new_hires
- Owner:** postgres
- Schema:** public
- Comment:** (empty text area)

At the bottom of the dialog, there are three buttons: "Save", "Cancel", and "Reset".

Use the fields in the *General* tab to identify the materialized view:


- Use the *Name* field to add a descriptive name for the materialized view. The name will be displayed in the *pgAdmin* tree control.
- Use the drop-down listbox next to *Owner* to select the role that will own the materialized view.
- Select the name of the schema in which the materialized view will reside from the drop-down listbox in the *Schema* field.
- Store notes about the materialized view in the *Comment* field.

Click the *Definition* tab to continue.



Use the text editor field in the *Definition* tab to provide the query that will populate the materialized view.

Click the *Storage* tab to continue.



The screenshot shows the 'Create - Materialized View' dialog box with the 'Storage' tab selected. The dialog has a blue title bar with a close button. Below the title bar are tabs for 'General', 'Definition', 'Storage', 'Parameter', 'Security', and 'SQL'. The 'Storage' tab is active and contains the following fields:

- With Data:** A toggle switch set to 'Yes'.
- Tablespace:** A drop-down listbox showing 'pg_default'.
- Fill Factor:** An empty text input field.

At the bottom of the dialog are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left side of the bottom bar.

Use the fields in the *Storage* tab to maintain the materialized view:

- Move the *With Data* switch to the *Yes* position to specify the materialized view should be populated at creation time. If not, the materialized view cannot be queried until you invoke `REFRESH MATERIALIZED VIEW`.
- Use the drop-down listbox next to *Tablespace* to select a location for the materialized view.
- Use the *Fill Factor* field to specify a fill factor for the materialized view. The fill factor for a table is a percentage between 10 and 100. 100 (complete packing) is the default.

Click the *Parameter* tab to continue.

Create - Materialized View
×

General
Definition
Storage
Parameter
Security
SQL

Table
Toast Table

Custom auto-vacuum? Yes

Enabled? No

Vacuum Table

Label	Value	Default value
ANALYZE scale factor		0.10
ANALYZE base threshold		50
FREEZE maximum age		200,000,000
VACUUM cost delay		20
VACUUM cost limit		-1
VACUUM scale factor		0.20
VACUUM base threshold		50
FREEZE minimum age		50,000,000
FREEZE table age		150,000,000

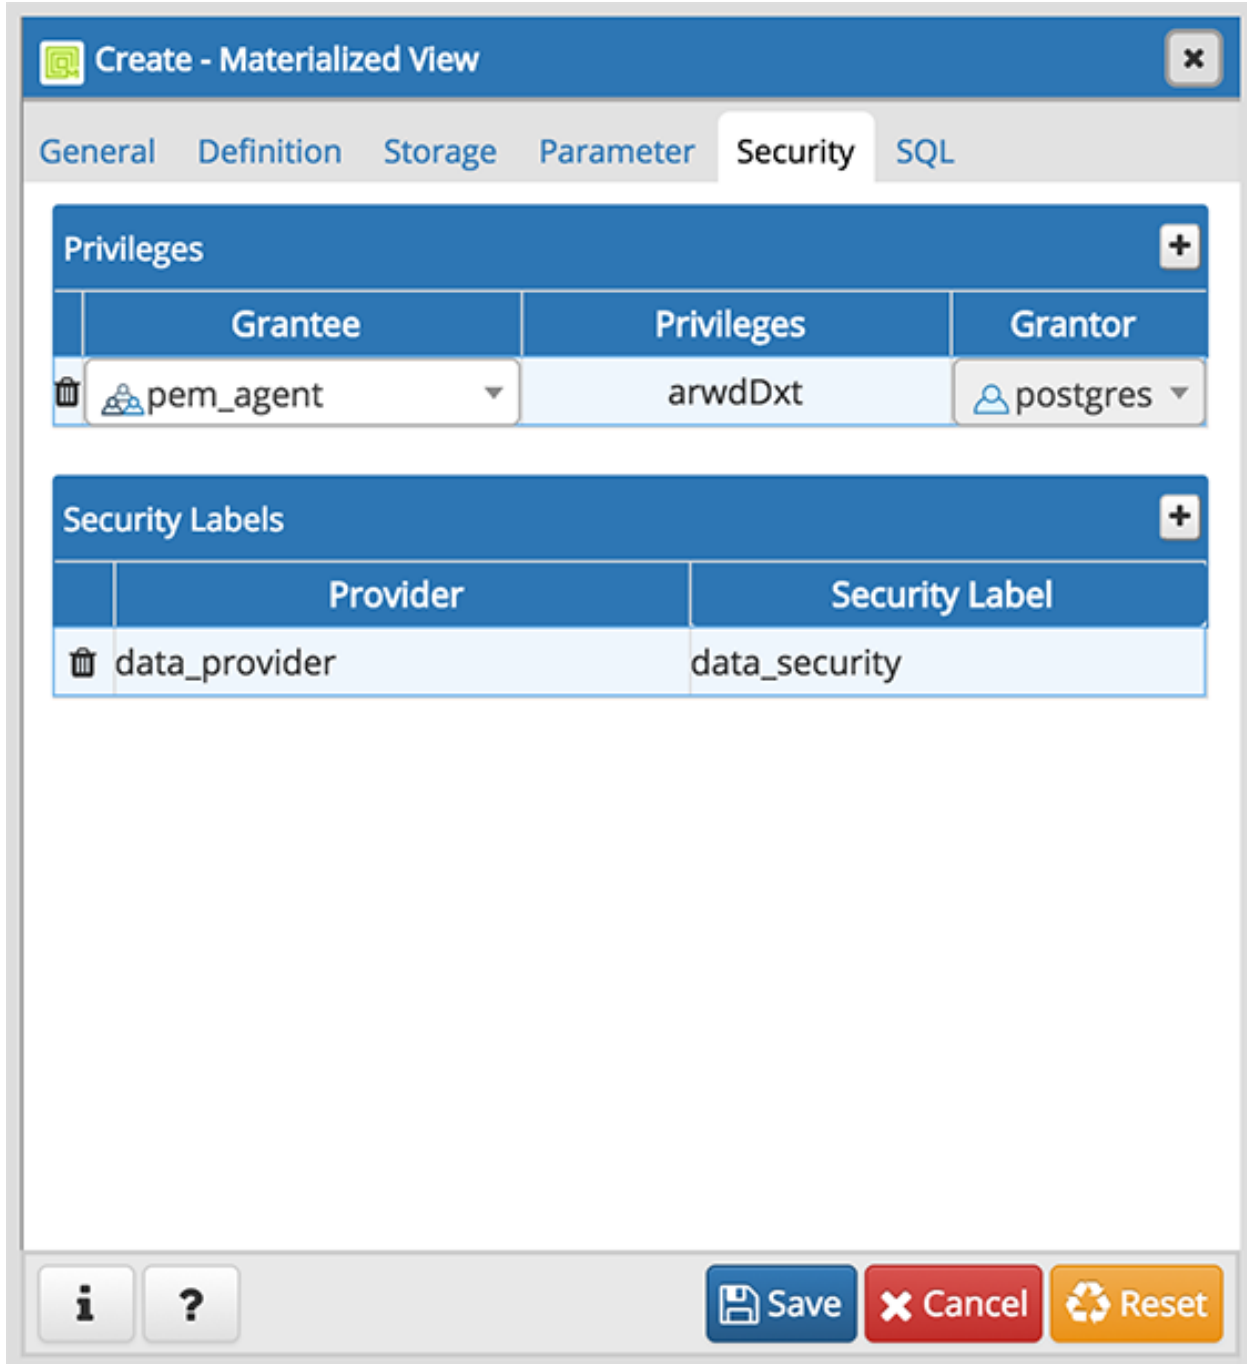
i
?

Save
 Cancel
 Reset

Use the tabs nested inside the *Parameter* tab to specify VACUUM and ANALYZE thresholds; use the *Table* tab and the *Toast Table* tab to customize values for the table and the associated toast table. To change the default values:

- Move the *Custom auto-vacuum?* switch to the *Yes* position to perform custom maintenance on the materialized view.
- Move the *Enabled?* switch to the *Yes* position to select values in the *Vacuum table*. Provide values for each row in the *Value* column.

Click the *Security* tab to continue.



Use the *Security* tab to assign privileges and define security labels.

Use the *Privileges* panel to assign privileges to a role. Click the *Add* icon (+) to set privileges for the materialized view:

- Select the name of the role from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.
- Select the name of the role from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click the *Add* icon (+) to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the materialized view. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

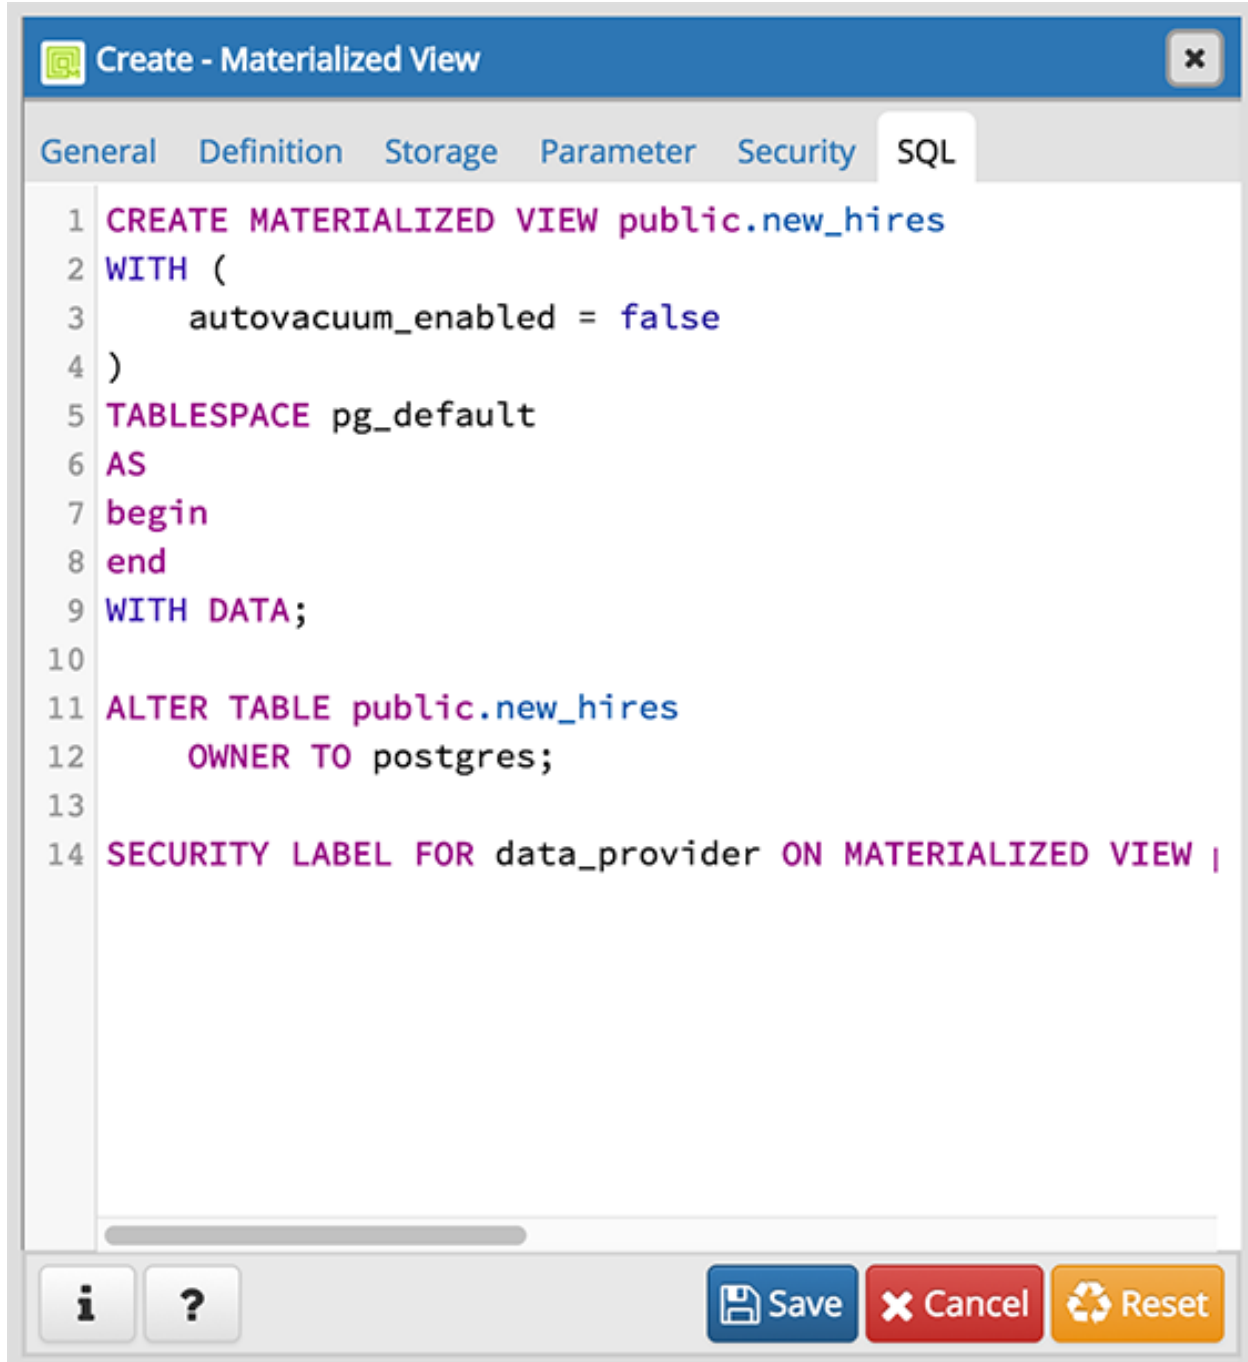
Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Materialized View* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Materialized View* dialog:




The example shown creates a query named *new_hires* that stores the result of the displayed query in the *pg_default* tablespace.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.17 Procedure Dialog

Use the *Procedure* dialog to create a procedure; procedures are supported by PostgreSQL v11+. The *Procedure* dialog allows you to implement options of the CREATE PROCEDURE command.

The *Procedure* dialog organizes the development of a procedure through the following dialog tabs: *General*, *Definition*, *Options*, *Arguments*, *Parameters*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Procedure' dialog box. The title bar is blue with a close button. The dialog has seven tabs: 'General', 'Definition', 'Options', 'Arguments', 'Parameters', 'Security', and 'SQL'. The 'General' tab is active. It contains the following fields:

- Name:** A text input field containing 'list_emp'.
- Owner:** A dropdown menu with a user icon, showing 'postgres' and a close button.
- Schema:** A dropdown menu with a database icon, showing 'public' and a close button.
- Comment:** A text area containing 'This procedure returns list of employees'.

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). On the left side of the bottom bar, there are two icons: an information icon (i) and a help icon (?).


Use the fields in the *General* tab to identify a procedure:

- Use the *Name* field to add a descriptive name for the procedure. The name will be displayed in the *pgAdmin* tree

control.

- Use the drop-down listbox next to *Owner* to select a role.
- Select the name of the schema in which the procedure will reside from the drop-down listbox in the *Schema* field.
- Store notes about the procedure in the *Comment* field.

Click the *Definition* tab to continue.



The screenshot shows the 'Create - Procedure' dialog box in pgAdmin 4, with the 'Definition' tab selected. The dialog has a blue header bar with the title 'Create - Procedure' and a close button. Below the header are several tabs: 'General', 'Definition', 'Options', 'Arguments', 'Parameters', 'Security', and 'SQL'. The 'Definition' tab is active, showing three fields: 'Return type' with the value '"char"[]', 'Language' with the value 'internal', and 'Code' with a text area containing SQL code. The SQL code is:

```
1 create table pem.test
2 (
3 coll integer,
4     primary key (coll)
5 )
```

 At the bottom of the dialog are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left side of the bottom bar.

Use the fields in the *Definition* tab to define the procedure:

- Use the drop-down listbox next to *Language* to select a language. The default is *sql*.
- Use the *Code* field to specify the code that will execute when the procedure is called.

Click the *Options* tab to continue.

The screenshot shows the 'Create - Procedure' dialog box with the 'Options' tab selected. The dialog contains the following fields and controls:

- Volatility:** A drop-down listbox with 'VOLATILE' selected.
- Returns a set?:** A checkbox with 'Yes' selected.
- Strict?:** A checkbox with 'Yes' selected.
- Security of definer?:** A checkbox with 'Yes' selected.
- Window?:** A checkbox with 'Yes' selected.
- Parallel:** A drop-down listbox with 'UNSAFE' selected.
- Estimated cost:** An empty text input field.
- Estimated rows:** An empty text input field.
- Leak proof?:** A checkbox with 'Yes' selected.

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left.

Use the fields in the *Options* tab to describe or modify the behavior of the procedure:

- Use the drop-down listbox under *Volatility* to select one of the following. *VOLATILE* is the default value.
 - *VOLATILE* indicates that the value can change even within a single table scan, so no optimizations can be made.

- *STABLE* indicates that the procedure cannot modify the database, and that within a single table scan it will consistently return the same result for the same argument values, but that its result could change across SQL statements.
- *IMMUTABLE* indicates that the procedure cannot modify the database and always returns the same result when given the same argument values.
- Move the *Strict?* switch to indicate if the procedure always returns NULL whenever any of its arguments are NULL. If *Yes*, the procedure is not executed when there are NULL arguments; instead a NULL result is assumed automatically. The default is *No*.
- Move the *Security of definer?* switch to specify that the procedure is to be executed with the privileges of the user that created it. The default is *No*.
- Use the *Estimated cost* field to specify a positive number representing the estimated execution cost for the procedure, in units of `cpu_operator_cost`. If the procedure returns a set, this is the cost per returned row.
- Move the *Leak proof?* switch to indicate whether the procedure has side effects — it reveals no information about its arguments other than by its return value. The default is *No*.

Click the *Arguments* tab to continue.

The screenshot shows the 'Create - Procedure' dialog box with the 'Arguments' tab selected. The dialog has a title bar 'Create - Procedure' and a close button. Below the title bar are tabs for 'General', 'Definition', 'Options', 'Arguments', 'Parameters', 'Security', and 'SQL'. The 'Arguments' tab contains a table with the following structure:

Arguments				
	Data Type	Mode	Argument Name	Default Value
	"char"			

At the bottom of the dialog are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields in the *Arguments* tab to define an argument. Click *Add* to set parameters and values for the argument:

- Use the drop-down listbox next to *Data type* to select a data type.
- Use the drop-down listbox next to *Mode* to select a mode. Select *IN* for an input parameter; select *OUT* for an output parameter; select *INOUT* for both an input and an output parameter; or, select *VARIADIC* to specify a VARIADIC parameter.
- Write a name for the argument in the *Argument Name* field.
- Specify a default value for the argument in the *Default Value* field.


Click *Add* to define another argument; to discard an argument, click the trash icon to the left of the row and confirm

deletion in the *Delete Row* popup.

Click the *Parameters* tab to continue.



The screenshot shows the 'Create - Procedure' dialog box in pgAdmin 4, with the 'Parameters' tab selected. The dialog has a title bar with a close button and a tabbed interface with 'General', 'Definition', 'Options', 'Arguments', 'Parameters', 'Security', and 'SQL'. The 'Parameters' tab contains a table with two columns: 'Name' and 'Value'. The 'Name' column has a dropdown menu with 'array_nulls' selected. The 'Value' column has a text input field containing 'False'. There is a trash icon to the left of the table row. At the bottom of the dialog, there are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

	Name	Value
	array_nulls	False

Use the fields in the *Parameters* tab to specify settings that will be applied when the procedure is invoked:

- Use the drop-down listbox next to *Parameter Name* in the *Parameters* panel to select a parameter.
- Click the *Add* button to add the variable to *Name* field in the table.
- Use the *Value* field to specify the value that will be associated with the selected variable. This field is context-sensitive.

Click the *Security* tab to continue.

The screenshot shows the 'Create - Procedure' dialog box in pgAdmin 4, with the 'Security' tab selected. The dialog contains two main sections:

Privileges

Grantee	Privileges	Grantor
pem_admin	X*	postgres

Security Labels

Provider	Security Label
myProvider	mySecurity

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left.

Use the *Security* tab to assign privileges and define security labels.

Use the *Privileges* panel to assign execute privileges for the procedure to a role:

- Select the name of the role from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.
- Select the name of the role from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click *Add* to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the procedure. Click *Add* to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

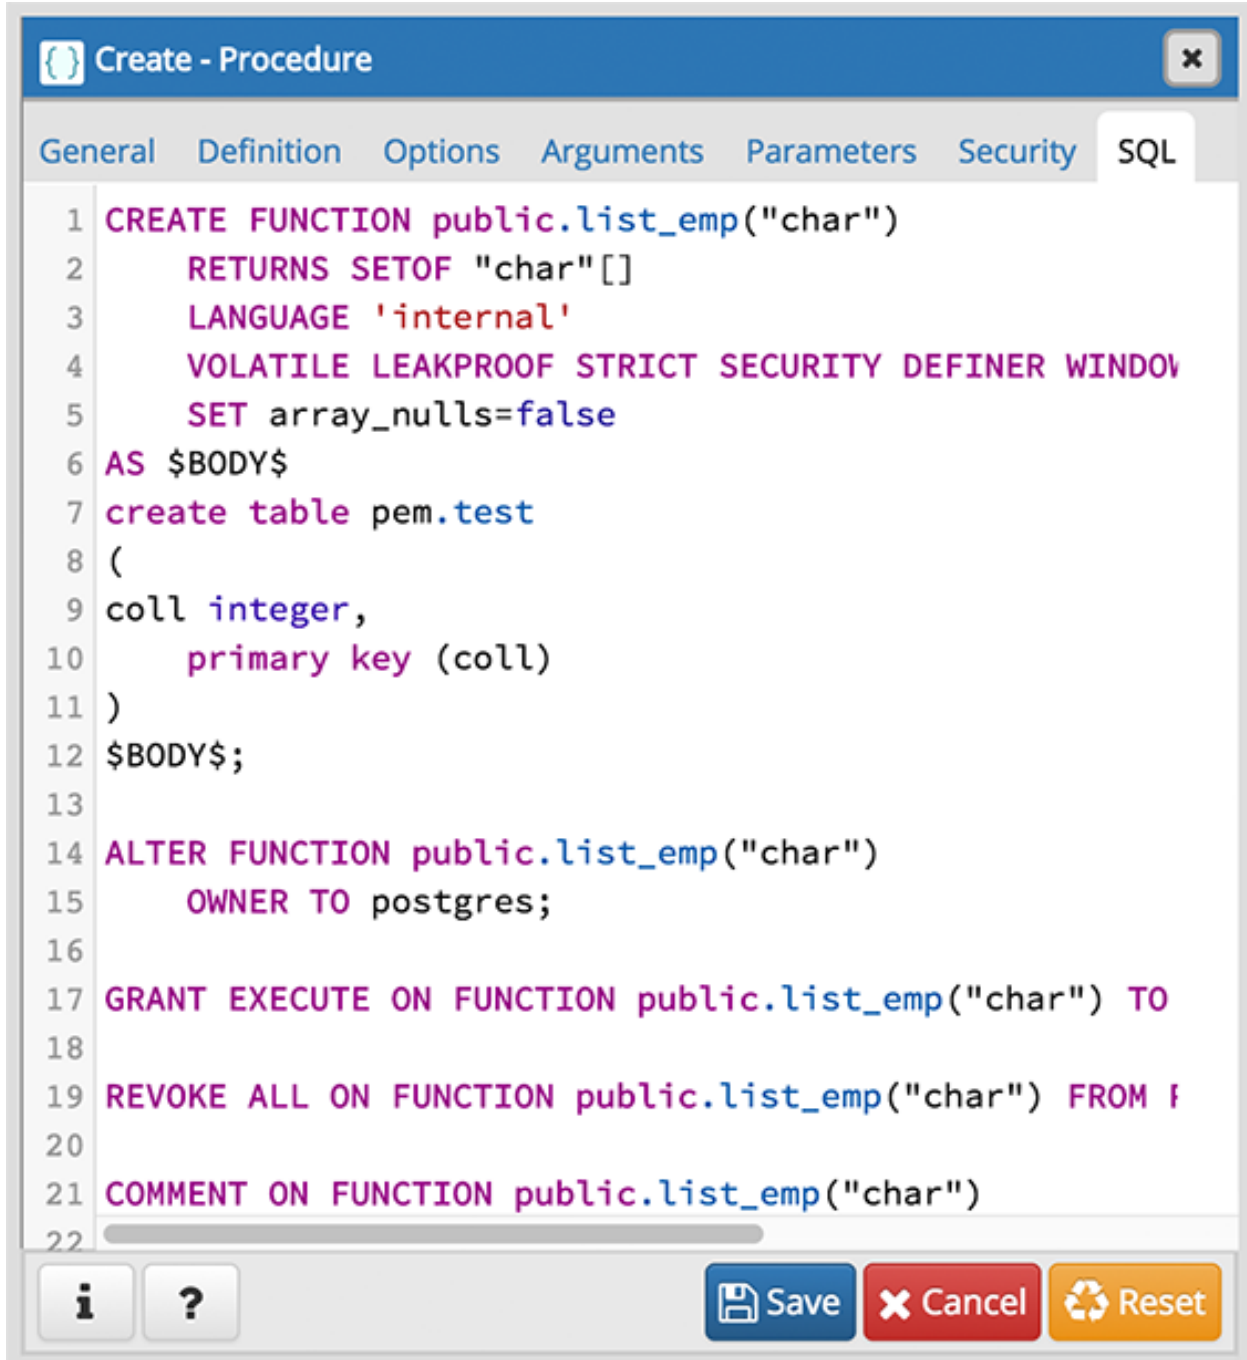
Click *Add* to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Procedure* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by selections made in the *Procedure* dialog:



The example demonstrates creating a procedure that returns a list of employees from a table named *emp*. The procedure is a SECURITY DEFINER, and will execute with the privileges of the role that defined the procedure.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.18 Schema Dialog

Use the *Schema* dialog to define a schema. A schema is the organizational workhorse of a database, similar to directories or namespaces. To create a schema, you must be a database superuser or have the CREATE privilege.

The *Schema* dialog organizes the development of schema through the following dialog tabs: *General* and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.

The screenshot shows the 'Create - Schema' dialog box. The title bar is blue and contains the PostgreSQL logo and the text 'Create - Schema'. Below the title bar are four tabs: 'General', 'Security', 'Default Privileges', and 'SQL'. The 'General' tab is selected and contains three fields: 'Name' with the value 'hr', 'Owner' with a dropdown menu showing 'postgres' and a user icon, and 'Comment' with a large empty text area. At the bottom of the dialog are four buttons: an information icon 'i', a question mark '?', a blue 'Save' button with a floppy disk icon, a red 'Cancel' button with an 'X' icon, and an orange 'Reset' button with a circular arrow icon.

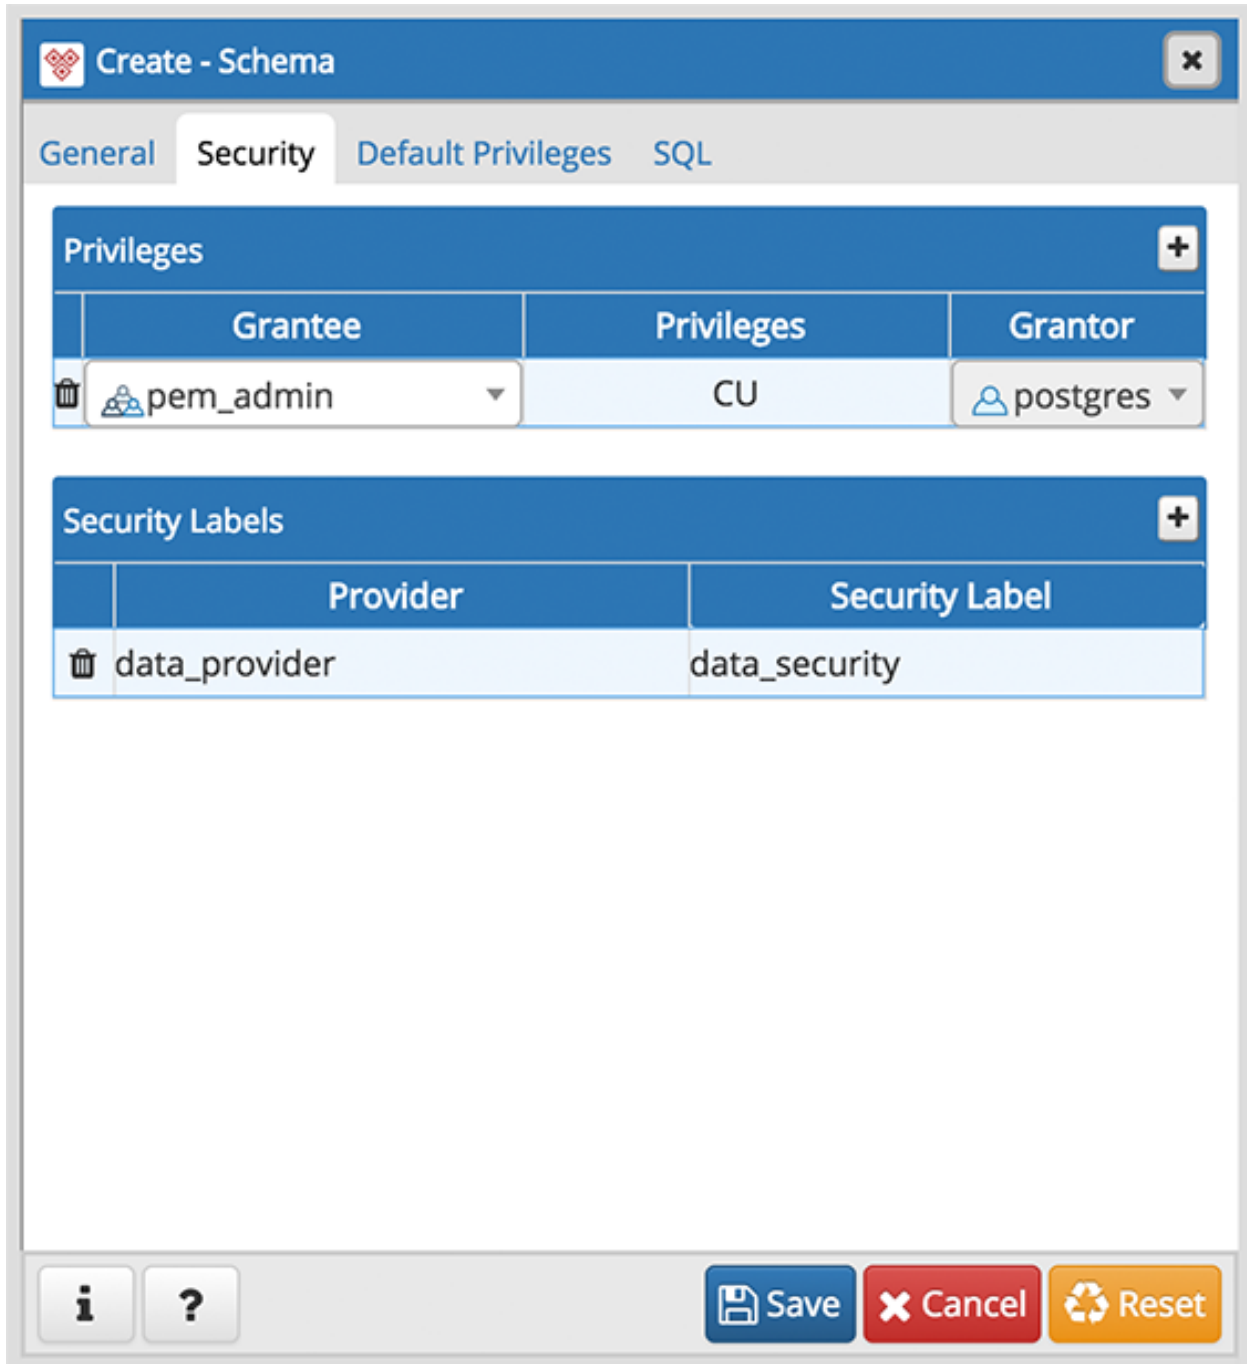
Use the fields on the *General* tab to identify the schema.

- Use the *Name* field to add a descriptive name for the schema. The name will be displayed in the *pgAdmin* tree

control.

- Select the owner of the schema from the drop-down listbox in the *Owner* field.
- Store notes about the schema in the *Comment* field.

Click the *Security* tab to continue.



Use the *Security* tab to assign privileges and security labels for the schema.

Click the *Add* icon (+) to assign a set of privileges in the *Privileges* panel:

- Select the name of the role from the drop-down listbox in the *Grantee* field.

- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privileges to the specified user.
- Select the name of the role that is granting the privilege from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the schema.

Click the *Add* icon (+) to assign additional sets of privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Add* icon (+) to assign a security label in the *Security Labels* panel:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Default Privileges* tab to continue.



Use the *Default Privileges* tab to grant privileges for tables, sequences, functions and types. Use the tabs nested inside the *Default Privileges* tab to specify the database object and click the *Add* icon (+) to assign a set of privileges:

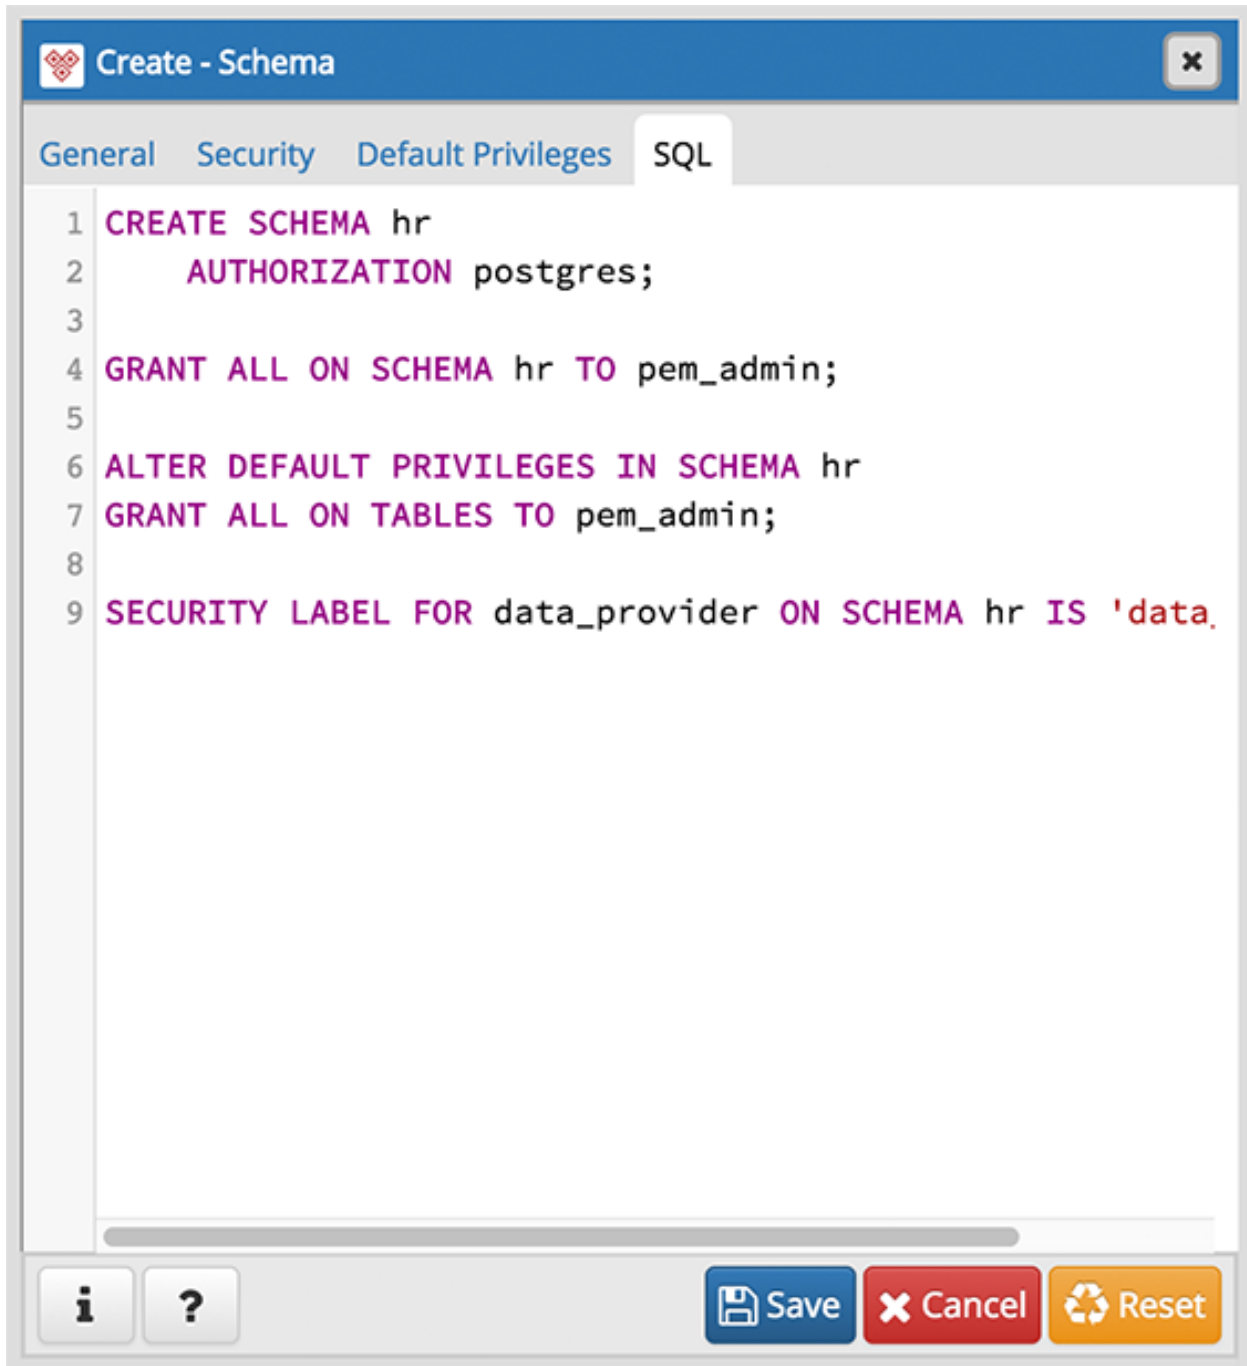
- Select the name of a role that will be granted privileges in the schema from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privileges to the specified user.
- Select the name of the role that is granting the privilege from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the schema.

Click the *SQL* tab to continue.

Your entries in the *Schema* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by selections made in the *Schema* dialog:



The example creates a schema named *hr*; the command grants *USAGE* privileges to *public* and assigns the ability to grant privileges to *alice*.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new


tab displays the PostgreSQL core documentation.

- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.19 Sequence Dialog

Use the *Sequence* dialog to create a sequence. A sequence generates unique values in a sequential order (not necessarily contiguous).

The *Sequence* dialog organizes the development of a sequence through the following dialog tabs: *General*, *Definition*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Sequence' dialog box in pgAdmin 4. The dialog has a blue title bar with the text '1..3 Create - Sequence' and a close button. Below the title bar are four tabs: 'General', 'Definition', 'Security', and 'SQL'. The 'General' tab is selected. The form contains the following fields:

- Name:** A text input field containing the text 'seconds'.
- Owner:** A dropdown menu showing 'postgres' with a user icon and a close button.
- Schema:** A dropdown menu showing 'public' with a schema icon and a close button.
- Comment:** A large, empty text area for entering notes.

At the bottom of the dialog, there are three buttons: an information icon, a question mark icon, a blue 'Save' button, a red 'Cancel' button, and an orange 'Reset' button.

Use the fields in the *General* tab to identify a sequence:

- Use the *Name* field to add a descriptive name for the sequence. The name will be displayed in the *pgAdmin* tree control. The sequence name must be distinct from the name of any other sequence, table, index, view, or foreign table in the same schema.
- Use the drop-down listbox next to *Owner* to select the name of the role that will own the sequence.
- Use the drop-down listbox next to *Schema* to select the schema in which the sequence will reside.
- Store notes about the sequence in the *Comment* field.

Click the *Definition* tab to continue.

1..3 Create - Sequence

General Definition Security SQL

Increment

Start

Minimum

Maximum

Cache

Cycled Yes

i **?** **Save** **Cancel** **Reset**

Use the fields in the *Definition* tab to define the sequence:

- Use the *Increment* field to specify which value is added to the current sequence value to create a new value.
- Provide a value in the *Start* field to specify the beginning value of the sequence. The default starting value is MINVALUE for ascending sequences and MAXVALUE for descending ones.
- Provide a value in the *Minimum* field to specify the minimum value a sequence can generate. If this clause is not supplied or NO MINVALUE is specified, then defaults will be used. The defaults are 1 and -263-1 for ascending and descending sequences, respectively.
- Provide a value in the *Maximum* field to specify the maximum value for the sequence. If this clause is not

supplied or NO MAXVALUE is specified, then default values will be used. The defaults are 263-1 and -1 for ascending and descending sequences, respectively.

- Provide a value in the *Cache* field to specify how many sequence numbers are to be preallocated and stored in memory for faster access. The minimum value is 1 (only one value can be generated at a time, i.e., no cache), and this is also the default.
- Move the *Cycled* switch to the *Yes* position to allow the sequence to wrap around when the MAXVALUE or the MINVALUE has been reached by an ascending or descending sequence respectively. If the limit is reached, the next number generated will be the MINVALUE or MAXVALUE, respectively. The default is *No*.

Click the *Security* tab to continue.

1..3 Create - Sequence
✕

General
Definition
Security
SQL

Privileges
+

	Grantee	Privileges	Grantor
✕	pem_admin ▾	rwU	postgres ▾

Security Labels
+

	Provider	Security Label
✕	data_provider	data_security

i
?

Save
Cancel
Reset

Use the *Security* tab to assign privileges and define security labels for the sequence.

Use the *Privileges* panel to assign privileges. Click the *Add* icon (+) to set privileges:

- Select the name of a role that will be granted privileges from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.
- Select the name of the role from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click the *Add* icon (+) to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the sequence. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

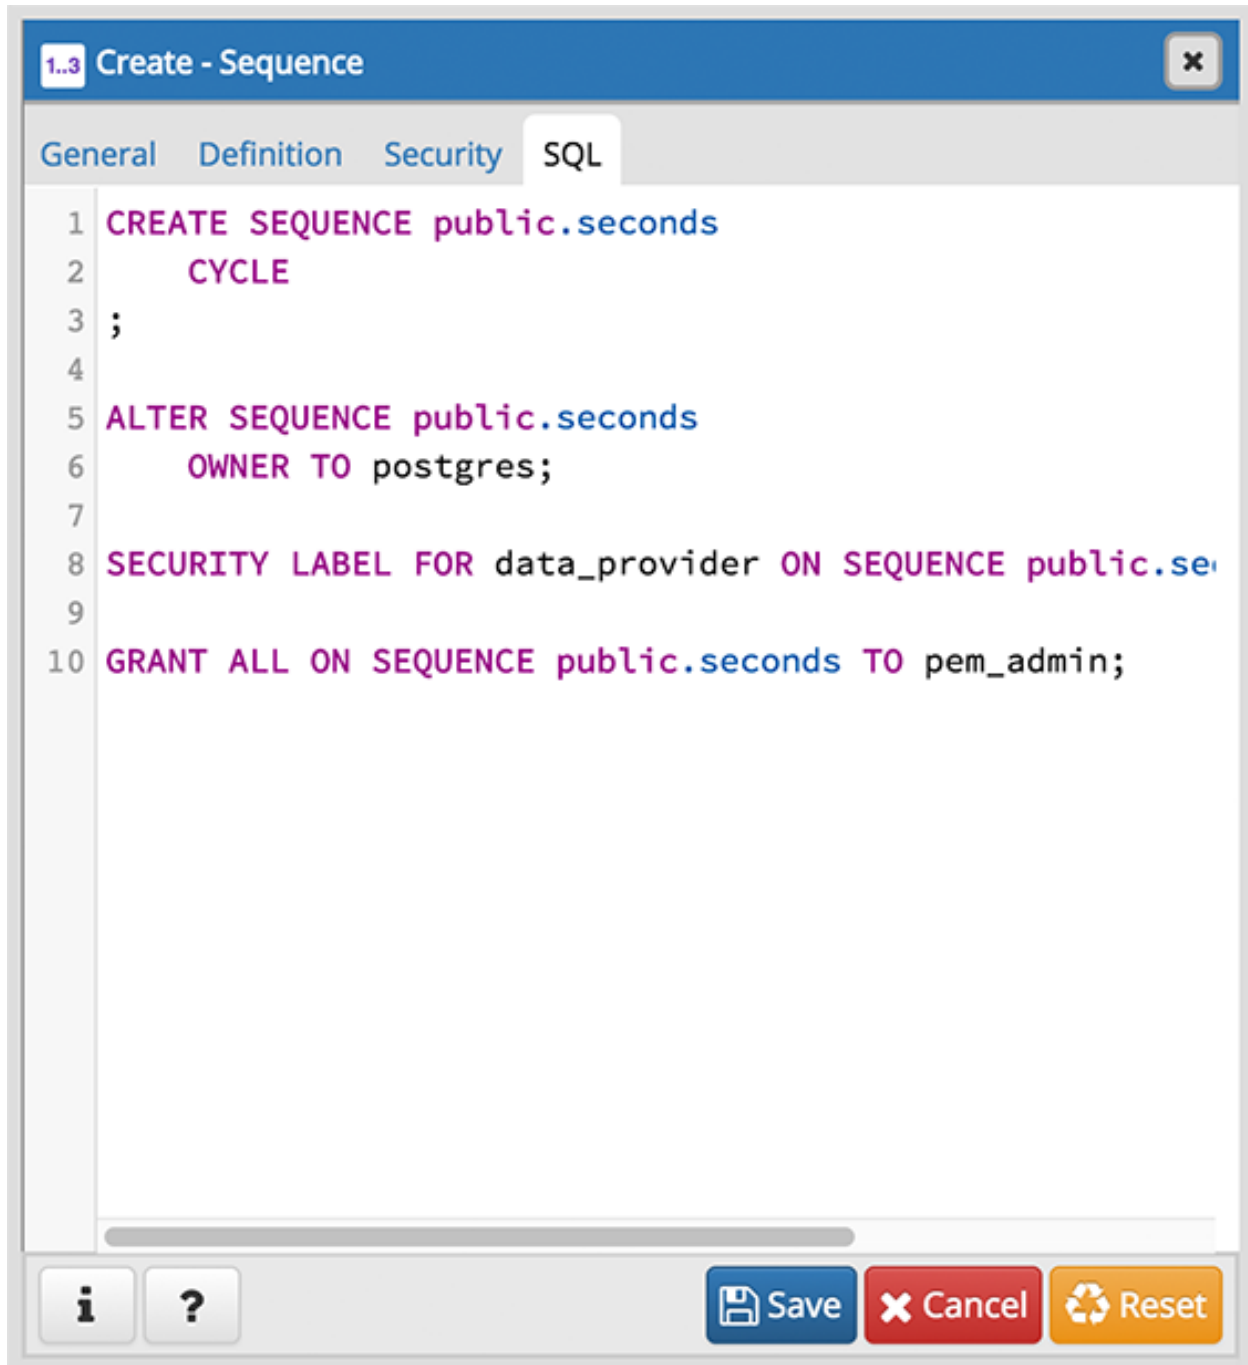
Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Sequence* dialog generate a generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Sequence* dialog:




The example shown demonstrates a sequence named *seconds*. The sequence will increase in 5 second increments, and stop when it reaches a maximum value equal of 60.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.20 Trigger function Dialog

Use the *Trigger function* dialog to create or manage a `trigger_function`. A trigger function defines the action that will be invoked when a trigger fires.

The *Trigger function* dialog organizes the development of a trigger function through the following dialog tabs: *General*, *Definition*, *Options*, *Parameters* and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the "Create - Trigger function" dialog box. The title bar is blue with a close button (X) on the right. Below the title bar are six tabs: "General", "Definition", "Options", "Parameters", "Security", and "SQL". The "General" tab is selected and contains the following fields:

- Name:** A text input field with a vertical cursor.
- Owner:** A dropdown menu showing "postgres" with a user icon and a close button (X).
- Schema:** A dropdown menu showing "public" with a database icon and a close button (X).
- Comment:** A large text area for entering a comment.

At the bottom of the dialog, there are five buttons: an information icon (i), a help icon (?), a "Save" button with a floppy disk icon, a "Cancel" button with a red background and a close icon (X), and a "Reset" button with a circular arrow icon.

Use the fields in the *General* tab to identify the trigger function:

- Use the *Name* field to add a descriptive name for the trigger function. The name will be displayed in the *pgAdmin*

tree control. Please note that trigger functions will be invoked in alphabetical order.

- Use the drop-down listbox next to *Owner* to select the role that will own the trigger function.
- Select the name of the schema in which the trigger function will reside from the drop-down listbox in the *Schema* field.
- Store notes about the trigger function in the *Comment* field.

Click the *Definition* tab to continue.



The screenshot shows a dialog box titled "Create - Trigger function" with a close button (X) in the top right corner. The dialog has five tabs: "General", "Definition", "Options", "Parameters", and "SQL". The "Definition" tab is selected and active. It contains three fields:


- Return type:** A dropdown menu with "trigger" selected.
- Language:** A dropdown menu with "Select from the list" selected.
- Code:** A text area containing the number "1".

At the bottom of the dialog, there are three buttons: "Save" (blue), "Cancel" (red), and "Reset" (orange). To the left of these buttons are two smaller buttons: "i" (information) and "?" (help).

Use the fields in the *Definition* tab to define the trigger function:

- Use the drop-down listbox next to *Return type* to specify the pseudotype that is associated with the trigger function:
 - Select *trigger* if you are creating a DML trigger.
 - Select *event_trigger* if you are creating a DDL trigger.
- Use the drop-down listbox next to *Language* to select the implementation language. The default is *plpgsql*.
- Use the *Code* field to write the code that will execute when the trigger function is called.

Click the *Options* tab to continue.



The screenshot shows the 'Create - Trigger function' dialog box with the 'Options' tab selected. The dialog has a title bar with a close button (X) and a tabbed interface with the following tabs: General, Definition, Options (selected), Parameters, Security, and SQL. The 'Options' tab contains the following fields:


- Volatility:** A drop-down menu with the text 'Select from the list' and a downward arrow.
- Returns a set?:** A checkbox with a blue 'No' button.
- Strict?:** A checkbox with a blue 'No' button.
- Security of definer?:** A checkbox with a blue 'No' button.
- Window?:** A checkbox with a blue 'No' button.
- Estimated cost:** An empty text input field.
- Estimated rows:** A greyed-out text input field.
- Leak proof?:** A checkbox with a blue 'No' button.

At the bottom of the dialog, there are three buttons: an information icon (i), a question mark icon (?), and three action buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange).

Use the fields in the *Options* tab to describe or modify the action of the trigger function:

- Use the drop-down listbox next to *Volatility* to select one of the following:
 - *VOLATILE* indicates that the trigger function value can change even within a single table scan.
 - *STABLE* indicates that the trigger function cannot modify the database, and that within a single table scan it will consistently return the same result for the same argument values.
 - *IMMUTABLE* indicates that the trigger function cannot modify the database and always returns the same result when given the same argument values.
- Move the *Returns a Set?* switch to indicate if the trigger function returns a set that includes multiple rows. The default is *No*.
- Move the *Strict?* switch to indicate if the trigger function always returns NULL whenever any of its arguments are NULL. If *Yes*, the function is not executed when there are NULL arguments; instead a NULL result is assumed automatically. The default is *No*.
- Move the *Security of definer?* switch to specify that the trigger function is to be executed with the privileges of the user that created it. The default is *No*.
- Move the *Window?* switch to indicate that the trigger function is a window function rather than a plain function. The default is *No*. This is currently only useful for trigger functions written in C.
- Use the *Estimated cost* field to specify a positive number representing the estimated execution cost for the trigger function, in units of `cpu_operator_cost`. If the function returns a set, this is the cost per returned row.
- Use the *Estimated rows* field to specify a positive number giving the estimated number of rows that the query planner should expect the trigger function to return. This is only allowed when the function is declared to return a set. The default assumption is 1000 rows.
- Move the *Leak proof?* switch to indicate whether the trigger function has side effects. The default is *No*. This option can only be set by the superuser.

Click the *Parameters* tab to continue.



The screenshot shows the 'Create - Trigger function' dialog box. The 'Parameters' tab is selected, displaying a table with two columns: 'Name' and 'Value'. The table is currently empty. An 'Add' icon (+) is located in the top right corner of the table area. At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two smaller buttons: an information icon (i) and a help icon (?).

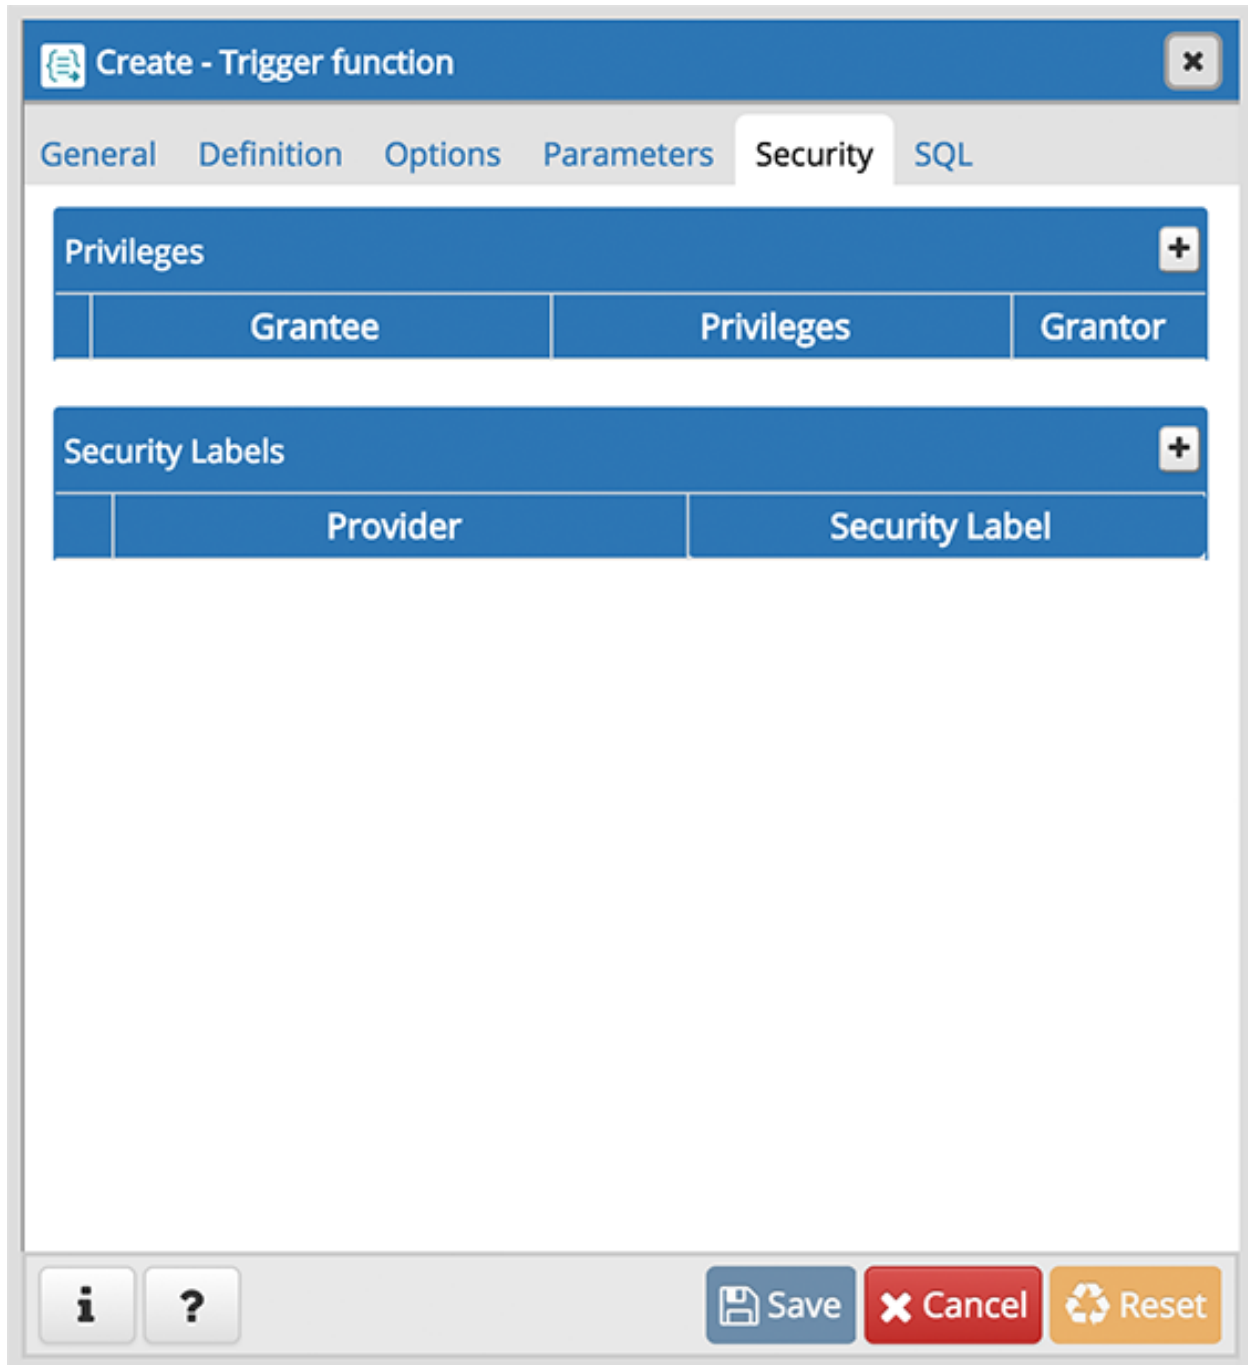
Parameters	Name	Value

Use the fields in the *Parameters* tab to specify settings that will be applied when the trigger function is invoked. Click the *Add* icon (+) to add a *Name/Value* pair to the table below.

- Use the drop-down listbox in the *Name* field to select a parameter.
- Use the *Value* field to specify the value that will be associated with the selected parameter. This field is context-sensitive.

Click the *Add* icon (+) to set additional parameters; to discard a parameter, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Security* tab to continue.



Use the *Security* tab to assign privileges and define security labels.

Use the *Privileges* panel to assign usage privileges for the trigger function to a role. Click the *Add* icon (+) to add a role to the table.

- Select the name of the role from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.
- Select the name of a role from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click the *Add* icon (+) to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the trigger function. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

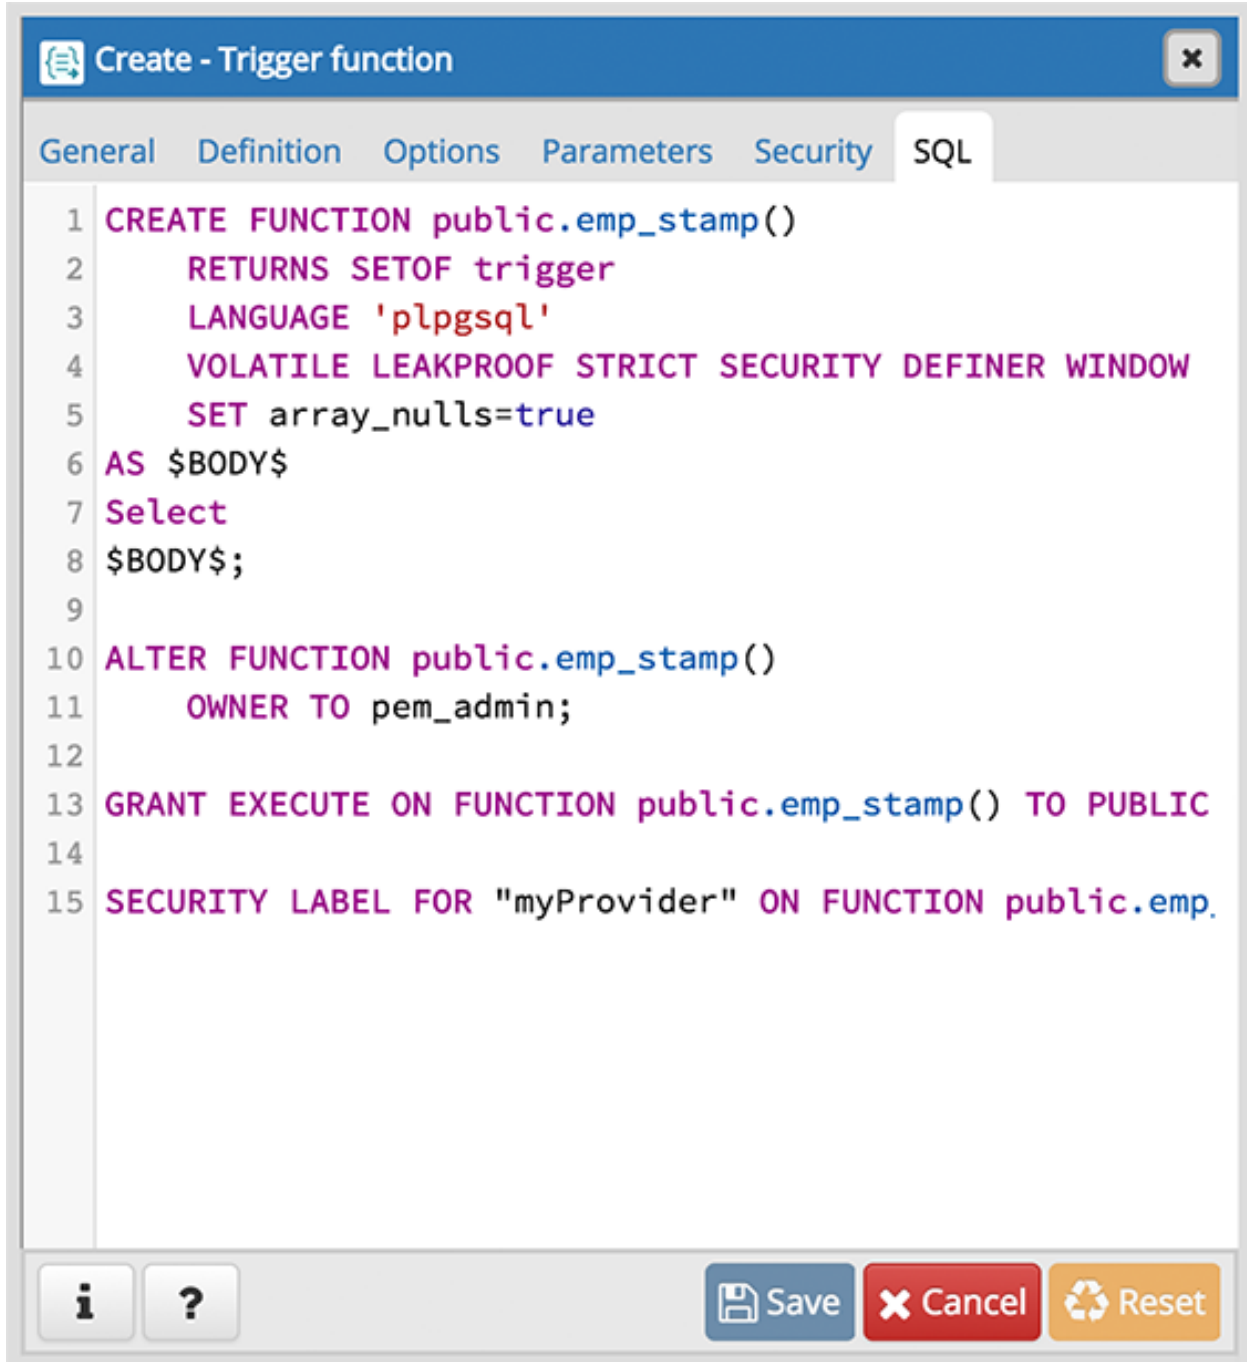
Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Trigger function* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit other tabs to modify the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Trigger function* dialog:



The example shown demonstrates creating a trigger function named *emp_stamp* that checks for a new employee's name, and checks that the employee's salary is a positive value.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.21 Type Dialog

Use the *Type* dialog to register a custom data type.

The *Type* dialog organizes the development of a data type through the following dialog tabs: *General*, *Definition*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Type' dialog box. The title bar is blue with a close button. Below the title bar are four tabs: 'General', 'Definition', 'Security', and 'SQL'. The 'General' tab is selected. The 'Name' field contains 'work_order'. The 'Owner' field is a dropdown menu with 'postgres' selected. The 'Schema' field is a dropdown menu with 'public' selected. The 'Comment' field is a large empty text area. At the bottom, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also information and help icons on the left.

Use the fields in the *General* tab to identify the custom data type:

- Use the *Name* field to add a descriptive name for the type. The name will be displayed in the *pgAdmin* tree control. The type name must be distinct from the name of any existing type, domain, or table in the same

schema.

- Use the drop-down listbox next to *Owner* to select the role that will own the type.
- Select the name of the schema in which the type will reside from the drop-down listbox in the *Schema* field.
- Store notes about the type in the *Comments* field.

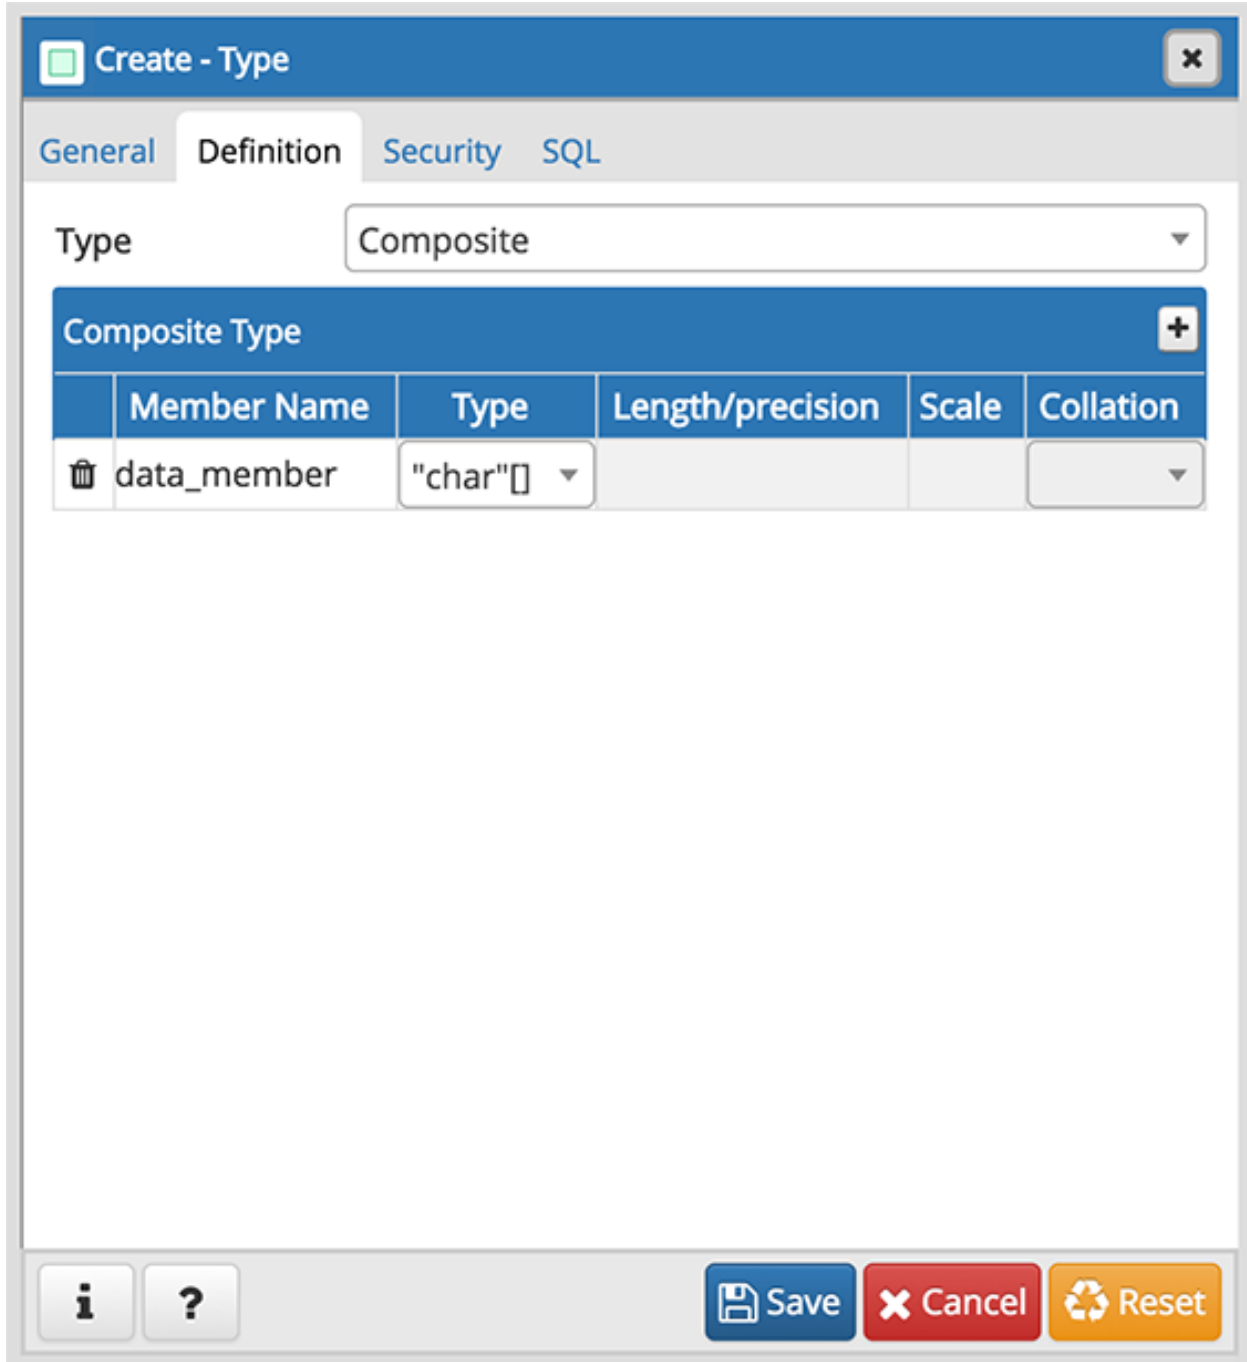
Click the *Definition* tab to continue.

Select a data type from the drop-down listbox next to *Type* on the *Definition* tab; the panel below changes to display the options appropriate for the selected data type. Use the fields in the panel to define the data type.

There are five data types:

- *Composite Type*
- *Enumeration Type*
- *Range Type*
- *External Type* (or *Base Type*)
- *Shell Type*

If you select *Composite* in the *Type* field, the *Definition* tab displays the *Composite Type* panel:

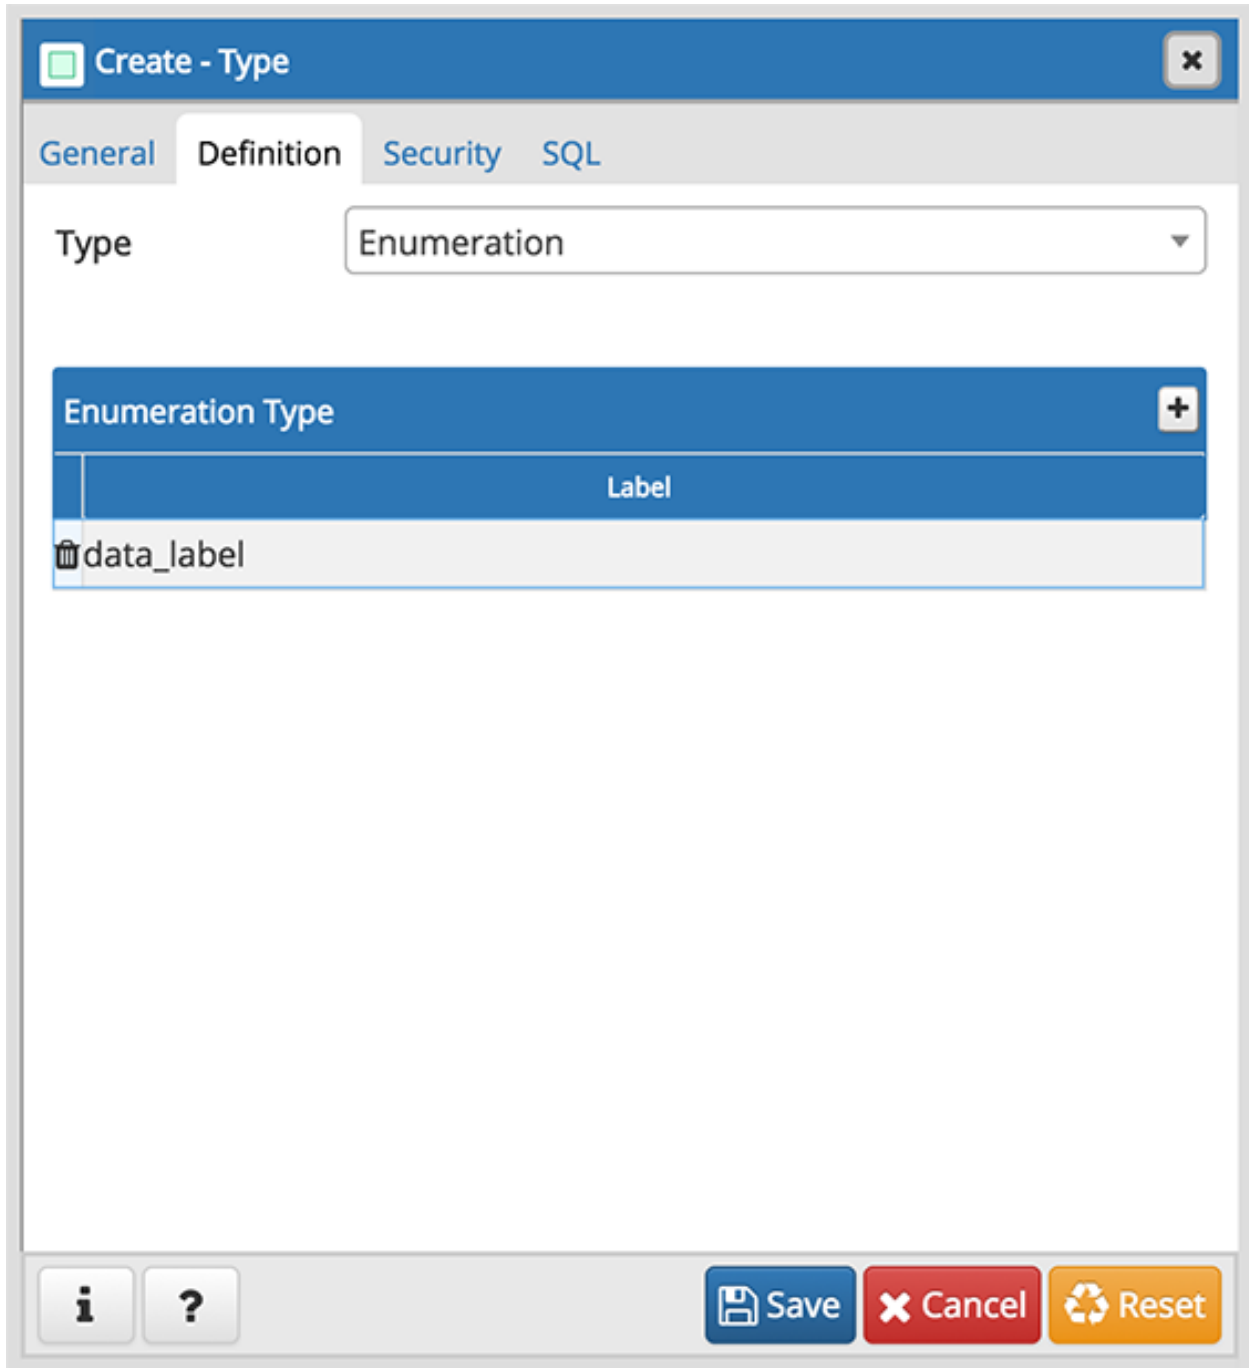


Click the *Add* icon (+) to provide attributes of the type. Fields on the *General* panel are context sensitive and may be disabled.

- Use the *Member Name* field to add an attribute name.
- Use the drop-down listbox in the *Type* field to select a datatype.
- Use the *Length/Precision* field to specify the maximum length of a non-numeric type, or the total count of significant digits in a numeric type.
- Use the *Scale* field to specify the number of digits to the right of the decimal point.
- Use the drop-down listbox in the *Collation* field to select a collation (if applicable).

Click the *Add* icon (+) to define an additional member; click the trash icon to the left of the row to discard a row.

If you select the *Enumeration* in the *Type* field, the *Definition* tab displays the *Enumeration Type* panel:



Click the *Add* icon (+) to provide a label for the type.

- Use the *Label* field to add a label, which must be less than 64 bytes long.

Click the *Add* icon (+) after each selection to create additional labels; to discard a label, click the trash icon to the left of the row.

If you select *External*, the *Definition* tab displays the *External Type* panel:

Create - Type

General Definition Security SQL

Type External

Required Optional-1 Optional-2

Input function pg_catalog.domain_in

Output function pg_catalog.numeric_sqrt

i ? Save Cancel Reset

On the *Required* tab:

- Use the drop-down listbox next to the *Input function* field to add an `input_function`. The `input_function` converts the type's external textual representation to the internal representation used by the operators and functions defined for the type.
- Use the drop-down listbox next to the *Output function* field to add an `output_function`. The `output_function` converts the type's internal representation used by the operators and functions defined for the type to the type's external textual representation.

On the *Optional-1* tab:

- Use the drop-down listbox next to the optional *Receive Function* field to select a `receive_function`. The optional `receive_function` converts the type's external binary representation to the internal representation. If this function is not supplied, the type cannot participate in binary input.
- Use the drop-down listbox next to the optional *Send function* field to select a `send_function`. The optional `send_function` converts from the internal representation to the external binary representation. If this function is not supplied, the type cannot participate in binary output.
- Use the drop-down listbox next to the optional *Typmod in function* field tab to select a `type_modifier_input_function`.
- Use the drop-down listbox next to the optional *Typmod out function* field tab to select a `type_modifier_output_function`. It is allowed to omit the `type_modifier_output_function`, in which case the default display format is the stored typmod integer value enclosed in parentheses.
- Use the optional *Internal length* to specify a value for internal representation.
- Move the *Variable?* switch to specify the internal representation is of variable length (VARIABLE). The default is a fixed length positive integer.
- Specify a default value in the optional *Default* field in cases where a column of the data type defaults to something other than the null value. Specify the default with the DEFAULT key word. (A default can be overridden by an explicit DEFAULT clause attached to a particular column.)
- Use the drop-down listbox next to the optional *Analyze function* field to select a function for performing type-specific statistics collection for columns of the data type.
- Use the drop-down listbox next to the optional *Category type* field to help control which implicit cast will be applied in ambiguous situations.
- Move the *Preferred?* switch to *Yes* to specify the selected category type is preferred. The default is *No*.

On the *Optional-2* tab:

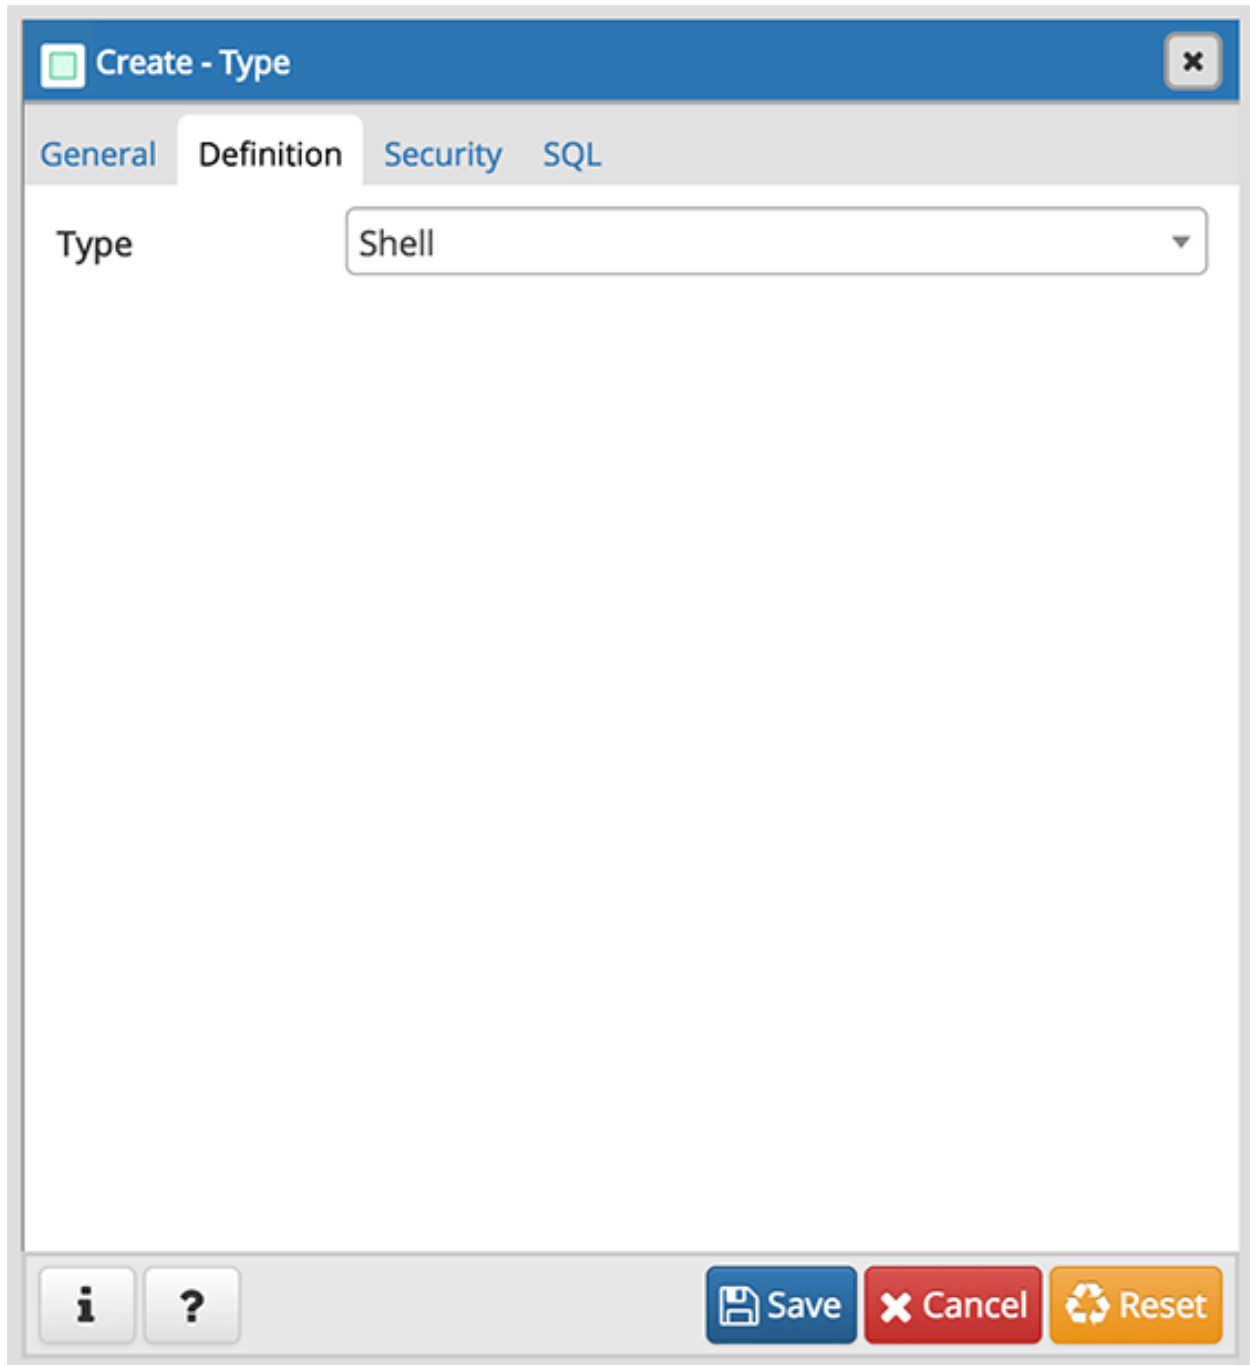
- Use the drop-down listbox next to the optional *Element type* field to specify a data type.
- Use the optional *Delimiter* field to indicate the delimiter to be used between values in the external representation of arrays for this data type. The default delimiter is the comma (.). Note that the delimiter is associated with the array element type, not the array type itself.
- Use the drop-down listbox next to *Alignment type* to specify the storage alignment required for the data type. The allowed values (char, int2, int4, and double) correspond with alignment on 1, 2, 4, or 8 byte boundaries.
- Use the drop-down listbox next to optional *Storage type* to select a strategy for storing data.
- Move the *Passed by value?* switch to *Yes* to override the existing data type value. The default is *No*.
- Move the *Collatable?* switch to *Yes* to specify column definitions and expressions of the type may carry collation information through use of the COLLATE clause. The default is *No*.

If you select *Range* in the *Type* field, the *Definition* tab displays the *Range* panel. Fields on the *Range* panel are context-sensitive and may be disabled.

The screenshot shows the 'Create - Type' dialog box with the 'Definition' tab selected. The 'Type' field is set to 'Range'. The 'Subtype' field contains 'abstime'. The 'Subtype operator class', 'Collation', 'Canonical function', and 'Subtype diff function' fields are currently empty. The 'Save', 'Cancel', and 'Reset' buttons are visible at the bottom.

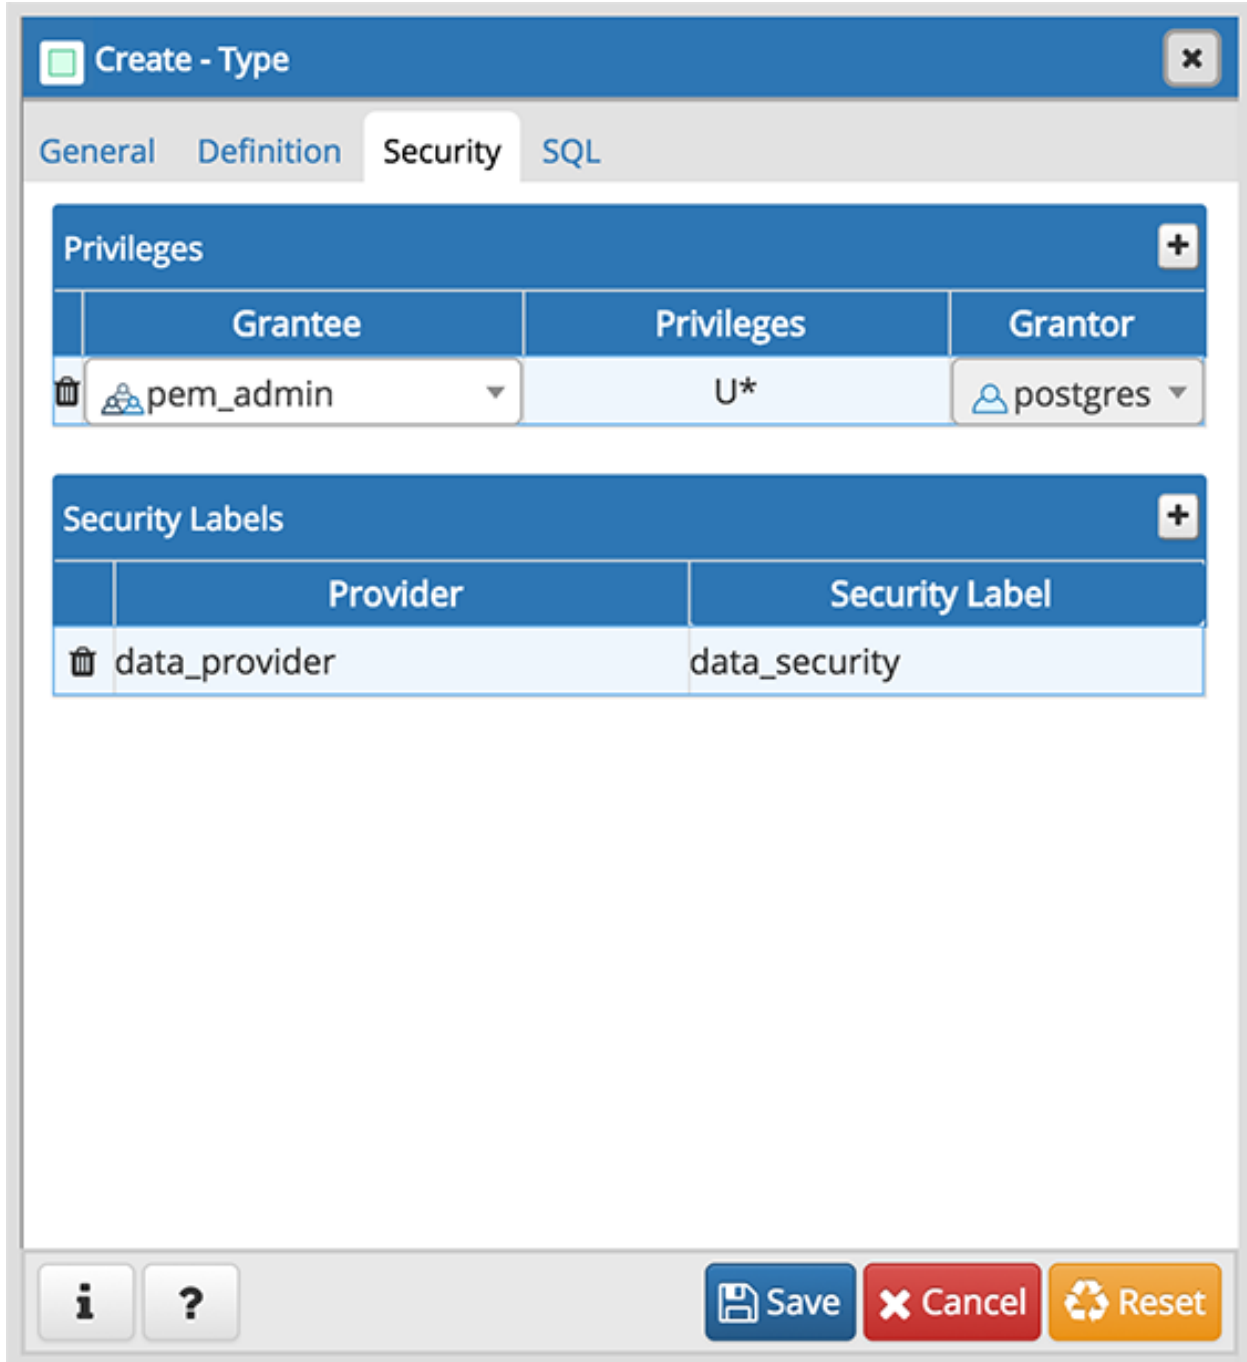
- Use the drop-down listbox next to *Sub-type* to select an associated b-tree operator class (to determine the ordering of values for the range type).
- Use the drop-down listbox next to *Sub-type operator class* to use a non-default operator class.
- Use the drop-down listbox next to *Collation* to use a non-default collation in the range's ordering if the sub-type is collatable.
- Use the drop-down listbox next to *Canonical function* to convert range values to a canonical form.
- Use the drop-down listbox next to *Sub-type diff function* to select a user-defined subtype_diff function.

If you select *Shell* in the *Type* field, the *Definition* tab displays the *Shell* panel:



A shell type is a placeholder for a type and has no parameters.

Click the *Security* tab to continue.



Use the *Security* tab to assign privileges and define security labels.

Use the *Privileges* panel to assign privileges for the type; click the *Add* icon (+) to grant privileges:

- Select the name of the role that will be granted privileges on the type from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.
- Select the name of the role that is granting privileges from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click the *Add* icon (+) to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the type. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

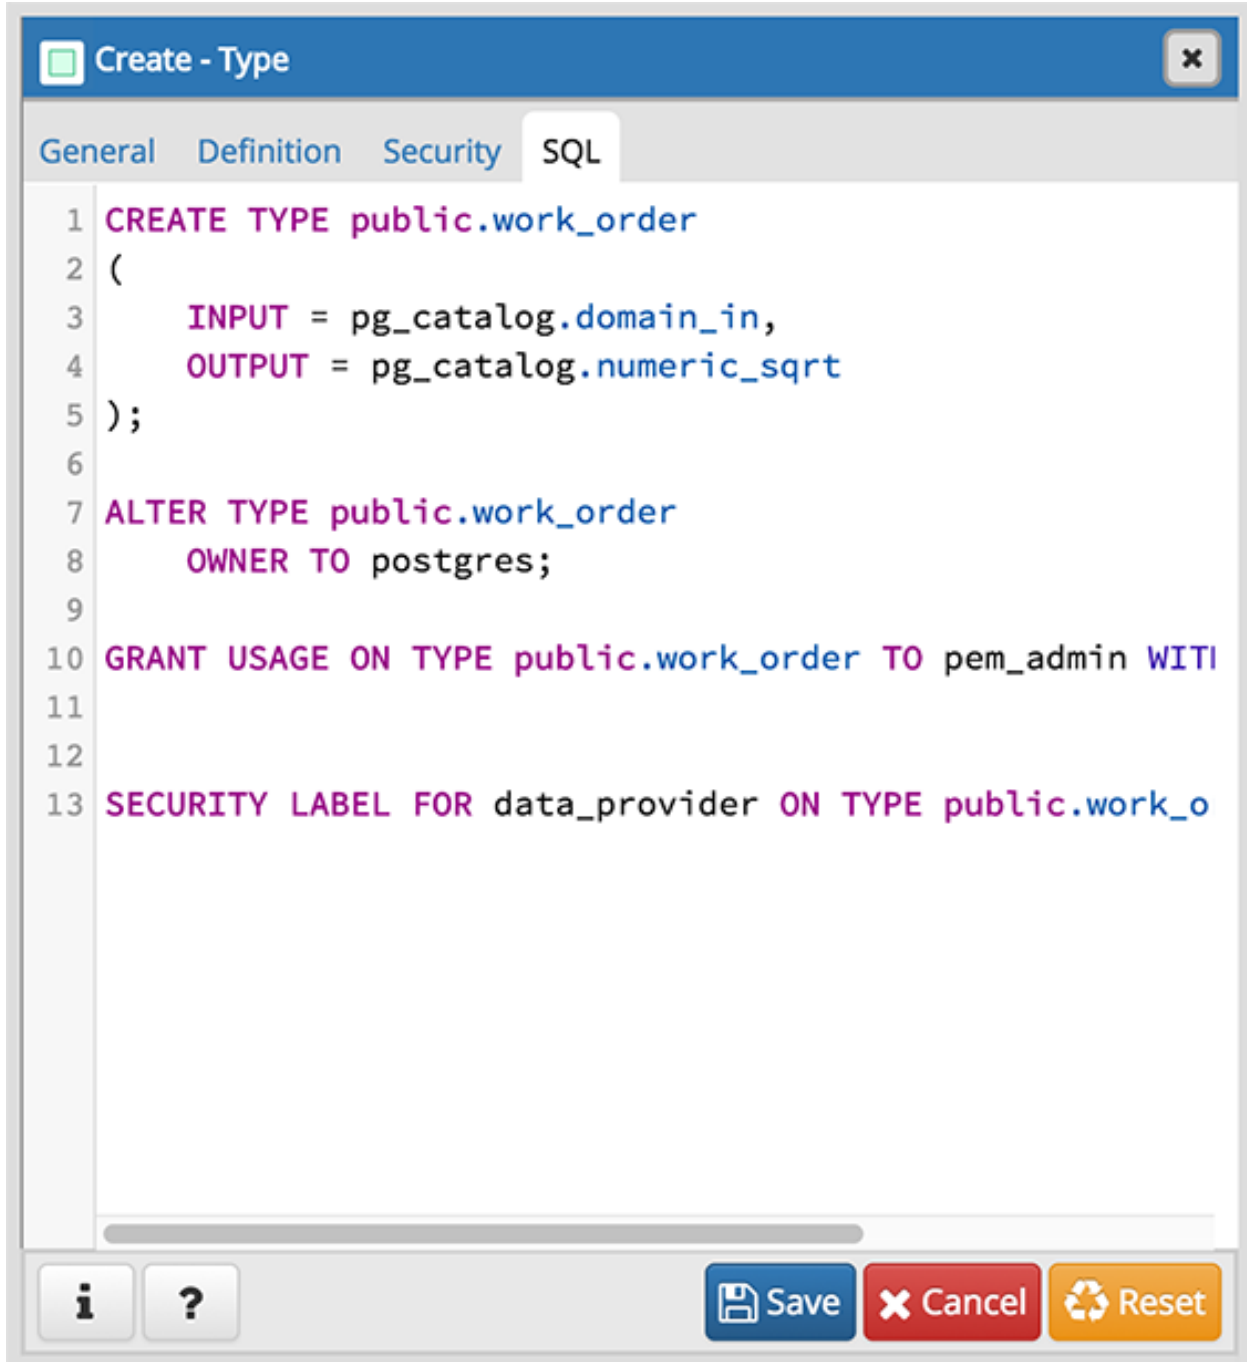
Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Type* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of a sql command generated by user selections made in the *Type* dialog:




The example shown demonstrates creating a data type named `work_order`. The data type is an enumerated type with three labels: `new`, `open` and `closed`.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.22 User Mapping Dialog

Use the *User Mapping* dialog to define a new mapping of a user to a foreign server.

The *User Mapping* dialog organizes the development of a user mapping through the following dialog tabs: *General* and *Options*. The *SQL* tab displays the SQL code generated by dialog selections.




Use the drop-down listbox in the *User* field in the *General* tab to identify the connecting role:

- Select *CURRENT_USER* to use the name of the current role.
- Select *PUBLIC* if no other user-specific mapping is applicable.

- Select a pre-defined role name to specify the name of an existing user.

Click the *Options* tab to continue.



The screenshot shows the 'Create - User Mapping' dialog box with the 'Options' tab selected. The dialog has a title bar with a close button. Below the title bar are three tabs: 'General', 'Options', and 'SQL'. The 'Options' tab is active and contains a table with two columns: 'Options' and 'Value'. The table has one row with the value 'data_options' in the 'Options' column and 'data_value' in the 'Value' column. There is a trash icon to the left of the row. At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two small icons: an information icon (i) and a help icon (?).

Options	Value
data_options	data_value

Use the fields in the *Options* tab to specify connection options; the accepted option names and values are specific to the foreign data wrapper associated with the server specified in the user mapping. Click the *Add* button to add an option/value pair.

- Specify the option name in the *Option* field.
- Provide a corresponding value in the *Value* field.

Click *Add* to specify each additional option/value pair; to discard an option, click the trash icon to the left of the row

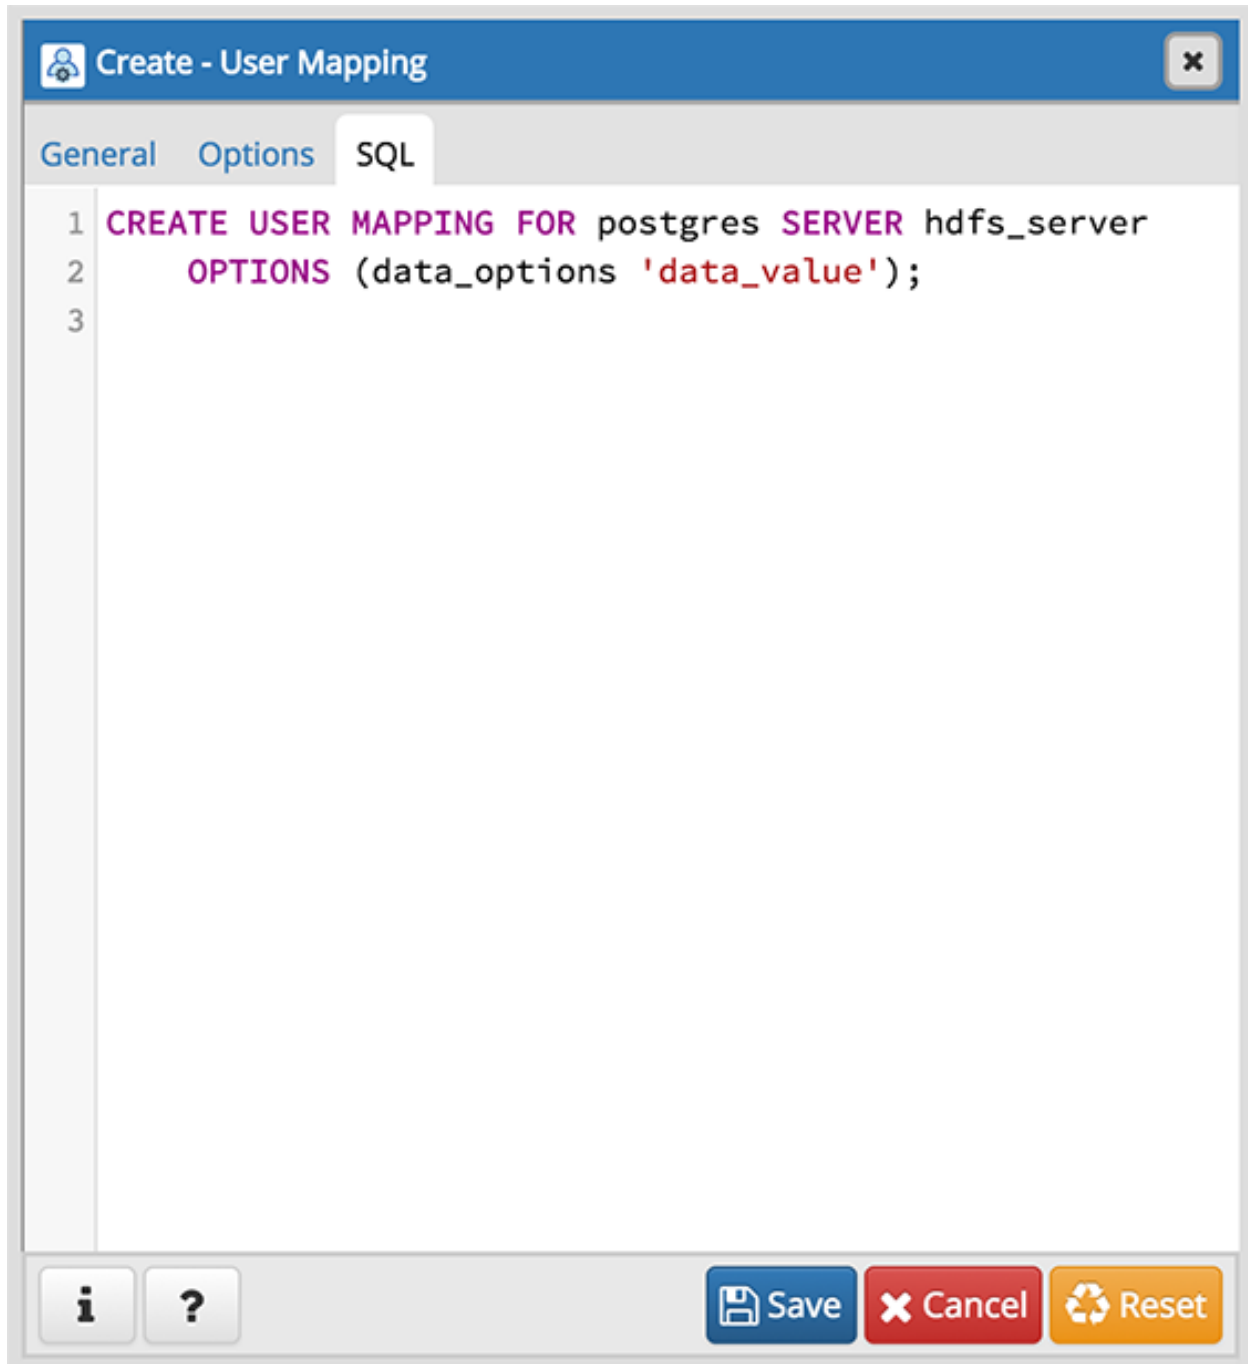
and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *User Mapping* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *User Mapping* dialog:



The example shown demonstrates a user mapping for the *hdfs_server*. The user is *CURRENT_USER* with a password *secret*.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

3.23 View Dialog

Use the *View* dialog to define a view. The view is not physically materialized; the query is executed each time the view is referenced in a query.

The *View* dialog organizes the development of a View through the following dialog tabs: *General*, *Definition*, and *Security*". The *SQL* tab displays the SQL code generated by dialog selections.

Click the *General* tab to begin.



The screenshot shows the 'Create - View' dialog box in pgAdmin 4. The dialog has a blue title bar with a close button (X) in the top right corner. Below the title bar are four tabs: 'General', 'Definition', 'Security', and 'SQL'. The 'General' tab is selected. The 'Name' field contains 'distributor_code'. The 'Owner' dropdown menu shows 'postgres' with a user icon. The 'Schema' dropdown menu shows 'public' with a schema icon. The 'Comment' field is a large empty text area. At the bottom of the dialog are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two smaller buttons: 'i' (info) and '?' (help).

Use the fields in the *General* tab to identify a view:

- Use the *Name* field to add a descriptive name for the view. The name of the view must be distinct from the name of any other view, table, sequence, index or foreign table in the same schema. The name will be displayed in the *pgAdmin* tree control.
- Use the drop-down listbox next to *Owner* to select the role that will own the view.
- If applicable, select the name of the schema in which the view will reside from the drop-down listbox in the *Schema* field.
- Store notes about the view in the *Comments* field.

Click the *Definition* tab to continue.



The screenshot shows the 'Create - View' dialog box in pgAdmin 4. The 'Definition' tab is selected. The 'Security barrier' switch is set to 'Yes'. The 'Check options' dropdown is set to 'Local'. The 'Definition' text area contains the SQL code: '1 begin', '2', '3 end;'. The dialog has tabs for 'General', 'Definition', 'Security', and 'SQL'. At the bottom, there are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields in the *Definition* tab to define properties of the view:

- Set the *Security Barrier* switch to *Yes* to indicate that the view is to act as a security barrier. For more information about defining and using a security barrier rule, see Section 38.5 of the PostgreSQL documentation.
- Use the **drop-down listbox next to *Check options* to select from *No*, *Local* or *Cascaded***. The *Local* option specifies that new rows are only checked against the conditions defined in the view. Any conditions defined on underlying base views are not checked (unless you specify the CHECK OPTION). The *Cascaded* option specifies new rows are checked against the conditions of the view and all underlying base views.

- Use the workspace in the *Definition* field to write a query to create a view.

Click the *Security* tab to continue.

Create - View

General Definition **Security** SQL

Privileges

	Grantee	Privileges	Grantor
	pem_admin	arwdDxt	postgres

Security labels

	Provider	Security Label
	data_provider	data_security

i **?** **Save** **Cancel** **Reset**

Use the *Security* tab to assign privileges and define security labels.

Use the *Privileges* panel to assign privileges to a role. Click the *Add* icon (+) to set privileges for the view:

- Select the name of the role that will be granted privileges from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.

- Select the name of a role with sufficient privileges to grant privileges on the view from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click the *Add* icon (+) to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the view. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

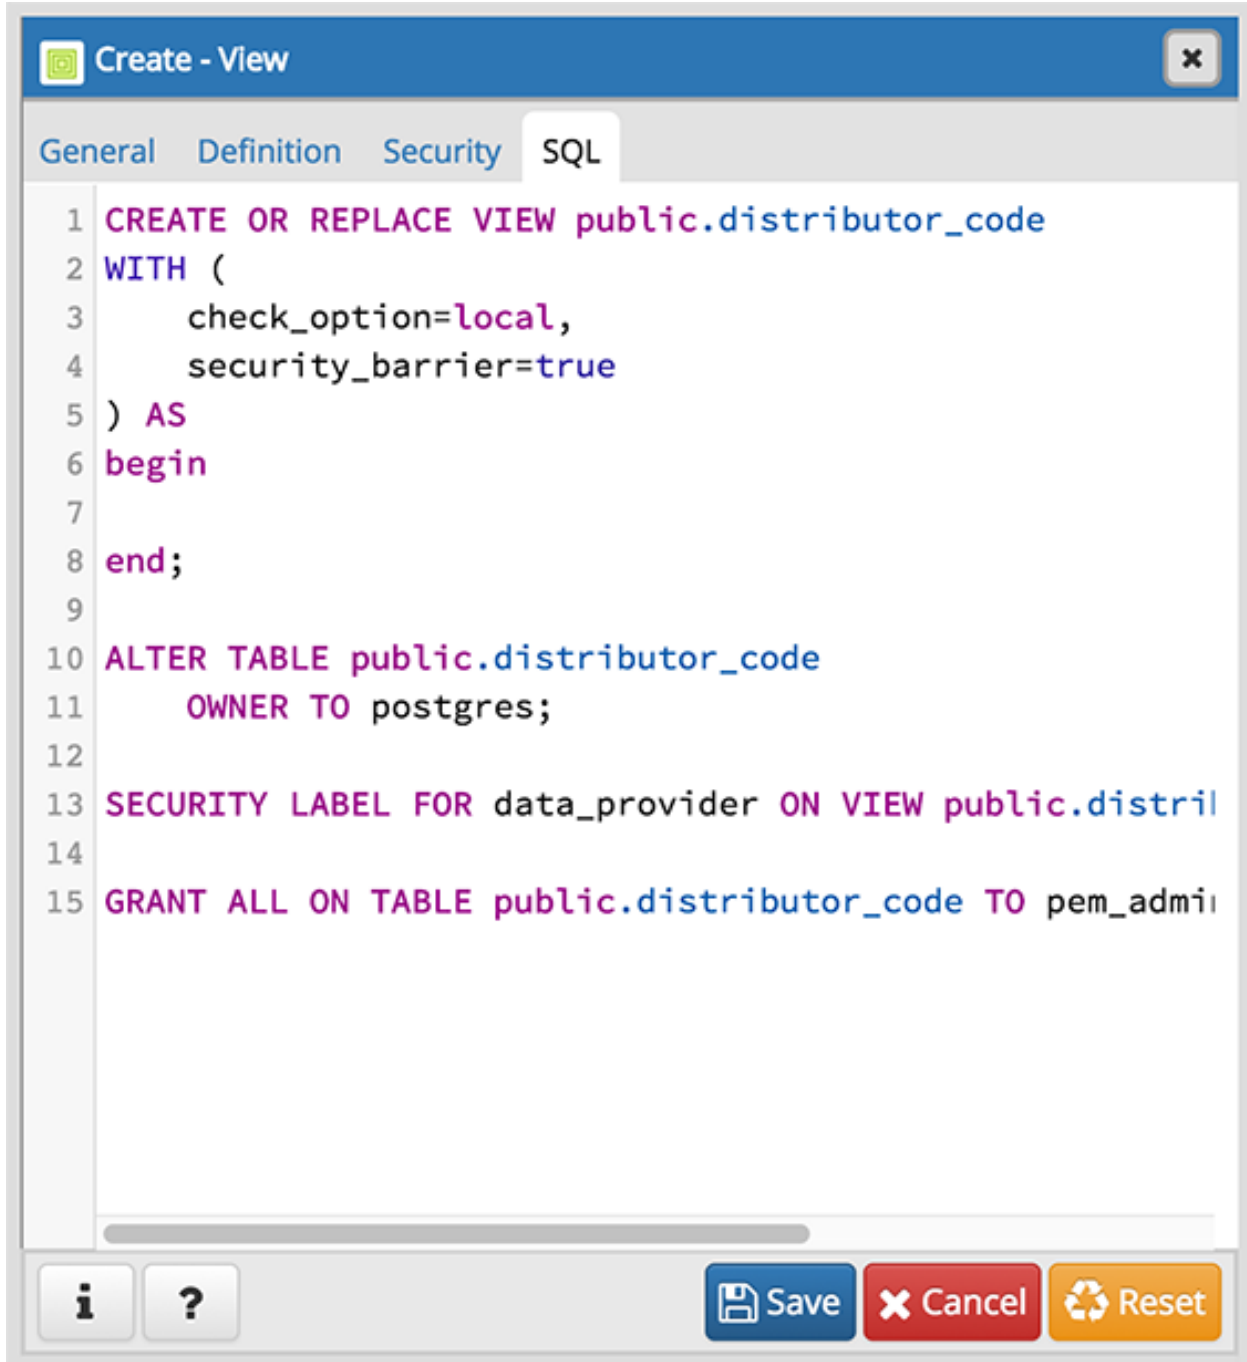
Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *View* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *View* dialog:



The example shown demonstrates creating a view named *distributor_codes* that includes the content of the *code* column from the *distributors* table.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

CREATING OR MODIFYING A TABLE

pgAdmin 4 provides dialogs that allow you to modify all table properties and attributes.


To access a dialog that allows you to create a database object, right-click on the object type in the pgAdmin tree control, and select the *Create* option for that object. For example, to create a new database, right-click on the *Casts* node, and select *Create Cast...*

Contents:

4.1 Check Dialog

Use the *Check* dialog to define or modify a check constraint. A check constraint specifies an expression that produces a Boolean result that new or updated rows must satisfy for an insert or update operation to succeed.

The *Check* dialog organizes the development of a check constraint through the *General* and *Definition* tabs. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Check' dialog box in pgAdmin 4. The window title is 'Create - Check' with a checkmark icon on the left and a close button on the right. Below the title bar are three tabs: 'General', 'Definition', and 'SQL'. The 'General' tab is selected. It contains two input fields: 'Name' (a single-line text box) and 'Comment' (a multi-line text area). At the bottom of the dialog, there are three buttons: an information icon (i), a help icon (?), and three action buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange).

Use the fields in the *General* tab to identify the check constraint:

- Use the *Name* field to provide a descriptive name for the check constraint that will be displayed in the *pgAdmin* tree control. With PostgreSQL 9.5 forward, when a table has multiple check constraints, they will be tested for each row in alphabetical order by name and after NOT NULL constraints.
- Store notes about the check constraint in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Check' dialog box in pgAdmin 4. The title bar is blue with a checkmark icon and the text 'Create - Check'. Below the title bar are three tabs: 'General', 'Definition', and 'SQL'. The 'Definition' tab is selected. The main area is divided into two sections. The top section is labeled 'Check' and contains a large, empty text area for entering the check constraint expression. The bottom section contains two toggle switches. The first is labeled 'No Inherit?' and is currently set to 'No'. The second is labeled 'Don't validate?' and is currently set to 'Yes'. At the bottom of the dialog are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). There are also two small icons, 'i' and '?', on the left side of the bottom bar.

Use the fields in the *Definition* tab to define the check constraint:

- Provide the expression that a row must satisfy in the *Check* field.
- Move the *No Inherit?* switch to the *Yes* position to specify this constraint is automatically inherited by a table's children. The default is *No*.
- Move the *Don't validate?* switch to the *No* position to skip validation of existing data; the constraint may not hold for all rows in the table. The default is *Yes*.

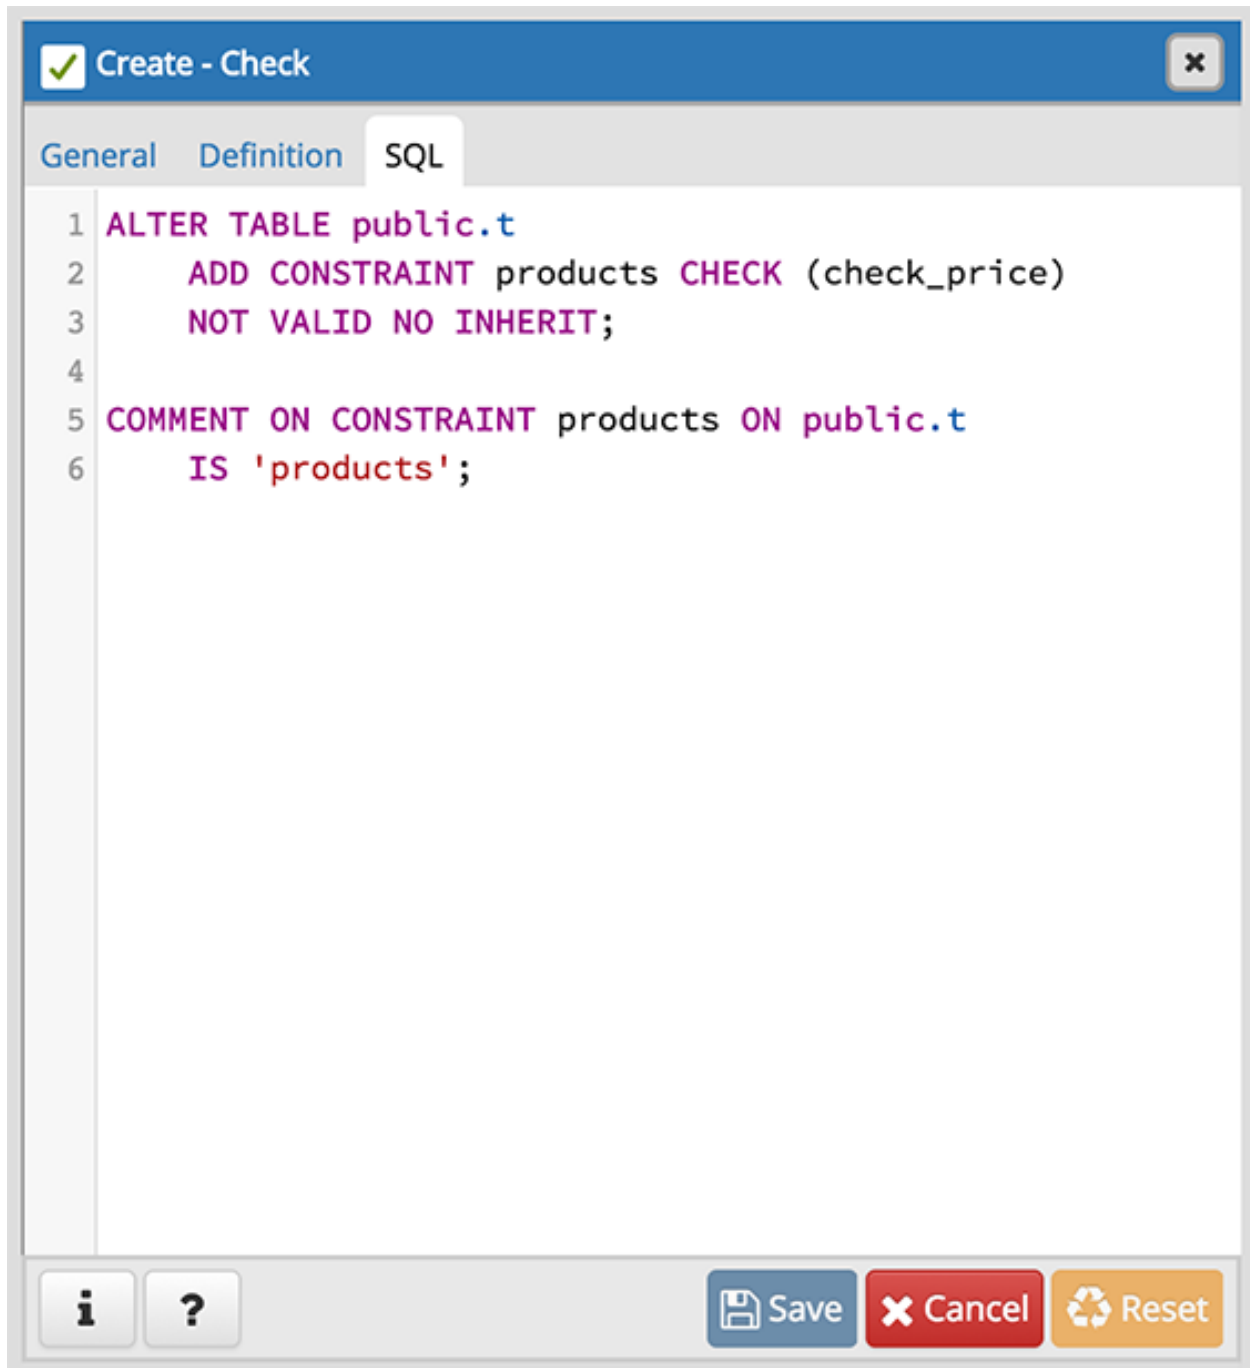
Click the *SQL* tab to continue.

Your entries in the *Check* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit

or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Check* dialog:



The example shown demonstrates creating a check constraint named *check_price* on the *price* column of the *products* table. The constraint confirms that any values added to the column are greater than 0.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.

- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

4.2 Column Dialog

Use the *Column* dialog to add a column to an existing table or modify a column definition.

The *Column* dialog organizes the development of a column through the following dialog tabs: *General*, *Definition*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.

The screenshot shows the 'Create - Column' dialog box. The title bar is blue with the text 'Create - Column' and a close button. Below the title bar are five tabs: 'General', 'Definition', 'Variables', 'Security', and 'SQL'. The 'General' tab is selected. In the 'General' tab, there is a 'Name' label followed by a text input field containing the text 'distributors'. Below the 'Name' field is a 'Comment' label followed by a large, empty text area. At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two small icons: an information icon (i) and a help icon (?).

Use the fields in the *General* tab to identify the column:

- Use the *Name* field to add a descriptive name for the column. The name will be displayed in the *pgAdmin* tree control. This field is required.
- Store notes about the column in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Column' dialog box in pgAdmin 4. The 'Definition' tab is selected, showing the following fields:

- Data type:** A dropdown menu with 'char' selected.
- Length:** A disabled text input field.
- Precision:** A disabled text input field.
- Collation:** A disabled dropdown menu.
- Default Value:** An empty text input field.
- Not NULL?:** A toggle switch currently set to 'Yes'.

At the bottom of the dialog, there are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields in the *Definition* tab to add parameters for the column. (Fields are disabled if inapplicable.)

- Use the drop-down listbox next to *Data Type* to select a data type for the column. For more information on the data types that are supported by PostgreSQL, refer to Chapter 8 of the Postgres core documentation. This field

is required.

- Use the *Length* and *Precision* fields to specify the maximum number of significant digits in a numeric value, or the maximum number of characters in a text value.
- Use the drop-down listbox next to *Collation* to apply a collation setting to the column.
- Use the *Default Value* field to specify a default data value.
- Move the *Not Null* switch to the *Yes* position to specify the column may not contain null values. The default is *No*.

Click the *Variables* tab to continue.

The screenshot shows the 'Create - Column' dialog box with the 'Variables' tab selected. The dialog has a title bar 'Create - Column' and a close button. Below the title bar are tabs for 'General', 'Definition', 'Variables', 'Security', and 'SQL'. The 'Variables' tab is active, showing a table with the following structure:

Variables		+
	Name	Value
	n_distinct	

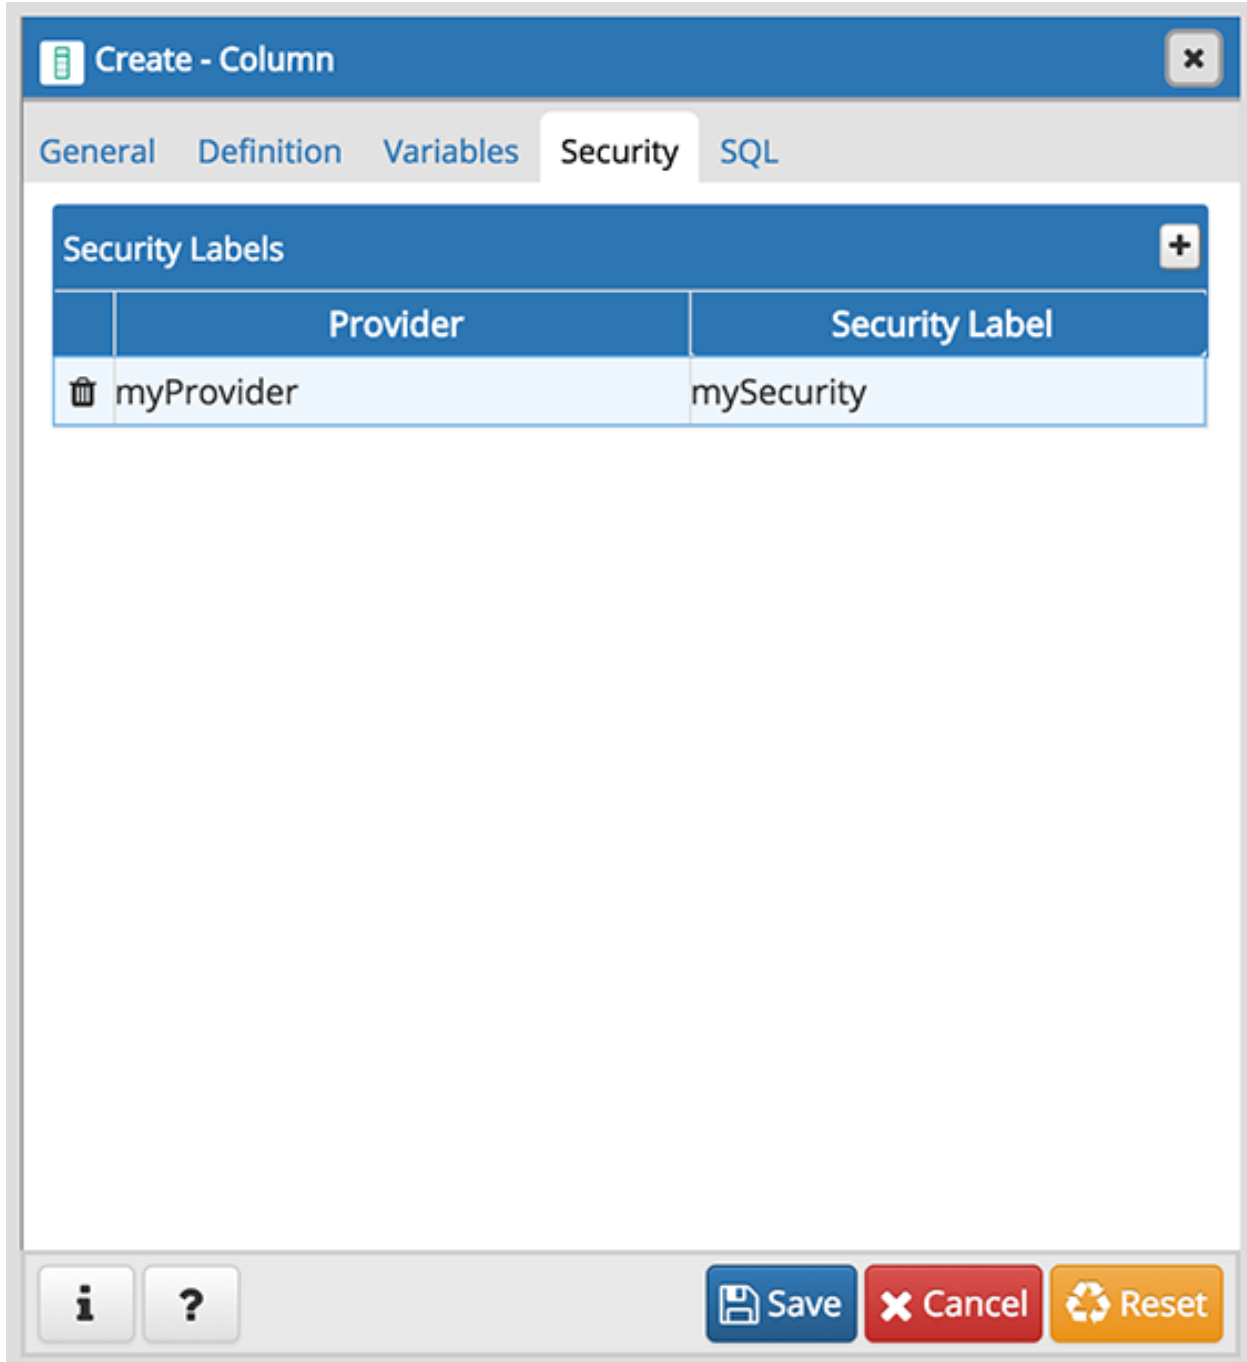
At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two small icons: an information icon 'i' and a help icon '?'.

Use the *Variables* tab to specify the number of distinct values that may be present in the column; this value overrides estimates made by the ANALYZE command. Click the *Add* icon (+) to add a *Name/Value* pair:

- Select the name of the variable from the drop-down listbox in the *Name* field.
 - Select *n_distinct* to specify the number of distinct values for the column.
 - Select *n_distinct_inherited* to specify the number of distinct values for the table and its children.
- Specify the number of distinct values in the *Value* field. For more information, see the documentation for [ALTER TABLE](#).

Click the *Add* icon (+) to specify each additional *Name/Value* pair; to discard a variable, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Security* tab to continue.



Use the *Security* tab to assign attributes and define security labels. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

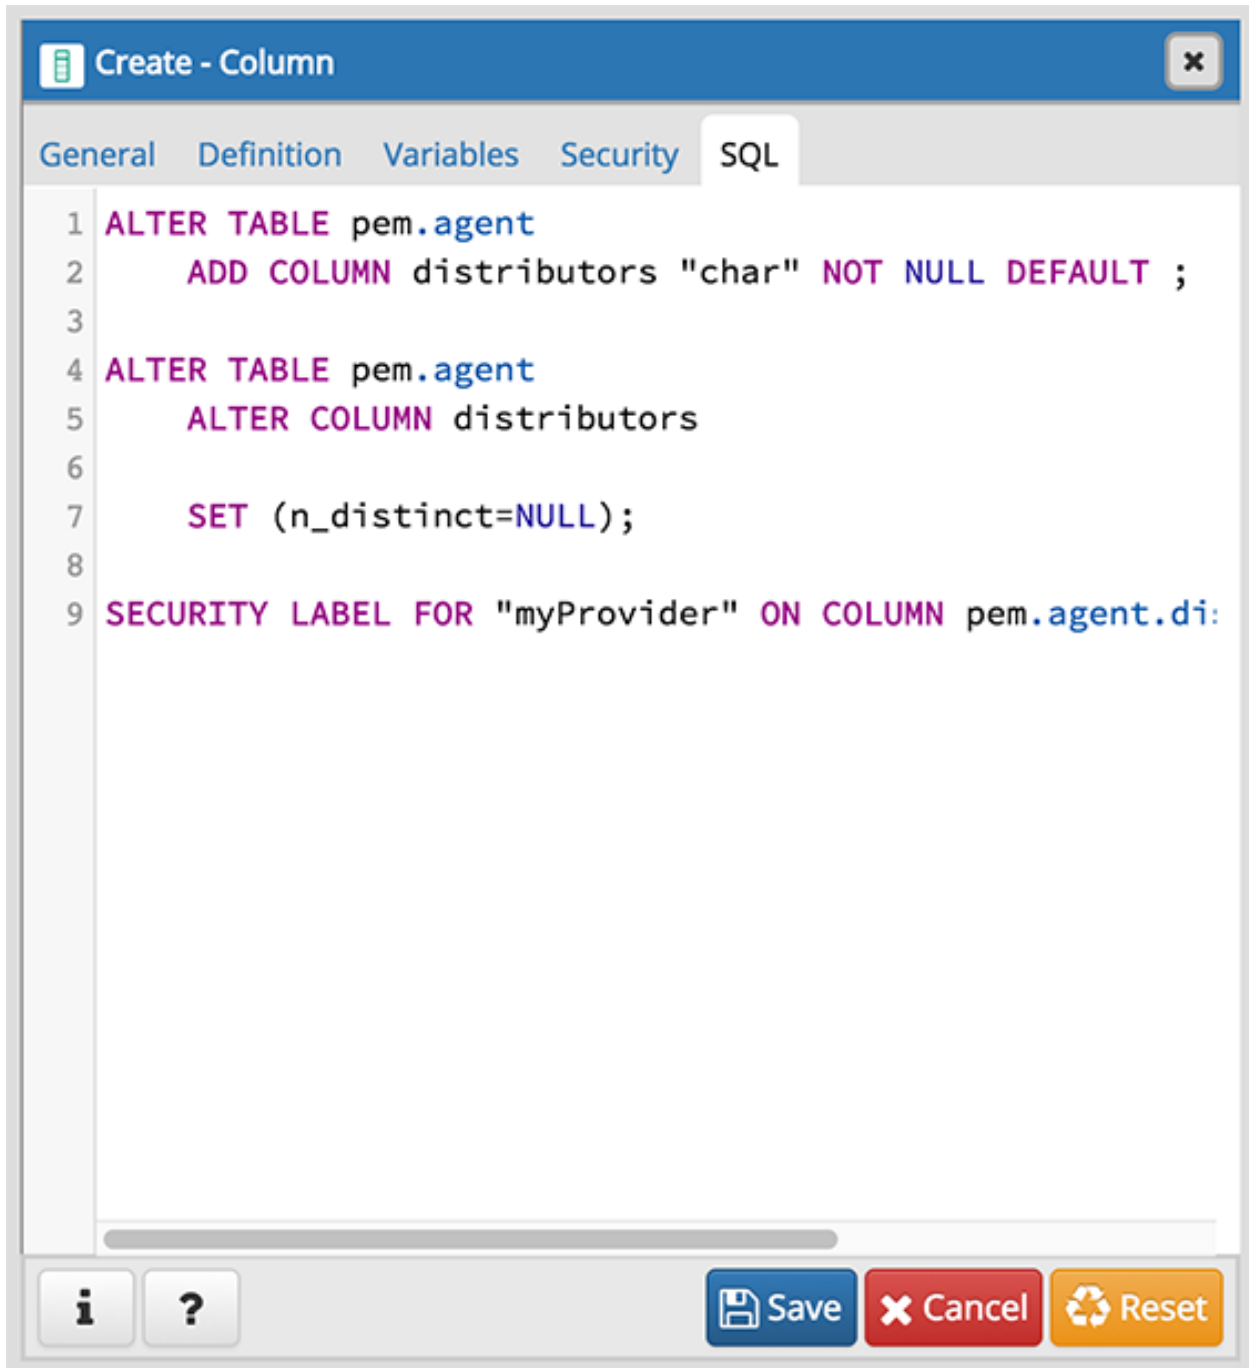
Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Column* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Column* dialog:



The example shown demonstrates creating a column named *territory* in the table named *distributors*.


- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.

- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

4.3 Exclusion constraint Dialog

Use the *Exclusion constraint* dialog to define or modify the behavior of an exclusion constraint. An exclusion constraint guarantees that if any two rows are compared on the specified column or expression (using the specified operator), at least one of the operator comparisons will return false or null.

The *Exclusion constraint* dialog organizes the development of an exclusion constraint through the following dialog tabs: *General*, *Definition*, and *Columns*. The *SQL* tab displays the SQL code generated by dialog selections.



Use the fields in the *General* tab to identify the exclusion constraint:

- Use the *Name* field to provide a descriptive name for the exclusion constraint. The name will be displayed in the *pgAdmin* tree control.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Exclusion constraint' dialog box in pgAdmin 4. The 'Definition' tab is selected, showing the following fields:

- Tablespace:** A dropdown menu with 'pg_default' selected.
- Access method:** A dropdown menu with 'btree' selected.
- Fill factor:** An empty text input field.
- Deferrable?:** A radio button with 'Yes' selected.
- Deferred?:** A radio button with 'Yes' selected.
- Constraint:** A large empty text area for defining the constraint.

At the bottom of the dialog, there are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields in the *Definition* tab to define the exclusion constraint:

- Use the drop-down listbox next to *Tablespace* to select the tablespace in which the index associated with the exclude constraint will reside.
- Use the drop-down listbox next to *Access method* to specify the type of index that will be used when implementing the exclusion constraint:
 - Select *gist* to specify a GiST index.
 - Select *spgist* to specify a space-partitioned GiST index.
 - Select *btree* to specify a B-tree index.

- Select *hash* to specify a hash index.
- Use the *Fill Factor* field to specify a fill factor for the table and associated index. The fill factor is a percentage between 10 and 100. 100 (complete packing) is the default.
- Move the *Deferrable?* switch to the *Yes* position to specify that the timing of the constraint is deferrable, and can be postponed until the end of the statement. The default is *No*.
- If enabled, move the *Deferred?* switch to the *Yes* position to specify the timing of the constraint is deferred to the end of the statement. The default is *No*.
- Use the *Constraint* field to provide a condition that a row must satisfy to be included in the table.

Click the *Columns* tab to continue.

The screenshot shows the 'Create - Exclusion constraint' dialog box with the 'Columns' tab selected. The dialog has four tabs: 'General', 'Definition', 'Columns', and 'SQL'. The 'Columns' tab is active, displaying a table with the following data:

Column	Operator class	DESC	NULLs order	Operator
col1	cidr_ops	ASC	FIRST	<=

Below the table, there is an 'Include columns' field containing 'col2' and 'col3'. At the bottom of the dialog are buttons for 'Save', 'Cancel', and 'Reset'.

Use the fields in the *Columns* tab to specify the column(s) to which the constraint applies. Use the drop-down listbox next to *Column* to select a column and click the *Add* icon (+) to provide details of the action on the column:

- The *Column* field is populated with the selection made in the *Column* drop-down listbox.
- If applicable, use the drop-down listbox in the *Operator class* to specify the operator class that will be used by the index for the column.
- Move the *DESC* switch to *DESC* to specify a descending sort order. The default is *ASC* which specifies an ascending sort order.
- Use the *NULLs order* column to specify the placement of NULL values (when sorted). Specify *FIRST* or *LAST*.
- Use the drop-down list next to *Operator* to specify a comparison or conditional operator.

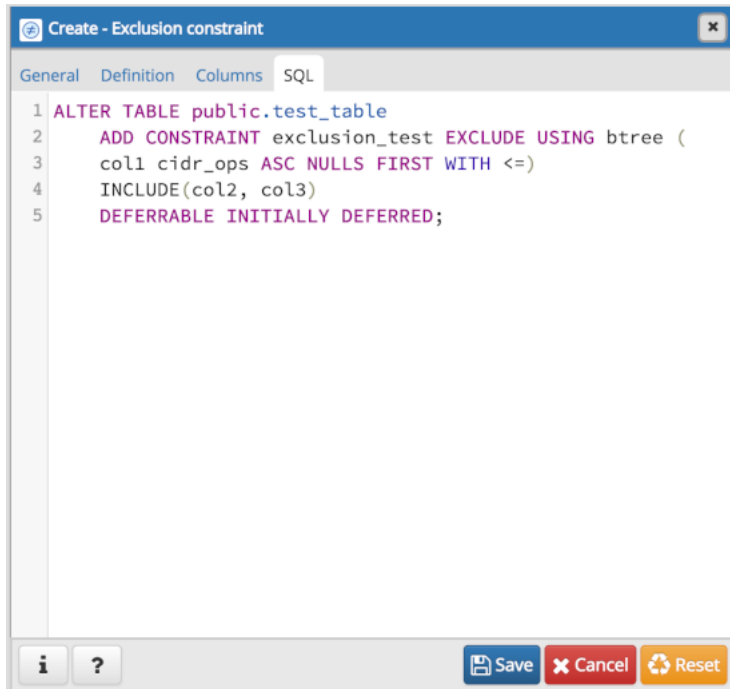
Use *Include columns* field to specify columns for *INCLUDE* clause of the constraint. This option is available in Postgres 11 and later.

Click the *SQL* tab to continue.

Your entries in the *Exclusion Constraint* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Exclusion Constraint* dialog:




The example shown demonstrates creating an exclusion constraint named *exclude_department* that restricts additions to the *dept* table to those additions that are not equal to the value of the *deptno* column. The constraint uses a btree index.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

4.4 Foreign key Dialog

Use the *Foreign key* dialog to specify the behavior of a foreign key constraint. A foreign key constraint maintains referential integrity between two tables. A foreign key constraint cannot be defined between a temporary table and a permanent table.

The *Foreign key* dialog organizes the development of a foreign key constraint through the following dialog tabs: *General*, *Definition*, *Columns*, and *Action*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Foreign key' dialog box. The title bar includes a key icon and a close button. The dialog has five tabs: 'General', 'Definition', 'Columns', 'Action', and 'SQL'. The 'General' tab is active and contains a 'Name' text input field and a 'Comment' text area. At the bottom, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange), along with information and help icons.

Use the fields in the *General* tab to identify the foreign key constraint:

- Use the *Name* field to add a descriptive name for the foreign key. The name will be displayed in the *pgAdmin* tree control.
- Store notes about the foreign key constraint in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Foreign key' dialog box with the 'Definition' tab selected. The configuration options are as follows:

Property	Value
Deferrable?	Yes
Deferred?	Yes
Match type	FULL
Validated?	Yes
Auto FK index?	Yes
Covering index	fki_distributors

Use the fields in the *Definition* tab to define the foreign key constraint:

- Move the *Deferrable?* switch to the *Yes* position to specify the timing of the constraint is deferrable and can be postponed until the end of the statement. The default is *No*.
- If enabled, move the *Deferred?* switch to the *Yes* position to specify the timing of the constraint is deferred to the end of the statement. The default is *No*.
- Move the *Match type* switch specify the type of matching that is enforced by the constraint:
 - Select *Full* to indicate that all columns of a multicolumn foreign key must be null if any column is null; if all columns are null, the row is not required to have a match in the referenced table.

- Select *Simple* to specify that a single foreign key column may be null; if any column is null, the row is not required to have a match in the referenced table.
- Move the *Validated* switch to the *Yes* position to instruct the server to validate the existing table content (against a foreign key or check constraint) when you save modifications to this dialog.
- Move the *Auto FK Index* switch to the *No* position to disable the automatic index feature.
- The field next to *Covering Index* generates the name of an index if the *Auto FK Index* switch is in the *Yes* position; or, this field is disabled.

Click the *Columns* tab to continue.

The screenshot shows the 'Create - Foreign key' dialog box with the 'Columns' tab selected. The dialog contains the following elements:

- Header:** 'Create - Foreign key' with a close button (X).
- Tabs:** 'General', 'Definition', 'Columns' (selected), 'Action', 'SQL'.
- Columns Section:**
 - Local column:** A dropdown menu showing 'agent_id' with a plus sign (+) to the right.
 - References:** A dropdown menu showing 'Select foreign table'.
 - Referencing:** A dropdown menu showing 'Select column'.
- Table:** A table with two columns: 'Local' and 'Referenced'.
- Warning Message:** A red banner at the bottom with a warning icon and the text 'Please specify columns for Foreign key.'
- Footer:** Information (i), Help (?), Save, Cancel, and Reset buttons.

Use the fields in the *Columns* tab to specify one or more reference column(s). A Foreign Key constraint requires that one or more columns of a table must only contain values that match values in the referenced column(s) of a row of a referenced table:

- Use the drop-down listbox next to *Local column* to specify the column in the current table that will be compared to the foreign table.
- Use the drop-down listbox next to *References* to specify the name of the table in which the comparison column(s) resides.
- Use the drop-down listbox next to *Referencing* to specify a column in the foreign table.

Click the *Add* icon (+) to add a column to the list; repeat the steps above and click the *Add* icon (+) to add additional columns. To discard an entry, click the trash icon to the left of the entry and confirm deletion in the *Delete Row* popup.

Click the *Action* tab to continue.

The screenshot shows the 'Create - Foreign key' dialog box in pgAdmin 4. The dialog has a title bar with a key icon and the text 'Create - Foreign key'. Below the title bar are five tabs: 'General', 'Definition', 'Columns', 'Action', and 'SQL'. The 'Action' tab is selected. In the 'Action' tab, there are two drop-down menus. The first is labeled 'On update' and the second is labeled 'On delete'. Both drop-down menus currently show 'NO ACTION'. At the bottom of the dialog, there is a red error bar with a warning icon and the text 'Please specify columns for Foreign key.'. Below the error bar are three buttons: an information icon, a help icon, and buttons for 'Save', 'Cancel', and 'Reset'.

Use the drop-down listboxes on the *Action* tab to specify behavior related to the foreign key constraint that will be performed when data within the table is updated or deleted:

- Use the drop-down listbox next to *On update* to select an action that will be performed when data in the table is updated.
- Use the drop-down listbox next to *On delete* to select an action that will be performed when data in the table is deleted.

The supported actions are:

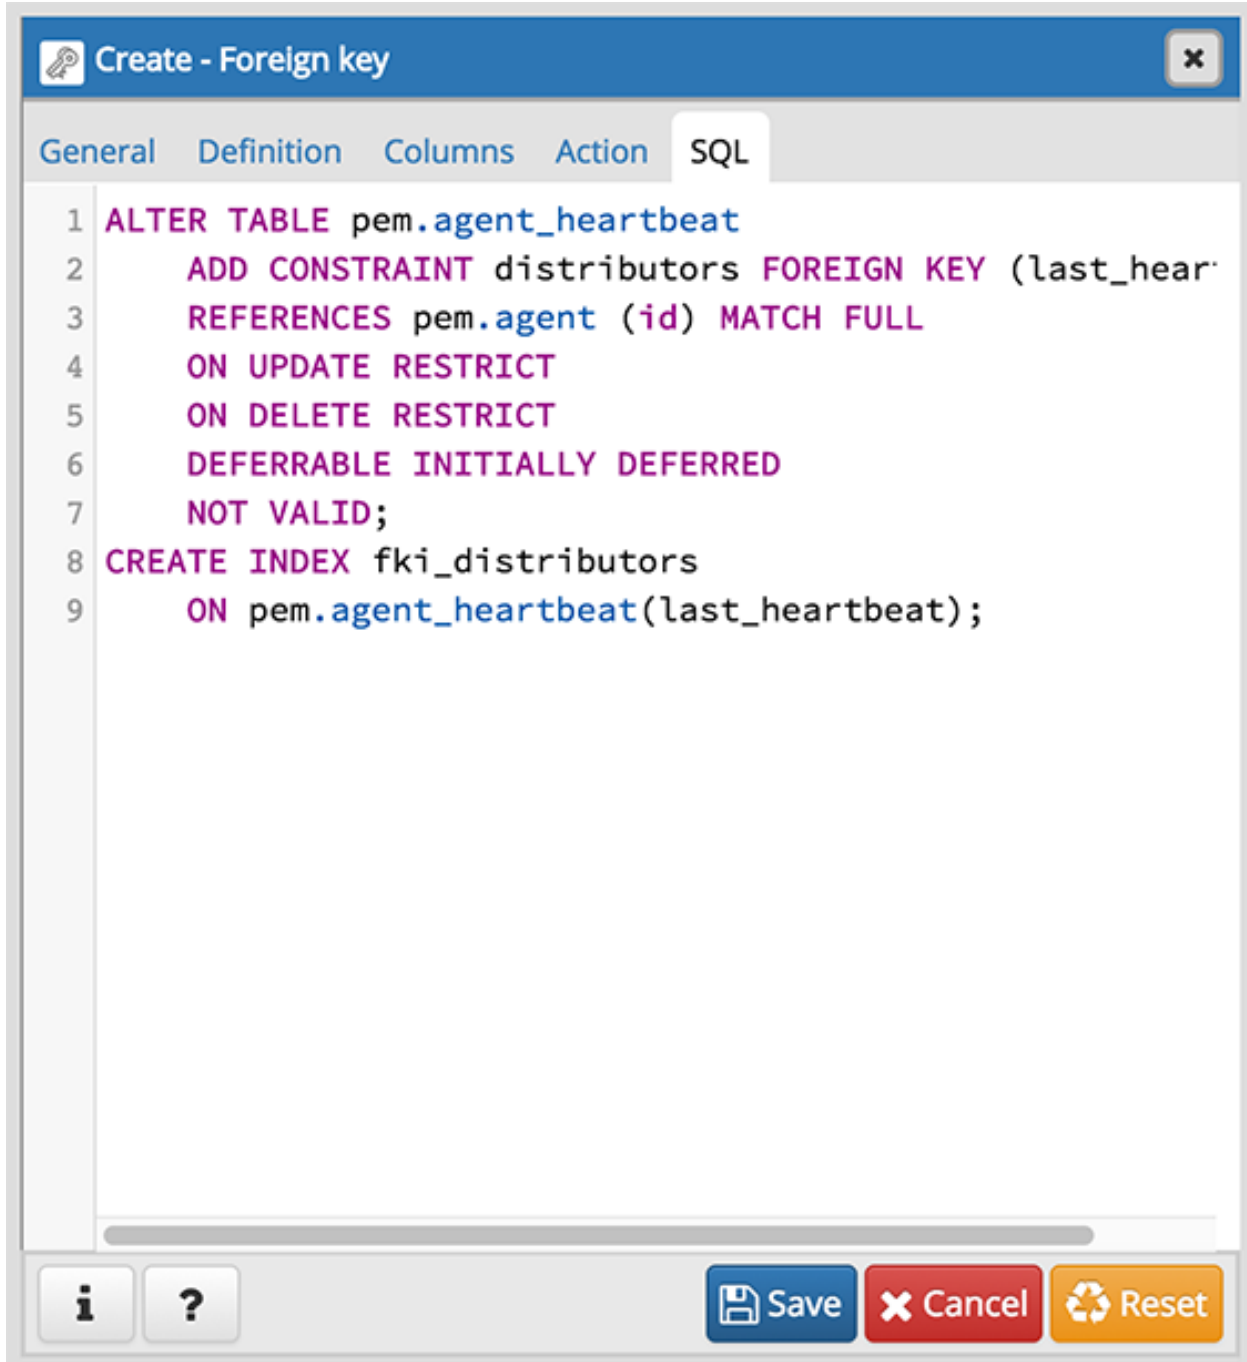
NO ACTION	Produce an error indicating that the deletion or update will create a foreign key constraint violation. If the constraint is deferred, this error will be produced at constraint check time if any referencing rows still exist. This is the default.
RESTRICT	Throw an error indicating that the deletion or update would create a foreign key constraint violation. This is the same as NO ACTION except that the check is not deferrable.
CASCADE	Delete any rows referencing the deleted row, or update the values of the referencing column(s) to the new values of the referenced columns, respectively.
SET NULL	Set the referencing column(s) to null.
SET DEFAULT	Set the referencing column(s) to their default values. There must be a row in the referenced table that matches the default values (if they are not null), or the operation will fail.

Click the *SQL* tab to continue.

Your entries in the *Foreign key* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Foreign key* dialog:



The example shown demonstrates creating a foreign key constraint named *territory_fkey* that matches values in the *distributors* table *territory* column with those of the *sales_territories* table *region* column.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

4.5 Index Dialog

Use the *Index* dialog to create an index on a specified table or materialized view.

The *Index* dialog organizes the development of a index through the following dialog tabs: *General* and *Definition*. The *SQL* tab displays the SQL code generated by dialog selections.



The screenshot shows the 'Create - Index' dialog box. The title bar reads 'Create - Index' with a close button. The dialog has three tabs: 'General', 'Definition', and 'SQL'. The 'General' tab is selected. It contains three input fields: 'Name' (a text box), 'Tablespace' (a dropdown menu with 'Select from the list' selected), and 'Comment' (a large text area). At the bottom, there are buttons for 'i' (info), '?' (help), 'Save', 'Cancel', and 'Reset'.

Use the fields in the *General* tab to identify the index:

- Use the *Name* field to add a descriptive name for the index. The name will be displayed in the *pgAdmin* tree control.

- Use the drop-down listbox next to *Tablespace* to select the tablespace in which the index will reside.
- Store notes about the index in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Index' dialog box with the 'Definition' tab selected. The 'Access Method' is set to 'btree'. The 'Unique?' switch is in the 'Yes' position. The 'Clustered?' switch is also in the 'Yes' position. The 'Concurrent build?' switch is in the 'Yes' position. The 'Constraint' field contains the number '1'. Below these options is a table with the following columns: Column, Operator class, Sort order, NULLs, and Collation. The table contains one row: 'col1', 'cidr_ops', 'ASC', 'LAST', and 'pg_catalog."C"'. Below the table is an 'Include columns' field containing 'col2'. At the bottom of the dialog are buttons for 'Save', 'Cancel', and 'Reset'.

Column	Operator class	Sort order	NULLs	Collation
col1	cidr_ops	ASC	LAST	pg_catalog."C"

Use the fields in the *Definition* tab to define the index:

- Use the drop-down listbox next to *Access Method* to select an index type:
 - Select *btree* to create a B-tree index. A B-tree index may improve performance when managing equality and range queries on data that can be sorted into some ordering (the default).
 - Select *hash* to create a hash index. A hash index may improve performance when managing simple equality comparisons.
 - Select *gist* to create a GiST index. A GiST index may improve performance when managing values with more than one key.
 - Select *gin* to create a GIN index. A GIN index may improve performance when managing two-dimensional geometric data types and nearest-neighbor searches.
 - Select *spgist* to create a space-partitioned GiST index. A SP-GiST index may improve performance when managing non-balanced data structures.
 - Select *brin* to create a BRIN index. A BRIN index may improve performance when managing minimum and maximum values and ranges.
- Use the *Fill Factor* field to specify a fill factor for the index. The fill factor specifies how full the selected method will try to fill each index page.
- Move the *Unique?* switch to the *Yes* position to check for duplicate values in the table when the index is created and when data is added. The default is *No*.
- Move the *Clustered?* switch to the *Yes* position to instruct the server to cluster the table.
- Move the *Concurrent build?* switch to the *Yes* position to build the index without taking any locks that prevent concurrent inserts, updates, or deletes on the table.
- Use the *Constraint* field to provide a constraint expression; a constraint expression limits the entries in the index to those rows that satisfy the constraint.

Use the context-sensitive fields in the *Columns* panel to specify which column(s) the index queries. Click the *Add* icon (+) to add a column:

- Use the drop-down listbox in *Column* field to select the name of the column from the table.
- If enabled, use the drop-down listbox to select an available *Operator class* to specify the type of action performed on the column.
- If enabled, move the *Sort order* switch to specify the sort order:
 - Select *ASC* to specify an ascending sort order (the default);
 - Select *DESC* to specify a descending sort order.
- If enabled, move the *Nulls* switch to specify the sort order of nulls:
 - Select *First* to specify nulls sort before non-nulls;
 - Select *Last* to specify nulls sort after non-nulls (the default).
- Use the drop-down listbox in the *Collation* field to select a collation to use for the index.


Use *Include columns* field to specify columns for *INCLUDE* clause of the index. This option is available in Postgres 11 and later.

Click the *SQL* tab to continue.

Your entries in the *Index* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Index* dialog:



```

1 CREATE UNIQUE INDEX CONCURRENTLY index_test
2   ON public.test_table USING btree
3   (col1 COLLATE pg_catalog.\"C\" cidr_ops ASC NULLS LAST)
4   INCLUDE(col2)
5   TABLESPACE pg_default;
6
7 ALTER TABLE public.test_table
8   CLUSTER ON index_test;
9
10 COMMENT ON INDEX public.index_test
11   IS 'Comment on index';

```

The example shown demonstrates creating an index named *dist_codes* that indexes the values in the *code* column of the *distributors* table.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.

- Click the *Reset* button to restore configuration parameters.

4.6 Primary key Dialog

Use the *Primary key* dialog to create or modify a primary key constraint. A primary key constraint indicates that a column, or group of columns, uniquely identifies rows in a table. This requires that the values in the selected column(s) be both unique and not null.

The *Primary key* dialog organizes the development of a primary key constraint through the *General* and *Definition* tabs. The *SQL* tab displays the SQL code generated by dialog selections.

The screenshot shows the 'Create - Primary key' dialog box in pgAdmin 4. The dialog has a blue title bar with a key icon and a close button. It features three tabs: 'General' (selected), 'Definition', and 'SQL'. The 'General' tab contains a 'Name' text input field and a 'Comment' text area. At the bottom, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange), along with information and help icons.

Use the fields in the *General* tab to identify the primary key:

- Use the *Name* field to add a descriptive name for the primary key constraint. The name will be displayed in the *pgAdmin* tree control.

Click the *Definition* tab to continue.

Use the fields in the *Definition* tab to define the primary key constraint:

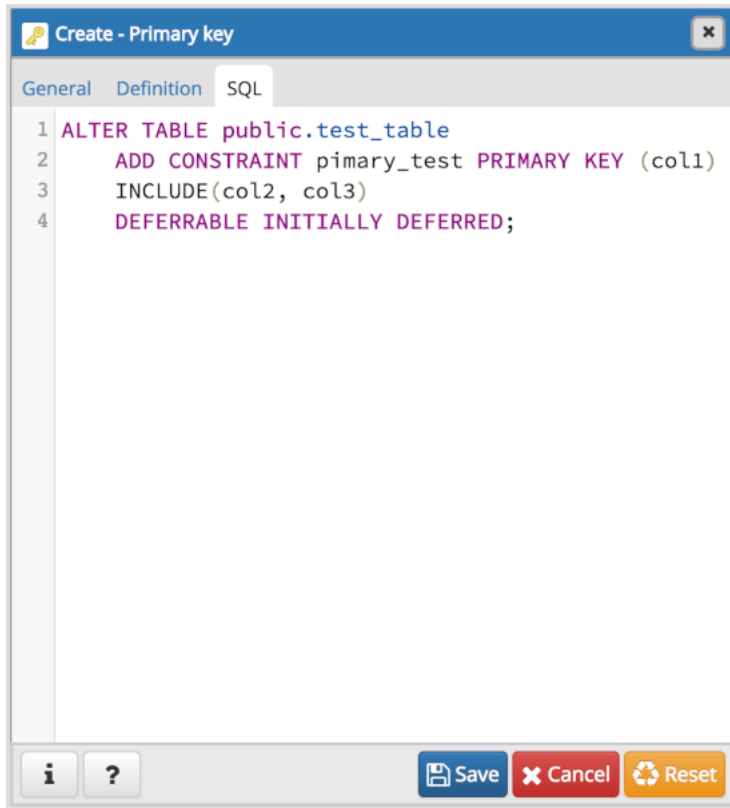
- Click inside the *Columns* field and select one or more column names from the drop-down listbox. To delete a selection, click the *x* to the left of the column name. The primary key constraint should be different from any unique constraint defined for the same table; the selected column(s) for the constraints must be distinct.
- Use *Include columns* field to specify columns for *INCLUDE* clause of the index. This option is available in Postgres 11 and later.
- Select the name of the tablespace in which the primary key constraint will reside from the drop-down listbox in the *Tablespace* field.
- Select the name of an index from the drop-down listbox in the *Index* field. This field is optional. Adding a primary key will automatically create a unique B-tree index on the column or group of columns listed in the primary key, and will force the column(s) to be marked NOT NULL.
- Use the *Fill Factor* field to specify a fill factor for the table and index. The fill factor for a table is a percentage between 10 and 100. 100 (complete packing) is the default.
- Move the *Deferrable?* switch to the *Yes* position to specify the timing of the constraint is deferrable and can be postponed until the end of the statement. The default is *No*.
- If enabled, move the *Deferred?* switch to the *Yes* position to specify the timing of the constraint is deferred to the end of the statement. The default is *No*.

Click the *SQL* tab to continue.

Your entries in the *Primary key* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Primary key* dialog:



The example shown demonstrates creating a primary key constraint named *dept_pkey* on the *dept_id* column of the *dept* table.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

4.7 Rule Dialog

Use the *Rule* dialog to define or modify a rule for a specified table or view. A PostgreSQL rule allows you to define an additional action that will be performed when a SELECT, INSERT, UPDATE, or DELETE is performed against a table.

The *Rule* dialog organizes the development of a rule through the *General*, and *Definition* tabs. The *SQL* tab displays the SQL code generated by dialog selections.

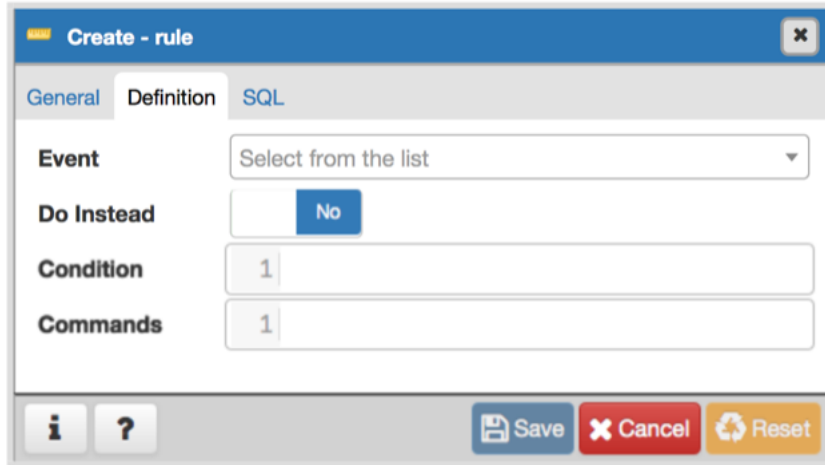


The screenshot shows a dialog box titled "Create - rule" with a close button in the top right corner. Below the title bar are three tabs: "General", "Definition", and "SQL". The "General" tab is selected. It contains two input fields: "Name" with the value "notify_me" and "Comment" with the value "This rule sends a notification.". At the bottom of the dialog, there are three buttons: "Save" (blue), "Cancel" (red), and "Reset" (orange). To the left of these buttons are two small icons: an information icon (i) and a help icon (?).

Use the fields in the *General* tab to identify the rule:

- Use the *Name* field to add a descriptive name for the rule. The name will be displayed in the *pgAdmin* tree control. Multiple rules on the same table are applied in alphabetical name order.
- Store notes about the rule in the *Comment* field.

Click the *Definition* tab to continue.



Use the fields in the *Definition* tab to write parameters:

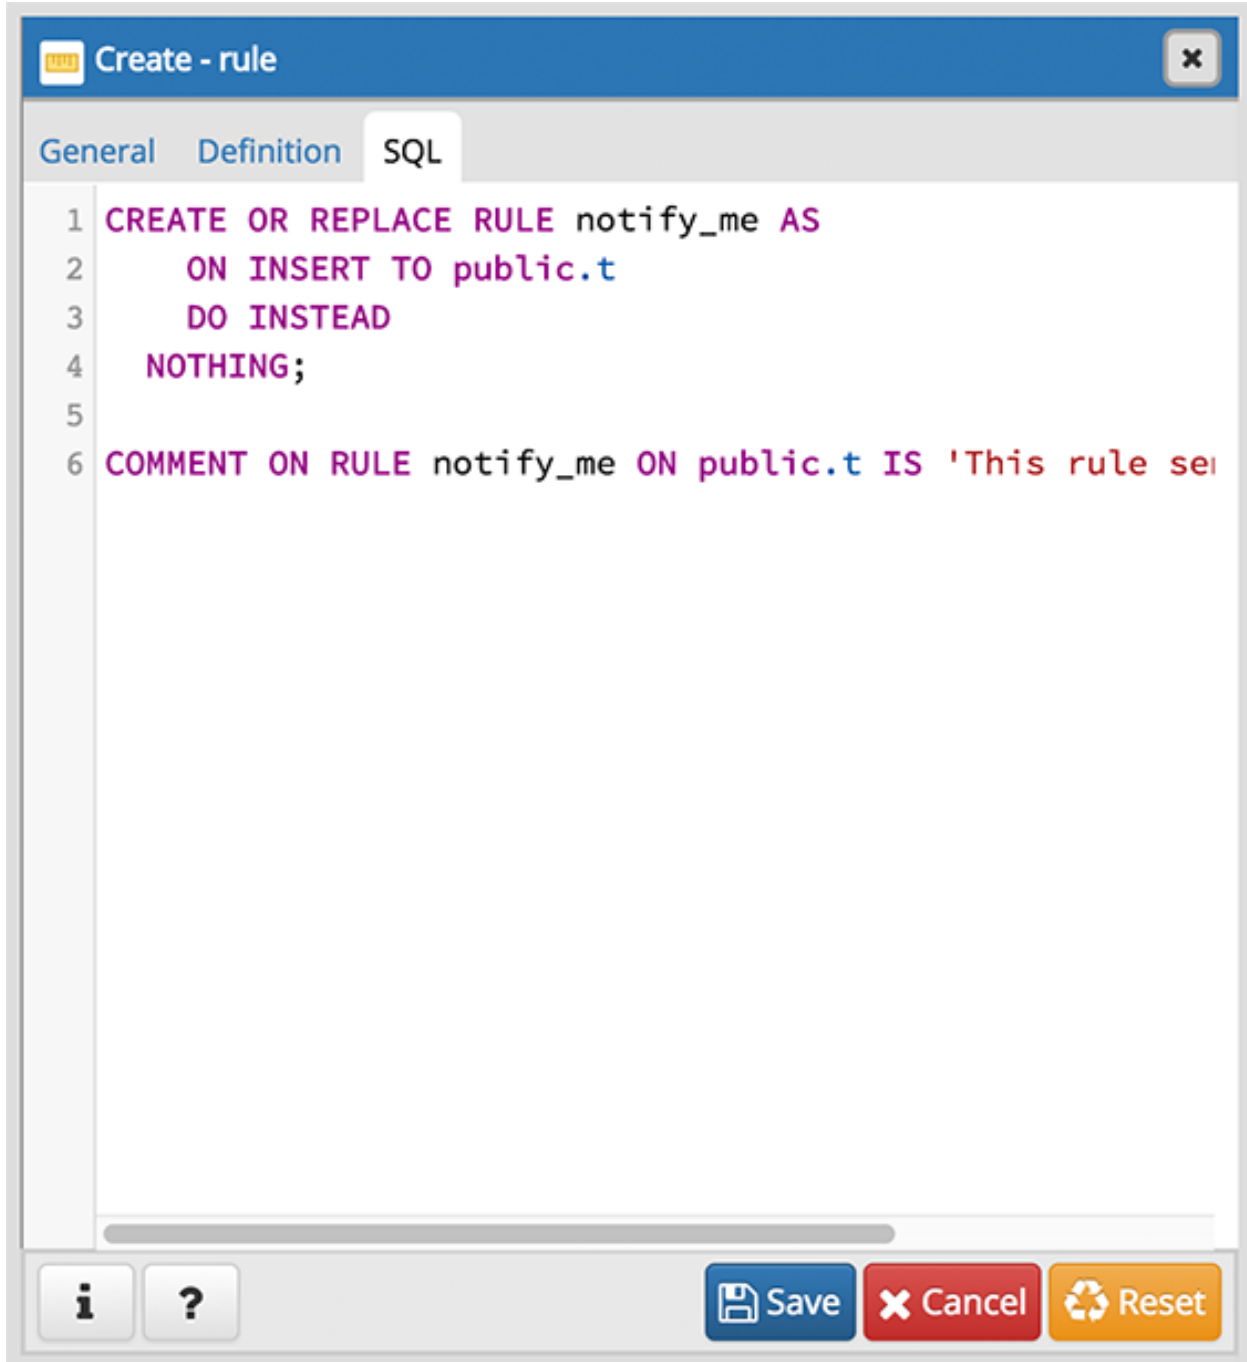
- Click inside the *Event* field to select the type of event that will invoke the rule; event may be *Select*, *Insert*, *Update*, or *Delete*.
- Move the *Do Instead* switch to *Yes* indicate that the commands should be executed instead of the original command; if *Do Instead* specifies *No*, the rule will be invoked in addition to the original command.
- Specify a SQL conditional expression that returns a boolean value in the *Condition* editor.
- Provide a command in the *Commands* editor that defines the action performed by the rule.

Click the *SQL* tab to continue.

Your entries in the *Rule* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Rule* dialog:




The example sends a notification when an UPDATE executes against a table.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

4.8 Table Dialog

Use the *Table* dialog to create or modify a table.


The *Table* dialog organizes the development of a table through the following dialog tabs: *General*, *Columns*, *Constraints*, *Advanced*, *Parameter*, and *Security*. The *SQL* tab displays the SQL code generated by dialog selections.



Use the fields in the *General* tab to identify the table:

- Use the *Name* field to add a descriptive name for the table. A table cannot have the same name as any existing table, sequence, index, view, foreign table, or data type in the same schema. The name specified will be displayed in the *pgAdmin* tree control. This field is required.
- Select the owner of the table from the drop-down listbox in the *Owner* field. By default, the owner of the table is the role that creates the table.
- Select the name of the schema in which the table will reside from the drop-down listbox in the *Schema* field.
- Use the drop-down listbox in the *Tablespace* field to specify the tablespace in which the table will be stored.
- Store notes about the table in the *Comment* field.

Click the *Columns* tab to continue.




Use the drop-down listbox next to *Inherited from table(s)* to specify any parent table(s); the table will inherit columns from the selected parent table(s). Click inside the *Inherited from table(s)* field to select a table name from a drop-down list. Repeat to add any other parent tables. Delete a selected table by clicking the *x* to the left of the parent name. Note that inherited column names and datatypes are not editable in the current dialog; they must be modified at the parent level.

Click the *Add* icon (+) to specify the names of columns and their datatypes in the *Columns* table:

- Use the *Name* field to add a descriptive name for the column.
- Use the drop-down listbox in the *Data type* field to select a data type for the column. This can include array specifiers. For more information on the data types supported by PostgreSQL, refer to Chapter 8 of the core documentation.
- If enabled, use the *Length* and *Precision* fields to specify the maximum number of significant digits in a numeric value, or the maximum number of characters in a text value.
- Move the *Not NULL?* switch to the *Yes* position to require a value in the column field.
- Move the *Primary key?* switch to the *Yes* position to specify the column is the primary key constraint.

Click the *Add* icon (+) to add additional columns; to discard a column, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *Constraints* tab to continue.



Use the fields in the *Constraints* tab to provide a table or column constraint. Optional constraint clauses specify constraints (tests) that new or updated rows must satisfy for an *INSERT* or *UPDATE* operation to succeed. Select the appropriate constraint type by selecting one of the following tabs on the *Constraints* panel:

Tab Name	Constraint
<i>Primary Key</i>	Provides a unique identifier for each row in the table.
<i>Foreign Key</i>	Maintains referential integrity between two tables.
<i>Check</i>	Requires data satisfies an expression or condition before insertion or modification.
<i>Unique</i>	Ensures that the data contained in a column, or a group of columns, is unique among all the rows in the table.
<i>Exclude</i>	Guarantees that if any two rows are compared on the specified column or expression (using the specified operator), at least one of the operator comparisons will return false or null.

To add a primary key for the table, select the *Primary Key* tab, and click the *Add* icon (+). To define the primary key, click the *Edit* icon to the left of the *Trash* icon. A dialog similar to the *Primary key* dialog (accessed by right clicking on *Constraints* in the *pgAdmin* tree control) opens.

Use the fields in the *General* tab to identify the primary key:

- Use the *Name* field to add a descriptive name for the primary key constraint. The name will be displayed in the

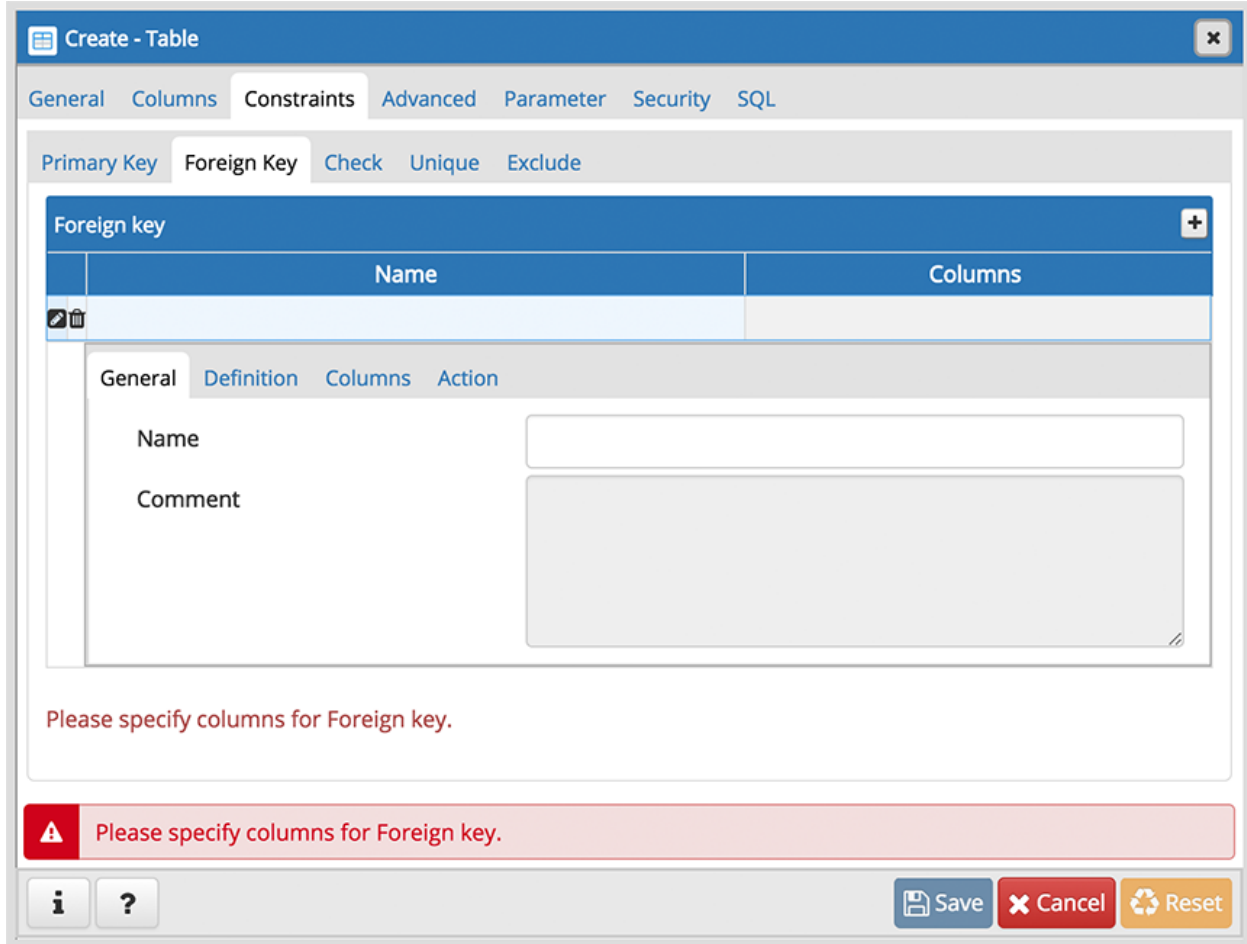
pgAdmin tree control.

- Provide notes about the primary key in the *Comment* field.

Click the *Definition* tab to continue.

Use the fields in the *Definition* tab to define the primary key constraint:

- Click inside the *Columns* field and select one or more column names from the drop-down listbox. To delete a selection, click the *x* to the left of the column name. The primary key constraint should be different from any unique constraint defined for the same table; the selected column(s) for the constraints must be distinct.
- Select the name of the tablespace in which the primary key constraint will reside from the drop-down listbox in the *Tablespace* field.
- Use the *Fill Factor* field to specify a fill factor for the table and index. The fill factor for a table is a percentage between 10 and 100. 100 (complete packing) is the default.
- Move the *Deferrable?* switch to the *Yes* position to specify the timing of the constraint is deferrable and can be postponed until the end of the statement. The default is *No*.
- If enabled, move the *Deferred?* switch to the *Yes* position to specify the timing of the constraint is deferred to the end of the statement. The default is *No*.



To add a foreign key constraint, select the *Foreign Key* tab, and click the *Add* icon (+). To define the constraint, click the *Edit* icon to the left of the *Trash* icon. A dialog similar to the *Foreign key* dialog (accessed by right clicking on *Constraints* in the *pgAdmin* tree control) opens.

Use the fields in the *General* tab to identify the foreign key constraint:

- Use the *Name* field to add a descriptive name for the foreign key constraint. The name will be displayed in the *pgAdmin* tree control.
- Provide notes about the foreign key in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Foreign key' dialog box with the 'Definition' tab selected. The 'Definition' tab contains the following fields and controls:

- Deferrable?**: A toggle switch set to 'Yes'.
- Deferred?**: A toggle switch set to 'Yes'.
- Match type**: A dropdown menu set to 'FULL'.
- Validated?**: A toggle switch set to 'Yes'.
- Auto FK index?**: A toggle switch set to 'Yes'.
- Covering index**: A text input field containing the value 'fki_distributors'.

At the bottom of the dialog, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). To the left of these buttons are two smaller buttons: an information icon ('i') and a help icon ('?').

Use the fields in the *Definition* tab to define the foreign key constraint:

- Move the *Deferrable?* switch to the *Yes* position to specify the timing of the constraint is deferrable and can be postponed until the end of the statement. The default is *No*.
- If enabled, move the *Deferred?* switch to the *Yes* position to specify the timing of the constraint is deferred to the end of the statement. The default is *No*.
- Move the *Match type* switch specify the type of matching that is enforced by the constraint:
 - Select *Full* to indicate that all columns of a multicolumn foreign key must be null if any column is null; if all columns are null, the row is not required to have a match in the referenced table.

- Select *Simple* to specify that a single foreign key column may be null; if any column is null, the row is not required to have a match in the referenced table.
- Move the *Validated* switch to the *Yes* position to instruct the server to validate the existing table content (against a foreign key or check constraint) when you save modifications to this dialog.
- Move the *Auto FK Index* switch to the *No* position to disable the automatic index feature.
- The field next to *Covering Index* generates the name of an index if the *Auto FK Index* switch is in the *Yes* position; or, this field is disabled.

Click the *Columns* tab to continue.

The screenshot shows the 'Create - Foreign key' dialog box in pgAdmin 4. The 'Columns' tab is selected, displaying a table with the following structure:

	Local	Referenced
Local column	agent_id	
References	Select foreign table	
Referencing	Select column	

A red warning message at the bottom of the dialog reads: "Please specify columns for Foreign key." The dialog also includes buttons for 'Save', 'Cancel', and 'Reset'.

Use the fields in the *Columns* tab to specify one or more reference column(s). A Foreign Key constraint requires that one or more columns of a table must only contain values that match values in the referenced column(s) of a row of a referenced table:

- Use the drop-down listbox next to *Local column* to specify the column in the current table that will be compared to the foreign table.
- Use the drop-down listbox next to *References* to specify the name of the table in which the comparison column(s) resides.
- Use the drop-down listbox next to *Referencing* to specify a column in the foreign table.

Click the *Add* icon (+) to add a column to the list; repeat the steps above and click the *Add* icon (+) to add additional columns. To discard an entry, click the trash icon to the left of the entry and confirm deletion in the *Delete Row* popup.

Click the *Action* tab to continue.

The screenshot shows the 'Create - Foreign key' dialog box in pgAdmin 4. The 'Action' tab is selected, showing two drop-down menus: 'On update' and 'On delete', both currently set to 'NO ACTION'. A red error bar at the bottom of the dialog reads: 'Please specify columns for Foreign key.' The dialog also features an information icon, a help icon, and buttons for 'Save', 'Cancel', and 'Reset'.

Use the drop-down listboxes on the *Action* tab to specify behavior related to the foreign key constraint that will be performed when data within the table is updated or deleted:

- Use the drop-down listbox next to *On update* to select an action that will be performed when data in the table is updated.
- Use the drop-down listbox next to *On delete* to select an action that will be performed when data in the table is deleted.

The supported actions are:

NO ACTION	Produce an error indicating that the deletion or update will create a foreign key constraint violation. If the constraint is deferred, this error will be produced at constraint check time if any referencing rows still exist. This is the default.
RESTRICT	Throw an error indicating that the deletion or update would create a foreign key constraint violation. This is the same as NO ACTION except that the check is not deferrable.
CASCADE	Delete any rows referencing the deleted row, or update the values of the referencing column(s) to the new values of the referenced columns, respectively.
SET NULL	Set the referencing column(s) to null.
SET DEFAULT	Set the referencing column(s) to their default values. There must be a row in the referenced table that matches the default values (if they are not null), or the operation will fail.

The screenshot shows the 'Create - Table' dialog in pgAdmin 4, specifically the 'Constraints' tab and the 'Check' sub-tab. The dialog has a header bar with 'Create - Table' and a close button. Below the header are tabs for 'General', 'Columns', 'Constraints', 'Advanced', 'Parameter', 'Security', and 'SQL'. Under the 'Constraints' tab, there are sub-tabs for 'Primary Key', 'Foreign Key', 'Check', 'Unique', and 'Exclude'. The 'Check' sub-tab is active, showing a table with two columns: 'Name' and 'Check'. Below the table is a 'General' tab with fields for 'Name' and 'Comment'. A red error message at the bottom reads 'Please specify columns for Foreign key.' Buttons for 'Save', 'Cancel', and 'Reset' are visible at the bottom right.

To add a check constraint, select the *Check* tab on the panel, and click the *Add* icon (+). To define the check constraint, click the *Edit* icon to the left of the *Trash* icon. A dialog similar to the *Check* dialog (accessed by right clicking on *Constraints* in the *pgAdmin* tree control) opens.

Use the fields in the *General* tab to identify the check constraint:

- Use the *Name* field to add a descriptive name for the check constraint. The name will be displayed in the *pgAdmin* tree control. With PostgreSQL 9.5 forward, when a table has multiple check constraints, they will be tested for each row in alphabetical order by name and after NOT NULL constraints.
- Provide notes about the check constraint in the *Comment* field.

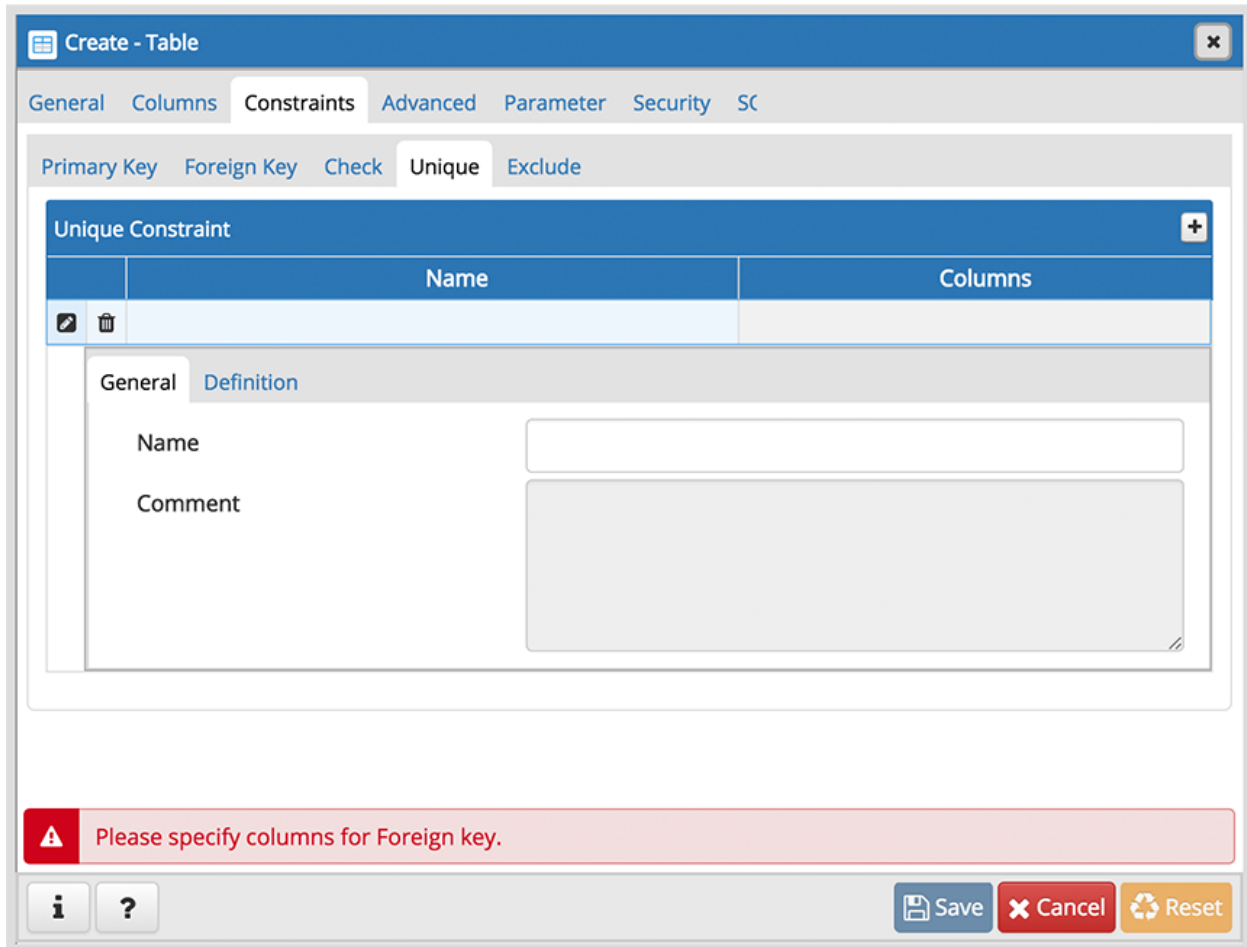
Click the *Definition* tab to continue.



The screenshot shows the 'Create - Check' dialog box in pgAdmin 4. The window title is 'Create - Check' with a checkmark icon and a close button. The dialog has three tabs: 'General', 'Definition', and 'SQL'. The 'Definition' tab is selected. It contains a large text area labeled 'Check' with a vertical cursor. Below it are two toggle switches: 'No Inherit?' with a blue 'No' button selected, and 'Don't validate?' with a green 'Yes' button selected. At the bottom, there are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (orange). On the left side of the bottom bar, there are two small buttons with an 'i' and a '?' icon.

Use the fields in the *Definition* tab to define the check constraint:

- Provide the expression that a row must satisfy in the *Check* field. This field is required.
- Move the *No Inherit?* switch to the *Yes* position to specify this constraint is automatically inherited by a table's children. The default is *No*.
- Move the *Don't validate?* switch to the *No* position to skip validation of existing data; the constraint may not hold for all rows in the table. The default is *Yes*.



To add a unique constraint, select the *Unique* tab on the panel, and click the *Add* icon (+). To define the constraint, click the *Edit* icon to the left of the *Trash* icon. A dialog similar to the *Unique constraint* dialog (accessed by right clicking on *Constraints* in the *pgAdmin* tree control) opens.

Use the fields in the *General* tab to identify the unique constraint:

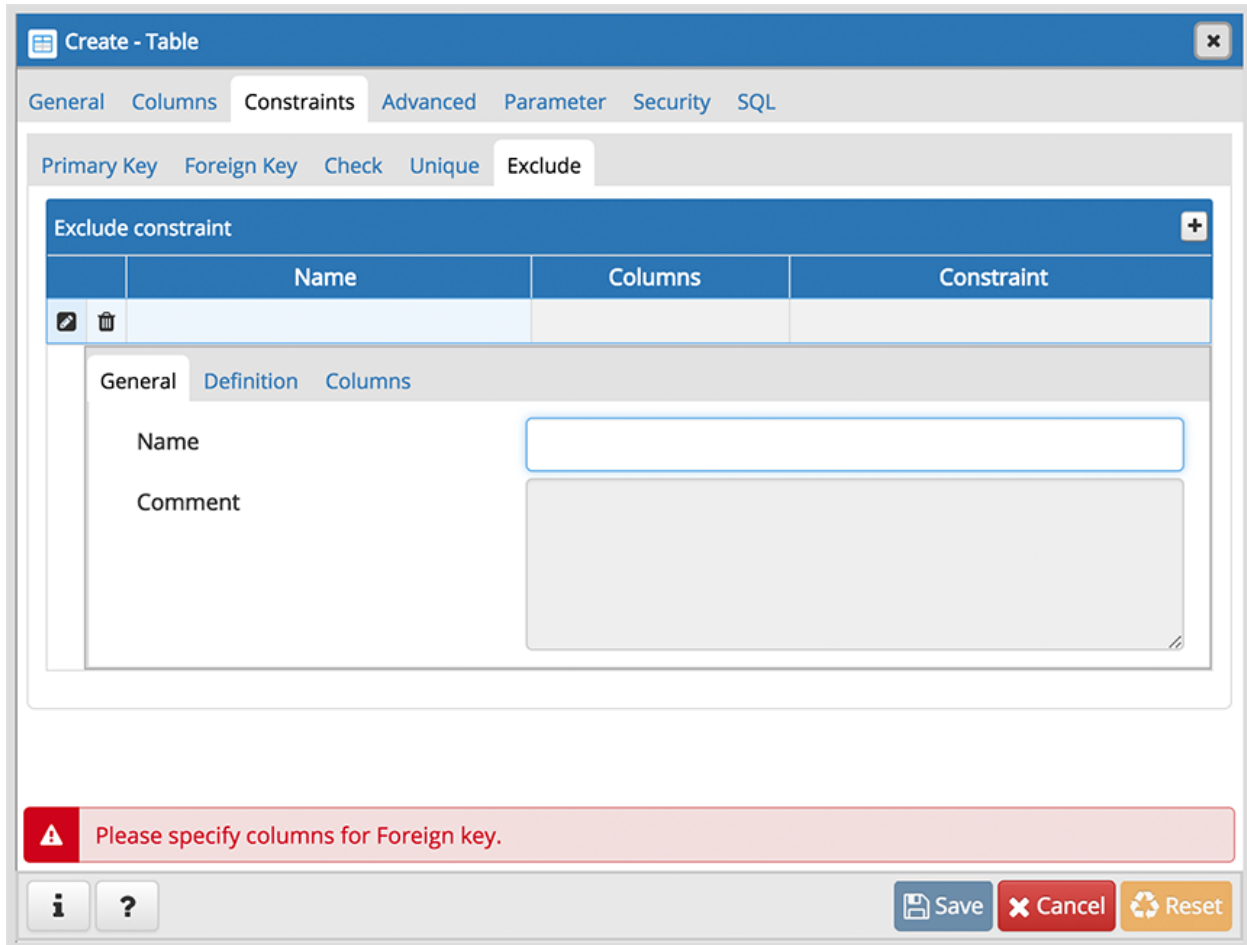
- Use the *Name* field to add a descriptive name for the unique constraint. The name will be displayed in the *pgAdmin* tree control.
- Provide notes about the unique constraint in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Unique constraint' dialog box. The 'Definition' tab is selected. The 'Columns' field contains 'x col1'. The 'Include columns' field contains 'x col2 x col3'. The 'Tablespace' and 'Index' dropdowns are set to 'Select from the list'. The 'Fill factor' field is empty. The 'Deferrable?' and 'Deferred?' toggle switches are both set to 'Yes'. The bottom of the dialog features 'Save', 'Cancel', and 'Reset' buttons, along with help and question mark icons.

Use the fields in the *Definition* tab to define the unique constraint:

- Click inside the *Columns* field and select one or more column names from the drop-down listbox. To delete a selection, click the *x* to the left of the column name. The unique constraint should be different from the primary key constraint defined for the same table; the selected column(s) for the constraints must be distinct.
- Select the name of the tablespace in which the unique constraint will reside from the drop-down listbox in the *Tablespace* field.
- Use the *Fill Factor* field to specify a fill factor for the table and index. The fill factor for a table is a percentage between 10 and 100. 100 (complete packing) is the default.
- Move the *Deferrable?* switch to the *Yes* position to specify the timing of the constraint is deferrable and can be postponed until the end of the statement. The default is *No*.
- If enabled, move the *Deferred?* switch to the *Yes* position to specify the timing of the constraint is deferred to the end of the statement. The default is *No*.




To add an exclusion constraint, select the *Exclude* tab on the panel, and click the *Add* icon (+). To define the constraint, click the *Edit* icon to the left of the *Trash* icon. A dialog similar to the *Exclusion constraint* dialog (accessed by right clicking on *Constraints* in the *pgAdmin* tree control) opens.

Use the fields in the *General* tab to identify the exclusion constraint:

- Use the *Name* field to provide a descriptive name for the exclusion constraint. The name will be displayed in the *pgAdmin* tree control.
- Provide notes about the exclusion constraint in the *Comment* field.

Click the *Definition* tab to continue.



The screenshot shows the 'Create - Exclusion constraint' dialog box in pgAdmin 4. The 'Definition' tab is selected, showing the following fields:

- Tablespace:** A drop-down menu with 'pg_default' selected.
- Access method:** A drop-down menu with 'btree' selected.
- Fill factor:** An empty text input field.
- Deferrable?:** A radio button with 'Yes' selected.
- Deferred?:** A radio button with 'Yes' selected.
- Constraint:** A large empty text area for defining the constraint.

At the bottom of the dialog, there are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields in the *Definition* tab to define the exclusion constraint:

- Use the drop-down listbox next to *Tablespace* to select the tablespace in which the index associated with the exclude constraint will reside.
- Use the drop-down listbox next to *Access method* to specify the type of index that will be used when implementing the exclusion constraint:
 - Select *gist* to specify a GiST index (the default).
 - Select *spgist* to specify a space-partitioned GiST index.
 - Select *btree* to specify a B-tree index.

- Select *hash* to specify a hash index.
- Use the *Fill Factor* field to specify a fill factor for the table and associated index. The fill factor is a percentage between 10 and 100. 100 (complete packing) is the default.
- Move the *Deferrable?* switch to the *Yes* position to specify that the timing of the constraint is deferrable, and can be postponed until the end of the statement. The default is *No*.
- If enabled, move the *Deferred?* switch to the *Yes* position to specify the timing of the constraint is deferred to the end of the statement. The default is *No*.
- Use the *Constraint* field to provide a condition that a row must satisfy to be included in the table.

Click the *Columns* tab to continue.

The screenshot shows the 'Create - Exclusion constraint' dialog box with the 'Columns' tab selected. The dialog has four tabs: General, Definition, Columns, and SQL. The 'Columns' tab is active, showing a table with the following data:

Column	Operator class	DESC	NULLs order	Operator
col1	cidr_ops	ASC	FIRST	<=

Below the table, there is an 'Include columns' field with two entries: 'x col2' and 'x col3'. At the bottom of the dialog, there are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields in the *Columns* tab to specify the column(s) to which the constraint applies. Use the drop-down listbox next to *Column* to select a column and click the *Add* icon (+) to provide details of the action on the column:

- The *Column* field is populated with the selection made in the *Column* drop-down listbox.
- If applicable, use the drop-down listbox in the *Operator class* to specify the operator class that will be used by the index for the column.
- Move the *DESC* switch to *DESC* to specify a descending sort order. The default is *ASC* which specifies an ascending sort order.
- Move the *NULLs order* switch to *LAST* to define an ascending sort order for NULLs. The default is *FIRST* which specifies a descending order.
- Use the drop-down list next to *Operator* to specify a comparison or conditional operator.

Click the *Advanced* tab to continue.

Use the fields in the *Advanced* tab to define advanced features for the table:

- Use the drop-down listbox next to *Of type* to copy the table structure from the specified composite type. Please note that a typed table will be dropped if the type is dropped (with `DROP TYPE ... CASCADE`).
- Use the *Fill Factor* field to specify a fill factor for the table. The fill factor for a table is a percentage between 10 and 100. 100 (complete packing) is the default.
- Move the *Has OIDs?* switch to the *Yes* position to specify that each row within a table has a system-assigned object identifier. The default is *No*.
- Move the *Unlogged?* switch to the *Yes* position to disable logging for the table. Data written to an unlogged table is not written to the write-ahead log. Any indexes created on an unlogged table are automatically unlogged as well. The default is *No*.

Use the fields in the **Like** box to specify which attributes of an existing table from which a table will automatically copy column names, data types, and not-null constraints; after saving the new or modified table, any changes to the original table will not be applied to the new table.

- Use the drop-down listbox next to *Relation* to select a reference table.
- Move the *With default values?* switch to the *Yes* position to copy default values.
- Move the *With constraints?* switch to the *Yes* position to copy table and column constraints.
- Move the *With indexes?* switch to the *Yes* position to copy indexes.
- Move the *With storage?* switch to the *Yes* position to copy storage settings.
- Move the *With comments?* switch to the *Yes* position to copy comments.

With PostgreSQL 10 forward, the *Partition* tab will be visible.

Click the *Partition* tab to continue.

The screenshot shows the 'Create - Table' dialog box with the 'Partition' tab selected. The 'Partition Type' dropdown is set to 'Range'. Below it, the 'Partition Keys' section contains a table with columns 'Key type', 'Column', and 'Expression', which is currently empty. A text box below the table provides instructions on how to define partition keys. The 'Partitions' section contains a table with columns 'Operation', 'Name', 'From', 'To', 'In', 'Modulus', and 'Remainder', which is also empty. A text box below this table provides instructions on how to define partitions. At the bottom of the dialog, there are 'Save', 'Cancel', and 'Reset' buttons.

Use the fields in the *partition* tab to create the partitions for the table:

- Select a partition type from the *Partition Type* selection box. There are 3 options available; Range, List and Hash. Hash option will only enable for PostgreSQL version ≥ 11 .

Use the *Partition Keys* panel to define the partition keys. Click the *Add* icon (+) to add each partition keys selection:

- Select a partition key type in the *Keytype* field.
- Select a partition column in the *Column* field if Column option selected for *Keytype* field .
- Specify the expression in the *Expression* field if Expression option selected for the *Keytype* field.

Use the *Partitions* panel to define the partitions of a table. Click the *Add* icon (+) to add each partition:

- Move the *Operation* switch to *attach* to attach the partition, by default it is *create*.
- Use the *Name* field to add the name of the partition.
- If partition type is Range then *From* and *To* fields will be enabled.
- If partition type is List then *In* field will be enabled.
- If partition type is Hash then *Modulus* and *Remainder* fields will be enabled.

Click the *Parameter* tab to continue.

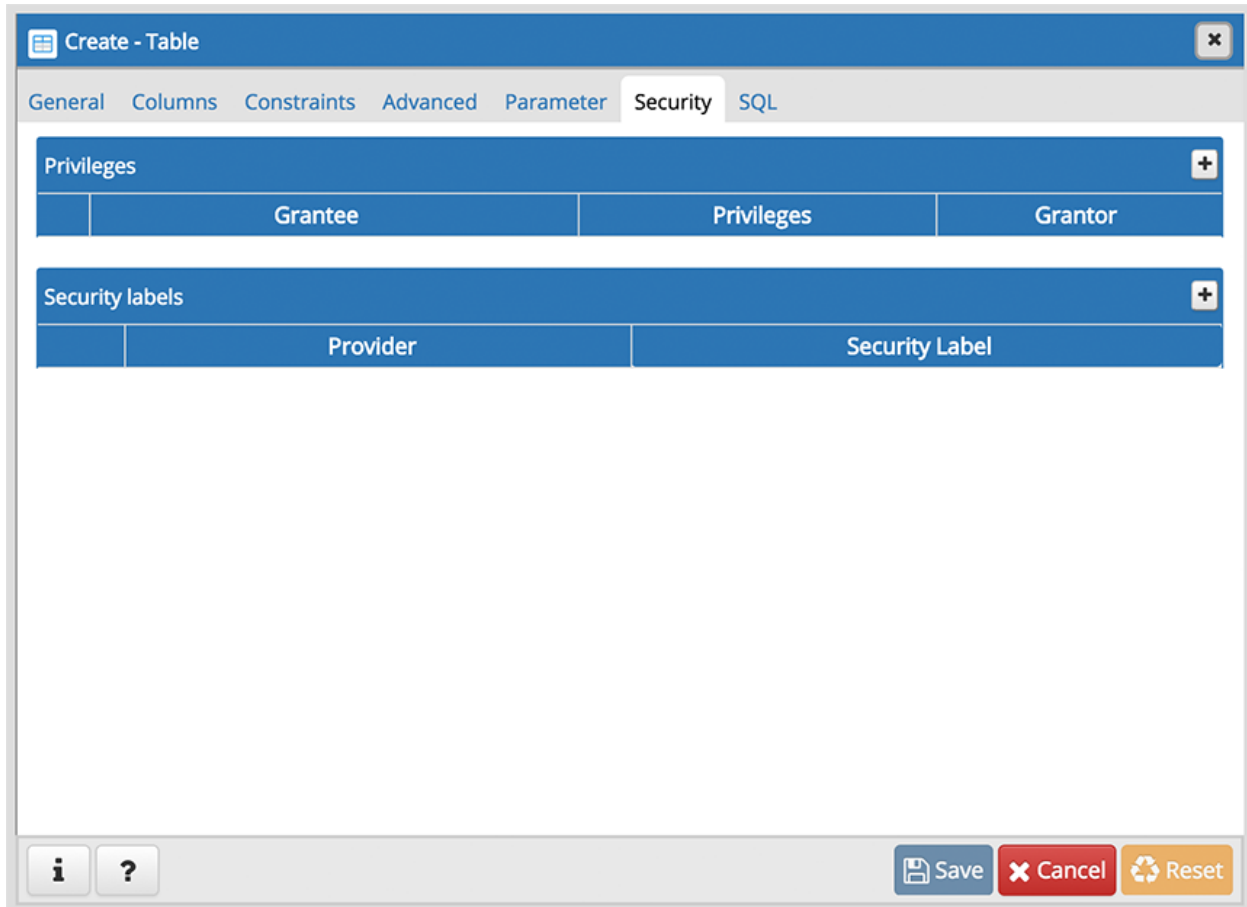
The screenshot shows the 'Create - Table' dialog box in pgAdmin 4, with the 'Parameter' tab selected. The 'Table' sub-tab is set to 'Toast Table'. There are two toggle switches: 'Custom auto-vacuum?' and 'Enabled?', both currently set to 'No'. Below these is a table titled 'Vacuum Table' with three columns: 'Label', 'Value', and 'Default value'. The bottom of the dialog features an information icon, a help icon, and three buttons: 'Save', 'Cancel', and 'Reset'.

Use the tabs nested inside the *Parameter* tab to specify VACUUM and ANALYZE thresholds; use the *Table* tab and the *Toast Table* tab to customize values for the table and the associated toast table:

- Move the *Custom auto-vacuum?* switch to the *Yes* position to perform custom maintenance on the table.
- Move the *Enabled?* switch to the *Yes* position to select values in the *Vacuum table*. The *Vacuum Table* provides default values for maintenance operations.

Provide a custom value in the *Value* column for each metric listed in the *Label* column.

Click the *Security* tab to continue.



Use the *Security* tab to assign privileges and define security labels.

Use the *Privileges* panel to assign privileges to a role. Click the *Add* icon (+) to set privileges for database objects:

- Select the name of the role from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privilege to the specified user.
- Select the name of the role from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.

Click the *Add* icon (+) to assign additional privileges; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Use the *Security Labels* panel to define security labels applied to the function. Click the *Add* icon (+) to add each security label selection:

- Specify a security label provider in the *Provider* field. The named provider must be loaded and must consent to the proposed labeling operation.
- Specify a security label in the *Security Label* field. The meaning of a given label is at the discretion of the label provider. PostgreSQL places no restrictions on whether or how a label provider must interpret security labels; it merely provides a mechanism for storing them.

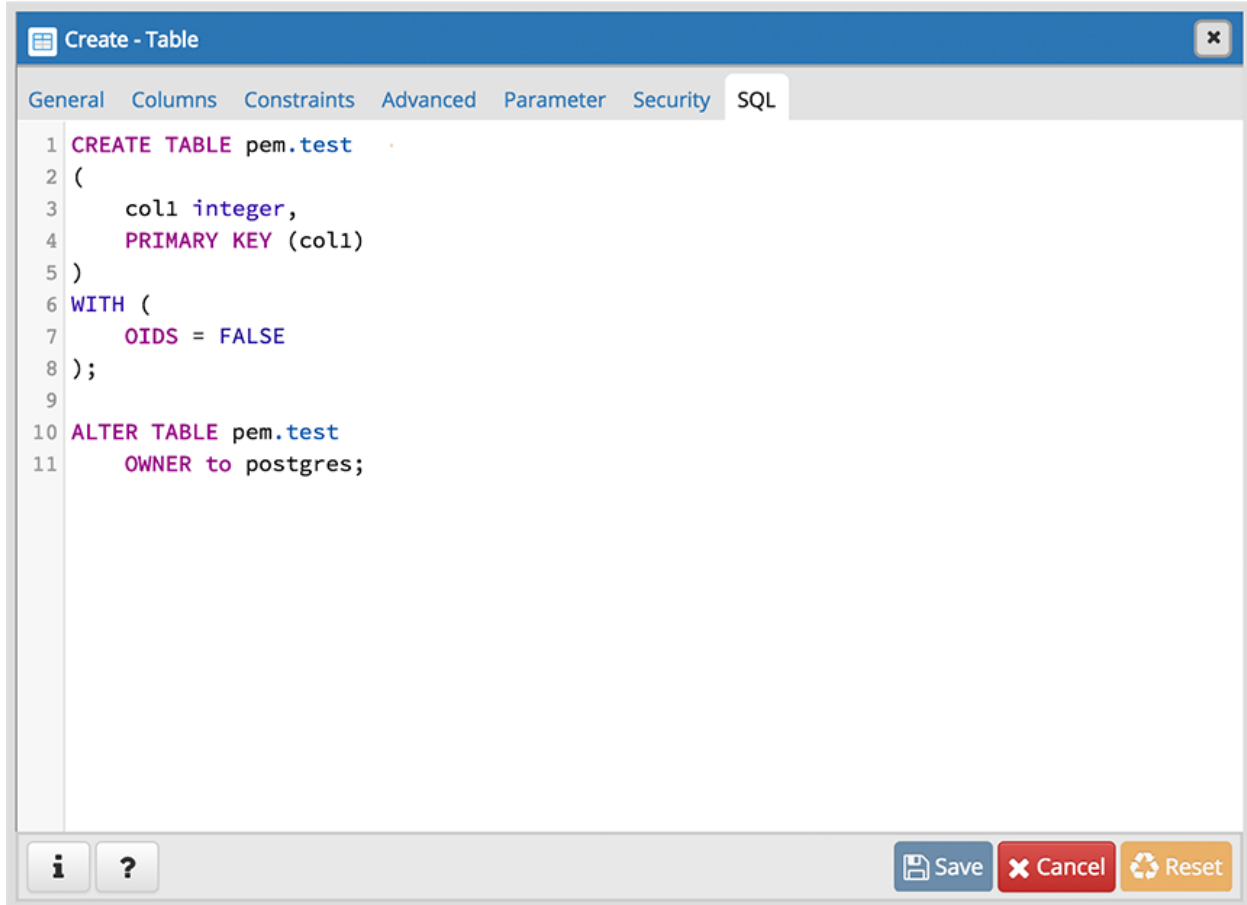
Click the *Add* icon (+) to assign additional security labels; to discard a security label, click the trash icon to the left of the row and confirm deletion in the *Delete Row* popup.

Click the *SQL* tab to continue.

Your entries in the *Table* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Table* dialog:



The example shown demonstrates creating a table named *product_category*. It has three columns and a primary key constraint on the *category_id* column.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

4.9 Trigger Dialog

Use the *Trigger* dialog to create a trigger or modify an existing trigger. A trigger executes a specified function when certain events occur.

The *Trigger* dialog organizes the development of a trigger through the following dialog tabs: *General*, *Definition*, *Events*, and *Code*. The *SQL* tab displays the SQL code generated by dialog selections.

The screenshot shows the 'Create - Trigger' dialog box. The title bar is blue and contains a close button. Below the title bar are five tabs: 'General', 'Definition', 'Events', 'Code', and 'SQL'. The 'General' tab is selected. The 'Name' field contains the text 'log_update'. The 'Comment' field contains the text 'Log any changes to data for the distributors table.'. At the bottom of the dialog are five buttons: an information button (i), a help button (?), a 'Save' button, a 'Cancel' button, and a 'Reset' button.

Use the fields in the *General* tab to identify the trigger:

- Use the *Name* field to add a descriptive name for the trigger. This must be distinct from the name of any other trigger for the same table. The name will be displayed in the *pgAdmin* tree control. Note that if multiple triggers of the same kind are defined for the same event, they will be fired in alphabetical order by name.
- Store notes about the trigger in the *Comment* field.

Click the *Definition* tab to continue.

The screenshot shows the 'Create - Trigger' dialog box in pgAdmin 4, with the 'Definition' tab selected. The dialog has a blue title bar with a close button. Below the title bar are tabs for 'General', 'Definition', 'Events', 'Code', and 'SQL'. The 'Definition' tab contains the following fields:

- 'Row trigger?' with a 'Yes' radio button selected.
- 'Constraint trigger?' with a 'Yes' radio button selected.
- 'Deferrable?' with a 'Yes' radio button selected.
- 'Deferred?' with a 'Yes' radio button selected.
- 'Trigger Function' with a text box containing 'public."ET"' and a dropdown arrow.
- 'Arguments' with an empty text box.

At the bottom of the dialog are buttons for 'i', '?', 'Save', 'Cancel', and 'Reset'.

Use the fields in the *Definition* tab to define the trigger:

- Move the *Row trigger?* switch to the *No* position to disassociate the trigger from firing on each row in a table. The default is *Yes*.
- Move the *Constraint trigger?* switch to the *Yes* position to specify the trigger is a constraint trigger.
- If enabled, move the *Deferrable?* switch to the *Yes* position to specify the timing of the constraint trigger is deferrable and can be postponed until the end of the statement. The default is *No*.
- If enabled, move the *Deferred?* switch to the *Yes* position to specify the timing of the constraint trigger is deferred to the end of the statement causing the triggering event. The default is *No*.
- Use the drop-down listbox next to *Trigger Function* to select a trigger function or procedure.
- Use the *Arguments* field to provide an optional (comma-separated) list of arguments to the function when the trigger is executed. The arguments are literal string constants.

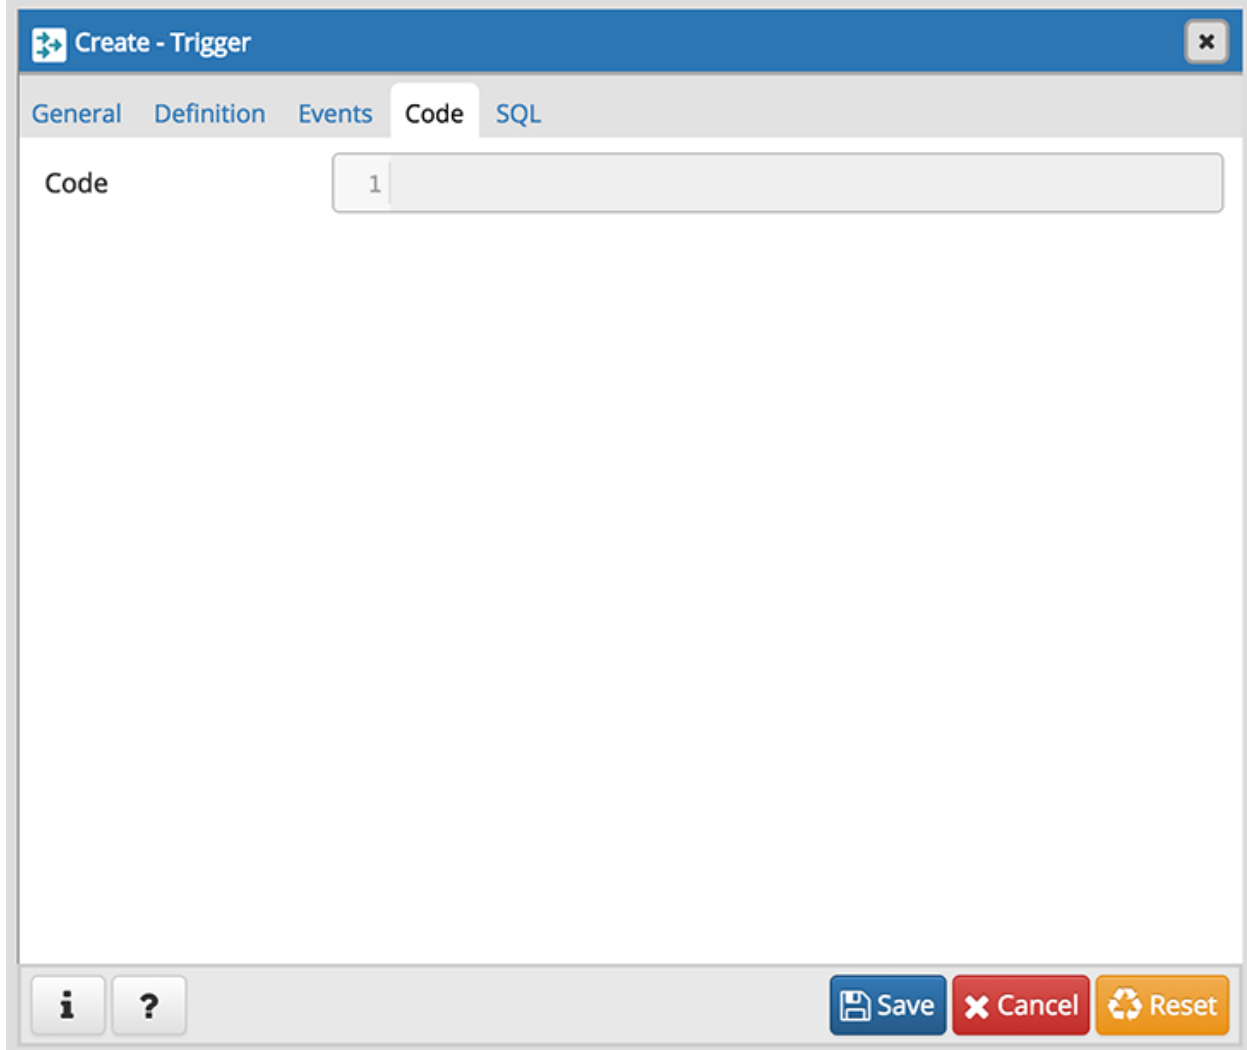
Click the *Events* tab to continue.

The screenshot shows the 'Create - Trigger' dialog box with the 'Events' tab selected. The 'Fires' dropdown is set to 'AFTER'. The 'Events' section contains four rows of event types with toggle switches: INSERT (Yes), UPDATE (Yes), DELETE (Yes), and TRUNCATE (No). The 'When' field contains the value '1'. The 'Columns' field contains a list with 'c'. The bottom of the dialog features buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields in the *Events* tab to specify how and when the trigger fires:

- Use the drop-down listbox next to the *Fires* fields to determine if the trigger fires *BEFORE* or *AFTER* a specified event. The default is *BEFORE*.
- Select the type of event(s) that will invoke the trigger; to select an event type, move the switch next to the event to the *YES* position. The supported event types are *INSERT*, *UPDATE*, *DELETE*, and *TRUNCATE*.
- Use the *When* field to provide a boolean condition that will invoke the trigger.
- If defining a column-specific trigger, use the *Columns* field to specify the columns or columns that are the target of the trigger.

Click the *Code* tab to continue.



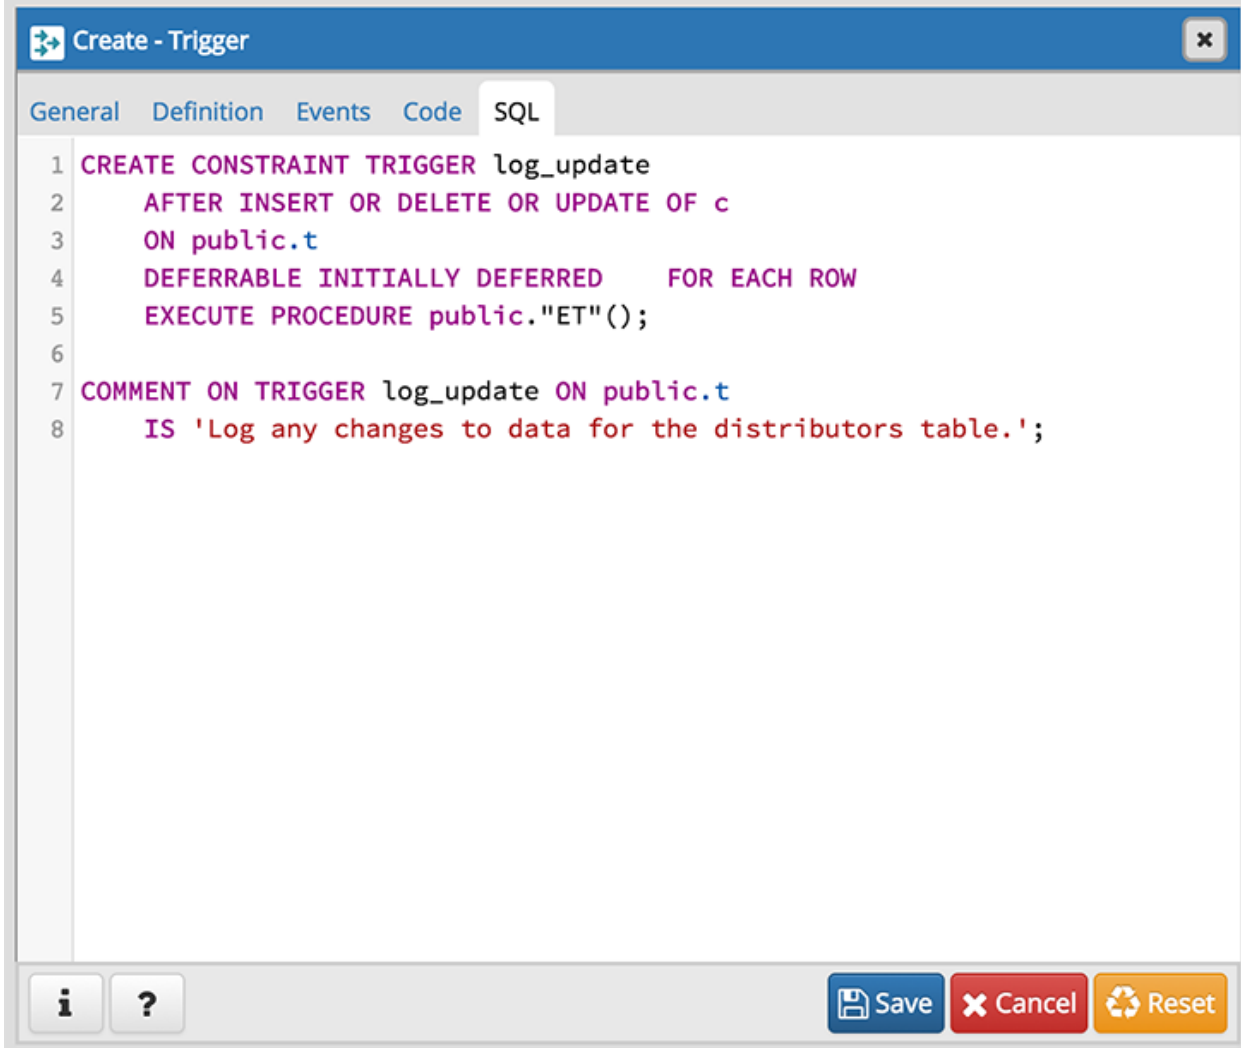
Use the *Code* field to specify any additional code that will be invoked when the trigger fires.

Click the *SQL* tab to continue.

Your entries in the *Trigger* dialog generate a *SQL* command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the *SQL* command.

Example

The following is an example of the *sql* command generated by user selections in the *Trigger* dialog:



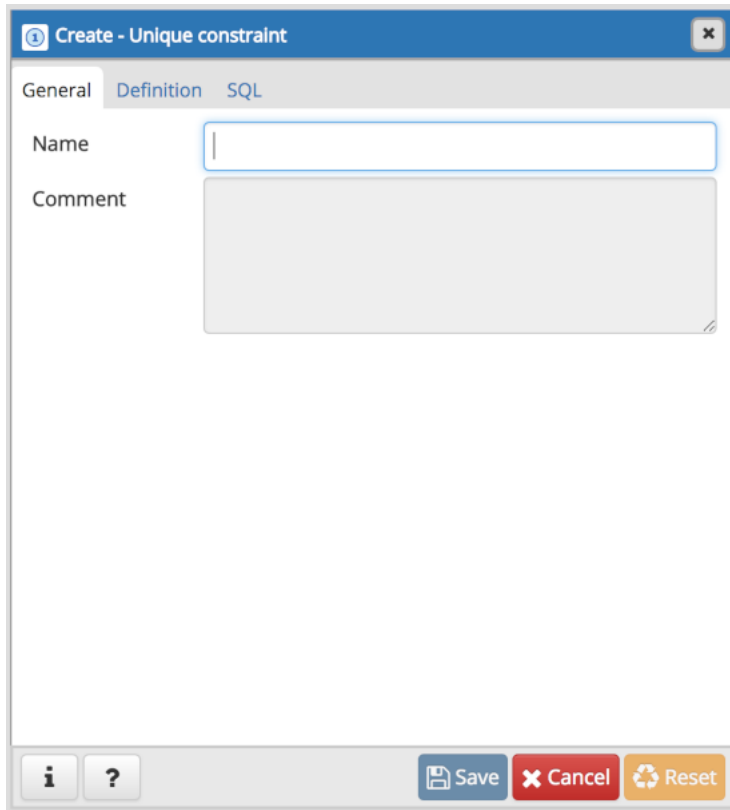
The example demonstrates creating a trigger named *log_update* that calls a procedure named *log_account_update* that logs any updates to the *distributors* table.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.

4.10 Unique Constraint Dialog

Use the *Unique constraint* dialog to define a unique constraint for a specified table. Unique constraints ensure that the data contained in a column, or a group of columns, is unique among all the rows in the table.

The *Unique constraint* dialog organizes the development of a unique constraint through the following dialog tabs: *General* and *Definition*. The *SQL* tab displays the SQL code generated by dialog selections.



Use the fields in the *General* tab to identify the unique constraint:

- Use the *Name* field to add a descriptive name for the unique constraint. The name will be displayed in the *pgAdmin* tree control.

Click the *Definition* tab to continue.

Use the fields in the *Definition* tab to define the unique constraint:

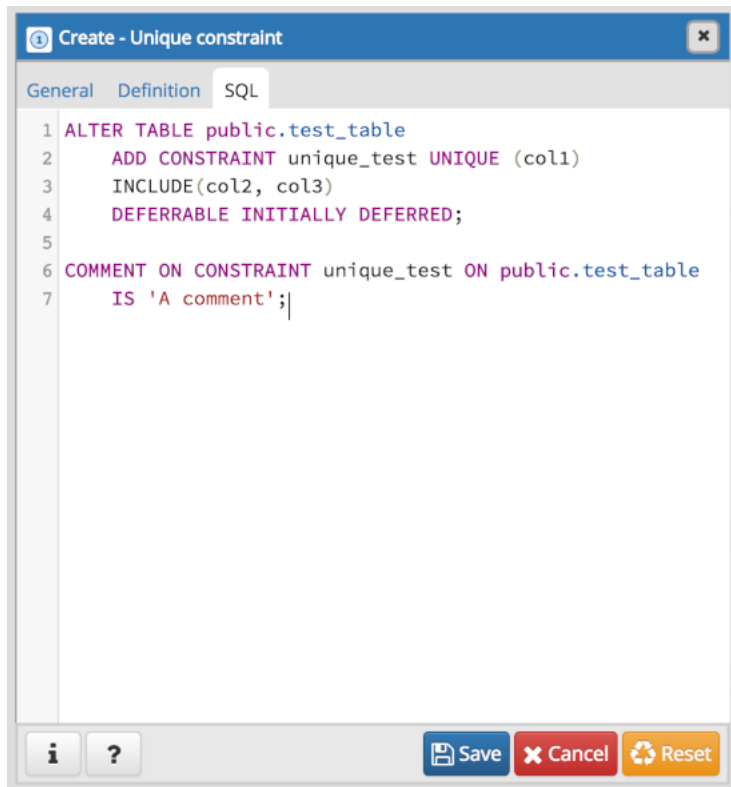
- Click inside the *Columns* field and select one or more column names from the drop-down listbox. To delete a selection, click the *x* to the left of the column name. The unique constraint should be different from the primary key constraint defined for the same table; the selected column(s) for the constraints must be distinct.
- Use *Include columns* field to specify columns for *INCLUDE* clause of the constraint. This option is available in Postgres 11 and later.
- Select the name of the tablespace in which the unique constraint will reside from the drop-down listbox in the *Tablespace* field.
- Select the name of an index from the drop-down listbox in the *Index* field. This field is optional. Adding a unique constraint will automatically create a unique B-tree index on the column or group of columns listed in the constraint, and will force the column(s) to be marked NOT NULL.
- Use the *Fill Factor* field to specify a fill factor for the table and index. The fill factor for a table is a percentage between 10 and 100. 100 (complete packing) is the default.
- Move the *Deferrable?* switch to the *Yes* position to specify the timing of the constraint is deferrable and can be postponed until the end of the statement. The default is *No*.
- If enabled, move the *Deferred?* switch to the *Yes* position to specify the timing of the constraint is deferred to the end of the statement. The default is *No*.

Click the *SQL* tab to continue.

Your entries in the *Unique constraint* dialog generate a SQL command (see an example below). Use the *SQL* tab for review; revisit or switch tabs to make any changes to the SQL command.

Example

The following is an example of the sql command generated by user selections in the *Unique constraint* dialog:



The example shown demonstrates creating a unique constraint named *name_con* on the *name* column of the *distributors* table.

- Click the *Info* button (i) to access online help. View context-sensitive help in the *Tabbed browser*, where a new tab displays the PostgreSQL core documentation.
- Click the *Save* button to save work.
- Click the *Cancel* button to exit without saving work.
- Click the *Reset* button to restore configuration parameters.


MANAGEMENT BASICS

pgAdmin provides point and click dialogs that help you perform server management functions. Dialogs simplify tasks such as managing named restore points, granting user privileges, and performing VACUUM, ANALYZE and REINDEX functions.

Contents:

5.1 Add named restore point Dialog

Use the *Add named restore point* dialog to take a named snapshot of the state of the server for use in a recovery file. To create a named restore point, the server's postgresql.conf file must specify a *wal_level* value of *archive*, *hot_standby*, or *logical*. You must be a database superuser to create a restore point.



When the *Restore point name* window launches, use the field *Enter the name of the restore point to add* to provide a descriptive name for the restore point.

For more information about using a restore point as a recovery target, please see the [PostgreSQL documentation](#).

- Click the *OK* button to save the restore point.
- Click the *Cancel* button to exit without saving work.


5.2 Change Password Dialog

It is a good policy to routinely change your password to protect data, even in what you may consider a 'safe' environment. In the workplace, failure to apply an appropriate password policy could leave you in breach of Data Protection laws.

Please consider the following guidelines when selecting a password:

- Ensure that your password is an adequate length; 6 characters should be the absolute minimum number of characters in the password.
- Ensure that your password is not open to dictionary attacks. Use a mixture of upper and lower case letters and numerics, and avoid words or names. Consider using the first letter from each word in a phrase that you will remember easily but is an unfamiliar acronym.
- Ensure that your password is changed regularly; at minimum, change it every ninety days.

The above should be considered a starting point: It is not a comprehensive list and it **will not guarantee security**.



The image shows a 'Change Password' dialog box with a blue title bar. It contains four text input fields: 'User' (with 'postgres' entered), 'Current Password', 'New Password', and 'Confirm Password'. At the bottom right, there are 'Ok' and 'Cancel' buttons.

Use the *Change Password* dialog to change your password:

- The name displayed in the *User* field is the role for which you are modifying the password; it is the role that is associated with the server connection that is highlighted in the tree control.
- Enter the password associated with the role in the *Current Password* field.
- Enter the desired password for in the *New Password* field.
- Re-enter the new password in the *Confirm Password* field.


Click the *OK* button to change your password; click *Cancel* to exit the dialog without changing your password.

5.3 Grant Wizard

The *Grant Wizard* tool is a graphical interface that allows you to manage the privileges of one or more database objects in a point-and-click environment. A search box, dropdown lists, and checkboxes facilitate quick selections of database objects, roles and privileges.

The wizard organizes privilege management through a sequence of windows: *Object Selection (step 1 of 3)*, *Privileges Selection (step 2 of 3)* and *Final (Review Selection) (step 3 of 3)*. The *Final (Review Selection)* window displays the SQL code generated by wizard selections.

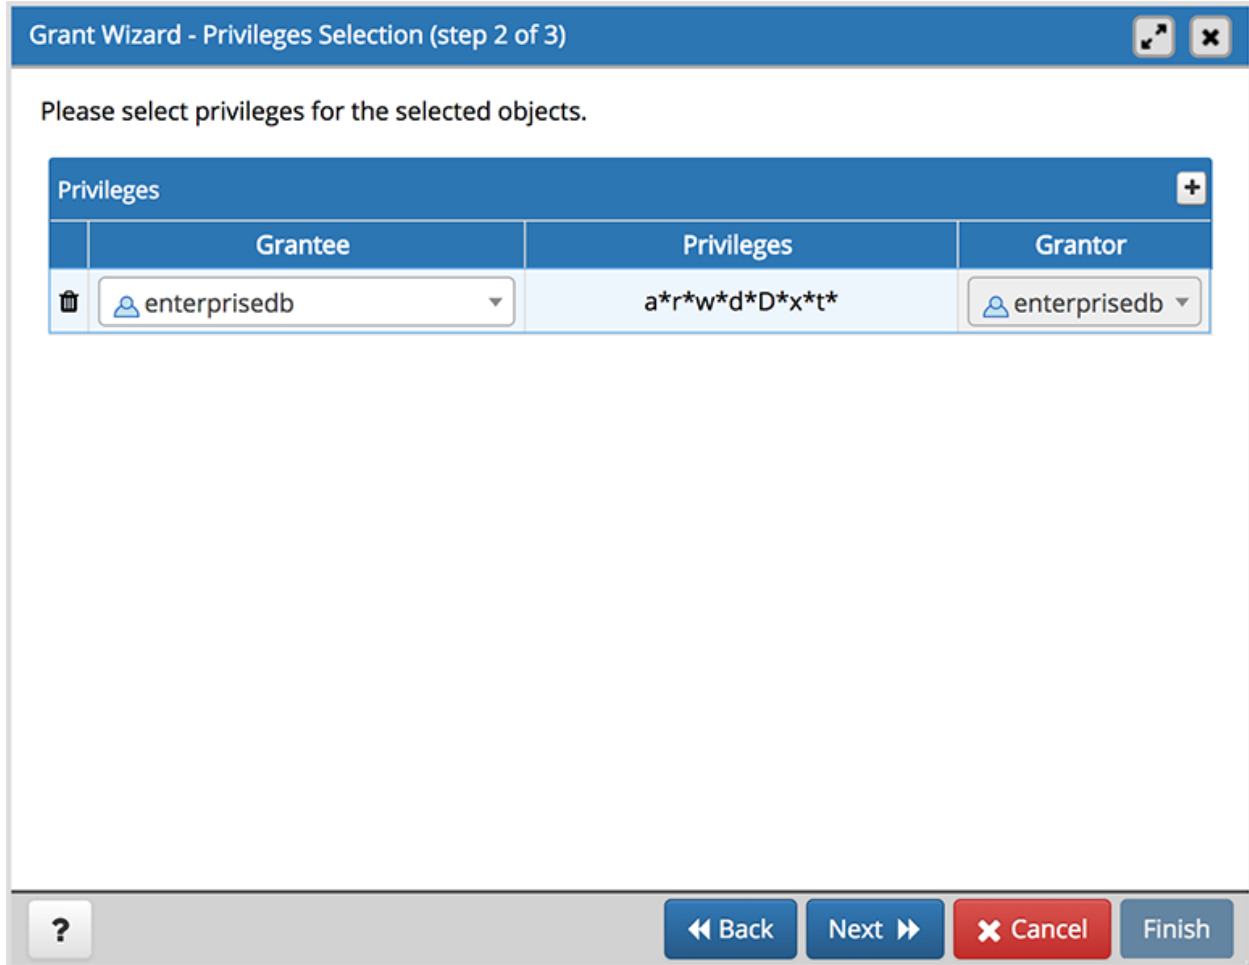
To launch the *Grant Wizard* tool, select a database object in the *pgAdmin* tree control, then navigate through *Tools* on the menu bar to click on the *Grant Wizard* option.



Use the fields in the *Object Selection (step 1 of 3)* window to select the object or objects on which you are modifying privileges. Use the *Search by object type or name* field to locate a database object, or use the scrollbar to scroll through the list of all accessible objects.

- Each row in the table lists object identifiers; check the checkbox in the left column to include an object as a target of the Grant Wizard. The table displays:
 - The object type in the *Object Type* field
 - The schema in which the object resides in the *Schema* field
 - The object name in the *Name* field.

Click the *Next* button to continue, or the *Cancel* button to close the wizard without modifying privileges.



Use the fields in the *Privileges Selection (step 2 of 3)* window to grant privileges. If you grant a privilege WITH GRANT OPTION, the Grantee will have the right to grant privileges on the object to others. If WITH GRANT OPTION is subsequently revoked, any role who received access to that object from that Grantee (directly or through a chain of grants) will lose their privileges on the object.

- Click the *Add* icon (+) to assign a set of privileges.
- Select the name of the role from the drop-down listbox in the *Grantee* field.
- Click inside the *Privileges* field. Check the boxes to the left of one or more privileges to grant the selected privileges to the specified user. If privileges have previously been granted on a database object, unchecking a privilege for a group or user will result in revoking that privilege.
- If enabled, select the name of the role from the drop-down listbox in the *Grantor* field. The default grantor is the owner of the database.
- Click the *Add* icon (+) to assign a set of privileges to another role; to discard a privilege, click the trash icon to the left of the row and confirm deletion in the *Delete Row* dialog.


For more information about granting privileges on database objects, see the [PostgreSQL core documentation](#).

Click the *Next* button to continue, the *Back* button to select or deselect additional database objects, or the *Cancel* button to close the wizard without modifying privileges.

Your entries in the *Grant Wizard* tool generate a SQL command; you can review the command in the *Final (Review Selection) (step 3 of 3)* window (see an example below).

Example

The following is an example of the sql command generated by user selections in the *Grant Wizard* tool:



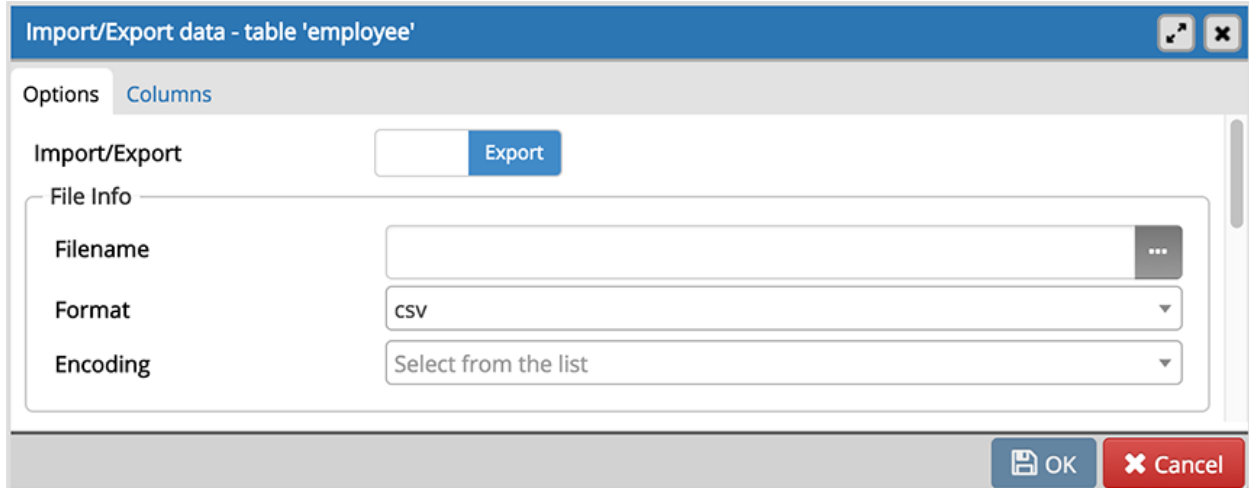
The commands displayed assign a role named *Bob* *INSERT* and *UPDATE* privileges *WITH GRANT OPTION* on the *sales_meetings* and the *sales_territories* tables.

- Click the *Back* button to select or deselect additional database objects, roles and privileges.
- Click the *Cancel* button to exit without saving work.
- Click the *Finish* button to save selections and exit the wizard.

5.4 Import/Export data Dialog

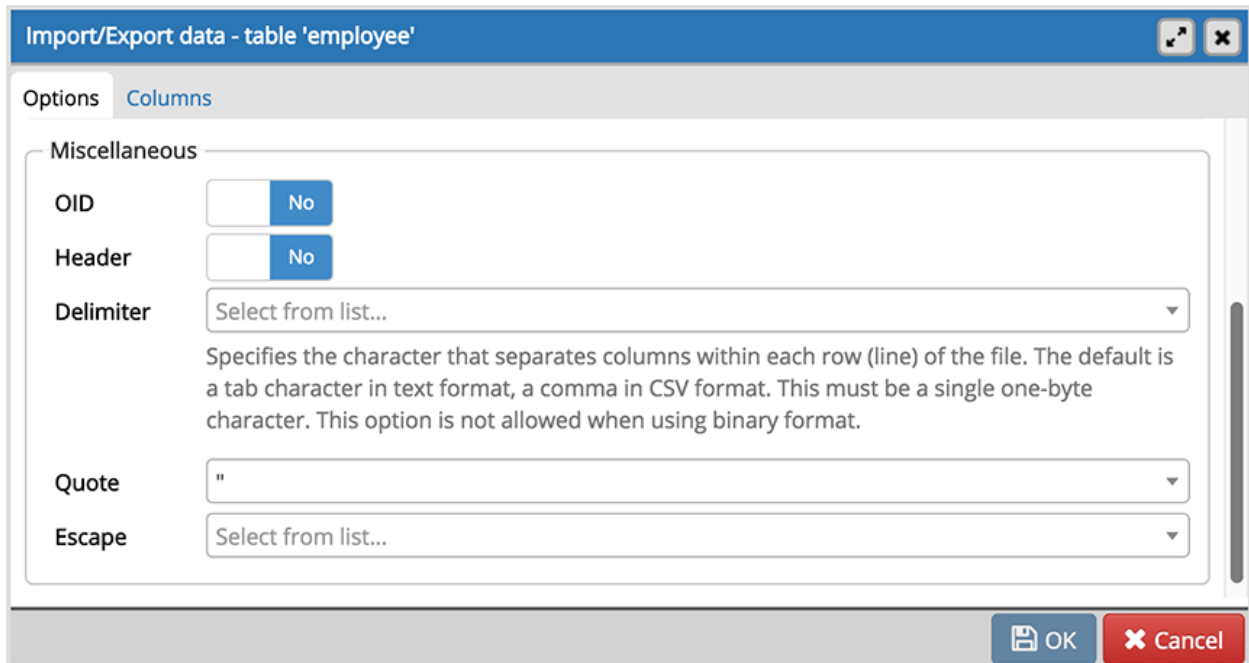
Use the *Import/Export data* dialog to copy data from a table to a file, or copy data from a file into a table.

The *Import/Export data* dialog organizes the import/export of data through the *Options* and *Columns* tabs.



Use the fields in the *Options* tab to specify import and export preferences:


- Move the *Import/Export* switch to the *Import* position to specify that the server should import data to a table from a file. The default is *Export*.
- Use the fields in the *File Info* field box to specify information about the source or target file:
 - Enter the name of the source or target file in the *Filename* field. Optionally, select the *Browser* icon (ellipsis) to the right to navigate into a directory and select a file.
 - Use the drop-down listbox in the *Format* field to specify the file type. Select:
 - * *binary* for a .bin file.
 - * *csv* for a .csv file.
 - * *text* for a .txt file.
 - Use the drop-down listbox in the *Encoding* field to specify the type of character encoding.



- Use the fields in the *Miscellaneous* field box to specify additional information:

- Move the *OID* switch to the *Yes* position to include the *OID* column. The *OID* is a system-assigned value that may not be modified. The default is *No*.
- Move the *Header* switch to the *Yes* position to include the table header with the data rows. If you include the table header, the first row of the file will contain the column names.
- If you are exporting data, specify the delimiter that will separate the columns within the target file in the *Delimiter* field. The separating character can be a colon, semicolon, a vertical bar, or a tab.
- Specify a quoting character used in the *Quote* field. Quoting can be applied to string columns only (i.e. numeric columns will not be quoted) or all columns regardless of data type. The character used for quoting can be a single quote or a double quote.
- Specify a character that should appear before a data character that matches the *QUOTE* value in the *Escape* field.

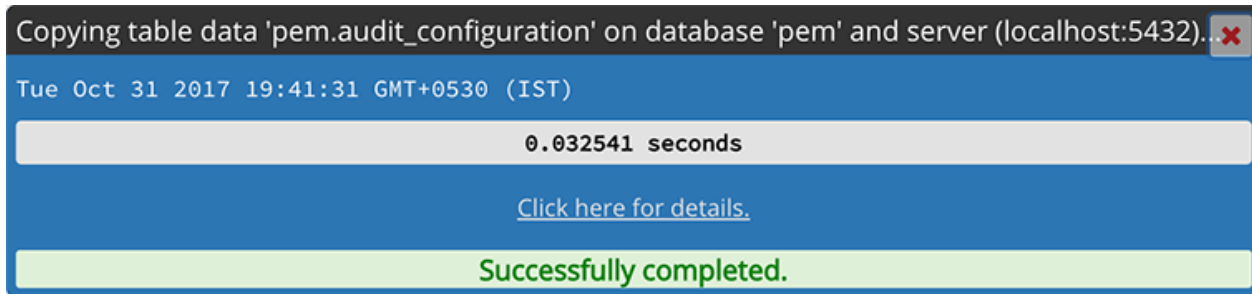
Click the *Columns* tab to continue.



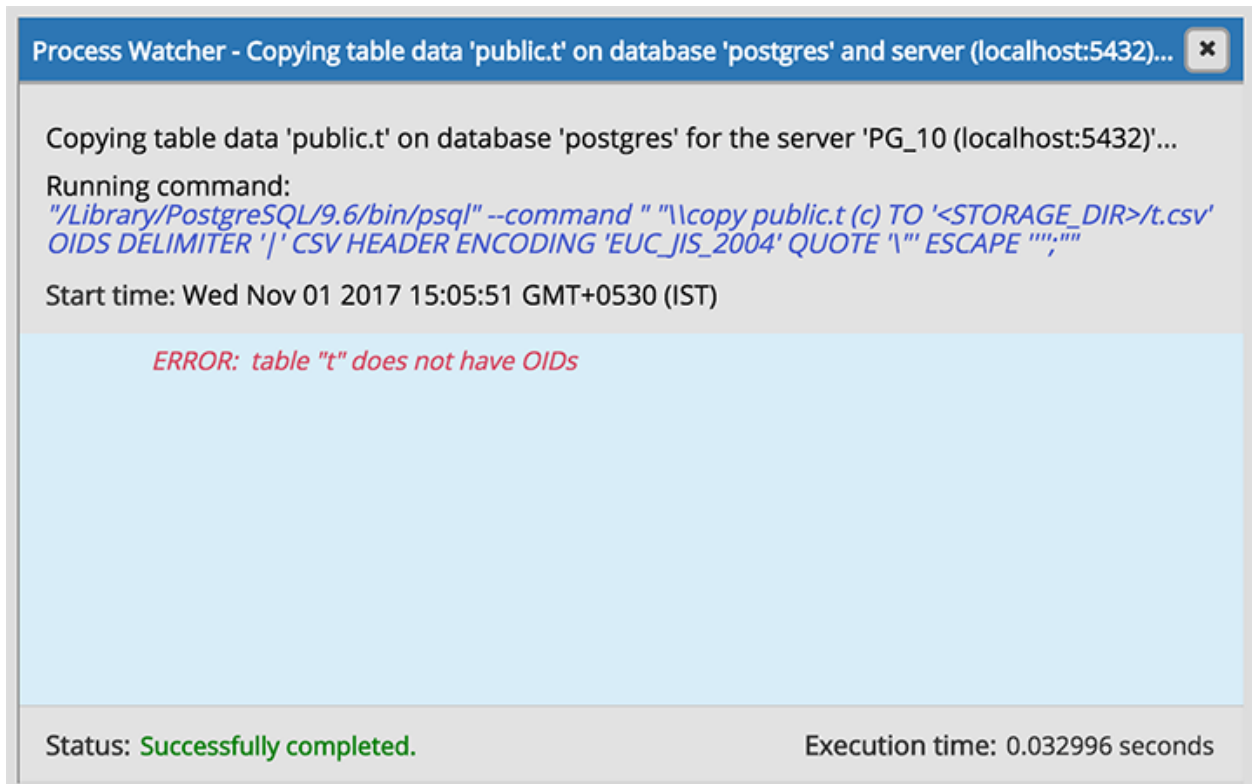
Use the fields in the *Columns* tab to select the columns that will be imported or exported:

- Click inside the *Columns to export/import* field to deselect one or more columns from the drop-down listbox. To delete a selection, click the *x* to the left of the column name. Click an empty spot inside the field to access the drop-down list.
- Use the *NULL Strings* field to specify a string that will represent a null value within the source or target file.
- If enabled, click inside the *Not null columns* field to select one or more columns that will not be checked for a NULL value. To delete a column, click the *x* to the left of the column name.

After completing the *Import/Export data* dialog, click the *OK* button to perform the import or export. pgAdmin will inform you when the background process completes:

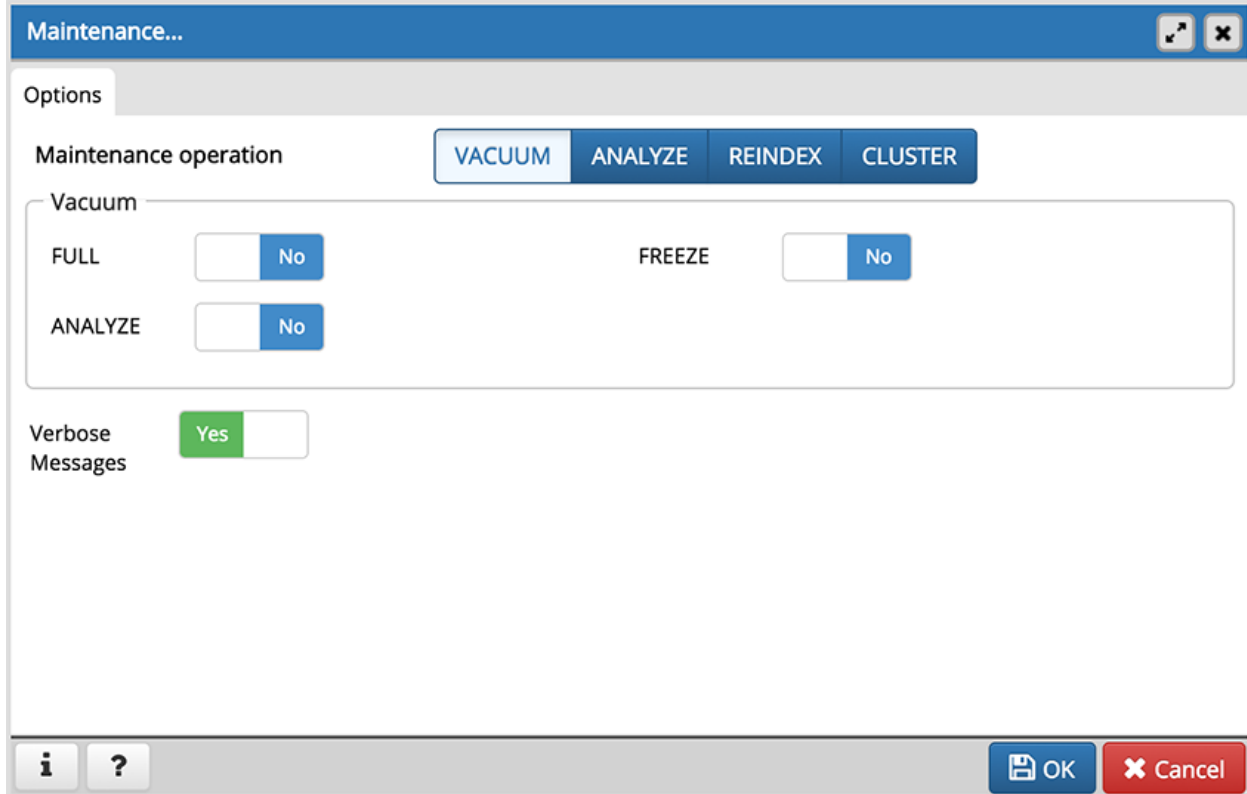


Use the *Click here for details* link on the notification to open the *Process Watcher* and review detailed information about the execution of the command that performed the import or export:



5.5 Maintenance Dialog

Use the *Maintenance* dialog to VACUUM, ANALYZE, REINDEX or CLUSTER a database or selected database objects.



While this utility is useful for ad-hoc maintenance purposes, you are encouraged to perform automatic VACUUM jobs on a regular schedule.

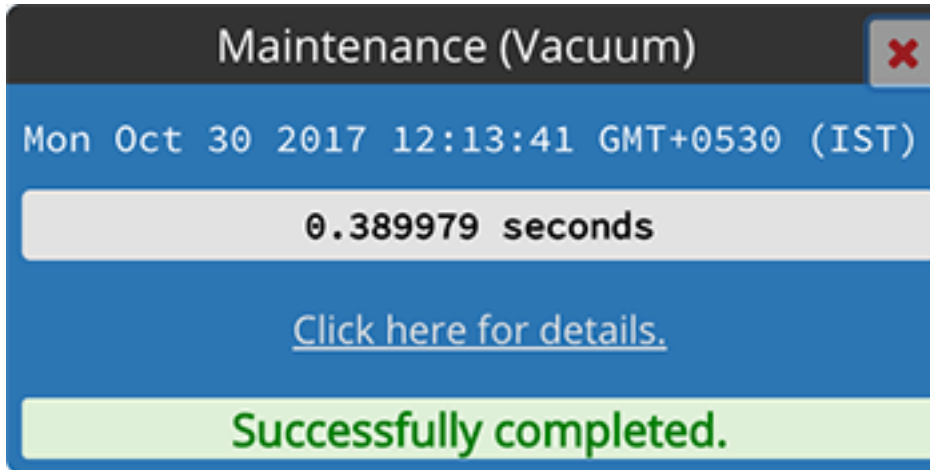
Select a button next to *Maintenance operation* to specify the type of maintenance:

- Click *VACUUM* to scan the selected database or table to reclaim storage used by dead tuples.
 - Move the *FULL* switch to the *Yes* position to compact tables by writing a completely new version of the table file without dead space. The default is *No*.
 - Move the *FREEZE* switch to the *Yes* position to freeze data in a table when it will have no further updates. The default is *No*.
 - Move the *ANALYZE* switch to the *Yes* position to issue *ANALYZE* commands whenever the content of a table has changed sufficiently. The default is *No*.
- Click *ANALYZE* to update the stored statistics used by the query planner. This enables the query optimizer to select the fastest query plan for optimal performance.
- Click *REINDEX* to rebuild any index in case it has degenerated due to the insertion of unusual data patterns. This happens, for example, if you insert rows with increasing index values, and delete low index values.
- Click *CLUSTER* to instruct PostgreSQL to cluster the selected table.

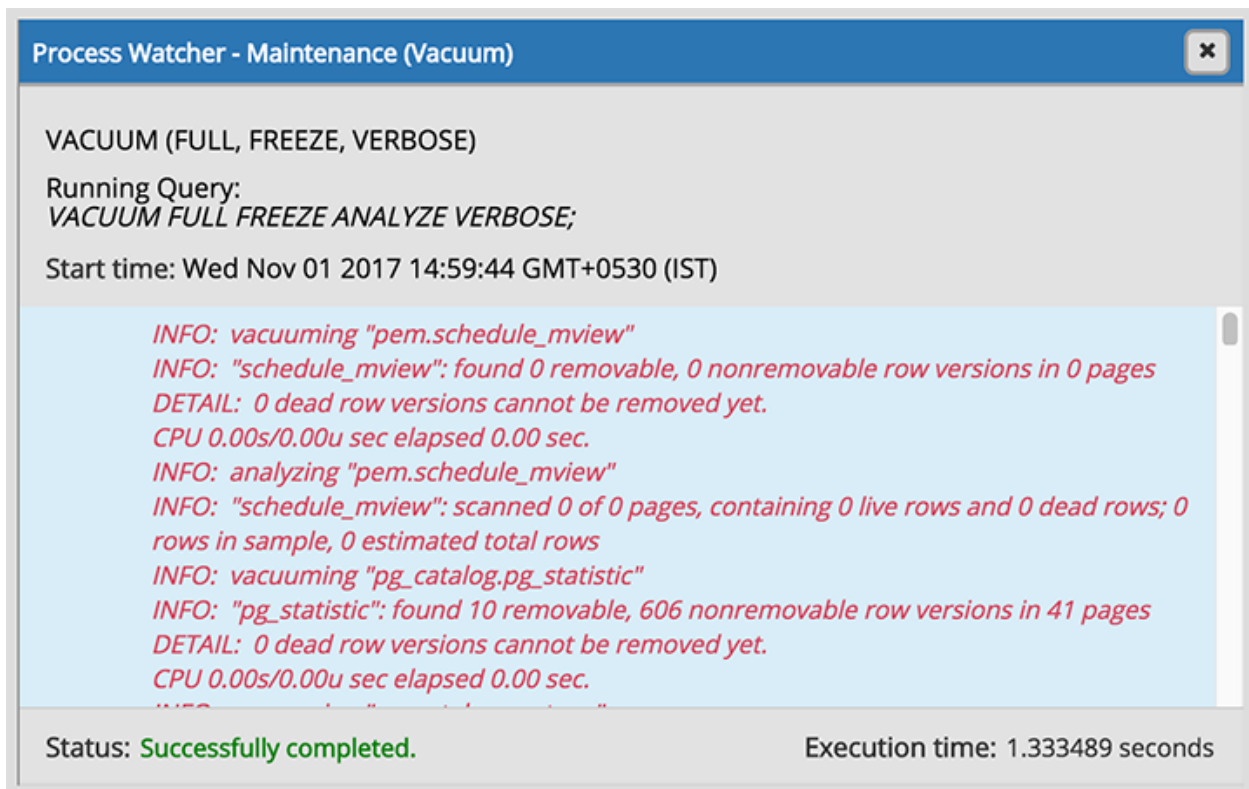
To exclude status messages from the process output, move the *Verbose Messages* switch to the *No* position; by default, status messages are included.

When you've completed the dialog, click *OK* to start the background process; to exit the dialog without performing maintenance operations, click *Cancel*.

pgAdmin will inform you when the background process completes:



Use the *Click here for details* link on the notification to open the *Process Watcher* and review detailed information about the execution of the command that performed the import or export:



BACKUP AND RESTORE

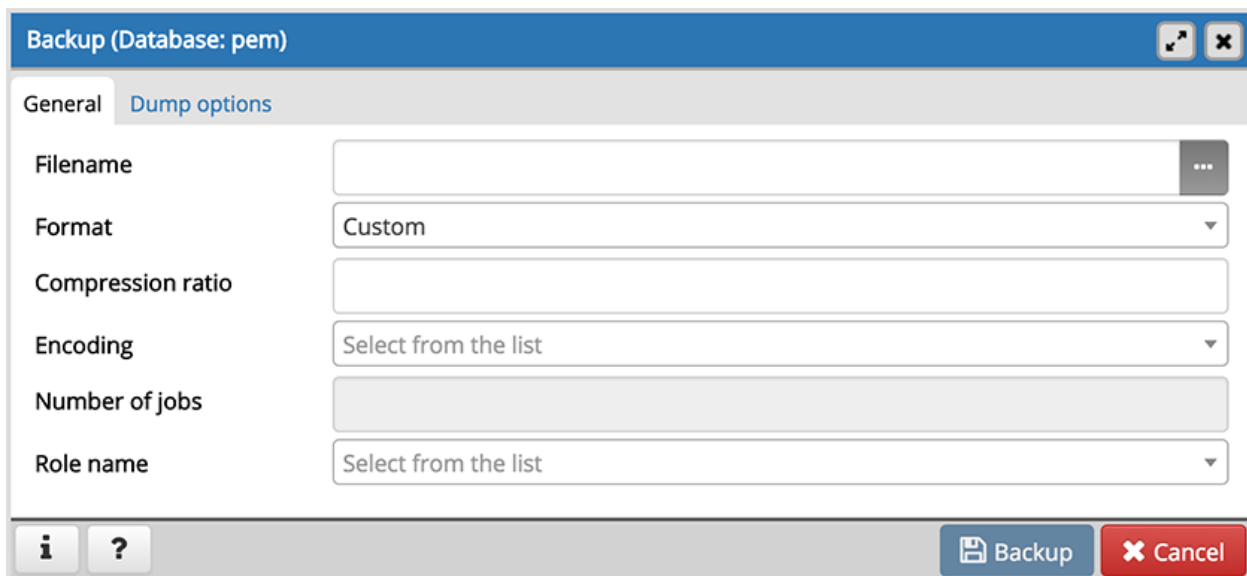
A powerful, but user-friendly Backup and Restore tool provides an easy way to use `pg_dump`, `pg_dumpall`, and `pg_restore` to take backups and create copies of databases or database objects for use in a development environment.

Contents:

6.1 Backup Dialog

Using the `pg_dump` utility, *pgAdmin* provides an easy way to create a backup in a plain-text or archived format. You can then use a client application (like `psql` or the *Query Tool*) to restore a plain-text backup file, or use the Postgres `pg_restore` utility to restore an archived backup. The `pg_dump` utility must have read access to all database objects that you want to back up.

You can backup a single table, a schema, or a complete database. Select the name of the backup source in the *pgAdmin* tree control, right click to open the context menu, and select *Backup...* to open the *Backup* dialog. The name of the object selected will appear in the dialog title bar.

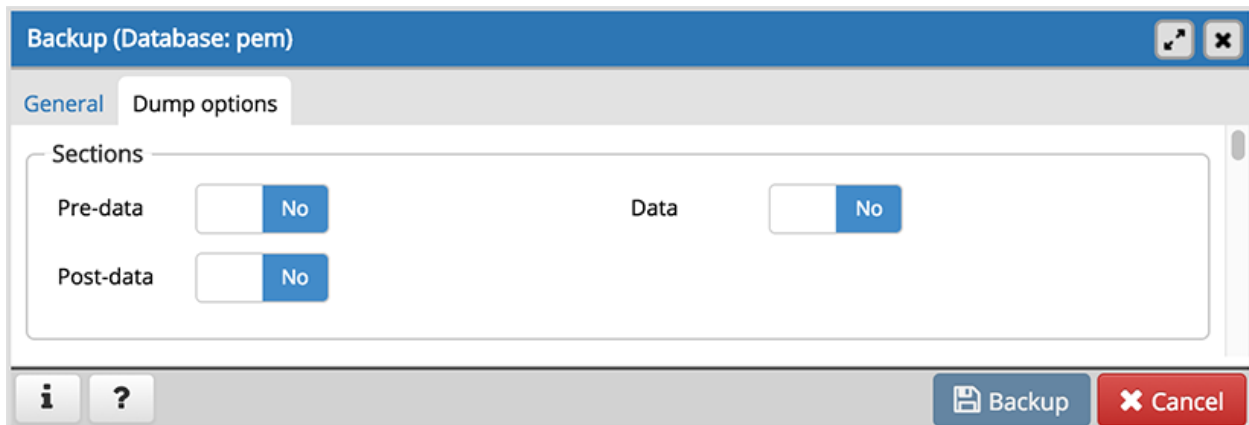


Use the fields in the *General* tab to specify parameters for the backup:

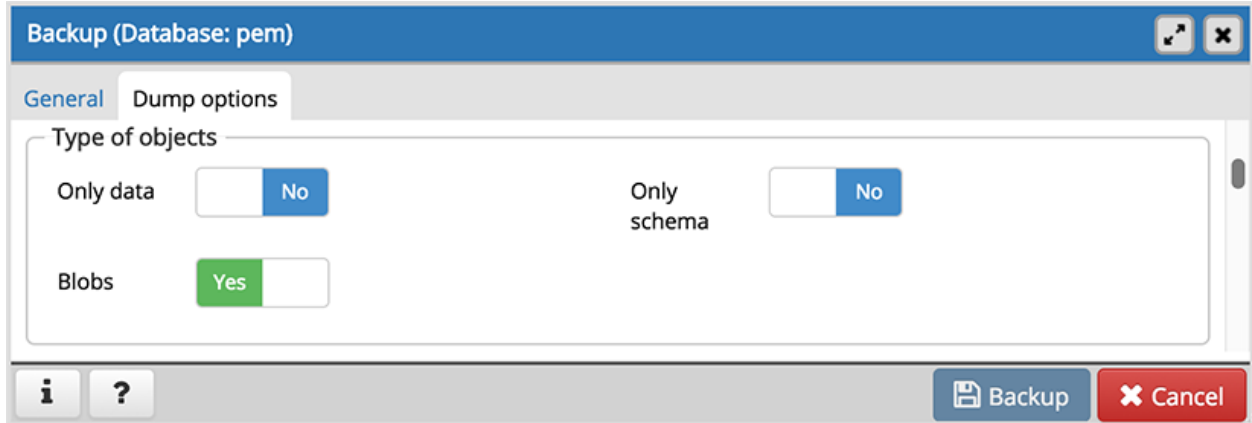
- Enter the name of the backup file in the *Filename* field. Optionally, select the *Browser* icon (...) to the right to navigate into a directory and select a file that will contain the archive.
- Use the drop-down listbox in the *Format* field to select the format that is best suited for your application. Each format has advantages and disadvantages:

- Select *Custom* to create a custom archive file that you can use with *pg_restore* to create a copy of a database. Custom archive file formats must be restored with *pg_restore*. This format offers the opportunity to select which database objects to restore from the backup file. *Custom* archive format is recommended for medium to large databases as it is compressed by default.
 - Select *Tar* to generate a tar archive file that you can restore with *pg_restore*. The tar format does not support compression.
 - Select *Plain* to create a plain-text script file. A plain-text script file contains SQL statements and commands that you can execute at the *psql* command line to recreate the database objects and load the table data. A plain-text backup file can be edited in a text editor, if desired, before using the *psql* program to restore database objects. *Plain* format is normally recommended for smaller databases; script dumps are not recommended for blobs. The SQL commands within the script will reconstruct the database to the last saved state of the database. A plain-text script can be used to reconstruct the database on another machine, or (with modifications) on other architectures.
 - Select *Directory* to generate a directory-format archive suitable for use with *pg_restore*. This file format creates a directory with one file for each table and blob being dumped, plus a *Table of Contents* file describing the dumped objects in a machine-readable format that *pg_restore* can read. This format is compressed by default.
- Use the *Compression Ratio* field to select a compression level for the backup. Specify a value of zero to mean use no compression; specify a maximum compression value of 9. Please note that tar archives do not support compression.
 - Use the *Encoding* drop-down listbox to select the character encoding method that should be used for the archive.
 - Use the *Number of Jobs* field (when applicable) to specify the number of tables that will be dumped simultaneously in a parallel backup.
 - Use the dropdown listbox next to *Rolename* to specify the role that owns the backup.

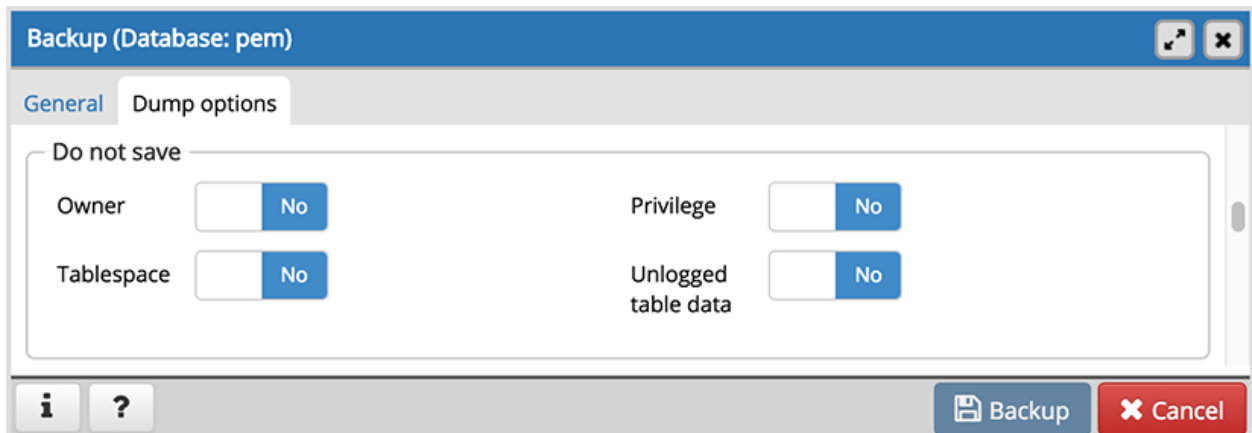
Click the *Dump options* tab to continue. Use the box fields in the *Dump options* tab to provide options for *pg_dump*.



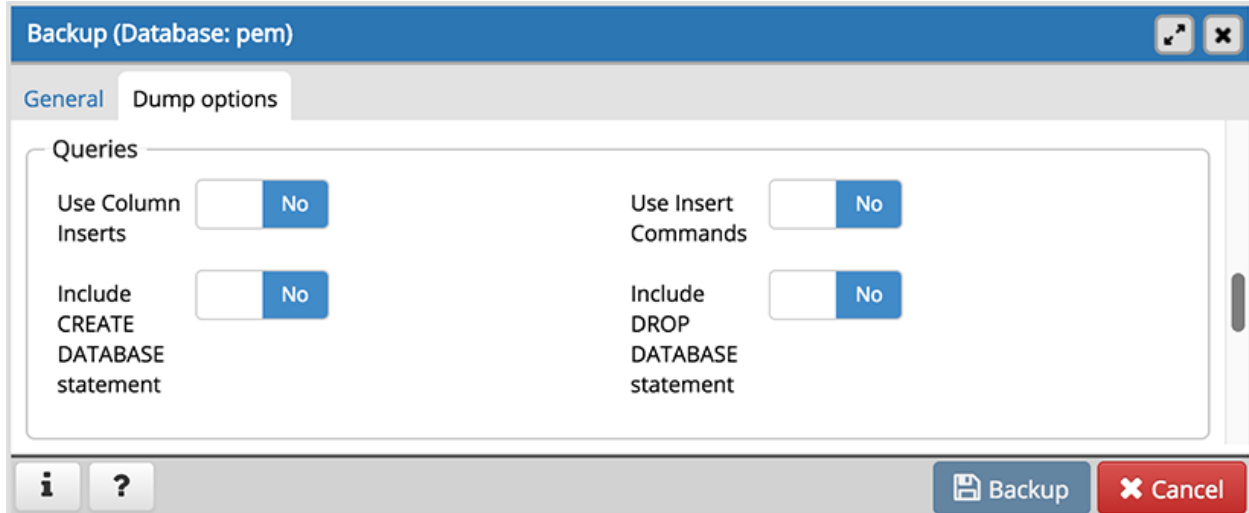
- Move switches in the **Sections** field box to select a portion of the object that will be backed up.
 - Move the switch next to *Pre-data* to the *Yes* position to include all data definition items not included in the data or post-data item lists.
 - Move the switch next to *Data* to the *Yes* position to backup actual table data, large-object contents, and sequence values.
 - Move the switch next to *Post-data* to the *Yes* position to include definitions of indexes, triggers, rules, and constraints other than validated check constraints.



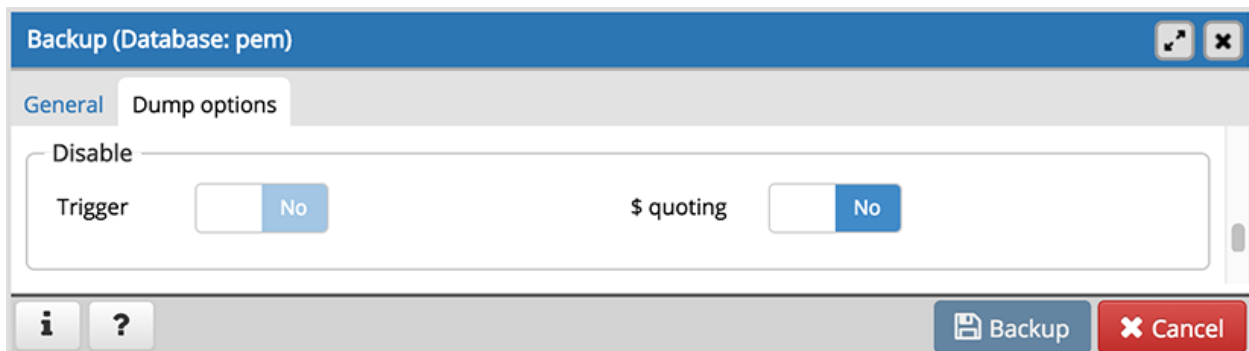
- Move switches in the **Type of objects** field box to specify details about the type of objects that will be backed up.
 - Move the switch next to *Only data* to the *Yes* position to limit the back up to data.
 - Move the switch next to *Only schema* to limit the back up to schema-level database objects.
 - Move the switch next to *Blobs* to the *No* position to exclude large objects in the backup.



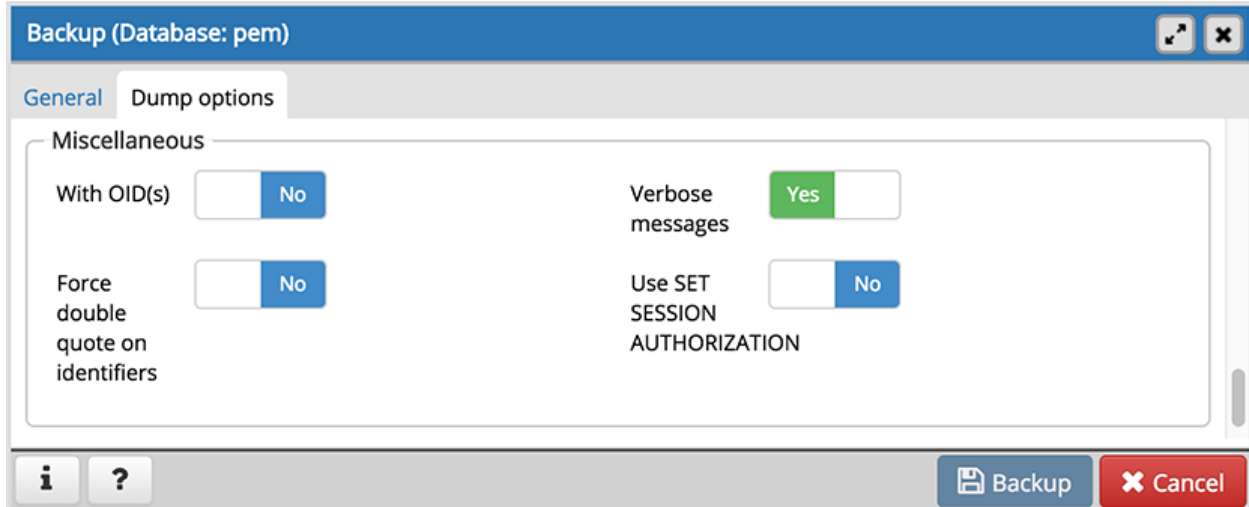
- Move switches in the **Do not save** field box to select the objects that will not be included in the backup.
 - Move the switch next to *Owner* to the *Yes* position to include commands that set object ownership.
 - Move the switch next to *Privilege* to the *Yes* position to include commands that create access privileges.
 - Move the switch next to *Tablespace* to the *Yes* position to include tablespaces.
 - Move the switch next to *Unlogged table data* to the *Yes* position to include the contents of unlogged tables.



- Move switches in the **Queries** field box to specify the type of statements that should be included in the backup.
 - Move the switch next to *Use Column Inserts* to the *Yes* position to dump the data in the form of INSERT statements and include explicit column names. Please note: this may make restoration from backup slow.
 - Move the switch next to *Use Insert commands* to the *Yes* position to dump the data in the form of INSERT statements rather than using a COPY command. Please note: this may make restoration from backup slow.
 - Move the switch next to *Include CREATE DATABASE statement* to the *Yes* position to include a command in the backup that creates a new database when restoring the backup.
 - Move the switch next to *Include DROP DATABASE statement* to the *Yes* position to include a command in the backup that will drop any existing database object with the same name before recreating the object during a backup.




- Move switches in the **Disable** field box to specify the type of statements that should be excluded from the backup.
 - Move the switch next to *Trigger* (active when creating a data-only backup) to the *Yes* position to include commands that will disable triggers on the target table while the data is being loaded.
 - Move the switch next to *\$ quoting* to the *Yes* position to enable dollar quoting within function bodies; if disabled, the function body will be quoted using SQL standard string syntax.



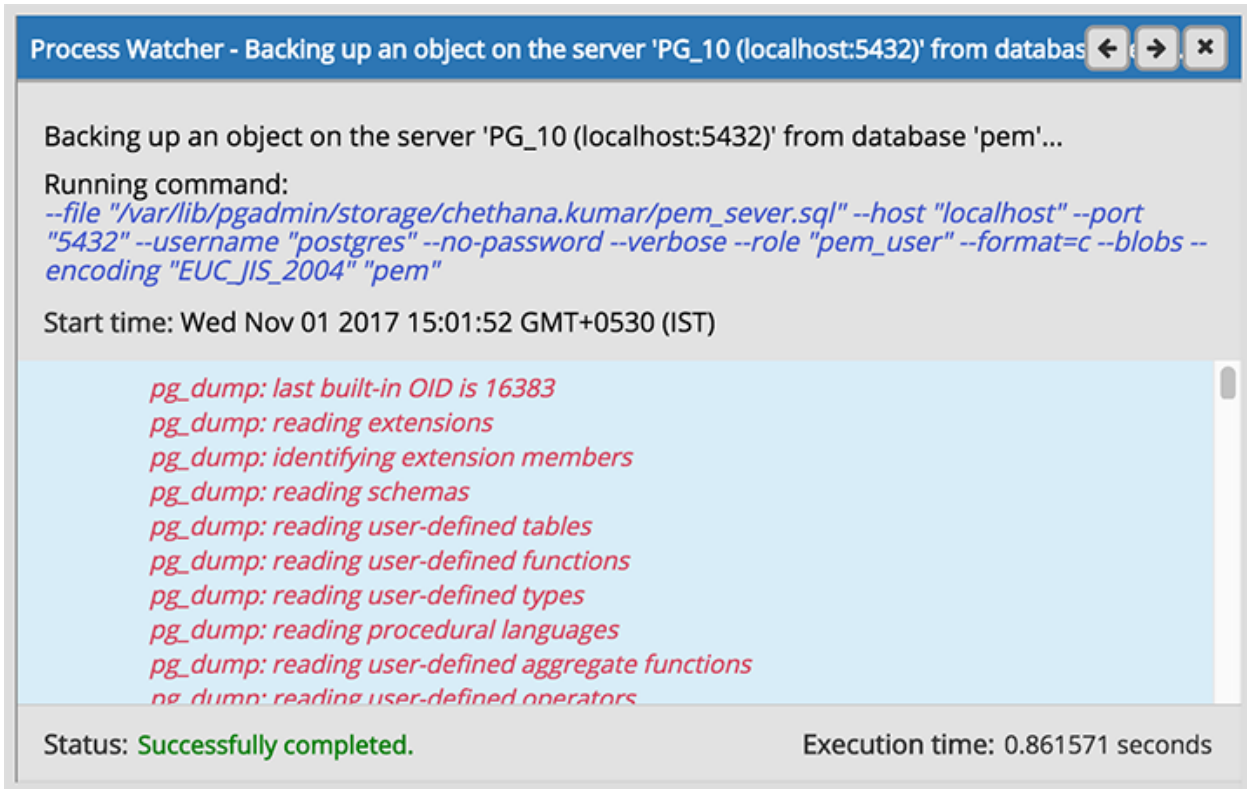
- Move switches in the **Miscellaneous** field box to specify miscellaneous backup options.
 - Move the switch next to *With OIDs* to the *Yes* position to include object identifiers as part of the table data for each table.
 - Move the switch next to *Verbose messages* to the *No* position to instruct *pg_dump* to exclude verbose messages.
 - Move the switch next to *Force double quotes on identifiers* to the *Yes* position to force the quoting of all identifiers.
 - Move the switch next to *Use SET SESSION AUTHORIZATION* to the *Yes* position to include a statement that will use a `SET SESSION AUTHORIZATION` command to determine object ownership (instead of an `ALTER OWNER` command).

When you've specified the details that will be incorporated into the `pg_dump` command:

- Click the *Backup* button to build and execute a command that builds a backup based on your selections on the *Backup* dialog.
- Click the *Cancel* button to exit without saving work.




If the backup is successful, a popup window will confirm success. Click *Click here for details* on the popup window to launch the *Process Watcher*. The *Process Watcher* logs all the activity associated with the backup and provides additional information for troubleshooting.



If the backup is unsuccessful, you can review the error messages returned by the backup command on the *Process Watcher*.

6.2 Backup Globals Dialog

Use the *Backup Globals* dialog to create a plain-text script that recreates all of the database objects within a cluster, and the global objects that are shared by those databases. Global objects include tablespaces, roles, and object properties. You can use the pgAdmin *Query Tool* to play back a plain-text script, and recreate the objects in the backup.



Use the fields in the *General* tab to specify the following:

- Enter the name of the backup file in the *Filename* field. Optionally, select the *Browser* icon (ellipsis) to the right to navigate into a directory and select a file that will contain the archive.
- Use the drop-down listbox next to *Role name* to specify a role with connection privileges on the selected server. The role will be used for authentication during the backup.

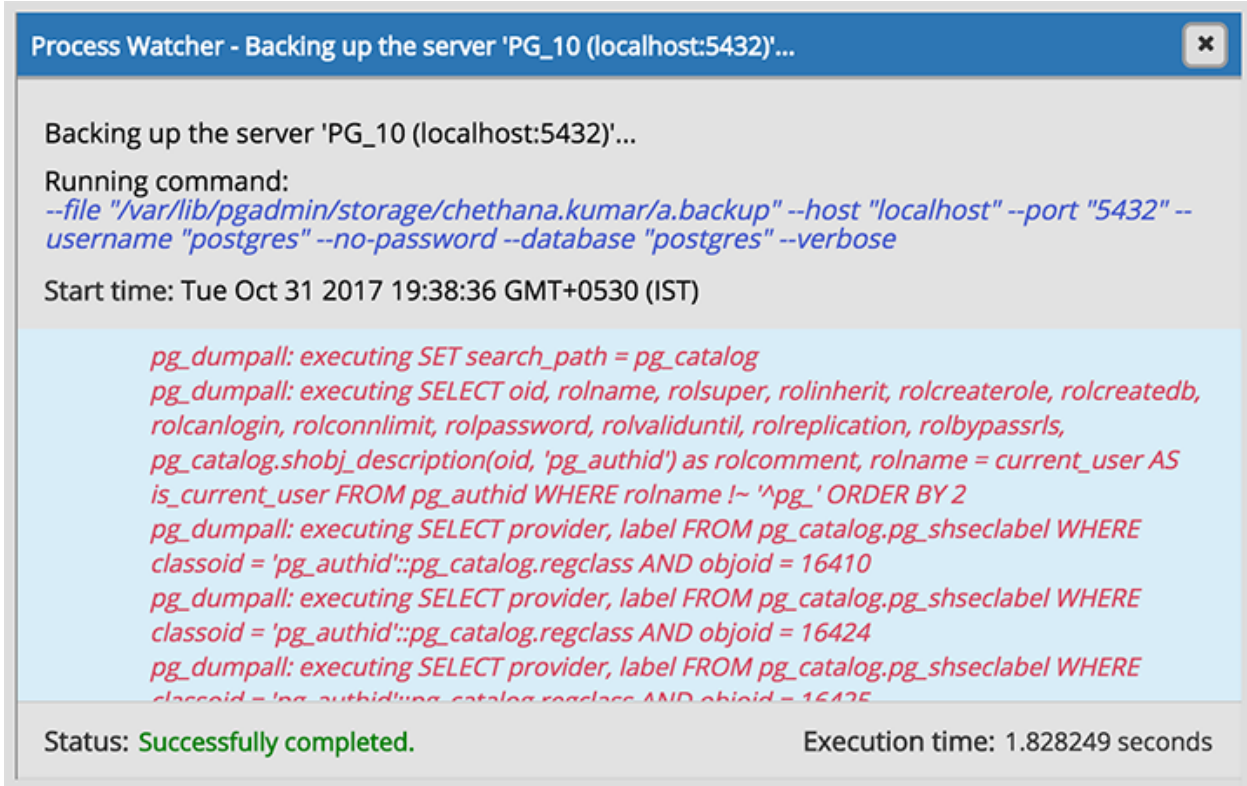
Move switches in the **Miscellaneous** field box to specify the type of statements that should be included in the backup.

- Move the *Verbose messages* switch to the *No* position to exclude status messages from the backup. The default is *Yes*.
- Move the *Force double quote on identifiers* switch to the *Yes* position to name identifiers without changing case. The default is *No*.

Click the *Backup* button to build and execute a command based on your selections; click the *Cancel* button to exit without saving work.



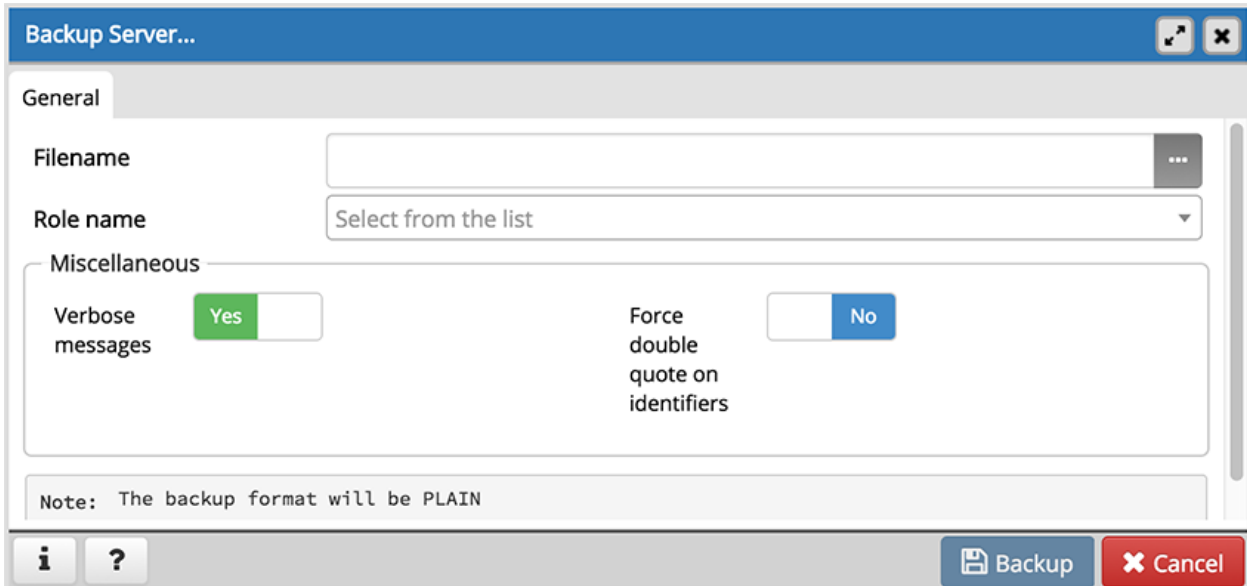
If the backup is successful, a popup window will confirm success. Click *Click here for details* on the popup window to launch the *Process Watcher*. The *Process Watcher* logs all the activity associated with the backup and provides additional information for troubleshooting.



If the backup is unsuccessful, review the error message returned by the *Process Watcher* to resolve any issue.

6.3 Backup Server Dialog

Use the *Backup Server* dialog to create a plain-text script that will recreate the selected server. You can use the pgAdmin *Query Tool* to play back a plain-text script, and recreate the server.




Use the fields in the *General* tab to specify the following:

- Enter the name of the backup file in the *Filename* field. Optionally, select the *Browser* icon (ellipsis) to the right to navigate into a directory and select a file that will contain the archive.
- Use the drop-down listbox next to *Role name* to specify a role with connection privileges on the selected server. The role will be used for authentication during the backup.

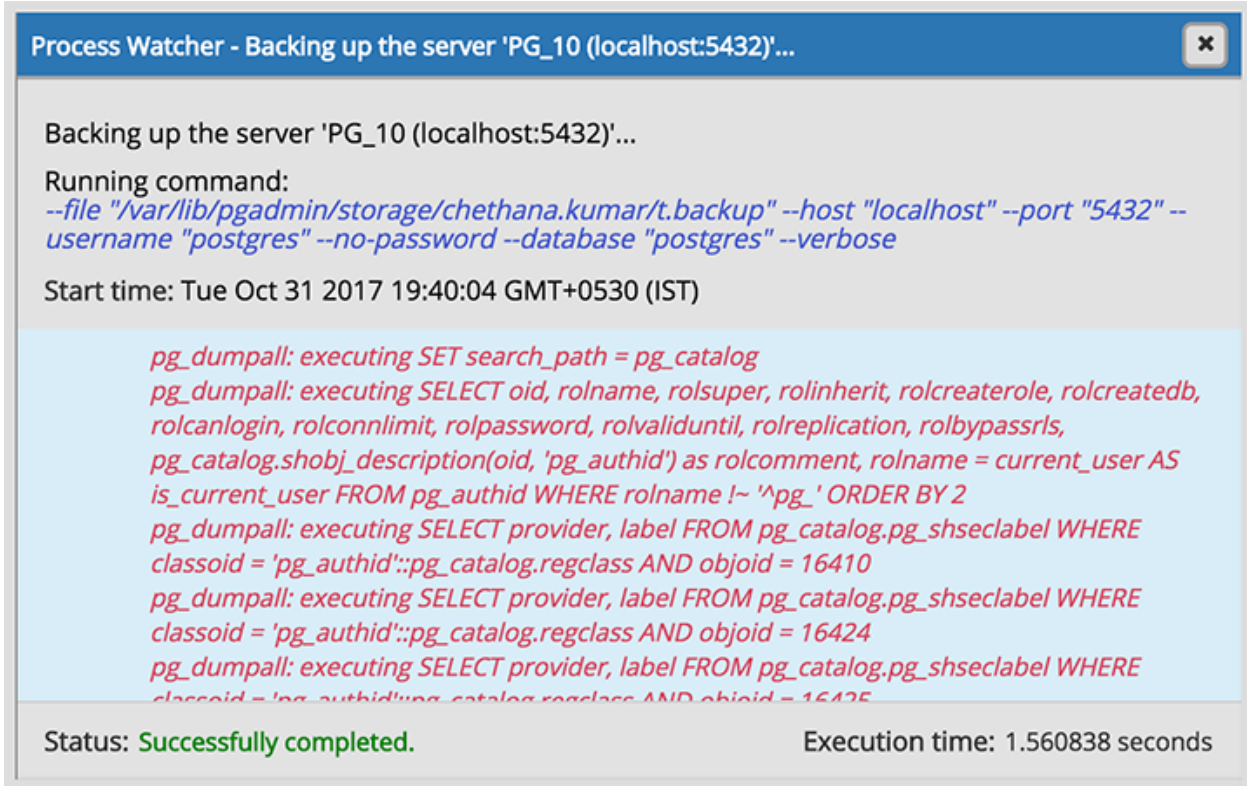
Move switches in the *Miscellaneous* box to specify the type of statements that should be included in the backup.

- Move the *Verbose messages* switch to the *No* position to exclude status messages from the backup. The default is *Yes*.
- Move the *Force double quote on identifiers* switch to the *Yes* position to name identifiers without changing case. The default is *No*.

Click the *Backup* button to build and execute a command based on your selections; click the *Cancel* button to exit without saving work.



If the backup is successful, a popup window will confirm success. Click *Click here for details* on the popup window to launch the *Process Watcher*. The *Process Watcher* logs all the activity associated with the backup and provides additional information for troubleshooting.

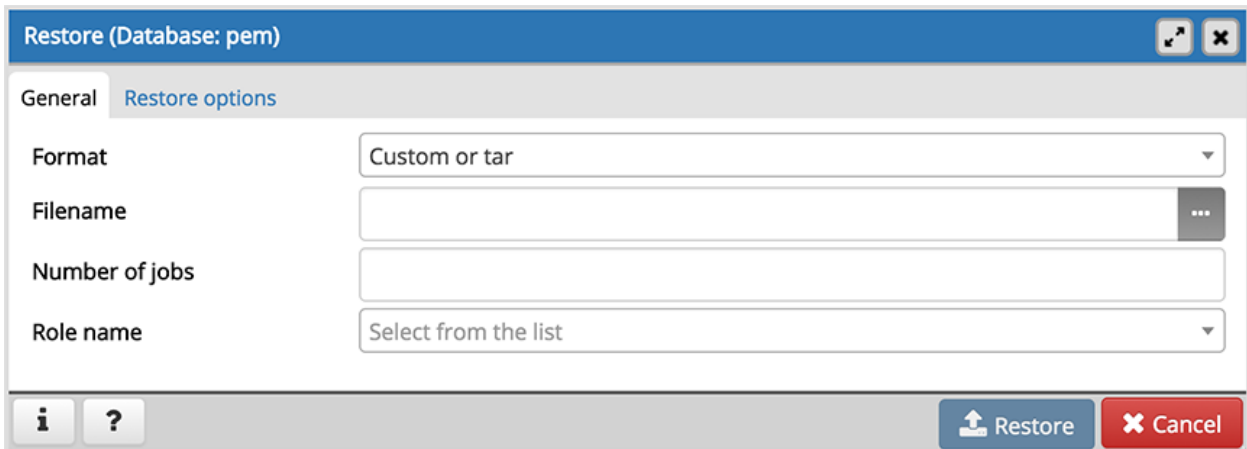


If the backup is unsuccessful, review the error message returned by the *Process Watcher* to resolve any issue.

6.4 Restore Dialog

The *Restore* dialog provides an easy way to use a Custom, tar, or Directory format backup taken with the *pgAdmin Backup* dialog to recreate a database or database object. The *Backup* dialog invokes options of the *pg_dump* client utility; the *Restore* dialog invokes options of the *pg_restore* client utility.

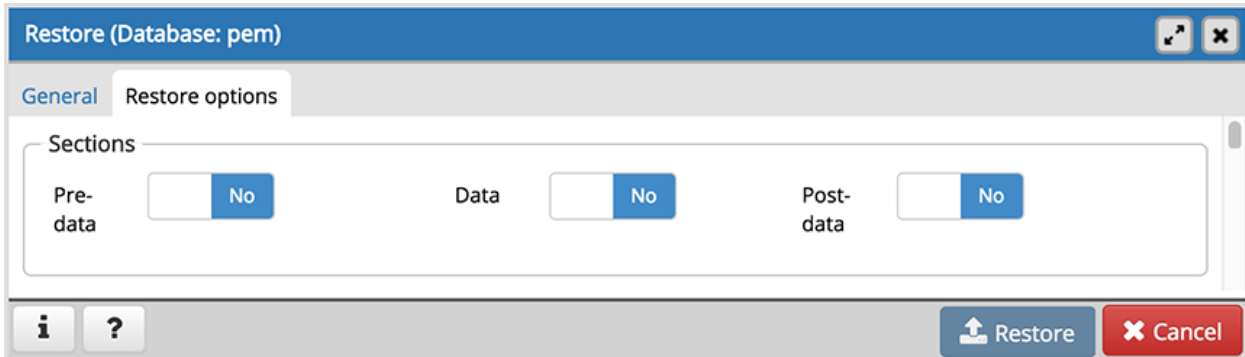
You can use the *Query Tool* to play back the script created during a plain-text backup made with the *Backup* dialog. For more information about backing up or restoring, please refer to the documentation for *pg_dump* or *pg_restore*.



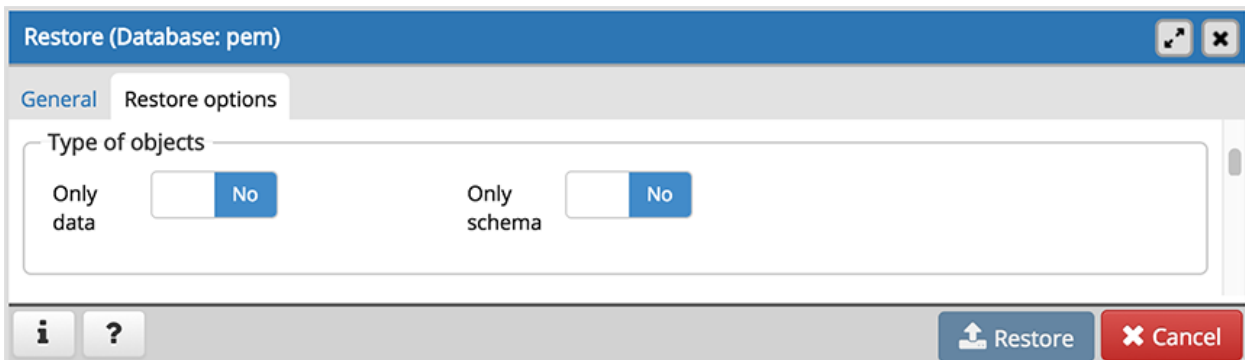
Use the fields on the *General* tab to specify general information about the restore process:

- Use the drop-down listbox in the *Format* field to select the format of your backup file.
 - Select *Custom or tar* to restore from a custom archive file to create a copy of the backed-up object.
 - Select *Directory* to restore from a compressed directory-format archive.
- Enter the complete path to the backup file in the *Filename* field. Optionally, select the *Browser* icon (ellipsis) to the right to navigate into a directory and select the file that contains the archive.
- Use the *Number of Jobs* field to specify if `pg_restore` should use multiple (concurrent) jobs to process the restore. Each job uses a separate connection to the server.
- Use the drop-down listbox next to *Rolename* to specify the role that will be used to authenticate with the server during the restore process.

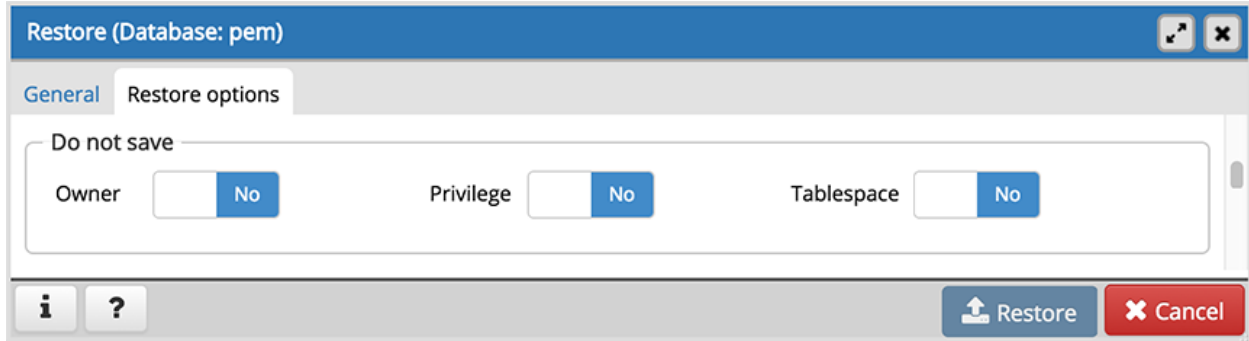
Click the *Restore options* tab to continue. Use the fields on the *Restore options* tab to specify options that correspond to `pg_restore` options.



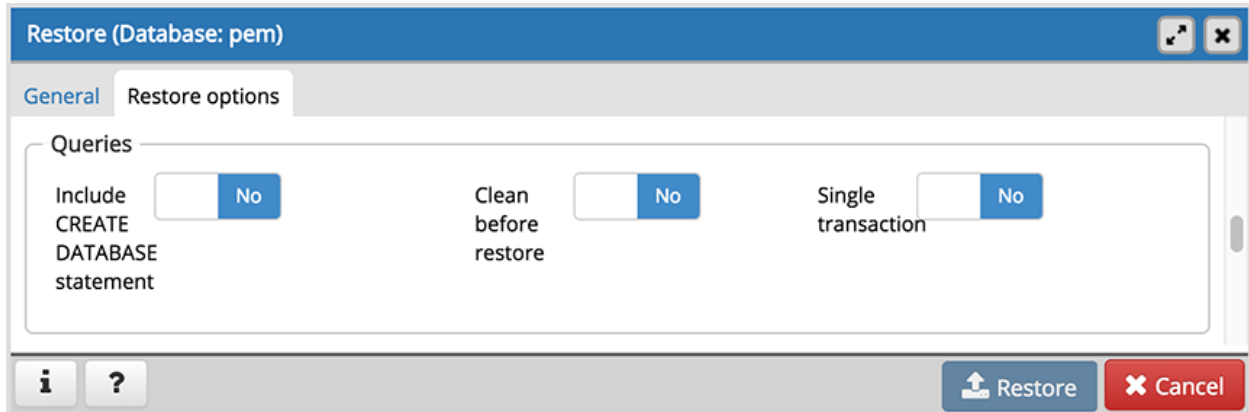
- Use the switches in the **Sections** box to specify the content that will be restored:
 - Move the switch next to *Pre-data* to the *Yes* position to restore all data definition items not included in the data or post-data item lists.
 - Move the switch next to *Data* to the *Yes* position to restore actual table data, large-object contents, and sequence values.
 - Move the switch next to *Post-data* to the *Yes* position to restore definitions of indexes, triggers, rules, and constraints (other than validated check constraints).



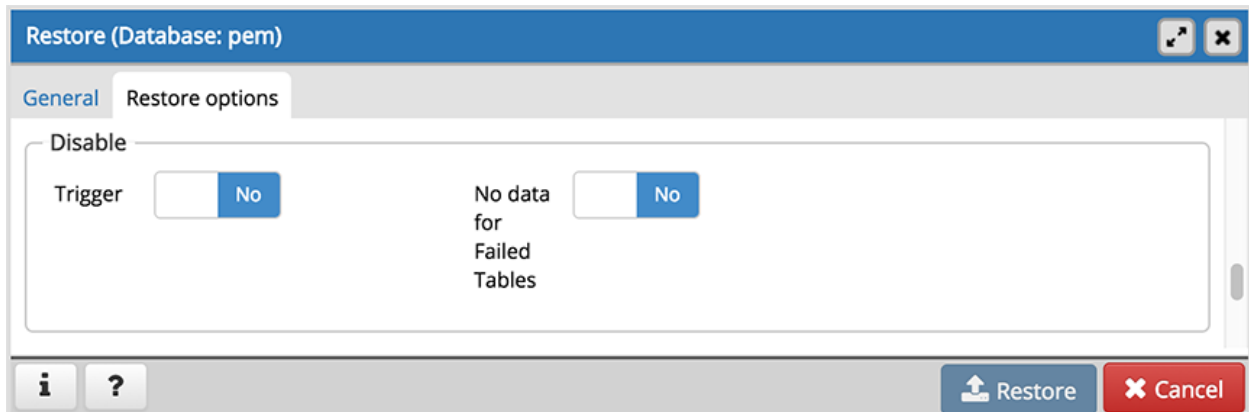
- Use the switches in the **Type of objects** box to specify the objects that will be restored:
 - Move the switch next to *Only data* to the *Yes* position to limit the restoration to data.
 - Move the switch next to *Only schema* to limit the restoration to schema-level database objects.



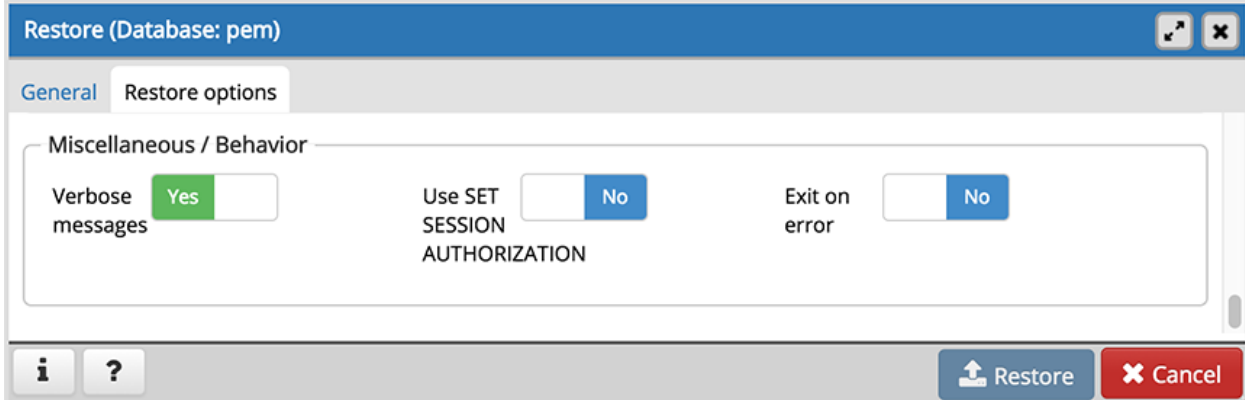
- Use the switches in the **Do not save** box to specify which objects will not be restored:
 - Move the switch next to *Owner* to the *Yes* position to exclude commands that set object ownership.
 - Move the switch next to *Privilege* to the *Yes* position to exclude commands that create access privileges.
 - Move the switch next to *Tablespace* to the *Yes* position to exclude tablespaces.



- Use the switches in the **Queries** box to specify the type of statements that should be included in the restore:
 - Move the switch next to *Include CREATE DATABASE statement* to the *Yes* position to include a command that creates a new database before performing the restore.
 - Move the switch next to *Clean before restore* to the *Yes* position to drop each existing database object (and data) before restoring.
 - Move the switch next to *Single transaction* to the *Yes* position to execute the restore as a single transaction (that is, wrap the emitted commands in *BEGIN/COMMIT*). This ensures that either all the commands complete successfully, or no changes are applied. This option implies *-exit-on-error*.




- Use the switches in the **Disable** box to specify the type of statements that should be excluded from the restore:
 - Move the switch next to *Trigger* (active when creating a data-only restore) to the *Yes* position to include commands that will disable triggers on the target table while the data is being loaded.
 - Move the switch next to *No data for Failed Tables* to the *Yes* position to ignore data that fails a trigger.



- Use the switches in the **Miscellaneous/Behavior** box to specify miscellaneous restore options:
 - Move the switch next to *Verbose messages* to the *No* position to instruct *pg_restore* to exclude verbose messages.
 - Move the switch next to *Use SET SESSION AUTHORIZATION* to the *Yes* position to include a statement that will use a `SET SESSION AUTHORIZATION` command to determine object ownership (instead of an `ALTER OWNER` command).
 - Move the switch next to *Exit on error* to the *Yes* position to instruct *pg_restore* to exit restore if there is an error in sending SQL commands. The default is to continue and to display a count of errors at the end of the restore.

When you've specified the details that will be incorporated into the `pg_restore` command, click the *Restore* button to start the process, or click the *Cancel* button to exit without saving your work. A popup will confirm if the restore is successful.



Click *Click here for details* on the popup to launch the *Process Watcher*. The *Process Watcher* logs all the activity associated with the restore, and provides additional information for troubleshooting should the restore command encounter problems.

Process Watcher - Copying table data 'public.t' on database 'postgres' and server (localhost:5432)... ✕

Copying table data 'public.t' on database 'postgres' for the server 'PG_10 (localhost:5432)'...

Running command:
`"/Library/PostgreSQL/9.6/bin/psql" --command " "\copy public.t (c) TO '<STORAGE_DIR>/t.csv' OIDS DELIMITER '|' CSV HEADER ENCODING 'EUC_JIS_2004' QUOTE '\"' ESCAPE '";"`

Start time: Wed Nov 01 2017 15:05:51 GMT+0530 (IST)

ERROR: table "t" does not have OIDs


Status: **Successfully completed.** Execution time: 0.032996 seconds

DEVELOPER TOOLS

The pgAdmin *Tools* menu displays a list of powerful developer tools that you can use to execute and analyze complex SQL commands, manage data, and debug PL/SQL code.

Contents:

7.1 pgAdmin Debugger



```

1
2 DECLARE
3     v_empno      NUMERIC(4);
4     v_ename      VARCHAR(10);
5     v_job        VARCHAR(9);
6     v_mgr        NUMERIC(4);
7     v_hiredate   DATE;
8     v_sal        NUMERIC(7,2);
9     v_comm       NUMERIC(7,2);
10    v_deptno     NUMERIC(2);
11 BEGIN
12    v_empno := new_empno();
13    INSERT INTO emp VALUES (v_empno, p_ename, 'SALESMAN', 7698,
14        CURRENT_DATE, p_sal, p_comm, 30);
15    SELECT INTO
16        v_empno, v_ename, v_job, v_mgr, v_hiredate, v_sal, v_comm, v_deptno
17        empno, ename, job, mgr, hiredate, sal, comm, deptno
18        FROM emp WHERE empno = v_empno;
19    RAISE INFO 'Department : %', v_deptno;
20    RAISE INFO 'Employee No: %', v_empno;
21    RAISE INFO 'Name       : %', v_ename;
22    RAISE INFO 'Job        : %', v_job;
23    RAISE INFO 'Manager   : %', v_mgr;
24    RAISE INFO 'Hire Date  : %', v_hiredate;
25    RAISE INFO 'Salary    : %', v_sal;
26    RAISE INFO 'Commission : %', v_comm;
27    RETURN v_empno;
28 EXCEPTION
29     WHEN OTHERS THEN
30         RAISE INFO 'The following is SQLERRM: ';
31         RAISE INFO '%', SQLERRM;
32         RAISE INFO 'The following is SQLSTATE: ';
33         RAISE INFO '%', SQLSTATE;
34     RETURN -1;
35 END;
36

```

Name	Type	Value
p_ename	character varying	ALBERT
p_sal	numeric	20
p_comm	numeric	30

The debugger may be used to debug PL/pgSQL functions in PostgreSQL. The Debugger is available as an extension for your PostgreSQL installation. You must have superuser privileges to use the debugger.

Before using the debugger, you must modify the *postgresql.conf* file, adding the server-side debugger components to the the value of the *shared_preload_libraries* parameter:

```
shared_preload_libraries = '$libdir/other_libraries/plugin_debugger'
```


After modifying the *shared_preload_libraries* parameter, restart the server to apply the changes.

The debugger may be used for either in-context debugging or direct debugging of a target function. When you use the debugger for in-context debugging, you set a breakpoint at the first line of a program; when a session invokes the target, control is transferred to the debugger. When using direct debugging, the debugger prompts you for any parameters required by the target, and then allows you to step through the code.

In-context Debugging

To set a breakpoint at the first line of a program, right-click the name of the object you would like to debug, and select

Set *breakpoint* from the *Debugging* sub-menu. The debugger window will open, waiting for another session to invoke the program.



When another session invokes the target, the debugger will display the code, allowing you to add break points, or step through line-by-line. The other session is suspended until the debugging completes; then control is returned to the session.

```

1
2 DECLARE
3     v_empno      NUMERIC(4);
4     v_ename      VARCHAR(10);
5     v_job        VARCHAR(9);
6     v_mgr        NUMERIC(4);
7     v_hiredate   DATE;
8     v_sal        NUMERIC(7,2);
9     v_comm       NUMERIC(7,2);
10    v_deptno     NUMERIC(2);
11 BEGIN
12     v_empno := new_empno();
13     INSERT INTO emp VALUES (v_empno, p_ename, 'SALESMAN', 7698,
14         CURRENT_DATE, p_sal, p_comm, 30);
15     SELECT INTO
16         v_empno, v_ename, v_job, v_mgr, v_hiredate, v_sal, v_comm, v_deptno
17         empno, ename, job, mgr, hiredate, sal, comm, deptno
18         FROM emp WHERE empno = v_empno;
19     RAISE INFO 'Department : %', v_deptno;
20     RAISE INFO 'Employee No: %', v_empno;
21     RAISE INFO 'Name       : %', v_ename;
22     RAISE INFO 'Job        : %', v_job;
23     RAISE INFO 'Manager    : %', v_mgr;
24     RAISE INFO 'Hire Date  : %', v_hiredate;
25     RAISE INFO 'Salary     : %', v_sal;
26     RAISE INFO 'Commission : %', v_comm;
27     RETURN v_empno;
28 EXCEPTION
29     WHEN OTHERS THEN
30         RAISE INFO 'The following is SQLERRM: ';
31         RAISE INFO '%', SQLERRM;
32         RAISE INFO 'The following is SQLSTATE: ';
33         RAISE INFO '%', SQLSTATE;
34     RETURN -1;
35 END;

```

Parameters Local variables Messages Results Stack

Name	Value	Line No
new_empno()		5
hire_salesman(character varying,numeric,numeric)	p_ename=ALBERT, p_sal=20, p_comm=30	12

Direct Debugging

To use the debugger for direct debugging, right click on the name of the object that you wish to debug in the pgAdmin tree control and select *Debug* from the *Debugging* sub-menu. The debugger window will open, prompting you for any values required by the program:

Name	Type	Null?	Expression?	Value	Use Default?	Default value
p_deptno	numeric	<input type="checkbox"/>	<input type="checkbox"/>		<input type="checkbox"/>	<No default value>
p_empno	numeric	<input type="checkbox"/>	<input type="checkbox"/>		<input type="checkbox"/>	<No default value>
p_ename	character varying	<input type="checkbox"/>	<input type="checkbox"/>		<input type="checkbox"/>	<No default value>

Use the fields on the *Debugger* dialog to provide a value for each parameter:

- The *Name* field contains the formal parameter name.
- The *Type* field contains the parameter data type.
- Check the *Null?* checkbox to indicate that the parameter is a NULL value.
- Check the *Expression?* checkbox if the *Value* field contains an expression.
- Use the *Value* field to provide the parameter value that will be passed to the program. When entering parameter values, type the value into the appropriate cell on the grid, or, leave the cell empty to represent NULL, enter '' (two single quotes) to represent an empty string, or to enter a literal string consisting of just two single quotes, enter '. PostgreSQL 8.4 and above supports variadic function parameters. These may be entered as a comma-delimited list of values, quoted and/or cast as required.
- Check the *Use default?* checkbox to indicate that the program should use the value in the Default Value field.
- The *Default Value* field contains the default value of the parameter.

Provide values required by the program, and click the *Debug* button to start stepping through the program.

```


1
2 DECLARE
3     v_empno      NUMERIC(4);
4     v_ename      VARCHAR(10);
5     v_job        VARCHAR(9);
6     v_mgr        NUMERIC(4);
7     v_hiredate   DATE;
8     v_sal        NUMERIC(7,2);
9     v_comm       NUMERIC(7,2);
10    v_deptno     NUMERIC(2);
11 BEGIN
12    v_empno := new_empno();
13    INSERT INTO emp VALUES (v_empno, p_ename, 'SALESMAN', 7698,
14        CURRENT_DATE, p_sal, p_comm, 30);
15    SELECT INTO
16        v_empno, v_ename, v_job, v_mgr, v_hiredate, v_sal, v_comm, v_deptno
17        empno, ename, job, mgr, hiredate, sal, comm, deptno
18        FROM emp WHERE empno = v_empno;
19    RAISE INFO 'Department : %', v_deptno;
20    RAISE INFO 'Employee No: %', v_empno;
21    RAISE INFO 'Name       : %', v_ename;
22    RAISE INFO 'Job        : %', v_job;
23    RAISE INFO 'Manager    : %', v_mgr;
24    RAISE INFO 'Hire Date  : %', v_hiredate;
25    RAISE INFO 'Salary     : %', v_sal;
26    RAISE INFO 'Commission : %', v_comm;
27    RETURN v_empno;
28 EXCEPTION
29    WHEN OTHERS THEN
30        RAISE INFO 'The following is SQLERRM: ';
31        RAISE INFO '%', SQLERRM;
32        RAISE INFO 'The following is SQLSTATE: ';
33        RAISE INFO '%', SQLSTATE;
34    RETURN -1;
35 END;

```

Name	Value	Line No.
new_empno()		5
hire_salesman(character varying,numeric,numeric)	p_ename=ALBERT, p_sal=20, p_comm=30	12

Using the Debugger

The main debugger window consists of two panels and a context-sensitive toolbar. Use toolbar icons to manage breakpoints and step into or through code; hover over an icon for a tooltip that identifies the option associated with the icon. The toolbar options are:



Option	Action
<i>Step into</i>	Click the <i>Step into</i> icon to execute the currently highlighted line of code.
<i>Step over</i>	Click the <i>Step over</i> icon to execute a line of code, stepping over any sub-functions invoked by the code. The sub-function executes, but is not debugged unless it contains a breakpoint.
<i>Continue/Start</i>	Click the <i>Continue/Start</i> icon to execute the highlighted code, and continue until the program encounters a breakpoint or completes.
<i>Toggle breakpoint</i>	Use the <i>Toggle breakpoint</i> icon to enable or disable a breakpoint (without removing the breakpoint).
<i>Clear all breakpoints</i>	Click the <i>Clear all breakpoints</i> icon to remove all breakpoints from the program.
<i>Stop</i>	Click the <i>Stop</i> icon to halt the execution of a program.

The top panel of the debugger window displays the program body; click in the grey margin next to a line number to add a breakpoint. The highlighted line in the top panel is the line that is about to execute.

The screenshot shows the pgAdmin 4 debugger interface. The top panel displays a PL/SQL procedure named `new_empno()`. The code is as follows:

```

1
2 DECLARE
3     v_empno      NUMERIC(4);
4     v_ename      VARCHAR(10);
5     v_job        VARCHAR(9);
6     v_mgr        NUMERIC(4);
7     v_hiredate   DATE;
8     v_sal        NUMERIC(7,2);
9     v_comm       NUMERIC(7,2);
10    v_deptno     NUMERIC(2);
11 BEGIN
12    v_empno := new_empno();
13    INSERT INTO emp VALUES (v_empno, p_ename, 'SALESMAN', 7698,
14        CURRENT_DATE, p_sal, p_comm, 30);
15    SELECT INTO
16        v_empno, v_ename, v_job, v_mgr, v_hiredate, v_sal, v_comm, v_deptno
17        empno, ename, job, mgr, hiredate, sal, comm, deptno
18        FROM emp WHERE empno = v_empno;
19    RAISE INFO 'Department : %', v_deptno;
20    RAISE INFO 'Employee No: %', v_empno;
21    RAISE INFO 'Name      : %', v_ename;
22    RAISE INFO 'Job       : %', v_job;
23    RAISE INFO 'Manager   : %', v_mgr;
24    RAISE INFO 'Hire Date  : %', v_hiredate;
25    RAISE INFO 'Salary    : %', v_sal;
26    RAISE INFO 'Commission : %', v_comm;
27    RETURN v_empno;
28 EXCEPTION
29     WHEN OTHERS THEN
30         RAISE INFO 'The following is SQLERRM: ';
31         RAISE INFO '%', SQLERRM;
32         RAISE INFO 'The following is SQLSTATE: ';
33         RAISE INFO '%', SQLSTATE;
34     RETURN -1;
35 END;
36

```

The bottom panel shows the execution parameters for the current step:

Name	Type	Value
p_ename	character varying	ALBERT
p_sal	numeric	20
p_comm	numeric	30

The lower panel of the debugger window provides a set of tabs that allow you to review information about the program:

- The *Parameters* tab displays the value of each parameter.
- The *Local variables* tab displays the current value of the program variables.
- The *Messages* tab displays any messages returned by the server (errors, warnings and informational messages).
- The *Results* tab displays the server message when the program completes.
- The *Stack* tab displays the list of functions that have been invoked, but which have not yet completed.

As you step through a program, the *Local variables* tab displays the current value of each variable:

Parameters Local variables Messages Results Stack			
Name	Type	Value	
v_empno	numeric	NULL	
v_ename	character varying	NULL	
v_job	character varying	NULL	
v_mgr	numeric	NULL	
v_hiredate	timestamp without time zone	NULL	
v_sal	numeric	NULL	
v_comm	numeric	NULL	
v_deptno	numeric	NULL	
sqlstate	text	NULL	

When you step into a subroutine, the *Stack* tab displays the call stack, including the name of each caller, the parameter values for each caller (if any), and the line number within each caller:


Parameters Local variables Messages Results Stack			
Name	Value	Line No.	
new_empno()		5	
hire_salesman(character varying,numeric,numeric)	p_ename=ALBERT, p_sal=20, p_comm=30	12	

Select a caller to change focus to that stack frame and display the state of the caller in the upper panel.

When the program completes, the *Results* tab displays the message returned by the server. If the program encounters an error, the *Messages* tab displays details:

ParametersLocal variablesMessagesResultsStack

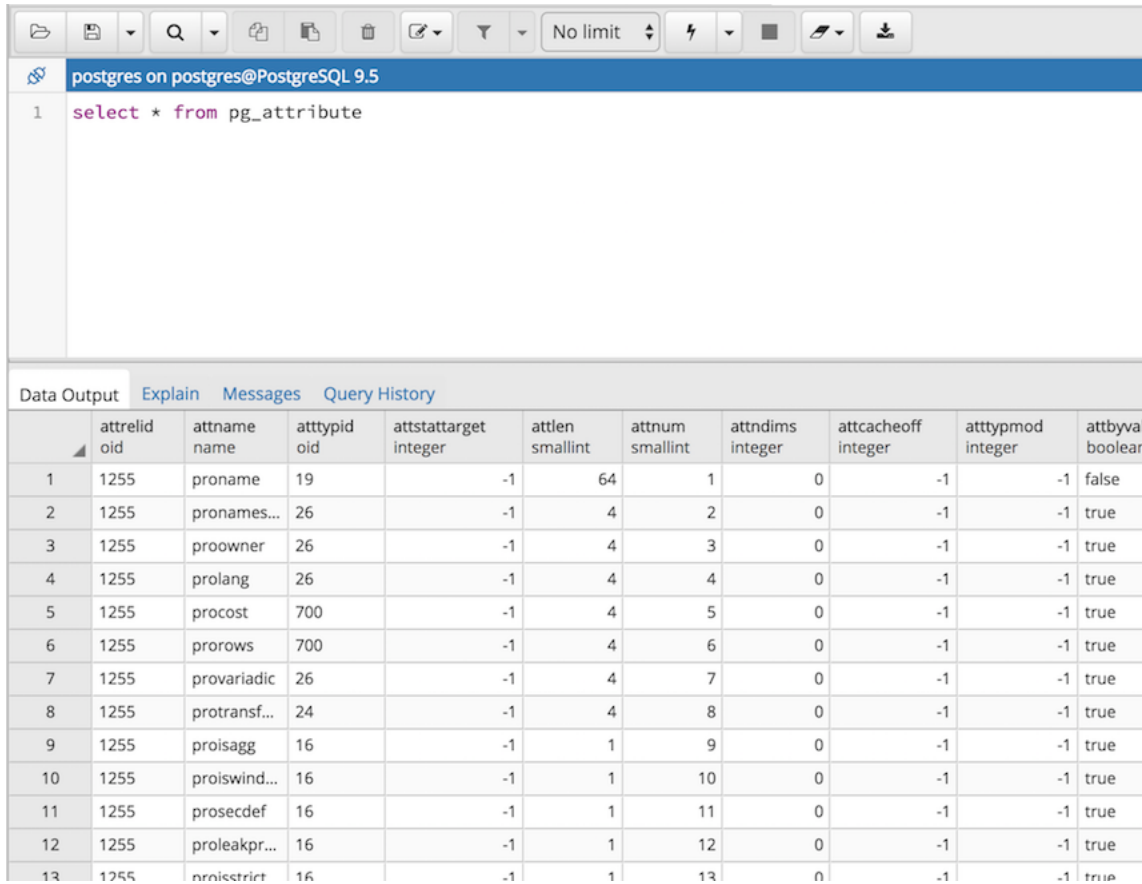
```
ERROR: column "deptno" does not exist
LINE 3:         WHERE deptno = p_deptno
                        ^
HINT:  Perhaps you meant to reference the column "emp.deptno".
QUERY: SELECT empno, ename, job, hiredate, sal
FROM emp WHERE deptno = p_deptno
        AND (empno = p_empno
            OR ename = UPPER(p_ename))
CONTEXT: edb-spl function emp_query(numeric,numeric,character varying) line 3 at SQL statement
edb-spl function emp_query_caller() line 12 at procedure/function invocation statement
```

 Execution completed with error

7.2 Query tool

The Query tool is a powerful, feature-rich environment that allows you to execute arbitrary SQL commands and review the result set. You can access the Query tool via the *Query Tool* menu option on the *Tools* menu, or through the context menu of select nodes of the Browser tree control. The Query Tool allows you to:

- Issue ad-hoc SQL queries.
- Execute arbitrary SQL commands.
- Displays current connection and transaction status as configured by the user.
- Save the data displayed in the output panel to a CSV file.
- Review the execution plan of a SQL statement in either a text or a graphical format.
- View analytical information about a SQL statement.



You can open multiple copies of the Query tool in individual tabs simultaneously. To close a copy of the Query tool, click the X in the upper-right hand corner of the tab bar.

The Query Tool features two panels:

- The upper panel displays the *SQL Editor*. You can use the panel to enter, edit, or execute a query.
- The lower panel displays the *Data Output* panel. The tabbed panel displays the result set returned by a query, information about a query’s execution plan, server messages related to the query’s execution, and a history of the queries invoked in the SQL Editor.

The Query Tool Toolbar

The *Query Tool* toolbar uses context-sensitive icons that provide shortcuts to frequently performed tasks. If an icon is highlighted, the option is enabled; if the icon is grayed-out, the task is disabled. Please note that disabled icons may support functionality accessed via the *data editor*.




Hover over an icon to display a tooltip that describes the icon’s functionality:


Icon	Behavior	Shortcut
<i>Open File</i>	Click the <i>Open File</i> icon to display a previously saved query in the SQL Editor.	
<i>Save</i>	Click the <i>Save</i> icon to perform a quick-save of a previously saved query, or to access the <i>Save</i> menu: <ul style="list-style-type: none"> • Select <i>Save</i> to save the selected content of the SQL Editor panel in a file. • Select <i>Save As</i> to open a new browser dialog and specify a new location to which to save the selected content of the SQL Editor panel. 	
<i>Find</i>	Use the <i>Find</i> menu to search, replace, or navigate the code displayed in the SQL Editor: <ul style="list-style-type: none"> • Select <i>Find</i> to provide a search target, and search the SQL Editor contents. • Select <i>Find next</i> to locate the next occurrence of the search target. • Select <i>Find previous</i> to move to the last occurrence of the search target. • Select <i>Persistent find</i> to identify all occurrences of the search target within the editor. • Select <i>Replace</i> to locate and replace (with prompting) individual occurrences of the target. • Select <i>Replace all</i> to locate and replace all occurrences of the target within the editor. • Select <i>Jump</i> to navigate to the next occurrence of the search target. 	Cmd+F Cmd+G Cmd+Shift+G Cmd+Shift+F Alt+G
<i>Copy</i>	Click the <i>Copy</i> icon to copy the content that is currently highlighted in the Data Output panel.	
<i>Edit</i>	Use options on the <i>Edit</i> menu to access text editing tools; the options operate on the text displayed in the SQL Editor panel: <ul style="list-style-type: none"> • Select <i>Indent Selection</i> to indent the currently selected text. • Select <i>Unindent Selection</i> to remove indentation from the currently selected text. 	Tab Shift+Tab Cmd+/ Cmd+. Shift+Cmd+/ Shift+Cmd+.
7.2. Query tool	<ul style="list-style-type: none"> • Select <i>Inline Comment Selection</i> to enclose any lines that contain the selection in SQL style comment notation. 	

The SQL Editor Panel

The *SQL editor* panel is a workspace where you can manually provide a query, copy a query from another source, or read a query from a file. The SQL editor features syntax coloring and autocompletion.



To use autocomplete, begin typing your query; when you would like the Query editor to suggest object names or commands that might be next in your query, press the Control+Space key combination. For example, type “*SELECT * FROM* ” (without quotes, but with a trailing space), and then press the Control+Space key combination to select from a popup menu of autocomplete options.



After entering a query, select the *Execute/Refresh* icon from the toolbar. The complete contents of the SQL editor panel will be sent to the database server for execution. To execute only a section of the code that is displayed in the SQL editor, highlight the text that you want the server to execute, and click the *Execute/Refresh* icon.

Interface showing a SQL query execution in pgAdmin 4. The query is:

```

1 SELECT GENERATE_SERIES(1,1000) AS ID, 'CHETHANA' AS NAME
2
3 SELECT GENERATE_SERIES(1,1000) AS ID, 'HARSHAL' AS NAME
4
5 SELECT GENERATE_SERIES(1,1000) AS ID, 'NEEL' AS NAME

```

The output is displayed in a table with columns 'id' (integer) and 'name' (unknown):

	id integer	name unknown
1	1	HARSHAL
2	2	HARSHAL
3	3	HARSHAL
4	4	HARSHAL
5	5	HARSHAL
6	6	HARSHAL
7	7	HARSHAL
8	8	HARSHAL
9	9	HARSHAL
10	10	HARSHAL
11	11	HARSHAL
12	12	HARSHAL
13	13	HARSHAL
14	14	HARSHAL

The message returned by the server when a command executes is displayed on the *Messages* tab. If the command is successful, the *Messages* tab displays execution details.

Interface showing a SQL query execution in pgAdmin 4. The query is:

```

1 SELECT * FROM pem.logexp_tagbasecharts
2 ORDER BY id ASC

```

The *Messages* tab displays the following execution details:

```

Successfully run. Total query runtime: 56 msec.
4 rows affected.

```

Options on the *Edit* menu offer functionality that helps with code formatting and commenting:

- The auto-indent feature will automatically indent text to the same depth as the previous line when you press the Return key.
- Block indent text by selecting two or more lines and pressing the Tab key.
- Implement or remove SQL style or toggle C style comment notation within your code.

The Data Output Panel

The *Data Output* panel displays data and statistics generated by the most recently executed query.

The screenshot shows the pgAdmin 4 interface. At the top, there is a toolbar with icons for file operations, search, and execution. Below the toolbar, a blue header bar indicates the user is logged in as 'pem on postgres@PG_10'. The main area contains a SQL editor with the query: `1 SELECT * FROM PG_DATABASE`. Below the editor, there are four tabs: 'Data Output', 'Explain', 'Messages', and 'Query History'. The 'Data Output' tab is active, displaying a table with the following data:

	datname name	datdba oid	encoding integer	datcollate name	datatype name	datistemplate boolean	datallowconn boolean	datconnlimit integer	datlastsysoid oid	dat xid
1	postgres	10	6	C	C	false	true	-1	12668	
2	template1	10	6	C	C	true	true	-1	12668	
3	template0	10	6	C	C	true	false	-1	12668	
4	pem	10	6	C	C	false	true	-1	12668	

The *Data Output* tab displays the result set of the query in a table format. You can:

- Select and copy from the displayed result set.
- Use the *Execute/Refresh* options to retrieve query execution information and set query execution options.
- Use the *Download as CSV* icon to download the content of the *Data Output* tab as a comma-delimited file.

All rowsets from previous queries or commands that are displayed in the *Data Output* panel will be discarded when you invoke another query; open another query tool browser tab to keep your previous results available.

Use the *Explain* tab to view a graphical representation of a query:

pem on postgres@PG_10

```

1 SELECT * FROM pg_catalog.pg_class t
2 LEFT JOIN pg_attribute a ON (t.oid = a.attrelid)
3 WHERE t.relkind = 't'

```

Data Output Explain Messages Query History

```

graph LR
    pg_class[pg_class] --> Hash[Hash]
    pg_attribute[pg_attribute] --> HashRightJoin[Hash Right Join]
    Hash --> HashRightJoin

```

To generate a graphical explain diagram, open the *Explain* tab, and select *Explain*, *Explain Analyze*, or one or more options from the *Explain options* menu on the *Execute/Refresh* drop-down. Please note that *EXPLAIN VERBOSE* cannot be displayed graphically. Hover over an icon on the *Explain* tab to review information about that item; a popup window will display information about the selected object:

```

1 SELECT * FROM pg_catalog.pg_class t
2 LEFT JOIN pg_attribute a ON (t.oid = a.attrelid)
3 WHERE t.relkind = 't'
    
```

Node Type	Hash
Actual Total Time	0.435
Peak Memory Usage	41
Shared Hit Blocks	33
Shared Read Blocks	0
Local Hit Blocks	0
Original Hash Batches	1
Local Dirtied Blocks	0
Temp Written Blocks	0
Plan Width	243
Actual Loops	1
Original Hash Buckets	1024
Hash Batches	1
Actual Startup Time	0.435
Temp Read Blocks	0
Local Read Blocks	0
Hash Buckets	1024
Startup Cost	47.46
Shared Dirtied Blocks	0
Shared Written Blocks	0
Local Written Blocks	0
Plan Rows	185

Note that the query plan that accompanies the *Explain analyze* is available on the *Data Output* tab.


Use the *Messages* tab to view information about the most recently executed query:

```

1 SELECT * FROM pem.logexp_tagbasecharts
2 ORDER BY id ASC
    
```

Successfully run. Total query runtime: 56 msec.
4 rows affected.

If the server returns an error, the error message will be displayed on the *Messages* tab, and the syntax that caused the error will be underlined in the SQL editor. If a query succeeds, the *Messages* tab displays how long the query took to complete and how many rows were retrieved:



Dashboard Properties SQL Statistics Dependencies Dependents Query - postgres...

PG_10 - postgres - pg_toast.pg_toast_1255

```


1 SELECT * FROM pg_toast.pg_toast_1255
2 LIMIT 100
3

```

Data Output Explain Messages Query History

Successfully run. Total query runtime: 75 msec.
0 rows affected.

Use the *Query History* tab to review activity for the current session:



pem on postgres@PG_10

```

1 SELECT * FROM pg_catalog.pg_class t
2 LEFT JOIN pg_attribute a ON (t.oid = a.attrelid)
3 WHERE t.relkind = 't'

```

Data Output Explain Messages Query History

Today - Oct 31 2017

Date	Rows Affected	Duration
10-31-17 18:59:50	1	27 msec
18:57:24		
18:57:06		
18:56:54		
18:56:16		
18:56:08		
18:53:52		
18:53:34		

Copy All

```

EXPLAIN (FORMAT JSON, ANALYZE ON,
WHERE t.relkind = 't'

```

Messages

Successfully run.
Total query runtime: 27 msec.
1 rows affected.

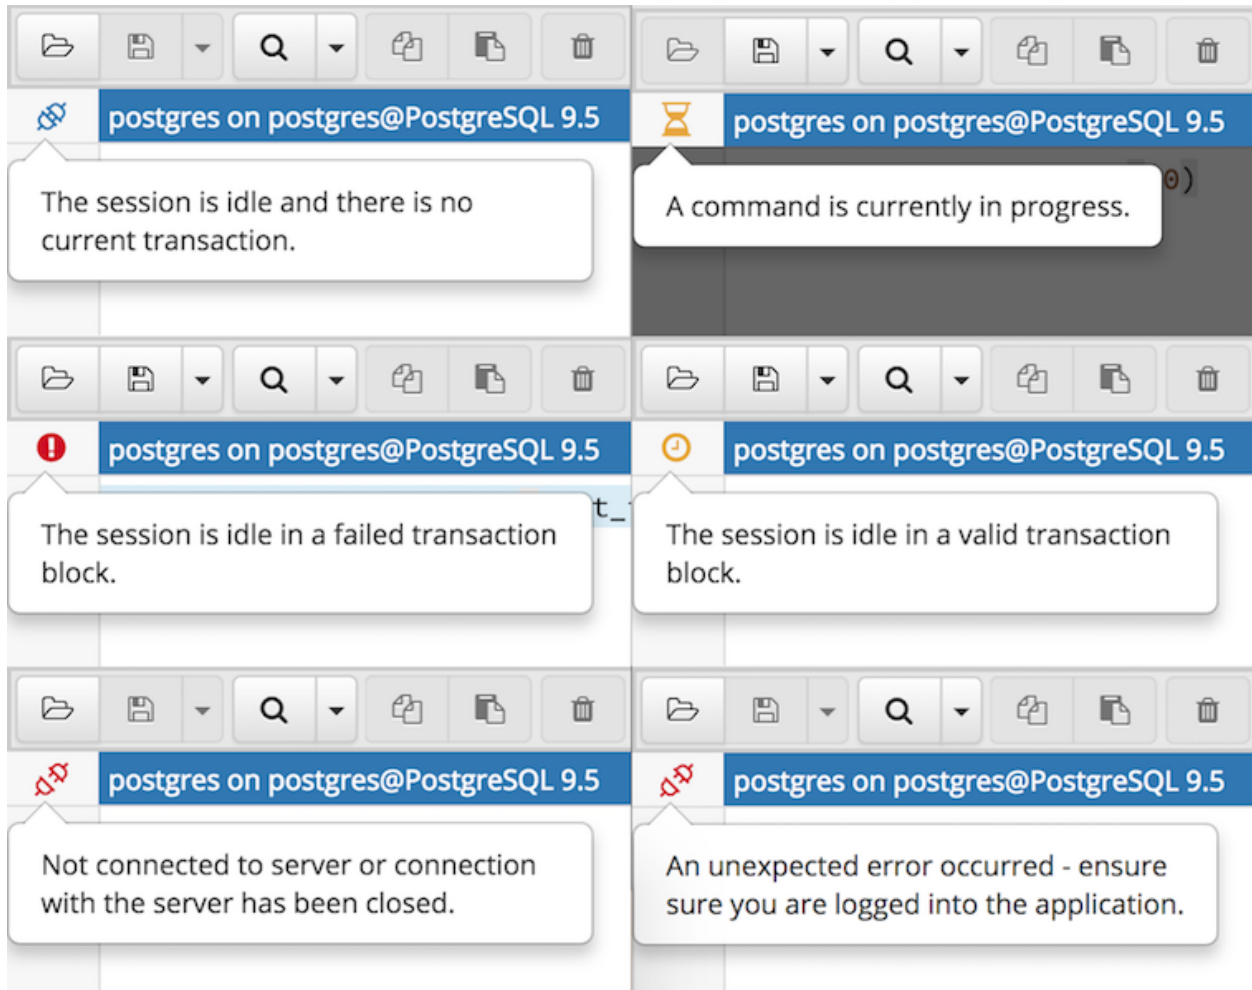
The Query History tab displays information about recent commands:

- The date and time that a query was invoked.
- The text of the query.

- The number of rows returned by the query.
- The amount of time it took the server to process the query and return a result set.
- Messages returned by the server (not noted on the *Messages* tab).

To erase the content of the *Query History* tab, select *Clear history* from the *Clear* drop-down menu.

Use the *Connection status* feature to view the current connection and transaction status by clicking on the status icon in query tool:



7.3 Reviewing and Editing Data

To review or modify data, right click on a table or view name in the *Browser* tree control. When the context menu opens, use the *View/Edit Data* menu to specify the number of rows you would like to display in the editor panel.

The screenshot shows the pgAdmin 4 interface. At the top, there is a toolbar with various icons for file operations, search, and execution. Below the toolbar, the title bar reads "EDB Postgres Advanced Server 10 - edb - public.emp". The main area contains a SQL query editor with the following text:

```

1 SELECT * FROM public.emp
2 ORDER BY empno ASC
    
```

Below the query editor, there are tabs for "Data Output", "Explain", "Messages", and "Query History". The "Data Output" tab is active, displaying a table with the following data:

	empno [PK] numeric (4)	ename character varying (10)	job character varying (9)	mgr numeric (4)	hiredate timestamp without time zone	sal numeric (7,2)	comm numeric (7,2)	deptno numeric (2)
1	7369	SMITH	SALESMAN	7902	1980-12-17 00:00:00	800.00	[null]	20
2	7499	ALLEN	SALESMAN	7698	1981-02-20 00:00:00	1600.00	300.00	30
3	7521	WARD	SALESMAN	7698	1981-02-22 00:00:00	1250.00	500.00	30
4	7566	JONES	MANAGER	7839	1981-04-02 00:00:00	2975.00	[null]	20
5	7654	MARTIN	SALESMAN	7698	1981-09-28 00:00:00	1250.00	1400.00	30
6	7698	BLAKE	MANAGER	7839	1981-05-01 00:00:00	2850.00	[null]	30
7	7782	CLARK	MANAGER	7839	1981-06-09 00:00:00	2450.00	[null]	10
8	7788	SCOTT	ANALYST	7566	1987-04-19 00:00:00	3000.00	[null]	20
9	7839	KING	PRESIDENT	[null]	1981-11-17 00:00:00	5000.00	[null]	10
10	7844	TURNER	SALESMAN	7698	1981-09-08 00:00:00	1500.00	0.00	30
11	7876	ADAMS	CLERK	7788	1987-05-23 00:00:00	1100.00	[null]	20
12	7900	JAMES	CLERK	7698	1981-12-03 00:00:00	950.00	[null]	30
13	7902	FORD	ANALYST	7566	1981-12-03 00:00:00	3000.00	[null]	20
14	7934	MILLER	CLERK	7782	1982-01-23 00:00:00	1300.00	[null]	10

To modify the content of a table, each row in the table must be uniquely identifiable. If the table definition does not include an OID or a primary key, the displayed data is read only. Note that views cannot be edited; updatable views (using rules) are not supported.

The editor features a toolbar that allows quick access to frequently used options, and a work environment divided into two panels:

- The upper panel displays the SQL command that was used to select the content displayed in the lower panel.
- The lower panel (the Data Grid) displays the data selected from the table or view.

The View/Edit Data Toolbar

The toolbar includes context-sensitive icons that provide shortcuts to frequently performed tasks.



Hover over an icon to display a tooltip that describes the icon's functionality.

Icon	Behavior	Shortcut
<i>Save</i>	Use the <i>Save</i> icon to save your changes to the currently displayed table contents.	
<i>Find</i>	Use options on the <i>Find</i> menu to access Search and Replace functionality or to Jump to another line.	Ctrl/Cmd +F
<i>Copy</i>	Click the <i>Copy</i> icon to copy the currently selected data.	Ctrl+C
<i>Paste Row</i>	Click the <i>Paste Row</i> icon to paste the content that is currently on the clipboard.	
<i>Delete Row</i>	Use the <i>Delete Row</i> icon to add a new row in the output panel.	
<i>Filter</i>	<p>Click the <i>Filter</i> icon to open a dialog that allows you to write and apply a filter for the content currently displayed in the output panel. Click the down arrow to open the <i>Filter</i> drop-down menu and select from pre-defined options:</p> <p>Use options on the <i>Filter</i> menu to quick-sort or quick-filter the data set:</p> <ul style="list-style-type: none"> • Filter: This option opens a dialog that allows you to define a filter. A filter is a condition that is supplied to an arbitrary WHERE clause that restricts the result set. • Remove Filter: This option removes all selection / exclusion filter conditions. • By Selection: This option refreshes the data set and displays only those rows whose column value matches the value in the cell currently selected. • Exclude Selection: This option refreshes the data set and ex- 	
316	<p>cludes those rows whose column value matches the value in the</p>	Chapter 7. Developer Tools

The Data Grid

The top row of the data grid displays the name of each column, the data type, and if applicable, the number of characters allowed. A column that is part of the primary key will additionally be marked with [PK].

To modify the displayed data:

- To change a numeric value within the grid, double-click the value to select the field. Modify the content in the square in which it is displayed.
- To change a non-numeric value within the grid, double-click the content to access the edit bubble. After modifying the content of the edit bubble, click the *Save* button to display your changes in the data grid, or *Cancel* to exit the edit bubble without saving.

To enter a newline character, click Ctrl-Enter or Shift-Enter. Newline formatting is only displayed when the field content is accessed via an edit bubble.

To add a new row to the table, enter data into the last (unnumbered) row of the table. As soon as you store the data, the row is assigned a row number, and a fresh empty line is added to the data grid.

To write a SQL NULL to the table, simply leave the field empty. When you store the new row, the will server fill in the default value for that column. If you store a change to an existing row, the value NULL will explicitly be written.


To write an empty string to the table, enter the special string '' (two single quotes) in the field. If you want to write a string containing solely two single quotes to the table, you need to escape these quotes, by typing ''

To delete a row, press the *Delete* toolbar button. A popup will open, asking you to confirm the deletion.

To commit the changes to the server, select the *Save* toolbar button. Modifications to a row are written to the server automatically when you select a different row.

Sort/Filter options dialog

You can access *Sort/Filter options dialog* by clicking on Sort/Filter button. This allows you to specify an SQL Filter to limit the data displayed and data sorting options in the edit grid window:



- Use *SQL Filter* to provide SQL filtering criteria. These will be added to the “WHERE” clause of the query used to retrieve the data. For example, you might enter:

```
id > 25 AND created > '2018-01-01'
```

- Use *Data Sorting* to sort the data in the output grid

To add new column(s) in data sorting grid, click on the [+] icon.

- Use the drop-down *Column* to select the column you want to sort.
- Use the drop-down *Order* to select the sort order for the column.

To delete a row from the grid, click the trash icon.

- Click the *Help* button (?) to access online help.
- Click the *Ok* button to save work.
- Click the *Close* button to discard current changes and close the dialog.

PGADMIN DEPLOYMENT

Pre-compiled and configured installation packages for pgAdmin 4 are available for a number of desktop environments; we recommend using an installer whenever possible. If you are interested in learning more about the project, or if a pgAdmin installer is not available for your environment, the pages listed below will provide detailed information about creating a custom deployment.

Contents:

PGADMIN PROJECT CONTRIBUTIONS

pgAdmin is an open-source project that invites you to get involved in the development process. For more information about contributing to the pgAdmin project, contact the developers on the pgAdmin mailing list pgadmin-hackers@postgresql.org to discuss any ideas you might have for enhancements or bug fixes.

In the sections listed below, you'll find detailed information about the development process used to develop, improve, and maintain the pgAdmin client.

Contents:

9.1 Submitting Patches

Before developing a patch for pgAdmin you should always contact the developers on the mailing list pgadmin-hackers@postgresql.org to discuss your plans. This ensures that others know if you're fixing a bug and can then avoid duplicating your work, and in the case of large patches, gives the community the chance to discuss and refine your ideas before investing too much time writing code that may later be rejected.

You should always develop patches against a checkout of the source code from the GIT source code repository, and not a release tarball. This ensures that you're working with the latest code on the branch and makes it easier to generate patches correctly. You can checkout the source code with a command like:

```
$ git clone git://git.postgresql.org/git/pgadmin4.git
```

Once you've made the changes you wish to make, commit them to a private development branch in your local repository. Then create a patch containing the changes in your development branch against the upstream branch on which your work is based. For example, if your current branch contains your changes, you might run:

```
$ git diff origin/master > my_cool_feature.diff
```

to create a patch between your development branch and the public master branch.

You can also create patches directly from the development tree, for example:

```
$ git diff > my_cool_feature.diff
```

If you are adding new files, you may need to stage them for commit, and then create your patch against the staging area. If any of the files are binary, for example, images, you will need to use the *-binary* option:

```
$ git add file1.py file2.py images/image1.png [...]
$ git diff --cached --binary > my_cool_feature.diff
```

Once you have your patch, check it thoroughly to ensure it meets the pgAdmin *Coding Standards*, and review it against the *Code Review Notes* to minimise the chances of it being rejected. Once you're happy with your work, mail it as an

attachment to the mailing list pgadmin-hackers@postgresql.org. Please ensure you include a full description of what the patch does, as well as the rationale for any important design decisions.

9.2 Code Overview

The bulk of pgAdmin is a Python web application written using the Flask framework on the backend, and HTML5 with CSS3, Bootstrap and jQuery on the front end. A desktop runtime is also included for users that prefer a desktop application to a web application, which is written in C++ using the QT framework.

9.2.1 Runtime

The runtime is essentially a Python webserver and browser in a box. Found in the `/runtime` directory in the source tree, it is a relatively simple QT application that is most easily modified using the **QT Creator** application.

9.2.2 Web Application

The web application forms the bulk of pgAdmin and can be found in the `/web` directory in the source tree. The main file is **pgAdmin4.py** which can be used to run the built-in standalone web server, or as a WSGI application for production use.

Configuration

The core application configuration is found in **config.py**. This file includes all configurable settings for the application, along with descriptions of their use. It is essential that various settings are configured prior to deployment on a web server; these can be overridden in **config_local.py** to avoid modifying the main configuration file.

User Settings

When running in desktop mode, pgAdmin has a single, default user account that is used for the desktop user. When running in server mode, there may be unlimited users who are required to login prior to using the application. pgAdmin utilised the **Flask-Security** module to manage application security and users, and provides options for self-service password reset and password changes etc.

Whether in desktop or server mode, each user's settings are stored in a SQLite database which is also used to store the user accounts. This is initially created using the **setup.py** script which will create the database file and schema within it, and add the first user account (with administrative privileges) and a default server group for them. A **settings** table is also used to store user configuration settings in a key-value fashion. Although not required, setting keys (or names) are typically formatted using forward slashes to artificially namespace values, much like the pgAdmin 3 settings files on Linux or Mac.

Note that the local configuration must be setup prior to **setup.py** being run. The local configuration will determine how the script sets up the database, particularly with regard to desktop vs. server mode.

9.2.3 pgAdmin Core

The heart of pgAdmin is the **pgadmin** package. This contains the globally available HTML templates used by the Jinja engine, as well as any global static files such as images, Javascript and CSS files that are used in multiple modules.

The work of the package is handled in it's constructor, **__init__.py**. This is responsible for setting up logging and authentication, dynamically loading other modules, and a few other tasks.

9.2.4 Modules

Units of functionality are added to pgAdmin through the addition of modules. These are Python object instance of classes, inherits the PgAdminModule class (a Flask Blueprint implementation), found in **web/pgadmin/utils.py**. It provide various hook points for other modules to utilise (primarily the default module - the browser).

To be recognised as a module, a Python package must be created. This must:

- 1) Be placed within the **web/pgadmin/** directory, and
- 2) Implements pgadmin.utils.PgAdminModule class
- 3) An instance variable (generally - named **blueprint**) representing that particular class in that package.

Each module may define a **template** and **static** directory for the Blueprint that it implements. To avoid name collisions, templates should be stored under a directory within the specified template directory, named after the module itself. For example, the **browser** module stores it's templates in **web/pgadmin/browser/templates/browser/**. This does not apply to static files which may omit the second module name.

In addition to defining the Blueprint, the **views** module is typically responsible for defining all the views that will be rendered in response to client requests, we must provide a REST API url(s) for these views. These must include appropriate route and security decorators. Take a look at the NodeView class, which uses the same approach as Flask's MethodView, it can be found in **web/pgadmin/browser/utils.py**. This specific class is used by browser nodes for creating REST API url(s) for different operation on them. i.e. list, create, update, delete, fetch children, get statistics/reversed SQL/dependencies/dependents list for that node, etc. We can use the same class for other purpose too. You just need to inherit that class, and overload the member variables operations, parent_ids, ids, node_type, and then register it as node view with PgAdminModule instance.

Most pgAdmin modules will also implement the **hooks** provided by the PgAdminModule class. This is responsible for providing hook points to integrate the module into the rest of the application - for example, a hook might tell the caller what CSS files need to be included on the rendered page, or what menu options to include and what they should do. Hook points need not exist if they are not required. It is the responsibility of the caller to ensure they are present before attempting to utilise them.

Hooks currently implemented are:

```
class MyModule(PgAdminModule):
    """
    This is class implements the pgadmin.utils.PgAdminModule, and
    implements the hooks
    """
    ...

    def get_own_stylesheets(self):
        """
        Returns:
            list: the stylesheets used by this module, not including any
                stylesheet needed by the submodules.
        """
        return [url_for('static', 'css/mymodule.css')]

    def get_own_javascripts(self):
        """
        Returns:
            list of dict:
            - contains the name (representation for this javascript
                module), path (url for it without .js suffix), deps (array of
                dependents), exports window object by the javascript module,
```

(continues on next page)

(continued from previous page)

```

        and when (would you like to load this javascript), etc
        information for this module, not including any script needed
        by submodules.
    """
    return [
        {
            'name': 'pgadmin.extension.mymodule',
            'path': url_for('static', filename='js/mymodule'),
            'exports': None,
            'when': 'server'
        }
    ]

    def get_own_menuitems(self):
        """
        Returns:
            dict: the menuitems for this module, not including
            any needed from the submodules.
        """
        return {
            'help_items': [
                MenuItem(
                    name='mnu_mymodule_help',
                    priority=999,
                    # We need to create javascript, which registers itself
                    # as module
                    module="pgAdmin.MyModule",
                    callback='about_show',
                    icon='fa fa-info-circle',
                    label=gettext('About MyModule')
                )
            ]
        }

    def get_panels(self):
        """
        Returns:
            list: a list of panel objects to add implemented in javascript
            module
        """
        return []
    ...

```

```
blueprint = MyModule('mymodule', __name__, static_url_path='/static')
```

pgAdmin Modules may include any additional Python modules that are required to fulfill their purpose, as required. They may also reference other dynamically loaded modules, but must use the defined hook points and fail gracefully in the event that a particular module is not present.

9.2.5 Nodes

Nodes are very similar to modules, it represents an individual node or, collection object on the browser treeview. To be recognised as a node module, a Python package (along with javascript modules) must be created. This must:

- 1) Be placed within the **web/pgadmin/browser/** directory, and

- 2) Implements the BrowserPluginModule, and registers the node view, which exposes required the REST APIs
- 3) An instance of the class object

9.2.6 Front End

pgAdmin uses javascript extensively for the front-end implementation. It uses require.js to allow the lazy loading (or, say load only when required), bootstrap for UI look and feel, Backbone for data manipulation of a node, Backform for generating properties/create dialog for selected node. We have divided each module in small chunks as much as possible. Not all javascript modules are required to be loaded (i.e. loading a javascript module for database will make sense only when a server node is loaded competely.) Please look at the the javascript files node.js, browser.js, menu.js, panel.js, etc for better understanding of the code.

9.3 Coding Standards

pgAdmin uses multiple technologies and multiple languages, each of which have their own coding standards.

9.3.1 General

In all languages, indentations should be made with 4 spaces, and excessively long lines wrapped where appropriate to ensure they can be read on smaller displays (80 characters is used in many places, but this is not a required maximum size as it's quite wasteful on modern displays). Typically lines should not be longer than 120 characters.

Comments should be included in all code where required to explain its purpose or how it works if not obvious from a quick review of the code itself.

9.3.2 CSS 3

CSS3 is used for styling and layout throughout the application. Extensive use is made of the Bootstrap Framework to aid in that process, however additional styles must still be created from time to time.

Most custom styling comes from individual modules which may advertise static stylesheets to be included in the module that is loading them via hooks.

Styling overrides (for example, to alter the Bootstrap look and feel) will typically be found in the **overrides.css** file in the main static file directory for the application.

Styling should never be applied inline in HTML, always through an external stylesheet, which should contain comments as appropriate to explain the usage or purpose for the style.

Styles should be specified clearly, one per line. For example:

```

/* iFrames should have no border */
iframe {
    border-width: 0;
}

/* Ensure the codemirror editor displays full height gutters when resized */
.CodeMirror, .CodeMirror-gutters {
    height: 100% !important;
}

```

All stylesheets must be CSS3 compliant.

9.3.3 HTML 5

HTML 5 is used for page structure throughout the application, in most cases being rendered from templates by the Jinja2 template engine in Flask.

All HTML must be HTML 5 compliant.

9.3.4 Javascript

Client-side code is written in Javascript using jQuery and various plugins. Whilst much of the code is rendered from static files, there is also code that is rendered from templates using Jinja2 (often to inject the users settings) or constructed on the fly from module hooks.

A typical Javascript function might be formatted like this (this snippet is from a template):

```
// Delete a server group
function delete_server_group(item) {
    alertify.confirm(
        'Delete server group?',
        'Are you sure you wish to delete the server group "{0}"?'.replace('{0}', tree.
↪getLabel(item)),
        function() {
            var id = tree.getId(item)
            $.post("{ url_for('NODE-server-group.delete') }", { id: id })
                .done(function(data) {
                    if (data.success == 0) {
                        report_error(data.errormsg, data.info);
                    } else {
                        var next = tree.next(item);
                        var prev = tree.prev(item);
                        tree.remove(item);
                        if (next.length) {
                            tree.select(next);
                        } else if (prev.length) {
                            tree.select(prev);
                        }
                    }
                })
        },
        null
    )
}
```

Note the use of a descriptive function name, using the underscore character to separate words in all lower case, and short but descriptive lower case variable names.

Note: From version 3.0 onwards, new or refactored code should be written using ES6 features and conventions.

C++

C++ code is used in the desktop runtime for the application, primarily with the QT framework and an embedded Python interpreter. Note the use of hanging braces, which may be omitted if on a single statement is present:

```

// Ping the application server to see if it's alive
bool PingServer(QUrl url)
{
    QNetworkAccessManager manager;
    QEventLoop loop;
    QNetworkReply *reply;
    QVariant redirectUrl;

    url.setPath("/utils/ping");

    do
    {
        reply = manager.get(QNetworkRequest(url));

        QObject::connect(reply, SIGNAL(finished()), &loop, SLOT(quit()));
        loop.exec();

        redirectUrl = reply->attribute(QNetworkRequest::RedirectionTargetAttribute);
        url = redirectUrl.toUrl();

        if (!redirectUrl.isNull())
            delete reply;
    } while (!redirectUrl.isNull());

    if (reply->error() != QNetworkReply::NoError)
        return false;

    QString response = reply->readAll();

    if (response != "PING")
    {
        qDebug() << "Failed to connect, server response: " << response;
        return false;
    }

    return true;
}

```

9.3.5 Python

Python is used for the backend web server. All code must be compatible with Python 2.7 and should include PyDoc comments whilst following the official Python coding standards defined in [PEP 8](#). An example function along with the required file header is shown below:

```

#####
#
# pgAdmin 4 - PostgreSQL Tools
#
# Copyright (C) 2013 - 2018, The pgAdmin Development Team
# This software is released under the PostgreSQL Licence
#
#####

"""Integration hooks for server groups."""

```

(continues on next page)

(continued from previous page)

```

from flask import render_template, url_for
from flask.ext.security import current_user

from pgadmin.settings.settings_model import db, ServerGroup

def get_nodes():
    """Return a JSON document listing the server groups for the user"""
    groups = ServerGroup.query.filter_by(user_id=current_user.id)

    value = ''
    for group in groups:
        value += '{"id":%d,"label": "%s", "icon": "icon-server-group", "inode": true}, ' \
                % (group.id, group.name)

    value = value[:-1]

    return value

```

9.4 Code Snippets

This document contains code for some of the important classes, listed as below:

- *PgAdminModule*
- *NodeView*
- *BaseDriver*
- *BaseConnection*

9.4.1 PgAdminModule

PgAdminModule is inherited from Flask.Blueprint module. This module defines a set of methods, properties and attributes, that every module should implement.

```

class PgAdminModule(Blueprint):
    """
    Base class for every PgAdmin Module.

    This class defines a set of method and attributes that
    every module should implement.
    """

    def __init__(self, name, import_name, **kwargs):
        kwargs.setdefault('url_prefix', '/' + name)
        kwargs.setdefault('template_folder', 'templates')
        kwargs.setdefault('static_folder', 'static')
        self.submodules = []
        self.parentmodules = []

        super(PgAdminModule, self).__init__(name, import_name, **kwargs)

    def create_module_preference():

```

(continues on next page)

(continued from previous page)

```

    # Create preference for each module by default
    if hasattr(self, 'LABEL'):
        self.preference = Preferences(self.name, self.LABEL)
    else:
        self.preference = Preferences(self.name, None)

    self.register_preferences()

    # Create and register the module preference object and preferences for
    # it just before the first request
    self.before_app_first_request(create_module_preference)

def register_preferences(self):
    pass

def register(self, app, options, first_registration=False):
    """
    Override the default register function to automagically register
    sub-modules at once.
    """
    if first_registration:
        self.submodules = list(app.find_submodules(self.import_name))

    super(PgAdminModule, self).register(app, options, first_registration)

    for module in self.submodules:
        if first_registration:
            module.parentmodules.append(self)
            app.register_blueprint(module)

def get_own_stylesheets(self):
    """
    Returns:
        list: the stylesheets used by this module, not including any
        stylesheet needed by the submodules.
    """
    return []

def get_own_messages(self):
    """
    Returns:
        dict: the i18n messages used by this module, not including any
        messages needed by the submodules.
    """
    return dict()

def get_own_javascripts(self):
    """
    Returns:
        list: the javascripts used by this module, not including
        any script needed by the submodules.
    """
    return []

def get_own_menuitems(self):
    """
    Returns:

```

(continues on next page)

(continued from previous page)

```

        dict: the menuitems for this module, not including
            any needed from the submodules.
    """
    return defaultdict(list)

def get_panels(self):
    """
    Returns:
        list: a list of panel objects to add
    """
    return []

def get_exposed_url_endpoints(self):
    """
    Returns:
        list: a list of url endpoints exposed to the client.
    """
    return []

@property
def stylesheets(self):
    stylesheets = self.get_own_stylesheets()
    for module in self.submodules:
        stylesheets.extend(module.stylesheets)
    return stylesheets

@property
def messages(self):
    res = self.get_own_messages()

    for module in self.submodules:
        res.update(module.messages)
    return res

@property
def javascripts(self):
    javascripts = self.get_own_javascripts()
    for module in self.submodules:
        javascripts.extend(module.javascripts)
    return javascripts

@property
def menu_items(self):
    menu_items = self.get_own_menuitems()
    for module in self.submodules:
        for key, value in module.menu_items.items():
            menu_items[key].extend(value)
    menu_items = dict((key, sorted(value, key=attrgetter('priority'))
                        for key, value in menu_items.items()))
    return menu_items

@property
def exposed_endpoints(self):
    res = self.get_exposed_url_endpoints()

    for module in self.submodules:
        res += module.exposed_endpoints

```

(continues on next page)

(continued from previous page)

```
return res
```

9.4.2 NodeView

NodeView class helps exposing basic REST APIs for different operations used by pgAdmin Browser. The basic idea has been taken from the Flask's MethodView class. Because - we need a lot more operations (not, just CRUD), we can not use it directly.

```
class NodeView(with_metaclass(MethodViewType, View)):
    """
    A PostgreSQL Object has so many operations/functions apart from CRUD
    (Create, Read, Update, Delete):
    i.e.
    - Reversed Engineered SQL
    - Modified Query for parameter while editing object attributes
      i.e. ALTER TABLE ...
    - Statistics of the objects
    - List of dependents
    - List of dependencies
    - Listing of the children object types for the certain node
      It will used by the browser tree to get the children nodes

    This class can be inherited to achieve the different routes for each of the
    object types/collections.

    OPERATION | URL | HTTP Method | Method
    -----+-----+-----+-----
    List       | /obj/[Parent URL]/ | GET | list
    Properties | /obj/[Parent URL]/id | GET | properties
    Create     | /obj/[Parent URL]/ | POST | create
    Delete     | /obj/[Parent URL]/id | DELETE | delete
    Update     | /obj/[Parent URL]/id | PUT | update

    SQL (Reversed | /sql/[Parent URL]/id | GET | sql
    Engineering) |
    SQL (Modified | /msql/[Parent URL]/id | GET | modified_sql
    Properties) |

    Statistics | /stats/[Parent URL]/id | GET | statistics
    Dependencies | /dependency/[Parent URL]/id | GET | dependencies
    Dependents | /dependent/[Parent URL]/id | GET | dependents

    Nodes | /nodes/[Parent URL]/ | GET | nodes
    Current Node | /nodes/[Parent URL]/id | GET | node

    Children | /children/[Parent URL]/id | GET | children

    NOTE:
    Parent URL can be seen as the path to identify the particular node.

    i.e.
    In order to identify the TABLE object, we need server -> database -> schema
    information.
    """
```

(continues on next page)

(continued from previous page)

```

operations = dict({
    'obj': [
        {'get': 'properties', 'delete': 'delete', 'put': 'update'},
        {'get': 'list', 'post': 'create'}
    ],
    'nodes': [{'get': 'node'}, {'get': 'nodes'}],
    'sql': [{'get': 'sql'}],
    'msql': [{'get': 'modified_sql'}],
    'stats': [{'get': 'statistics'}],
    'dependency': [{'get': 'dependencies'}],
    'dependent': [{'get': 'dependents'}],
    'children': [{'get': 'children'}]
})

@classmethod
def generate_ops(cls):
    cmds = []
    for op in cls.operations:
        idx = 0
        for ops in cls.operations[op]:
            meths = []
            for meth in ops:
                meths.append(meth.upper())
            if len(meths) > 0:
                cmds.append({
                    'cmd': op, 'req': (idx == 0),
                    'with_id': (idx != 2), 'methods': meths
                })
            idx += 1
    return cmds

# Inherited class needs to modify these parameters
node_type = None
# This must be an array object with attributes (type and id)
parent_ids = []
# This must be an array object with attributes (type and id)
ids = []

@classmethod
def get_node_urls(cls):
    assert cls.node_type is not None, \
        "Please set the node_type for this class ({0})".format(
            str(cls.__class__.__name__))
    common_url = '/'
    for p in cls.parent_ids:
        common_url += '<{0}:{1}>/'.format(str(p['type']), str(p['id']))

    id_url = None
    for p in cls.ids:
        id_url = '{0}<{1}:{2}>'.format(
            common_url if not id_url else id_url,
            p['type'], p['id'])

    return id_url, common_url

def __init__(self, **kwargs):
    self.cmd = kwargs['cmd']

```

(continues on next page)

(continued from previous page)

```

# Check the existence of all the required arguments from parent_ids
# and return combination of has parent arguments, and has id arguments
def check_args(self, **kwargs):
    has_id = has_args = True
    for p in self.parent_ids:
        if p['id'] not in kwargs:
            has_args = False
            break

    for p in self.ids:
        if p['id'] not in kwargs:
            has_id = False
            break

    return has_args, has_id and has_args

def dispatch_request(self, *args, **kwargs):
    http_method = flask.request.method.lower()
    if http_method == 'head':
        http_method = 'get'

    assert self.cmd in self.operations, \
        'Unimplemented command ({0}) for {1}'.format(
            self.cmd,
            str(self.__class__.__name__)
        )

    has_args, has_id = self.check_args(**kwargs)

    assert (
        self.cmd in self.operations and
        (has_id and len(self.operations[self.cmd]) > 0 and
         http_method in self.operations[self.cmd][0]) or
        (not has_id and len(self.operations[self.cmd]) > 1 and
         http_method in self.operations[self.cmd][1]) or
        (len(self.operations[self.cmd]) > 2 and
         http_method in self.operations[self.cmd][2])
    ), \
        'Unimplemented method ({0}) for command ({1}), which {2} ' \
        'an id'.format(http_method,
                       self.cmd,
                       'requires' if has_id else 'does not require')

    meth = None
    if has_id:
        meth = self.operations[self.cmd][0][http_method]
    elif has_args and http_method in self.operations[self.cmd][1]:
        meth = self.operations[self.cmd][1][http_method]
    else:
        meth = self.operations[self.cmd][2][http_method]

    method = getattr(self, meth, None)

    if method is None:
        return make_json_response(
            status=406,
            success=0,

```

(continues on next page)

(continued from previous page)

```

        errormsg=gettext(
            'Unimplemented method ({0}) for this url ({1})'.format(
                meth, flask.request.path)
        )
    )

    return method(*args, **kwargs)

    @classmethod
    def register_node_view(cls, blueprint):
        cls.blueprint = blueprint
        id_url, url = cls.get_node_urls()

        commands = cls.generate_ops()

        for c in commands:
            cmd = c['cmd'].replace('.', '-')
            if c['with_id']:
                blueprint.add_url_rule(
                    '{0}{1}'.format(
                        c['cmd'], id_url if c['req'] else url
                    ),
                    view_func=cls.as_view(
                        '{0}{1}'.format(
                            cmd, '_id' if c['req'] else ''
                        ),
                        cmd=c['cmd']
                    ),
                    methods=c['methods']
                )
            else:
                blueprint.add_url_rule(
                    '{0}'.format(c['cmd']),
                    view_func=cls.as_view(
                        cmd, cmd=c['cmd']
                    ),
                    methods=c['methods']
                )

    def module_js(self, **kwargs):
        """
        This property defines (if javascript) exists for this node.
        Override this property for your own logic.
        """
        return flask.make_response(
            flask.render_template(
                "{0}/js/{0}.js".format(self.node_type)
            ),
            200, {'Content-Type': 'application/x-javascript'}
        )

    def children(self, *args, **kwargs):
        """Build a list of treeview nodes from the child nodes."""
        children = []

        for module in self.blueprint.submodules:
            children.extend(module.get_nodes(*args, **kwargs))

```

(continues on next page)

(continued from previous page)

```

# Return sorted nodes based on label
return make_json_response(
    data=sorted(
        children, key=lambda c: c['label']
    )
)

```

9.4.3 BaseDriver

```

class BaseDriver(object):
    """
    class BaseDriver(object):

    This is a base class for different server types.
    Inherit this class to implement different type of database driver
    implementation.

    (For PostgreSQL, we will be using psycopg2)

    Abstract Properties:
    -----
    * Version (string):
        Current version string for the database server

    * libpq_version (string):
        Current version string for the used libpq library

    Abstract Methods:
    -----
    * get_connection(*args, **kwargs)
    - It should return a Connection class object, which may/may not be
      connected to the database server.

    * release_connection(*args, **kwargs)
    - Implement the connection release logic

    * gc()
    - Implement this function to release the connections assigned in the
      session, which has not been pinged from more than the idle timeout
      configuration.
    """

    @abstractproperty
    def Version(cls):
        pass

    @abstractproperty
    def libpq_version(cls):
        pass

    @abstractmethod
    def get_connection(self, *args, **kwargs):
        pass

    @abstractmethod

```

(continues on next page)

(continued from previous page)

```

def release_connection(self, *args, **kwargs):
    pass

@abstractmethod
def gc(self):
    pass

```

9.4.4 BaseConnection

```

class BaseConnection(object):
    """
    class BaseConnection(object)

    It is a base class for database connection. A different connection
    drive must implement this to expose abstract methods for this server.

    General idea is to create a wrapper around the actual driver
    implementation. It will be instantiated by the driver factory
    basically. And, they should not be instantiated directly.

    Abstract Methods:
    -----
    * connect(**kwargs)
      - Define this method to connect the server using that particular driver
        implementation.

    * execute_scalar(query, params, formatted_exception_msg)
      - Implement this method to execute the given query and returns single
        datum result.

    * execute_async(query, params, formatted_exception_msg)
      - Implement this method to execute the given query asynchronously and
        returns result.

    * execute_void(query, params, formatted_exception_msg)
      - Implement this method to execute the given query with no result.

    * execute_2darray(query, params, formatted_exception_msg)
      - Implement this method to execute the given query and returns the result
        as a 2 dimensional array.

    * execute_dict(query, params, formatted_exception_msg)
      - Implement this method to execute the given query and returns the result
        as an array of dict (column name -> value) format.

    * def async_fetchmany_2darray(records=-1, formatted_exception_msg=False):
      - Implement this method to retrieve result of asynchronous connection and
        polling with no_result flag set to True.
        This returns the result as a 2 dimensional array.
        If records is -1 then fetchmany will behave as fetchall.

    * connected()
      - Implement this method to get the status of the connection. It should
        return True for connected, otherwise False

```

(continues on next page)

(continued from previous page)

```

* reset()
  - Implement this method to reconnect the database server (if possible)

* transaction_status()
  - Implement this method to get the transaction status for this
    connection. Range of return values different for each driver type.

* ping()
  - Implement this method to ping the server. There are times, a connection
    has been lost, but - the connection driver does not know about it. This
    can be helpful to figure out the actual reason for query failure.

* _release()
  - Implement this method to release the connection object. This should not
    be directly called using the connection object itself.

NOTE: Please use BaseDriver.release_connection(...) for releasing the
      connection object for better memory management, and connection pool
      management.

* _wait(conn)
  - Implement this method to wait for asynchronous connection to finish the
    execution, hence - it must be a blocking call.

* _wait_timeout(conn, time)
  - Implement this method to wait for asynchronous connection with timeout.
    This must be a non blocking call.

* poll(formatted_exception_msg, no_result)
  - Implement this method to poll the data of query running on asynchronous
    connection.

* cancel_transaction(conn_id, did=None)
  - Implement this method to cancel the running transaction.

* messages()
  - Implement this method to return the list of the messages/notices from
    the database server.

* rows_affected()
  - Implement this method to get the rows affected by the last command
    executed on the server.
"""

ASYNC_OK = 1
ASYNC_READ_TIMEOUT = 2
ASYNC_WRITE_TIMEOUT = 3
ASYNC_NOT_CONNECTED = 4
ASYNC_EXECUTION_ABORTED = 5
ASYNC_TIMEOUT = 0.2
ASYNC_NOTICE_MAXLENGTH = 100000

@abstractmethod
def connect(self, **kwargs):
    pass

```

(continues on next page)

(continued from previous page)

```
@abstractmethod
def execute_scalar(self, query, params=None,
                  formatted_exception_msg=False):
    pass

@abstractmethod
def execute_async(self, query, params=None,
                 formatted_exception_msg=True):
    pass

@abstractmethod
def execute_void(self, query, params=None,
                formatted_exception_msg=False):
    pass

@abstractmethod
def execute_2darray(self, query, params=None,
                   formatted_exception_msg=False):
    pass

@abstractmethod
def execute_dict(self, query, params=None,
                formatted_exception_msg=False):
    pass

@abstractmethod
def async_fetchmany_2darray(self, records=-1,
                            formatted_exception_msg=False):
    pass

@abstractmethod
def connected(self):
    pass

@abstractmethod
def reset(self):
    pass

@abstractmethod
def transaction_status(self):
    pass

@abstractmethod
def ping(self):
    pass

@abstractmethod
def _release(self):
    pass

@abstractmethod
def _wait(self, conn):
    pass

@abstractmethod
def _wait_timeout(self, conn, time):
    pass
```

(continues on next page)

(continued from previous page)

```
@abstractmethod
def poll(self, formatted_exception_msg=True, no_result=False):
    pass

@abstractmethod
def status_message(self):
    pass

@abstractmethod
def rows_affected(self):
    pass

@abstractmethod
def cancel_transaction(self, conn_id, did=None):
    pass
```

9.5 Code Review Notes

This document lists a number of standard items that will be checked during the review process for any patches submitted for inclusion in pgAdmin.

- Ensure all code follows the pgAdmin *Coding Standards*.
- Ensure all code has unit test coverage and API/feature test coverage where appropriate.
- Copyright years must be correct and properly formatted (to make it easy to make bulk updates every year). The start date should always be 2013, and the end year the current year, e.g.

Copyright (C) 2013 - 2018, The pgAdmin Development Team

- Ensure there's a blank line immediately following any copyright headers.
- Include PyDoc comments for functions, classes and modules. Node modules should be ""Implements the XXXX node"".
- Ensure that any generated SQL does not have any leading or trailing blank lines and consistently uses 4 space indents for nice formatting.
- Don't special-case any Slony objects. pgAdmin 4 will have no direct knowledge of Slony, unlike pgAdmin 3.
- If you copy/paste modules, please ensure any comments are properly updated.
- Read all comments, and ensure they make sense and provide useful commentary on the code.
- Ensure that field labels both use PostgreSQL parlance, but also are descriptive. A good example is the "Init" field on an FTS Template - Init is the PG term, but adding the word "Function" after it makes it much more descriptive.
- Re-use code wherever possible, but factor it out into a suitably central location - don't copy and paste it unless modifications are required!
- Format code nicely to make it readable. Break up logical chunks of code with blank lines, and comment well to describe what different sections of code are for or pertain to.
- Ensure that form validation works correctly and is consistent with other dialogues in the way errors are displayed.

- On dialogues with Schema or Owner fields, pre-set the default values to the current schema/user as appropriate. In general, if there are common or sensible default values available, put them in the fields for the user.
- 1 patch == 1 feature. If you need to fix/update existing infrastructure in your patch, it's usually easier if it's in a separate patch. Patches containing multiple new features or unrelated changes are likely to be rejected.
- Ensure the patch is fully functional, and works! If a patch is being sent as a work in progress, not intended for commit, clearly state that it's a WIP, and note what does or does not yet work.

9.6 Translations

pgAdmin supports multiple languages using the [Flask-Babel](#) Python module. A list of supported languages is included in the `web/config.py` configuration file and must be updated whenever languages are added or removed with ISO 639-1 (two letter) language codes. The codes are named `$LANG` in this document.

9.6.1 Translation Marking

Strings can be marked for translation in either Python code (using `gettext()`) or Jinja templates (using `_()`). Here are some examples that show how this is achieved.

Python:

```
errmsg = gettext('No server group name was specified')
```

Jinja:

```
<input type="submit" value="{{ _('Change Password') }}">
```

```
<title>{{ _('%(appname)s Password Change', appname=config.APP_NAME) }}</title>
```

```
define(['sources/gettext', ...], function(gettext, ...){
    ...
    var alert = alertify.prompt(
        gettext('Password Change'),
        gettext('New password for %(userName)s', {userName: 'jsmith' }),
        ...
    )
})
```

9.6.2 Updating and Merging

Whenever new strings are added to the application, the template catalogue (`web/pgadmin/messages.pot`) and the existing translation catalogues (`web/pgadmin/translations/$LANG/LC_MESSAGES/messages.po`) must be updated and compiled. This can be achieved using the following commands from the `web` directory in the Python virtual environment for pgAdmin:

```
(pgadmin4) user$ pybabel extract -F babel.cfg -o pgadmin/messages.pot pgadmin
```

Once the template has been updated it needs to be merged into the existing message catalogues:

```
(pgadmin4) user$ pybabel update -i pgadmin/messages.pot -d pgadmin/translations
```

Finally, the message catalogues can be compiled for use:


```
(pgadmin4) user$ pybabel compile -d pgadmin/translations
```

9.6.3 Adding a New Language

Adding a new language is simple. First, add the language name and identifier to **web/config.py**:

```
# Languages we support in the UI
LANGUAGES = {
    'en': 'English',
    'zh': 'Chinese (Simplified)',
    'de': 'German',
    'pl': 'Polish'
}
```

Then, create the new message catalogue from the **web** directory in the source tree in the Python virtual environment for pgAdmin:

```
(pgadmin4) user$ pybabel init -i pgadmin/messages.pot -d pgadmin/translations -l $LANG
```


pgAgent is a job scheduling agent for Postgres databases, capable of running multi-step batch or shell scripts and SQL tasks on complex schedules.

pgAgent is distributed independently of pgAdmin. You can download pgAgent from the [download area](#) of the pgAdmin website.

Contents:

10.1 Using pgAgent

pgAgent is a scheduling agent that runs and manages jobs; each job consists of one or more steps and schedules. If two or more jobs are scheduled to execute concurrently, pgAgent will execute the jobs in parallel (each with its own thread).

A step may be a series of SQL statements or an operating system batch/shell script. Each step in a given job is executed when the previous step completes, in alphanumeric order by name. Switches on the *pgAgent Job* dialog (accessed through the *Properties* context menu) allow you to modify a job, enabling or disabling individual steps as needed.

Each job is executed according to one or more schedules. Each time the job or any of its schedules are altered, the next runtime of the job is re-calculated. Each instance of pgAgent periodically polls the database for jobs with the next runtime value in the past. By polling at least once every minute, all jobs will normally start within one minute of the specified start time. If no pgAgent instance is running at the next runtime of a job, it will run as soon as pgAgent is next started, following which it will return to the normal schedule.

When you highlight the name of a defined job in the pgAdmin tree control, the *Properties* tab of the main pgAdmin window will display details about the job, and the *Statistics* tab will display details about the job's execution.

10.1.1 Security concerns

pgAgent is a very powerful tool, but does have some security considerations that you should be aware of:

Database password - *DO NOT* be tempted to include a password in the pgAgent connection string - on Unix systems it may be visible to all users in 'ps' output, and on Windows systems it will be stored in the registry in plain text. Instead, use a libpq `~/.pgpass` file to store the passwords for every database that pgAgent must access. Details of this technique may be found in the [PostgreSQL documentation on .pgpass file](#).

System/database access - all jobs run by pgAgent will run with the security privileges of the pgAgent user. SQL steps will run as the user that pgAgent connects to the database as, and batch/shell scripts will run as the operating system user that the pgAgent service or daemon is running under. Because of this, it is essential to maintain control over the users that are able to create and modify jobs. By default, only the user that created the pgAgent database objects will be able to do this - this will normally be the PostgreSQL superuser.

10.2 Installing pgAgent

pgAgent runs as a daemon on Unix systems, and a service on Windows systems. In most cases it will run on the database server itself - for this reason, pgAgent is not automatically configured when pgAdmin is installed. In some cases however, it may be preferable to run pgAgent on multiple systems, against the same database; individual jobs may be targeted at a particular host, or left for execution by any host. Locking prevents execution of the same instance of a job by multiple hosts.

10.2.1 Database setup

Before using pgAdmin to manage pgAgent, you must create the pgAgent extension in the maintenance database registered with pgAdmin. To install pgAgent on a PostgreSQL host, connect to the *postgres* database, and navigate through the *Tools* menu to open the Query tool. For server versions 9.1 or later, and pgAgent 3.4.0 or later, enter the following command in the query window, and click the *Execute* icon:

```
CREATE EXTENSION pgagent;
```

This command will create a number of tables and other objects in a schema called 'pgagent'.

The database must also have the pl/pgsql procedural language installed - use the PostgreSQL CREATE LANGUAGE command to install pl/pgsql if necessary. To install pl/pgsql, enter the following command in the query window, and click the *Execute* icon:

```
CREATE LANGUAGE plpgsql;
```

If you are using an earlier version of PostgreSQL or pgAgent, use the *Open file* icon on the Query Tool toolbar to open a browse

- On Windows, it is usually located under *C:Program filespgAdmin III* (or *C:Program filesPostgreSQL8.xpgAdmin III* if installed with the PostgreSQL server installer).
- On Linux, it is usually located under */usr/local/pgadmin3/share/pgadmin3* or */usr/share/pgadmin3*.

After loading the file into the Query Tool, click the *Execute* icon to execute the script. The script will create a number of tables and other objects in a schema named *pgagent*.

10.2.2 Daemon installation on Unix

To install the pgAgent daemon on a Unix system, you will normally need to have root privileges to modify the system startup scripts. Modifying system startup scripts is quite system-specific so you should consult your system documentation for further information.

The program itself takes few command line options, most of which are only needed for debugging or specialised configurations:

```
Usage:
  /path/to/pgagent [options] <connect-string>

options:
  -f run in the foreground (do not detach from the terminal)
  -t <poll time interval in seconds (default 10)>
  -r <retry period after connection abort in seconds (>=10, default 30)>
  -s <log file (messages are logged to STDOUT if not specified)>
  -l <logging verbosity (ERROR=0, WARNING=1, DEBUG=2, default 0)>
```

The connection string is a standard PostgreSQL libpq connection string (see the [PostgreSQL documentation on the connection string](#) for further details). For example, the following command line will run pgAgent against a server listening on the localhost, using a database called 'pgadmin', connecting as the user 'postgres':

```
/path/to/pgagent hostaddr=127.0.0.1 dbname=postgres user=postgres
```

10.2.3 Service installation on Windows

pgAgent can install itself as a service on Windows systems. The command line options available are similar to those on Unix systems, but include an additional parameter to tell the service what to do:

```
Usage:
  pgAgent REMOVE <serviceName>
  pgAgent INSTALL <serviceName> [options] <connect-string>
  pgAgent DEBUG [options] <connect-string>

options:
  -u <user or DOMAIN\user>
  -p <password>
  -d <displayname>
  -t <poll time interval in seconds (default 10)>
  -r <retry period after connection abort in seconds (>=10, default 30)>
  -l <logging verbosity (ERROR=0, WARNING=1, DEBUG=2, default 0)>
```

The service may be quite simply installed from the command line as follows (adjust the path as required):

```
"C:\Program Files\pgAdmin III\pgAgent" INSTALL pgAgent -u postgres -p secret_
↪hostaddr=127.0.0.1 dbname=postgres user=postgres
```

You can then start the service at the command line using `net start pgAgent`, or from the *Services* control panel applet. Any logging output or errors will be reported in the Application event log. The DEBUG mode may be used to run pgAgent from a command prompt. When run this way, log messages will output to the command window.

10.3 Creating a pgAgent Job

pgAgent is a scheduling agent that runs and manages jobs; each job consists of steps and schedules.

To create or manage a job, use the pgAdmin tree control to browse to the server on which the pgAgent database objects were created. The tree control will display a *pgAgent Jobs* node, under which currently defined jobs are displayed. To add a new job, right click on the *pgAgent Jobs* node, and select *Create pgAgent Job...* from the context menu.

When the pgAgent dialog opens, use the tabs on the *pgAgent Job* dialog to define the steps and schedule that make up a pgAgent job.

The screenshot shows a 'Create - pgAgent job' dialog box with the following fields and options:

- Name:** eom_batch_process
- Enabled?:** Yes (checked)
- Job class:** Data Import (dropdown menu)
- Host agent:** (empty text field)
- Comment:** (empty text area)

Buttons at the bottom: Save, Cancel, Reset.

Use the fields on the *General* tab to provide general information about a job:


- Provide a name for the job in the *Name* field.
- Move the *Enabled* switch to the *Yes* position to enable a job, or *No* to disable a job.
- Use the *Job Class* drop-down to select a class (for job categorization).
- Use the *Host Agent* field to specify the name of a machine that is running pgAgent to indicate that only that machine may execute the job. Leave the field blank to specify that any machine may perform the job.

Note: It is not always obvious what value to specify for the Host Agent in order to target a job step to a specific machine. With pgAgent running on the required machines and connected to the scheduler database, you can use the following query to view the hostnames as reported by each agent:

```
SELECT jagstation FROM pgagent.pga_jobagent
```



Use the hostname exactly as reported by the query in the Host Agent field.






- Use the *Comment* field to store notes about the job.



Create - pgAgent Job

General Steps Schedules SQL

	Name	Enabled?	Kind	Connection type	On error
  ID		True	SQL	Local	Fail

Use the *Steps* tab to define and manage the steps that the job will perform. Click the Add icon (+) to add a new step; then click the compose icon (located at the left side of the header) to open the step definition dialog:

The screenshot shows the 'Create - pgAgent job' dialog box with the 'Steps' tab selected. At the top, there is a table listing job steps:

	Name	Enabled?	Kind	Connection type	On error
	ID	<input checked="" type="checkbox"/>	SQL	Local	Success

Below the table, the configuration form for the selected step 'ID' is shown. It includes the following fields and options:

- Name:** ID
- Enabled?:** Yes (checked)
- Kind:** SQL
- Connection type:** Local
- Database:** postgres
- Connection string:** (empty field)
- On error:** Success

At the bottom of the dialog, there are buttons for 'Save', 'Cancel', and 'Reset', along with an information icon and a help icon.

Use fields on the step definition dialog to define the step:

- Provide a name for the step in the *Name* field; please note that steps will be performed in alphanumeric order by name.
- Use the *Enabled* switch to include the step when executing the job (*True*) or to disable the step (*False*).
- Use the *Kind* switch to indicate if the job step invokes SQL code (*SQL*) or a batch script (*Batch*).
- If you select *SQL*, use the *Code* tab to provide SQL code for the step.
- If you select *Batch*, use the *Code* tab to provide the batch script that will be executed during the step.
- Use the *Connection type* switch to indicate if the step is performed on a local server (*Local*) or on a remote host (*Remote*). If you specify a remote connection should be used for the step, the *Connection string* field will be enabled, and you must provide a libpq-style connection string.
- Use the *Database* drop-down to select the database on which the job step will be performed.

- Use the *Connection string* field to specify a libpq-style connection string to the remote server on which the step will be performed. For more information about writing a connection string, please see the [PostgreSQL documentation](#).
- Use the *On error* drop-down to specify the behavior of pgAgent if it encounters an error while executing the step. Select from:
 - *Fail* - Stop the job if you encounter an error while processing this step.
 - *Success* - Mark the step as completing successfully, and continue.
 - *Ignore* - Ignore the error, and continue.
- Use the *Comment* field to provide a comment about the step.

The screenshot shows the 'Create - pgAgent job' dialog box with the 'Steps' tab selected. A table lists the step configuration:

ID	Name	Enabled?	Kind	Connection type	On error
		True	SQL	Local	Success

Below the table, the 'Code' sub-tab is active, showing the 'SQL query' field with the following content:

```

1 begin
2 end;

```


At the bottom of the dialog, there are buttons for 'Save', 'Cancel', and 'Reset', along with an information icon and a help icon.

Use the context-sensitive field on the step definition dialog's *Code* tab to provide the SQL code or batch script that will be executed during the step:

- If the step invokes SQL code, provide one or more SQL statements in the *SQL query* field.

- If the step performs a batch script, provide the script in the *Script* field. If you are running on a Windows server, standard batch file syntax must be used. When running on a Linux server, any shell script may be used, provided that a suitable interpreter is specified on the first line (e.g. `#!/bin/sh`).

When you've provided all of the information required by the step, click the compose icon to close the step definition dialog. Click the add icon (+) to add each additional step, or select the *Schedules* tab to define the job schedule.



Click the Add icon (+) to add a schedule for the job; then click the compose icon (located at the left side of the header) to open the schedule definition dialog:

Name	Enabled?	Start	End
schedule	<input type="checkbox"/> False	2017-11-06 14:27:22 +05:30	2017-11-06 14:26:35 +05:30

General Repeat Exceptions

Name

Enabled? No

Start

End

Comment

Use the fields on the schedule definition tab to specify the days and times at which the job will execute.

- Provide a name for the schedule in the *Name* field.
- Use the *Enabled* switch to indicate that pgAgent should use the schedule (*Yes*) or to disable the schedule (*No*).
- Use the calendar selector in the *Start* field to specify the starting date and time for the schedule.
- Use the calendar selector in the *End* field to specify the ending date and time for the schedule.
- Use the *Comment* field to provide a comment about the schedule.

Select the *Repeat* tab to define the days on which the schedule will execute.

Create - pgAgent job

General Steps Schedules SQL

Name	Enabled?	Start	End
schedule	<input type="checkbox"/> False	2017-11-06 14:27:22 +05:30	2017-11-06 14:26:35 +05:30

General Repeat Exceptions

Schedules are specified using a cron-style format.
 For each selected time or date element, the schedule will execute.
 e.g. To execute at 5 minutes pas every hour, simply select '05' in the Minutes list box.
 Value from more than one field may be specified in order to further control the schedule. e.g.
 To execute at 12:05 and 14:05 every Monday and Thursday, you would click minute 05, hours 12 and 14 and weekdays Monday and Thursday.

For additional flexibility, the Month Days check list includes an extra Last Day option. This matches the last day of the month, whether it happens to be the 28th, 29th, 30th or 31st.

Days

Week Days x

Month Days x

Months x

Times

Hours x

Minutes x

i ? Save Cancel Reset

Use the fields on the *Repeat* tab to specify the details about the schedule in a cron-style format. The job will execute on each date or time element selected on the *Repeat* tab.

Click within a field to open a list of valid values for that field; click on a specific value to add that value to the list of selected values for the field. To clear the values from a field, click the X located at the right-side of the field.

Use the fields within the *Days* box to specify the days on which the job will execute:

- Use the *Week Days* field to select the days on which the job will execute.
- Use the *Month Days* field to select the numeric days on which the job will execute. Specify the *Last Day* to indicate that the job should be performed on the last day of the month, irregardless of the date.
- Use the *Months* field to select the months in which the job will execute.

Use the fields within the *Times* box to specify the times at which the job will execute:

- Use the *Hours* field to select the hour at which the job will execute.
- Use the *Minutes* field to select the minute at which the job will execute.

Select the *Exceptions* tab to specify any days on which the schedule will *not* execute.

The screenshot shows the 'Create - pgAgent Job' dialog box with the 'Schedules' tab selected. The 'Exceptions' sub-tab is active, displaying a table with one exception row.

Name	Enabled?	Start	End
schedule	<input type="checkbox"/> False	2017-11-06 14:27:22 +05:30	2017-11-06 14:26:35 +05:30

Date	Time
2017-11-06	09:27

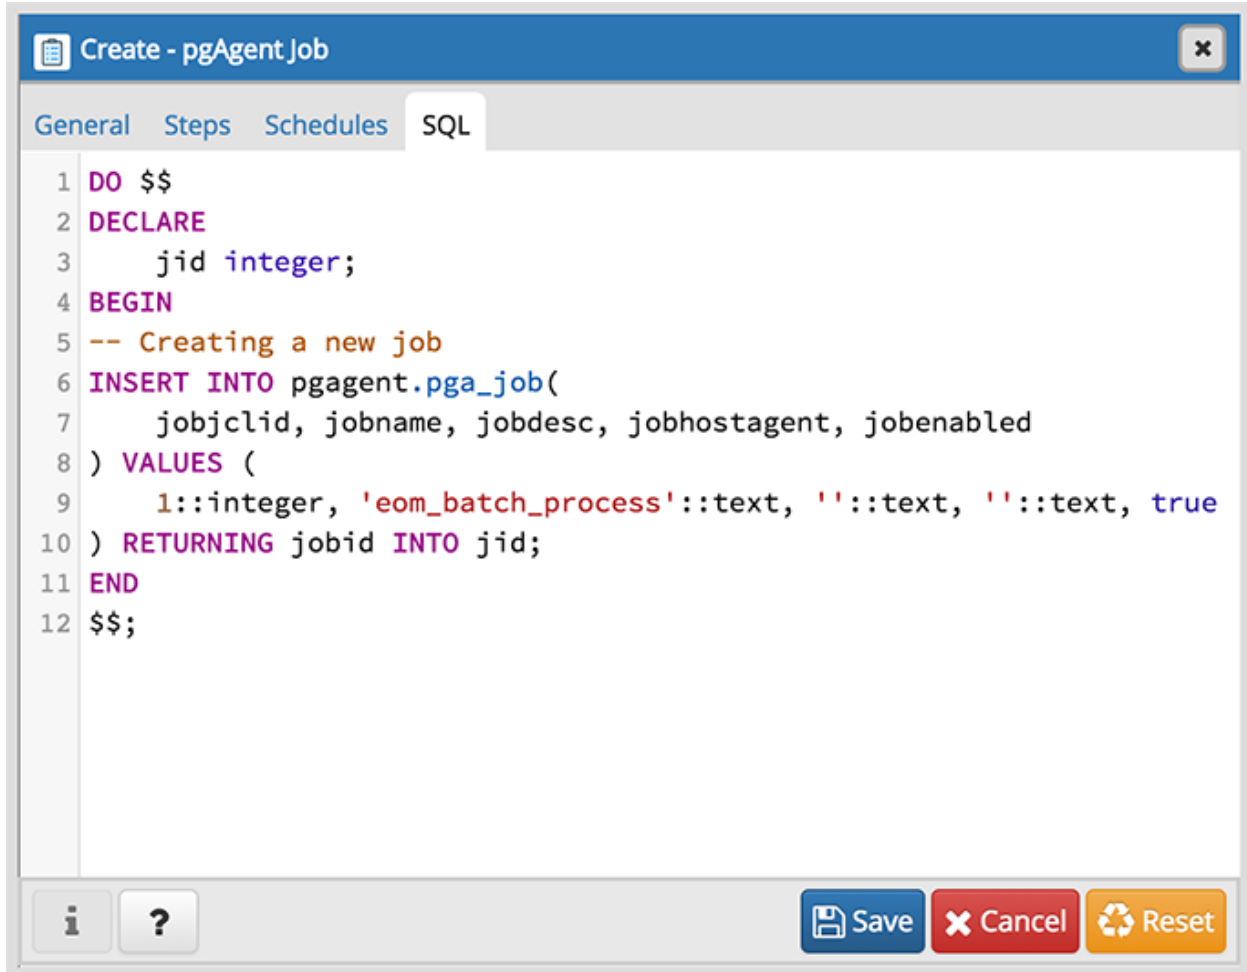
At the bottom of the dialog, there are buttons for 'Save', 'Cancel', and 'Reset', along with information and help icons.

Use the fields on the *Exceptions* tab to specify days on which you wish the job to not execute; for example, you may wish for jobs to not execute on national holidays.

Click the Add icon (+) to add a row to the exception table, then:


- Click within the *Date* column to open a calendar selector, and select a date on which the job will not execute. Specify *<Any>* in the *Date* column to indicate that the job should not execute on any day at the time selected.
- Click within the *Time* column to open a time selector, and specify a time on which the job will not execute. Specify *<Any>* in the *Time* column to indicate that the job should not execute at any time on the day selected.

When you've finished defining the schedule, you can use the *SQL* tab to review the code that will create or modify your job.



Click the *Save* button to save the job definition, or *Cancel* to exit the job without saving. Use the *Reset* button to remove your unsaved entries from the dialog.

After saving a job, the job will be listed under the *pgAgent Jobs* node of the pgAdmin tree control of the server on which it was defined. The *Properties* tab in the main pgAdmin window will display a high-level overview of the selected job, and the *Statistics* tab will show the details of each run of the job.



The screenshot shows the pgAdmin 4 interface with the 'Properties' window open for a job. The 'General' tab is selected, displaying the following details:

Property	Value
Name	test_batch__
ID	2
Enabled?	<input checked="" type="checkbox"/> Yes
Job class	Routine Maintenance
Host agent	
Created	2017-08-11 11:58:37.523716+05:30
Changed	2017-08-11 11:58:37.523716+05:30
Next run	2017-10-23 13:21:00+05:30
Last run	2017-10-23 13:20:01.8437+05:30
Last result	s
Running at	Not running currently.
Comment	

To modify an existing job or to review detailed information about a job, right-click on a job name, and select *Properties* from the context menu.

**CHAPTER
ELEVEN**

LICENCE

pgAdmin is released under the [PostgreSQL Licence](#), which is a liberal Open Source licence similar to BSD or MIT, and approved by the Open Source Initiative. The copyright for the project source code, website and documentation is attributed to the [pgAdmin Development Team](#).

RELEASE NOTES

pgAdmin release notes provide information on the features and improvements in each release. This page includes release notes for major releases and minor (bugfix) releases. Select your version from the list below to see the release notes for it.

12.1 Version 3.2

Release date: 2018-08-09

This release contains a number of features and fixes reported since the release of pgAdmin4 3.1

12.1.1 Features

Feature #2136 - Added version number for URL's to ensure that files are only cached on a per-version basis.

Feature #2214 - Add support for SCRAM password changes (requires psycopg2 >= 2.8).

Feature #3074 - Add support for reset saved password.

Feature #3397 - Add support for Trigger and JIT stats in the graphical query plan viewer.

Feature #3412 - Add support for primary key, foreign key, unique key, indexes and triggers on partitioned tables for PG/EPAS 11.

Feature #3506 - Allow the user to specify a fixed port number in the runtime to aid cookie whitelisting etc.

Feature #3510 - Add a menu option to the runtime to copy the appserver URL to the clipboard.

12.1.2 Bug fixes

Bug #3185 - Fix the upgrade check on macOS.

Bug #3191 - Fix a number of debugger execution issues.

Bug #3294 - Infrastructure (and changes to the Query Tool, Dashboards and Debugger) for realtime preference handling.

Bug #3309 - Fix Directory format support for backups.

Bug #3316 - Support running on systems without a system tray.

Bug #3319 - Cleanup and fix handling of Query Tool Cancel button status.

Bug #3363 - Fix restoring of restore options for sections.

Bug #3371 - Don't create a session when the /misc/ping test endpoint is called.

Bug #3446 - Various procedure/function related fixes for EPAS/PG 11.

Bug #3448 - Exclude system columns in Import/Export.

Bug #3457 - Fix debugging of procedures in EPAS packages.

- Bug #3458 - pgAdmin4 should work with python 3.7.
- Bug #3468 - Support SSH tunneling with keys that don't have a passphrase.
- Bug #3471 - Ensure the SSH tunnel port number is honoured.
- Bug #3511 - Add support to save and clear SSH Tunnel password.
- Bug #3526 - COST statement should not be automatically duplicated after creating trigger function.
- Bug #3527 - View Data->Filtered Rows dialog should be displayed.

12.2 Version 3.1

Release date: 2018-06-28

This release contains a number of features and fixes reported since the release of pgAdmin4 3.0

12.2.1 Features

- Feature #1447 - Add support for SSH tunneled connections
- Feature #2686 - Add an option to auto-complete keywords in upper case
- Feature #3204 - Add support for LISTEN/NOTIFY in the query tool
- Feature #3273 - Allow sorting in the file dialogue
- Feature #3362 - Function and procedure support for PG11
- Feature #3388 - Allow the connection timeout to be configured on a per-server basis

12.2.2 Bug fixes

- Bug #1220 - Backup and Restore should not be started if database name contains "=" symbol
- Bug #1221 - Maintenance should not be started if database name contains "=" symbol
- Bug #3179 - Fix an error generating SQL for trigger functions
- Bug #3238 - Standardise the error handling for parsing of JSON response messages from the server
- Bug #3250 - Fix handling of SQL_ASCII data in the Query Tool
- Bug #3257 - Catch errors when trying to EXPLAIN an invalid query
- Bug #3277 - Ensure server cleanup on exit only happens if the server actually started up
- Bug #3284 - F5 key should work to refresh Browser tree
- Bug #3289 - Fix handling of SQL_ASCII data in the Query Tool
- Bug #3290 - Close button added to the alertify message box, which pops up in case of backend error
- Bug #3295 - Ensure the debugger gets focus when loaded so shortcut keys work as expected
- Bug #3298 - Fixed query tool keyboard issue where arrow keys were not behaving as expected for execute options dropdown
- Bug #3303 - Fix a Japanese translation error that could prevent the server starting up
- Bug #3307 - Allow connections to servers with port numbers < 1024 which may be seen in container environments
- Bug #3310 - Fixed layout of the alertify error message in the query tool
- Bug #3333 - Ensure the runtime core application is setup before trying to access any settings
- Bug #3342 - Set SESSION_COOKIE_SAMESITE='Lax' per Flask recommendation to prevents sending cookies with CSRF-prone requests from external sites, such as submitting a form
- Bug #3353 - Handle errors properly if they occur when renaming a database
- Bug #3356 - Include the schema name on RE-SQL for packages
- Bug #3374 - Fix autocomplete

Bug #3392 - Fix IPv6 support in the container build

Bug #3411 - Fix a French translation error that could prevent the server starting up

12.3 Version 3.0

Release date: 2018-03-22

This release contains a number of features and fixes reported since the release of pgAdmin4 2.1

12.3.1 Features

Feature #1894 - Allow sorting when viewing/editing data

Feature #1978 - Add the ability to enable/disable UI animations

Feature #2895 - Add keyboard navigation options for the main browser windows

Feature #2896 - Add keyboard navigation in Query tool module via Tab/Shift-Tab key

Feature #2897 - Support keyboard navigation in the debugger

Feature #2898 - Support tab navigation in dialogs

Feature #2899 - Add configurable shortcut keys for various common options in the main window

Feature #2901 - Configurable shortcuts in the Debugger

Feature #2904 - Ensure clickable images/buttons have appropriate tooltips for screen readers

Feature #2950 - Add a marker (*/pga4dash/*) to the dashboard queries to allow them to be more easily filtered from server logs

Feature #2951 - Allow dashboard tables and charts to be enabled/disabled

Feature #3086 - Rewrite the runtime as a tray-based server which can launch a web browser

Feature #3098 - Unvendorize REACT so no longer required in our source tree

Feature #3140 - Add support for connecting using pg_service.conf files

Feature #3182 - Update Jasmine to v3

Feature #3184 - Add a French translation

Feature #3195 - Pass the service name to external processes

Feature #3246 - Update container build to use Alpine Linux and Gunicorn instead of CentOS/Apache

In addition, various changes were made for PEP8 compliance

12.3.2 Bug fixes

Bug #1173 - Add a comment to the existing node

Bug #1925 - Fix issue resizing column widths not resizable in Query Tool after first query

Bug #2104 - Runtime update display file version and copyright year under installers properties

Bug #2249 - Application no longer hangs after reload in runtime

Bug #2251 - Runtime fixed OSX html scroll direction ignored in MacOS setup

Bug #2309 - Allow text selection/copying from disabled CodeMirror instances

Bug #2480 - Runtime update fix to Context Menus on Mac that do not work

Bug #2578 - Runtime update fix to HTML access keys that don't work

Bug #2581 - Fix keyboard shortcut for text selection

Bug #2677 - Update Elephant icon for pgAdmin4 on Windows

- Bug #2776 - Fix unreadable font via Remote Desktop
- Bug #2777 - Fix spacing issue on server tree
- Bug #2783 - Runtime update fixed blank screen on Windows Desktop
- Bug #2906 - Correct display issues on HiDPI screens
- Bug #2961 - Issues when creating a pgAgent Schedule
- Bug #2963 - Fix unicode handling in the external process tools and show the complete command in the process viewer
- Bug #2980 - Copy text from the Query tool into the clipboard adds invisible characters
- Bug #2981 - Support keyboard navigation in the debugger
- Bug #2983 - Fix intermittent specified_version_number ValueError issue on restart
- Bug #2985 - Fix drag and drop issues
- Bug #2998 - Don't listen on port 443 if TLS is not enabled when launching the container
- Bug #3001 - Runtime update fix scrolling with mouse wheel on mac pgAdmin 4.2.1
- Bug #3002 - Fix block indent/outdent with configurable width
- Bug #3003 - Runtime update fix copy to clipboard
- Bug #3005 - Runtime update fix unable to select tabs in pgAdmin 4.2.1
- Bug #3013 - Fix a minor UI issue on dashboard while displaying subnode control in Backgrid
- Bug #3014 - Fix validation of sequence parameters
- Bug #3016 - Ensure debug messages are available in "messages" window when error occurs
- Bug #3021 - Update scan and index scan EXPLAIN icons for greater clarity
- Bug #3027 - Ensure we capture notices raised by queries
- Bug #3031 - Runtime issue causing double and single quotes not to work
- Bug #3039 - Runtime issue causing wrong row counts on count column
- Bug #3042 - Runtime issue causing empty dialog box when refreshing
- Bug #3043 - Runtime issue causing word sizing in macOS High Sierra
- Bug #3045 - Runtime issue causing copy cells issues copying cells for key binding
- Bug #3046 - Fix connection status indicator on IE/FF
- Bug #3052 - Don't include sizes on primitive data types that shouldn't have them when modifying columns
- Bug #3054 - Ensure the user can use keyboard shortcuts after using button controls such as Cancel, Open and Save
- Bug #3057 - Update the regression tests to fix issues with Python 3.5 and PG 9.2
- Bug #3058 - Fix on-click handling of treeview nodes that wasn't refreshing SQL/Dependencies/Dependents in some circumstances
- Bug #3060 - Fix quoting of function names in RE-SQL
- Bug #3066 - Ensure column names on indexes on views are properly quoted in RE-SQL
- Bug #3067 - Prevent the filter dialog CodeMirror from overflowing onto the button bar of the dialog
- Bug #3072 - Add a (configurable) limit to the number of pgAgent job history rows displayed on the statistics tab
- Bug #3073 - Ensure the pgAgent job start/end time grid fields synchronise with the subnode control and validate correctly
- Bug #3075 - Runtime issue causing Select, Update, and Insert script generation for a table fails to load
- Bug #3077 - Remove dependency on standards_conforming_strings being enabled
- Bug #3079 - Fix handling of tie/datetime array types when adding columns to a table
- Bug #3080 - Fix alignment issues in keyboard shortcut options
- Bug #3081 - Add missing reverse-engineered SQL header and drop statement for sequences
- Bug #3090 - Ensure message severity is decoded when necessary by the driver
- Bug #3094 - Ensure all messages are retrieved from the server in the Query Tool
- Bug #3105 - Ensure we can properly update rows with upper-case primary key columns
- Bug #3135 - Insert rows correctly when a table has OIDs and a Primary Key in uppercase

Bug #3122 - Ensure SSL options are pushed down to external tools like pg_dump
Bug #3129 - Handle opening of non-UTF8 compatible files
Bug #3137 - Allow copying of SQL from the dashboard tables
Bug #3138 - Fix tablespace tests for Python 3.x
Bug #3157 - Fix unicode handling in the external process tools and show the complete command in the process viewer
Bug #3171 - Runtime issue causing inability to scroll in File Selector with trackpad on OSX
Bug #3180 - Ensure Indexes are displayed on PG 10 tables
Bug #3196 - Ensure the file manager properly escapes file & directory names
Bug #3197 - Appropriately set the cookie path
Bug #3200 - Ensure the host parameter is correctly pickup up from the service file
Bug #3219 - Update required ChromeDriver version for current versions of Chrome
Bug #3226 - Move the field error indicators in front of the affected fields so they don't obscure spinners or drop downs etc.
Bug #3244 - Show more granular timing info in the query tool history panel
Bug #3248 - Ensure Alertify dialogues are modal to prevent them being closed by mis-click

12.4 Version 2.1

Release date: 2018-01-11

This release contains a number of features and fixes reported since the release of pgAdmin4 2.0

12.4.1 Features

Feature #1383 - Allow connections to be coloured in the treeview and query tool
Feature #1489 - Improve user interface for selection query in Data Filter window
Feature #2368 - Improve data entry in Query Tool
Feature #2781 - Allow configuration of CSV and clipboard formatting of query results
Feature #2802 - Allow connections to be coloured in the treeview and query tool.
Feature #2810 - Allow files to be opened by double clicking on them within Query Tool
Feature #2845 - Make the “Save Changes” prompts in the query tool optional
Feature #2849 - Add support for editing data in tables with OIDs but no primary keys and updates the editor to retrieve all row values on save, thus immediately showing default values and allowing subsequent editing without a refresh

12.4.2 Bug fixes

Bug #1365 - Prevent the Windows installer accepting paths containing invalid characters
Bug #1366 - Fix /NOICONS switch in the windows installer
Bug #1749 - Fixes in pgAgent module including; 1) allowing start date earlier than end date when scheduling job, 2) Datetime picker not displaying in grid and 3) validation error not displaying properly for Datetime control
Bug #2094 - Display relevant error messages when access is denied creating a schema
Bug #2098 - Cleanup some inconsistent error dialog titles
Bug #2258 - Fix handling of DATERANGE[] type
Bug #2278 - Display long names appropriately in dialogue headers

- Bug #2443 - Confirm with the user before exiting the runtime
- Bug #2524 - Fix debugging of self-referencing functions
- Bug #2566 - Fix the Pause/Resume Replay of WAL files for PostgreSQL 10
- Bug #2624 - Ensure the switch animation is consistent on the table dialogue and avoid displaying an error incorrectly
- Bug #2651 - Ensure estimated rows are included correctly in CREATE script for functions
- Bug #2679 - Getting started links does not open second time if User open any URL and Click on Close button with cross bar
- Bug #2705 - User can add expiry date on Windows
- Bug #2715 - Ensure we can download large files and keep the user informed about progress
- Bug #2720 - Ensure password changes are successful if authenticating using a pgsync file
- Bug #2726 - Ensure the auto-complete selection list can display longer names
- Bug #2738 - Ensure line numbers from CodeMirror don't appear on top of menus
- Bug #2748 - Format JSON/JSONB nicely when displaying it in the grid editor pop-up
- Bug #2760 - When selecting an SSL cert or key, update only the expected path in the UI, not all of them
- Bug #2765 - Do not decrypt the password when the password is 'None'. This should avoid the common but harmless exception "ValueError: IV must be 16 bytes long while decrypting the password."
- Bug #2768 - Only allow specification of a pgsync file if libpq >= 10
- Bug #2769 - Correct keyboard shortcut. Don't un-comment code with alt+. in the query tool. It's only supposed to respond to ctrl/cmd+
- Bug #2772 - Remove external links from Panel's context menu
- Bug #2778 - Ensure the datatype cache is updated when a domain is added
- Bug #2779 - Ensure column collation isn't lost when changing field size
- Bug #2780 - Ensure auto-indent honours the spaces/tabs config setting
- Bug #2782 - Re-hash the way that we handle rendering of special types such as arrays
- Bug #2787 - Quote the owner name when creating types
- Bug #2806 - Attempt to decode database errors based on lc_messages
- Bug #2811 - Display process output as it happens
- Bug #2820 - Logs available when executing backup and restore
- Bug #2821 - Attempt to decode database errors based on lc_messages
- Bug #2822 - Re-hash the way that we handle rendering of special types such as arrays.
- Bug #2824 - Fix a number of graphical explain rendering issues
- Bug #2836 - Fix counted rows display in table properties
- Bug #2842 - Fix a number of graphical explain rendering issues
- Bug #2846 - Add an option to manually count rows in tables to render the properties
- Bug #2854 - Fix utility output capture encoding
- Bug #2859 - Allow form validation messages to be close in case the eclipse anything on the form
- Bug #2866 - Ensure we don't show the full path on the server when using virtual filesystem roots in server mode for SSL certs
- Bug #2875 - Ensure the scroll location is retains in the query tool data grid if the user changes tab and then returns
- Bug #2877 - Remove the artificial limit of 4000 characters from text areas
- Bug #2880 - Honour whitespace properly in the data grid
- Bug #2881 - Fix support for time without timezone
- Bug #2886 - Resolve issue where Insert failed when tried with default primary key value
- Bug #2891 - Allow changing of the users password without leaving the app
- Bug #2892 - Refuse password changes (and tell the user) if the notification email cannot be sent
- Bug #2908 - Fix bundle creation on Windows which was failing due to rn line endings in code mirror
- Bug #2918 - Add missing init.py to backports.csv when building the MSVC windows build

Bug #2920 - Push HTTPD logs to container stdout/stderr as appropriate
Bug #2921 - Fixes in pgAgent module including; 1) allowing start date earlier than end date when scheduling job, 2) Datetime picker not displaying in grid and 3) validation error not displaying properly for Datetime control
Bug #2922 - Don't login the user with every request in desktop mode. Just do it once
Bug #2923 - Prevent the user pressing the select button in the file manager when it is supposed to be disabled
Bug #2924 - Cleanup the layout of the filter data dialogue
Bug #2928 - Prevent multiple connections to new slow-to-respond servers being initiated in error
Bug #2934 - Fix a reference before assignment error in the file dialogue
Bug #2937 - Prevent attempts to select directories as files in the file dialogue
Bug #2945 - Ensure invalid options can't be selected on triggers on views
Bug #2949 - Display complete SQL for FTS dictionaries
Bug #2952 - Don't try to render security URLs in desktop mode
Bug #2954 - Allow selection of validation error text
Bug #2974 - Clear the messages tab when running EXPLAIN/EXPLAIN ANALYZE
Bug #2993 - Fix view data for views/mat views

12.5 Version 2.0

Release date: 2017-10-05

This release contains a number of features and fixes reported since the release of pgAdmin4 1.6

12.5.1 Features

Feature #1918 - Add a field to the Server Dialogue allowing users to specify a subset of databases they'd like to see in the treeview
Feature #2135 - Significantly speed up loading of the application
Feature #2556 - Allow for slow vs. fast connection failures
Feature #2579 - Default the file browser view to list, and make it configurable
Feature #2597 - Allow queries to be cancelled from the dashboard and display additional info in the subnode control
Feature #2649 - Support use of SSL certificates for authentication
Feature #2650 - Support use of pgsync files
Feature #2662 - Ship with pre-configured paths that can work in both Server and Desktop modes out of the box
Feature #2689 - Update icons with new designs and remove from menus to de-clutter the UI

12.5.2 Bug fixes

Bug #1165 - Prevent continual polling for graph data on the dashboard if the server is disconnected
Bug #1697 - Update CodeMirror version
Bug #2043 - Properly handle trigger functions with parameters
Bug #2074 - Make \$ quoting consistent
Bug #2080 - Fix issue where Browser hangs/crashes when loading data (using sql editor) from table which contains large blob data
Bug #2153 - Fix handline of large file uploads and properly show any errors that may occur
Bug #2168 - Update CodeMirror version
Bug #2170 - Support SSL in the regression tests

Bug #2324 - Fix PostGIS Datatypes in SQL tab, Create / Update dialogues for Table, Column, Foreign Table and Type node

Bug #2447 - Update CodeMirror version

Bug #2452 - Install pgadmin4-v1 1.5 on Centos7

Bug #2501 - Fix collation tests on Windows, replace use of default 'POSIX' collation with 'C' collation for testing

Bug #2541 - Fix issues using special keys on MacOS

Bug #2544 - Correct malformed query generated when using custom type

Bug #2551 - Show tablespace on partitions

Bug #2555 - Fix issue in query tool where messages were not displaying from functions/procedures properly

Bug #2557 - Tidy up tab styling

Bug #2558 - Prevent the tab bar being hidden when detached tabs are being closed

Bug #2559 - Stop tool buttons from changing their styling unexpectedly

Bug #2560 - Fix View 'CREATE Script' Problem

Bug #2562 - Update CodeMirror version

Bug #2563 - Fix paths under non-standard virtual directories

Bug #2566 - Fix Pause/Resume Replay of WAL files for PostgreSQL 10

Bug #2567 - Use the proper database connection to fetch the default privileges in the properties tab of the database

Bug #2582 - Unset compression ratio if it is an empty string in Backup module

Bug #2586 - Cleanup feature tests

Bug #2590 - Allow navigation of query history using the arrow keys

Bug #2592 - Stop Flask from initialising service twice in Debug mode

Bug #2593 - Ensure babel-polyfill is loaded in older qWebKits

Bug #2594 - Fix disconnection of new databases

Bug #2596 - Define the proper NODE_ENV environment during running the webpack

Bug #2606 - Ensure role names are escaped in the membership control

Bug #2616 - Domain create dialog do not open and Font size issue in Security label control

Bug #2617 - Add missing pgagent file in webpack.config.js

Bug #2619 - Fix quoting of index column names on tables

Bug #2620 - Set database name to blank('') when job type is set to batch, while creating pgAgent job

Bug #2631 - Change mapping of cell from 'numeric' to 'integer' for integer control as numeric cell has been removed from the code

Bug #2633 - Fix pgAgent job step issues

Bug #2634 - Add New Server through Quick links

Bug #2637 - Fix Copy so it still works after query results have been copied

Bug #2641 - User management issues - styling and inability to edit users properly

Bug #2644 - Fix alertify notification messages where checkmark box disconnected from frame

Bug #2646 - Fix the path reference of load-node.gif which was referencing to vendor directory

Bug #2654 - Update datetime picker

Bug #2655 - Fix connection string validation for pgAgent jobs

Bug #2656 - Change Datetimepicker to expand from bottom in pgAgent so calendar does not get hidden

Bug #2657 - Fix syntax error while saving changes for start/end time, weekdays, monthdays, month, hours, minutes while updating the pgAgent Job

Bug #2659 - Fix issue where unable to add/update variables for columns of a table

Bug #2660 - Not able to select rows in History Tab

Bug #2668 - Fix RE-SQL for triggers with a single arg

Bug #2670 - Improve datamodel validations for default Validator if user (developer) does not implement validate function in datamodel

Bug #2671 - Fix array data type formatting for bigint, real, float, double precision
Bug #2681 - Reset query tool options before running tests
Bug #2684 - Fix layout of password prompt dialogue
Bug #2691 - View data option is missing from pgAdmin4 2.0 version
Bug #2692 - Base type is missing for Domain on pgAdmin4
Bug #2693 - User list is not available on User mapping pgAdmin4
Bug #2698 - User can not create function due to missing return type
Bug #2699 - Filtered Rows issue on pgAdmin4
Bug #2700 - Cancel button is visible after query executed successfully
Bug #2707 - Disable trigger button does not work on pgAdmin4
Bug #2708 - Tablespace name should displayed instead of %s(new_tablespace)s with Move Objects to another tablespace
Bug #2709 - Display user relations in schema prefixed by 'pg'
Bug #2713 - Fix an exception seen sometimes when the server is restarted
Bug #2742 - Ensure using an alternate role to connect to a database doesn't cause an error when checking recovery state.

12.6 Version 1.6

Release date: 2017-07-13

This release contains a number of features and fixes reported since the release of pgAdmin4 1.5

12.6.1 Features

Feature #1344 - Allow the Query Tool, Debugger and web browser tabs to be moved to different monitors as desired
Feature #1533 - Set focus on the first enabled field when a dialogue is opened
Feature #1535 - Teach dialogues about Escape to cancel, Enter to Save/OK, and F1 for help
Feature #1971 - Retain column sizing in the query tool results grid when the same query is re-run multiple times in a row
Feature #1972 - Prompt the user to save dirty queries rather than discard them for a more natural workflow
Feature #2137 - On-demand loading for the query tool results
Feature #2191 - Add support for the hostaddr connection parameter. This helps us play nicely with Kerberos/SSPI and friends
Feature #2282 - Overhaul the query history tab to allow browsing of the history and full query text
Feature #2379 - Support inserting multiple new rows into a table without clicking Save for each row
Feature #2485 - Add a shortcut to reset the zoom level in the runtime
Feature #2506 - Allow the user to close the dashboard panel
Feature #2513 - Add preferences to enable brace matching and brace closing in the SQL editors

12.6.2 Bug fixes

Bug #1126 - Various FTS dictionary cleanups
Bug #1229 - Fix default values and SQL formatting for event triggers
Bug #1525 - Make \$ quoting consistent
Bug #1575 - Properly display security labels on EPAS 9.2+

- Bug #1795 - Fix validation for external and range types
- Bug #1831 - Fix server stats display for EPAS 9.2, where inet needs casting to text for concatenation
- Bug #1851 - Reverse engineer SQL for table-returning functions correctly
- Bug #1860 - Ensure default values are honoured when adding/editing columns
- Bug #1888 - Fix various issues with pgAgent job steps and schedules
- Bug #1889 - Fix various issues with pgAgent job steps and schedules
- Bug #1890 - Fix various issues with pgAgent job steps and schedules
- Bug #1920 - Ensure saved passwords are effective immediately, not just following a restart when first saved
- Bug #1928 - Fix the handling of double precision[] type
- Bug #1934 - Fix import/export to work as expected with TSV data
- Bug #1999 - Handle warning correctly when saving query results to an unmounted USB drive
- Bug #2013 - Increase the default size of the Grant Wizard to enable it to properly display privileges at the default size on smaller displays
- Bug #2043 - Properly handle trigger functions with parameters
- Bug #2078 - Refresh the SQL editor view on resize to ensure the contents are re-rendered for the new viewport
- Bug #2086 - Allow editing of the WITH ADMIN option of role membership
- Bug #2113 - Correct the validation logic when modifying indexes/exclusion constraints
- Bug #2116 - Enable dialogue help buttons on Language and Foreign Table dialogues
- Bug #2142 - Fix canceling of Grant Wizard on Windows
- Bug #2155 - Fix removal of sizes from column definitions
- Bug #2162 - Allow non-superusers to debug their own functions and prevent them from setting global breakpoints
- Bug #2242 - Fix an issue in NodeAjaxControl caching with cache-node field and add cache-node field in Trigger & Event trigger node so that whenever the user creates new Trigger Function we get new data from server in NodeAjaxControl
- Bug #2324 - Fix the PostGIS Datatypes in SQL tab, Create / Update dialogues for Table, Column, Foreign Table and Type node
- Bug #2344 - Fix issue with ctrl-c / ctrl-v not working in query tool
- Bug #2348 - Fix issue when resizing columns in Query Tool/View Data where all row/columns will select/deselect
- Bug #2355 - Properly refresh the parent node when renaming children
- Bug #2357 - Cache statistics more reliably
- Bug #2381 - Fix the RE-SQL for for views to properly qualify trigger function names
- Bug #2386 - Display and allow toggling of trigger enable/disable status from the trigger dialogue
- Bug #2398 - Bypass the proxy server for local addresses on Windows
- Bug #2400 - Cleanup handling of default/null values when data editing
- Bug #2414 - Improve error handling in cases where the user tries to rename or create a server group that would duplicate an existing group
- Bug #2417 - Order columns in multi-column pkeys correctly
- Bug #2422 - Fix RE-SQL for rules which got the table name wrong in the header and DROP statement
- Bug #2425 - Handle composite primary keys correctly when deleting rows in the Edit Grid
- Bug #2426 - Allow creation of ENUM types with no members
- Bug #2427 - Add numerous missing checks to ensure objects really exist when we think they do
- Bug #2435 - Pass the database ID to the query tool when using the Script options
- Bug #2436 - Ensure the last placeholder is included when generating UPDATE scripts for tables
- Bug #2448 - Ensure that boolean checkboxes cycle values in the correct order
- Bug #2450 - Fix error on the stats tab with PG10. Also, rename the 10.0_plus template directory to 10_plus to match the new versioning
- Bug #2461 - Allow users to remove default values from columns properly
- Bug #2468 - Fix issue where function create script won't compile

- Bug #2470 - Fix an intermittent error seen during result polling
- Bug #2476 - Improvements to the Query Results grid including improvements to the UI and allow copy/paste from sets of rows, columns or arbitrary blocks of cells
- Bug #2477 - Ensure text editors render in an appropriate place on the results grid
- Bug #2479 - No need for the menu icon to link to the homepage, as pgAdmin is a SPA
- Bug #2482 - Use a more sensible name for Query Tool tabs
- Bug #2486 - Ensure the feature tests use the correct test settings database
- Bug #2487 - Maintain a client-side cache of preference values, populated using an async call
- Bug #2489 - Fix clipboard handling with large datasets
- Bug #2492 - Ensure the initial password is properly hashed during setup in web mode
- Bug #2498 - Properly handle bytea[], and 'infinity'::real/real[]
- Bug #2502 - Properly handle bytea[], and 'infinity'::real/real[]
- Bug #2504 - Update MatView and pgAgent modules to work with recent integer/numeric changes
- Bug #2507 - Ensure revoked public privileges are displayed in the RE-SQL for functions
- Bug #2518 - Fix encoding issue when saving servers
- Bug #2522 - Improve speed of Select All in the results grid
- Bug #2527 - Fix deletion of table rows with the column definition having NOT NULL TRUE and HAS NO DEFAULT VALUE
- Bug #2528 - Allow breakpoints to be set on triggers on views
- Bug #2529 - Resolve a number of issues with domains and domain constraints
- Bug #2532 - Refresh nodes correctly when there is a single child that is updated
- Bug #2534 - Fix handling of CREATE TABLE OF <type>
- Bug #2535 - Fix clear history functionality
- Bug #2540 - Ensure the save password option is enabled when creating a server

12.7 Version 1.5

Release date: 2017-05-19

This release contains a number of features and fixes reported since the release of pgAdmin4 1.4.

12.7.1 Features

Feature #2216 - Allow column or row selection in the query tool

12.7.2 Bug fixes

- Bug #2225 - Hide menu options for creating objects, if the object type is set to hidden. Includes Jasmine tests
- Bug #2253 - Fix various issues in CSV file download feature
- Bug #2257 - Improve handling of nulls and default values in the data editor
- Bug #2271 - Don't change the trigger icon back to "enabled" when the trigger is updated when it's disabled
- Bug #2284 - Allow creation of tables with pure numeric names
- Bug #2292 - Only reconnect to databases that were previously connected
- Bug #2314 - Fix various issues in CSV file download feature
- Bug #2315 - Fix sorting of sizes on the statistics views by sorting raw values and prettifying on the client side. Includes Jasmine tests for the prettyfying function

Bug #2318 - Order foreign table columns correctly
Bug #2331 - Fix binary search algorithm so new treeview nodes are added in the correct position
Bug #2336 - Update inode info when refreshing treeview nodes.
Bug #2339 - Ensure the treeview can be scrolled horizontally
Bug #2350 - Fix handling of default parameters ordering in functions
Bug #2354 - Fix the Backup module where it was not working if user changes its preference language other than English
Bug #2356 - Ensure errors thrown when deleting rows in the query tool in edit mode are shown properly
Bug #2360 - Fix various issues in CSV file download feature
Bug #2369 - Support loading files with Unicode BOMs
Bug #2377 - Update psycopg2 version for PostgreSQL 10 compatibility
Bug #2379 - Make various improvements to the NULL/DEFAULT handling in the data editor
Bug #2405 - Ensure object names are properly escaped for external process management
Bug #2410 - Fix PostgreSQL 10.0 compatibility issues

12.8 Version 1.4

Release date: 2017-04-13

This release contains a number of features and fixes reported since the release of pgAdmin4 1.3.

12.8.1 Features

Feature #2232 - Add the ability to gray-out/disable the “Save Password” option when creating a connection to a server
Feature #2320 - Added German translation

12.8.2 Bug fixes

Bug #2077 - Add missing “Run Now” option for pgAdmin jobs
Bug #2105 - Fix validation on the table dialogue so the Save button isn’t enabled if the name is removed and autovac custom settings are enabled
Bug #2145 - Resolve the issue for restoring the table from the backup
Bug #2187 - Ensure the web/ directory is cleared before upgrading Windows installations
Bug #2190 - Allow users to select UI language at login or from Preferences rather than unpredictable behaviour from browsers
Bug #2226 - Show tooltips for disabled buttons to help user learning
Bug #2241 - Fix numeric control validation in nested schemas
Bug #2243 - Fix dropping of databases with Unicode names
Bug #2244 - Prevent an error being displayed if the user views data on a table with no columns
Bug #2246 - Add missing braces to reverse engineered SQL header block for Functions
Bug #2258 - Fix handling of DATERANGE[] type
Bug #2264 - Resolve error message *ExtDeprecationWarning* displayed on new pgAdmin4 setup for Python 3.4 on ubuntu 14.04 Linux 64
Bug #2265 - Resolved import/Export issue for a table
Bug #2274 - Properly handle truncated table names
Bug #2277 - Resolved various file-system encoding/decoding related cases

- Bug #2281 - Ensure menus are updated after disconnecting a server
- Bug #2283 - Check if cell is in multiselect mode before setting default selection of multiple values
- Bug #2287 - Properly handle EXPLAIN queries entered directly by the user in the query tool
- Bug #2291 - Fix error highlighting in the query tool
- Bug #2299 - Fix usage of QString
- Bug #2303 - Fix ascending/descending sort order in backgrid while clicking on the headers
- Bug #2304 - Resolve the issue for restoring the table from the backup
- Bug #2305 - Resolve the issue where Generic function qtLiteral was not adapting values properly when they contain non ascii characters
- Bug #2310 - Fix Dialog Help where query tool/Debugger opens in new browser tab
- Bug #2319 - Resolve issue where Click on pgAdmin4 logo leads to unauthorized error
- Bug #2321 - Improved functionality of browser tree when adding new nodes if parent collection node has not loaded
- Bug #2330 - Ensure the query tool displays but does not render HTML returned by the server in the results grid

12.9 Version 1.3

Release date: 2017-03-10

This release contains a number of features and fixes reported since the release of pgAdmin4 1.2.

12.9.1 Features

- Feature #2036 - Query tool efficiency - SlickGrid result set format efficiency
- Feature #2038 - Query tool efficiency - Incremental back off when polling
- Feature #2163 - Make syntax highlighting more visible
- Feature #2210 - Build a universal Python wheel instead of per-python-version ones
- Feature #2215 - Improve visibility of syntax highlighting colours

12.9.2 Bug fixes

- Bug #1796 - Add missing “Run Now” option for pgAdmin jobs
- Bug #1797 - Resolve encoding issues with DATA_DIR
- Bug #1914 - Resolved error utf8’ codec can’t decode byte
- Bug #1983 - Fix bug in Sql query contains Arabic Charaters
- Bug #2089 - Add PARALLEL SAFE|UNSAFE|RESTRICTED support
- Bug #2115 - Fix exclusion constraint reverse engineered SQL
- Bug #2119 - Fix display of long integers and decimals
- Bug #2151 - Display un-sized varlen column types correctly in the query tool
- Bug #2154 - Fix display of long integers and decimals
- Bug #2159 - Resolve issue where Query editor is not working with Python2.6
- Bug #2160 - Various encoding fixes to allow ‘ascii’ codec to decode byte 0xc3 in position 66: ordinal not in range(128)
- Bug #2166 - Resolved import/Export issue for a table
- Bug #2173 - Resolved issues where Sequences API test cases are not working in PG9.2
- Bug #2174 - Resolved various file-system encoding/decoding related cases
- Bug #2185 - Removed sorting columns on the treeview

Bug #2192 - Fix startup complete tests to ensure we properly poll the server for completed startup
Bug #2198 - Fix function arguments when generating create SQL
Bug #2200 - Properly handle event trigger functions in different schemas
Bug #2201 - Fix renaming of check constraints when the table name is changed at the same time
Bug #2202 - Fix issue where Dependents query fails due to non ascii characters
Bug #2204 - Fixed issue where pgadmin 4 jobs not showing any activity
Bug #2205 - Fix display of boolean nulls in the query tool
Bug #2208 - Ensure primary key column names are quoted in View Data mode of the Query Tool
Bug #2212 - Ensure servers are deleted when their parent group is deleted
Bug #2213 - Enable right click on browser tree
Bug #2218 - Show the correct indeterminate state when editing new boolean values
Bug #2228 - Authenticate the runtime to the server
Bug #2230 - Prevent the Slonik logo obscuring the login dialogue on small displays in server mode

12.10 Version 1.2

Release date: 2017-02-10

This release contains a number of features and fixes reported since the release of pgAdmin4 1.1.

12.10.1 Features

Feature #1375 - Migrate the runtime to QtWebEngine from QtWebKit
Feature #1765 - Find and replace functionality with regexp and group replacement
Feature #1789 - Column width of data output panel should fit to data (as pgAdmin III)
Feature #1790 - [Web] Support setting a field's value to "null"
Feature #1848 - macOS appbundle is missing postgresql binaries for import etc.
Feature #1910 - Remember last used directory in the file manager
Feature #1911 - Direct path navigation in the file manager
Feature #1922 - Improve handling of corrupt configuration databases
Feature #1963 - Add a Chinese (Simplified) translation
Feature #1964 - Create a docs tarball along with the source tarball
Feature #2025 - Allow the SQL Editors to word-wrap
Feature #2124 - Create a template loader to simplify SQL template location, and remove duplicate templates

12.10.2 Bug fixes

Bug #1227 - Display improved error message for Debugger listener starting error and reset between executions
Bug #1267 - Fix issue where MINIFY_HTML doesn't work with the docs
Bug #1364 - Ensure dialogue control buttons are consistent
Bug #1394 - Fix Table dialogue column specification issues
Bug #1432 - Enhanced OSX File Browser
Bug #1585 - Cannot save scripts to the network
Bug #1599 - Ensure the grant wizard works with objects with special characters in the name
Bug #1603 - Fix quoting of objects names for external utilities.

- Bug #1679 - Re-engineer the background process executor to avoid using sqlite as some builds of components it relies on do not support working in forked children
- Bug #1680 - Render column headers at the correct width in the query tool under Firefox
- Bug #1729 - Improve display of role options
- Bug #1730 - Improve the display of role membership on both the properties panel and role dialogue
- Bug #1745 - Ensure breakpoints are cleared properly when working with Debugger
- Bug #1747 - Add newly created triggers to the treeview
- Bug #1780 - Properly size the SQL Editor gutter as the width of the line numbers increases
- Bug #1792 - List files and folders alphabetically
- Bug #1800 - Handle the template property on databases appropriately
- Bug #1801 - Handle databases with datallowconn == false
- Bug #1807 - Properly detect when files have changed in the query tool and set flag accordingly
- Bug #1830 - Fix a SQL error when reverse-engineering ROLE SQL on EPAS servers
- Bug #1832 - Prevent attempts to access what may be an empty list in Dependancies tab
- Bug #1840 - Enable/disable NULLs and ASC/DESC options for index columns and exclusion constraints appropriately
- Bug #1842 - Show index columns in the correct order in RE-SQL
- Bug #1855 - Ensure dialogue panels show their errors themselves, and not in the properties panel when creating Trigger Function
- Bug #1865 - Properly schema qualify domains when reverse engineering SQL
- Bug #1874 - Add file resources to the windows runtime
- Bug #1893 - Fix refreshing of Unique constraints
- Bug #1896 - Use the correct OID for retrieving properties of freshly created exclusion constraints
- Bug #1899 - Properly quote role names when specifying function ownership
- Bug #1909 - Handle startup errors more gracefully in the runtime
- Bug #1912 - Properly format arguments passed by triggers to functions
- Bug #1919 - Ensure all changes to rows are stored in the data editor
- Bug #1924 - Ensure the check_option is only set when editing views when appropriate
- Bug #1936 - Don't strip rn from "Download as CSV" batches of rows, as it leads to malformed data
- Bug #1937 - Generate mSQL for new schemas correctly
- Bug #1938 - Fix sorting of numerics in the statistics grids
- Bug #1939 - Updated dynamic default for the window size (90% x 90%)
- Bug #1949 - Ensure trigger function names are schema qualified in trigger RE-SQL
- Bug #1951 - Fix issue where nnable to browse table columns when oid values exceed max int
- Bug #1953 - Add display messages and notices received in the query tool
- Bug #1961 - Fix upgrade check on Python 3
- Bug #1962 - Ensure treeview collection nodes are translated in the UI
- Bug #1967 - Store layout changes on each adjustment
- Bug #1976 - Prevent users selecting elements of the UI that shouldn't be selectable
- Bug #1979 - Deal with Function arguments correctly in the properties dialogue
- Bug #1986 - Fix various encoding issues with multibyte paths and filenames resulting in empty file save
- Bug #1992 - Quote identifiers correctly in auto-complete
- Bug #1994 - Update to show modifications in edit grid
- Bug #2000 - Allow setting of effective_io_concurrency on tablespaces in 9.6+
- Bug #2005 - Fix various mis-spellings of VACUUM
- Bug #2006 - Fix error when modifying table name or set schema on tables with postgis geometry column
- Bug #2007 - Correctly sort rows by the pkey when viewing first/last 100

- Bug #2009 - Reset the column list properly if the access method is changed on an index to ensure error handling works correctly
- Bug #2012 - Prevent attempts to create server groups with no name
- Bug #2015 - Enable trigger option when user tries to change Row trigger value through properties section
- Bug #2024 - Properly handle setting comments and other options on databases with allowconn = False
- Bug #2026 - Improve detection of the plpgsql extension and functions before allowing debugging
- Bug #2027 - Fix inconsistent table styling
- Bug #2028 - Fix display of double scrollbars on the grant wizard
- Bug #2032 - Fix time formatting on dashboards
- Bug #2033 - Show icons for unique and exclusion constraints in the dependency/dependents panels
- Bug #2045 - Update copyright year on doc page
- Bug #2046 - Fix error when setting up regression on Windows for pgadmin4
- Bug #2047 - Ensure dialogues cannot be moved under the navbar
- Bug #2061 - Enable/disable NULLs and ASC/DESC options for index columns and exclusion constraints appropriately
- Bug #2065 - Improve display of columns of exclusion constraints and foreign keys in the properties lists
- Bug #2069 - Correct tablespace displayed in table properties
- Bug #2076 - Handle sized time/timestamp columns correctly
- Bug #2109 - Update copyright year
- Bug #2110 - Handle saved directories that no longer exist gracefully
- Bug #2112 - Enable comments on Initial database through right Click
- Bug #2133 - Fix display of graphical query plans for UPDATE/DELETE queries
- Bug #2138 - Fix display of zeros in read-only grid editors
- Bug #2139 - Fixed issue causing Message (Connection to the server has been lost.) displayed with Materialized view and view under sql tab
- Bug #2152 - Fix handling of “char” columns
- Bug #2156 - Added compatibility fixes for newer versions of Jinja2 (e.g. 2.9.5+)

12.11 Version 1.1

Release date: 2016-10-27

This release contains a number of features and fixes reported since the release of pgAdmin4 1.0;

12.11.1 Features

Feature #1328 - Add Python 3.5 Support

Feature #1859 - Include wait information on the activity tab of the dashboards

12.11.2 Bug fixes

Bug #1155 - Display the start value when the user creates sequence

Bug #1531 - Fix to update privileges for Views and Materials Views where “string indices must be integers error” displayed

Bug #1574 - Display SQL in SQL pane for security label in PG and EPAS server

Bug #1596 - Correct spelling error from evnt_turncate to evnt_truncate

Bug #1599 - Ensure the grant wizard works with objects with special characters in the name

Bug #1728 - Properties refreshing after objects are edited

Bug #1739 - Prevent the user from trying to. . . .

Bug #1785 - Correctly identify server type upon first connection

Bug #1786 - Ensure errorModel unset property is set correctly when adding a new server

Bug #1808 - Set seconds to valid value in pgAgent job schedule

Bug #1817 - Display message “server does not support ssl” if server with ca-cert or ca-full added

Bug #1821 - Optionally sign both the Mac app bundle and the disk image

Bug #1822 - Handle non-ascii responses from the server when connecting

Bug #1823 - Attempt to sign the Windows installer, failing with a warning if there’s no cert available

Bug #1824 - Add documentation for pgAgent

Bug #1835 - Allow users to choose SELECT permissions for tables and sequences in the grant wizard

Bug #1837 - Fix refreshing of FTS dictionaries which was causing error “Connection to the server has been lost”

Bug #1838 - Don’t append new objects with the wrong parent in tree browser if the correct one isn’t loaded

Bug #1843 - Function definition matches value returned from pg_get_functiondef()

Bug #1847 - Identify the collation correctly when reverse engineering table SQL. ERROR: schema “default” does not exist no longer displayed

Bug #1849 - Remove security keys from config.py/config_local.py

Bug #1857 - Fix error while renaming FTS dictionary and FTS template nodes

Bug #1858 - Ensure the File Manager honours the file type while traversing the directories.

Bug #1861 - Properly generate exclusion constraint SQL.

Bug #1863 - Correctly quote type names in reverse engineered SQL for tables

Bug #1864 - Fix layout of DateTimePicker control help message.

Bug #1868 - Resolved issue where Integer type of preferences are not updated

Bug #1872 - Fix the file manager when used under Python 3.

Bug #1877 - Ensure preferences values are stored properly.

Bug #1878 - Ensure steps and schedules can be created in empty jobs. ProgrammingError: can’t adapt type ‘Undefined’ was displayed

Bug #1880 - Add new indexes to the correct parent on the treeview.

12.12 Version 1.0

Release date: 2016-09-29

The first major release of pgAdmin 4. With a more modern look and feel, this release includes the following features;

- Multiplatform
- Designed for multiple PostgreSQL versions and derivatives
- Extensive documentation
- Multiple deployment models
- Tools
- Routine maintenance
- Create, view and edit all common PostgreSQL objects
- Multibyte support