
OmniDB

Release 2.15.0

Oct 22, 2020

Contents:

1	1. Introduction	1
2	2. Installation	3
3	3. Creating Users and Connections	9
4	4. Managing Databases	19
5	5. Creating, Changing and Removing Tables	33
6	6. Managing Table Data	43
7	7. Writing SQL Queries	51
8	8. Visualizing Query Plans	57
9	9. Visualizing Data	63
10	10. Managing other Elements	67
11	11. Additional Features	73
12	12. OmniDB Config Tool	81
13	13. Writing and Debugging PL/pgSQL Functions	83
14	14. Monitoring Dashboard	101
15	15. Logical Replication	119
16	16. pglogical	127
17	17. Postgres-BDR	143
18	18. Postgres-XL	157
19	19. Deploying omnidb-server	173
20	20. Console Tab	179

21	21. Plugin System	183
22	22. Advanced Object Search	185
23	23. Debugger Plugin Installation	189
24	23.1. Linux Installation	191
25	23.2. Windows Installation	195
26	23.3. FreeBSD Installation	197
27	23.4. MacOSX Installation	199
28	23.5. Post-installation steps ** REQUIRED **	201
29	Indices and tables	203

1. Introduction

OmniDB is an open source browser-based app designed to access and manage many different Database Management systems, e.g. PostgreSQL, Oracle and MySQL. OmniDB can run either as an App or via Browser, combining the flexibility needed for various access paths with a design that puts security first. OmniDB is actively developed, automatically tested on a variety of databases and browsers and comes with full documentation.

Since early development, OmniDB was designed as an browser-based app. Consequently, it runs in any browser, from any operational system. It can be accessed by several computers and multiple users, each one of them with his/her own group of connections. It also can be hosted in any operational system, without the need of install any dependencies. We will see further details on installation in the next chapters.

OmniDB's main objective is to offer an unified workspace with all functionalities needed to manipulate different DBMS. DBMS specific tools aren't required: in OmniDB, the context switch between different DBMS is done with a simple connection switch, without leaving the same page. The end-user's sensation is that there is no difference when he/she manipulates different DBMS, it just feels like different connections.

Despite this, OmniDB is built with simplicity in mind, designed to be a fast and lightweight browser-based application. OmniDB is also powered by the WebSocket technology, allowing the user to execute multiple queries and procedures in multiple databases in multiple hosts in background.

OmniDB is also secure. All OmniDB user data are stored encrypted, and no database password is stored at all. When the user first connects to a database, OmniDB asks for the password. This password is encrypted and stored in memory for a specific amount of time. When this time expires, OmniDB asks the password again. This ensures maximum security for the database OmniDB is connecting to.

1.1 History

OmniDB's creators, Rafael Thofehn Castro and William Ivanski, worked in a company where they needed to deal with several different databases from customers on a daily basis. These databases were from different DBMS technologies, and so they needed to keep switching between database management tools (typically one for each DBMS). As they were not keen of the existing unified database management tools (that could manage different DBMS), they came up with OmniDB's main idea.

OmniDB's first version was presented as an undergrad final project in the Computer Science Course from the Federal University of Paraná, in Brazil. The objective was to trace a common line between popular DBMS, and to study deeply their *metadata*. The result was a tool written in ASP.NET/C# capable of connecting and identifying the main structures (tables, keys, indexes and constraints), in a generic way, from several DBMS:

- Firebird
- MariaDB / MySQL
- Oracle
- PostgreSQL
- SQLite
- Microsoft SQL Server

OmniDB's first version also allowed the conversion between all DBMSs supported by the tool. This feature was developed to be user friendly, requiring just a few steps: the user needs to select a source connection, the structures that will be converted (just tables and all their structures, along with their data) and the target connection.

2. Installation

OmniDB provides 2 kinds of packages to fit every user needs:

- **OmniDB Application:** Runs a web server on a random port behind, and provides a simplified web browser window to use OmniDB interface without any additional setup. Just feels like a desktop application.
- **OmniDB Server:** Runs a web server on a random port, or a port specified by the user. User needs to connect to it through a web browser. Provides user management, ideal to be hosted on a server on users' networks.

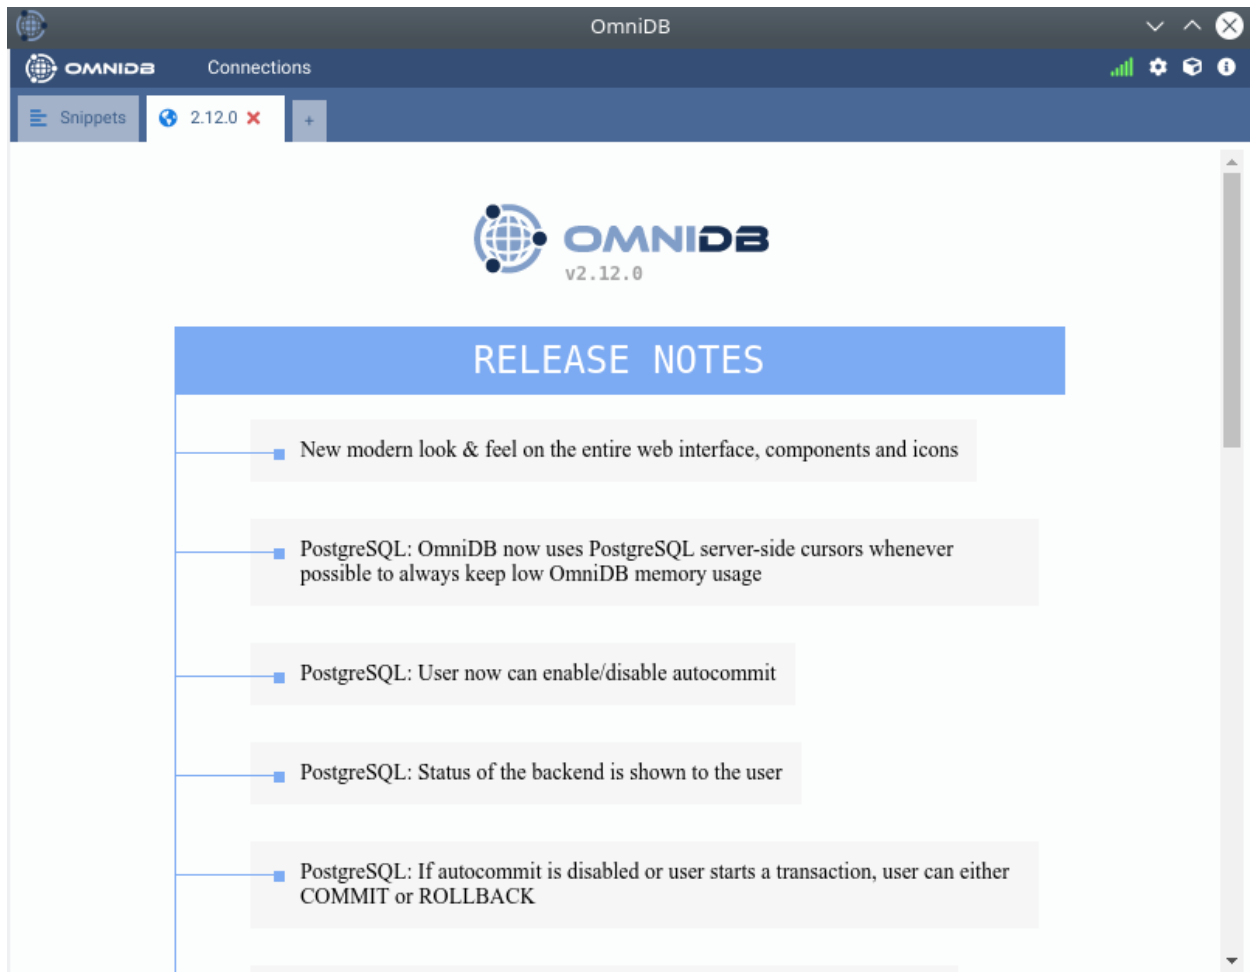
Both application and server can be installed on the same machine.

2.1 OmniDB Application

In order to run OmniDB app, you don't need to install any additional piece of software. Just head to omnidb.org and download the latest package for your specific operating system and architecture:

- Linux 64 bits
 - DEB installer
 - RPM installer
- Windows 64 bits
 - EXE installer
- Mac OSX
 - DMG installer

Use the specific installer for your Operational System and it will be available through your desktop environment application menu or via command line with `omnidb-app`.



2.2 OmniDB Server

Like OmniDB app, OmniDB server doesn't require any additional piece of software and the same options for operating system and architecture are provided.

Use the specific installer for your Operational System and it will be available through command line with `omnidb-server`:

```
user@machine:~$ omnidb-server
Starting OmniDB websocket...
Checking port availability...
Starting websocket server at port 25482.
Starting OmniDB server...
Checking port availability...
Starting server OmniDB 2.4.0 at 0.0.0.0:8000.
Starting migration of user database from version 0.0.0 to version 2.4.0
OmniDB successfully migrated user database from version 0.0.0 to version 2.4.0
Press Ctrl+C to exit
```

Note how OmniDB starts a *websocket server* in port 25482 and a *web server* in port 8000. You can also specify both ports and listening address:

```

user@machine:~$ omnidb-server -p 8080 -w 25000 -H 127.0.0.1
Starting OmniDB websocket...
Checking port availability...
Starting websocket server at port 25000.
Starting OmniDB server...
Checking port availability...
Starting server OmniDB 2.4.0 at 0.0.0.0:8080.
Starting migration of user database from version 0.0.0 to version 2.4.0
OmniDB successfully migrated user database from version 0.0.0 to version 2.4.0
Press Ctrl+C to exit

```

2.3 OmniDB with Oracle

OmniDB app and server does not require any piece of additional software, as explained above. But if you are going to connect to an *Oracle* database, then you need to download and install *Oracle Instant Client* (or extract it to a specific folder, depending on the operating system you use):

- **MacOSX:** Download Oracle Instant Client (64-bit) and extract in ~/lib;
- **Linux:** Download Oracle Instant Client (32-bit) (64-bit) and install it on your system, then set LD_LIBRARY_PATH;
- **Windows:** Download Oracle Instant Client (32-bit) (64-bit) and extract it into OmniDB's folder.

Note for Windows users using OmniDB app: For OmniDB 2.8 and above, you will need to extract Oracle Instant Client libraries inside of folder OMNIDBAPPINSTALLFOLDER\resources\app\omnidb-server.

2.4 OmniDB User Database

Since version 2.4.0, upon initialization both server and app will create a file ~/.omnidb/omnidb-app/omnidb.db (for OmniDB app) or ~/.omnidb/omnidb-server/omnidb.db (for OmniDB server) in the user home directory, if it does not exist. That can be confirmed by the message *OmniDB successfully migrated user database from version 0.0.0 to version 2.4.0* you saw above. This file is also called **user database** and contains user data. If it already exists, then OmniDB will check whether the version of the server matches the version of the user database:

```

user@machine:~$ omnidb-server
Starting OmniDB websocket...
Checking port availability...
Starting websocket server at port 25482.
Starting OmniDB server...
Checking port availability...
Starting server OmniDB 2.4.0 at 0.0.0.0:8000.
User database version 2.4.0 is already matching server version.
Press Ctrl+C to exit

```

Future releases of OmniDB will contain the **user database migration** SQL commands required to upgrade the user database, if necessary. This way user data is not lost by upgrading OmniDB. Imagine the following scenario: you use OmniDB 2.4.0 now and you decide to upgrade it to newest release 2.5.0, for example. After the upgrade, when you start OmniDB server, it will apply the changes version 2.5.0 requires. So you will see something like that:

```

user@machine:~$ omnidb-server
Starting OmniDB websocket...
Checking port availability...

```

(continues on next page)

(continued from previous page)

```
Starting websocket server at port 25482.  
Starting OmniDB server...  
Checking port availability...  
Starting server OmniDB 2.5.0 at 0.0.0.0:8000.  
Starting migration of user database from version 2.4.0 to version 2.5.0  
OmniDB successfully migrated user database from version 2.4.0 to version 2.5.0  
Press Ctrl+C to exit
```

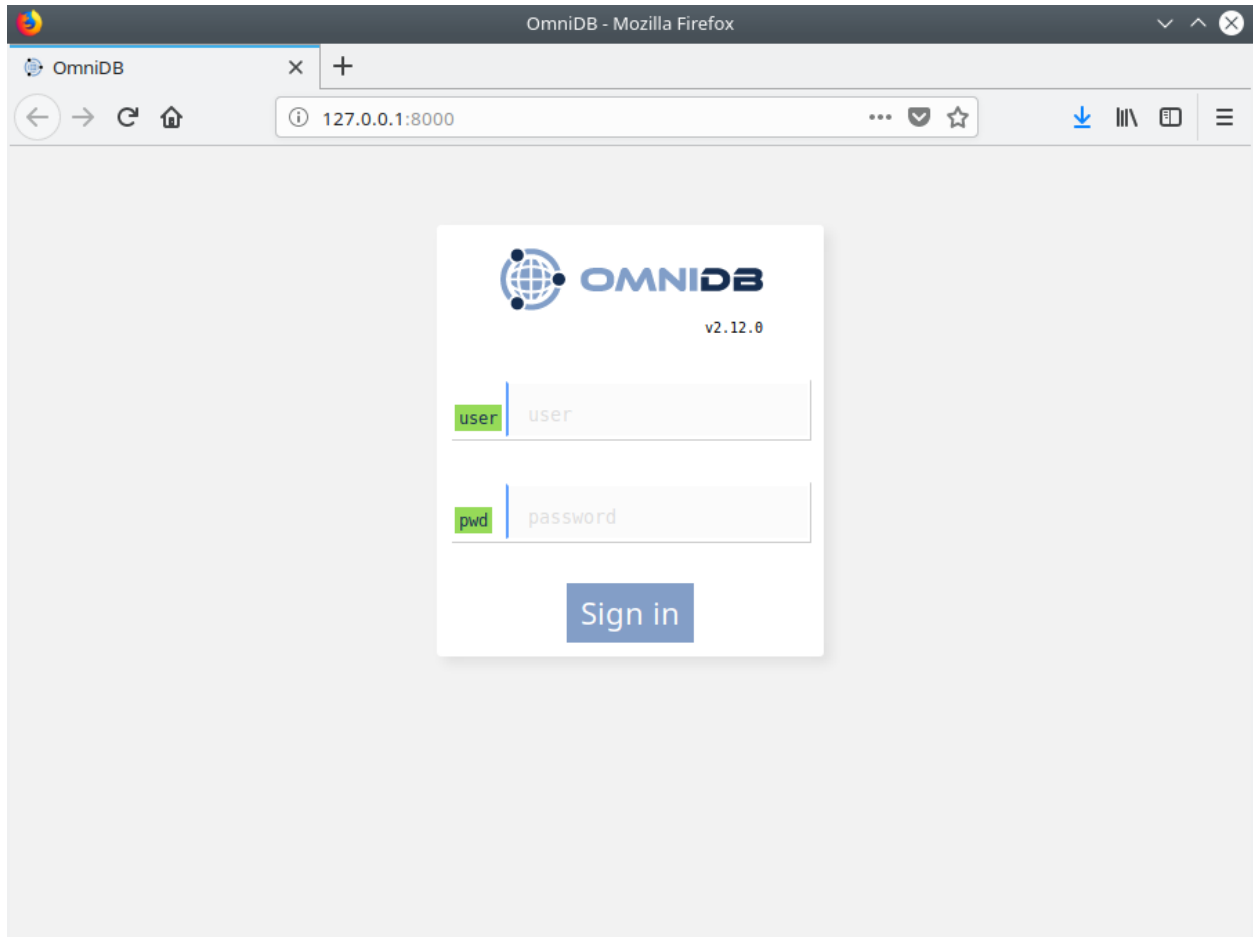
2.5 OmniDB configuration file

Starting on version 2.1.0, OmniDB server comes with a configuration file `omnidb.conf` that enables the user to specify parameters such as port and listening address. Also, 2.1.0 enables us to start the server with SSL, this requires a certificate and is configured in the same configuration file. For more details about how to deploy the OmniDB server, please read Chapter 19.

Starting on version 2.4.0, this file is located in `~/.omnidb/omnidb-server/omnidb.conf` in the user home directory.

2.6 OmniDB in the browser

Now that the web server is running, you may access OmniDB browser-based app on your favorite browser. Type in address bar: `localhost:8000` and hit Enter. If everything went fine, you shall see a page like this:



Now you know that OmniDB is running correctly. In the next chapters, we will see how to login for the first time, how to create an user and to utilize OmniDB.

3. Creating Users and Connections

3.1 Logging in as user *admin*

OmniDB comes only with the user *admin*. If you are using the server version, the first thing to do is sign in as *admin*, the default password is *admin*. You don't need to login in the app version.

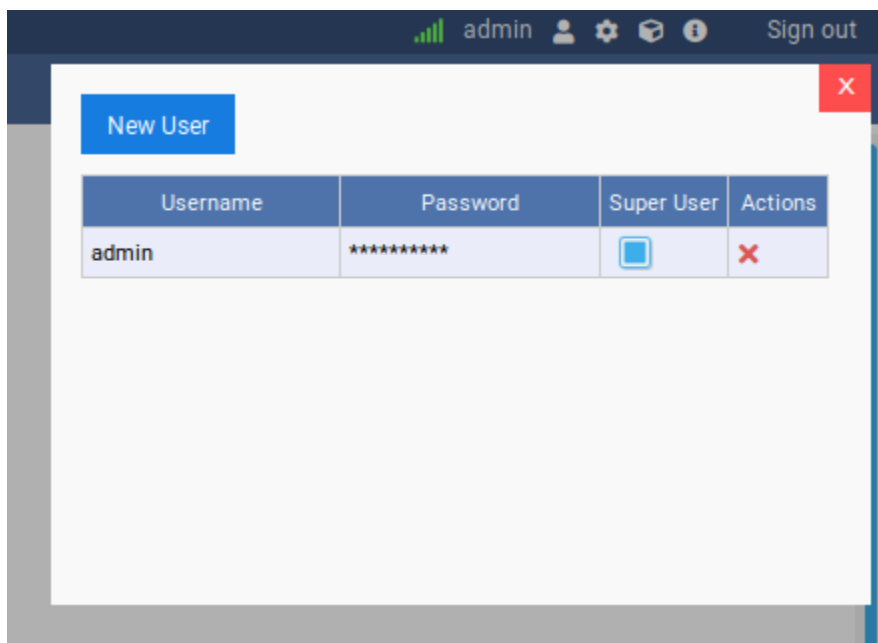


The next window is the initial window.

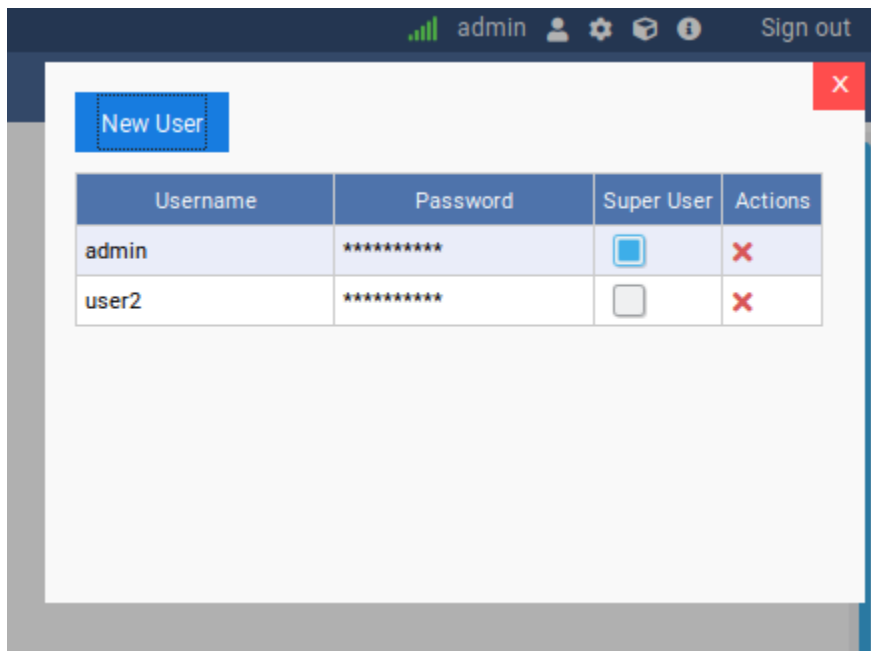


3.2 Creating another user

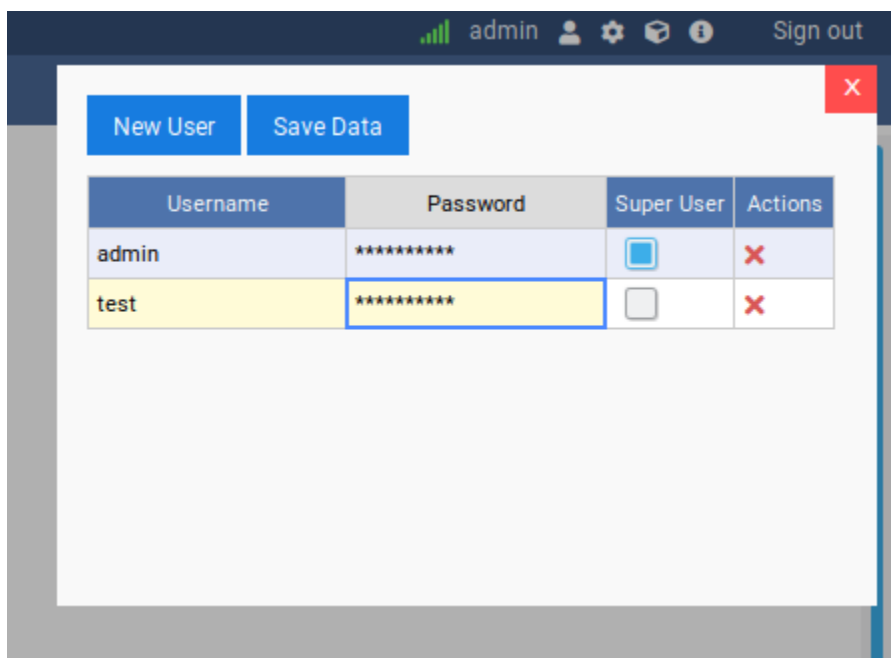
Click on the *Users* icon on the upper right corner. It will open a popup that allows the current OmniDB super user to create a new OmniDB user.



After clicking on the *Users* icon the tool inserts a new user called *user2* (if that is the first user after *admin*).



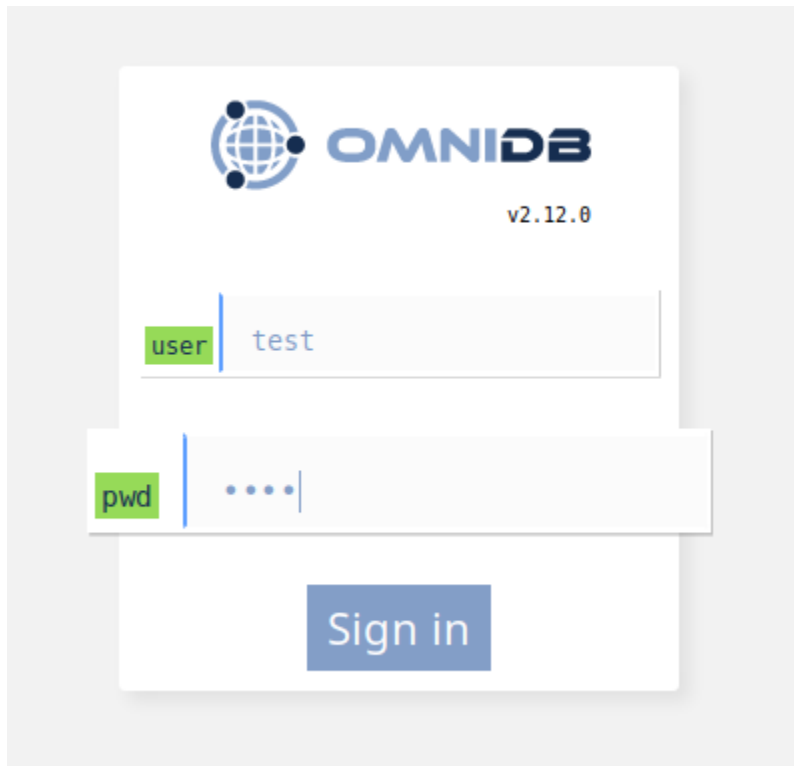
You will have to change the *username* and *password*. Check if you want this new user to be a *super user*. This user management window is only seen by super users. When you are done, click on the *Save Data* button inside the popup.



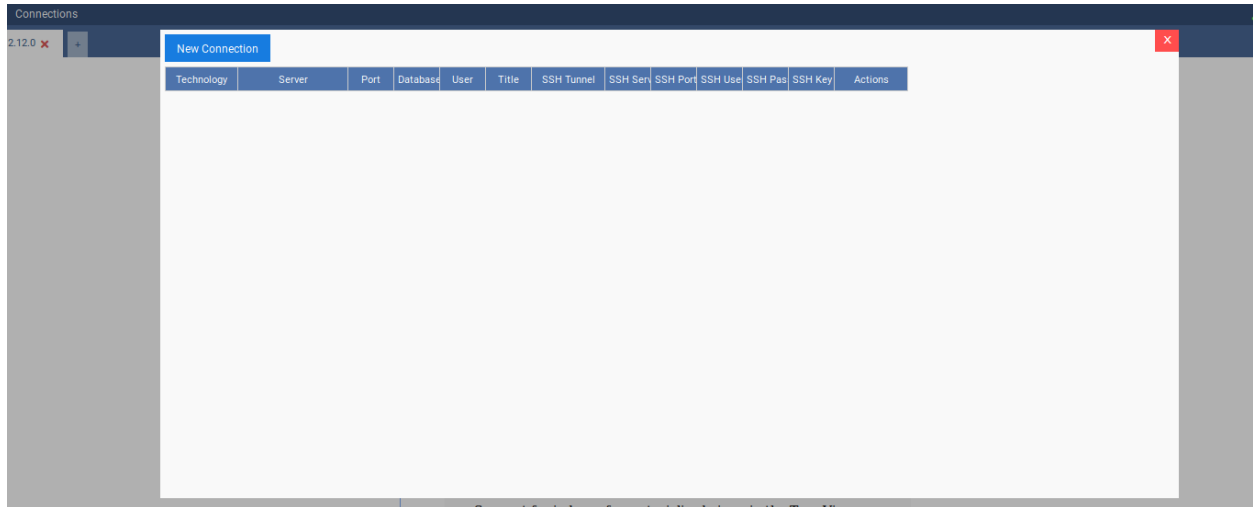
You can create as many users as you want, edit existing users and also delete users by clicking on the red cross at the actions column. Now you can logout by clicking in the *Sign Out* button in the top right corner.

3.3 Signing in as the new user

Let us sign in as the user we just created.



And we can see the window again. Note that now there is no *Users* icon, because the *test* user is not a super user. Go ahead and click on **Connections** on the upper left corner. You will see a popup like this:



3.4 Creating connections


At the moment, OmniDB supports PostgreSQL, Oracle, MySQL and MariaDB. More DBMS support is being added as you read this.

We will now create one connection to a PostgreSQL database, one connection to an Oracle database and one connection to a MariaDB database. To create the connections you have to click on the button *New Connection* and then choose the connection and fill the other fields. After filling all the fields for both connections, click on the *Save Data* button.

New Connection												
Technology	Server	Port	Database	User	Title	SSH Tunnel	SSH Server	SSH Port	SSH User	SSH Password	SSH Key	Actions
postgresql	127.0.0.1	5432	testdb	williamivanski	PostgreSQL	<input type="checkbox"/>		22				✗ ↕ ✓
oracle	127.0.0.1	1521	XE	SYSTEM	Oracle	<input type="checkbox"/>		22				✗ ↕ ✓
mariadb	127.0.0.1	3306	employees	root	MariaDB	<input type="checkbox"/>		22				✗ ↕ ✓

For each connection there is an *Actions* column where you can delete, test and select them. Go ahead and test the PostgreSQL connection.

Technology	Server	Port	Database	User	Title	SSH Tunnel	SSH Server	SSH Port	SSH User	SSH Password	SSH Key	Actions
postgresql	127.0.0.1	5432	testdb	williamivanski	PostgreSQL	<input type="checkbox"/>		22				✗ ↕ ✓
oracle	127.0.0.1	1521	XE	SYSTEM	Oracle	<input type="checkbox"/>		22				✗ ↕ ✓
mariadb	127.0.0.1	3306	employees	root	MariaDB	<input type="checkbox"/>		22				✗ ↕ ✓



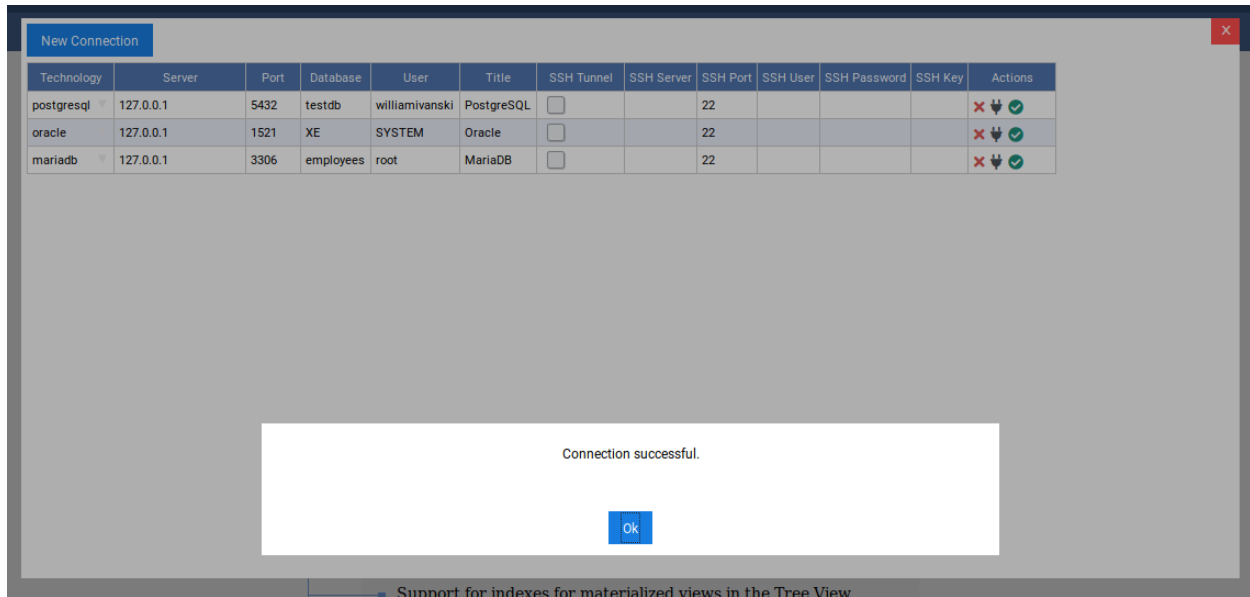
fe_sendauth: no password supplied

Password

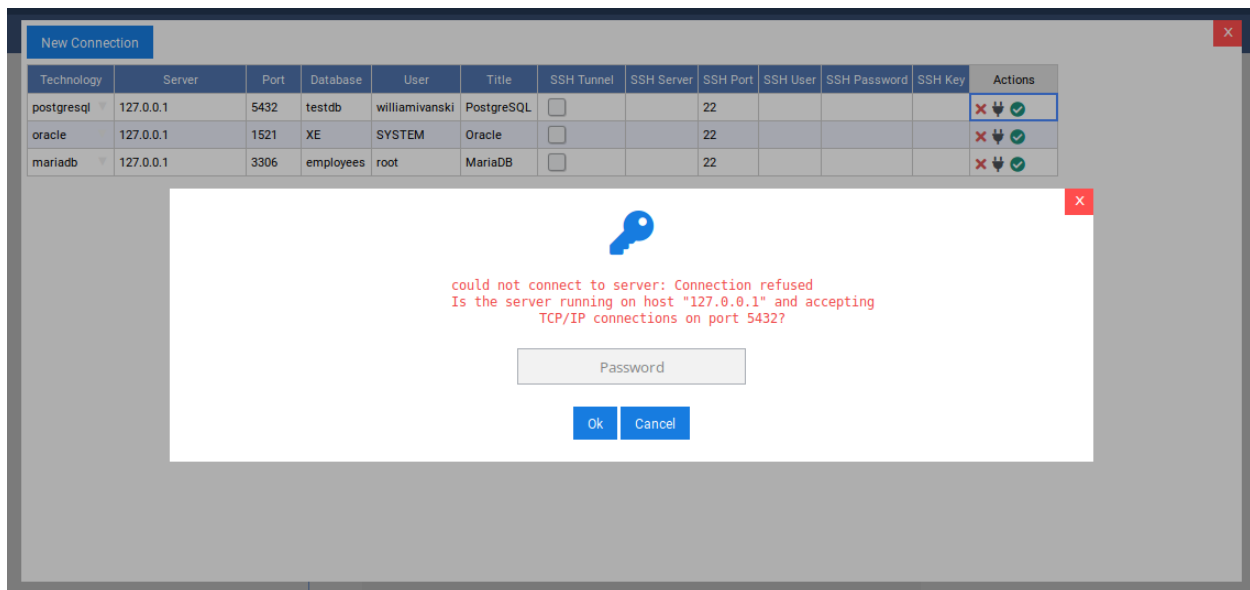
OK
Cancel

Notice a pop-up appears with the message *fe_sendauth: no password supplied*. This is happening because OmniDB does not store the database user password on disk. Not having any password at hand, OmniDB will try to connect without one, thus trying to take advantage of automatic authentication methods that might be in place: `trust` method, `.pgpass` file, and so on. As the database server replies with an error not allowing the user to connect, then OmniDB understands a password is required and asks it to the user. When the user types a password in this popup, the password is encrypted and stored in memory.

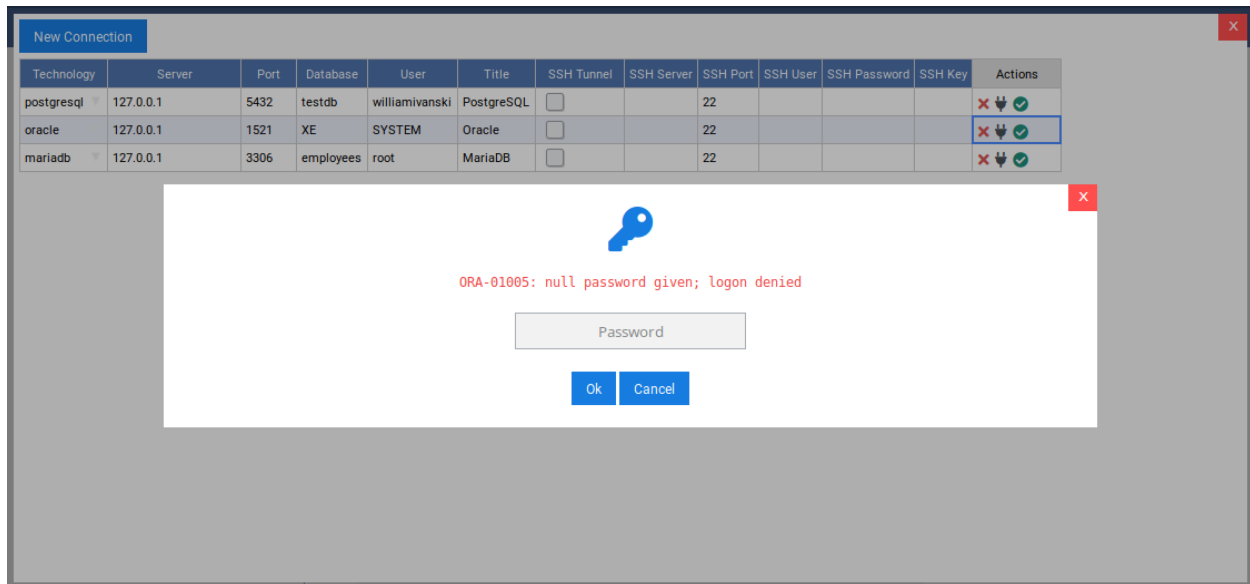
After you type the password and hit *Enter*, if the connection to the database is successful you will see a confirmation pop-up.



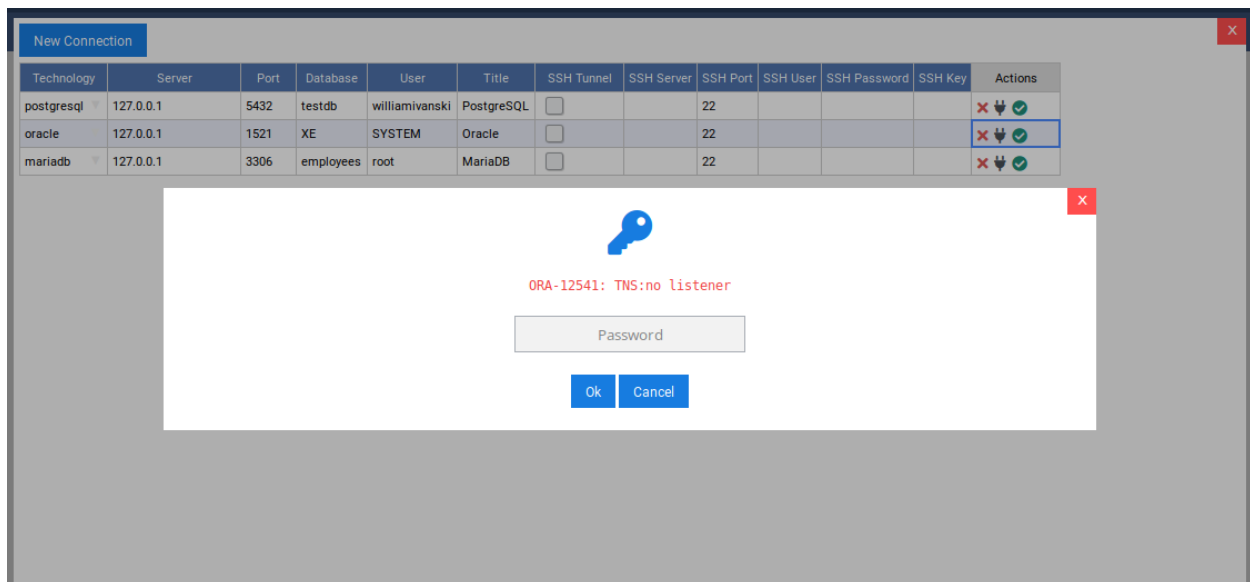
But, if you have trouble of any kind connecting to your PostgreSQL database, the same popup will remain showing the error OmniDB got.



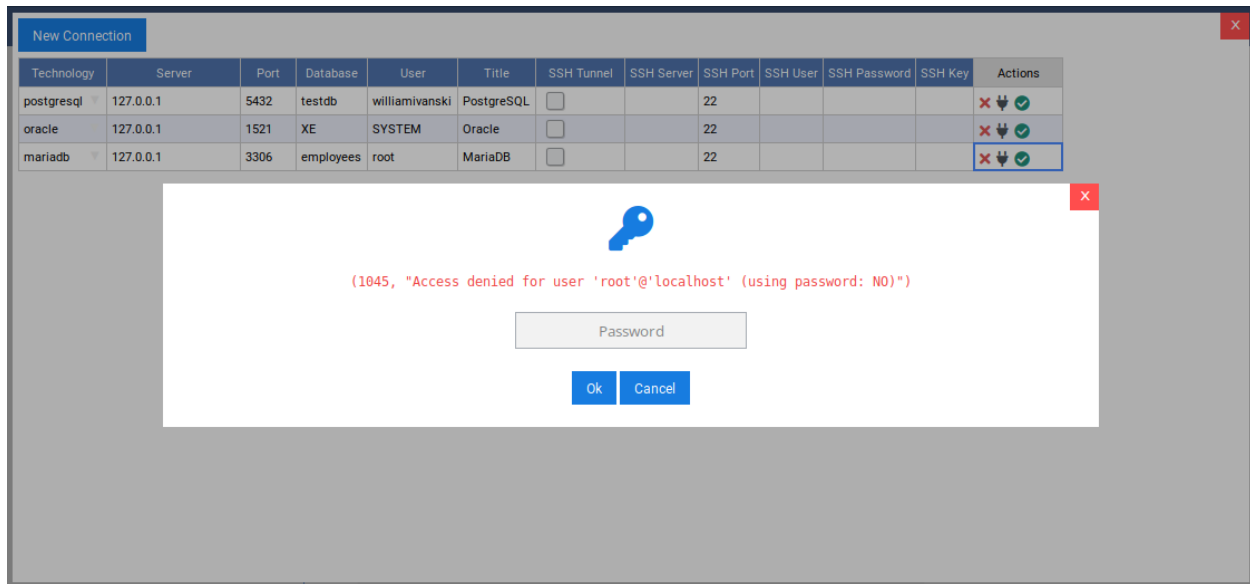
For Oracle, the behavior is similar. When OmniDB first tries to connect to an Oracle database without a password, you will see a message like this:



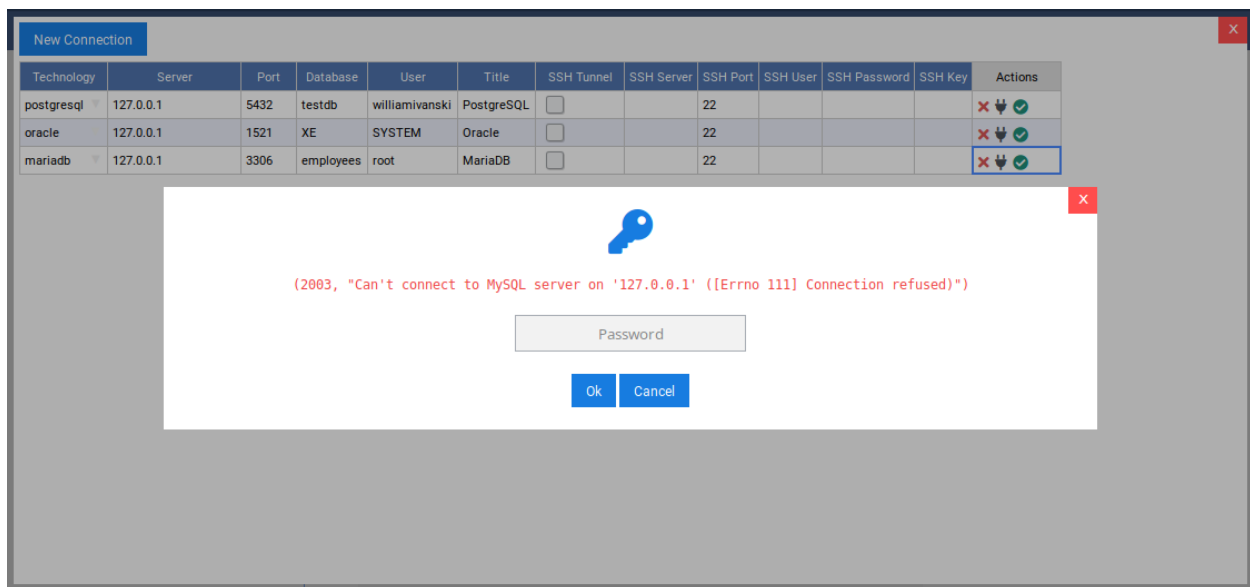
If you have any trouble connection to your Oracle database, the same popup will remain showing the error OmniDB got:



MariaDB and MySQL databases also works in the same way. First time, no password was given:



But if you have any problems, such as database server down:



Finally, in the connections grid, if you click on the *Select Connection* action, OmniDB will open it in a new **Connection Outer Tab** as we can see in the next chapter.

3.5 Using SSH tunnels

Starting from 2.8, OmniDB allows the user to connect to any remote database through SSH tunnels. The user needs to fill SSH tunnel information in each connection in the *Connections Grid*.

New Connection

X


Technology	Server	Port	Database	User	Title	SSH Tunnel	SSH Server	SSH Port	SSH User	SSH Password	SSH Key	Actions
postgresql	127.0.0.1	5432	testdb	williamivanski	PostgreSQL	<input type="checkbox"/>		22				✗ 🔑 ✓
oracle	127.0.0.1	1521	XE	SYSTEM	Oracle	<input type="checkbox"/>		22				✗ 🔑 ✓
mariadb	127.0.0.1	3306	employees	root	MariaDB	<input type="checkbox"/>		22				✗ 🔑 ✓
postgresql	127.0.0.1	5432	william	william	OmniDB	<input checked="" type="checkbox"/>	omnodb.org	22	omnodb	*****		✗ 🔑 ✓

- *SSH Server*: The server you are connecting to via SSH;
- *SSH Port*: The port of the SSH server (default is 22, but it can be any port number);
- *SSH User*: The operating system user name you use to connect to the SSH server;
- *SSH Password*: The password of the operating system user. If you fill the field *SSH Key*, then this is optional;
- *SSH Key*: The contents of the local private SSH key you can use to connect to the SSH server. If you fill this field, then you can also fill the field *SSH Password*, but in this case it will be the password for the SSH private key.

Please note that all information is stored encrypted in your local OmniDB *User Database*.


While using SSH tunnels, you also need to fill all database fields accordingly. But instead of being relative to the OmniDB server, they will be relative to the SSH Server. This can be done in 2 scenarios as explained below.

If the database is inside the same server as you are connecting to via SSH, then you will have a situation like this:



In this scenario, the database *Server* will be `127.0.0.1`, as the database is in the same machine as the *SSH Server*.

But the database can be outside the SSH server, like this:

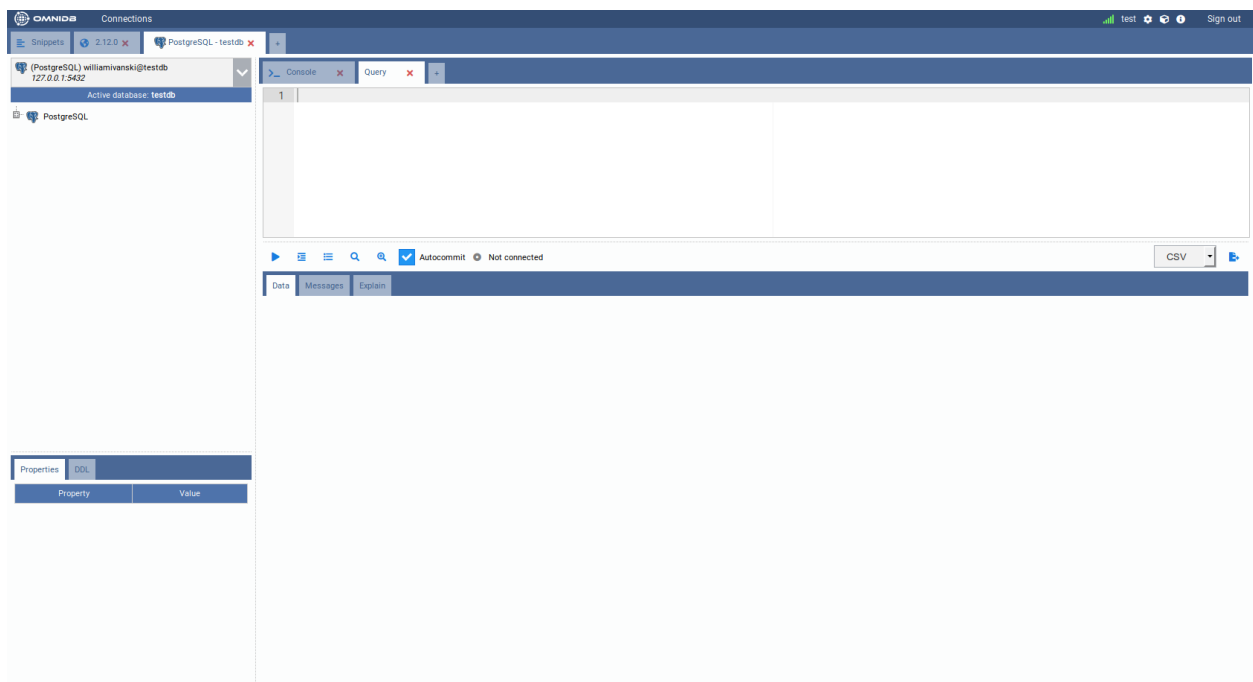


Here the database *Server* needs to be `192.168.0.10`, as it is the relative address for the SSH server to connect to the database server.

CHAPTER 4


4. Managing Databases

After creating a connection you can select it by clicking in the *Select Connection* action in the connections grid. You will see that the connection will be represented by a kind of outer tab called a *Connection Tab*. And this whole area is called the *Workspace Window*.



4.1 Sections of the *Workspace* window

This interface has several elements:

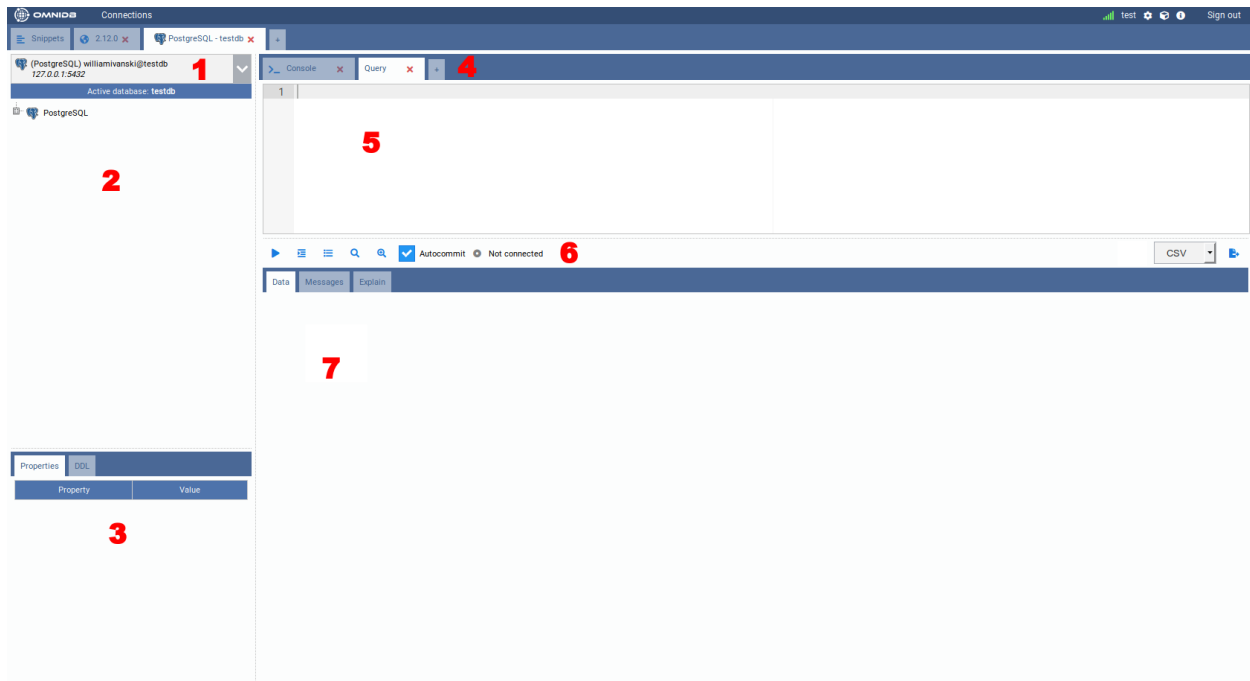


- **1) Connections:** Opens a popup with the *Connections* grid
- **2) Outer Tabs:** OmniDB lets you work with several databases at the same time. Each database will be accessible through an *outer tab*. Outer tabs also can host miscellaneous connection-independent features, like the *Snippets* feature
- **3) Options:** Shows the current user logged in, and if user is a superuser, also shows a link for *user management*. Also shows links for *user settings*, *installed plugins*, *query history*, *information* and *logout*.

4.2 Connection Outer Tab

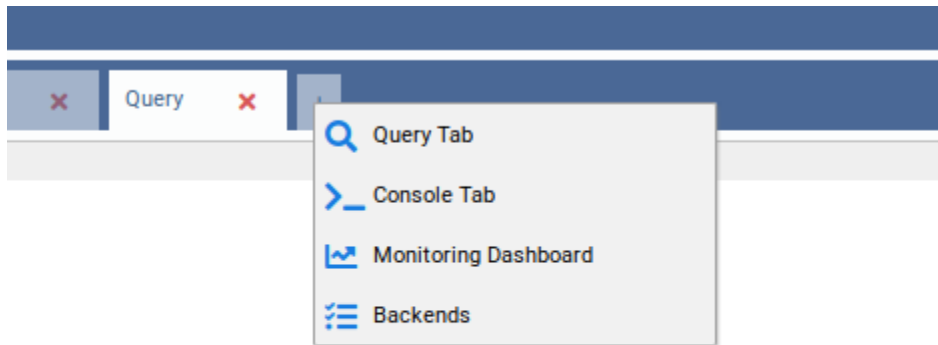
The outer table named *PostgreSQL - testdb* has this name because of the alias (*PostgreSQL*) we put in the connection to the *testdb* database. This tab is a *Connection Outer Tab*. Notice the little tab with a cross besides the *PostgreSQL - testdb* outer tab. This allows you to create a new outer tab that will automatically be a *Connection Outer Tab*. However, the *Snippet Outer Tab* is fixed and will always be the first.

A new *Connection Outer Tab* will always automatically point to the first connection on your list of database connections. Or, if you clicked on the *Select Connection* action, it will point to the selected connection. Observe the elements inside of this tab:



- **1) Connection Selector:** Shows all connections and lets the user select the current one
- **2) Tree of Structures:** Displays a hierarchical tree where you can navigate through the database elements
- **3) Properties and DDL Panels:** Display Properties and DDL about the currently selected node in the tree view
- **4) Inner Tabs:** Allows the user to execute actions in the current database. There are several kinds of inner tabs for the current database. By clicking on the last small tab with a cross, you can add a new tab. A new tab can be

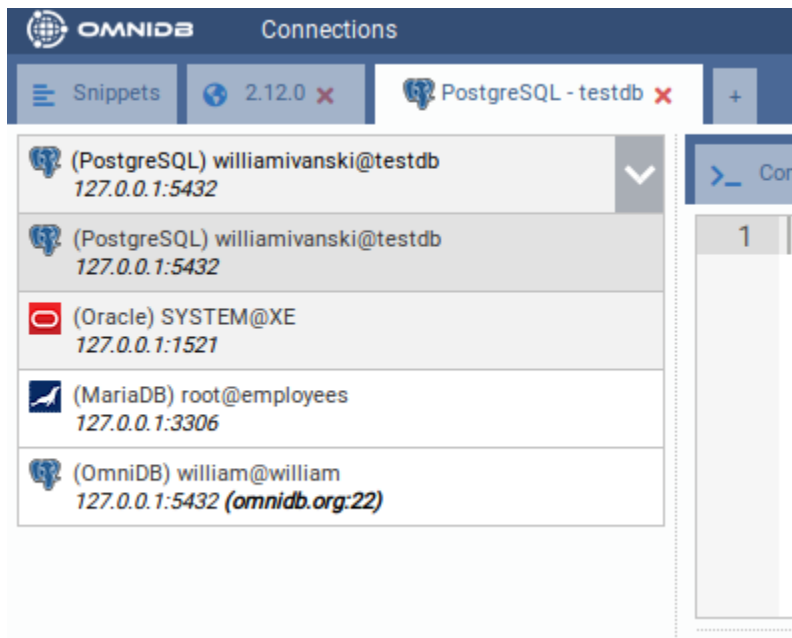
a *Query Tab*, *Console Tab*, *Monitoring Dashboard* or *Backends*



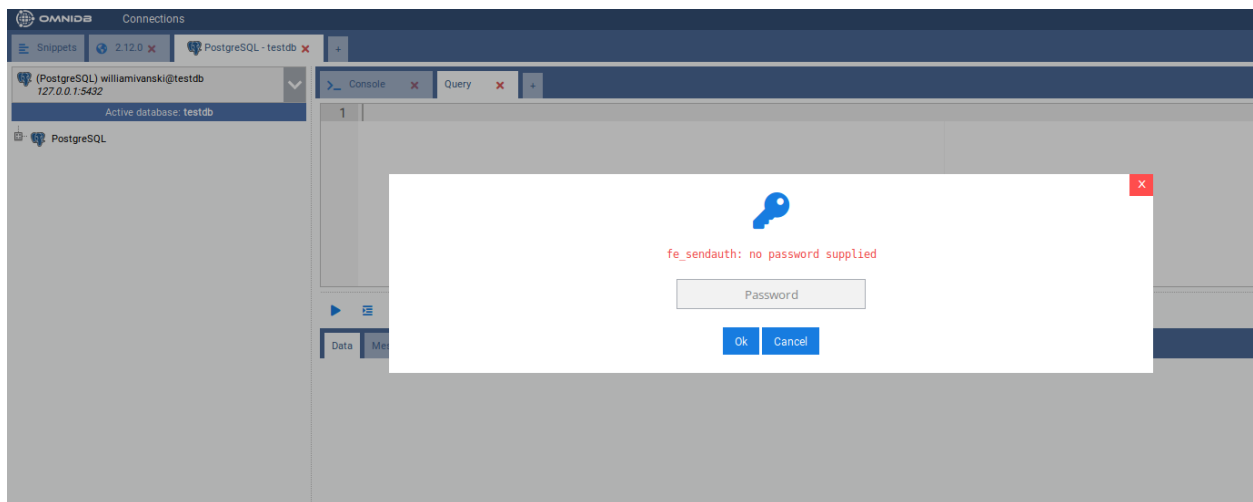
- **5) Inner Tab Content:** Can vary depending on the kind of inner tab. The figure shows a *Query Tab* and in this case the content will be an *SQL Editor*, with syntax highlight, autocomplete and find & replace
- **6) Inner Tab Actions:** Can vary depending on the kind of inner tab. For a *Query Tab*, they are *Run*, *Indent SQL*, *Command History*, *Explain*, *Explain Analyze*, *Autocommit* and *Export to File*
- **7) Inner Tab Results:** A *Query Tab*, after you click in the *Execute Button* or type the run shortcut (**Alt-Q**), will show a grid with the query results in the *Data* subtab. If the query calls a function that raises messages, those will be shown in the *Messages* subtab. If instead of *Run* you clicked in *Explain* or *Explain Analyze*, the explain plan for the query will be shown in the *Explain* subtab.

4.3 Working with databases

Take a look at your connections selector. OmniDB always points to the first available connection but you can change it by clicking on the selector.

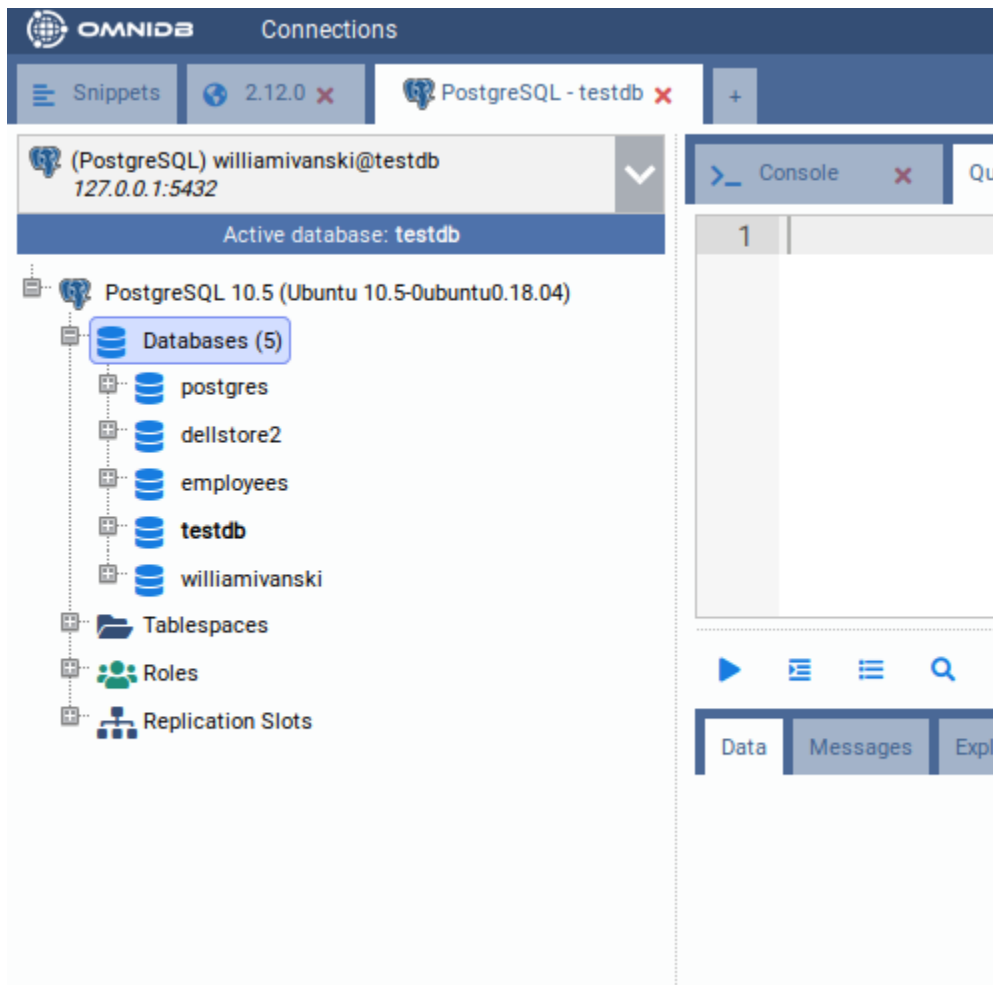


Select the *PostgreSQL* connection. Now go to the tree right below the selector and click to expand the root node *PostgreSQL*.

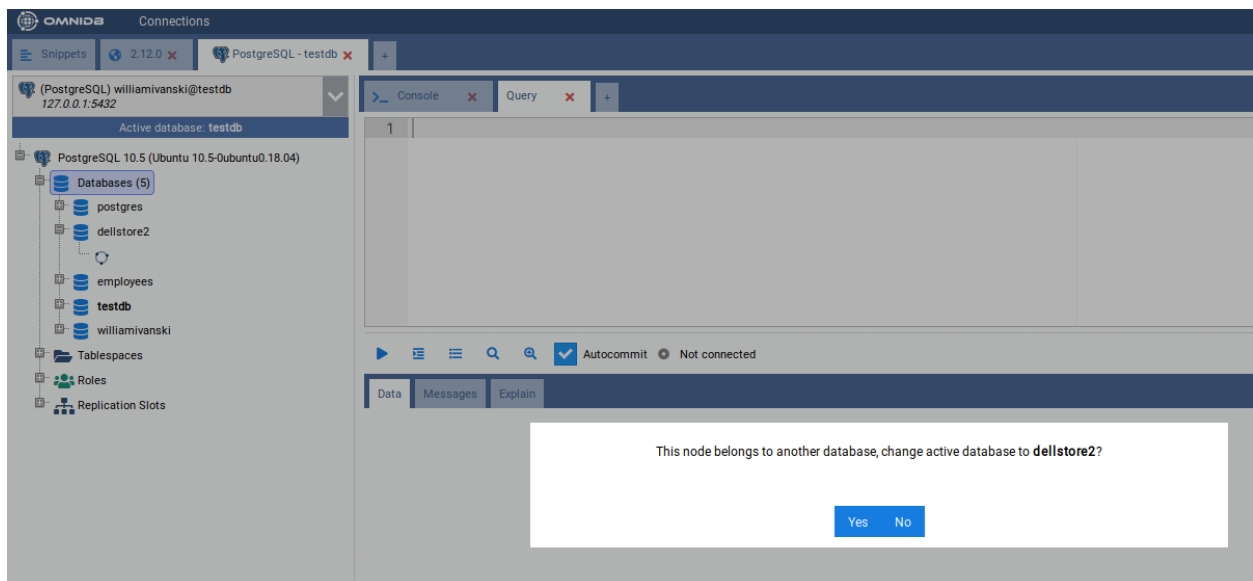


Bear in mind that every 30 minutes you keep without performing actions on the database, will trigger an *Authentication* popup, meaning that the password that OmniDB has encrypted and stored in memory is now expired. As explained before, this is important for your database security. After you type the correct password, you will see the PostgreSQL node now shows the PostgreSQL version and also was expanded, showing the current database connection and also instance wide elements: *Databases*, *Tablespaces*, *Roles* and *Replication Slots*.

You can connect to a single PostgreSQL database, and using the same connection you can connect to other databases in the same PostgreSQL instance. The currently active database will be indicated below the connection selector.

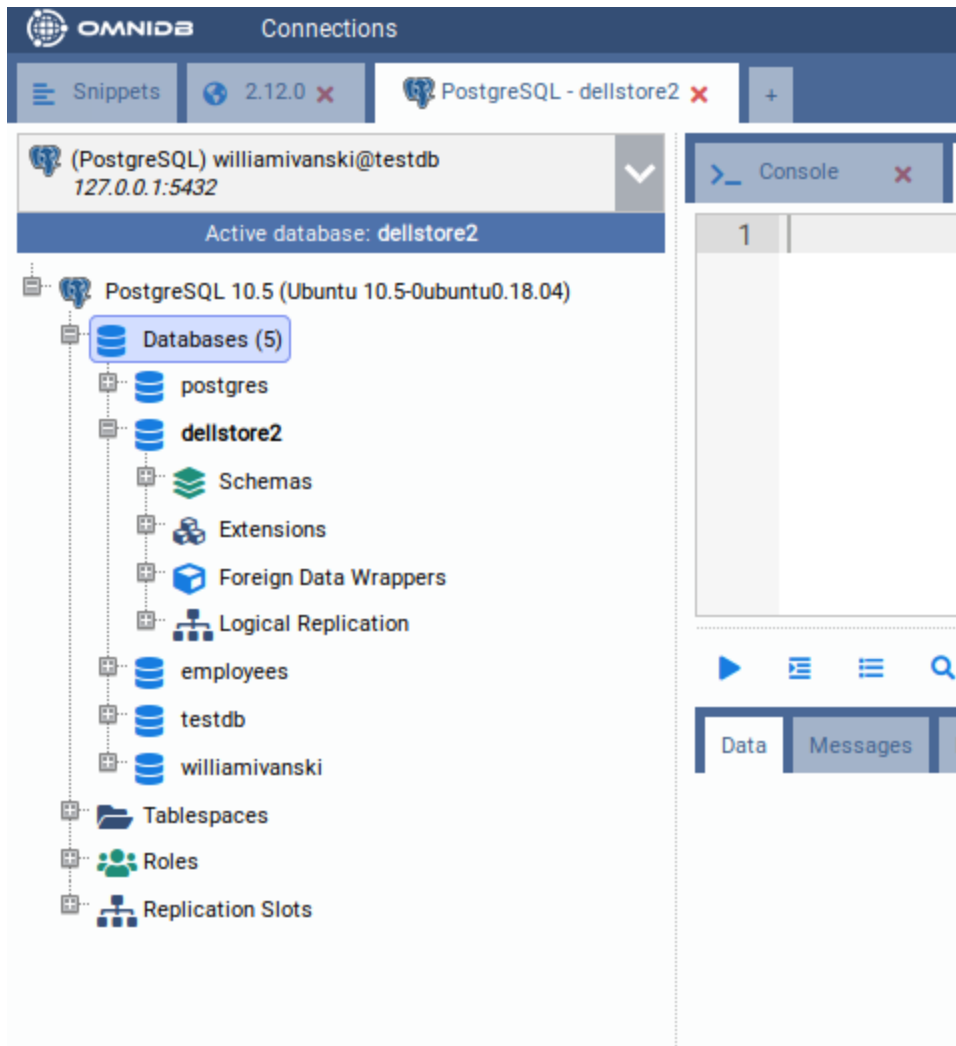


To connect to a different database, expand the node corresponding to that database. A popup will appear asking if you really want to change the active database.

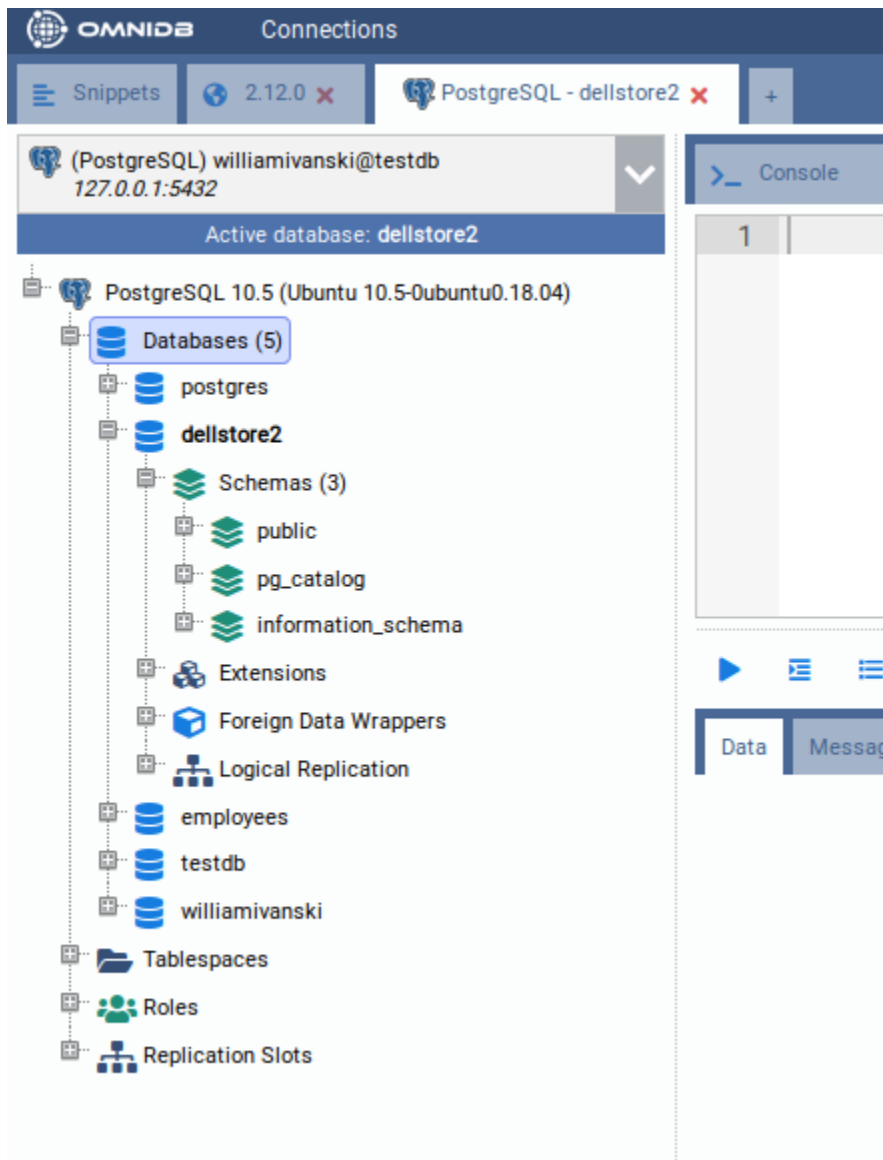


Click on *Yes* and OmniDB will change the active database to the database you choose. It will be reflected on the *Active*

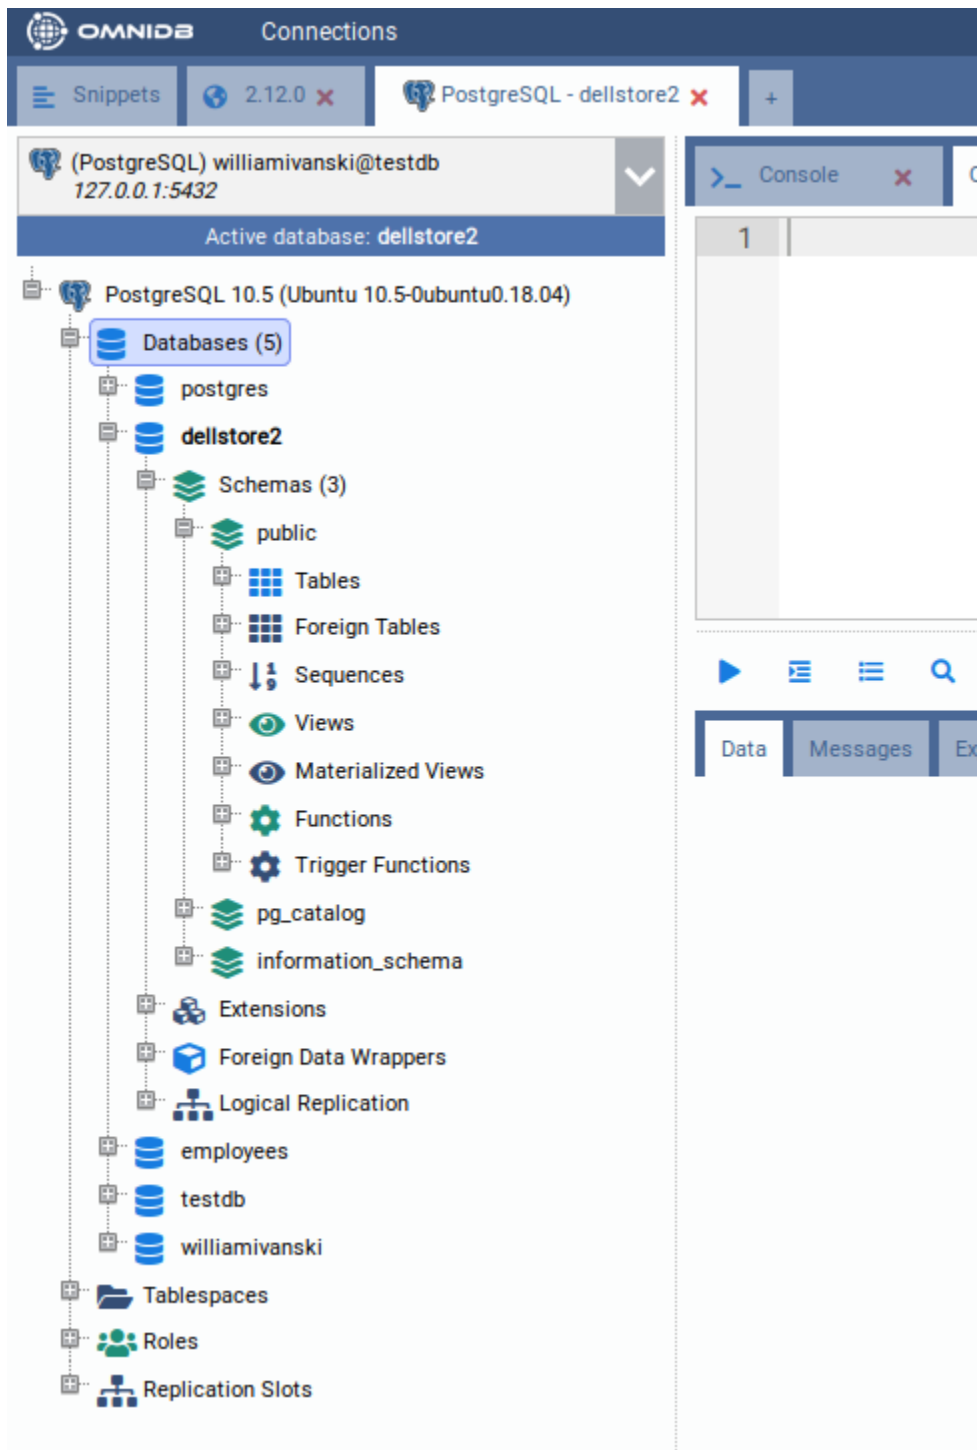
database indicator, and also on the outer tab name.



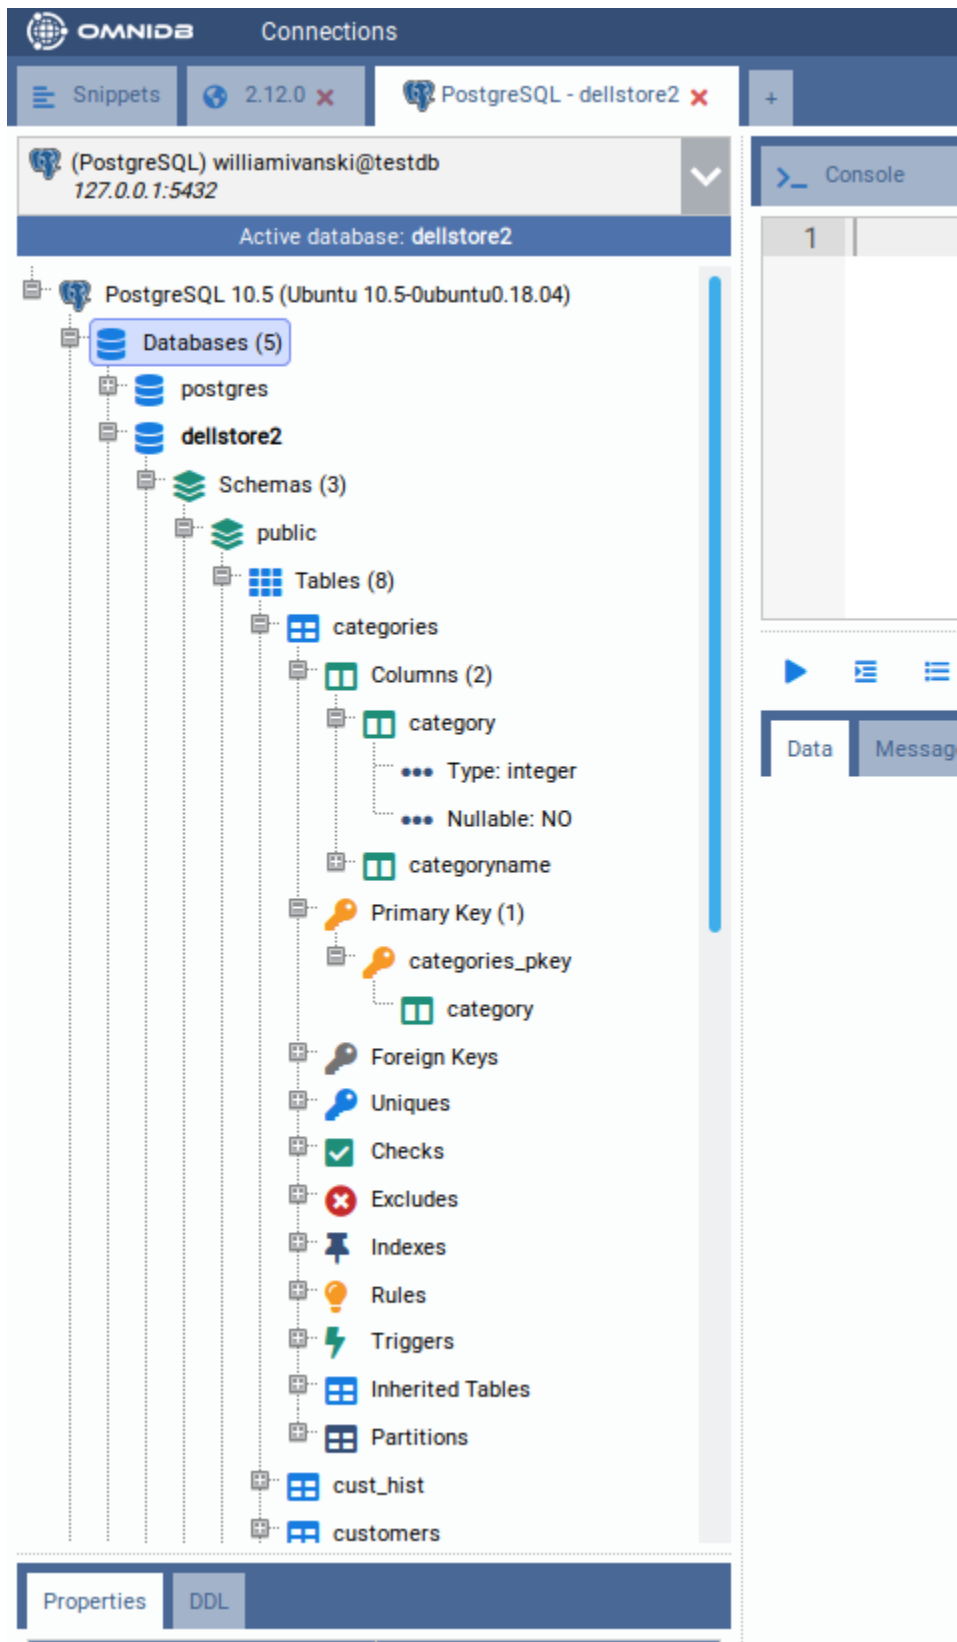
Go ahead and expand the *Schemas* node. You will see all schemas in the current database (in case of PostgreSQL, TOAST and temp schemas are not shown).



Now click to expand the schema `public`. You will see different kinds of elements contained in this schema.

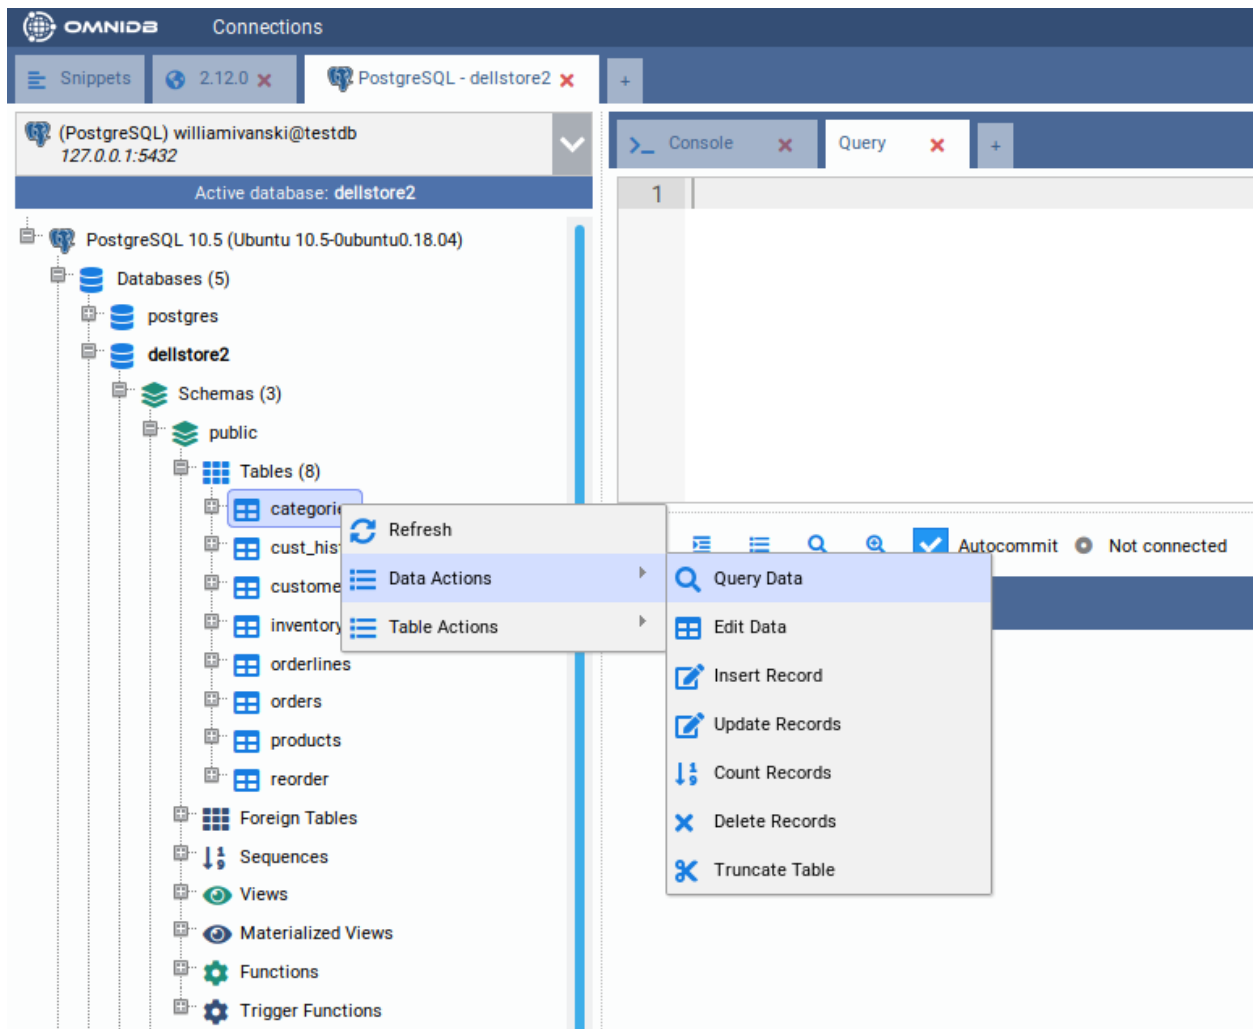


Now click to expand the node *Tables*, and you will see all tables contained in the schema `public`. Expand any table and you will see its columns, primary key, foreign keys, constraints, indexes, rules, triggers and partitions.



In order to view records inside a table, right click it and choose *Data Actions

Query Data*.



Notice that OmniDB opens a new SQL editor with a simple query to list table records. The records are displayed in a grid right below the editor. This grid can be controlled with keyboard as if you were using a spreadsheet manager. You can also copy data from single cells or block of cells (that can be selected with the keyboard or mouse) and paste on any spreadsheet manager.

PostgreSQL 10.5 (Ubuntu 10.5-0ubuntu0.18.04)

Databases (5)

- postgres
- dellstore2
 - Schemas (3)
 - public
 - Tables (8)
 - categories
 - cust_hist
 - customers
 - inventory
 - orderlines
 - orders
 - products
 - reorder
 - Foreign Tables
 - Sequences
 - Views
 - Materialized Views
 - Functions
 - Trigger Functions
 - pg_catalog
 - information_schema
 - Extensions
 - Foreign Data Wrappers
 - Logical Replication
 - employees

Active database: dellstore2

```

1 SELECT t.category
2       , t.categoryname
3 FROM public.categories t
4 ORDER BY t.category

```

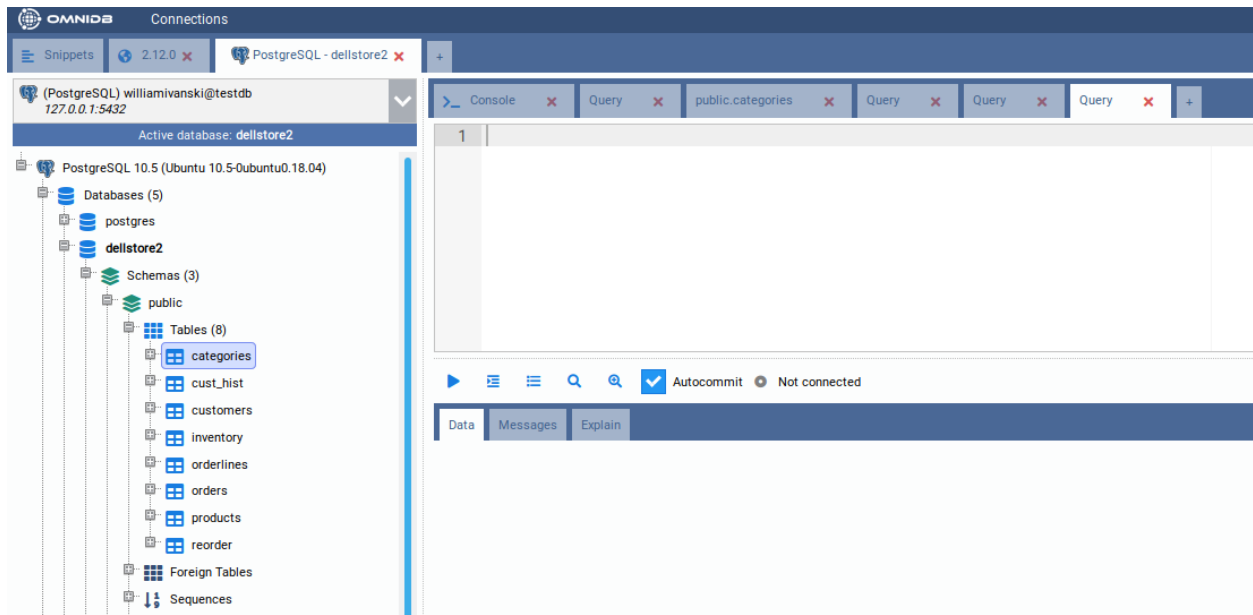
Autocommit Idle Number of records: 16 Start time: 10/21/2018 15:43:58 Duration: 40.141 ms

	category	categoryname
1	1	Action
2	2	Animation
3	3	Children
4	4	Classics
5	5	Comedy
6	6	Documentary
7	7	Drama
8	8	Family
9	9	Foreign
10	10	Games
11	11	Horror
12	12	Music
13	13	New
14	14	Sci-Fi
15	15	Sports
16	16	Travel

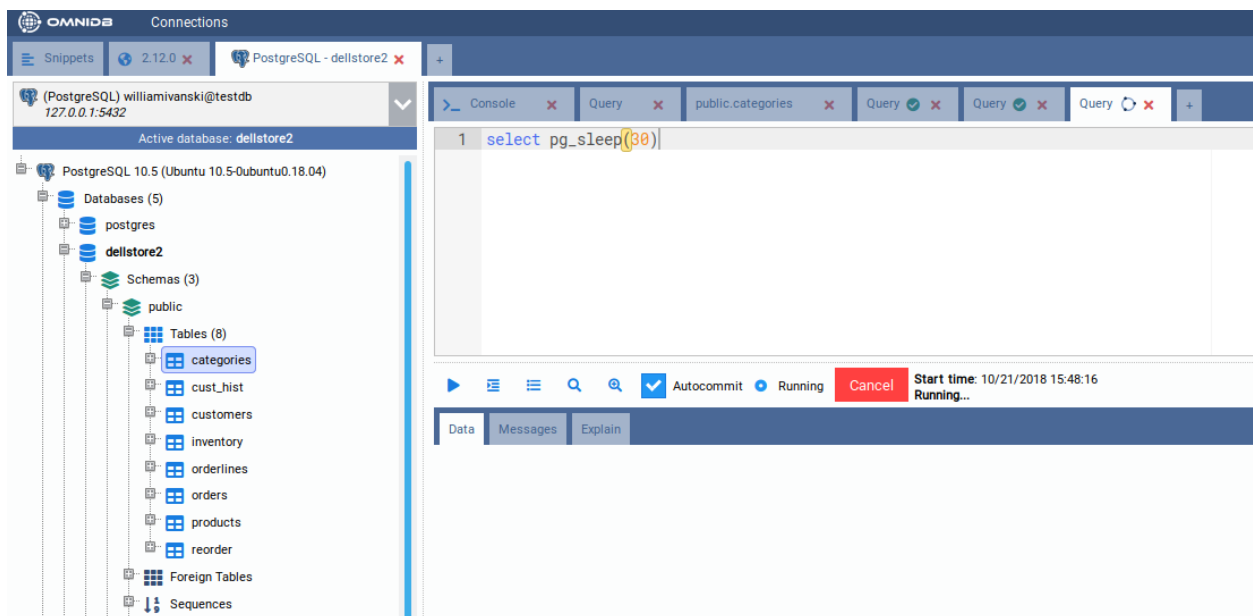
You can edit the query on the SQL editor, writing simple or more complex queries. To execute, click on the action button or hit the keystroke `Alt-Q`. If the results exceed 50 registers, then extra buttons *Fetch More* and *Fetch All* will appear. More details in the next chapters.

4.4 Working with multiple tabs inside the same connection

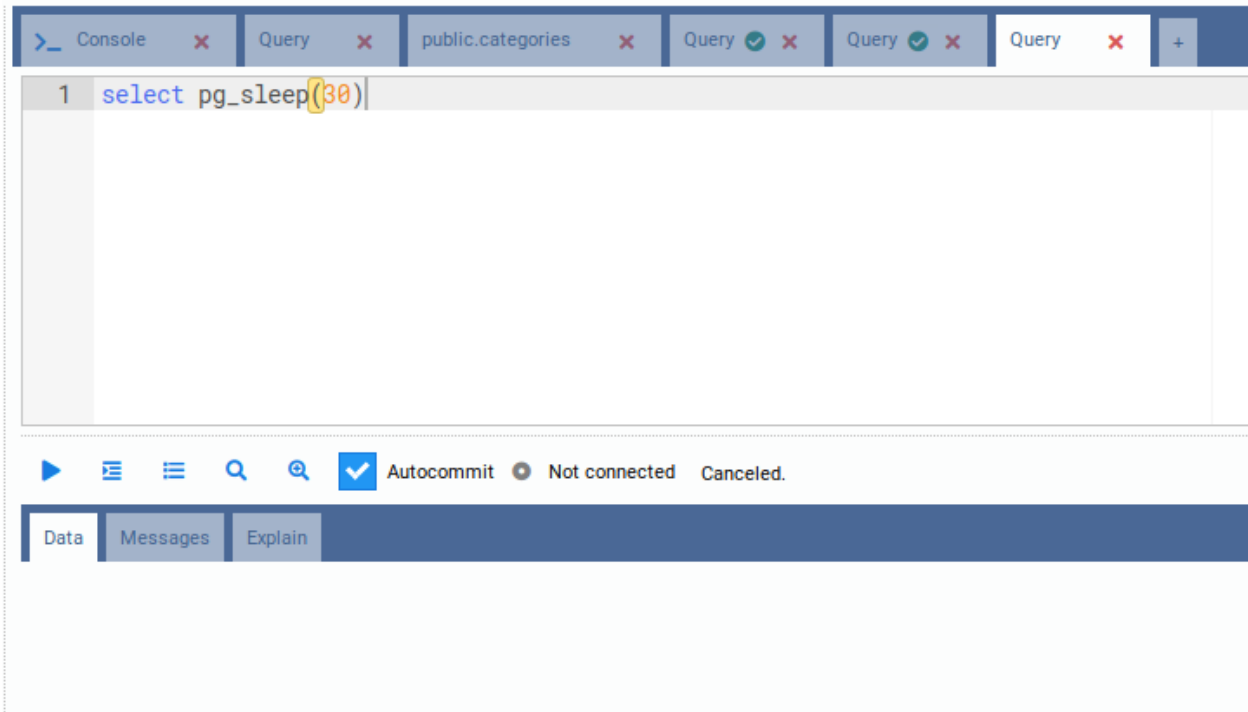
Inside a single connection, you can create several inner query tabs by clicking on the last little tab with a cross, and then choosing *Query Tab*.



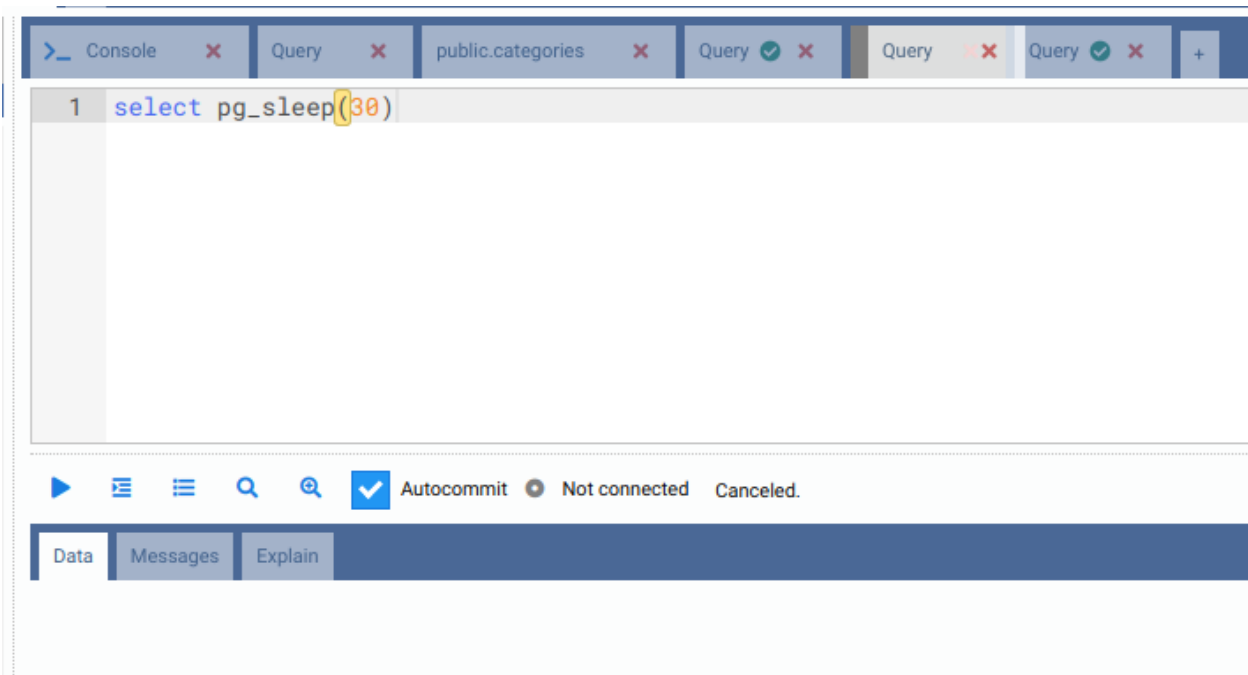
On OmniDB, you can execute several SQL statements and procedures in parallel. When it is executing, an icon will be shown in the tab to indicate its current state. If some process is finished and it is not in the current tab, that tab will show a green icon indicating the routine being executed there is now finished.



By clicking in the *Cancel button*, you can cancel a process running inside the database.



You can also drag and drop a tab to change its order. This works with both inner and outer tabs.



Additionally, you can use keyboard shortcuts to manage inner tabs (SQL Query) and outer tabs (Connection):

- **Ctrl-Insert:** Insert a new inner tab
- **Ctrl-Delete:** Removes an inner tab
- **Ctrl-<:** Change focus to inner tab at left
- **Ctrl->:** Change focus to inner tab at right

- **Ctrl-Shift-Insert:** Insert a new outer tab
- **Ctrl-Shift-Delete:** Removes an outer tab
- **Ctrl-Shift-<:** Change focus to outer tab at left
- **Ctrl-Shift->:** Change focus to outer tab at right

Starting from OmniDB version 2.3.0, all SQL Query tabs are automatically saved whenever you execute them. Even if you close OmniDB window or browser tab, they are already stored in OmniDB *User Database*. They will be automatically restored when you open OmniDB again (if you are using app), open it in another browser window (if you are using server), or even if you clicked in the *Connections* window or logged out. Removing an outer tab or inner tab by the interface makes it permanently deleted, so it will not be restored.

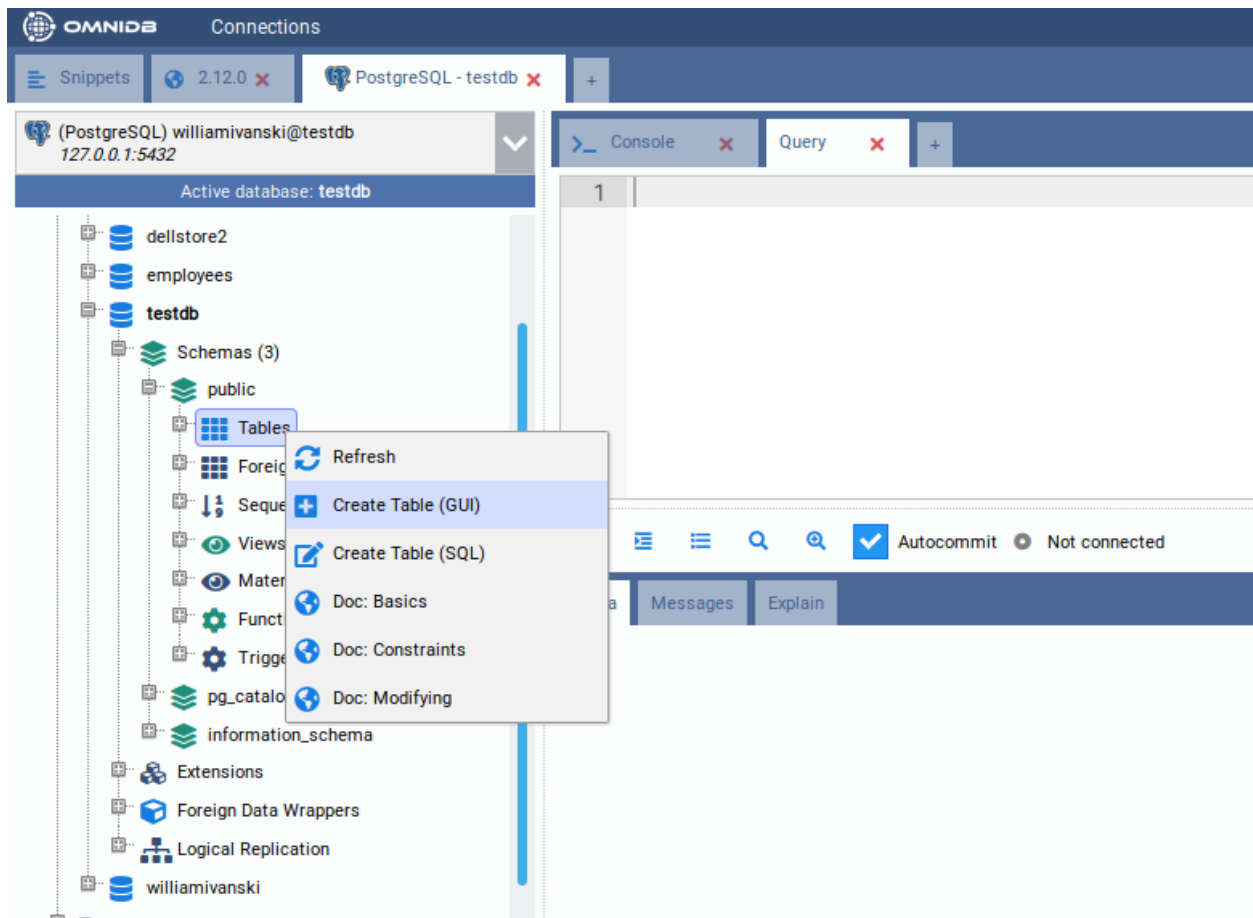
5. Creating, Changing and Removing Tables

5.1 Creating tables

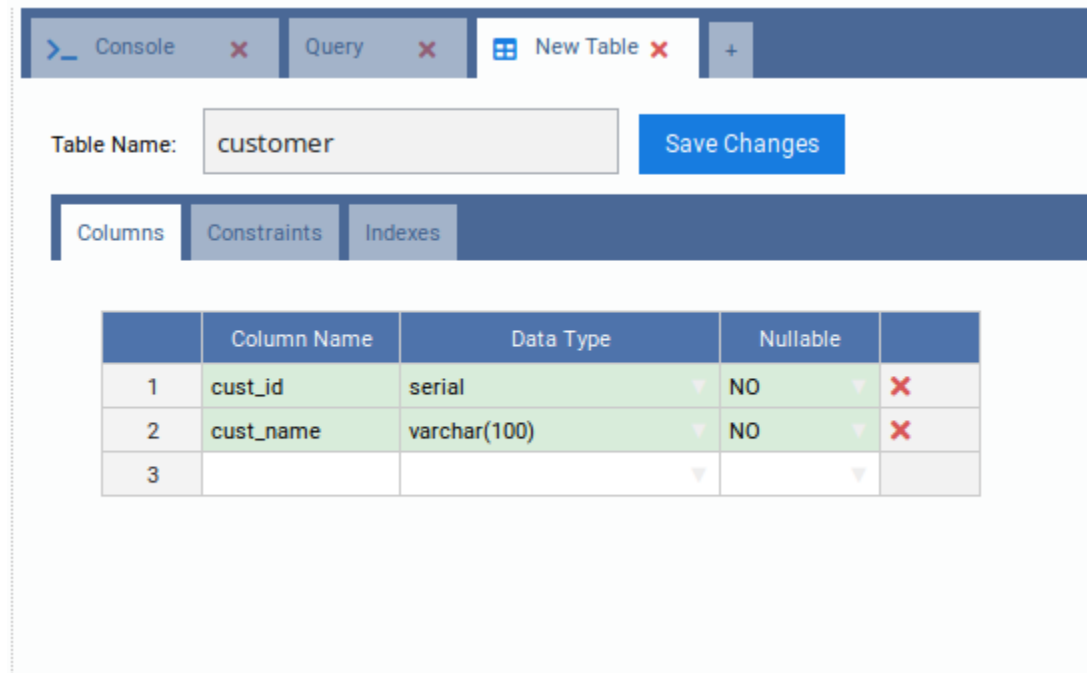
OmniDB has a table creation interface that lets you configure columns, constraints and indexes. A couple of observations should be mentioned:

- Most DBMS automatically create indexes when primary keys and unique constraints are created. Because of that, the indexes tab is only available after creating the table.
- Each DBMS has its unique characteristics and limitations regarding table creation and the OmniDB interface reflects these limitations. For instance, SQLite does not allow us to change existing columns and constraints. Because of that, the interface lets us change only table name and add new columns when dealing with SQLite databases (it is still not the case in OmniDB Python version, as it currently supports only PostgreSQL databases).

We will create example tables (*customers* and *addresses*) in the `testdb` database we connected to earlier. Right click on the **Tables** node and select the **Create Table (GUI)** action:

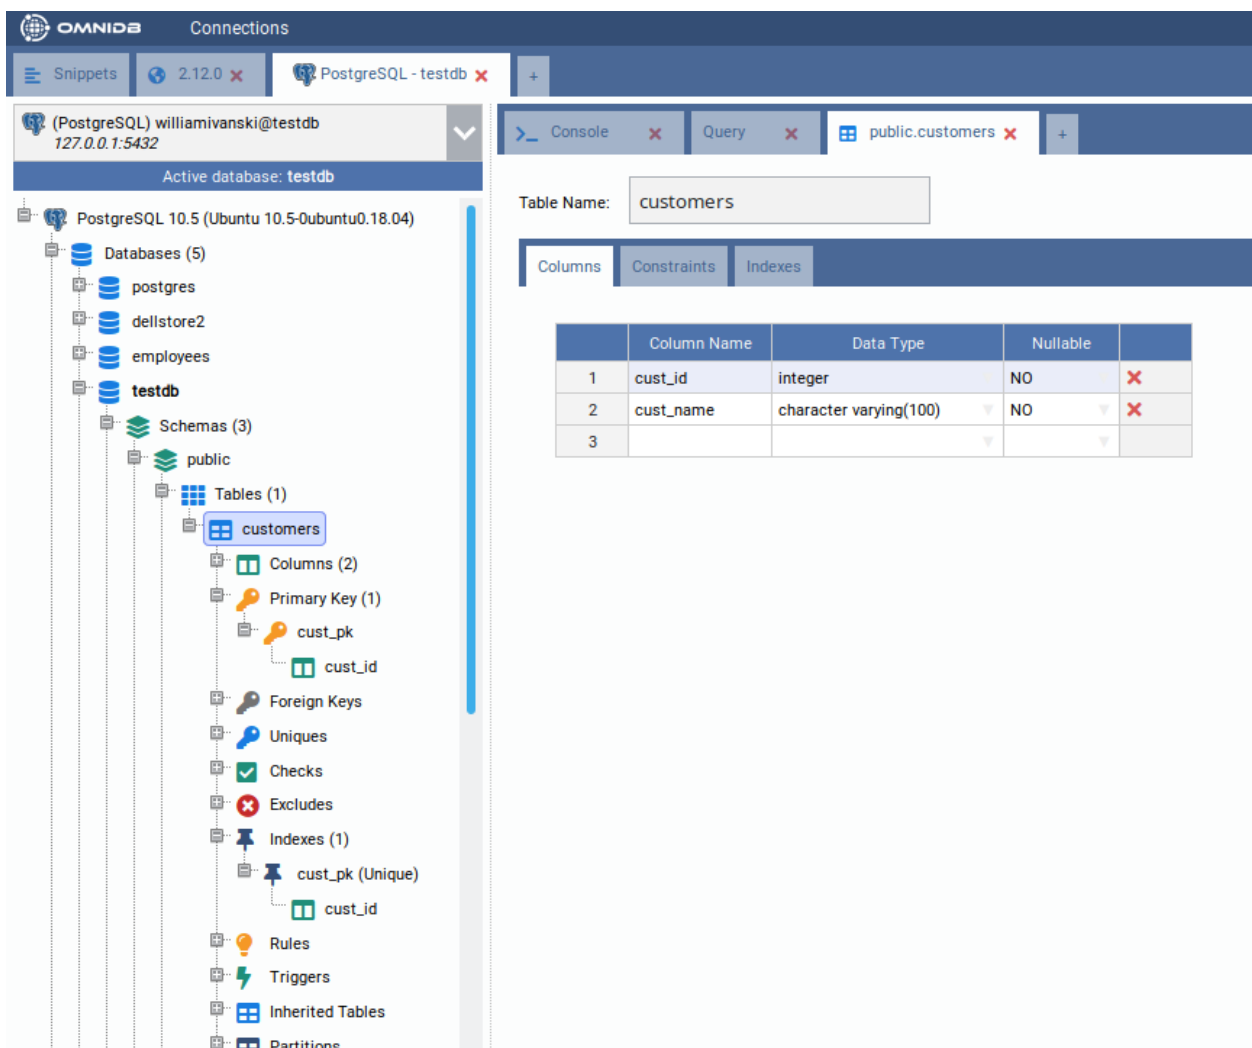


We will create the table *customers* with a primary key that will be referenced by the table *addresses*:





Click on the *Save Changes* button. Right-click the *Tables* tree node and click *Refresh*. Note how the table appears in the *Tables* tree node:



By keeping the table *customers* selected in the treeview, check its properties and DDL:

The screenshot displays the OmniDB interface for a PostgreSQL database named 'testdb'. The top bar shows the connection details: (PostgreSQL) williamivanski@testdb, 127.0.0.1:5432. The active database is 'testdb'. On the left, a tree view shows the database structure: public > Tables (1) > customers. Below this, the 'Properties' tab is selected, showing a table of properties for the 'customers' table. To the right, the 'Table Name' is set to 'customers', and the 'Columns' tab is active, displaying the table's schema.

Property	Value
Database	testdb
Schema	public
Table	customers
OID	64989
Owner	williamivanski
Size	0 bytes
Tablespace	pg_default
ACL	
Options	
Filenode	base/64978/64989
Estimate Count	0
Has Index	true
Persistence	Permanent
Number of Attributes	2
Number of Checks	0
Has OIDs	false
Has Primary Key	true
Has Rules	false
Has Triggers	false
Has Subclass	false
Is Partitioned	false
Is Partition	false
Partition Of	

	Column Name	Data Type	Nullable	
1	cust_id	integer	NO	✗
2	cust_name	character varying(100)	NO	✗
3				

The screenshot shows the OmniDB interface with a PostgreSQL connection to a database named 'testdb'. The left sidebar displays a tree view of the database structure, including schemas (public), tables (customers), and other database objects. The main panel shows the DDL (Data Definition Language) for the 'customers' table, including its creation and subsequent alterations.

```

1  --
2  -- Type: TABLE ; Name: customers; Owner: williamivanski
3  --
4
5  CREATE TABLE customers (
6      cust_id integer NOT NULL,
7      cust_name character varying(100) NOT NULL
8  );
9
10
11 ALTER TABLE public.customers ALTER cust_id SET DEFAULT nextval('customers_cust_id_seq'::regclass);
12
13 ALTER TABLE customers ADD CONSTRAINT cust_pk
14     PRIMARY KEY (cust_id);
15
16 ALTER TABLE customers OWNER TO williamivanski;
17
18

```

Now create the table *addresses* with a primary key and a foreign key:

Console
Query
public.customers
New Table
+

Table Name:

Columns
Constraints
Indexes

	Column Name	Data Type	Nullable	
1	add_id	serial	NO	✗
2	add_street	varchar(200)	NO	✗
3	add_number	integer	YES	✗
4	cust_id	integer	NO	✗
5				

Console
Query
public.customers
New Table
+

Table Name:

Columns
Constraints
Indexes

Constraint Name	Type	Columns	Referenced Table	Referenced Columns	Delete Rule	Update Rule	
add_pk	Primary Key	add_id					✗
add_fk1	Foreign Key	cust_id	public.customers	cust_id	CASCADE	CASCADE	✗

Don't forget to click on the *Save Changes* button when done. At this point we have two tables in schema `public`. The schema structure can be seen with the graph feature by right clicking on the schema `public` node of the tree and selecting *Render Graph > Simple Graph*:

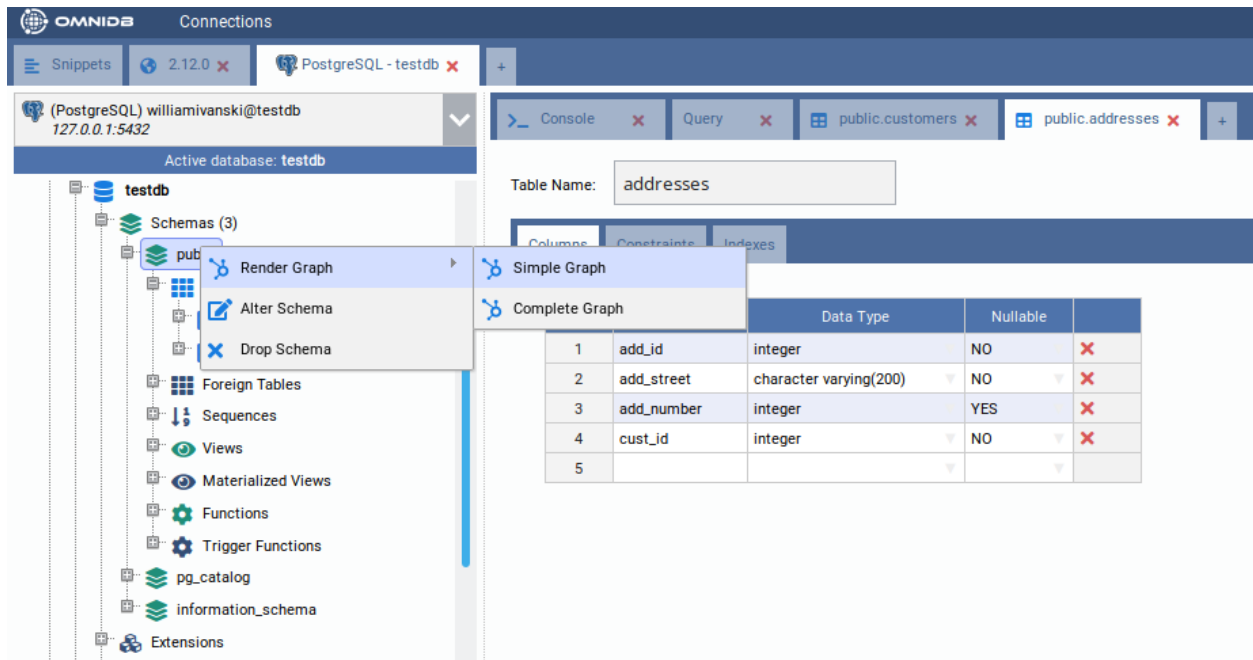


Table Name: addresses

		Data Type	Nullable	
1	add_id	integer	NO	✓
2	add_street	character varying(200)	NO	✓
3	add_number	integer	YES	✓
4	cust_id	integer	NO	✓
5				

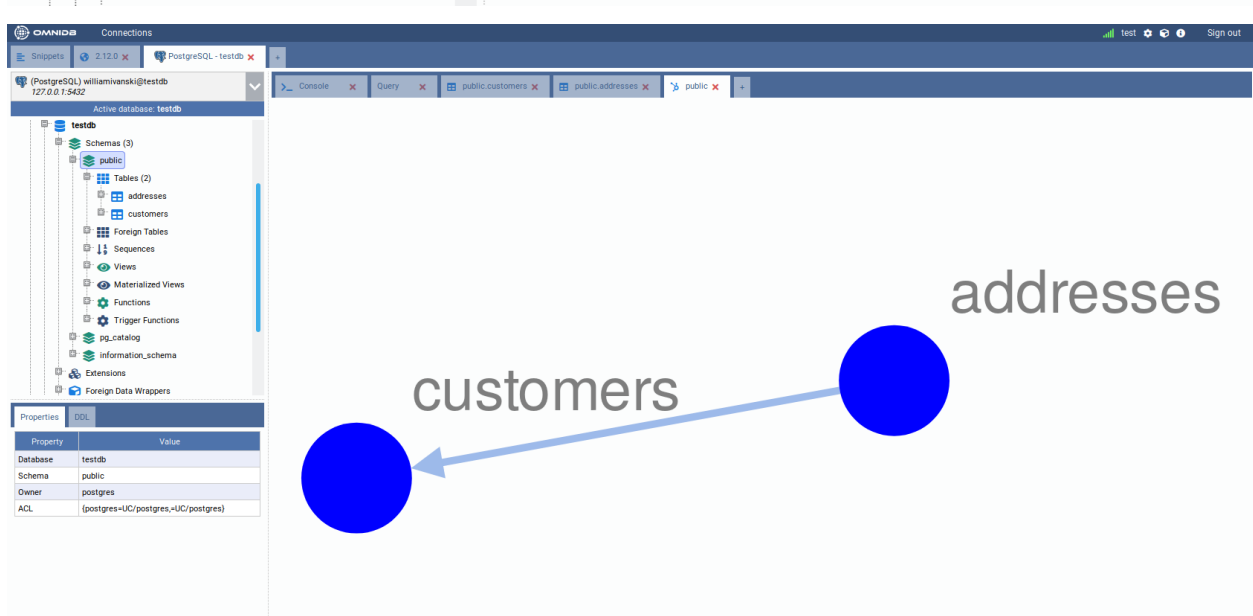
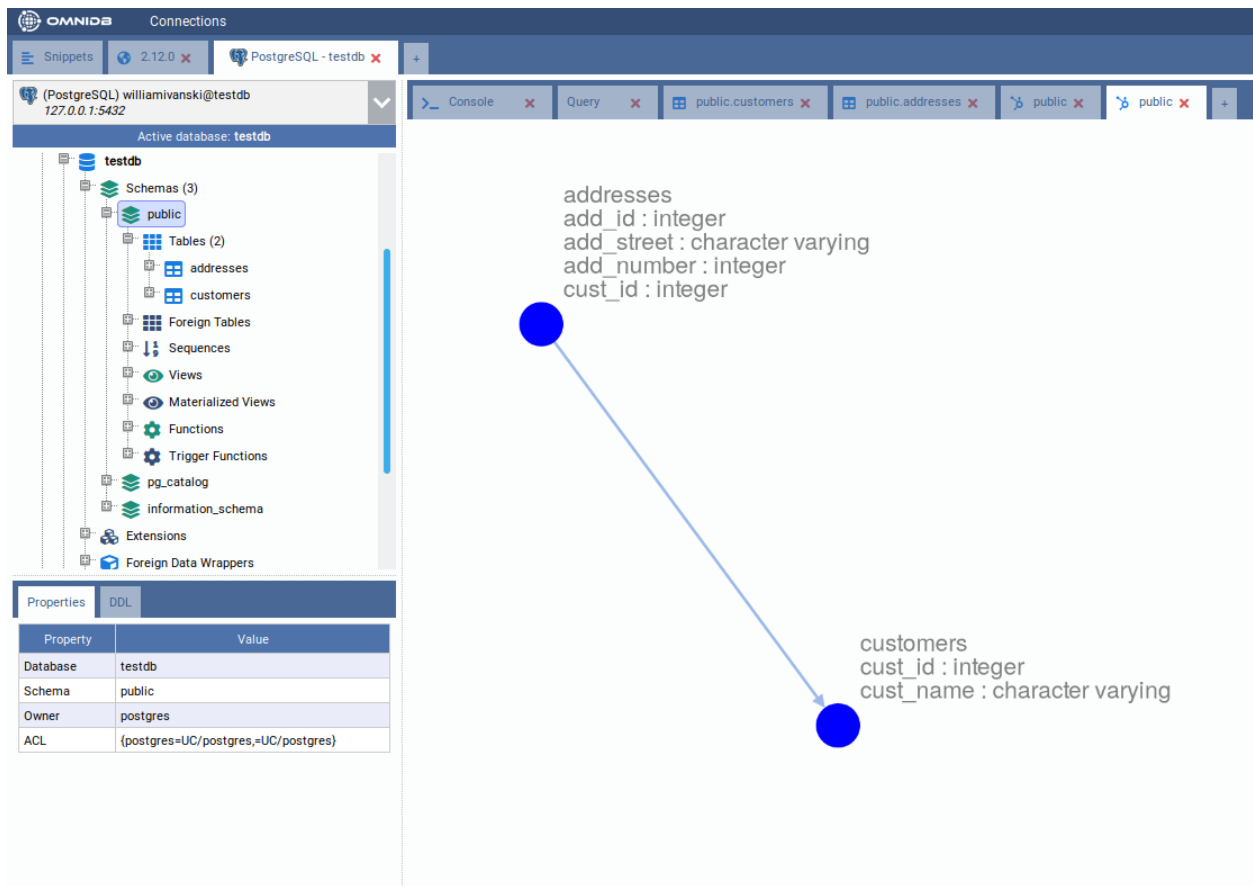


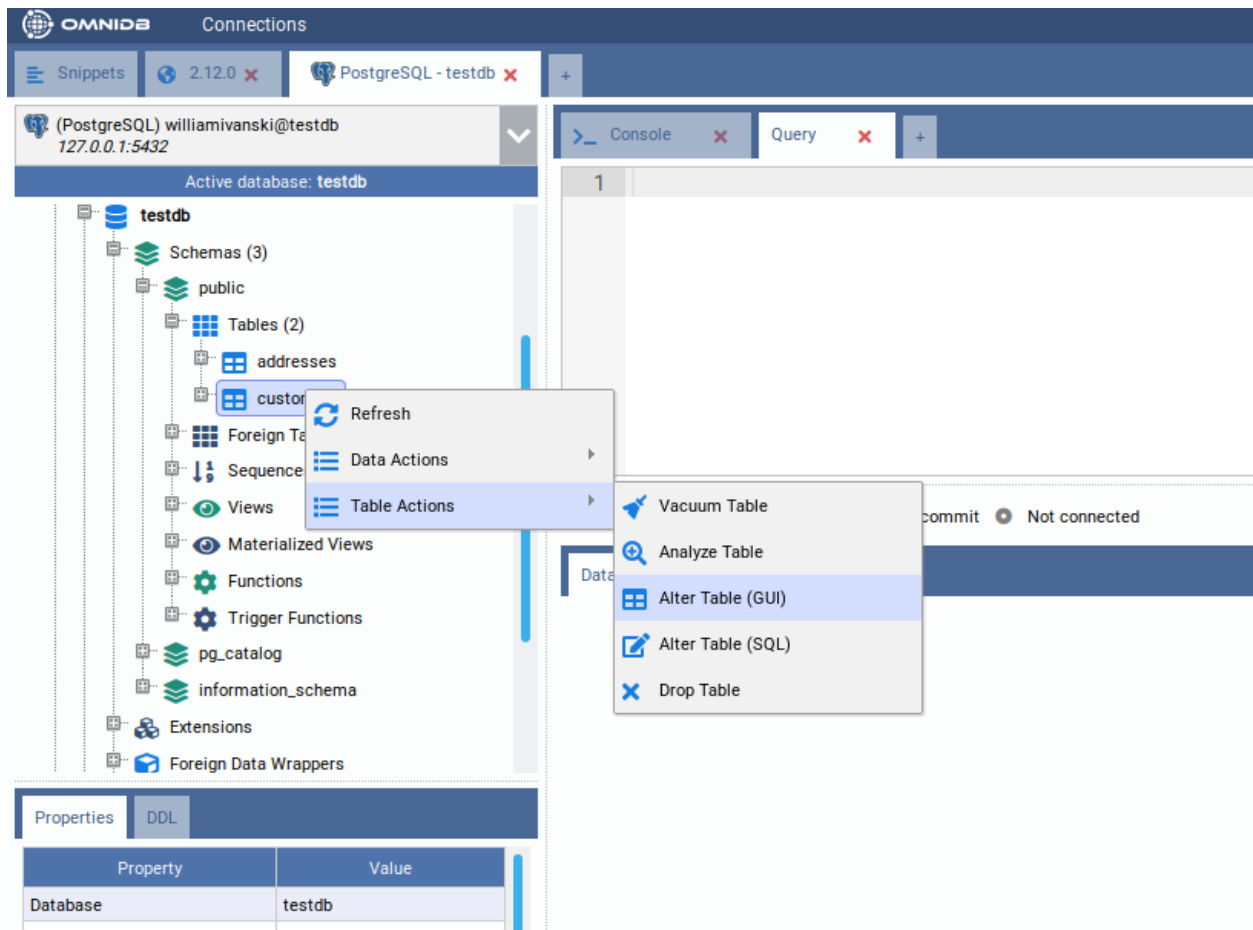
Diagram illustrating the relationship between 'customers' and 'addresses' tables. The 'addresses' table is connected to the 'customers' table via a relationship line.

And this is what the *Complete Graph* looks like:

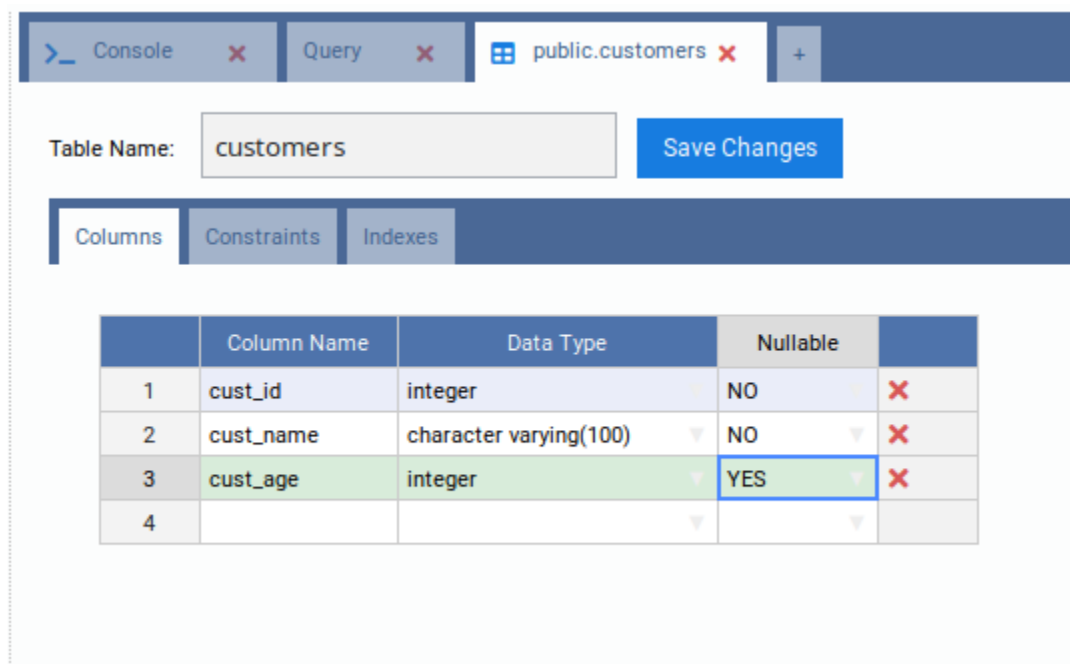


5.2 Editing tables

OmniDB also lets you edit existing tables (always following DBMS limitations). To test this feature we will add a new column to the table *customers*. To access the alter table interface just right click the table node and select the action *Table Actions > Alter Table*:

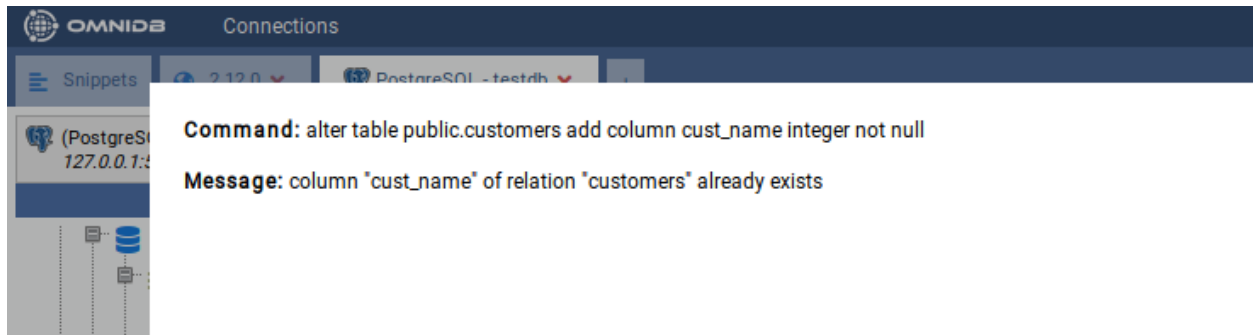


Add the column *cust_age* and save:



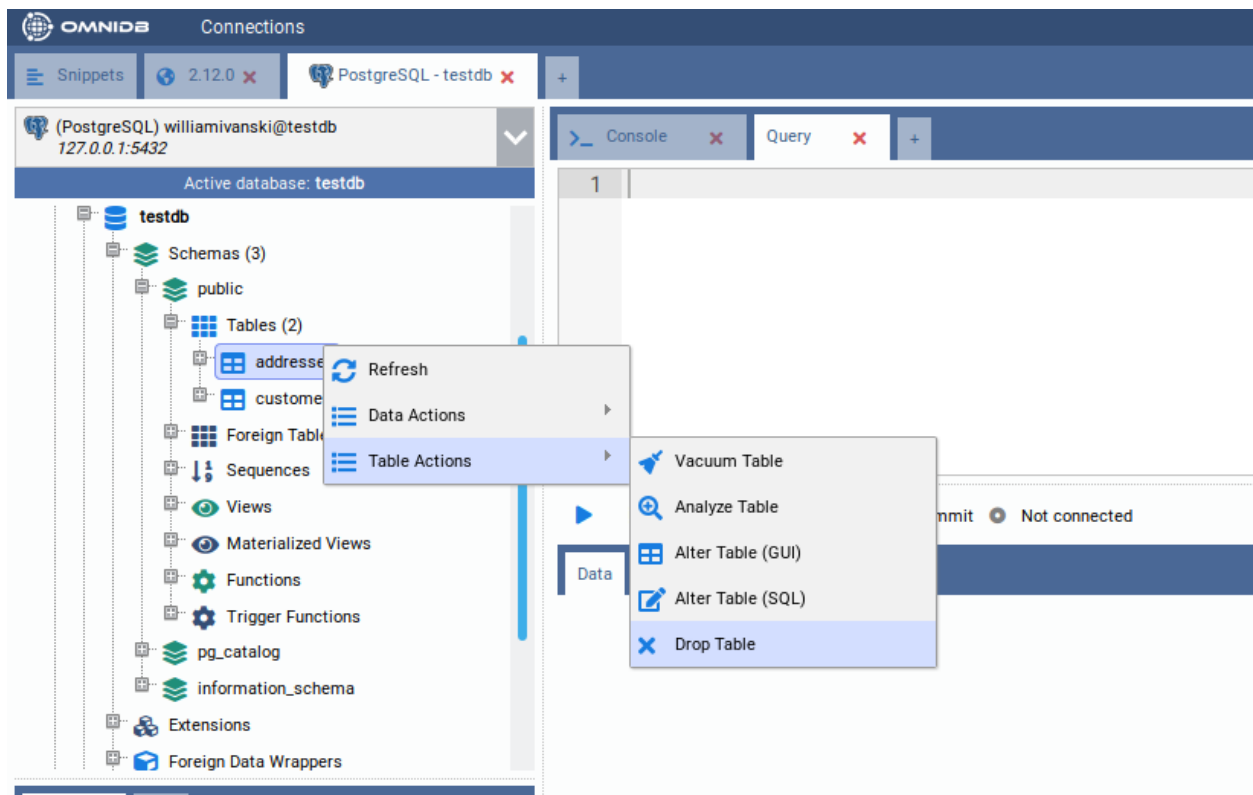
The interface is capable of detecting errors that may occur during alter table operations, showing the command and the

error that occurred. To demonstrate it we will try to add the column `cust_name`, which already belongs to this table:



5.3 Removing tables

In order to remove a table just right click the table node and select the action *Table Actions > Drop Table*:



6. Managing Table Data

The tool allows us to edit records contained in tables through a very simple and intuitive interface. Given that only a few DBMS have unique identifiers for table records, we opted to allow data editing and removal only for tables that have a primary key. Tables that do not have it can only receive new records.

To access the record editing interface, right click the table node and select the action *Data Actions > Edit Data*:

The screenshot shows the OmniDB interface with a PostgreSQL connection to a testdb database. The left sidebar displays a tree view of the database structure, including databases, schemas, tables, and views. The 'customer' table is selected, and a context menu is open, showing options like 'Refresh', 'Data Actions', and 'Table Actions'. The 'Data' tab is active, and the 'Edit Data' option is highlighted. The bottom panel shows the 'Properties' tab with the database name 'testdb'.

Connections

Snippets 2.12.0 PostgreSQL - testdb

(PostgreSQL) williamivanski@testdb 127.0.0.1:5432

Active database: testdb

PostgreSQL 10.5 (Ubuntu 10.5-0ubuntu0.18.04)

Databases (5)

- postgres
- dellstore2
- employees
- testdb

Schemas (3)

- public

Tables (2)

- addresses
- customer

Foreign Table

Sequences

Views

Materialized Views

Functions

Trigger Functions

pg_catalog

information_schema

Extensions

Refresh

Data Actions

Table Actions

Query Data

Edit Data

Insert Record

Update Records

Count Records

Delete Records

Truncate Table

Autocommit Not connected

Data Messages Explain

Properties DDL

Property	Value
Database	testdb

The screenshot shows the OmniDB interface with a PostgreSQL connection to a testdb database. The left sidebar displays a tree view of the database structure, including databases (postgres, dellstore2, employees, testdb), schemas (public), tables (addresses, customers), foreign tables, sequences, views, materialized views, functions, and trigger functions. The main panel shows a SQL editor with the query `select * from public.customers t order by t.cust_id`. Below the editor, a query execution button is visible. The results section shows "Number of records: 0" and "Response time: 0.029 seconds". A table with 3 columns (cust_id, cust_name, cust_age) and 1 row (1, +) is displayed.

The interface has a SQL editor where you can filter and order records. To prevent that the interface requests too many records, there is a field that limits the number of records to be displayed. The records grid has column names and data types. Columns that belong to the primary key have a key icon next to their names.

The row of the grid that have the symbol + is the row to add new records. Let us insert some records in the table customers:

Console
Query
public.customers
+

```
select * from public.customers t
1 order by t.cust_id
```

▶
Query 10 rows
Number of records: 0
Response time: 0.029 seconds
Save Changes

		cust_id (integer)	cust_name (character varying)	cust_age (integer)
1	✗	0	Pedro	22
2	✗	1	Ryan	18
3	✗	2	William	23
4	✗	3	Susan	31
5	✗	4	Nicole	19
6	✗	5	Ricardo	45
7	✗	6	Ademar	60
8	✗	7	Felipe	29
9	✗	8	Rafael	21
10	+			

After saving, the records will be inserted and can be edited (only because this table has a primary key). Let's change the *cust_name* of some of the existing records and, at the same time, let's remove one of the rows:

> Console
Query
public.customers
+

```
select * from public.customers t
```

1
order by t.cust_id

▶
Query 10 rows
Save time: 0.152 seconds
Save Changes

		cust_id (integer)	cust_name (character varying)	cust_age (integer)
1	✗	0	Pedro	22
2	✗	1	Ryan	18
3	✗	2	William Changed	23
4	✗	3	Susan	31
5	✗	4	Nicole	19
6	✗	5	Ricardo	45
7	✗	6	Ademar Changed	60
8	✗	7	Felipe	29
9	✗	8	Rafael	21
10	+			

Tables can have fields with values represented by very long strings. To help edit these fields, OmniDB has an interface that can be accessed by right clicking the specific cell:

Console Query public.customers

```
select * from public.customers t
1 order by t.cust_id
```

Query 10 rows Save time: 0.046 seconds

		cust_id (integer)	cust_name (character varying)	cust_age (integer)
1	✗	0	Pedro	22
2	✗	2	William Changed	23
3	✗	3	Susan	31
4	✗	4	Nicole	19
5	✗	5	Ricardo	45
6	✗	6	Ademar Changed	60
7	✗	7	Felipe	
8	✗	8	Rafael	
9	+			

Copy Edit Content

Connections

PostgreSQL 127.0.0.1

Properties

Database: PostgreSQL

Schema: public

Table: customers

OID: 16384

Owner: postgres

Size: 8192 B

Tablespace: pg_default

ACL: postgres=U

Options: (empty)

FileNode: (empty)

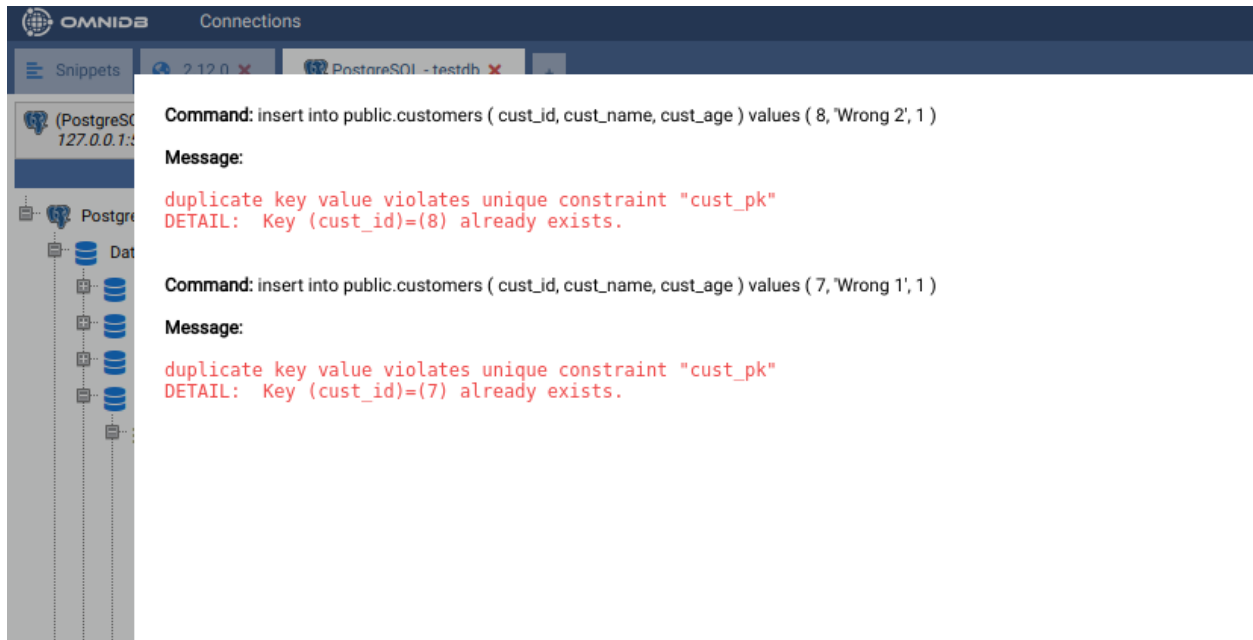
Estimate Count: 0

Has Index: true

Ademar Changed

The interface detects errors that may occur during operations related to records. To demonstrate, let us insert two

records with existing `cust_id` (primary key):



It shows which commands tried to be executed and the respective errors.

To complete this chapter, let's add some records to the *Address* table:

> Console

Query

public.addresses

+

select * from public.addresses t

1 order by t.add_id

▶

Query 10 rows

Number of records: 0
Response time: 0.028 seconds

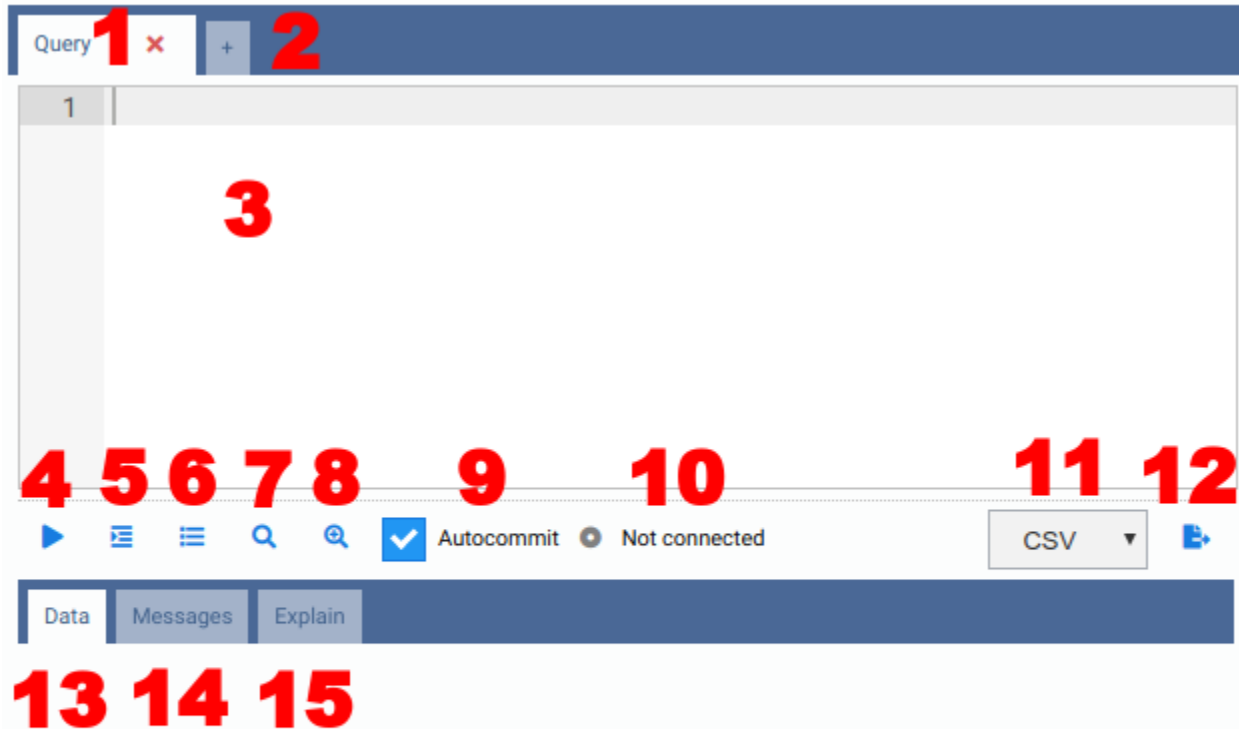
Save Changes

		add_id (integer)	add_street (character varying)	add_number (integer)	cust_id (integer)
1	✗	0	Blue Street	114	0
2	✗	2	Red Street	471	2
3	✗	3	Black Street	355	3
4	✗	4	Black Street	1002	4
5	✗	5	White Street	1056	5
6	✗	6	Green Street	19	6
7	✗	7	Yellow Street	47	7
8	✗	8	Purple Street	33	8
9	+				

CHAPTER 7

7. Writing SQL Queries

The most common kind of inner tab is the *Query Tab*, containing the following elements:



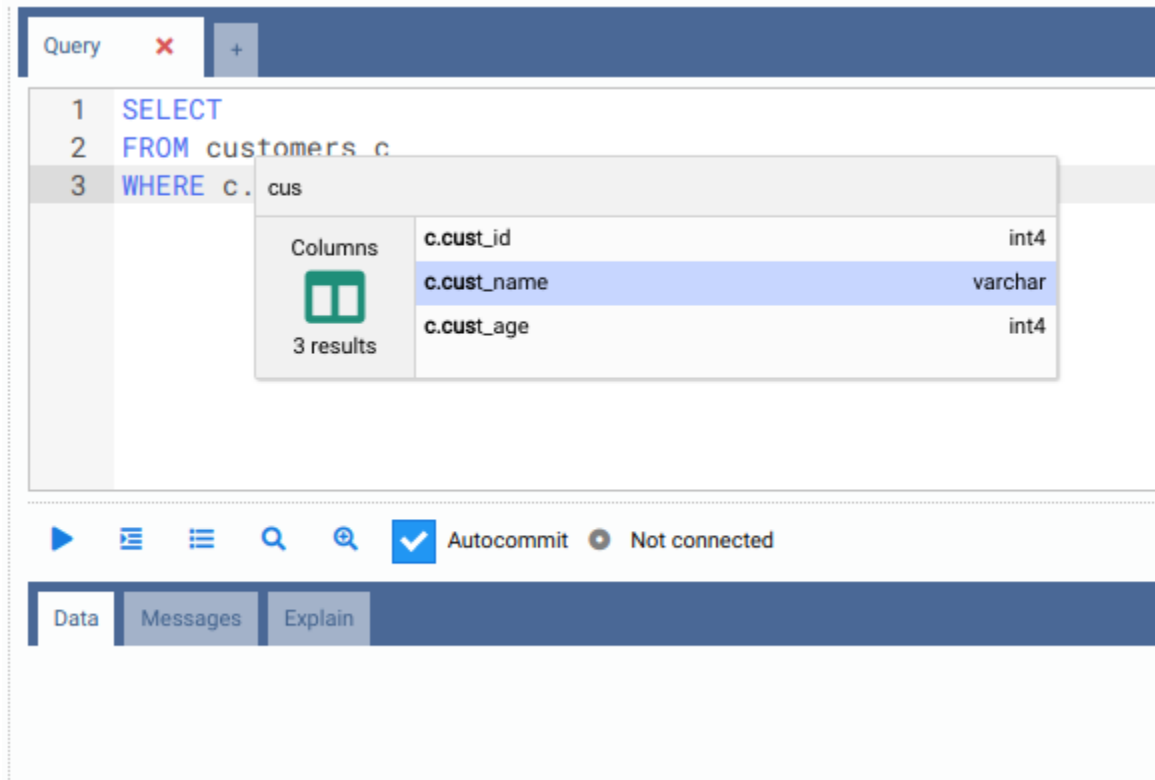
- **1) Tab Header:** You can see the name of the tab and an icon to close it. If there is a query running in the tab, you will see an indicator. If the tab finishes running and you are working on a different tab, a green indicator you be shown. By double-clicking on the tab name, you will be able to rename the tab.
- **2) Add Tab:** You can quickly add another inner tab by clicking on the plus icon.
- **3) SQL Editor:** Full-featured SQL editor with SQL syntax highlighting, Find & Replace (Ctrl-F and Ctrl-H) and an autocomplete component, explained below.
- **4) Execute:** The text contained in the SQL Editor will be executed against the current active database when you click on this button (or hit the shortcut, Alt-Q by default). If there is some selected text in the SQL Editor, it will execute only the selected text. Once the command is running, a red *Cancel* button will be shown, allowing you to cancel the execution (or using the shortcut, Alt-C by default).
- **5) Indent SQL:** This button will prettify any SQL code written in the SQL editor (shortcut Alt-D by default).
- **6) Query History:** All commands executed against the current database are stored in the Query History, which can be accessed by clicking on this button. You also will be able to filter by date and text to find a SQL command you need.
- **7) Explain (PostgreSQL only):** Call your SQL query against PostgreSQL by putting EXPLAIN in front of it. The results will be shown in a textual and graphical form in the Explain tab (please see Chapter 8 for more details).
- **8) Explain Analyze (PostgreSQL only):** Same as Explain button, but call the SQL query with EXPLAIN ANALYZE, which will effectively execute the query.

- **9) Autocommit (PostgreSQL only):** When enabled, every query executed will be committed to the database. When disabled, OmniDB starts a transaction and upon execution of a query, the interface will show buttons allowing the user to *Commit* or *Rollback*. The user can also keep the transaction open and execute other commands.
- **10) Backend Status (PostgreSQL only):** When you open a new Query Tab, the status is “Not Connected”, because OmniDB didn’t start a PostgreSQL backend yet. When you execute the first query, OmniDB starts a new backend and keep it linked to the Query Tab (each Query Tab will be assigned its own backend). The status of the backend (*idle*, *active*, *idle in transaction*, etc) will be shown in this field. When you close the Query Tab, OmniDB terminates the backend.
- **11) Export File Type:** Can be either CSV or XLSX.
- **12) Export To File:** By clicking on this button, OmniDB executes the current query and saves it to a file in OmniDB’s temp folder. After the file is saved, the interface allows the user to download it.
- **13) Data Results:** If the query is a `SELECT`, then it will show a grid with the results. If the query is a DML or DDL, it will show the message returned by the RDBMS.
- **14) Messages (PostgreSQL only):** Any messages (such as the ones given by the command `RAISE NOTICE`) will be shown here.
- **15) Explain View (PostgreSQL only):** Shows a full-featured component to view the PostgreSQL execution plan in textual or graphical form.

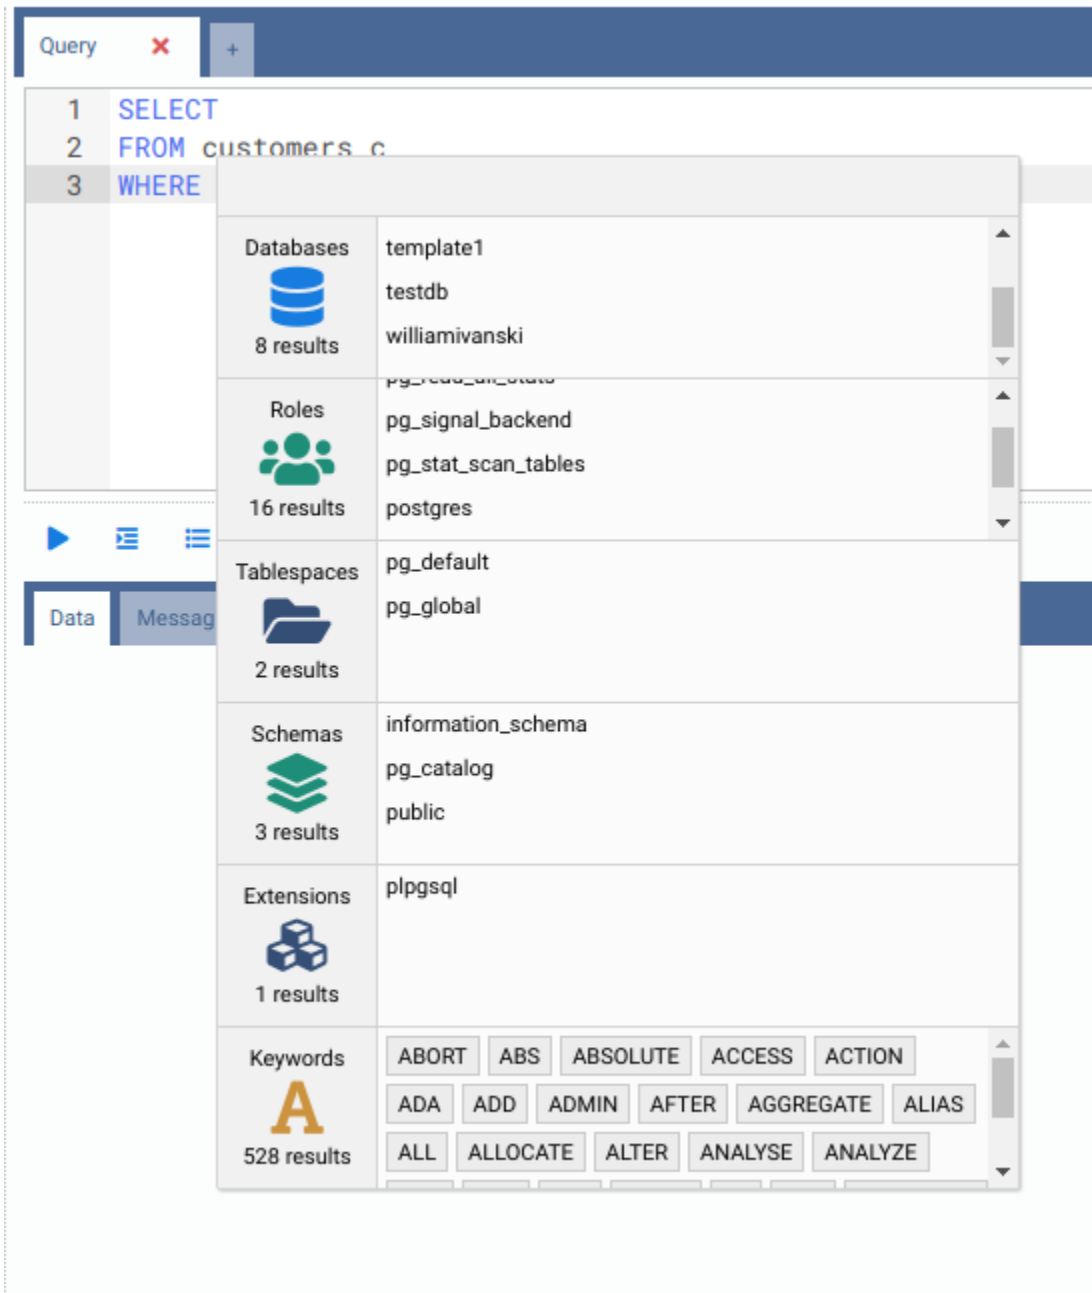
Once executed, the tabs are also saved in OmniDB user database (title and contents), so the next time you open OmniDB, you will see them all open. Also, every command you execute in a Query Tab is saved to your Query History and to the `omnldb.log` file too.

7.1 SQL Autocomplete

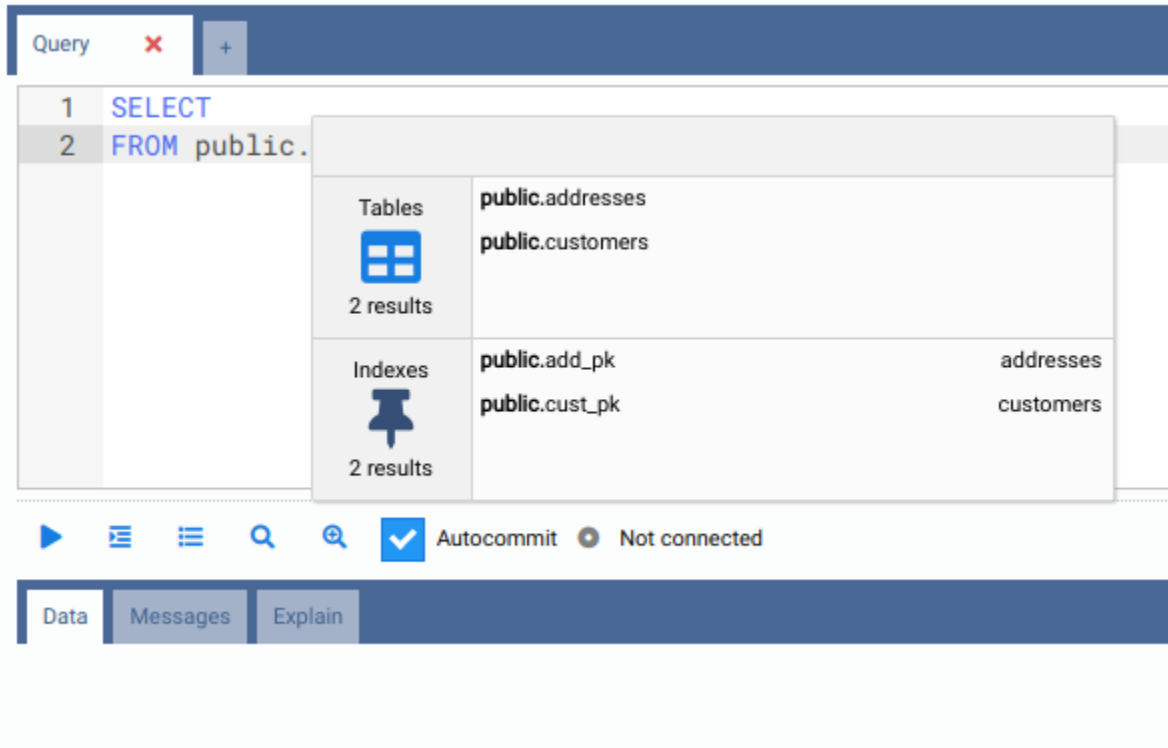
The SQL editor has a feature that helps a lot when creating new queries: SQL code completion. With this feature it is possible to autocomplete columns contained in a table referenced by an alias. To open the autocomplete interface you just have to type the alias, the character `.` and then hit `Ctrl-Space`:



If the user does not start the autocomplete with the cursor close to a table alias, the component will show multiple categories of data. By typing in the filter textbox, elements in all categories will be filtered:



The autocomplete component is also able to identify some contextual information. For example, if you type a name of a schema, then type the character `.`, then hit `Ctrl-Space`, you will be able to filter among objects contained only in that schema:



Please note that for RDBMS other than PostgreSQL, the autocomplete component only works for table columns.

8. Visualizing Query Plans

OmniDB 2.2.0 introduced a very useful feature: graphical query plan visualization. This may come in handy when writing or optimizing queries, since it allows you to easily identify performance bottlenecks in your SQL query.

For this feature, *SQL Query* inner tab shows 2 buttons: *Explain* and *Explain Analyze*.

8.1 Textual visualization

When you click the *Explain* button, OmniDB will execute an `EXPLAIN` command in your query. Initial visualization is *textual* and will show exactly the output of the `EXPLAIN` command, but with colored bars representing the estimated cost. The higher the cost, the darker and wider the bar.

The screenshot shows the OmniDB interface. At the top, there's a 'Query' tab with a red 'X' and a '+' button. Below it, a SQL query is entered in a text area:

```

1 select *
2 from (
3     select cust_name,
4         (select count(*)
5          from addresses addr
6          where addr.cust_id = cust.cust_id) as num_addresses
7     from customers cust
8 ) subquery
9 where subquery.cust_name = 'Rafael'

```

Below the query area, there's a toolbar with icons for execution, settings, and search. A status bar shows 'Autocommit' (checked), 'Idle', and 'Start time: 11/01/2018 15:02:24 Duration: 8.899 ms'.

Below the toolbar, there's a tabbed interface with 'Data', 'Messages', and 'Explain' tabs. The 'Explain' tab is selected, showing the 'QUERY PLAN' section. The plan is as follows:

```

# QUERY PLAN
1 Seq Scan on customers cust (cost=0.00..38.27 rows=2 width=226)
2 Filter: ((cust_name)::text = 'Rafael'::text)
3 SubPlan 1
4   Aggregate (cost=12.13..12.14 rows=1 width=8)
5     Seq Scan on addresses addr (cost=0.00..12.12 rows=1 width=0)
6       Filter: (cust_id = cust.cust_id)

```

When you click the *Explain Analyze* button, OmniDB will execute an `EXPLAIN ANALYZE` command in your query. Beware that this command will really execute the query. Also, the textual visualization will show much more information, and the costs are not estimated as in those provided by the `EXPLAIN` command; they are real costs.

The screenshot shows the OmniDB SQL editor interface. At the top, there's a 'Query' tab with a red 'X' and a '+' icon. Below it, a SQL query is entered in a text area:

```

1 select *
2 from (
3     select cust_name,
4         (select count(*)
5          from addresses addr
6          where addr.cust_id = cust.cust_id) as num_addresses
7     from customers cust
8 ) subquery
9 where subquery.cust_name = 'Rafael'

```

Below the query editor, there's a toolbar with icons for running, saving, and other actions. A status bar shows 'Autocommit' is checked, 'Idle' status, and 'Start time: 11/01/2018 15:02:49 Duration: 14.903 ms'.

Below the status bar, there are tabs for 'Data', 'Messages', and 'Explain'. The 'Explain' tab is selected, showing the 'QUERY PLAN' for the executed query. The plan is as follows:

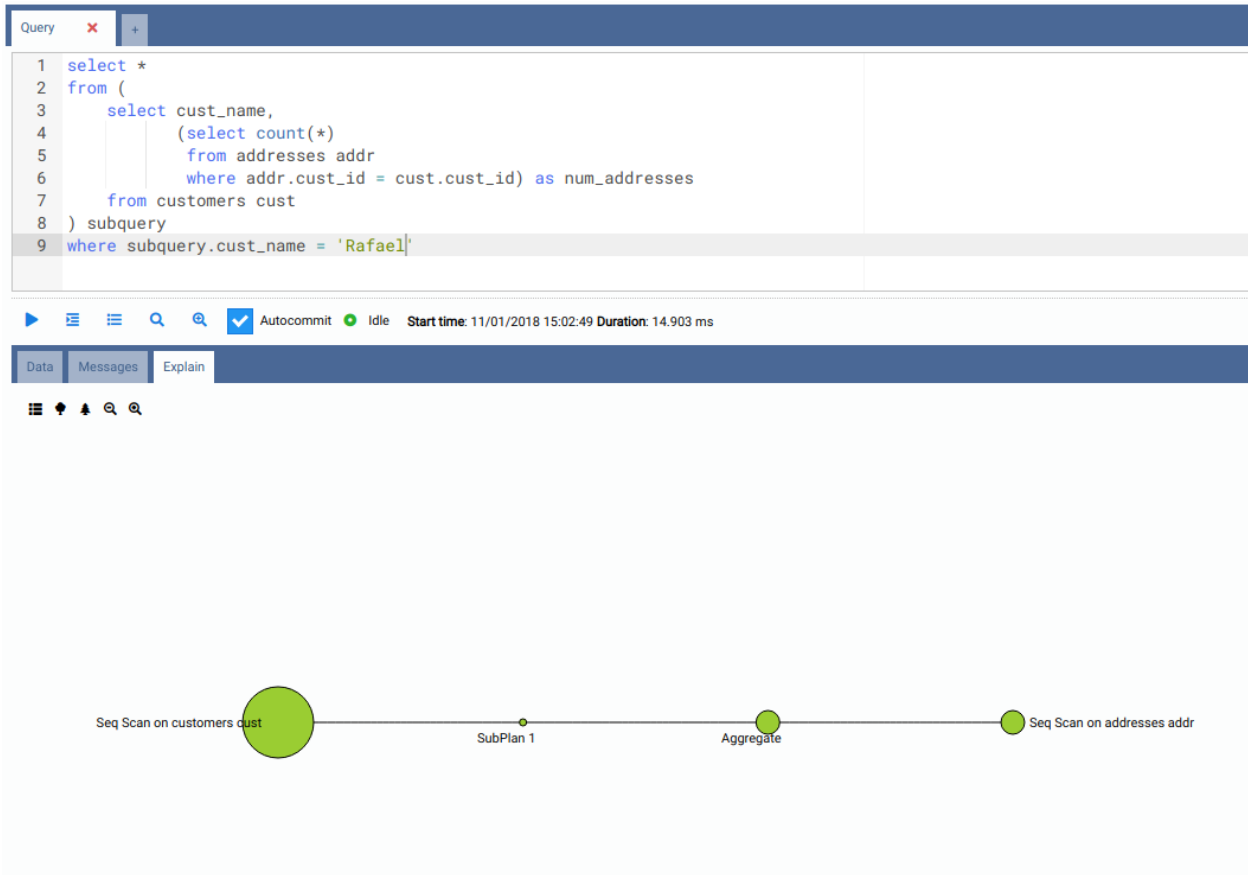
```

# QUERY PLAN
1 Seq Scan on customers cust (cost=0.00..38.27 rows=2 width=226) (actual time=0.029..0.030 rows=1 loops=1)
2   Filter: ((cust_name)::text = 'Rafael'::text)
3   Rows Removed by Filter: 7
4   SubPlan 1
5     Aggregate (cost=12.13..12.14 rows=1 width=8) (actual time=0.011..0.012 rows=1 loops=1)
6       Seq Scan on addresses addr (cost=0.00..12.12 rows=1 width=0) (actual time=0.006..0.006 rows=1 loops=1)
7         Filter: (cust_id = cust.cust_id)
8         Rows Removed by Filter: 7
9   Planning time: 0.155 ms
10  Execution time: 0.070 ms

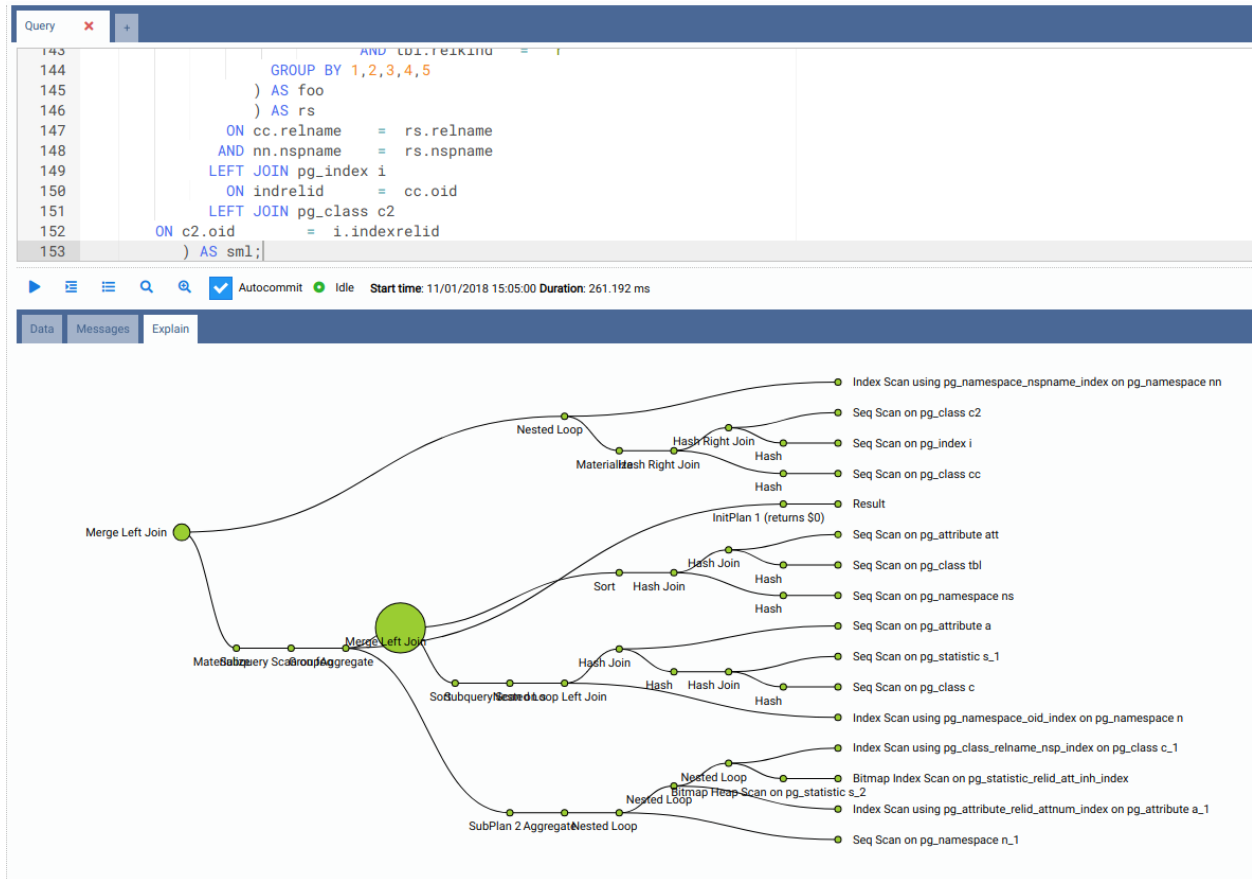
```

8.2 Tree visualization

Both *Explain* and *Explain Analyze* modes also can graphically represent the textual output into a *tree* diagram. Each circle represent a node executed by the query plan, and the larger the circle, the higher the cost.



When queries become more and more complex, also its query plan can be very complex. With such queries (like the *check bloat* query we executed below) the tree visualization can be very interesting:



The query plan visualization component allows you to easily switch between textual and 2 tree visualizations, which can be zoomed in and out.

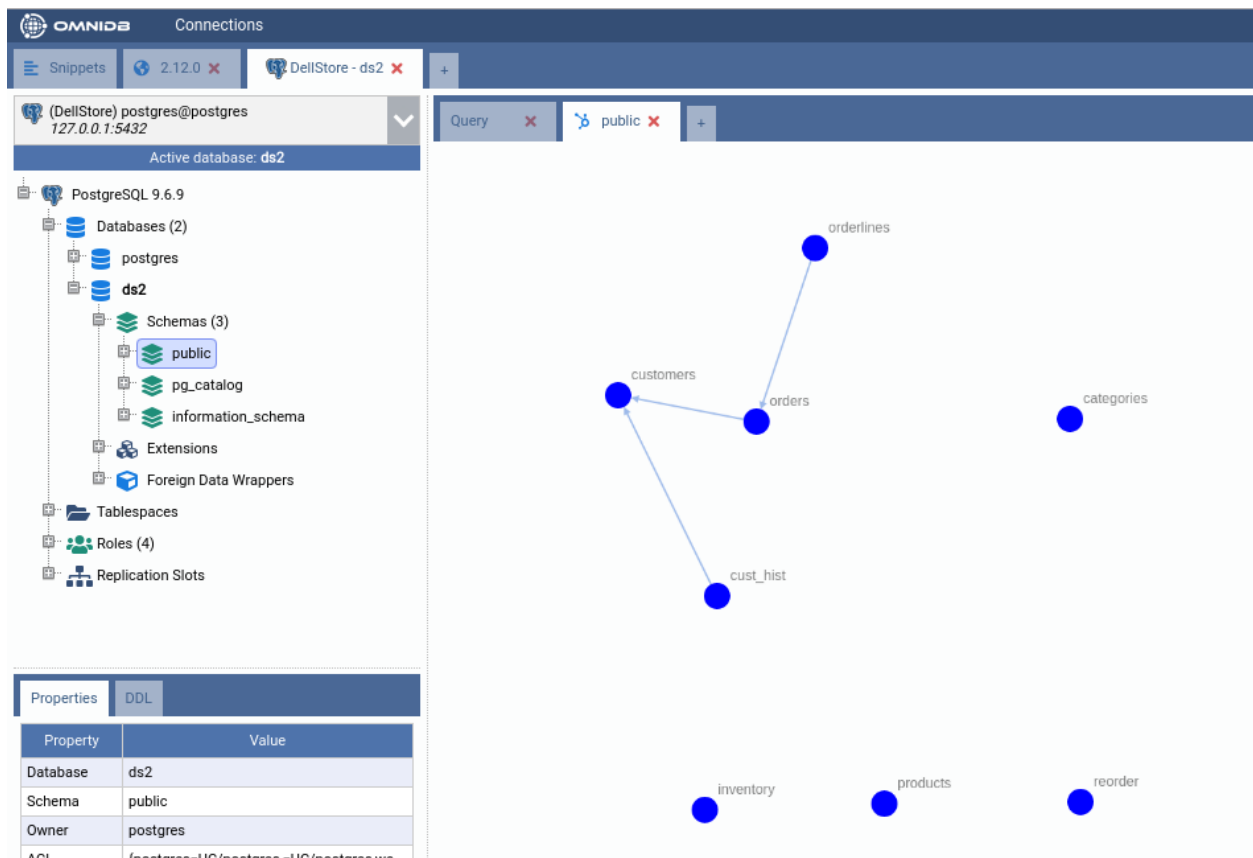
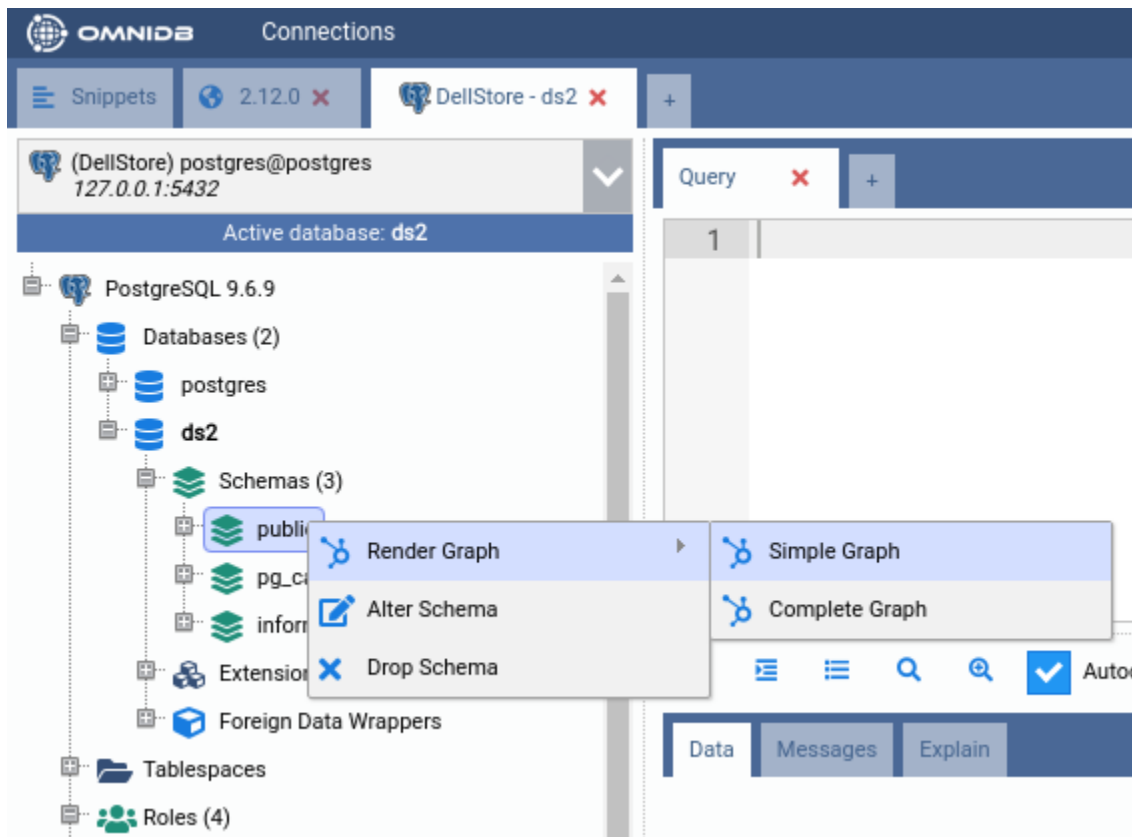
9. Visualizing Data

This feature displays a graph with nodes representing tables and edges representing table relationships with foreign keys. Using the mouse, the user is able to zoom in, zoom out, and drag and drop nodes to change its position.

There are two types of graphs: *Simple Graph* and *Complete Graph*.

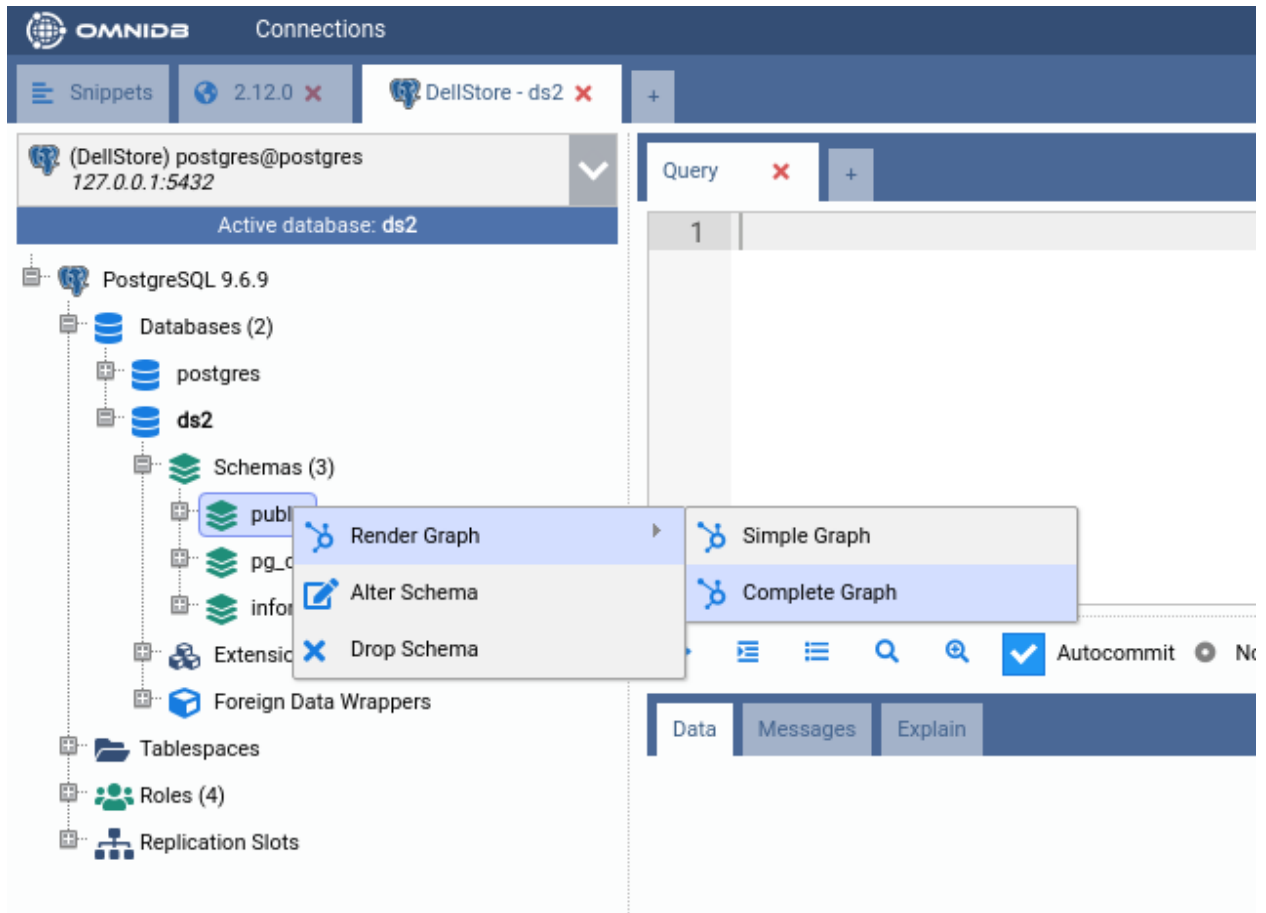
9.1 Simple graph

This one display simple table nodes and their relationships. To access it just right click the schema node you want in the tree and then select the action *Render Graph > Simple Graph*:



9.2 Complete graph

This graph displays tables with all its columns and respective data types. To access it just right click the schema you want in the tree and then select the action *Render Graph > Complete Graph*:



The image shows the OmniDB interface with a PostgreSQL 9.6.9 database connection. The left sidebar displays the database structure, including the 'ds2' database and its schemas. The main area shows a diagram of database relationships between 'orderlines', 'orders', and 'customers' tables.

Database Structure:

- PostgreSQL 9.6.9
 - Databases (2)
 - postgres
 - ds2
 - Schemas (3)
 - public
 - pg_catalog
 - information_schema
 - Extensions
 - Foreign Data Wrappers
 - Tablespaces
 - Roles (4)
 - Replication Slots

Table Definitions:

- orderlines**
 - orderid : integer
 - prod_id : integer
 - quantity : smallint
 - orderdate : date
- orders**
 - orderid : integer
 - orderdate : date
 - customerid : integer
 - netamount : numeric
 - tax : numeric
 - totalamount : numeric
- customers**
 - customerid : integer
 - firstname : character varying
 - lastname : character varvint

Properties Panel:

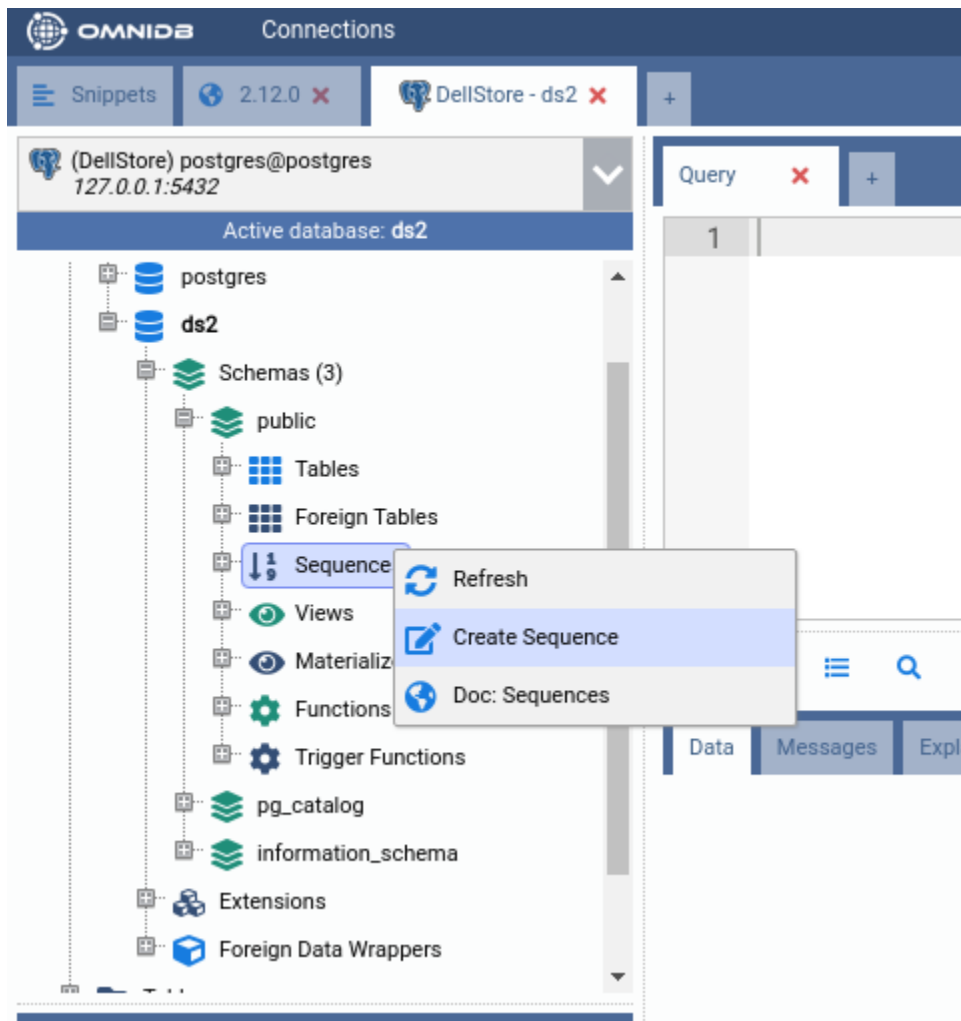
Property	Value
Database	ds2
Schema	public
Owner	postgres
ACL	{postgres=UC/postgres,=UC/postgres,we...

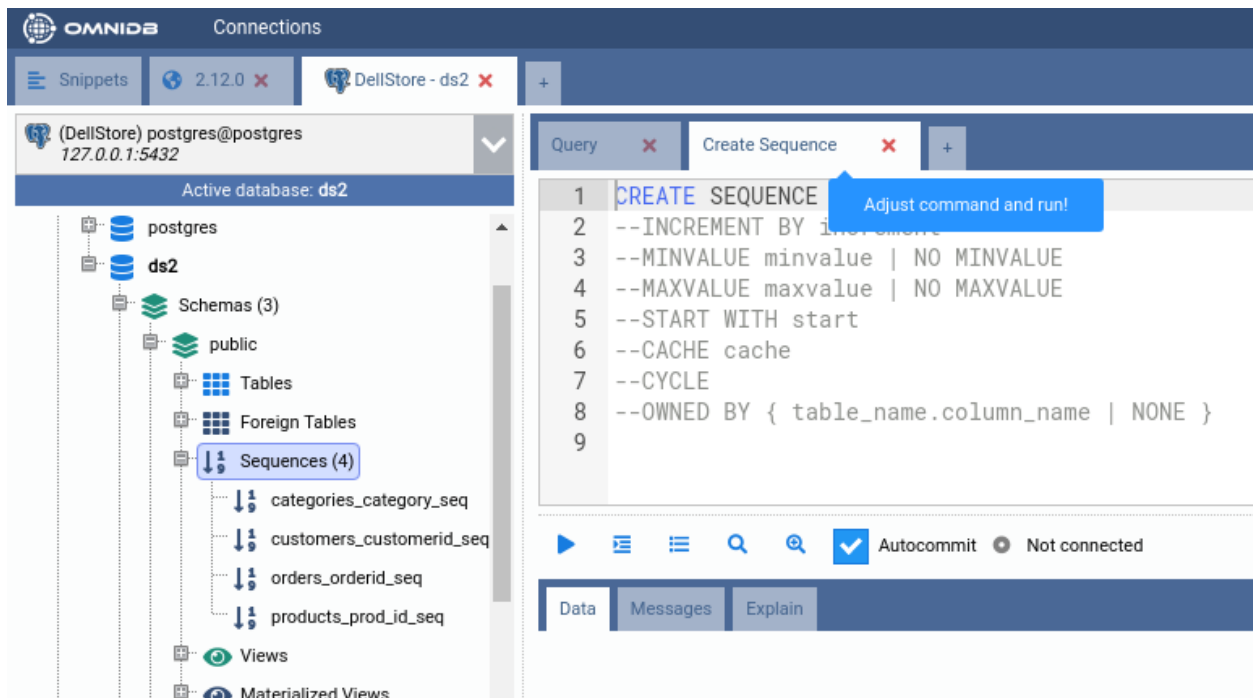
CHAPTER 10

10. Managing other Elements

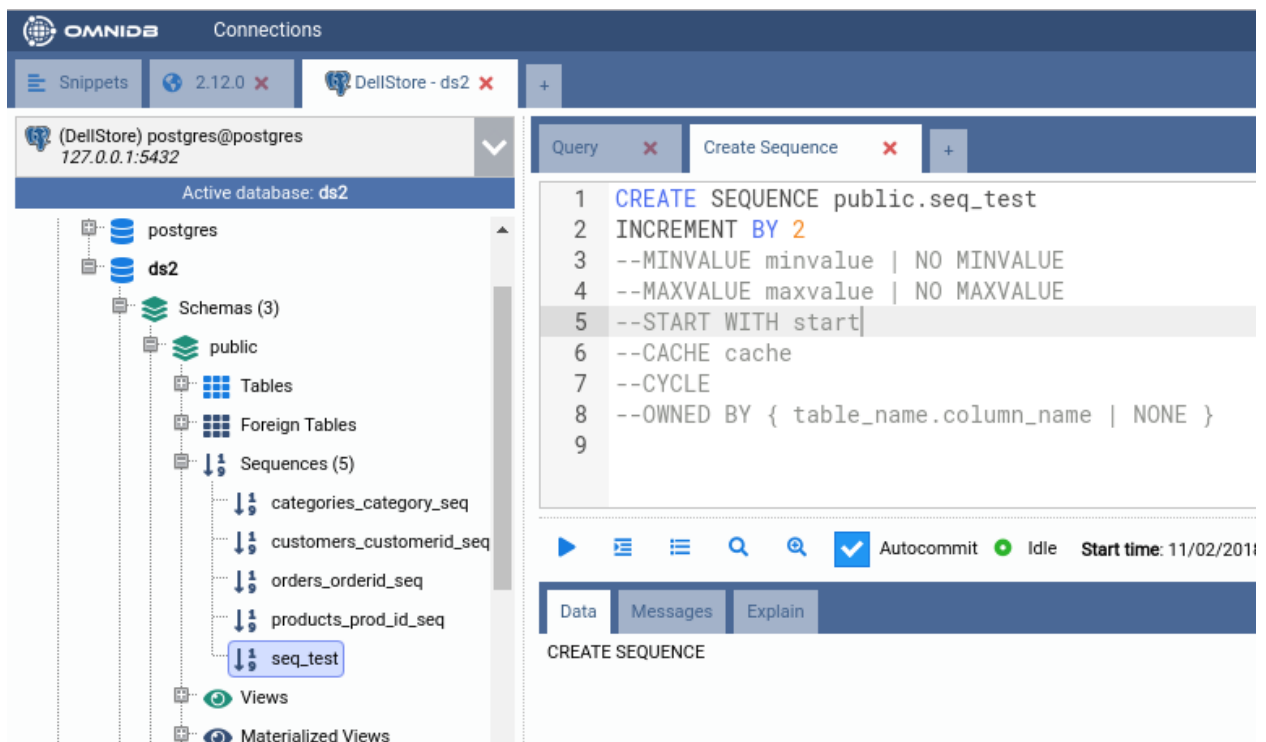
All PostgreSQL structures are possible to be managed with the use of *SQL templates*. This gives the user more power than using graphical forms to manipulate structures.

For example, let's consider the sequences inside the schema `public` of the database `ds2`. To create a new sequence, right click on the *Sequences* node, and choose *Create Sequence*.

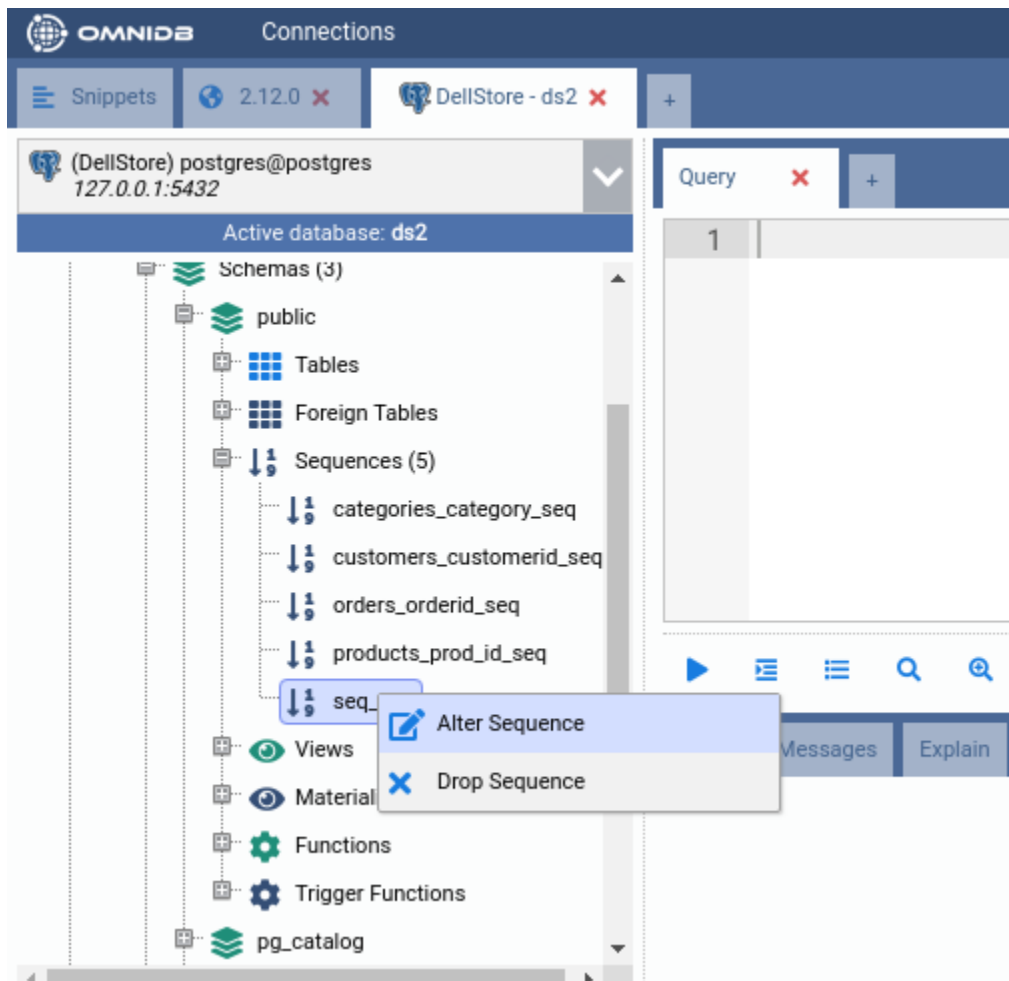




After you change the name of the sequence, you can uncomment other command options and set them accordingly to your needs. When the entire command looks fine, just execute it and the new sequence will be created:



With right click on an existing sequence, you can alter or drop it. It will work the same way as the creation, by using a SQL template for the user to change.



Connections: Snippets, 2.12.0, DellStore - ds2

(DellStore) postgres@postgres
127.0.0.1:5432

Active database: ds2

Schemas (3)

- public
 - Tables
 - Foreign Tables
 - Sequences (5)
 - categories_category_seq
 - customers_customerid_seq
 - orders_orderid_seq
 - products_prod_id_seq
 - seq_test
 - Views
 - Materialized Views
 - Functions
 - Trigger Functions

Query Editor: Alter Sequence

```
1 ALTER SEQUENCE
2 --INCREMENT BY increment
3 --MINVALUE minvalue | NO MINVALUE
4 --MAXVALUE maxvalue | NO MAXVALUE
5 --START WITH start
6 --RESTART
7 --RESTART WITH restart
8 --CACHE cache
9 --CYCLE
10 --NO CYCLE
11 OWNED BY ( table_name column_name | NONE )
```

Adjust command and run!

Autocommit: ☒ Autocommit ☐ Not connected

Data Messages Explain

Connections: Snippets, 2.12.0, DellStore - ds2

(DellStore) postgres@postgres
127.0.0.1:5432

Active database: ds2

Schemas (3)

- public
 - Tables
 - Foreign Tables
 - Sequences (5)
 - categories_category_seq
 - customers_customerid_seq
 - orders_orderid_seq
 - products_prod_id_seq
 - seq_test
 - Views
 - Materialized Views
 - Functions
 - Trigger Functions

Query Editor: Drop Sequence

```
1 DROP SEQUENCE p
2 --CASCADE
3
```

Adjust command and run!

Autocommit: ☒ Autocommit ☐ Not connected

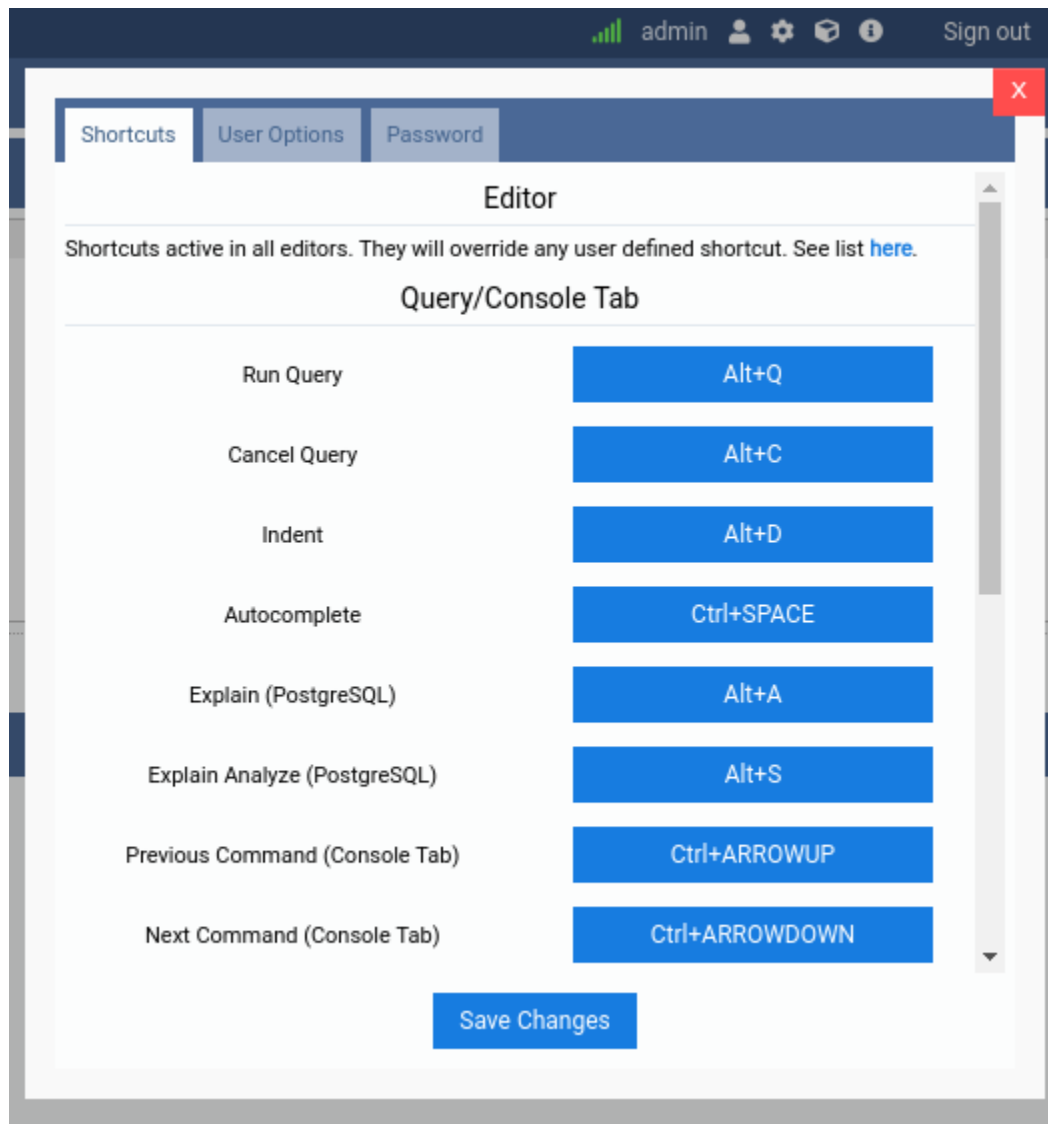
Data Messages Explain

11. Additional Features

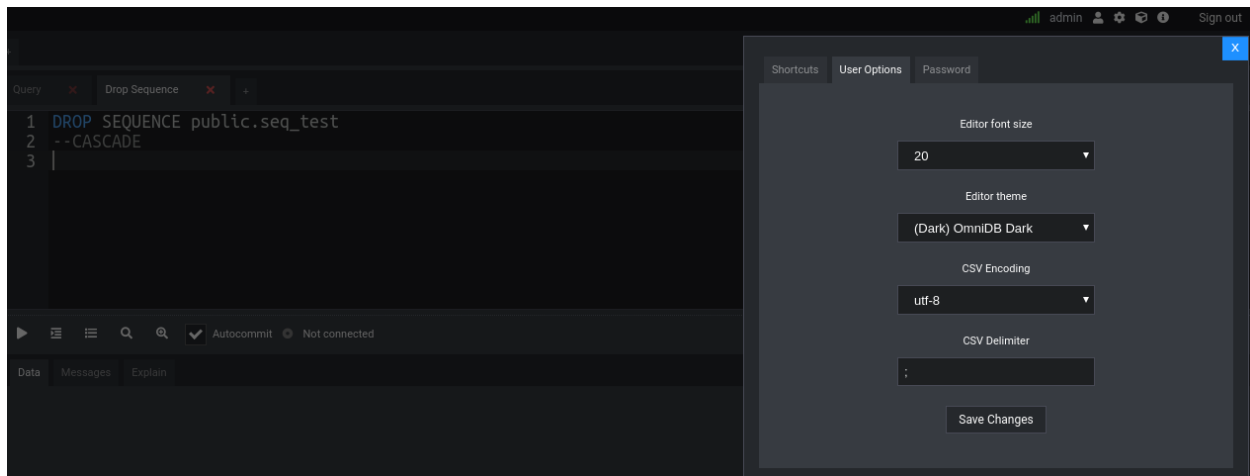
11.1 User Settings

Also in the upper right corner, by clicking in the gear-like icon, OmniDB will open the *User Settings* pop-up. It is composed by three tabs:

- **Shortcuts:** Allows the user to change its shortcuts in OmniDB.



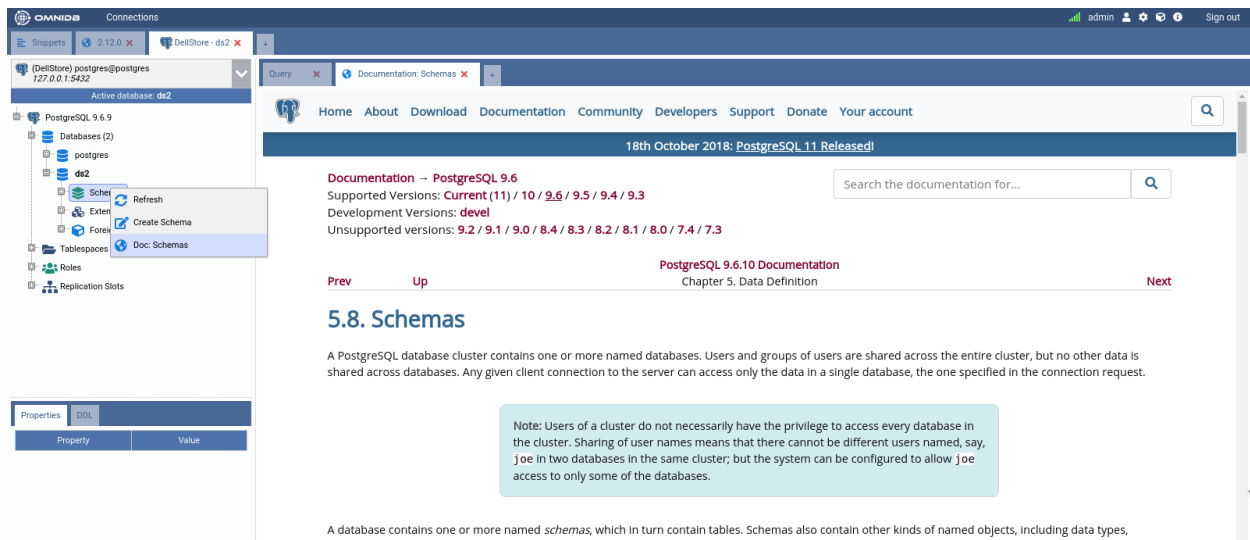
- **User Options:** Allows the user to change the font size of the SQL Editor, change the entire OmniDB theme and configurate CSV related options. There are a lot of OmniDB themes, each of them also change the syntax highlight color of the editor. They are also categorized in light and dark themes. A light theme is the default; a dark theme will change the entire interface of OmniDB.



- **Password:** Allows the user to change its password.

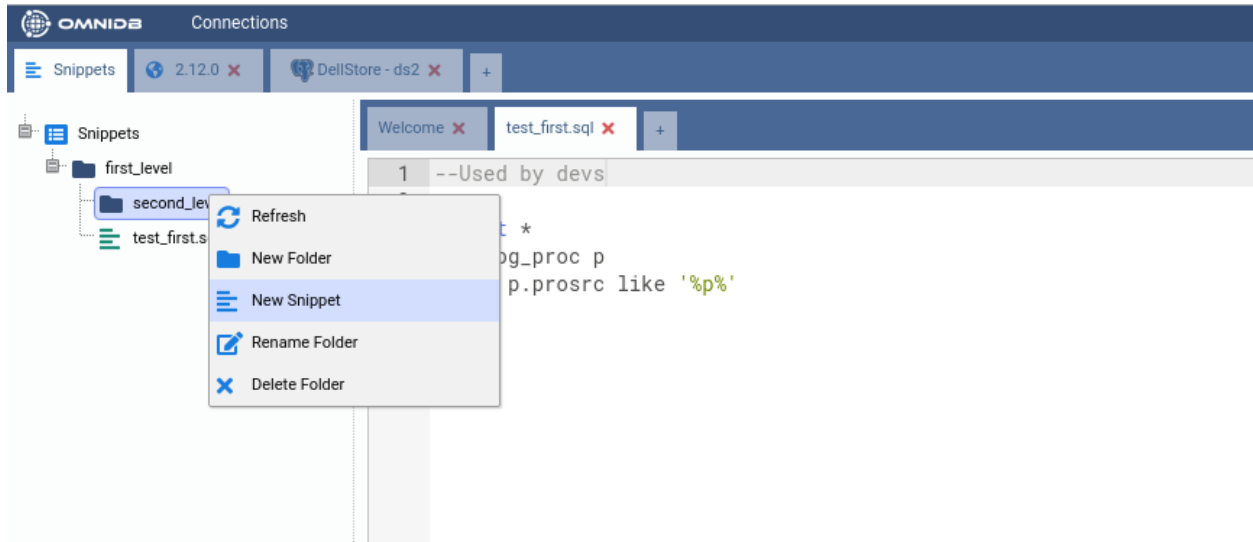
11.2 Contextual Help

Most of tree nodes (generally grouping ones like *Schemas* or *Tables*) offer contextual help. This feature can be accessed by right-clicking the tree node. When you click in the *Doc: ...* option, OmniDB will open an inner tab showing a web browser pointing to the specific page in the online **PostgreSQL Documentation**. Also, it will redirect to the specific page considering the PostgreSQL version you are connected to.



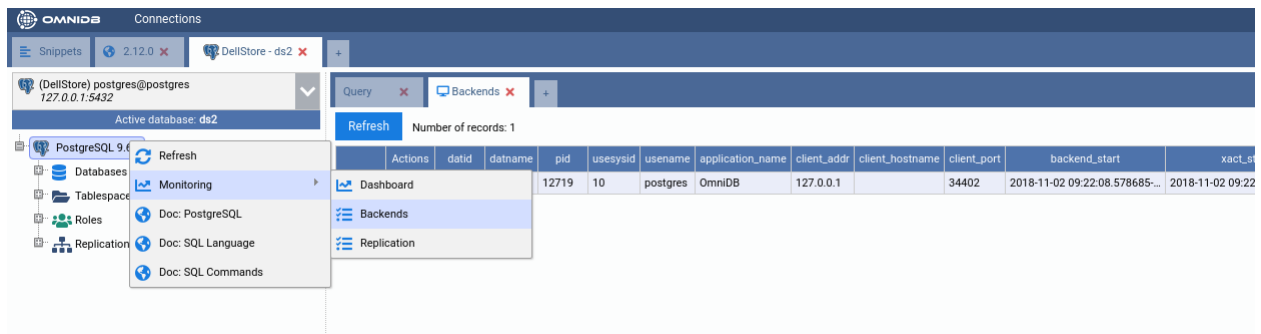
11.3 Snippets

Workspace Window has a fixed outer tab with an useful feature called *Snippets*. With this feature you can store queries, command instructions and any other kinds of text you want. You can also structure the snippets in a directory tree the way you want. All directories and snippets you create are stored inside of `omnodb.db` user database and persist when you upgrade OmniDB.

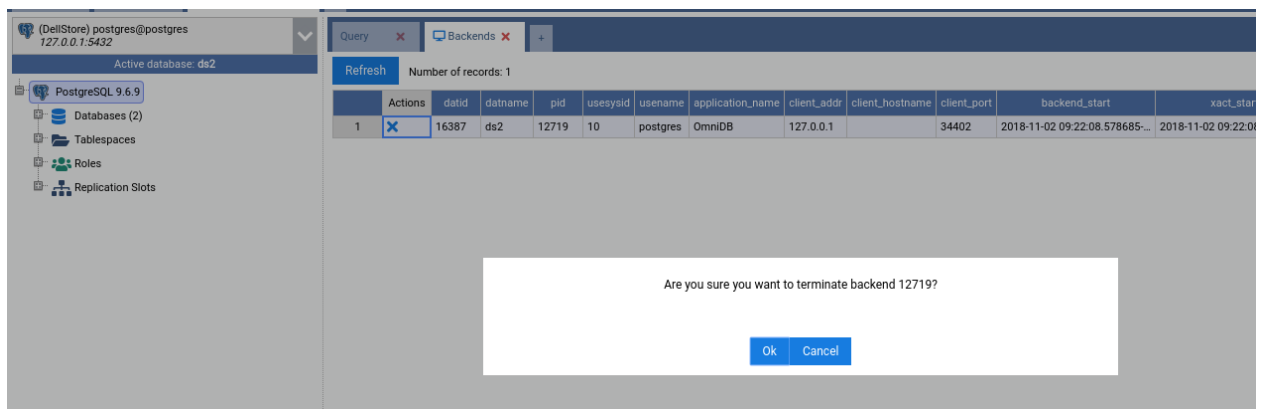


11.4 Backend Management

By right-clicking in the tree root node, then moving mouse pointer to *Monitoring* and then clicking on *Backends*, the user can see all activities going on in the database. Some information are hidden for normal users, only database superusers are allowed to see.



By clicking in the *X* in the *Actions* column, you can terminate the backend. A confirmation popup will appear.



11.5 Properties and DDL

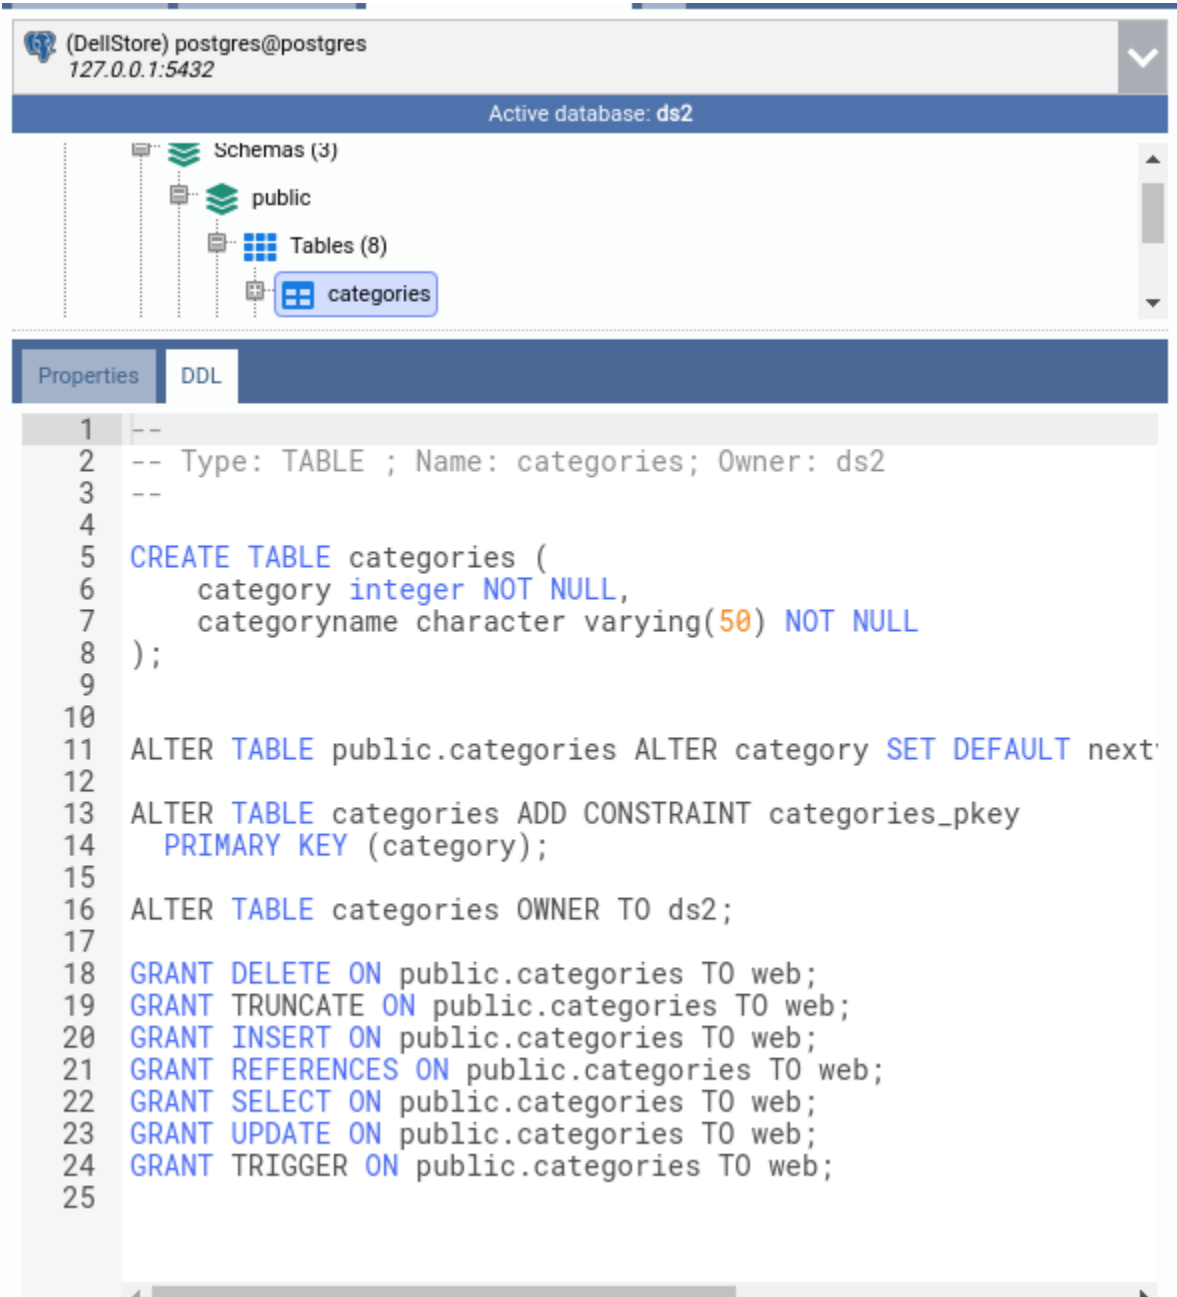
By clicking on most of objects in the tree view (tables, sequences, views, roles, databases, etc), the user will be able to see a very comprehensive list of properties of the object.



The screenshot shows the OmniDB interface. At the top, the connection is identified as '(DellStore) postgres@postgres' with IP '127.0.0.1:5432'. The active database is 'ds2'. The tree view on the left shows the hierarchy: Schemas (3) > public > Tables (8) > categories. The 'categories' table is selected. Below the tree view, there are two tabs: 'Properties' (active) and 'DDL'. The 'Properties' tab displays a table with the following data:

Property	Value
Database	ds2
Schema	public
Table	categories
OID	16434
Owner	ds2
Size	8192 bytes
Tablespace	pg_default
ACL	{ds2=arwdDxt/ds2,web=arwdDxt/ds2}
Options	
Filenode	base/16387/16434
Estimate Count	16
Has Index	true
Persistence	Permanent
Number of Attributes	2
Number of Checks	0
Has OIDs	false
Has Primary Key	true
Has Rules	false
Has Triggers	false

In the other panel called *DDL*, the user will be able to see the SQL DDL source code that can be used to re-create the object. The user can copy this text and paste it wherever he/she wants.



(DellStore) postgres@postgres
127.0.0.1:5432

Active database: ds2

Schemas (3)
public
Tables (8)
categories

Properties DDL

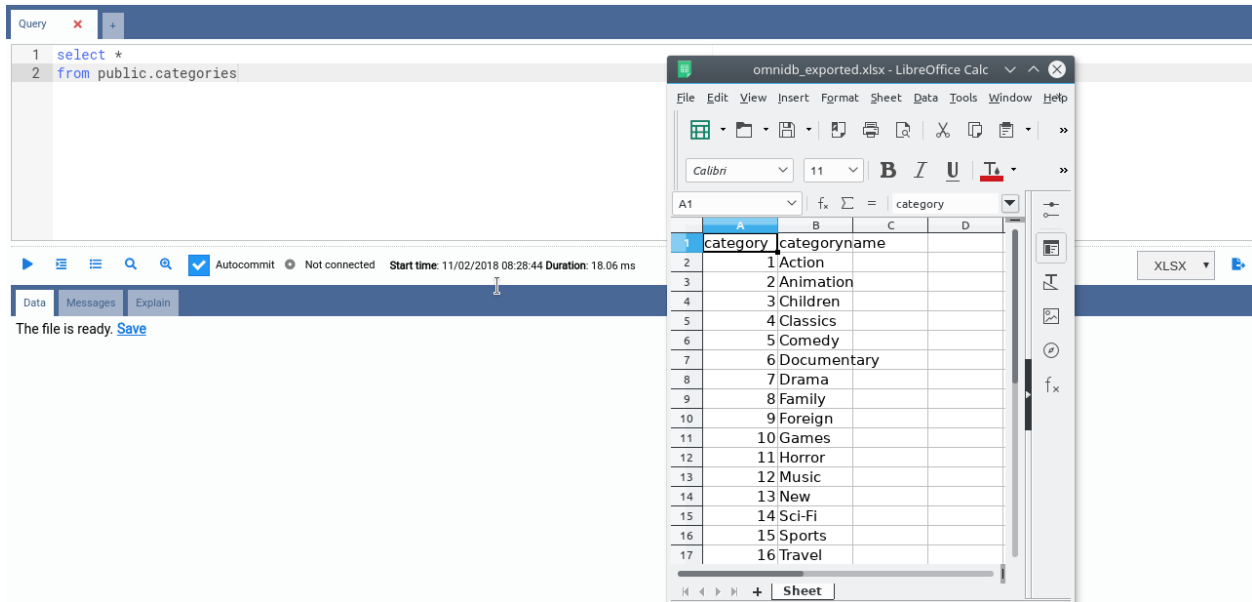
```

1  --
2  -- Type: TABLE ; Name: categories; Owner: ds2
3  --
4
5  CREATE TABLE categories (
6      category integer NOT NULL,
7      categoryname character varying(50) NOT NULL
8  );
9
10
11 ALTER TABLE public.categories ALTER category SET DEFAULT next
12
13 ALTER TABLE categories ADD CONSTRAINT categories_pkey
14     PRIMARY KEY (category);
15
16 ALTER TABLE categories OWNER TO ds2;
17
18 GRANT DELETE ON public.categories TO web;
19 GRANT TRUNCATE ON public.categories TO web;
20 GRANT INSERT ON public.categories TO web;
21 GRANT REFERENCES ON public.categories TO web;
22 GRANT SELECT ON public.categories TO web;
23 GRANT UPDATE ON public.categories TO web;
24 GRANT TRIGGER ON public.categories TO web;
25

```

11.6 Export Data

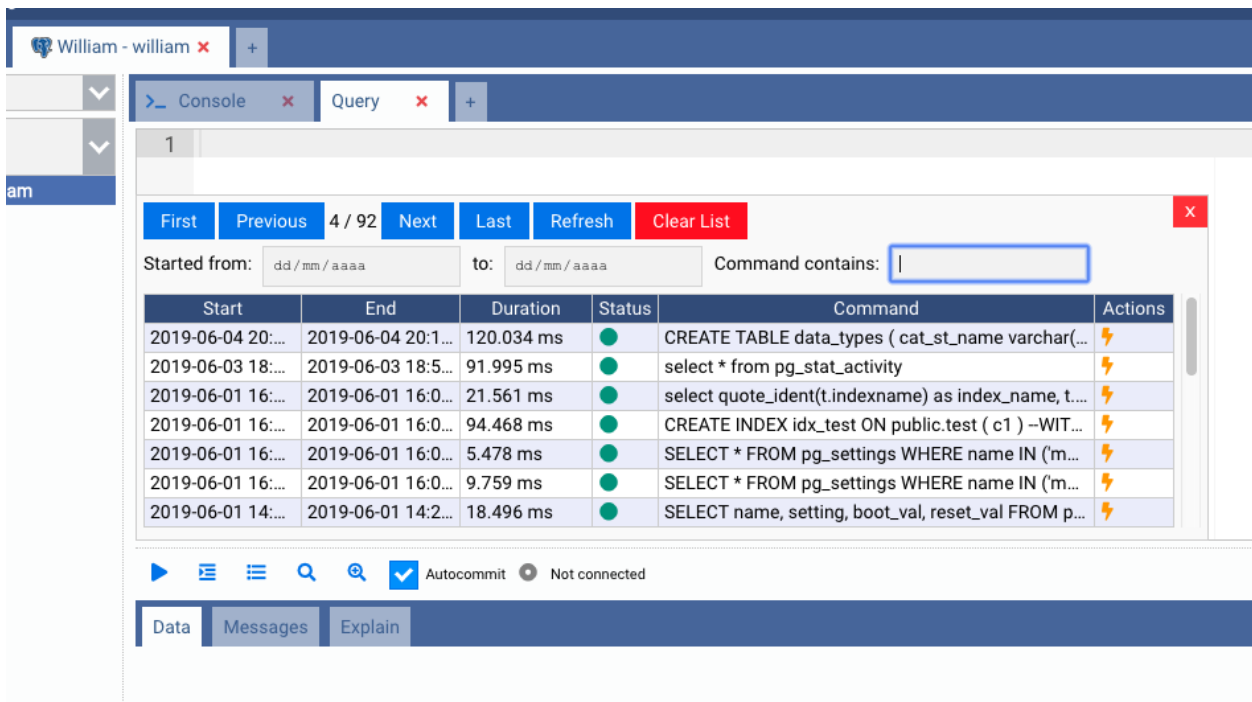
The *Query Tab* provides a way to save data from query results into a CSV or XLSX file. Once you click the *Export Data* button, a cancellable backend starts to save data into the file. Once it is done, OmniDB provides a link called *Save*, so the user can download the file.



All files are stored in a temporary folder inside OmniDB folder. OmniDB regularly cleans this folder, keeping only files newer than 24 hours.

11.7 Query History

From the *Query Tab* you can click on the *Command History* button to see a full, browsable and searchable query tab.



11.8 SSH Console

OmniDB also provides a full-featured SSH Console.

```

The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
Last login: Wed Jun 12 14:56:51 2019 from 45.4.238.41
omnidb@omnidb:~$ stty rows 38 cols 157
omnidb@omnidb:~$ cd static/OmnidB_app/
omnidb@omnidb:~/static/OmnidB_app$ ls -lh
total 24K
drwxrwxr-x 6 omnidb omnidb 4.0K Nov 19 2018 css
drwxrwxr-x 4 omnidb omnidb 4.0K Nov 19 2018 fa
drwxrwxr-x 4 omnidb omnidb 4.0K Nov 19 2018 fonts
drwxrwxr-x 4 omnidb omnidb 4.0K Nov 19 2018 images
drwxrwxr-x 3 omnidb omnidb 4.0K Nov 19 2018 js
drwxrwxr-x 11 omnidb omnidb 4.0K Nov 19 2018 lib
omnidb@omnidb:~/static/OmnidB_app$ cd lib/
omnidb@omnidb:~/static/OmnidB_app/lib$ ls -lh
total 44K
drwxrwxr-x 3 omnidb omnidb 12K Nov 19 2018 ace
drwxrwxr-x 4 omnidb omnidb 4.0K Nov 19 2018 aimaraJS
drwxrwxr-x 2 omnidb omnidb 4.0K Nov 19 2018 chart
drwxrwxr-x 3 omnidb omnidb 4.0K Nov 19 2018 cytoscape
drwxrwxr-x 2 omnidb omnidb 4.0K Nov 19 2018 emojiarea
drwxrwxr-x 3 omnidb omnidb 4.0K Nov 19 2018 jqplot
drwxrwxr-x 4 omnidb omnidb 4.0K Nov 19 2018 jquery-ui
drwxrwxr-x 2 omnidb omnidb 4.0K Nov 19 2018 popupJS
drwxrwxr-x 4 omnidb omnidb 4.0K Nov 19 2018 tabs
omnidb@omnidb:~/static/OmnidB_app/lib$

```

12. OmniDB Config Tool

Every installation of OmniDB also comes with a small CLI utility called *OmniDB Config*. It will have a different file name, depending on the way you installed OmniDB:

- If you are using a tarball or zip package, it is called **omnidb-config**, for both server and app versions;
- If you used an installer (like the .deb file) of server version, it is called **omnidb-config-server**;
- If you used an installer of app version, it is called **omnidb-config-app**.

Despite having different names, the utility does exactly the same. If you used an installer, it will be put in your \$PATH.

```
user@machine:~$ omnidb-config-app --help
Usage: omnidb-config-app [options]

Options:
  --version            show program\'s version number and exit
  -h, --help           show this help message and exit
  -d HOMEDIR, --homedir=HOMEDIR
                        home directory containing local databases config and
                        log files
  -c username password, --createsuperuser=username password
                        create super user: -c username password
  -a, --vacuum          databases maintenance
  -r, --resetdatabase  reset user and session databases
  -t, --deletetemp      delete temporary files
```

12.1 Set home directory

Option `-d` allows you to set the path to the OmniDB folder that contains the config and database files where you want to execute other options, like creating a new super user (`-c`).

12.2 Create super user

Option `-c` allows you to create a new super user, without needing to open OmniDB interface.

```
user@machine:~$ omnidb-config-app -c william password
Creating superuser...
Superuser created.
```

12.3 Vacuum

OmniDB has two databases:

- `omnidb.db`: Stores all users and connections, and other OmniDB related stuff;
- *Sessions database*: Stores Django user sessions.

Both databases are SQLite, so it can be useful to vacuum them sometimes to reduce file size. This can be done with the `-a` option.

```
user@machine:~$ omnidb-config-app -a
Vacuuming OmniDB database...
Done.
Vacuuming Sessions database...
Done.
```

12.4 Reset database

If you wish to wipe out all OmniDB information and get a clean database as it was just installed, you can use the `-r` option. Use it with caution!

```
user@machine:~$ omnidb-config-app -r
*** ATENTION *** ALL USERS DATA WILL BE LOST
Would you like to continue? (y/n) y
Cleaning users...
Done.
Cleaning sessions...
Vacuuming OmniDB database...
Done.
Vacuuming Sessions database...
Done.
```

12.5 Delete temporary files

If you desire to remove temporary files that OmniDB creates along its execution, like exported queries in CSV/XLSX format, you can use the `-t` option.

```
user@machine:~$ omnidb-config-app -t
Cleaning temp folder...
Done.
```

13. Writing and Debugging PL/pgSQL Functions

13.1 Introduction

PostgreSQL is more than a RDBMS engine. It is a developing platform. It provides a very powerful and flexible programming language called PL/pgSQL. Using this language you can write your own *user-defined functions* to achieve abstraction levels and procedural calculations that would be difficult to achieve with plain SQL (and sometimes impossible to achieve without context-switching with the application). While you always could develop and manage your own functions within OmniDB, it is a recent feature that allows you to also *debug* your own functions.

OmniDB 2.3.0 introduced this great feature: a debugger for PL/pgSQL functions. It was implemented by scratch and takes advantage of hooks, an extensibility in PostgreSQL's source code that allows us to perform custom actions when specific events are triggered in the database. For the debugger we use hooks that are triggered when PL/pgSQL functions are called, and each statement is executed.

This requires the user to install a binary library called `omnidb_plugin` and enable it in PostgreSQL's config file. The debugger also uses a special schema with special tables to control the whole debugging process. This can be manually created or with an extension.

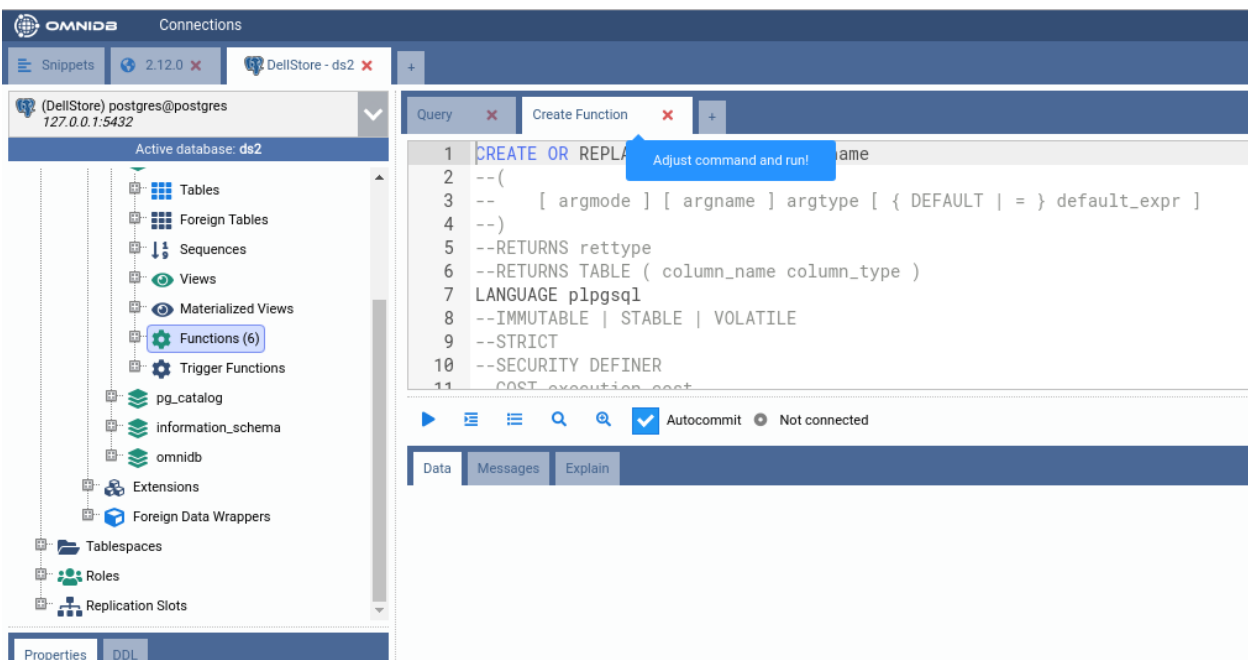
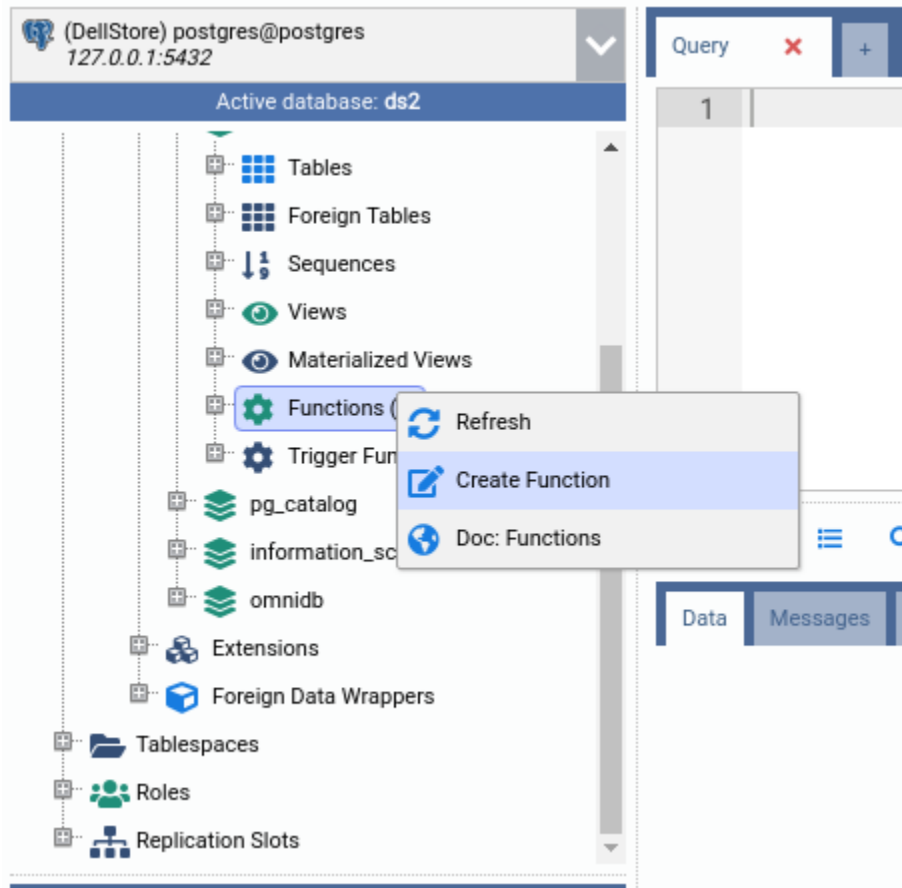
For more details on the installation, please refer to the [instructions](#), also available in Chapter 23. Also please read the notes in this document, to be aware that currently there are some limitations.

After successfully installing the debugger, you will see a schema called `omnidb` in your database. Also, if you compiled the debugger yourself, you can install it as a PostgreSQL extension, and in this case it will appear under the *Extensions* tree node.

Property	Value
Database	ds2
Schema	omnidb
Owner	postgres
ACL	{postgres=UC/postgres,ds2=UC/postgres...

13.2 Writing functions

In the `public` schema, right-click the `Functions` node and click on *Create Function*. It will open a *SQL Query* inner tab, already containing a SQL Template to help you create your first PL/pgSQL function.



You can refer to PostgreSQL documentation on how to write user-defined functions. No need to open a new browser tab: just right-click the *Functions* node and click on *Doc: Functions* to view the documentation inside OmniDB.

For now, let us replace this SQL template entirely for the source code below:

```
CREATE OR REPLACE FUNCTION public.fnc_count_vowels (p_input text)
RETURNS integer LANGUAGE plpgsql AS
$function$
DECLARE
    str text;
    ret integer;
    i integer;
    len integer;
    tmp text;
BEGIN
    str := upper(p_input);
    ret := 0;
    i := 1;
    len := length(p_input);
    WHILE i <= len LOOP
        IF substr(str, i, 1) in ('A', 'E', 'I', 'O', 'U') THEN
            SELECT pg_sleep(1) INTO tmp;
            ret := ret + 1;
        END IF;
        i := i + 1;
    END LOOP;
    RETURN ret;
END;
$function$
```

This will create a function called `fnc_count_vowels` inside the schema `public`. This function takes a text argument called `p_input` and counts how many vowels there are in this *string*. Then returns this count.

To create the function, execute the command in the SQL Query inner tab. If successful, the function will appear under the *Functions* tree node (you can refresh it by right-clicking and then clicking in *Refresh*). By expanding the function node as well, you can see its return type and its argument.

Connections: 2.12.0 x DellStore - ds2 x

(DellStore) postgres@postgres 127.0.0.1:5432

Active database: ds2

Foreign Tables, Sequences, Views, Materialized Views, Functions (7): browse_by_actor, browse_by_category, browse_by_title, **fnc_count_vowels** (returns integer, p_input text), login, new_customer, purchase, Trigger Functions

Properties | DDL

Property	Value
Database	ds2
Schema	public
Function	fnc_count_vowels

```

14 len := length(p_input),
15 WHILE i <= len LOOP
16   IF substr(str, i, 1) in ('A', 'E', 'I', 'O', 'U') THEN
17     SELECT pg_sleep(1) INTO tmp;
18     ret := ret + 1;
19   END IF;
20   i := i + 1;
21 END LOOP;
22 RETURN ret;
23 END;
24 $function$

```

Autocommit: ☒ Idle Start time: 11/03/2018 06:26:02 Duration: 21.899 ms

Data | Messages | Explain

CREATE FUNCTION

Now let us execute this new function for the first time. Open a simple SQL Query inner tab and execute the following SQL query:

```
SELECT public.fnc_count_vowels('The quick brown fox jumps over the lazy dog.')
```

Connections: 2.12.0 x DellStore - ds2 x

(DellStore) postgres@postgres 127.0.0.1:5432

Active database: ds2

Foreign Tables, Sequences, Views, Materialized Views, Functions (7): browse_by_actor, browse_by_category, browse_by_title, **fnc_count_vowels** (returns integer, p_input text), login, new_customer, purchase, Trigger Functions

Properties | DDL

Property	Value
Database	ds2
Schema	public
Function	fnc_count_vowels

```

1 SELECT public.fnc_count_vowels('The quick brown fox jumps over the lazy dog.')

```

Autocommit: ☒ Idle Number of records: 1 Start time: 11/03/2018 06:27:27 Duration: 00:00:11

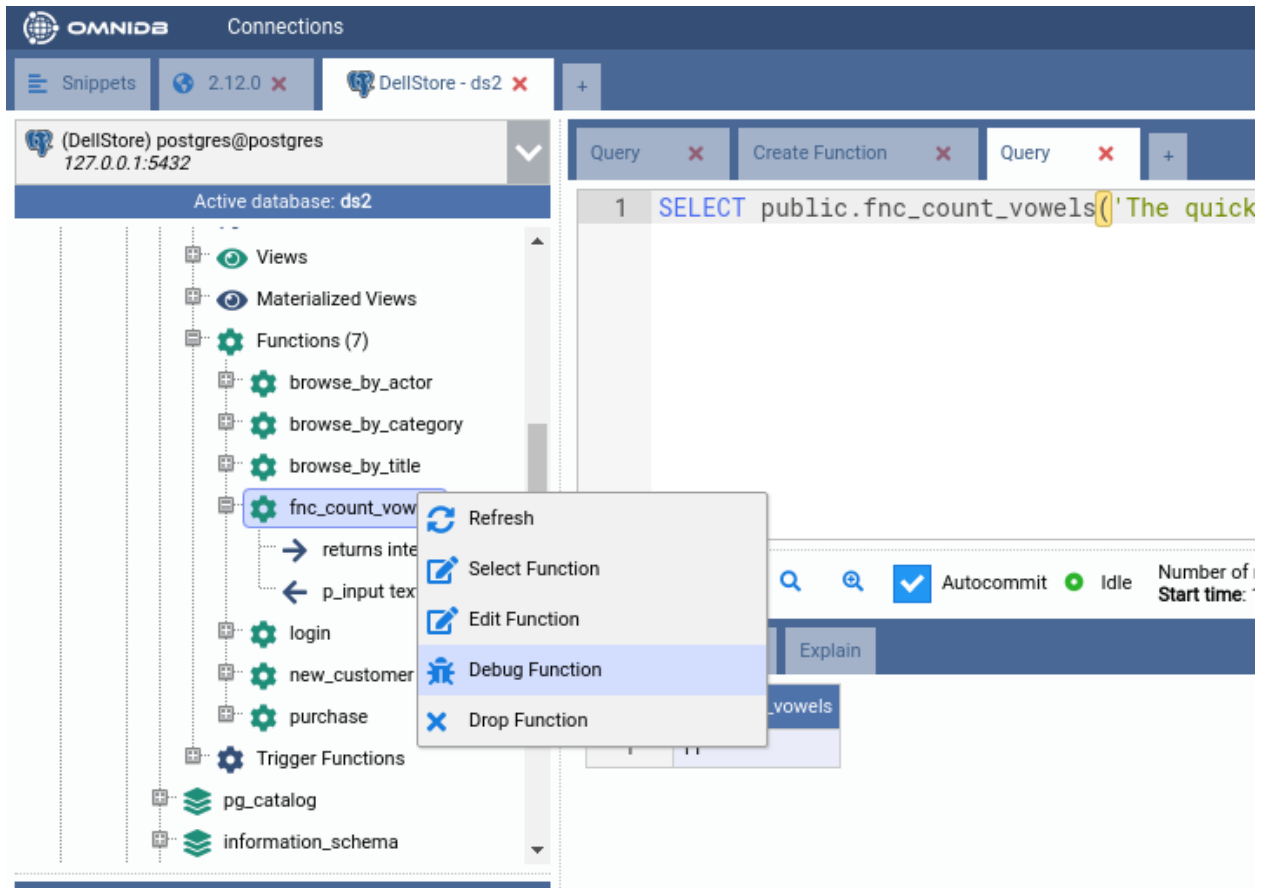
Data | Messages | Explain

fnc_count_vowels
11

Note how the query returns a single value, containing the number of vowels in the text. Note also how the query took several seconds to finish; this is caused by the `pg_sleep` we put in the source code of the function `fnc_count_vowels`.

By right-clicking the function node, you can see there are actions to edit, select and drop it. As you probably guessed, each action will open SQL Query inner tabs with handy SQL templates in them. But the most interesting action right

now is *Debug Function*. Go ahead and click it!



13.3 Debugging functions

The debugger is a specific inner tab composed of a SQL editor that will show the process step by step on top of the function source code, and 5 tabs to manage and view different parts of the debugger.

The screenshot shows the OmniDB interface with a SQL editor and a debugger panel. The SQL editor contains the following code:

```

1
2 DECLARE
3     str text;
4     ret integer;
5     i integer;
6     len integer;
7     tmp text;
8 BEGIN
9     str := upper(p_input);
10    ret := 0;
11    i := 1;

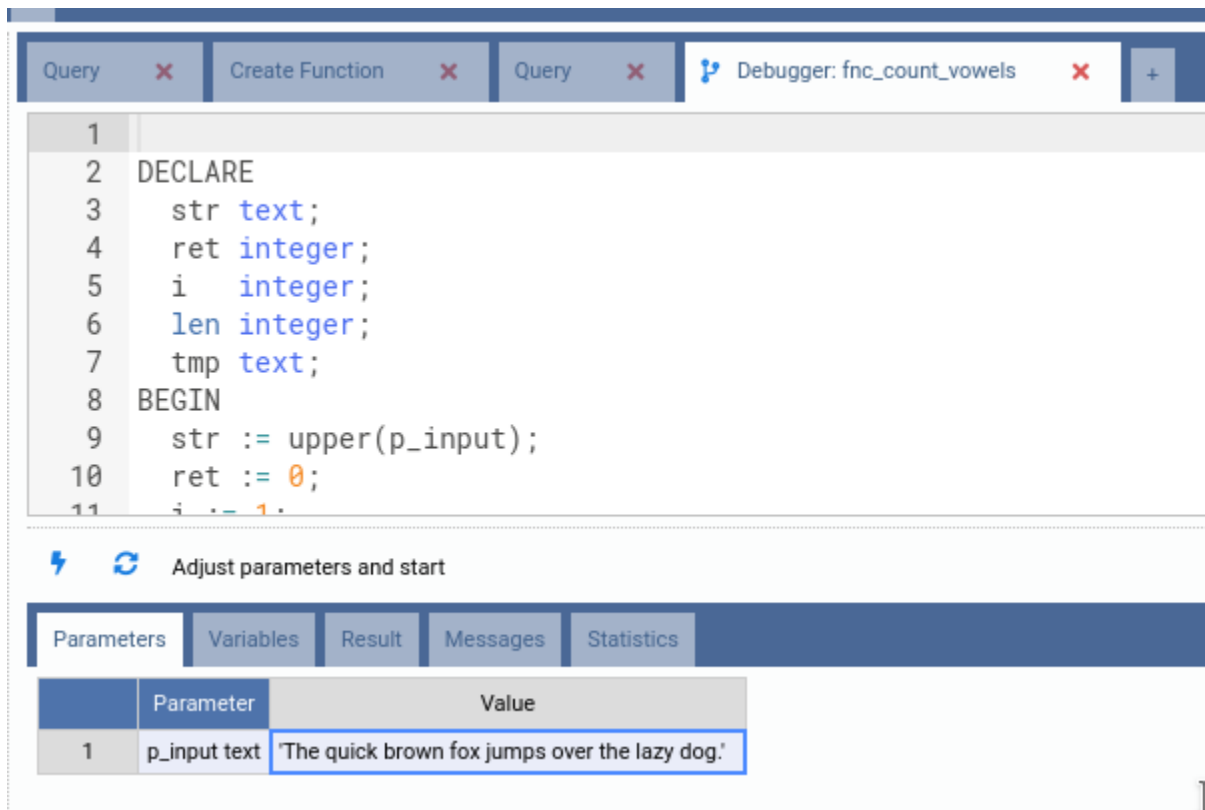
```

Below the editor is a button labeled "Adjust parameters and start" with a lightning bolt icon. The debugger panel has tabs for "Parameters", "Variables", "Result", "Messages", and "Statistics". The "Parameters" tab is active, showing a table with one parameter:

	Parameter	Value
1	p_input text	

- **Parameters:** Before the debugging process starts, the user must provide all the parameters in this tab. Parameters must be provided exactly the same way you would provide them if you were executing the function in plain SQL, quoting strings for instance;
- **Variables:** This grid displays the current value of each variable that exists in the current execution context, it will be updated with every step;
- **Result:** When the function ends, this tab will show the result of the function call. It could be empty, a single value or even a set of rows;
- **Messages:** Messages returned explicitly by RAISE commands or even automatic messages from PostgreSQL will be presented in this tab;
- **Statistics:** At the end of the debugging process, a chart depicting execution times for each line in the function body will be presented in this tab. Additionally, the SQL editor will be updated with a set of colors representing a heat map, from blue to red, according to the max duration of each line.

Now let us start debugging this function. First thing to do is to fill *every* parameter in the *Parameters* tab:



Then click on the *Start* button. Note how OmniDB automatically goes to the *Variables* tab, which is the interesting tab now that the function is being debugged. The argument `p_input` is now called `$1`, indicating the first argument of the function. Also note the variable `found`, which is a PostgreSQL reserved variable that indicates whether or not a query has returned values inside of the function.

Also note that OmniDB points to the first line of the source code of the function, highlighting it in green. This is the line that is about to be executed.

```

5  i integer;
6  len integer;
7  tmp text;
8  BEGIN
9  str := upper(p_input);
10 ret := 0;
11 i := 1;
12 len := length(p_input);
13 WHILE i <= len LOOP
14     IF substr(str, i, 1) in ('A', 'E', 'I', 'O', 'U') THEN

```

> >> Cancel Ready

	Variable	Attribute	Type	Value
1	\$1		text	The quick brown fox jumps o...
2	found		bool	f
3	str		text	NULL
4	ret		int4	NULL
5	i		int4	NULL
6	len		int4	NULL
7	tmp		text	NULL

Now click in the first button below the SQL editor. It is the *Step Over* button, and it means that OmniDB will execute the next statement and stop right after it.

Debugger: fnc_count_vowels

```

6  len integer;
7  tmp text;
8  BEGIN
9  str := upper(p_input);
10 ret := 0;
11 i := 1;
12 len := length(p_input);
13 WHILE i <= len LOOP
14     IF substr(str, i, 1) in ('A', 'E', 'I', 'O', 'U') THEN
15         SELECT pg_sleep(1) INTO tmp;

```

> >> Cancel Ready

	Variable	Attribute	Type	Value
1	\$1		text	The quick brown fox jumps o...
2	found		bool	f
3	str		text	THE QUICK BROWN FOX JU...
4	ret		int4	NULL
5	i		int4	NULL
6	len		int4	NULL
7	tmp		text	NULL

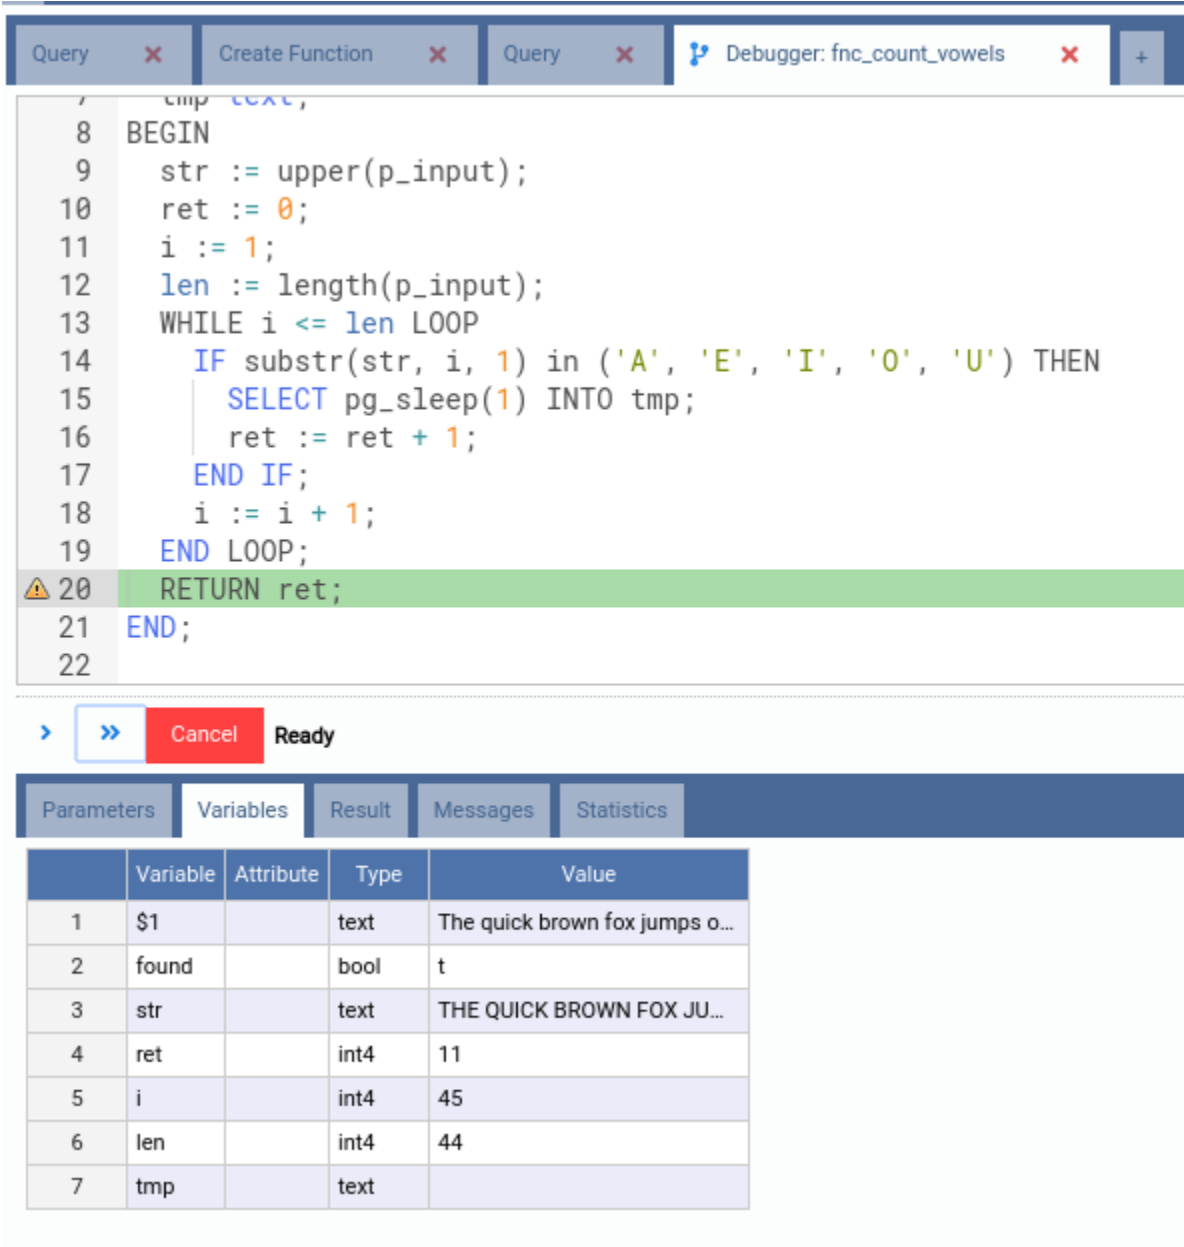
Note how the variable `str` has the value assigned to it during execution of line 9. Right now OmniDB is about to execute line 10, showing the current execution state.

Now that you know how to step over, let us speed up things a little bit. Click on the header of the line 20, the last line of code. By doing this, you just placed a *breakpoint*. The debugger interface allows you to place one breakpoint at a time.

The screenshot shows the OmniDB debugger interface. At the top, there are tabs for 'Query', 'Create Function', 'Query', and 'Debugger: fnc_count_vowels'. The main area displays the SQL code for the function 'fnc_count_vowels'. A breakpoint is set at line 20, which is highlighted with a yellow warning icon. Below the code, there is a control bar with a play button, a pause button, and buttons for 'Cancel' and 'Ready'. Below the control bar, there are tabs for 'Parameters', 'Variables', 'Result', 'Messages', and 'Statistics'. The 'Variables' tab is active, showing a table of variables and their values.

	Variable	Attribute	Type	Value
1	\$1		text	The quick brown fox jumps o...
2	found		bool	f
3	str		text	THE QUICK BROWN FOX JU...
4	ret		int4	NULL
5	i		int4	NULL
6	len		int4	NULL
7	tmp		text	NULL

After setting a breakpoint, you can click in the second button, *Resume*. OmniDB will carry on with the debugging process until it reaches the line of code with the breakpoint. This may take a while because of the `pg_sleep` commands we put in the source code. Note that if you click this button without previously setting a breakpoint, OmniDB will execute the entire function to the end.



The screenshot shows the OmniDB debugger interface. At the top, there are tabs for 'Query', 'Create Function', 'Query', and 'Debugger: fnc_count_vowels'. The main window displays the PL/pgSQL code for the function 'fnc_count_vowels'. A breakpoint is set at line 20, which is highlighted in green. Below the code, there is a 'Ready' status bar with a 'Cancel' button. At the bottom, there is a table showing the current state of variables.

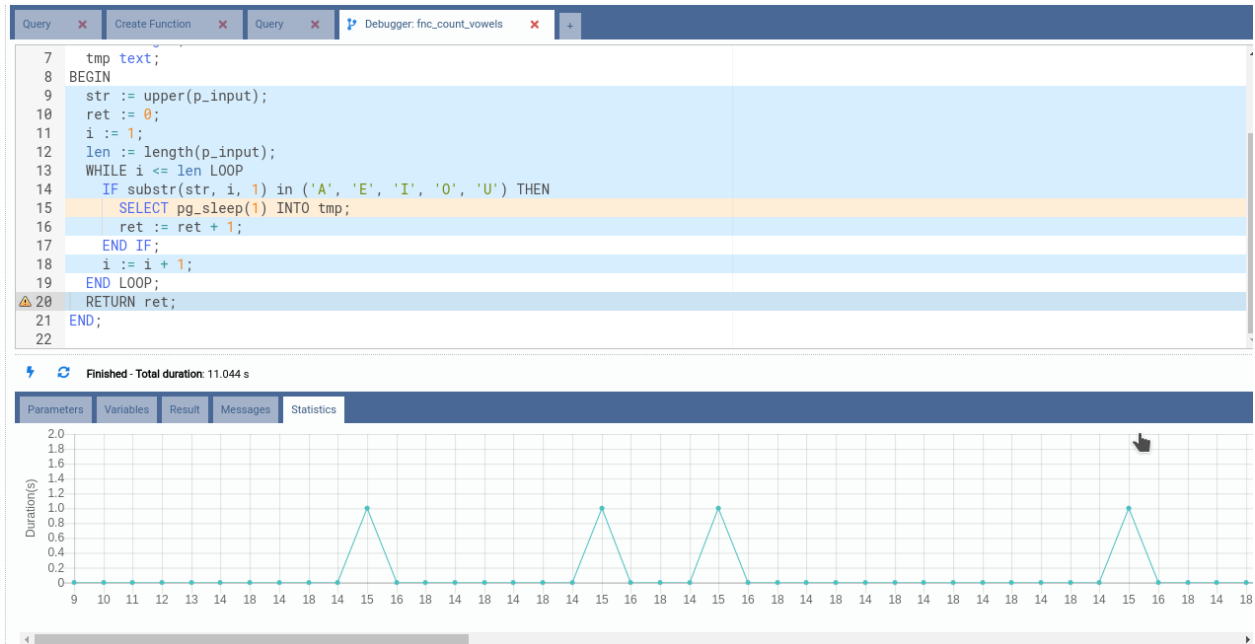
```

7  tmp text,
8  BEGIN
9    str := upper(p_input);
10   ret := 0;
11   i := 1;
12   len := length(p_input);
13   WHILE i <= len LOOP
14     IF substr(str, i, 1) in ('A', 'E', 'I', 'O', 'U') THEN
15       SELECT pg_sleep(1) INTO tmp;
16       ret := ret + 1;
17     END IF;
18     i := i + 1;
19   END LOOP;
20   RETURN ret;
21 END;
22

```

	Variable	Attribute	Type	Value
1	\$1		text	The quick brown fox jumps o...
2	found		bool	t
3	str		text	THE QUICK BROWN FOX JU...
4	ret		int4	11
5	i		int4	45
6	len		int4	44
7	tmp		text	

Observe the values for each variable. We can see that the value of `ret` is 11 even before the function finishes. Also note that OmniDB does not remove the breakpoint you placed. To do that, you can click in the breakpoint little icon. Now hit *Resume* again. Let us see now what happens when the function finishes.



OmniDB will go automatically to the *Statistics* tab, which shows 2 interesting features:

- *Sum of Duration per Line of Code Chart*: in the bottom, a chart represents total duration of the function distributed in the lines of code. With this chart, you can easily spot bottlenecks in your code. In our example, it was line 15, which we deliberately put a `pg_sleep(1)` call;
- *Colored lines of source code*: OmniDB colors the lines accordingly to the numbers seen in the chart. Colors vary from blue (small duration), passing through yellow (medium duration) until red (high duration), as in a *temperature diagram*.

Also note the *Total duration* message, which shows execution time of the function, without considering the time you spent analyzing it.

13.4 Inspecting record attribute values

An interesting feature that we do not usually see in other debuggers is the ability to inspect each attribute of a variable of type `record`. OmniDB debugger does that as it is split into different variables, allowing you to see the value and type of each attribute.

To illustrate that, let us create another function, similar to the previous one, but now called `fnc_count_vowels2`:

```

CREATE OR REPLACE FUNCTION public.fnc_count_vowels2 (p_input text)
RETURNS integer LANGUAGE plpgsql AS
$function$
DECLARE
    str text;
    i integer;
    len integer;
    rec record;
BEGIN
    str := upper(p_input);
    i := 1;
    len := length(p_input);
    SELECT 0 AS a, 0 AS e, 0 AS i, 0 AS o, 0 AS u INTO rec;

```

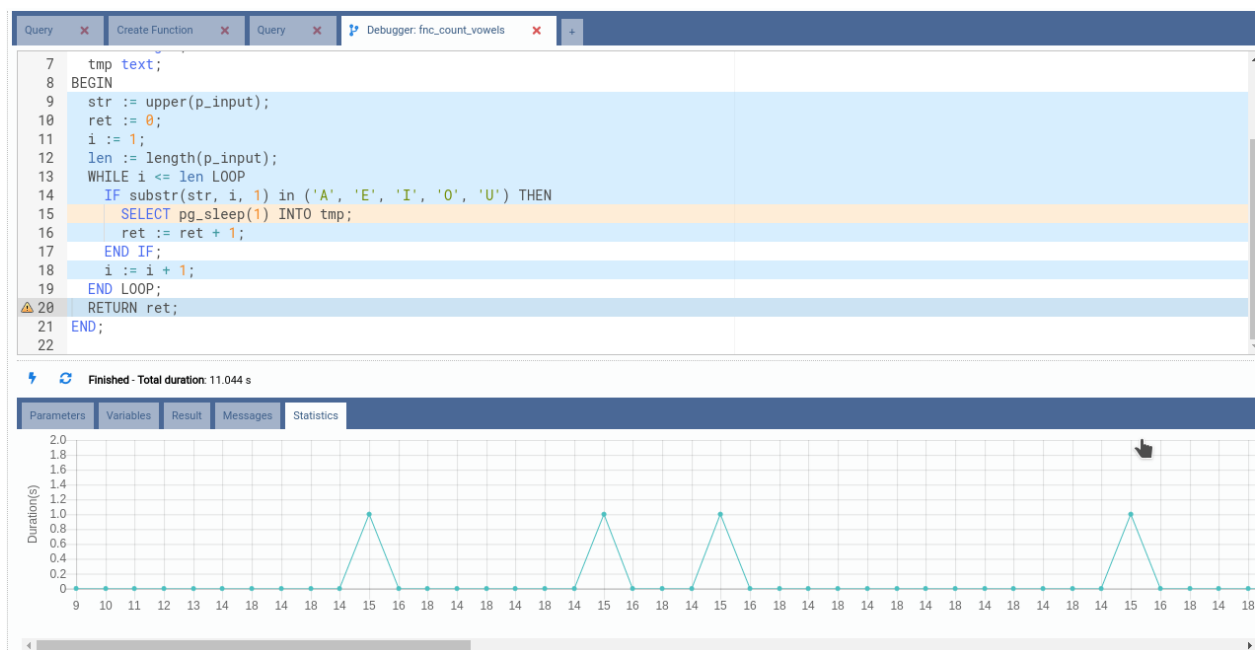
(continues on next page)

(continued from previous page)

```

WHILE i <= len LOOP
  CASE substr(str, i, 1)
    WHEN 'A' then rec.a := rec.a + 1;
    WHEN 'E' then rec.e := rec.e + 1;
    WHEN 'I' then rec.i := rec.i + 1;
    WHEN 'O' then rec.o := rec.o + 1;
    WHEN 'U' then rec.u := rec.u + 1;
    ELSE NULL;
  END CASE;
  i := i + 1;
END LOOP;
RETURN rec.a + rec.e + rec.i + rec.o + rec.u;
END;
$function$

```



Observe how we keep track of every vowel count individually. Now let us start debugging it, using the same text as before ('The quick brown fox jumps over the lazy dog.');

Debugger: fnc_count_vowels2

```

1
2 DECLARE
3   str text;
4   i integer;
5   len integer;
6   rec record;
7 BEGIN
8   str := upper(p_input);
9   i := 1;
10  len := length(p_input);
11  SELECT 0 AS a, 0 AS e, 0 AS i, 0 AS o, 0 AS u INTO rec;
12  WHILE i <= len LOOP
13    CASE substr(str, i, 1)
14      WHEN 'A' then rec.a := rec.a + 1;
15      WHEN 'E' then rec.e := rec.e + 1;
16      WHEN 'I' then rec.i := rec.i + 1;
17      WHEN 'O' then rec.o := rec.o + 1;
18      WHEN 'U' then rec.u := rec.u + 1;
19      ELSE NULL;

```

Adjust parameters and start

Parameter	Value
p_input text	'The quick brown fox jumps over the lazy dog.'

Debugger: fnc_count_vowels2

```

1
2 DECLARE
3   str text;
4   i integer;
5   len integer;
6   rec record;
7 BEGIN
8   str := upper(p_input);
9   i := 1;
10  len := length(p_input);
11  SELECT 0 AS a, 0 AS e, 0 AS i, 0 AS o, 0 AS u INTO rec;
12  WHILE i <= len LOOP
13    CASE substr(str, i, 1)
14      WHEN 'A' then rec.a := rec.a + 1;
15      WHEN 'E' then rec.e := rec.e + 1;
16      WHEN 'I' then rec.i := rec.i + 1;
17      WHEN 'O' then rec.o := rec.o + 1;
18      WHEN 'U' then rec.u := rec.u + 1;
19      ELSE NULL;

```

Cancel Ready

Variable	Attribute	Type	Value
\$1		text	The quick brown fox jumps o...
found		bool	f
str		text	NULL
i		int4	NULL
len		int4	NULL
__Case_Variable_16__		int4	NULL

Note from the picture above that PostgreSQL created an internal *Case Variable*. Also note that the variable `rec` is not shown in the list of known variables. This is because PostgreSQL still does not know what attributes `rec` will contain. Let's step over some more steps.

```

3  str text;
4  i integer;
5  len integer;
6  rec record;
7  BEGIN
8  str := upper(p_input);
9  i := 1;
10 len := length(p_input);
11 SELECT 0 AS a, 0 AS e, 0 AS i, 0 AS o, 0 AS u INTO rec;
12 WHILE i <= len LOOP
13 CASE substr(str, i, 1)
14 WHEN 'A' then rec.a := rec.a + 1;
15 WHEN 'E' then rec.e := rec.e + 1;

```

Parameters Variables Result Messages Statistics

	Variable	Attribute	Type	Value
1	\$1		text	The quick brown fox jumps o...
2	found		bool	t
3	str		text	THE QUICK BROWN FOX JU...
4	i		int4	1
5	len		int4	44
6	rec	a	int4	0
7	rec	e	int4	0
8	rec	i	int4	0
9	rec	o	int4	0
10	rec	u	int4	0
11	_Case_Variable_16_...		int4	NULL

Right after the execution of line 11, `rec` variable comes to life and we can see it has 5 attributes: `a`, `e`, `i`, `o` and `u`, all of the type `int` and having initial value 0.

Now set a breakpoint in line 23 and click the *Resume* button.

```

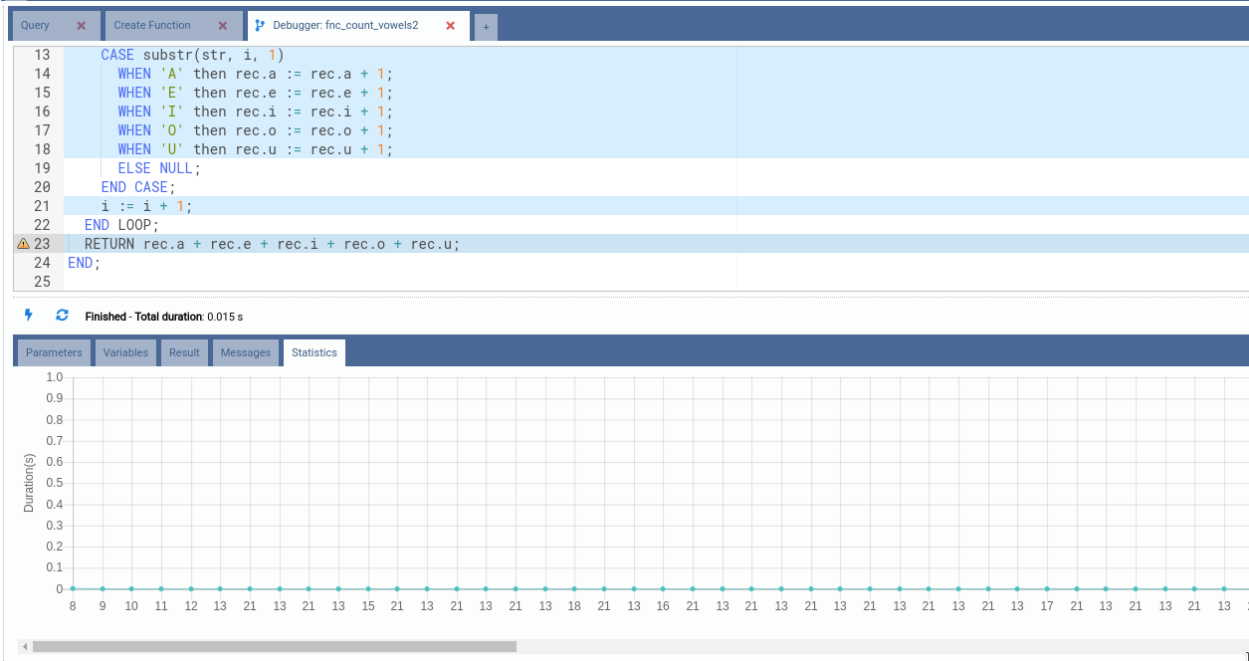
13 CASE substr(str, i, 1)
14 WHEN 'A' then rec.a := rec.a + 1;
15 WHEN 'E' then rec.e := rec.e + 1;
16 WHEN 'I' then rec.i := rec.i + 1;
17 WHEN 'O' then rec.o := rec.o + 1;
18 WHEN 'U' then rec.u := rec.u + 1;
19 ELSE NULL;
20 END CASE;
21 i := i + 1;
22 END LOOP;
23 RETURN rec.a + rec.e + rec.i + rec.o + rec.u;
24 END;
25

```

Parameters Variables Result Messages Statistics

	Variable	Attribute	Type	Value
1	\$1		text	The quick brown fox jumps o...
2	found		bool	t
3	str		text	THE QUICK BROWN FOX JU...
4	i		int4	45
5	len		int4	44
6	rec	a	int4	1
7	rec	e	int4	3
8	rec	i	int4	1
9	rec	o	int4	4
10	rec	u	int4	2
11	_Case_Variable_16_...		text	NULL

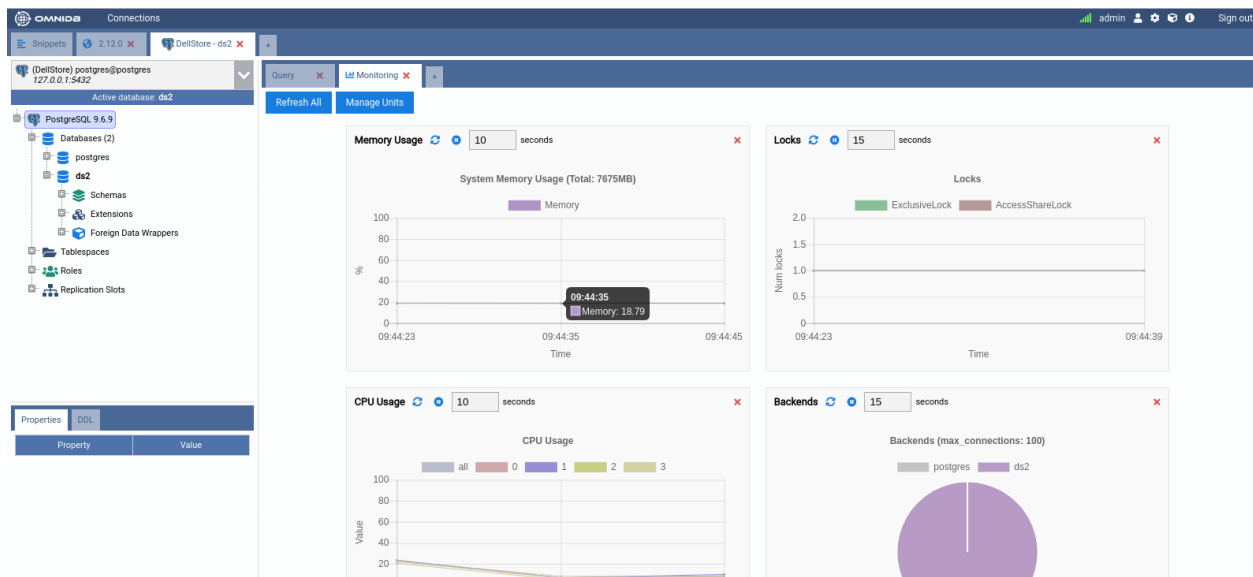
See how we can inspect every attribute, observing how many of each vowel the text contain. Now let's finish this function.



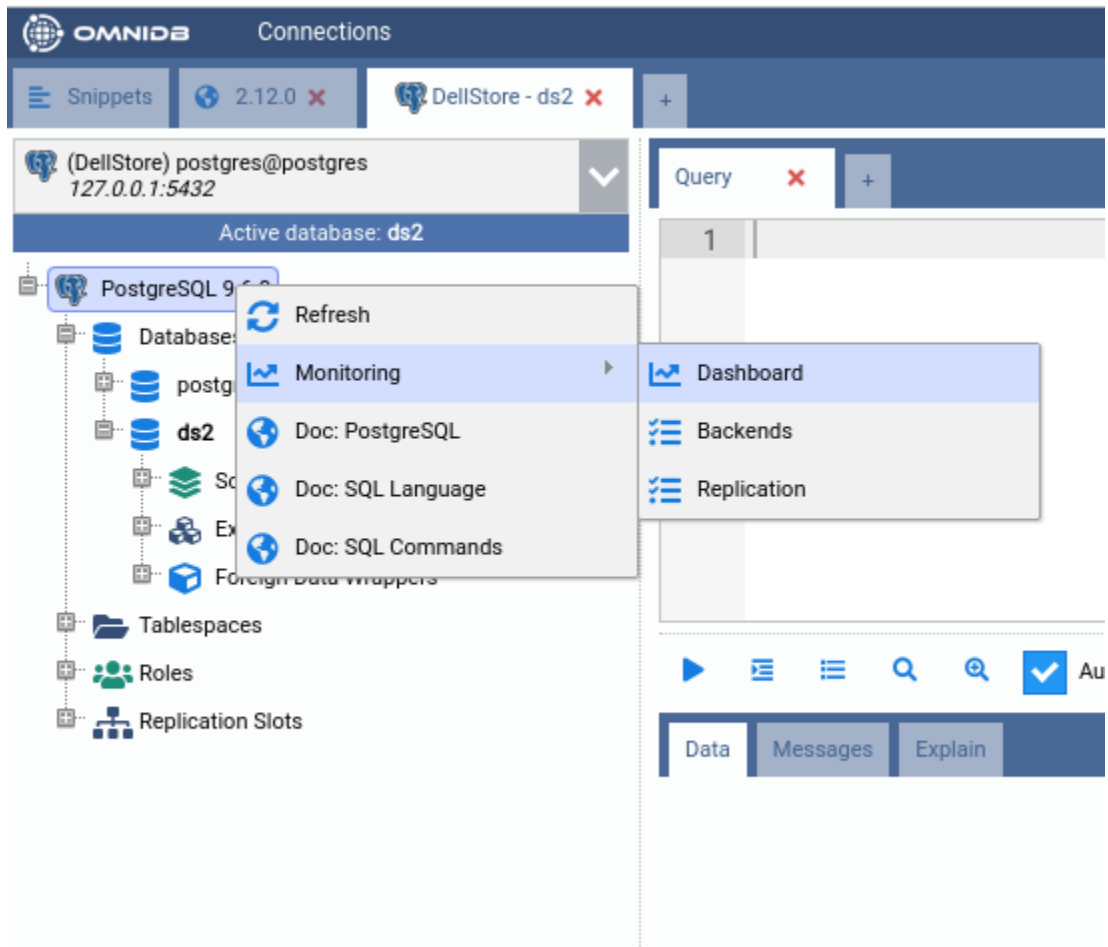
CHAPTER 14

14. Monitoring Dashboard

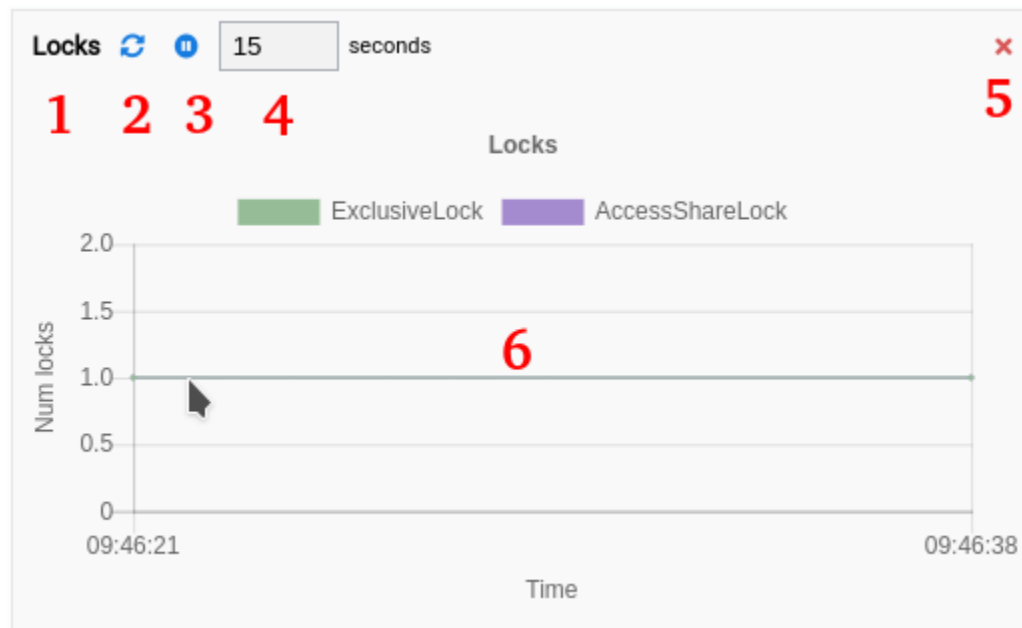
OmniDB 2.4.0 introduced a new cool feature called *Monitoring Dashboard*. We know a picture is worth a thousand words, so please take a look:



As you can see, this is a new kind of inner tab showing some charts and grids. This *Monitoring* inner tab is automatically opened once you expand the tree root node (the *PostgreSQL* node). You can keep it open or close it at any time. To open it again, right-click the root node and click on *Dashboard*.



The dashboard is composed of handy information rectangles called *Monitoring Units*. Here is an example of Monitoring Unit and its interface elements:



- **1:** Title of the Monitoring Unit;
- **2:** Refresh the Monitoring Unit. Depending on the type, clicking on this button will refresh the entire drawing area or just make the chart acquire a new set of values;
- **3:** Pause the Monitoring Unit;
- **4:** Interval in seconds for automatic refreshing;
- **5:** Remove the Monitoring Unit of the Monitoring Dashboard;
- **6:** Drawing area, that will be different depending on the type of the Monitoring Unit.

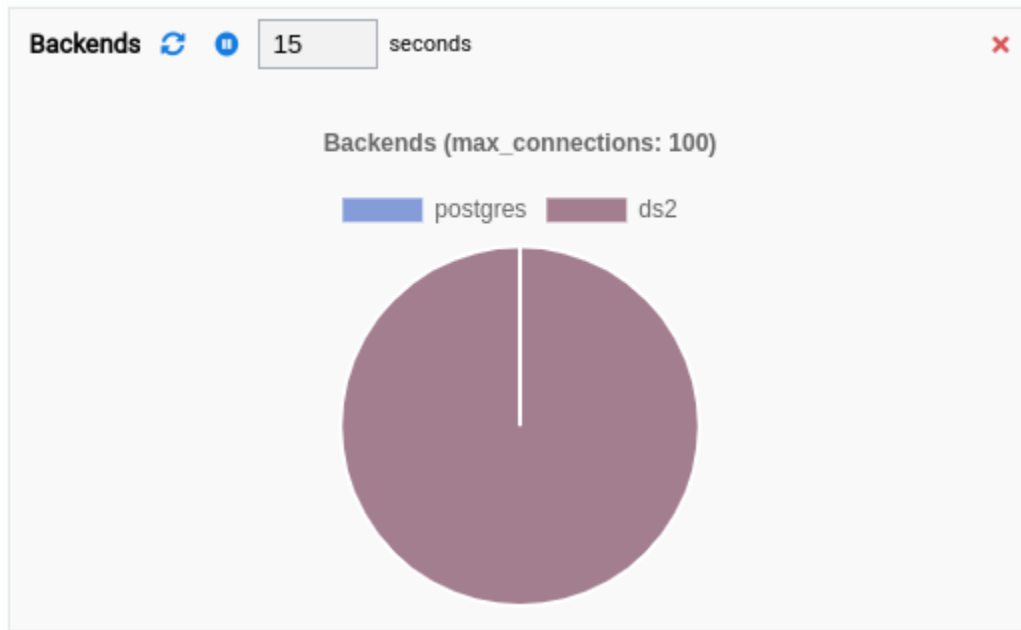
14.1 Types of Monitoring Units

Currently there are 3 types of Monitoring Units:

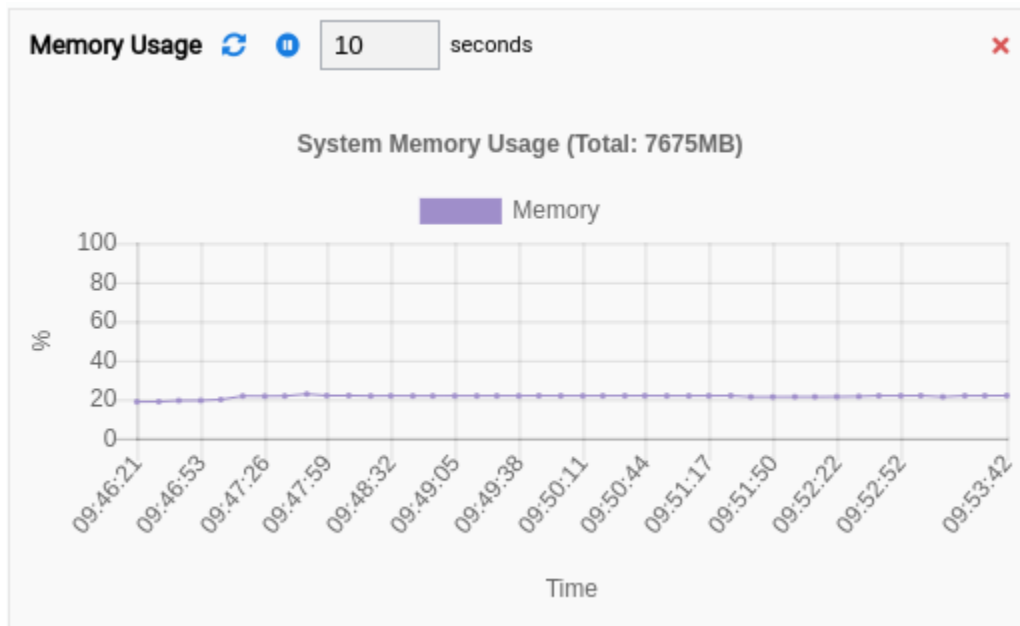
- **Grid:** The most simple kind, just executes a query from time to time and shows the results in a data grid.

	datid	datname	pid	usesysid	username	application_name	client_addr	client
1	16387	ds2	8037	10	postgres	OmniDB	127.0.0.1	
2	16387	ds2	8036	10	postgres	OmniDB	127.0.0.1	

- **Chart:** Every time it refreshes, it renders a new complete chart. The old set of values is lost. This is most useful for pie charts, but other kind of charts can be used too.

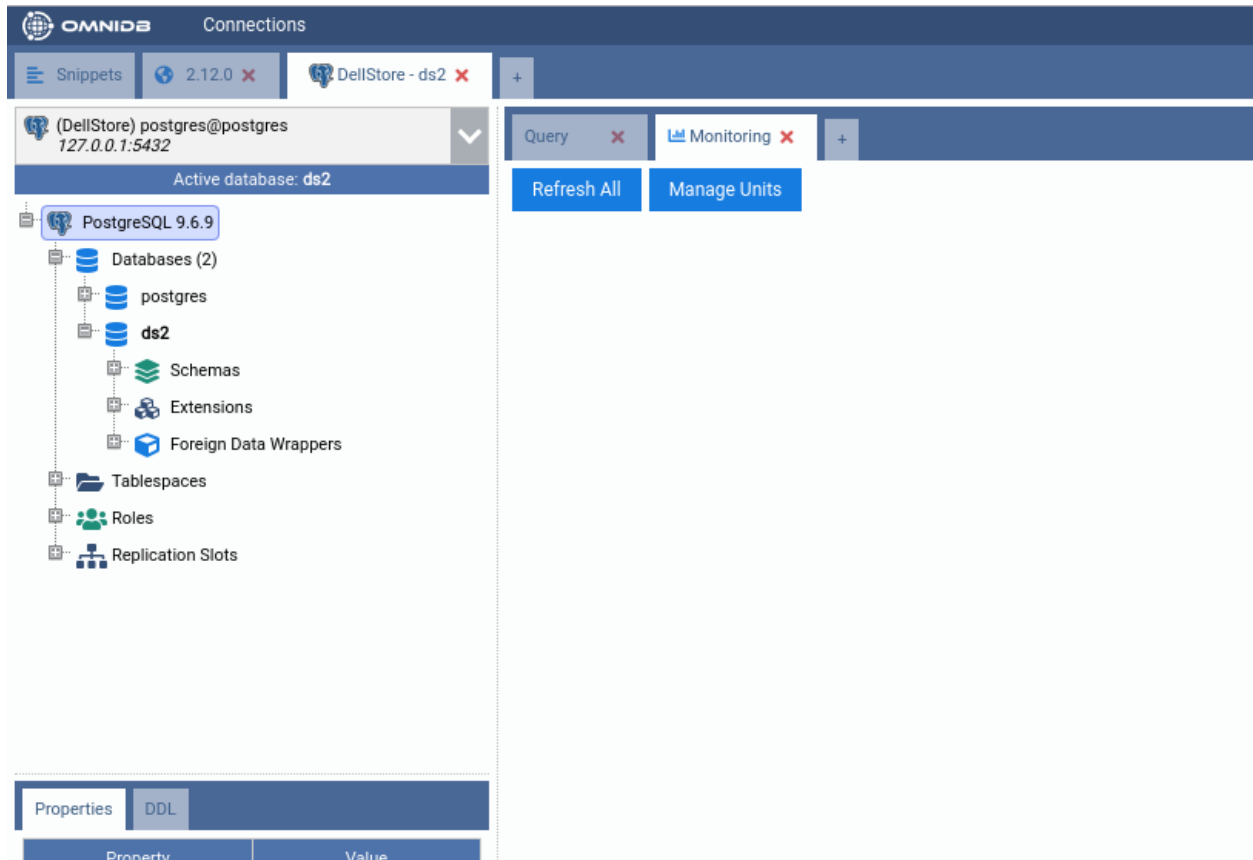


- **Chart-Append:** Perhaps this is the most useful kind of Monitoring Unit. It is a chart that appends a new set of values every time it refreshes. Line or bar charts fit best for this type. The last 50 set of values are kept by the component client-side to be viewed by the user.



14.2 Showing and hiding units in the dashboard

If you click in the button *Refresh All*, then all Monitoring Units will be refreshed at once. You can also remove undesired Monitoring Units by clicking in the *Remove* button. Let us go ahead and remove all units from the dashboard, making it empty:



All Monitoring Units that come with OmniDB are open source and available in this [repository](#) (feel free to contribute). But be aware that some Monitoring Units require the `plpythonu` script to be installed in the database. Please refer to the instructions specific to your operating system on how to install `plpythonu` if you desire to use and create Monitoring Units that use `plpythonu`.

The screenshot displays the OmniDB interface with a PostgreSQL 9.6.9 connection named 'ds2'. The left sidebar shows the database hierarchy: Databases (2) including 'postgres' and 'ds2', Schemas, Extensions (2) including 'plpgsql' and 'plpythonu', Foreign Data Wrappers, Tablespaces, Roles, and Replication Slots. The 'Active database: ds2' is selected. The main panel shows a SQL query editor with the following code:

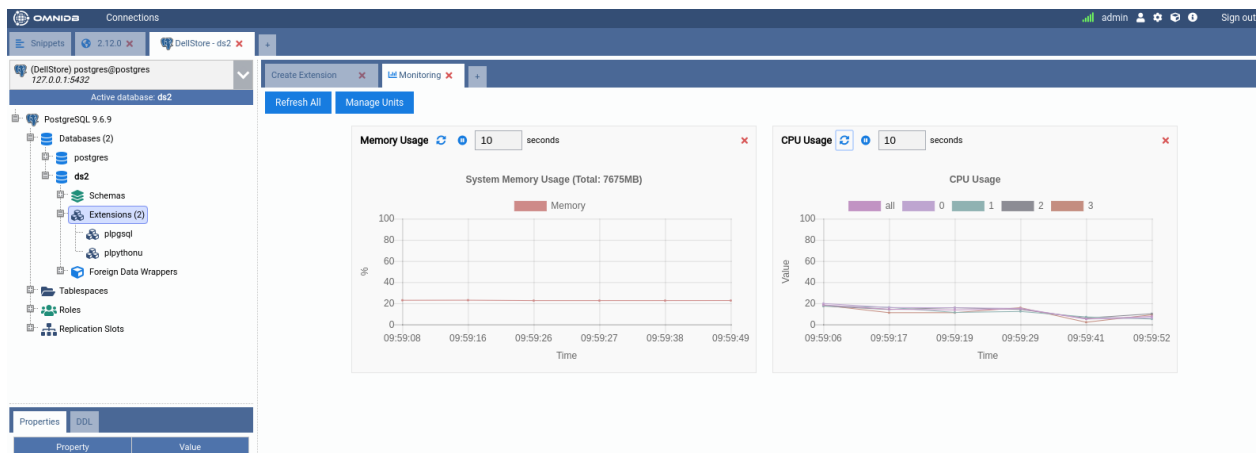
```
1 CREATE EXTENSION plpythonu
2 --SCHEMA schema_name
3 --VERSION VERSION
4 --FROM old_version
5
```

Below the query editor, there are buttons for 'Data', 'Messages', and 'Explain'. The 'Data' button is active, and the text 'CREATE EXTENSION' is displayed below it. The interface also includes a top bar with 'Connections' and a bottom bar with 'Autocommit' status.

Now that our dashboard is empty, let us add some units. Click on the *Manage Units* button.

	Actions	Title	Type	Interval(s)
1	✓	Backends	Chart	15
2	✓	Database Size	Chart	30
3	✓	Backends	Chart (Append)	5
4	✓	Database Size	Chart (Append)	30
5	✓	Size: Top 5 Tables	Chart (Append)	15
6	✓	CPU Usage	Chart (Append)	10
7	✓	Locks	Chart (Append)	15
8	✓	Bloat: Top 5 Tables	Chart (Append)	45
9	✓	Master: Replication Lag	Chart (Append)	15
10	✓	Standby: Replication Lag (Size)	Chart (Append)	15
11	✓	Standby: Replication Lag (Time)	Chart (Append)	15
12	✓	Memory Usage	Chart (Append)	10
13	✓	Longest Active Query	Chart (Append)	5
14	✓	WAL Folder Size	Chart (Append)	30
15	✓	Activity	Grid	15

Click on the green action to add the Monitoring Units called *CPU Usage* and *Memory Usage*. Bear in mind that both units require `plpythonu` extension in the database. *CPU Usage* also requires that the tool `mpstat` should be installed in the server. Also both units are of type *Chart-Append*. Wait for some seconds and you will have a dashboard like this:



In a similar way, you can add and remove any unit you want to customize the dashboard the way you want.

14.3 Writing custom Monitoring Units: Grid

OmniDB provides you the power to write your own units and customize existing ones. Everything is done through Python scripts that run inside a sandbox. Beware that units powered by `plpythonu` can have access to the file system the database user also has access to, and any Monitoring Unit have the same permission as the database user you configured in the Connection.

To create a new Monitoring Unit, click on the *Manage Units* button in the dashboard, then click on the *New Unit* button. It will open a new kind of inner tab like this:

The easiest way to write a custom unit is to use an existing one as template. Go ahead and select the *(Grid) Activity* template:

Note how OmniDB fills the *Data Script* source code. This script is responsible for generating data for the unit every time it refreshes. As a grid unit is nothing else but a grid of data, we can rely on only this script for now.

Now let us take a look at the source code of this template:

```
from datetime import datetime

data = connection.Query('''
    SELECT *
    FROM pg_stat_activity
''')
```

(continues on next page)

(continued from previous page)

```
result = {
    "columns": data.Columns,
    "data": data.Rows
}
```

It is simple enough. It executes an SQL query into the current connection using the reserved `connection` variable. Also, the grid unit type expects its results in a JSON variable that must be called `result` and must have the attributes "columns" (an array of column names) and "data" (an array of rows, each row being an array of values). The `connection.Query()` function already does the job pretty well, so let us just change the SQL query this way:

```
from datetime import datetime

data = connection.Query('''
    SELECT random() as "Random Number"
''')

result = {
    "columns": data.Columns,
    "data": data.Rows
}
```

Copy and paste the above Python code into the *Data Script* text field and then click on the *Test* (lightning) button:

The screenshot shows the OmniDB interface for creating a monitoring unit. At the top, there are buttons for 'Create Extension', 'Monitoring', and 'Monitor Unit'. Below these, the 'Name' field is empty, 'Type' is set to 'Grid', and 'Refresh Interval' is set to '15 seconds'. The 'Template' dropdown is set to '(Grid) Activity'. The 'Data Script' field contains the following Python code:

```
1 from datetime import datetime
2
3 data = connection.Query('''
4     SELECT random() as "Random Number"
5 ''')
6
7 result = {
8     "columns": data.Columns,
9     "data": data.Rows
10 }
```

The 'Test' button (lightning bolt) is highlighted. To the right of the 'Data Script' field is a 'Chart Script' field, which is currently empty. Below the 'Data Script' field, a preview table is shown with the following data:

	Random Number
1	0.310025388840586

Note how the grid was rendered in the preview drawing area. You can click the *Test* button as many times as you want. Now we will give the unit a name, set a refresh interval and then hit the *Save* button:

Monitor unit saved.

Ok

	Random Number
1	0.310025388840586

Click the *OK* button and then close the edit tab. Our new Monitoring Unit will be in the list of available units. As we created this unit, we can either add it to the dashboard, edit it or remove it. Let us add it to the dashboard (green action):

Random Number 3 seconds 1 rows

	Random Number
1	0.123685032594949

Memory Usage 10 seconds

System Memory Usage (Total: 7675MB)

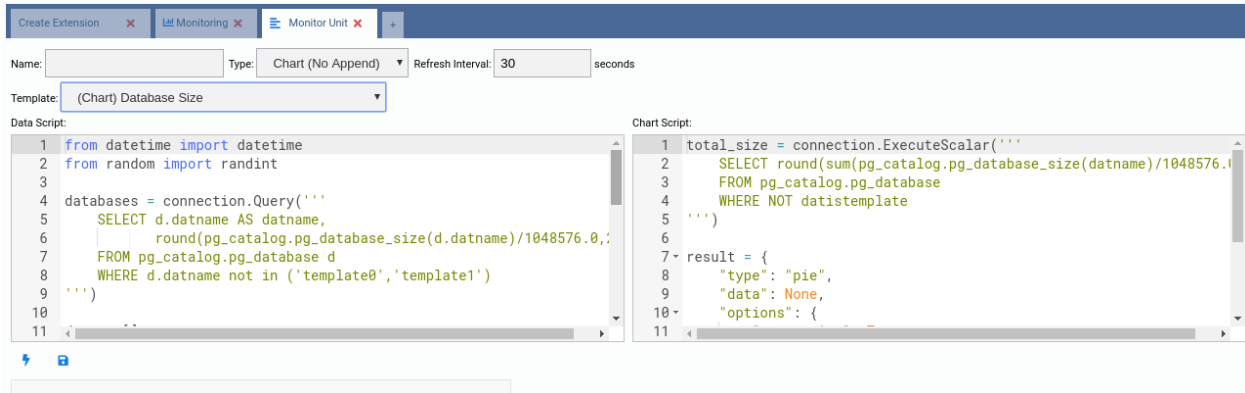
CPU Usage 10 seconds

CPU Usage

all 0 1 2 3

14.4 Writing custom Monitoring Units: Chart

Click in the *Manage Units* button and then in the *New Unit* button. This time we will create a Chart Monitoring Unit. So choose *(Chart) Database Size* as a template.



The source code of this kind of unit is more complex. There are two scripts:

- **Data Script:** Executed every time the unit is refreshed;
- **Chart Script:** Executed only at the beginning to build the chart.

The chart units are based in the component [Chart.js](#) and each chart type contains a specific JSON structure. The best approach to build new chart units is to start from a template and also check the [Chart.js docs](#) to see every property that can be added to make the output even better for each situation.

Let us take a look at the **Data Script**:

```
from datetime import datetime
from random import randint

databases = connection.Query('''
    SELECT d.datname AS datname,
           round(pg_catalog.pg_database_size(d.datname)/1048576.0,2) AS size
    FROM pg_catalog.pg_database d
    WHERE d.datname not in ('template0','template1')
''')

data = []
color = []
label = []

for db in databases.Rows:
    data.append(db["size"])
    color.append("rgb(" + str(randint(125, 225)) + "," + str(randint(125, 225)) + "," + str(
        randint(125, 225)) + ")")
    label.append(db["datname"])

result = {
    "labels": label,
    "datasets": [
        {
            "data": data,
            "backgroundColor": color,
            "label": "Dataset 1"
        }
    ]
}
```

Here we can see that the reserved variable `connection` is still being used to retrieve data from the database. Bear in mind that this variable is always pointing to the current `Connection`.

This template is for a Pie chart, which contains only one dataset and three arrays for the data:

- data: One value per slice;
- color: One color per slice;
- label: One label per slice.

This way, `data[0]`, `color[0]` and `label[0]` refer to the first slice, while `data[1]`, `color[1]` and `label[1]` refer to the second slice, and so on.

This script must return a variable called `result` and also needs to be a JSON like in the above script.

So right now you are probably guessing that you just need to change the SQL query to make the chart behave different. Well, in terms of data and datasets, you guessed right. So let's change the SQL query of this chart to compare sizes of tables of schema `public`. Also change the references from `datname` to `tablename`, as we have changed the column name.

```
from datetime import datetime
from random import randint

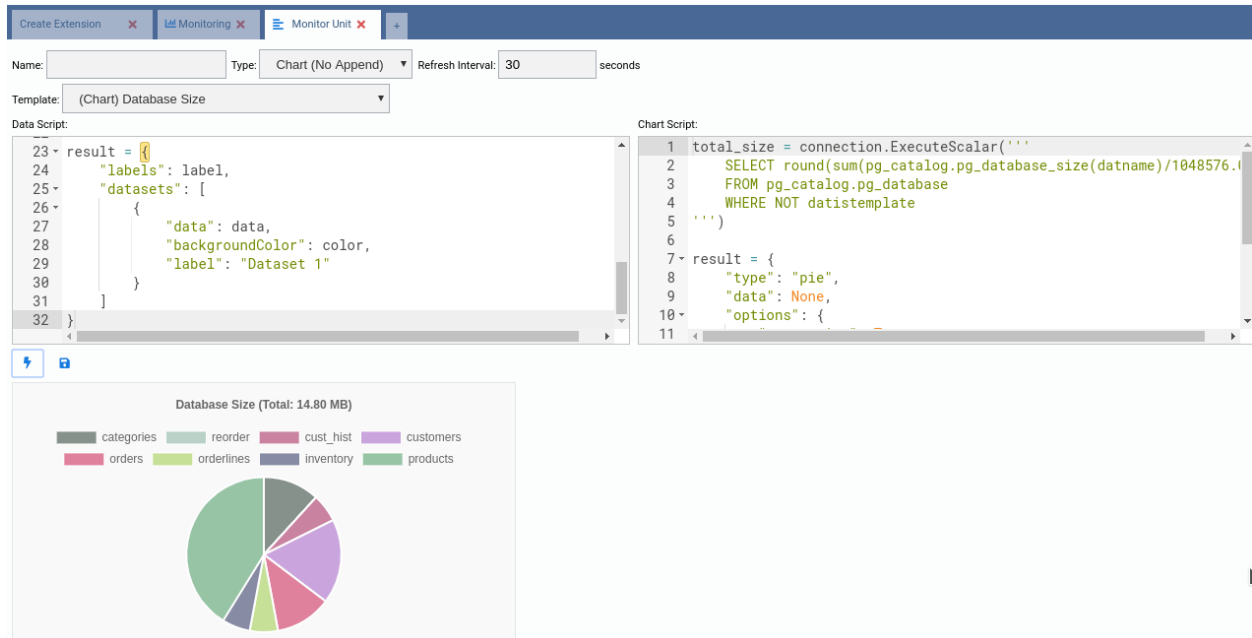
databases = connection.Query('''
    SELECT c.relname as tablename,
           round(pg_catalog.pg_total_relation_size(c.oid)/1048576.0,2) AS size
    FROM pg_catalog.pg_class c
    INNER JOIN pg_catalog.pg_namespace n
    ON n.oid = c.relnamespace
    WHERE n.nspname = 'public'
       AND c.relkind = 'r'
''')

data = []
color = []
label = []

for db in databases.Rows:
    data.append(db["size"])
    color.append("rgb(" + str(randint(125, 225)) + "," + str(randint(125, 225)) + "," + str(
        randint(125, 225)) + ")")
    label.append(db["tablename"])

result = {
    "labels": label,
    "datasets": [
        {
            "data": data,
            "backgroundColor": color,
            "label": "Dataset 1"
        }
    ]
}
```

Copy and paste the above script into the *Data Script* field and then hit the *Test* button:



Apparently the chart is almost done. We need to fix the title, it still says *Database Size*, when this chart is about table size. Any information about the format of the chart itself is defined in the *Chart Script* text field. Let us understand the current source code:

```
total_size = connection.ExecuteScalar('''
    SELECT round(sum(pg_catalog.pg_database_size(datname))/1048576.0),2)
    FROM pg_catalog.pg_database
    WHERE NOT datistemplate
''')

result = {
    "type": "pie",
    "data": None,
    "options": {
        "responsive": True,
        "title": {
            "display": True,
            "text": "Database Size (Total: " + str(total_size) + ")"
        }
    }
}
```

Easy enough. We can make use of the reserved variable `connection` to retrieve data in the *Chart Script* too. This is mainly used to put information in the chart title. The variable `result` must be defined here. Note how its JSON value defines a pie chart and the title. So we just need to change the query and adjust the title, this way:

```
total_size = connection.ExecuteScalar('''
    SELECT round(sum(pg_catalog.pg_total_relation_size(c.oid))/1048576.0),2) AS size
    FROM pg_catalog.pg_class c
    INNER JOIN pg_catalog.pg_namespace n
    ON n.oid = c.relnamespace
    WHERE n.nspname = 'public'
    AND c.relkind = 'r'
''')
```

(continues on next page)

(continued from previous page)

```

result = {
    "type": "pie",
    "data": None,
    "options": {
        "responsive": True,
        "title": {
            "display": True,
            "text": "Table Size (Total: " + str(total_size) + ")"
        }
    }
}

```

Copy and paste the above Python code into the *Chart Script*. Then click in the *Test* button:

The screenshot shows the OmniDB interface with the following components:

- Top Bar:** Contains tabs for "Create Extension", "Monitoring", and "Monitor Unit".
- Form Fields:**
 - Name:** A text input field.
 - Type:** A dropdown menu set to "Chart (No Append)".
 - Refresh Interval:** A text input field set to "30" with the unit "seconds".
 - Template:** A dropdown menu set to "(Chart) Database Size".
- Data Script Editor:** A code editor on the left showing a Python script that defines a chart configuration. The script includes a dictionary for "result" with "type", "data", and "options" (including "responsive" and "title").
- Chart Script Editor:** A code editor on the right showing the same Python script as the Data Script editor.
- Test Button:** A blue button with a lightning bolt icon, used to execute the script.
- Chart Preview:** A pie chart titled "Table Size (Total: 0.17)". The chart is divided into several segments, each representing a different database table. A legend above the chart identifies the segments: categories, reorder, cust_hist, customers, orders, orderlines, inventory, and products.

Now that the chart finally works the way we want, we can give it a title, adjust the refresh interval and then click in the *Save* button. After that we can add it to the dashboard.

The screenshot shows the 'Monitor Unit' configuration window in OmniDB. The 'Name' is 'Table Size', 'Type' is 'Chart (No Append)', and 'Refresh Interval' is '5 seconds'. The 'Template' is '(Chart) Database Size'. The 'Data Script' is:

```
23 result = {}
24   "labels": label,
25   "datasets": [
26     {
27       "data": data,
28     }
29   ]
30 }
31 ]
32 }
```

The 'Chart Script' is:

```
11 "type": "pie",
12 "data": None,
13 "options": {
14   "responsive": True,
15   "title": {
16     "display": True,
17     "text": "Table Size (Total: " + str(total_size) + ")"
18   }
19 }
```

A dialog box in the center says 'Monitor unit saved.' with an 'Ok' button. Below the scripts, a pie chart titled 'Table Size (Total: 0.17)' is shown, with a legend for categories: categories, reorder, cust_hist, customers, orders, orderlines, inventory, and products.

The screenshot shows the 'Monitoring' tab in OmniDB with four monitoring units:

- Table Size**: Refresh interval 5 seconds. Shows a pie chart titled 'Table Size (Total: 0.17)' with a legend for categories, reorder, cust_hist, customers, orders, orderlines, inventory, and products.
- Random Number**: Refresh interval 3 seconds, 1 row. Shows a table with one row: 1, 0.57640914991498.
- Memory Usage**: Refresh interval 10 seconds. Shows a line chart titled 'System Memory Usage (Total: 7675MB)' with a legend for Memory. The y-axis is labeled '%' and ranges from 0 to 100. The x-axis shows time from 12:49 to 1:23.
- CPU Usage**: Refresh interval 10 seconds. Shows a line chart titled 'CPU Usage' with a legend for all, 0, 1, 2, 3. The y-axis is labeled 'Value' and ranges from 0 to 100. The x-axis shows time from 12:49 to 1:23.

14.5 Writing custom Monitoring Units: Chart-Append

Now for the last, but most interesting kind of Monitoring Unit: *Chart-Append*. It is interesting because there is a wide range of applications for these units, since they keep recent historic data that allows us to see a comparison of values.

Go ahead and add a new chart using (*Chart (Append)*) Size: *Top 5 Tables* as template:

Name: Type: Refresh Interval: seconds

Template:

Data Script:

```

1 from datetime import datetime
2 from random import randint
3
4 tables = connection.Query('''
5     SELECT nsname || '.' || relname AS relation,
6           round(pg_relation_size(c.oid)/1048576.0,2) AS size
7     FROM pg_class c
8     LEFT JOIN pg_namespace n ON (n.oid = c.relnamespace)
9     WHERE nsname NOT IN ('pg_catalog', 'information_schema')
10    AND c.relkind <> 'i'
11 ''')

```

Chart Script:

```

1 result = {
2     "type": "line",
3     "data": None,
4     "options": {
5         "responsive": True,
6         "title": {
7             "display": True,
8             "text": "Size: Top 5 Tables"
9         },
10    "tooltips": {
11        "mode": "index",

```

Now take a look at the source code of both *Data Script* and *Chart Script*. It is not too different from the Chart units. The dataset creation is a bit more complex as it involves other JSON settings, but that's all.

As an exercise, based on this chart, create another one called *Size: Top 20 Tables*. It should look like this:

Name: Type: Refresh Interval: seconds

Template:

Data Script:

```

1 from datetime import datetime
2 from random import randint
3
4 tables = connection.Query('''
5     SELECT nsname || '.' || relname AS relation,
6           round(pg_relation_size(c.oid)/1048576.0,2) AS size
7     FROM pg_class c
8     LEFT JOIN pg_namespace n ON (n.oid = c.relnamespace)
9     WHERE nsname NOT IN ('pg_catalog', 'information_schema')
10    AND c.relkind <> 'i'
11 ''')

```

Chart Script:

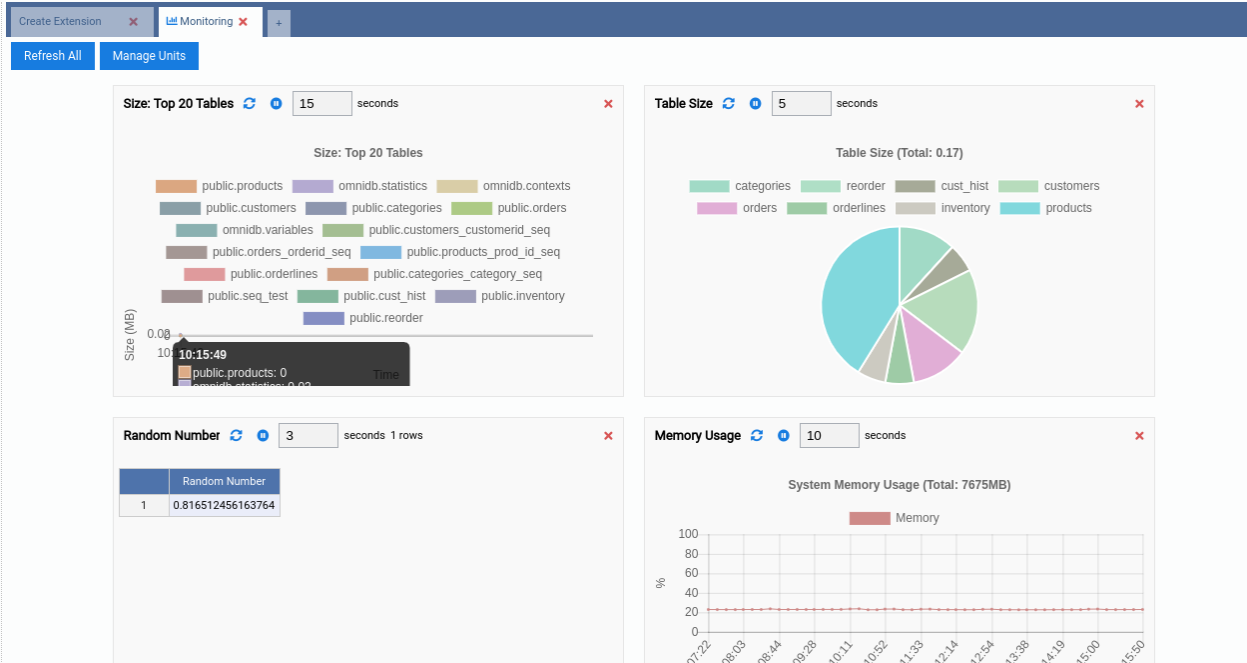
```

1 result = {
2     "type": "line",
3     "data": None,
4     "options": {
5         "responsive": True,
6         "title": {
7             "display": True,
8             "text": "Size: Top 20 Tables"
9         },
10    "tooltips": {
11        "mode": "index",

```

Preview:

Now save it and add it to your dashboard:



15. Logical Replication

PostgreSQL 10 introduces native logical replication, which uses a publish/subscribe model and so we can create publications on the upstream (or publisher) and subscriptions on downstream (or subscriber). For more details about it, please refer to the [PostgreSQL documentation](#).

In this chapter, we will use a 2-node cluster to demonstrate PostgreSQL 10 native logical replication. Note that on each PostgreSQL instance, you need to configure `wal_level = logical` and also make sure to adjust file `pg_hba.conf` to grant access to replication between the 2 nodes.

15.1 Creating a test environment

OmniDB repository provides a 2-node Vagrant test environment. If you want to use it, please do the following:

```
git clone --depth 1 https://github.com/OmniDB/OmniDB
cd OmniDB/OmniDB_app/tests/vagrant/postgresql-10-2nodes/
vagrant up
```

It will take a while, but once finished, 2 virtual machines with IP addresses `10.33.2.114` and `10.33.2.115` will be up and each of them will have PostgreSQL 10 listening to port `5432`, with all settings needed to configure native logical replication. A new database called `omnidb_tests` is also created on both machines. To connect, user is `omnidb` and password is `omnidb`.

15.2 Connecting to both nodes

Let's use OmniDB to connect to both PostgreSQL nodes. First of all, fill out connection info in the connection grid:

New Connection

Technology	Server	Port	Database	User	Title	SSH Tunnel	SSH Server	SSH Port	SSH User	SSH Password	SSH Key	Actions
postgresql ▾	10.33.2.114	5432	omnidb_tests	omnidb	Node 1	<input type="checkbox"/>		22				
postgresql	10.33.2.115	5432	omnidb_tests	omnidb	Node 2	<input type="checkbox"/>		22				

Then select both connections. Note how OmniDB understands it is connected to PostgreSQL 10 and enables a new node in the current connection tree view: it is called *Logical Replication*. Inside of it, we can see *Publications* and *Subscriptions*.

The screenshot displays the OmniDB interface. At the top, there's a 'Connections' bar with tabs for 'Snippets', '2.12.0', 'Node 1 - omnidb_tests', and 'Node 2 - omnidb_tests'. Below this, the 'Active database: omnidb_tests' is shown. The left sidebar contains a tree view of the database structure, including 'Databases (2)', 'postgres', 'omnidb_tests', 'Schemas', 'Extensions', 'Foreign Data Wrappers', 'Logical Replication', 'Publications', 'Subscriptions', 'Tablespaces', 'Roles', and 'Replication Slots'. The main area shows the 'Console' tab with a query editor and a 'Query' tab. The bottom status bar shows 'Autocommit' checked and 'Not connected'.

15.3 Creating a test table on both nodes

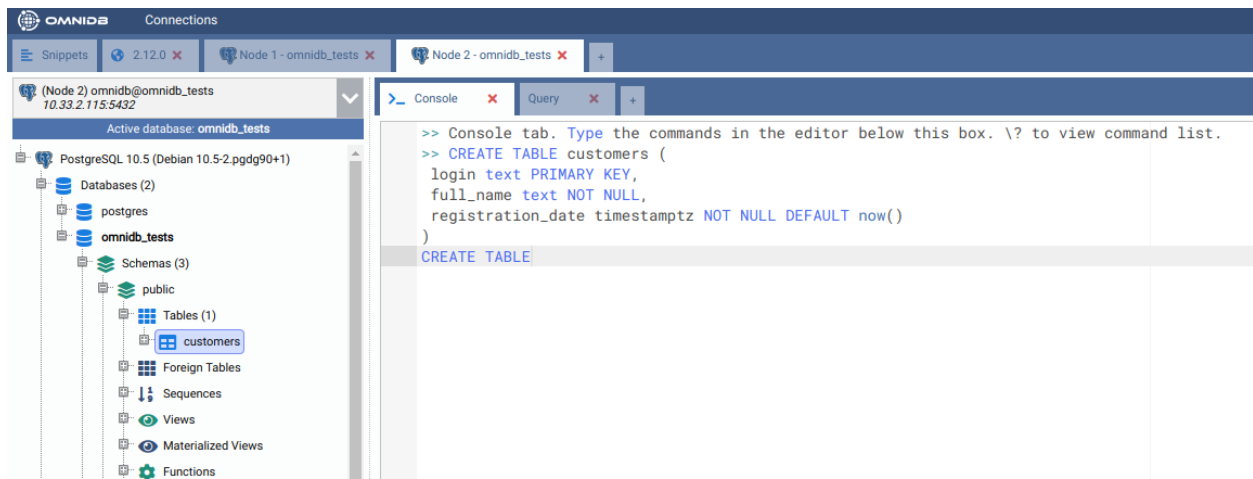
On both nodes, create a table like this:

```
CREATE TABLE customers (
  login text PRIMARY KEY,
  full_name text NOT NULL,
```

(continues on next page)

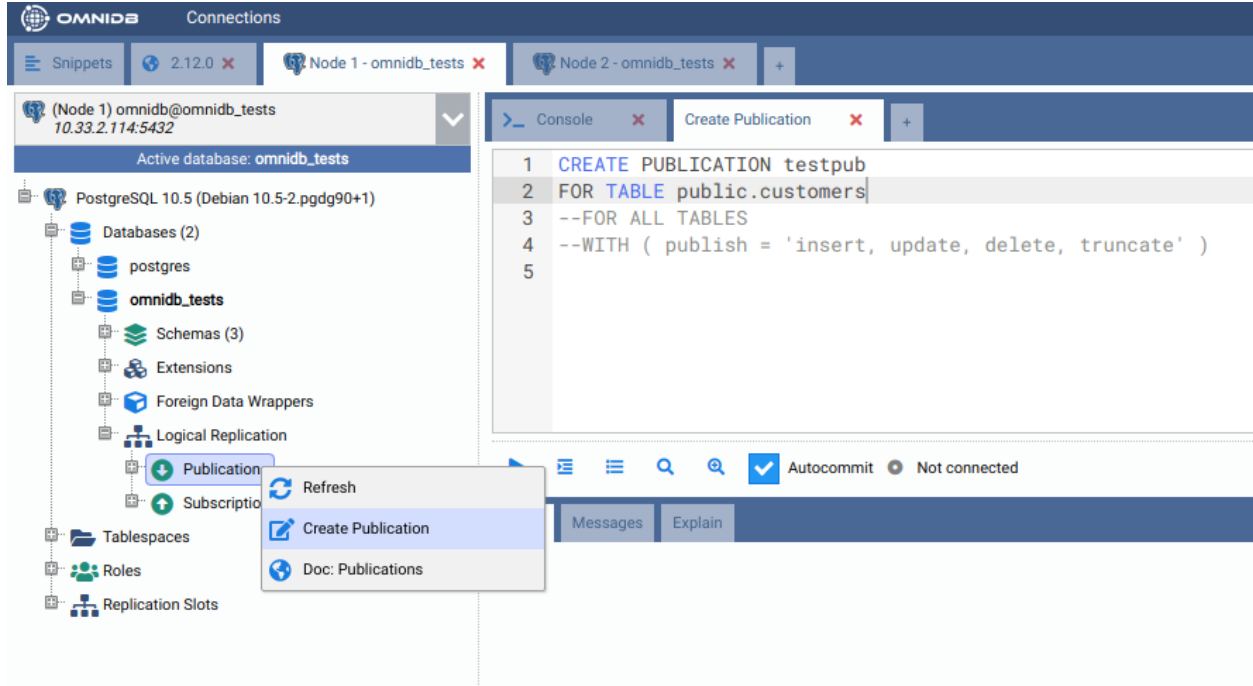
(continued from previous page)

```
registration_date timestampz NOT NULL DEFAULT now()
)
```

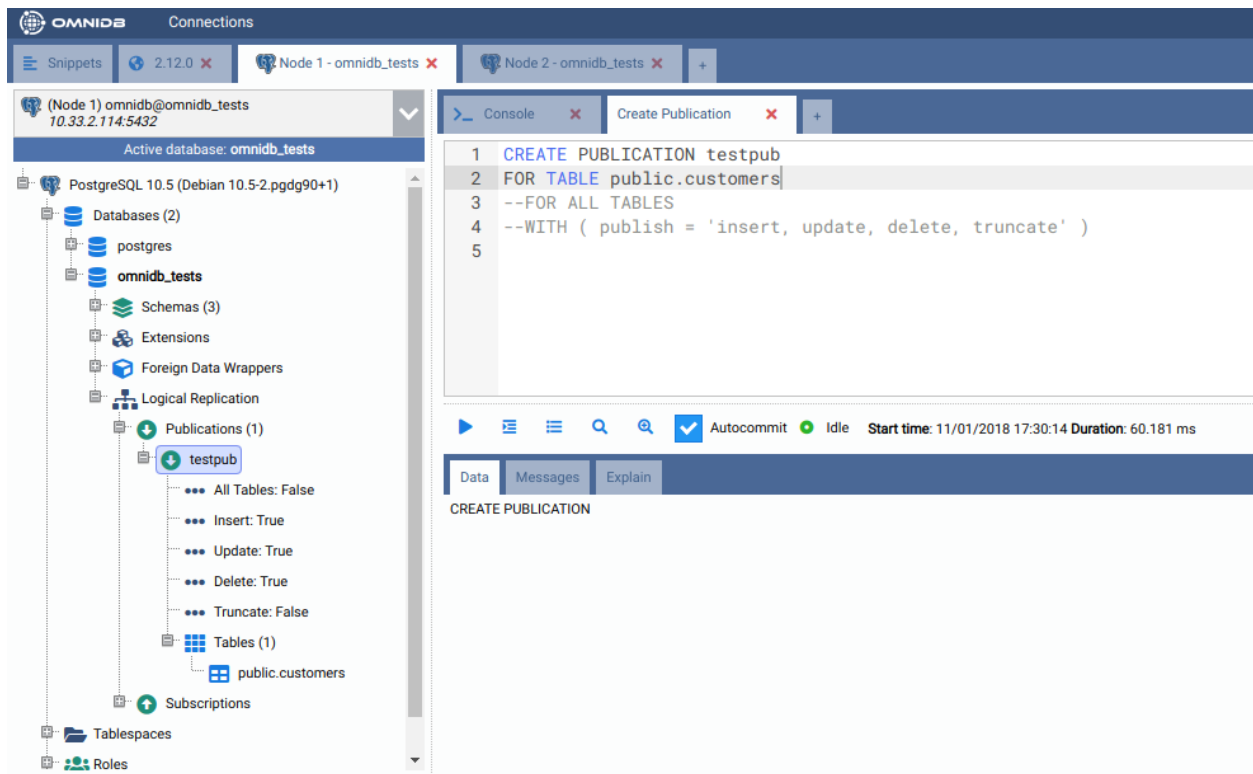


15.4 Create a publication on the first machine

Inside the connection node, expand the *Logical Replication* node, then right click in the *Publications* node, and choose the action *Create Publication*. OmniDB will open a SQL template tab with the `CREATE PUBLICATION` command ready for you to make some adjustments and run:

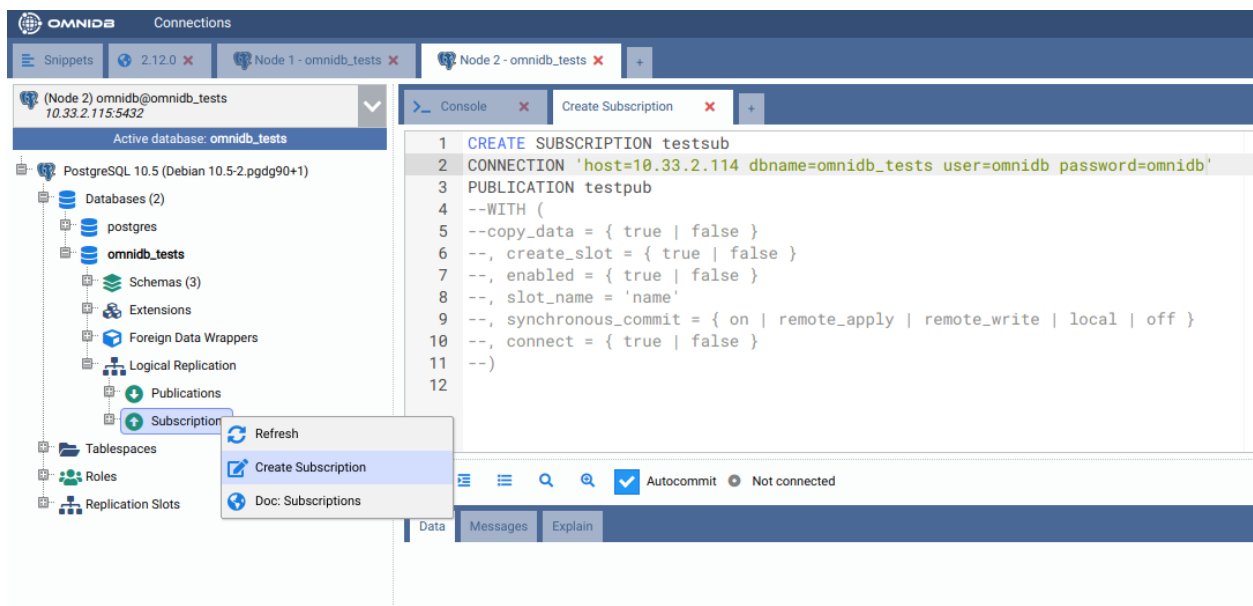


After adjusting and executing the command, you can right click the *Publications* node again and click on the *Refresh* action. You will see that will be created a new node with the same name you gave to the publication. Expanding this node, you will see the details and the tables for the publication:



15.5 Create a subscription on the second machine

Inside the connection node, expand the *Logical Replication* node, then right click in the *Subscriptions* node, and choose the action *Create Subscription*. OmniDB will open a SQL template tab with the `CREATE SUBSCRIPTION` command ready for you to make some adjustments and run:



After adjusting and executing the command, you can right click the *Subscriptions* node again and click on the *Refresh* action. You will see that will be created a new node with the same name you gave to the subscription. Expanding this

node, you will see the details, the referenced publications and the tables for the subscription:

The screenshot shows the OmniDB interface with the 'CREATE SUBSCRIPTION' command executed in the console. The command is as follows:

```

1 CREATE SUBSCRIPTION testsub
2 CONNECTION 'host=10.33.2.114 dbname=omnidb_tests user=omnidb password=omnidb'
3 PUBLICATION testpub
4 --WITH (
5 --copy_data = { true | false }
6 --, create_slot = { true | false }
7 --, enabled = { true | false }
8 --, slot_name = 'name'
9 --, synchronous_commit = { on | remote_apply | remote_write | local | off }
10 --, connect = { true | false }
11 --)
12

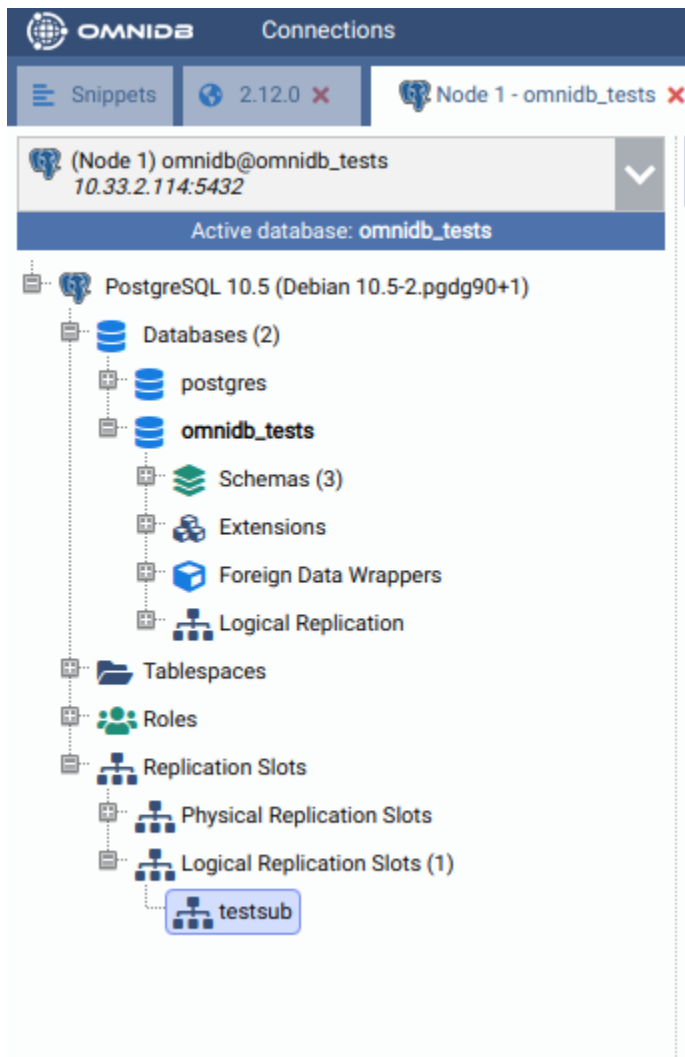
```

The interface also shows the database structure on the left pane, highlighting the 'Subscriptions (1)' section under 'Logical Replication'. The subscription 'testsub' is listed with the following details:

- Enabled: True
- ConnInfo: host=10.33.2.114 dbname=om
- Referenced Publications: testpub
- Tables (1): public.customers

The bottom pane shows the 'Data' tab with the command 'CREATE SUBSCRIPTION'.

Also, the `CREATE SUBSCRIPTION` command created a *logical replication slot* called `testsub` (the same name as the subscription) in the first machine:



15.6 Testing the logical replication

To test the replication is working, let's create some data on the node 1. Right click on the table `public.customers`, then point to *Data Actions*, then click on the action *Edit Data*. In this grid, you are able to add, edit and remove data from the table. Add 2 sample rows, like this:

The screenshot shows the OmniDB interface with a PostgreSQL 10.5 connection. The left sidebar displays the database structure, including the 'public' schema and the 'customers' table. A context menu is open over the 'customers' table, showing options like 'Refresh', 'Data Actions', and 'Table Actions'. The 'Data Actions' menu is expanded, showing options like 'Query Data', 'Edit Data', 'Insert Record', 'Update Records', 'Count Records', 'Delete Records', and 'Truncate Table'. The main console area shows a SQL query: `select * from public.customers t order by t.login`. Below the query, the results are displayed in a table with 3 rows and 4 columns: `login` (text), `full_name` (text), and `registration_date` (timestamp). The 'Query 10 rows' button is visible, and the save time is 0.057 seconds.

Active database: omnidb_tests

PostgreSQL 10.5 (Debian 10.5-2.pgdg90+1)

Databases (2)

- postgres
- omnidb_tests
 - Schemas (3)
 - public
 - Tables (1)
 - customers
 - Foreign Tables
 - Sequences
 - Views
 - Materialized Views
 - Functions
 - Trigger Functions
 - pg_catalog
 - information_schema
 - Extensions
 - Foreign Data Wrappers
 - Logical Replication
 - Tablespaces

Properties DDL

Property	Value
Database	omnidb_tests
Schema	public
Table	customers

select * from public.customers t
1 order by t.login

Query 10 rows Save time: 0.057 seconds

		login (text)	full_name (text)	registration_date (timestamp)
1	✗	william	William Ivanski	2018-11-01
2	✗	rafael	Rafael Thofehn Castro	2018-11-01
3	+			

Then, on the other node, check if the table `public.customers` was automatically populated. Right click on the table `public.customers`, then point to *Data Actions*, then click on the action *Query Data*:

The screenshot shows the OmniDB interface with a PostgreSQL 10.5 connection. The left sidebar displays the database structure, including the 'customers' table in the 'public' schema. A context menu is open over the 'customers' table, showing options like 'Query Data', 'Edit Data', 'Insert Record', 'Update Records', 'Count Records', 'Delete Records', and 'Truncate Table'. The main window displays a SQL query:

```
1 SELECT t.login
2       , t.full_name
3       , t.registration_date
4 FROM public.customers t
5 ORDER BY t.login
```

Below the query, the execution status is shown: 'Autocommit' is checked, 'Idle' status, 'Number of records: 2', 'Start time: 11/01/2018 17:39:11', and 'Duration: 40.221 ms'. The 'Data' tab is selected, showing the following table:

	login	full_name	registration_date
1	rafael	Rafael Thofehrn Castro	2018-11-01 00:00:00+00:00
2	william	William Ivanski	2018-11-01 00:00:00+00:00

As we can see, both rows created in the first machine were replicated into the second machine. This tells us that the logical replication is working.

Now you can perform other actions, such as adding/removing tables to the publication and creating a new publication that publishes all tables.

CHAPTER 16

16. pglogical

pglogical is a PostgreSQL extension that provides an advanced logical replication system that serves as a highly efficient method of replicating data as an alternative to physical replication.

In this chapter, we will use a 2-node cluster to demonstrate pglogical with PostgreSQL 10. Note that on each PostgreSQL instance, you need to configure:

```
wal_level = 'logical'
track_commit_timestamp = on
max_worker_processes = 10    # one per database needed on provider node
                               # one per node needed on subscriber node
max_replication_slots = 10   # one per node needed on provider node
max_wal_senders = 10         # one per node needed on provider node
shared_preload_libraries = 'pglogical'
```

Also make sure to adjust file `pg_hba.conf` to grant access to replication between the 2 nodes.

16.1 Creating a test environment

OmniDB repository provides a 2-node Vagrant test environment. If you want to use it, please do the following:

```
git clone --depth 1 https://github.com/OmniDB/OmniDB
cd OmniDB/OmniDB_app/tests/vagrant/pglogical-2-postgresql-10-2nodes/
vagrant up
```

It will take a while, but once finished, 2 virtual machines with IP addresses `10.33.3.114` and `10.33.3.115` will be up and each of them will have PostgreSQL 10 listening to port `5432`, with all settings needed to configure pglogical replication. A new database called `omnidb_tests` is also created on both machines. To connect, user is `omnidb` and password is `omnidb`.

16.2 Install OmniDB pglogical plugin

OmniDB core does not support pglogical by default. You will need to download and install pglogical plugin. If you are using OmniDB server, these are the steps:

```
wget https://omnidb.org/dist/plugins/omnidb-pglogical_1.0.0.zip
unzip omnidb-pglogical_1.0.0.zip
sudo cp -r plugins/ static/ /opt/omnidb-server/OmnidB_app/
sudo systemctl restart omnidb
```

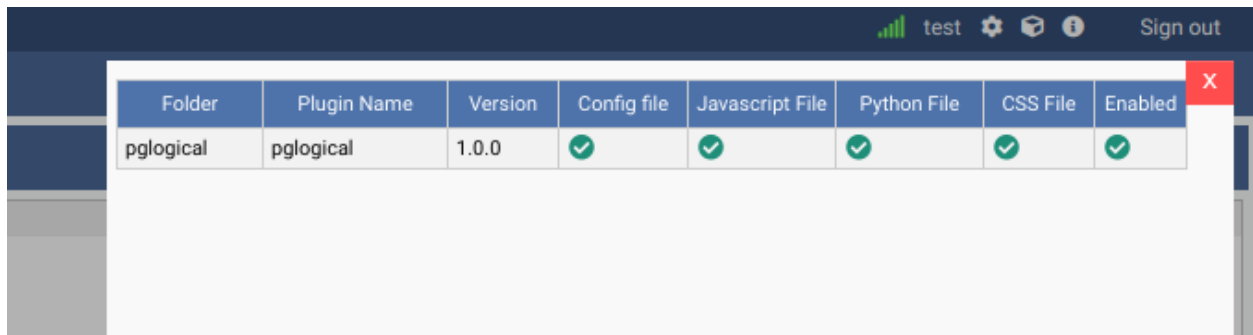
And then refresh the OmniDB web page in the browser.

For OmniDB app, these are the steps:

```
wget https://omnidb.org/dist/plugins/omnidb-pglogical_1.0.0.zip
unzip omnidb-pglogical_1.0.0.zip
sudo cp -r plugins/ static/ /opt/omnidb-app/resources/app/omnidb-server/OmnidB_app/
```

And then restart OmniDB app.

If everything worked correctly, by clicking on the “plugins” icon in the top right corner, you will see the plugin installed and enabled:

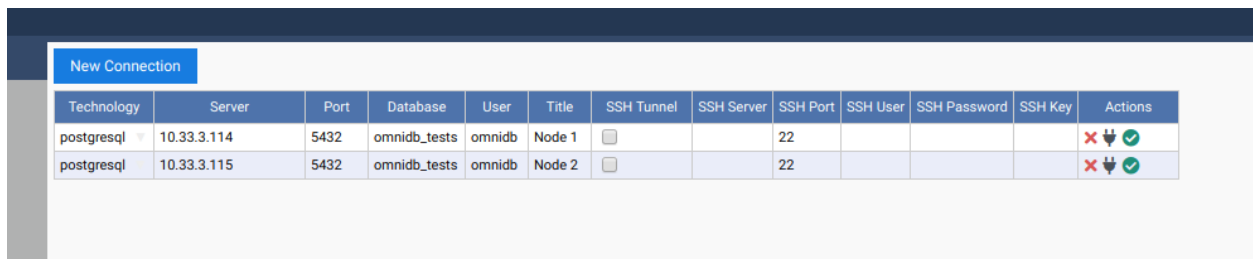


The screenshot shows the OmniDB web interface with a top navigation bar containing a signal icon, the word 'test', a settings gear icon, an information icon, and a 'Sign out' button. Below the navigation bar is a table listing installed plugins. The table has columns: Folder, Plugin Name, Version, Config file, Javascript File, Python File, CSS File, and Enabled. A single row is shown for the 'pglogical' plugin, with all status icons (green checkmarks) indicating it is installed and enabled. A red 'X' icon is visible in the top right corner of the table area.

Folder	Plugin Name	Version	Config file	Javascript File	Python File	CSS File	Enabled
pglogical	pglogical	1.0.0	✓	✓	✓	✓	✓

16.3 Connecting to both nodes

Let's use OmniDB to connect to both PostgreSQL nodes. First of all, fill out connection info in the connection grid:



The screenshot shows the 'New Connection' dialog in the OmniDB web interface. It contains a table with columns: Technology, Server, Port, Database, User, Title, SSH Tunnel, SSH Server, SSH Port, SSH User, SSH Password, SSH Key, and Actions. Two connections are listed, both for 'postgresql' technology, with servers '10.33.3.114' and '10.33.3.115', port '5432', database 'omnidb_tests', and user 'omnidb'. The titles are 'Node 1' and 'Node 2'. The 'SSH Tunnel' column has checkboxes that are currently unchecked. The 'Actions' column contains icons for delete (red X), refresh (circular arrow), and connect (green checkmark).

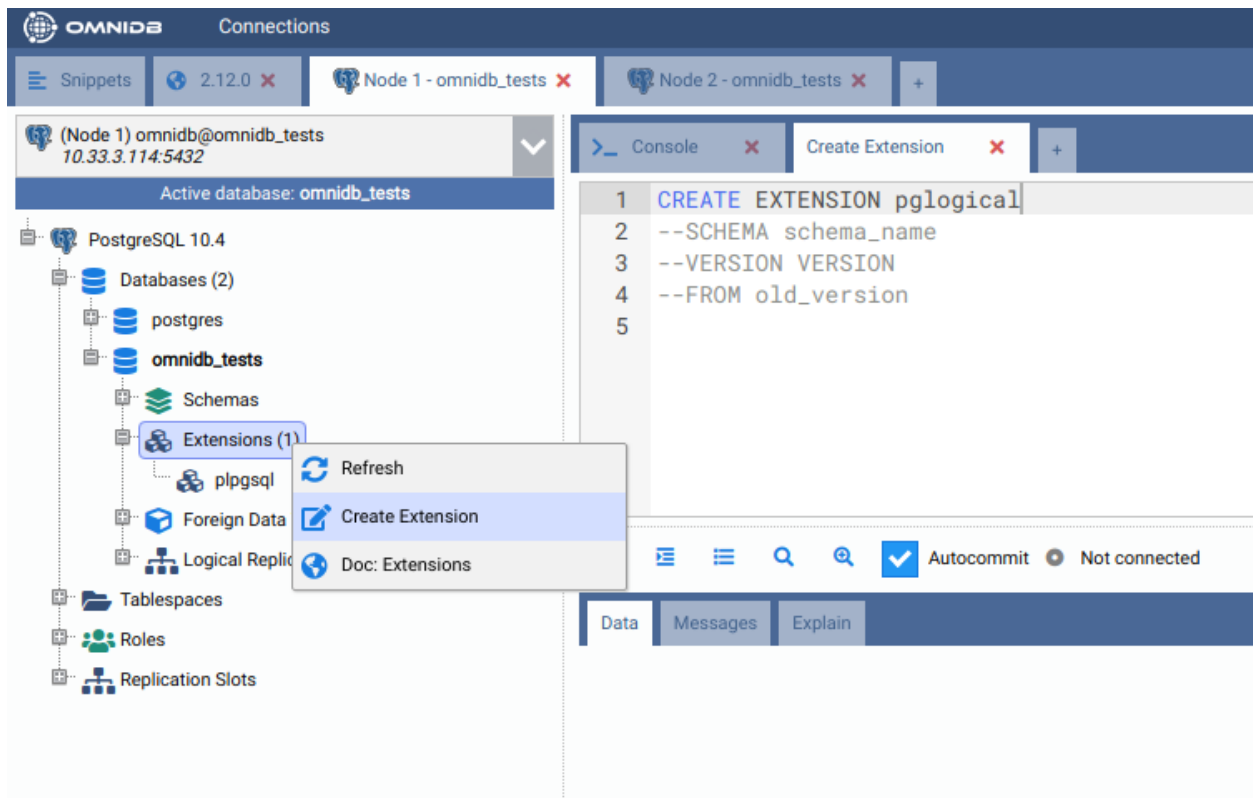
Technology	Server	Port	Database	User	Title	SSH Tunnel	SSH Server	SSH Port	SSH User	SSH Password	SSH Key	Actions
postgresql	10.33.3.114	5432	omnidb_tests	omnidb	Node 1	<input type="checkbox"/>		22				✗ ↻ ✓
postgresql	10.33.3.115	5432	omnidb_tests	omnidb	Node 2	<input type="checkbox"/>		22				✗ ↻ ✓

Then select both connections.

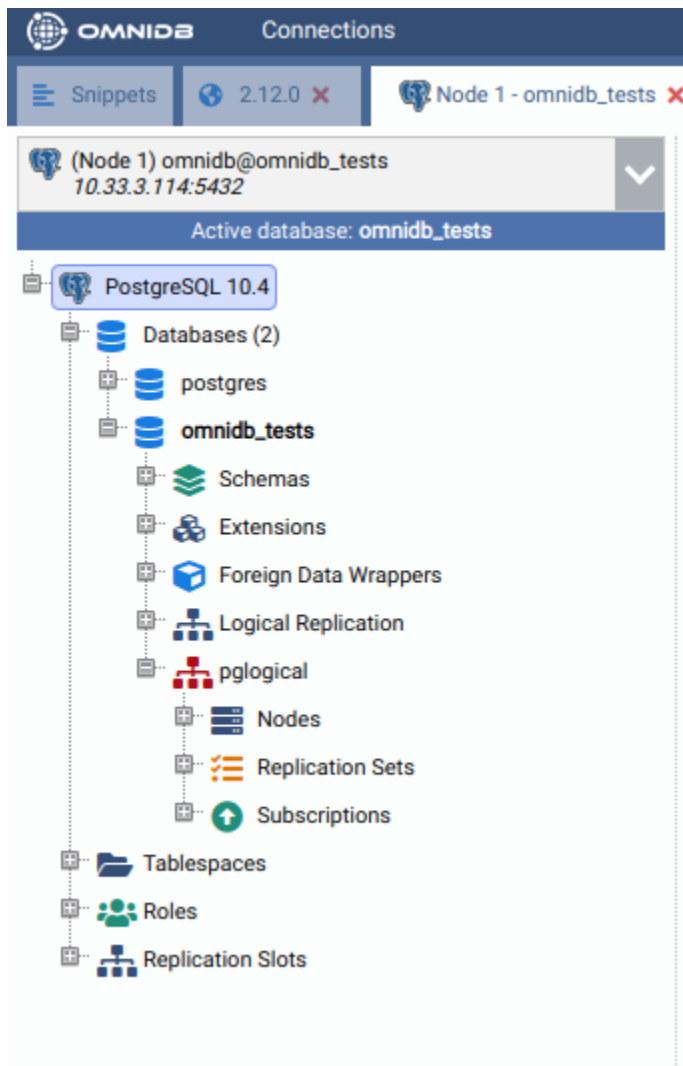
16.4 Create pglogical extension in both nodes

pglogical requires an extension to be installed in both nodes. Inside OmniDB, you can create the extension by right clicking on the *Extensions* node, and choosing the action *Create Extension*. OmniDB will open a SQL template tab

with the `CREATE EXTENSION` command ready for you to make some adjustments and run:

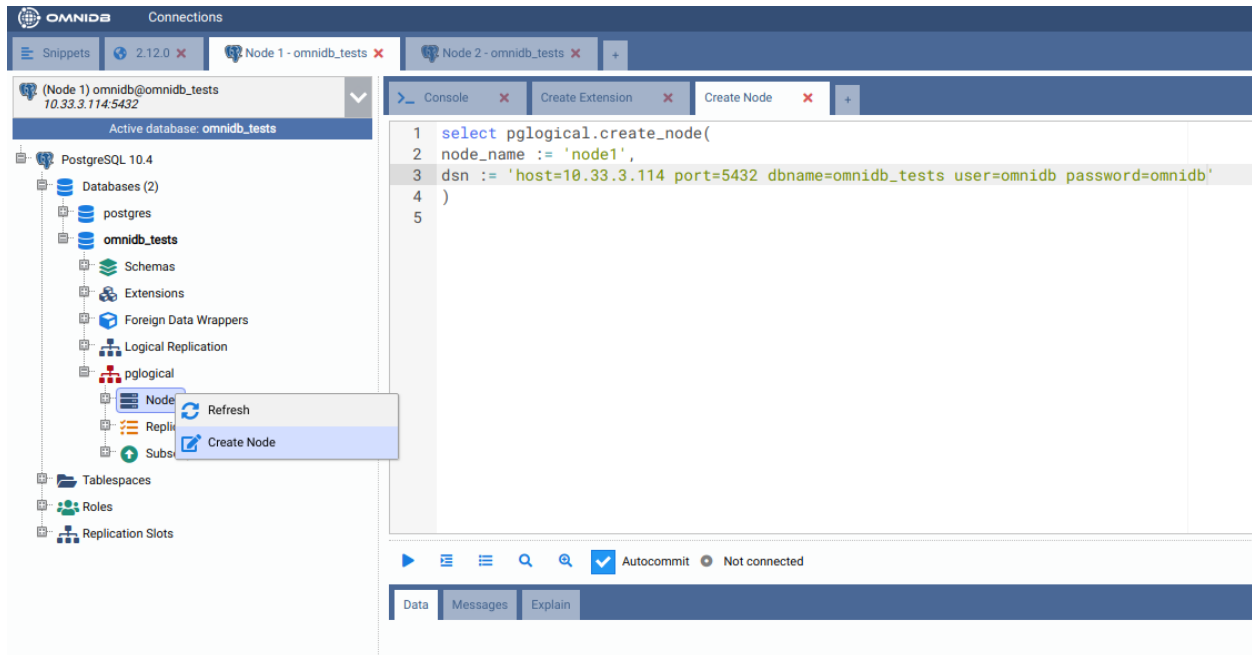


After you have created the extension, you need to refresh the root node of the treeview, by right-clicking on it and choosing *Refresh*. Then you will see that OmniDB already acknowledges the existence of `pglogical` in this database. However, `pglogical` is not active yet.



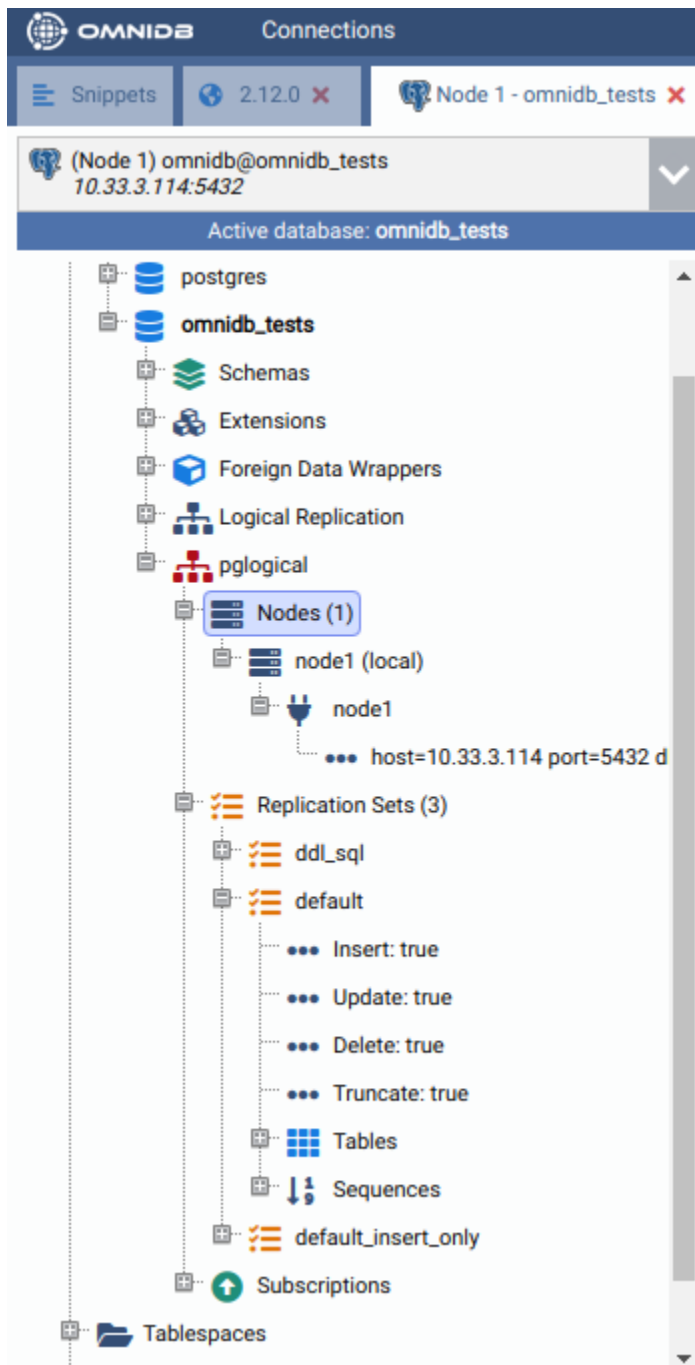
16.5 Create pglogical nodes

To activate pglogical in this database, we need to create a pglogical node on each machine. Inside the *pglogical* node of the treeview, right click *Nodes*, then choose *Create Node*. In the SQL template that will open, adjust the node name and the DSN and run the command.



Then right click *Nodes* again, but this time choose *Refresh*. You will see the node you just created. Note how OmniDB understands that this node is local. Expand the local node to see its interface inside. You can manage the interfaces of the nodes using OmniDB too.

Go ahead and expand the *Replication Sets* node. You can see pglogical default replication sets are already created: *ddl_sql*, *default* and *default_insert_only*. You can also manage replication sets using OmniDB.



Now create a node on the other machine too. Choose a different name for the node.

16.6 Create a table on the first machine

In the first machine, under the *Schemas* node, expand the *public* node, then right-click the *Tables* node and choose *Create Table*. In the form tab that will open, give the new table a name and some columns. Also add a primary key in the *Constraints* tab. When done, click in the *Save Changes* button.

The screenshot shows the OmniDB interface with two panels. The top panel displays the 'New Table' dialog for a table named 'test_table'. The 'Columns' tab is active, showing a table with 3 columns: 'id' (integer, NOT NULL), 'name' (text, NOT NULL), and an empty row. The bottom panel displays the 'New Constraint' dialog for a primary key constraint named 'pk_test' on the 'id' column.

Table Name: test_table **Save Changes**

	Column Name	Data Type	Nullable	
1	id	integer	NO	✗
2	name	text	NO	✗
3				

New Constraint

Constraint Name	Type	Columns	Referenced Table	Referenced Columns	Delete Rule	Update Rule	
pk_test	Primary Key	id					✗

16.7 Add the new table to a replication set on the first machine

In the first machine, under the *default_insert_only* replication set, right click the *Tables* node and choose *Add Table*. In the SQL template tab that will open, change the table name in the *relation* argument and then execute the command.

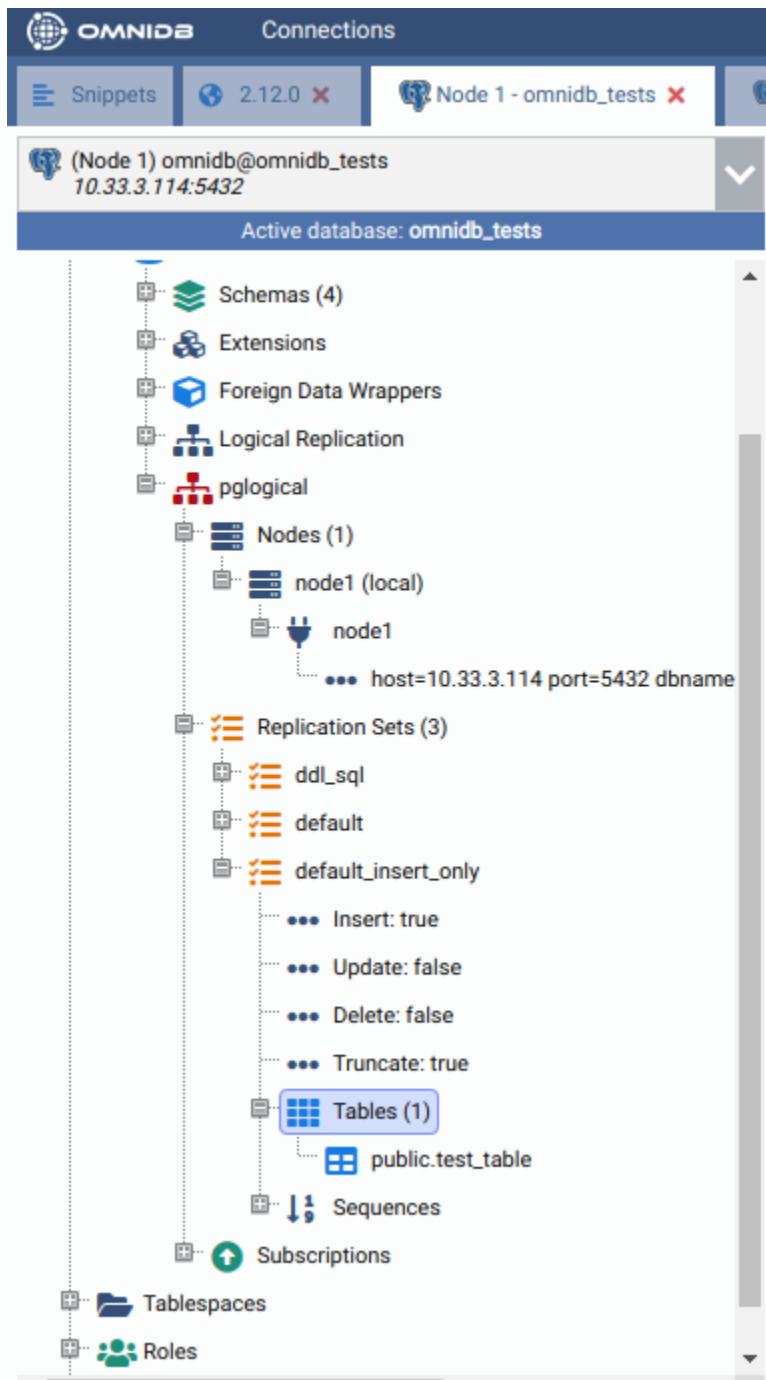
The screenshot shows the OmniDB interface with the following components:

- Top Bar:** Connections tab, showing Node 1 - omnidb_tests and Node 2 - omnidb_tests.
- Left Sidebar:** Active database: omnidb_tests. The tree structure includes:
 - Schemas (4)
 - Extensions
 - Foreign Data Wrappers
 - Logical Replication
 - pglogical
 - Nodes (1)
 - node1 (local)
 - node1
 - host=10.33.3.114 port=5432 dbname
 - Replication Sets (3)
 - ddl_sql
 - default
 - default_insert_only
 - Insert: true
 - Update: false
 - Delete: false
 - Truncate: true
 - Tables
 - Sequences
 - Subscriptions
 - Tablespaces
 - Roles
 - Destination State

- Right Pane:** SQL console showing the command:

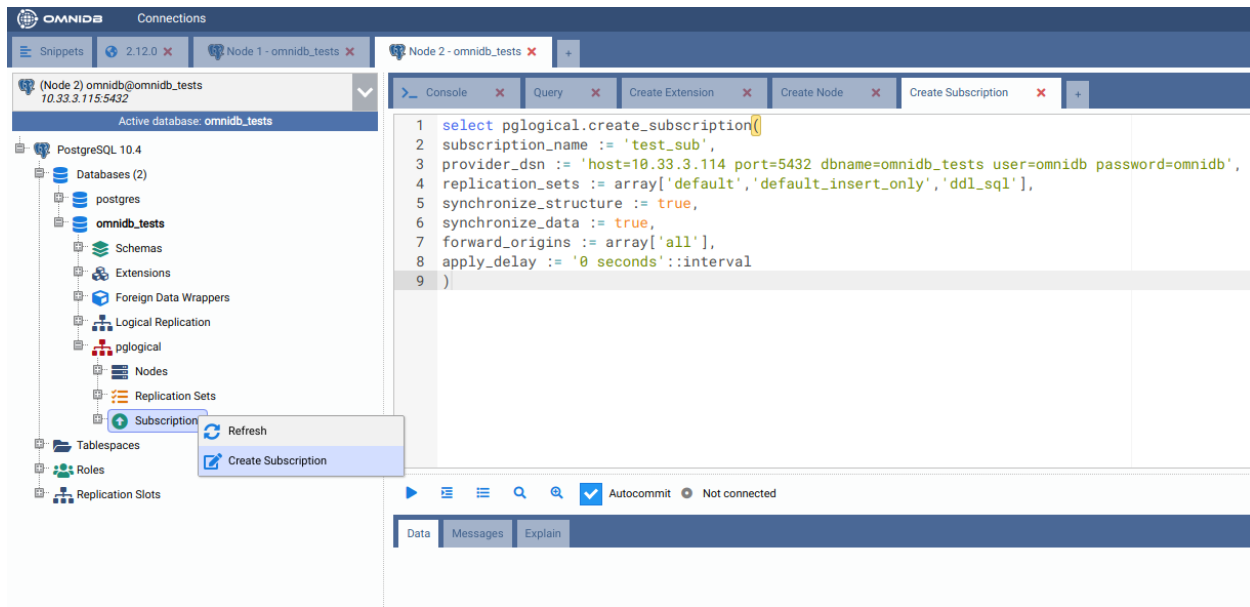
```
1 select pglogical.replication_set_add_table(  
2   set_name := 'default_insert_only',  
3   relation := 'public.test_table'::regclass,  
4   synchronize_data := true,  
5   columns := null,  
6   row_filter := null  
7 )
```
- Bottom Bar:** Data, Messages, Explain tabs. Status: Autocommit (checked), Not connected.

Refresh the *Tables* node to check the table was added to the replication set.



16.8 Add a subscription on the second machine

In the second machine, right-click the *Subscriptions* node and choose *Create Subscription*. In the SQL template tab that will open, change DSN of the first machine and then execute the command.



Refresh and expand both *Nodes* and *Subscriptions* nodes of the treeview. Note how now the second machine knows about the first machine. Also check the information OmniDB shows about the subscription we just created.

The screenshot displays the OmniDB interface with the following components:

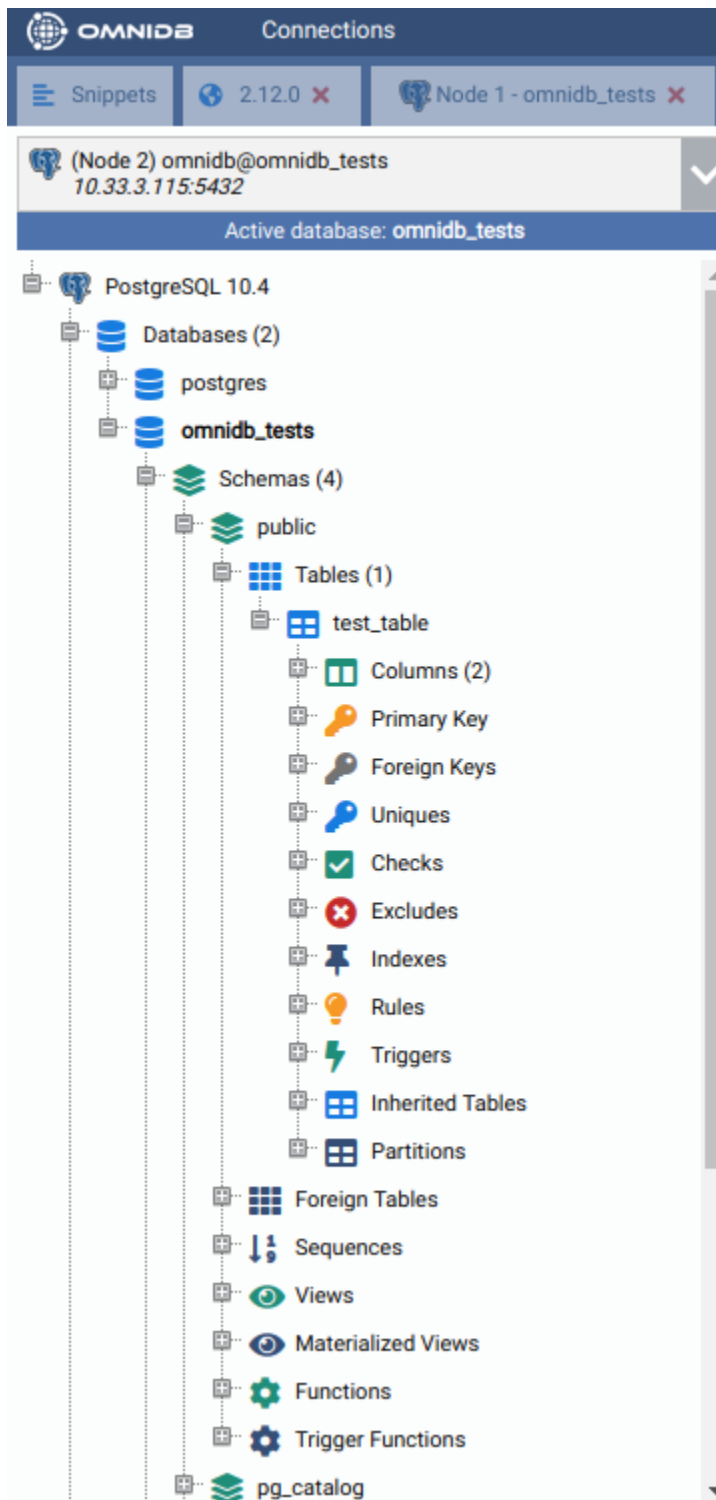
- Top Bar:** Shows the OmniDB logo, the title "Connections", and tabs for "Snippets", "2.12.0", "Node 1 - omnidb_tests", and "Node 2 - omnidb_tests".
- Left Pane:** A tree view of the database structure. The "pglogical" extension is expanded, showing "Nodes (2)". Under "Nodes (2)", there are two nodes: "node1" and "node2 (local)". Each node has a corresponding "node1" and "node2" entry with details like "host=10.33.3.114 port=5432 dbna". Below the nodes, there are "Replication Sets" and "Subscriptions (1)". The "test_sub" subscription is highlighted, showing details: "Status: replicating", "Provider: node1", "Enabled: true", and "Apply Delay: 00:00:00". Under "Replication Sets (3)", there are "default", "default_insert_only", and "ddl_sql".
- Right Pane:** A SQL query in the Console and its results in the Data tab.


```

1 select pglogical.create_subscription(
2   subscription_name :=
3   provider_dsn := 'host=10.33.3.114 port=5432 dbna'
4   replication_sets :=
5   synchronize_structure := true
6   synchronize_data := true
7   forward_origins := all
8   apply_delay := '0 seconds'
9 )
      
```

create_subscription	
1	2769129391

Also verify that the table *public.test_table* was created automatically in the second machine:



16.9 Add some data in the table on the first machine

In the first machine, under the *Schemas* node, expand the *public* node and the *Tables* node. Right-click in our table, *test_table*, move the mouse pointer to *Data Actions* and then click on *Edit Data*. Insert some data to the table. When

finished, click on the *Save Changes* button.

The screenshot shows the OmniDB interface with the following components:

- Connections:** Node 1 - omnidb_tests (2.12.0), Node 2 - omnidb_tests.
- Active database:** omnidb_tests.
- Database Structure:** PostgreSQL 10.4, Databases (2), postgres, omnidb_tests, Schemas (4), public, Tables (1), test_table.
- Context Menu:** Refresh, Data Actions, Table Actions.
- Data Actions Sub-menu:** Query Data, Edit Data, Insert Record, Update Records, Count Records, Delete Records, Truncate Table.
- Console:** select * from public.test_table t, 1 order by t.id.
- Table Data:**

		id (integer)	name (text)
1	✗	1	john
2	✗	2	paul
3	✗	3	george
4	✗	4	stuart
5	✗	5	pete
6	✗	6	yoko
7	+		
- Properties/DDL:**

Property	Value
Database	omnidb_tests
Schema	public
- Query Results:** Query 10 rows, Number of records: 0, Response time: 0.033 seconds, Save Changes button.

Now let us check the data was replicated. Go to the second machine and right-click the table, move the mouse pointer to *Data Actions* and then click on *Query Data*.

Active database: omnidb_tests

PostgreSQL 10.4

Databases (2)

- postgres
- omnidb_tests
 - Schemas (4)
 - public
 - Tables (1)
 - test_table
 - Foreign Tables
 - Sequences
 - Views
 - Materialized Views
 - Functions
 - Trigger Functions
 - pg_catalog
 - information_schema
 - pglogical
 - Extensions
 - Foreign Data Wrappers
 - Logical Replication
 - pglogical
- Tablespaces
- Roles
- Replication Slots

Refresh

Data Actions

Table Actions

Query Data

Edit Data

Insert Record

Update Records

Count Records

Delete Records

Truncate Table

```
1 SELECT t.id
2    , t.name
3 FROM public.test_table t
4 ORDER BY t.id
```

Autocommit Idle Number of records: 6
Start time: 11/08/2018 15:34:50 Duration: 46.643 ms

	id	name
1	1	john
2	2	paul
3	3	george
4	4	stuart
5	5	pete
6	6	yoko

Property	Value
Database	omnidb_tests
Schema	public

16.10 Check if delete is being replicated

In the *Edit Data* tab in the first machine, remove Pete and Stuart. Click on the button *Save Changes* when done.

The screenshot shows the OmniDB web interface. At the top, there is a tab bar with 'Console', 'public.test_table', 'Add Table', and another 'public.test_table' tab. The main area displays a SQL query: `select * from public.test_table t order by t.id`. Below the query, there is a button 'Query 10 rows' and a status bar indicating 'Save time: 0.107 seconds' and a 'Save Changes' button. The results are shown in a table with columns 'id (integer)' and 'name (text)'. The table contains 7 rows, with the first 6 rows having a red 'X' in the first column and the 7th row having a '+' sign.

		id (integer)	name (text)
1	✗	1	john
2	✗	2	paul
3	✗	3	george
4	✗	4	stuart
5	✗	5	pete
6	✗	6	yoko
7	+		

Check if these 2 rows were deleted in the second machine.

The screenshot shows the OmniDB SQL console interface. At the top, there are tabs for 'Console', 'Create Subscription', and 'public.test_table'. The SQL query is as follows:

```
1 SELECT t.id
2    , t.name
3 FROM public.test_table t
4 ORDER BY t.id
```

Below the query, there is a toolbar with icons for execution, settings, and search. The status bar indicates 'Autocommit' is checked, the connection is 'Idle', and the query executed successfully. The results are displayed in a table with 6 records.

Number of records: 6
Start time: 11/08/2018 15:37:04 Duration: 4.892 ms

	id	name
1	1	john
2	2	paul
3	3	george
4	4	stuart
5	5	pete
6	6	yoko

They were not removed in the second machine because the table *public.test_table* is in the replication set *default_insert_only*, that does not replicate *updates* and *deletes*.

CHAPTER 17

17. Postgres-BDR

Postgres-BDR (or just **BDR**, for short) is an open source project from 2ndQuadrant that provides multi-master features for PostgreSQL.

In this chapter, we will use a 2-node cluster to demonstrate Postgres-BDR 9.4. Note that on each PostgreSQL instance, you need to configure:

```
wal_level = 'logical'
track_commit_timestamp = on
max_worker_processes = 10    # one per database needed on provider node
                               # one per node needed on subscriber node
max_replication_slots = 10   # one per node needed on provider node
max_wal_senders = 10         # one per node needed on provider node
shared_preload_libraries = 'bdr'
```

Also make sure to adjust file `pg_hba.conf` to grant access to replication between the 2 nodes.

17.1 Creating a test environment

OmniDB repository provides a 2-node Vagrant test environment. If you want to use it, please do the following:

```
git clone --depth 1 https://github.com/OmniDB/OmniDB
cd OmniDB/OmniDB_app/tests/vagrant/postgresql-bdr-9.4-2nodes/
vagrant up
```

It will take a while, but once finished, 2 virtual machines with IP addresses `10.33.4.114` and `10.33.4.115` will be up and each of them will have PostgreSQL 10 listening to port 5432, with all settings needed to configure BDR multi-master replication. A new database called `omnidb_tests` is also created on both machines. To connect, user is `omnidb` and password is `omnidb`.

17.2 Install OmniDB BDR plugin

OmniDB core does not support BDR by default. You will need to download and install BDR plugin. If you are using OmniDB server, these are the steps:

```
wget https://omnidb.org/dist/plugins/omnidb-bdr_1.0.0.zip
unzip omnidb-bdr_1.0.0.zip
sudo cp -r plugins/ static/ /opt/omnidb-server/OmnidB_app/
sudo systemctl restart omnidb
```

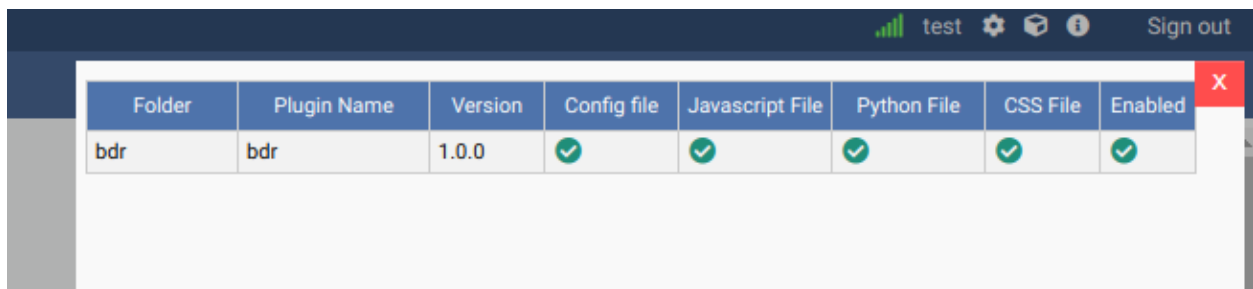
And then refresh the OmniDB web page in the browser.

For OmniDB app, these are the steps:

```
wget https://omnidb.org/dist/plugins/omnidb-bdr_1.0.0.zip
unzip omnidb-bdr_1.0.0.zip
sudo cp -r plugins/ static/ /opt/omnidb-app/resources/app/omnidb-server/OmnidB_app/
```

And then restart OmniDB app.

If everything worked correctly, by clicking on the “plugins” icon in the top right corner, you will see the plugin installed and enabled:

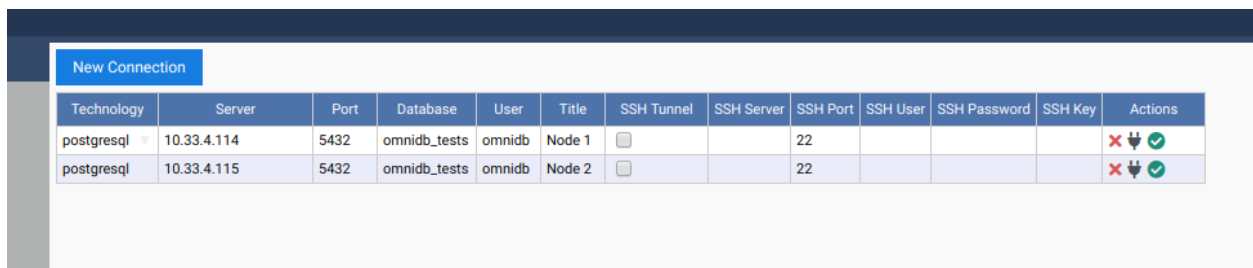


The screenshot shows the OmniDB web interface with a top navigation bar containing a signal icon, 'test', a settings gear, a mail icon, an information icon, and a 'Sign out' link. Below the navigation bar is a table listing installed plugins. The table has columns: Folder, Plugin Name, Version, Config file, Javascript File, Python File, CSS File, and Enabled. A single row is shown for the 'bdr' plugin, version 1.0.0, with all file types and the 'Enabled' status marked with green checkmarks. A red 'X' icon is visible in the top right corner of the table area.

Folder	Plugin Name	Version	Config file	Javascript File	Python File	CSS File	Enabled
bdr	bdr	1.0.0	✓	✓	✓	✓	✓

17.3 Connecting to both nodes

Let's use OmniDB to connect to both PostgreSQL nodes. First of all, fill out connection info in the connection grid:



The screenshot shows the 'New Connection' dialog box in OmniDB. It contains a table with columns: Technology, Server, Port, Database, User, Title, SSH Tunnel, SSH Server, SSH Port, SSH User, SSH Password, SSH Key, and Actions. Two connections are listed, both for PostgreSQL on server 10.33.4.114 and 10.33.4.115, port 5432, database 'omnidb_tests', user 'omnidb'. The first connection is titled 'Node 1' and the second 'Node 2'. Both have the 'SSH Tunnel' checkbox unchecked and the 'SSH Port' set to 22. The 'Actions' column for each connection shows a red 'X', a plug icon, and a green checkmark.

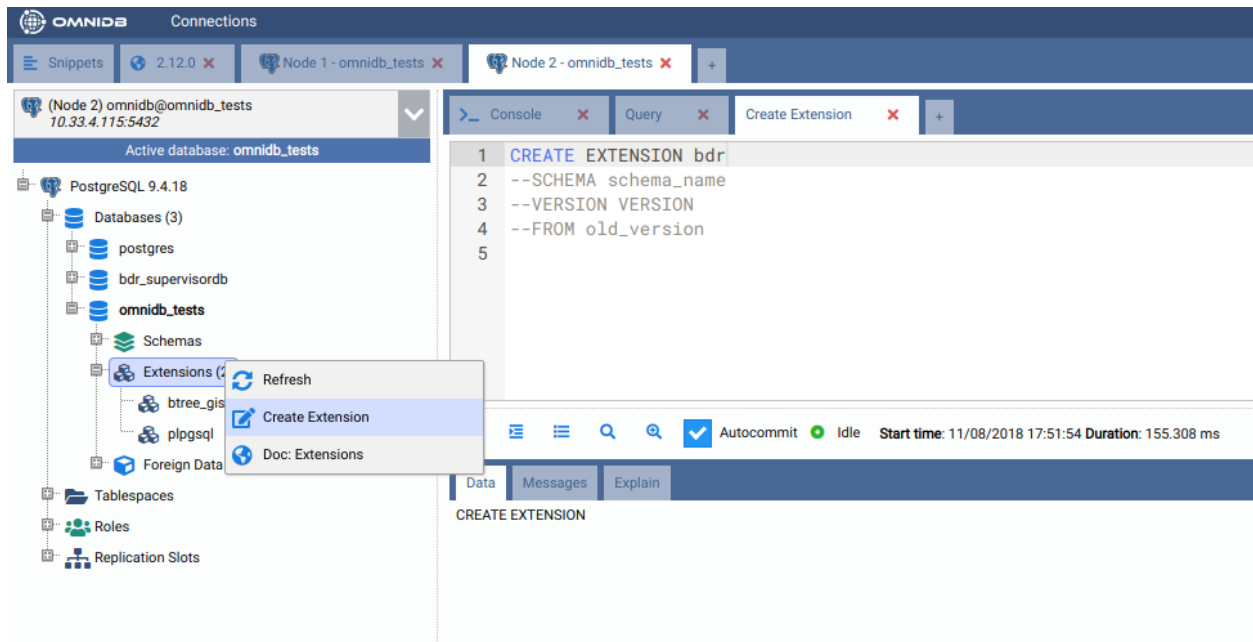
Technology	Server	Port	Database	User	Title	SSH Tunnel	SSH Server	SSH Port	SSH User	SSH Password	SSH Key	Actions
postgresql	10.33.4.114	5432	omnidb_tests	omnidb	Node 1	<input type="checkbox"/>		22				✗ ⚡ ✓
postgresql	10.33.4.115	5432	omnidb_tests	omnidb	Node 2	<input type="checkbox"/>		22				✗ ⚡ ✓

Then select both connections.

17.4 Create required extensions

BDR requires 2 extensions to be installed on each database that should have multi-master capabilities: `btree_gist` and `bdr`. Inside OmniDB, you can create both extensions by right clicking on the *Extensions* node, and choosing the

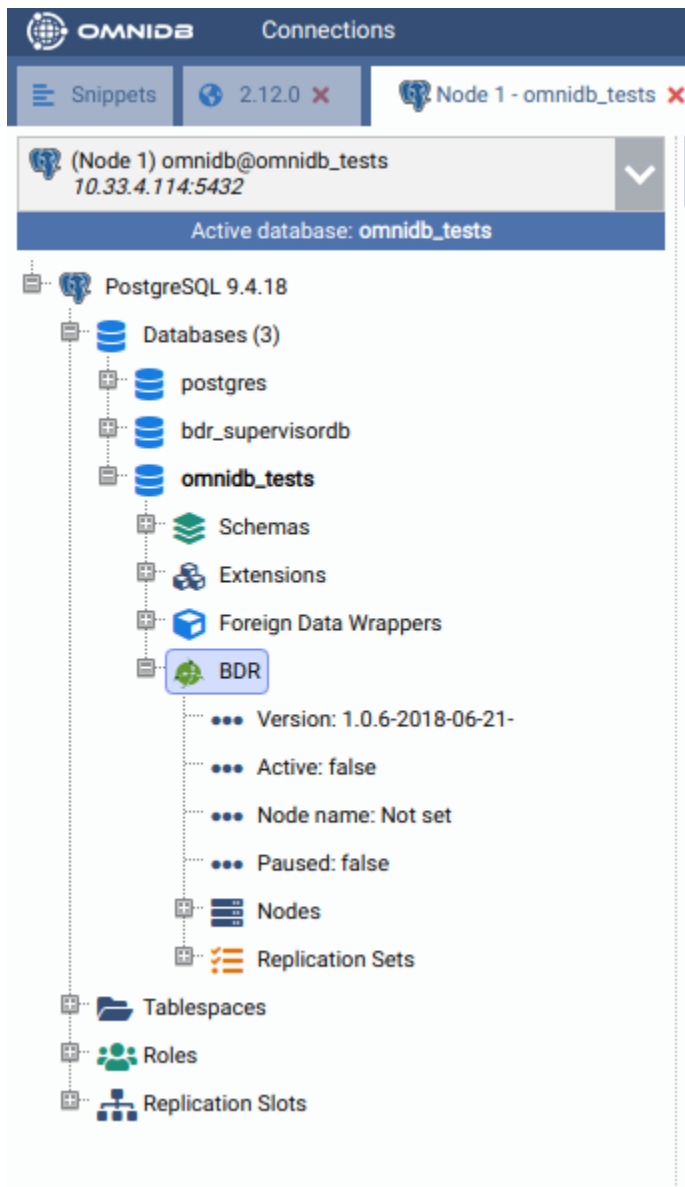
action *Create Extension*. OmniDB will open a SQL template tab with the `CREATE EXTENSION` command ready for you to make some adjustments and run:



You need to create both extensions `btree_gist` and `bdr` on both nodes.

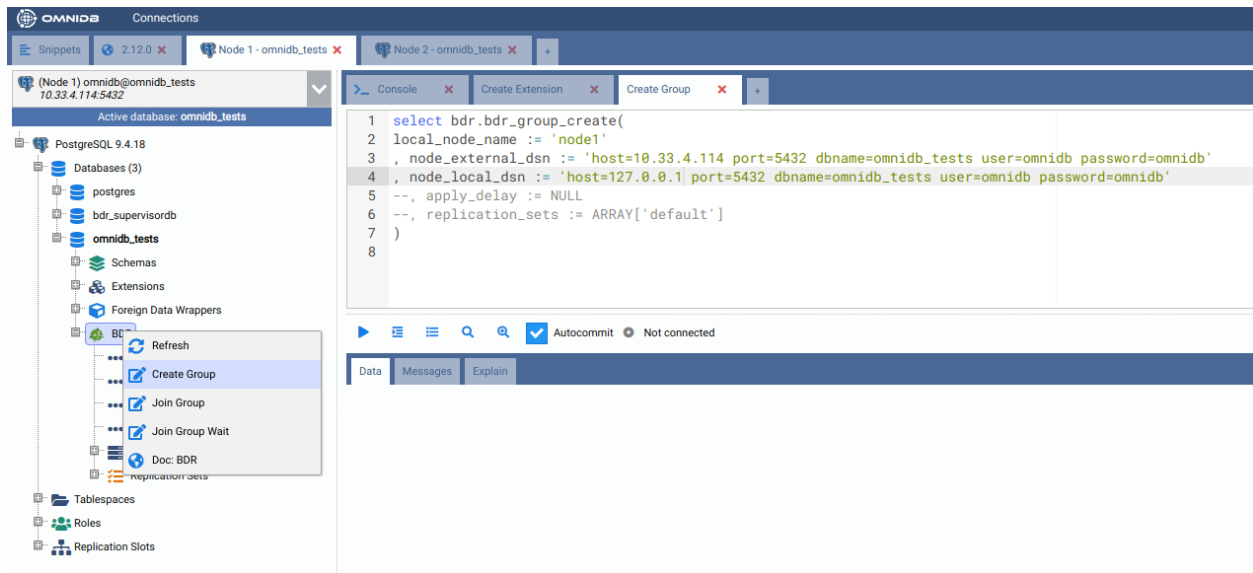
17.5 Create the BDR group in the first node

With both extensions installed, you can refresh the root node of the OmniDB tree view. A new *BDR* node will appear just inside your database. You can expand this node to see some informations about BDR:

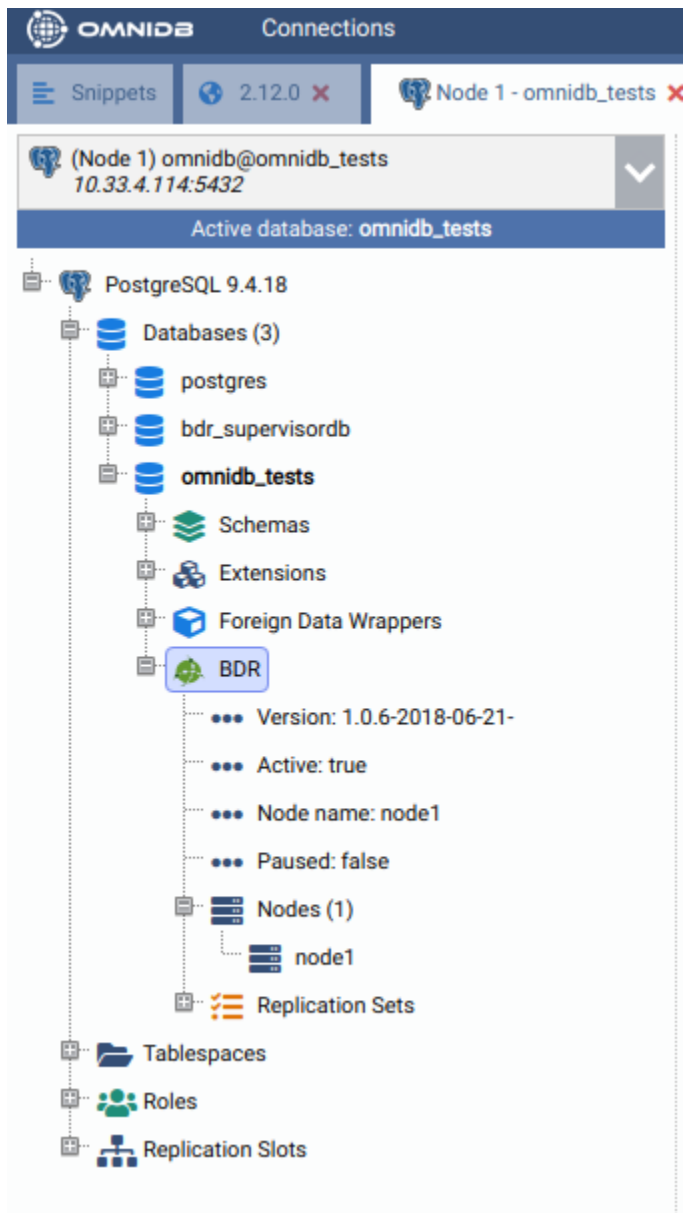


As you can see, BDR is not active yet. In the first node, we need to create a *BDR group*. The other nodes will join this group later.

To create a BDR group, right click in the *BDR* node. In the SQL template, adjust the node name and the node external connection info (the way other nodes will use to connect to this node):

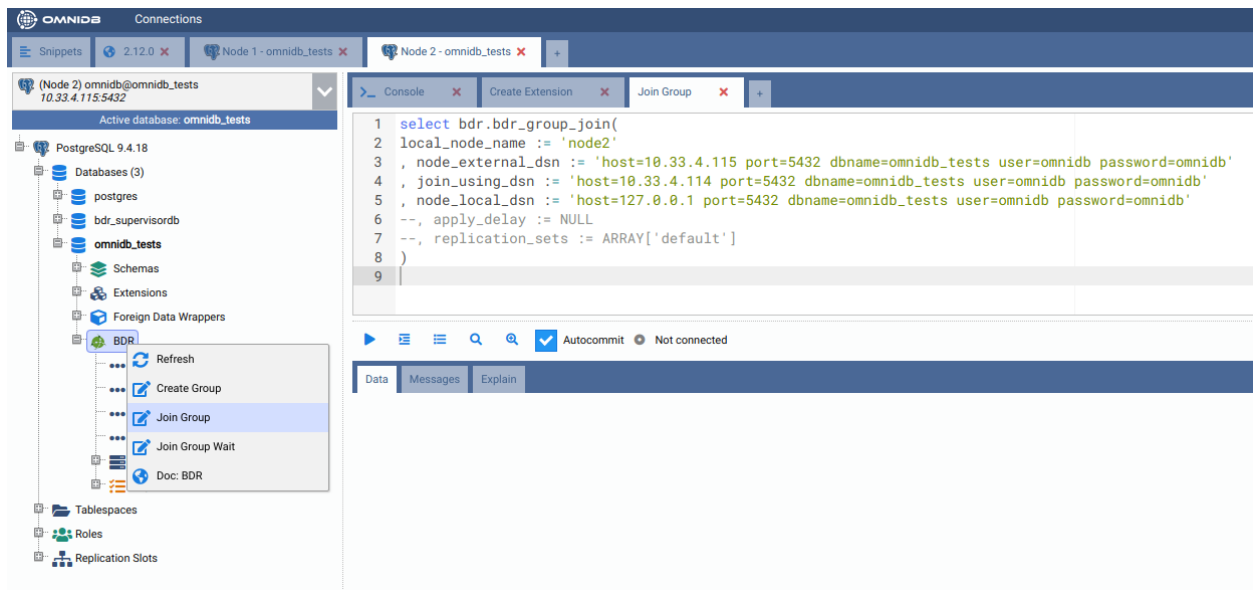


After you execute the above command, right click the *BDR* node and choose *Refresh*. You will see that now BDR is active in this node, now called *node1*. If you expand *Nodes*, you will see that this BDR group has only 1 node:

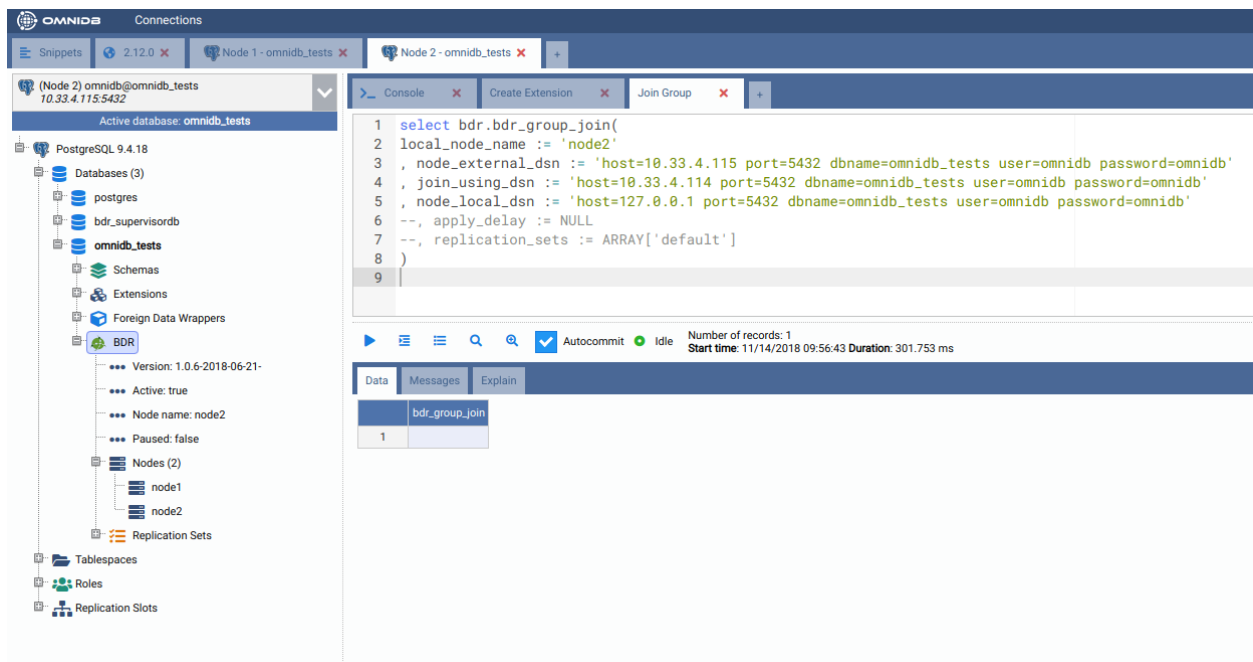


17.6 Join the BDR group in the second node

Now let's move to the other node. You can see that BDR is installed but not active yet. To link the two nodes, we will need to make this node join the BDR group that was previously created in the first node:

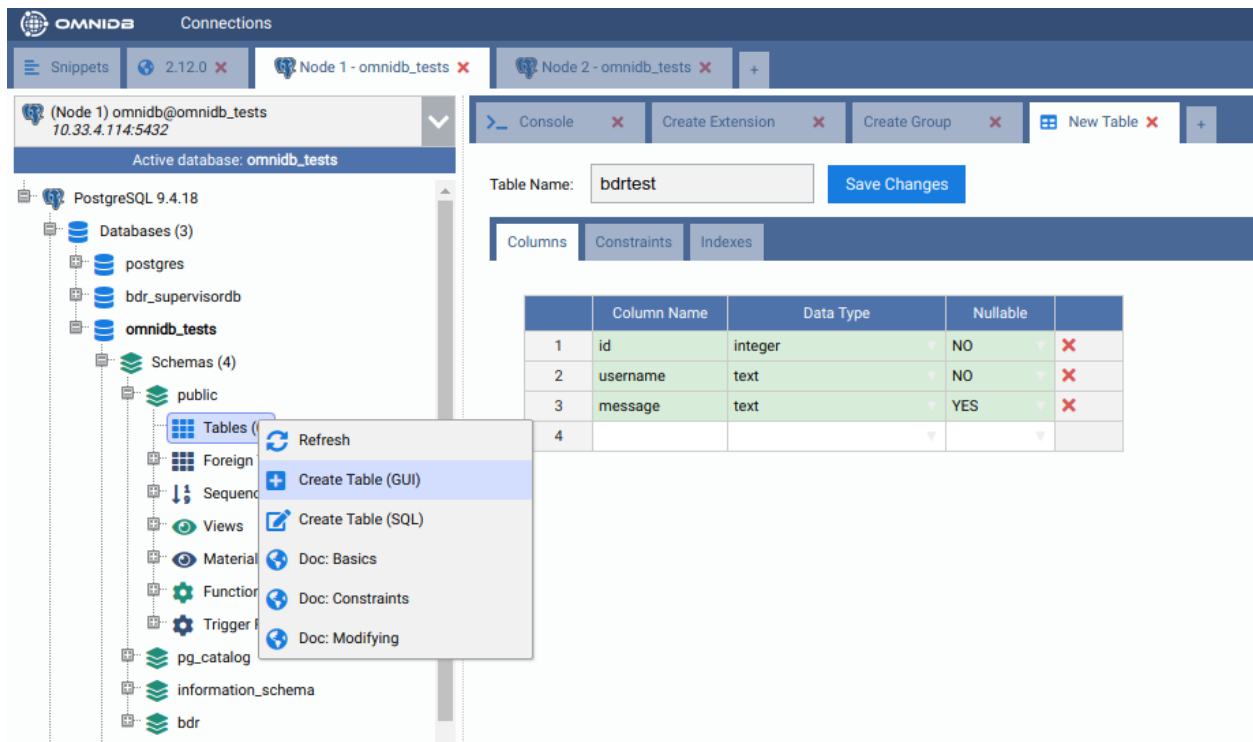


And now we can see that the second node has BDR active, his name in the BDR group is `node2`, and now the BDR group has 2 nodes:



17.7 Creating a table in the first node

Let's create a table in the first node. Expand the `public` schema, right click the *Tables* node and choose *Create Table*. Give the new table a name and add some columns. When done, click in the button *Save Changes*:



Now confirm that the table has been created in the first node by right clicking the *Tables* node and choosing *Refresh*. Go to the second node, expand the schema `public`, then expand the *Tables* node. Note that the table has been replicated from node1 to node2. If the table was created in the second node, it would have been created in the first node as well, because in BDR all nodes are masters.

OMNIDB Connections

Snippets 2.12.0 Node 1 - omnidb_tests Node 2 - omnidb_tests

(Node 2) omnidb@omnidb_tests 10.33.4.115:5432

Active database: omnidb_tests

PostgreSQL 9.4.18

Databases (3)

- postgres
- bdr_supervisordb
- omnidb_tests

Schemas (4)

- public

Tables (1)

- bdrtest

Columns (3)

- id
- username
- message

Primary Key

Foreign Keys

Uniques

Checks

Excludes

Indexes

Rules

Triggers

Inherited Tables

BDR

```

1 select bdr.bdr_group_join
2 local_node_name := 'node1'
3 , node_external_dsn := 'node2'
4 , join_using_dsn := 'node1'
5 , node_local_dsn := 'node1'
6 --, apply_delay := NULL
7 --, replication_sets := 'node1'
8 )
9

```

Console Create Extension

▶ ⌵ 🔍 🔍 Aut

Data Messages Explain

	bdr_group_join
1	

17.8 Adding some data in the second node

While you are at the second node, right click the table `bdrtest`, point to *Data Actions* and then click in *Edit Data*. Add some rows to this table. When finished, click in the *Save Changes* button.

The screenshot shows the OmniDB interface with a PostgreSQL 9.4.18 connection. The left sidebar displays the database structure, including the 'omnidb_tests' database, 'public' schema, and the 'bdrtest' table. A context menu is open over the 'bdrtest' table, showing options like 'Refresh', 'Data Actions', and 'Table Actions'. The 'Data Actions' menu is expanded, showing options like 'Query Data', 'Edit Data', 'Insert Record', 'Update Records', 'Count Records', 'Delete Records', and 'Truncate Table'. The 'Query Data' option is selected. The main window shows the query 'select * from public.bdrtest t' and its results. The results table has columns 'id (integer)', 'username (text)', and 'message (text)'. The first two rows show data inserted in Node 2, and the third row is a placeholder for a new record.

Active database: omnidb_tests

PostgreSQL 9.4.18

Databases (3)

- postgres
- bdr_supervisordb
- omnidb_tests

Schemas (4)

- public

Tables (1)

- bdrtest

Refresh

Data Actions

- Query Data
- Edit Data
- Insert Record
- Update Records
- Count Records
- Delete Records
- Truncate Table

Table Actions

message

Primary Key

Foreign Keys

Uniques

Checks

Excludes

Indexes

Rules

Triggers

Inherited Tables

BDR

Query 10 rows

Number of records: 0
Response time: 0.026 seconds

Save Changes

		id (integer)	username (text)	message (text)
1	✗	1	john	I was inserted in Node 2
2	✗	2	paul	I was inserted in Node 2 too
3	+			

Properties

DDL

Property	Value
Database	omnidb_tests
Schema	public
Table	bdrtest

Now go to the first node, right click the table, point to *Data Actions* and then click in *Query Data*. See how the rows created in node2 were automatically replicated into node1.

Active database: omnidb_tests

PostgreSQL 9.4.18

Databases (3)

- postgres
- bdr_supervisordb
- omnidb_tests

Schemas (4)

- public

Tables (1)

- bdrtest

Refresh

Data Actions

Table Actions

- Query Data
- Edit Data
- Insert Record
- Update Records
- Count Records
- Delete Records
- Truncate Table

```

1 SELECT t.id
2     , t.username
3     , t.message
4 FROM public.bdrtest t

```

Properties DDL

Property	Value
Database	omnidb_tests
Schema	public
Table	bdrtest
OID	24581
Owner	omnidb

Autocommit Idle Number of records: 2 Start time: 11/14/2018 10:16:07 Duration: 46.939 ms

	id	username	message
1	1	john	I was inserted in Node 2
2	2	paul	I was inserted in Node 2 too

17.9 Adding some data in the first node

Let's repeat the same procedure above, but instead of inserting rows from the second node, let's insert some rows while connected to the first node. Note how they replicate into the second node in the same way.

The screenshot displays the OmniDB web interface. At the top, there's a navigation bar with 'Connections' and tabs for 'Node 1 - omnidb_tests' and 'Node 2 - omnidb_tests'. Below this, a sidebar shows the database hierarchy: PostgreSQL 9.4.18 > Databases (3) > postgres > bdr_supervisordb > omnidb_tests > Schemas (4) > public > Tables (1) > bdrtest. A context menu is open over the 'bdrtest' table, showing options like 'Refresh', 'Data Actions', 'Table Actions', 'Uniques', 'Checks', 'Excludes', 'Indexes', 'Rules', 'Triggers', 'Inherited Tables', and 'BDR'. The 'Data Actions' menu is further expanded, showing options like 'Query Data', 'Edit Data', 'Insert Record', 'Update Records', 'Count Records', 'Delete Records', and 'Truncate Table'. The main panel shows a query result for 'select * from public.bdrtest t'. The query returned 2 records. Below the query result, there's a table with columns 'id (integer)', 'username (text)', and 'message (text)'. The table contains 5 rows of data. At the bottom, there's a 'Properties' tab showing the table's metadata, including 'Database: omnidb_tests', 'Schema: public', 'Table: bdrtest', 'OID: 24581', 'Owner: omnidb', 'Size: 8192 bytes', 'Tablespace: pg_default', and 'ACL'.

Active database: omnidb_tests

PostgreSQL 9.4.18

Databases (3)

- postgres
- bdr_supervisordb
- omnidb_tests

Schemas (4)

- public

Tables (1)

- bdrtest

Refresh

Data Actions

Table Actions

Uniques

Checks

Excludes

Indexes

Rules

Triggers

Inherited Tables

BDR

select * from public.bdrtest t

1

Query 10 rows

Number of records: 2
Response time: 0.035 seconds

Save Changes

	id (integer)	username (text)	message (text)
1	1	john	I was inserted in Node 2
2	2	paul	I was inserted in Node 2 too
3	3	ringo	I was inserted in Node 1
4	4	george	I am from Node 1 too
5	5	yoko	Node 1 too
6	+		

Property	Value
Database	omnidb_tests
Schema	public
Table	bdrtest
OID	24581
Owner	omnidb
Size	8192 bytes
Tablespace	pg_default
ACL	

OMNiDB

Connections

Snippets

2.12.0

Node 1 - omnidb_tests

Node 2 - omnidb_tests

+

(Node 2) omnidb@omnidb_tests
10.33.4.115:5432

Active database: omnidb_tests

PostgreSQL 9.4.18

Databases (3)

postgres

bdr_supervisordb

omnidb_tests

Schemas (4)

public

Tables (1)

bdrtest

Refresh

Data Actions

Table Actions

message

Primary Key

Foreign Keys

Uniques

Checks

Excludes

Indexes

Rules

Triggers

Inherited Tables

BDR

Console

Create Extension

Join Group

public.bdrtest

public.bdrtest

+

1 SELECT t.id

2 , t.username

3 , t.message

4 FROM public.bdrtest t

Autocommit

Idle

Number of records: 5

Start time: 11/14/2018 10:20:55

Duration: 47.932 ms

Properties

DDL

Property	Value
Database	omnidb_tests
Schema	public
Table	bdrtest
OID	24576
Owner	omnidb
Size	8192 bytes

Data

Messages

Explain

	id	username	message
1	1	john	I was inserted in Node 2
2	2	paul	I was inserted in Node 2 too
3	3	ringo	I was inserted in Node 1
4	4	george	I am from Node 1 too
5	5	yoko	Node 1 too

CHAPTER 18

18. Postgres-XL

Postgres-XL (or just **XL**, for short) is an open source project from 2ndQuadrant. It is a massively parallel database built on top of PostgreSQL, and it is designed to be horizontally scalable and flexible enough to handle various workloads.

In this chapter, we will use a cluster with 4 virtual machines: 1 GTM, 1 coordinator and 2 data nodes.

Machine	IP	Role
xlgtm	10.33.1.114	GTM
xlcoord	10.33.1.115	coordinator
xldata1	10.33.1.116	data node
xldata2	10.33.1.117	data node

On each machine, you need to clone Postgres-XL repository and compile it. You also need to set specific XL parameters on file `postgresql.conf` and make sure all machines are communicating to each other by adjusting file `pg_hba.conf`. More information on how Postgres-XL works and how to install it on [Postgres-XL documentation](#). You can also refer to [this blog post](#).

18.1 Creating a test environment

OmniDB repository provides a 4-node Vagrant test environment. If you want to use it, please do the following:

```
git clone --depth 1 https://github.com/OmniDB/OmniDB
cd OmniDB/OmniDB_app/tests/vagrant/xl-9.5/
vagrant up
```

It will take a while, but once finished, 4 virtual machines with IP addresses 10.33.1.114, 10.33.1.115, 10.33.1.116 and 10.33.1.117 will be up and each of them will have Postgres-XL 9.5 up and listening to port 5432, with all settings needed. To create all nodes, please do:

```
vagrant ssh xlcoord -c '/vagrant/setup.sh 10.33.1.115 10.33.1.116 10.33.1.117'
vagrant ssh xldata1 -c '/vagrant/setup.sh 10.33.1.115 10.33.1.116 10.33.1.117'
vagrant ssh xldata2 -c '/vagrant/setup.sh 10.33.1.115 10.33.1.116 10.33.1.117'
```

Then connect to the coordinator and define a password for the `postgres` user:

```
$ vagrant ssh xlcoord -c 'sudo su - postgres -c /usr/local/pgsql/bin/psql'
psql (PGXL 9.5r1.6, based on PG 9.5.12 (Postgres-XL 9.5r1.6))
Type "help" for help.

postgres=# ALTER USER postgres PASSWORD 'omnidb';
ALTER ROLE
postgres=#
```

Now the XL cluster will be ready for you to use.

18.2 Install OmniDB XL plugin

OmniDB core does not support XL by default. You will need to download and install XL plugin. If you are using OmniDB server, these are the steps:

```
wget https://omnidb.org/dist/plugins/omnidb-xl_1.0.0.zip
unzip omnidb-xl_1.0.0.zip
sudo cp -r plugins/ static/ /opt/omnidb-server/OmnidB_app/
sudo systemctl restart omnidb
```

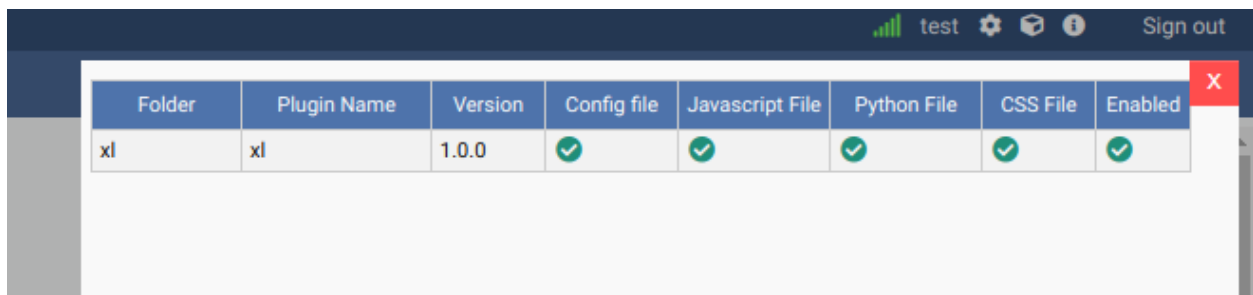
And then refresh the OmniDB web page in the browser.

For OmniDB app, these are the steps:

```
wget https://omnidb.org/dist/plugins/omnidb-xl_1.0.0.zip
unzip omnidb-xl_1.0.0.zip
sudo cp -r plugins/ static/ /opt/omnidb-app/resources/app/omnidb-server/OmnidB_app/
```

And then restart OmniDB app.

If everything worked correctly, by clicking on the “plugins” icon in the top right corner, you will see the plugin installed and enabled:



The screenshot shows the OmniDB web interface with a dark blue header. In the top right corner, there are icons for signal strength, a 'test' button, a settings gear, a cube icon, an information 'i' icon, and a 'Sign out' button. Below the header is a table with 8 columns: Folder, Plugin Name, Version, Config file, Javascript File, Python File, CSS File, and Enabled. The table contains one row for the 'xl' plugin, where all fields show a green checkmark, indicating it is installed and enabled. A red 'X' icon is visible in the top right corner of the table area.

Folder	Plugin Name	Version	Config file	Javascript File	Python File	CSS File	Enabled
xl	xl	1.0.0	✓	✓	✓	✓	✓

18.3 Connecting to the cluster

Let's use OmniDB to connect to the coordinator node. First of all, fill out connection info in the connection grid:

New Connection												
Technology	Server	Port	Database	User	Title	SSH Tunnel	SSH Server	SSH Port	SSH User	SSH Password	SSH Key	Actions
postgresql	10.33.1.115	5432	postgres	postgres		<input type="checkbox"/>		22				✖ ⚡ ✔

Then select the connection. You will see OmniDB workspace window. Expand the tree root node. Note that OmniDB identifies it is connected to a Postgres-XL cluster and shows a specific node called *Postgres-XL* just inside the tree root node. Expand this node to see all the nodes we have in our cluster:

The screenshot shows the OmniDB workspace window with the 'Connections' tab selected. The active connection is 'postgres@postgres' on '10.33.1.115:5432'. The active database is 'postgres'. The tree view shows the following structure:

- PostgreSQL 9.5.12 (Postgres-XL 9.5r1.6)
 - Databases
 - Tablespaces
 - Roles
 - Replication Slots
 - Postgres-XL
 - Nodes (3)
 - xlcoord
 - xldata1
 - Type: datanode
 - Host: 10.33.1.116
 - Port: 5432
 - Primary: false
 - Preferred: false
 - xldata2
 - Groups

18.4 Creating a HASH table

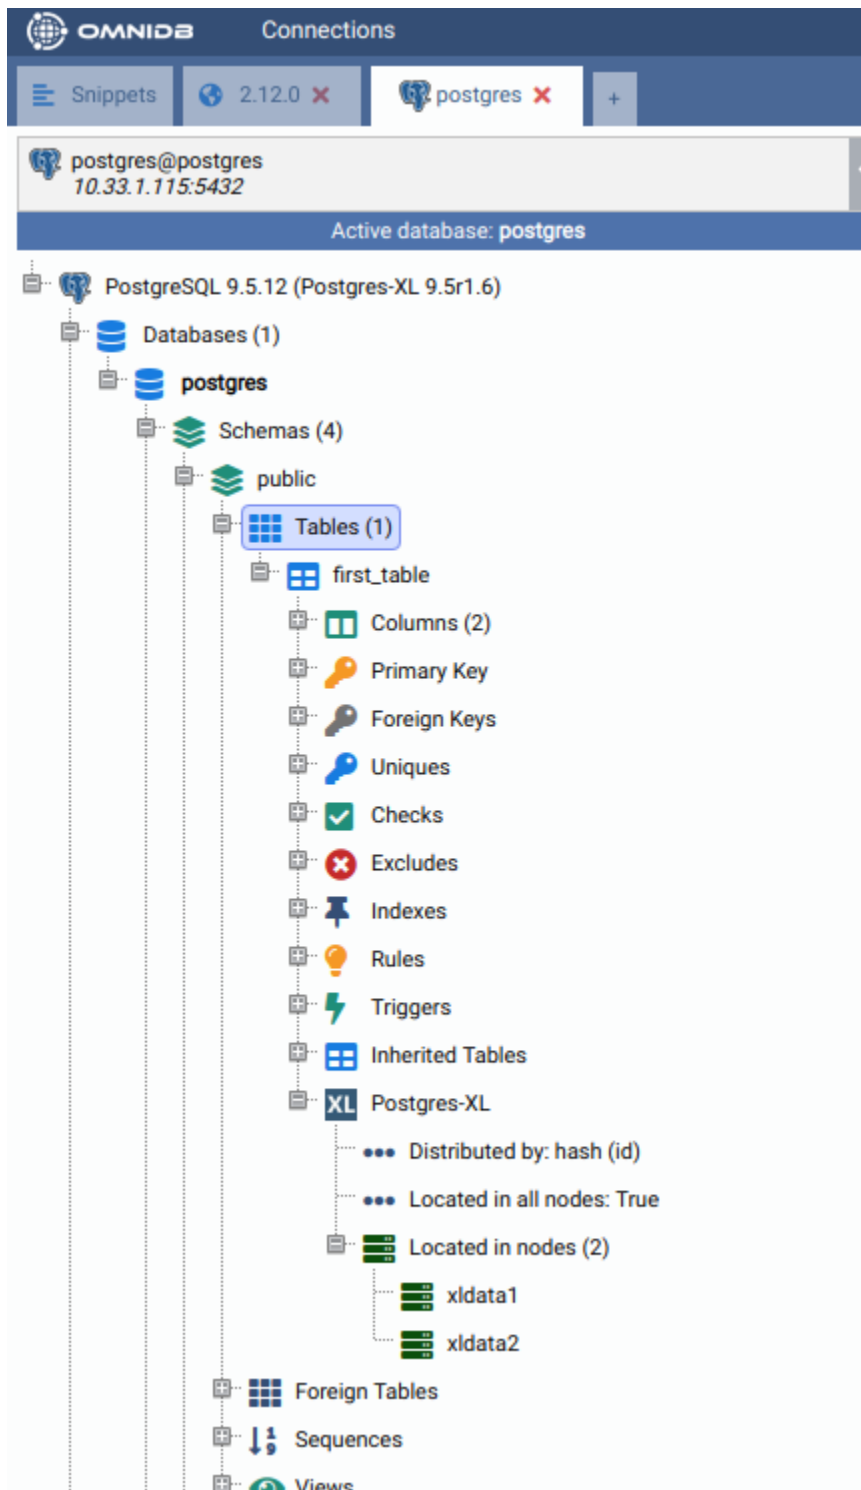
From the root node, expand *Schemas*, then *public*, then right click on the *Tables* node. Click on *Create Table*. Name your new table, add some columns to it and do not forget to add a primary key too:

The screenshot shows the OmniDB interface with a PostgreSQL 9.5.12 connection. The left sidebar shows the database structure, including the *public* schema and the *Tables* node. A context menu is open over the *Tables* node, with *Create Table (GUI)* selected. The main panel shows the 'New Table' dialog with the table name 'first_table' and a 'Save Changes' button. Below the dialog, the 'Columns' tab is active, showing a table with 3 columns: 'id' (integer, nullable), 'name' (text, nullable), and an empty row. The 'Constraints' tab is also visible, showing a 'New Constraint' dialog with a table for defining constraints. The table has 7 columns: Constraint Name, Type, Columns, Referenced Table, Referenced Columns, Delete Rule, and Update Rule. A primary key constraint 'pk_first_table' is defined on the 'id' column.

	Column Name	Data Type	Nullable	
1	id	integer	NO	✗
2	name	text	YES	✗
3				

Constraint Name	Type	Columns	Referenced Table	Referenced Columns	Delete Rule	Update Rule	
pk_first_table	Primary Key	id					✗

When done, click on the *Save Changes* button. Now right click on the *Tables* node and click on *Refresh*. You will see the new table created. Expand it to see that there is also a *Postgres-XL* node inside of it. Check its properties.



By default, Postgres-XL always try to create a table distributed by HASH. It means that the data will be split into the nodes regularly, through a hash function applied on the specified column. If present, it will use the primary key, or a unique constraint otherwise. If there is no primary key nor unique constraint, Postgres-XL uses the first eligible column. If not possible to distribute by HASH, then Postgres-XL will create the table distributed by ROUNDROBIN, which means that the data will be split in a way that every new row will be added to a different data node.

Now let's add some rows in our new table. Right click on the table, then go to *Data Actions* and then click on *Edit*

Data. Add some rows and then click on the *Save Changes* button:

The screenshot shows the OmniDB interface with a PostgreSQL connection. The left sidebar displays the database structure: Databases (1) - postgres - Schemas (4) - public - Tables (1) - first_table. A context menu is open over the 'first_table' table, showing options like Refresh, Data Actions, and Table Actions. The 'Data Actions' menu is expanded, showing options like Query Data, Edit Data, Insert Record, Update Records, Count Records, Delete Records, and Truncate Table. The 'Query Data' option is selected. The main console shows the SQL query: `select * from public.first_table t order by t.id`. Below the console, a table displays the query results:

		id (integer)	name (text)
1	✗	1	John
2	✗	2	Paul
3	✗	3	Ringo
4	✗	4	George
5	✗	5	Yoko
6	+		

At the bottom right, there is a 'Query 10 rows' dropdown, 'Number of records: 0', 'Response time: 0.051 seconds', and a 'Save Changes' button.

Right click on the table again, *Data Actions*, *Query Data*. You will see that cluster-wide the table has all data inside.

OmniDB Connections

Active database: postgres

PostgreSQL 9.5.12 (Postgres-XL 9.5r1.6)

Databases (1)

postgres

Schemas (4)

public

Tables (1)

first_table

Refresh

Data Actions

Table Actions

Query Data

Edit Data

Insert Record

Update Records

Count Records

Delete Records

Truncate Table

```

1 SELECT t.id
2     , t.name
3 FROM public.first_table t
4 ORDER BY t.id

```

Autocommit Idle Number of records: 5 Start time: 11/14/2018 15:12:28 Duration: 59.519 ms

	id	name
1	1	John
2	2	Paul
3	3	Ringo
4	4	George
5	5	Yoko

Properties DDL

Property	Value
Database	postgres

But how the data was distributed in the data nodes? In the *Postgres-XL* main node, right click on each node and click on *Execute Direct*. Adjust the query that will be executed directly into the data node, as you can see below.

The screenshot shows the OmniDB interface with a PostgreSQL connection. The left sidebar displays the database structure, including the 'Postgres-XL' section with nodes 'xlcoord', 'xldata1', and 'xldata2'. A context menu is open over 'xldata1', showing options like 'Refresh', 'Execute Direct', 'Pool Reload', 'Alter Node', and 'Drop Node'. The main console shows the execution of a query: `EXECUTE DIRECT ON (xldata1) 'SELECT * FROM first_table'`. The result is displayed in a table with 3 records.

Active database: postgres

PostgreSQL 9.5.12 (Postgres-XL 9.5r1.6)

Databases (1)

- postgres
 - Schemas (4)
 - Extensions
 - Foreign Data Wrappers
 - Tablespaces
 - Roles
 - Replication Slots
 - Postgres-XL
 - Nodes (3)
 - xlcoord
 - xldata1 (selected)
 - Refresh
 - Execute Direct
 - Pool Reload
 - Alter Node
 - Drop Node
 - xldata2
 - Groups

Console

```
1 EXECUTE DIRECT ON (xldata1)
2 'SELECT * FROM first_table'
3
```

Autocommit Idle Number of records: 3 Start time: 11/14/2018 15:19:12 Duration: 52.605 ms

	id	name
1	1	John
2	2	Paul
3	5	Yoko

The screenshot shows the OmniDB interface with a PostgreSQL connection. The left sidebar displays the database structure, including the 'Postgres-XL' section with nodes 'xlcoord', 'xldata1', and 'xldata2'. A context menu is open over 'xldata1', showing options like 'Refresh', 'Execute Direct', 'Pool Reload', 'Alter Node', and 'Drop Node'. The main console shows the execution of a query: `EXECUTE DIRECT ON (xldata2) 'SELECT * FROM first_table'`. The result is displayed in a table with 2 records.

Active database: postgres

PostgreSQL 9.5.12 (Postgres-XL 9.5r1.6)

Databases (1)

- postgres
 - Schemas (4)
 - Extensions
 - Foreign Data Wrappers
 - Tablespaces
 - Roles
 - Replication Slots
 - Postgres-XL
 - Nodes (3)
 - xlcoord
 - xldata1
 - xldata2 (selected)
 - Refresh
 - Execute Direct
 - Pool Reload
 - Alter Node
 - Drop Node
 - Groups

Console

```
1 EXECUTE DIRECT ON (xldata2)
2 'SELECT * FROM first_table'
3
```

Autocommit Idle Number of records: 2 Start time: 11/14/2018 15:20:50 Duration: 52.994 ms

	id	name
1	3	Ringo
2	4	George

18.5 Creating a REPLICATION table

While HASH distribution is great for write-only and write-mainly tables, REPLICATION distribution is great for read-only and read-mainly tables. However, a table distributed by REPLICATION will store all data in all nodes it is located.

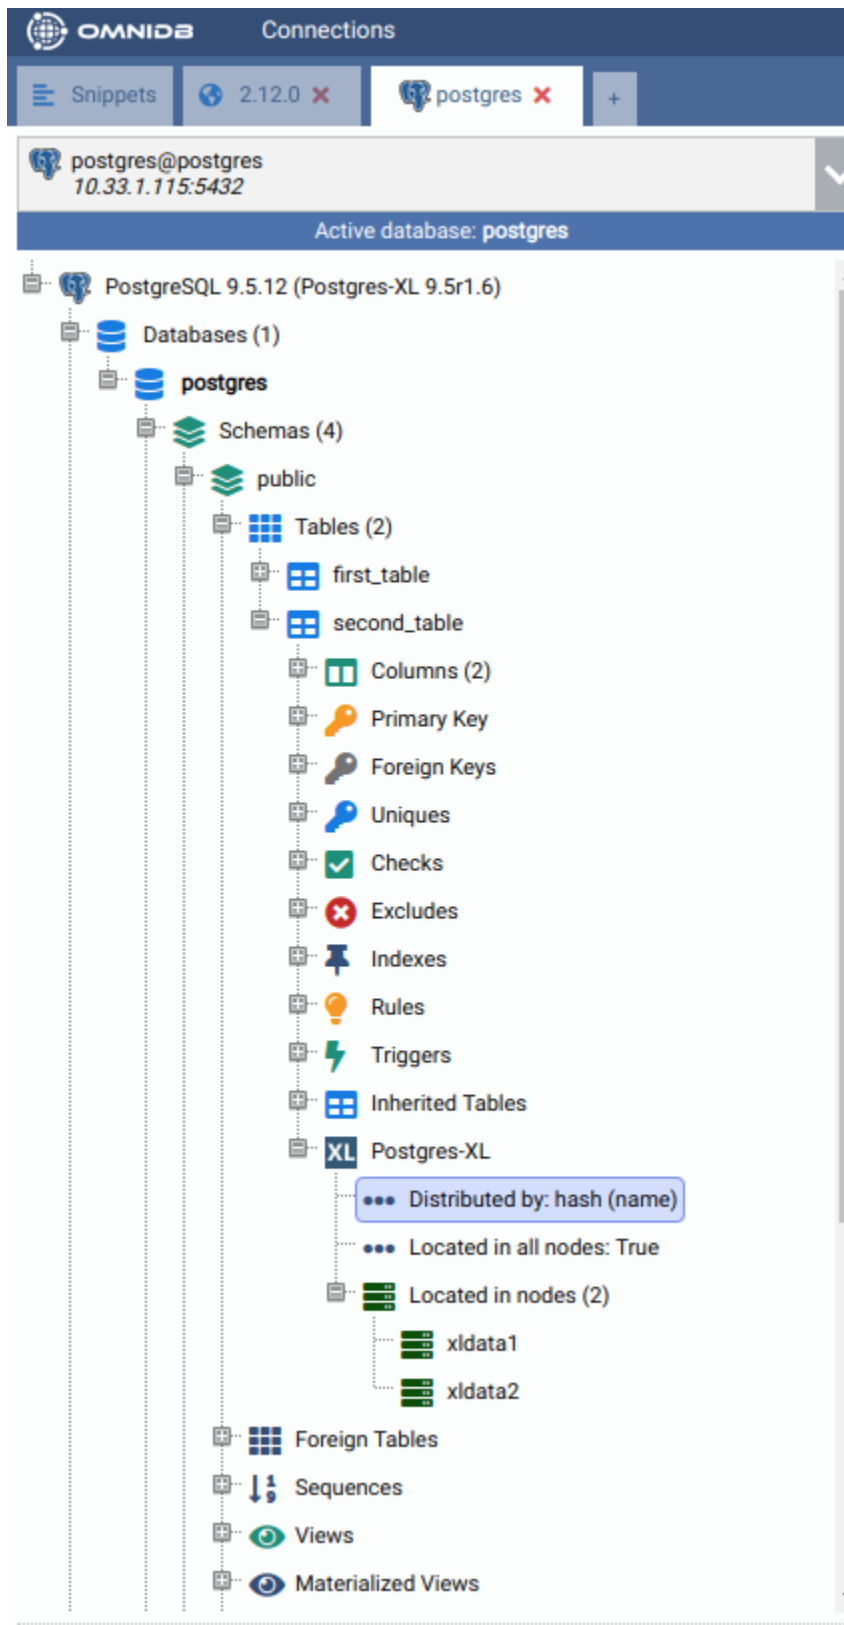
In order to create a REPLICATION table, let us create a new table like we did before:

The screenshot shows the OmniDB interface with a PostgreSQL 9.5.12 connection. The left sidebar displays the database structure, including the 'public' schema and a list of tables. A context menu is open over the 'Tables (1)' list, showing options like 'Refresh', 'Create Table (GUI)', and 'Create Table (SQL)'. The main panel shows the 'New Table' dialog for 'second_table'. The 'Columns' tab is active, displaying a table with columns 'id' (integer, nullable) and 'name' (text, nullable). The 'Constraints' tab is also visible, showing a 'New Constraint' dialog with a table for defining constraints. The table has columns: Constraint Name, Type, Columns, Referenced Table, Referenced Columns, Delete Rule, Update Rule, and a status column. The first row shows 'pk_second_table' as a Primary Key on the 'name' column.

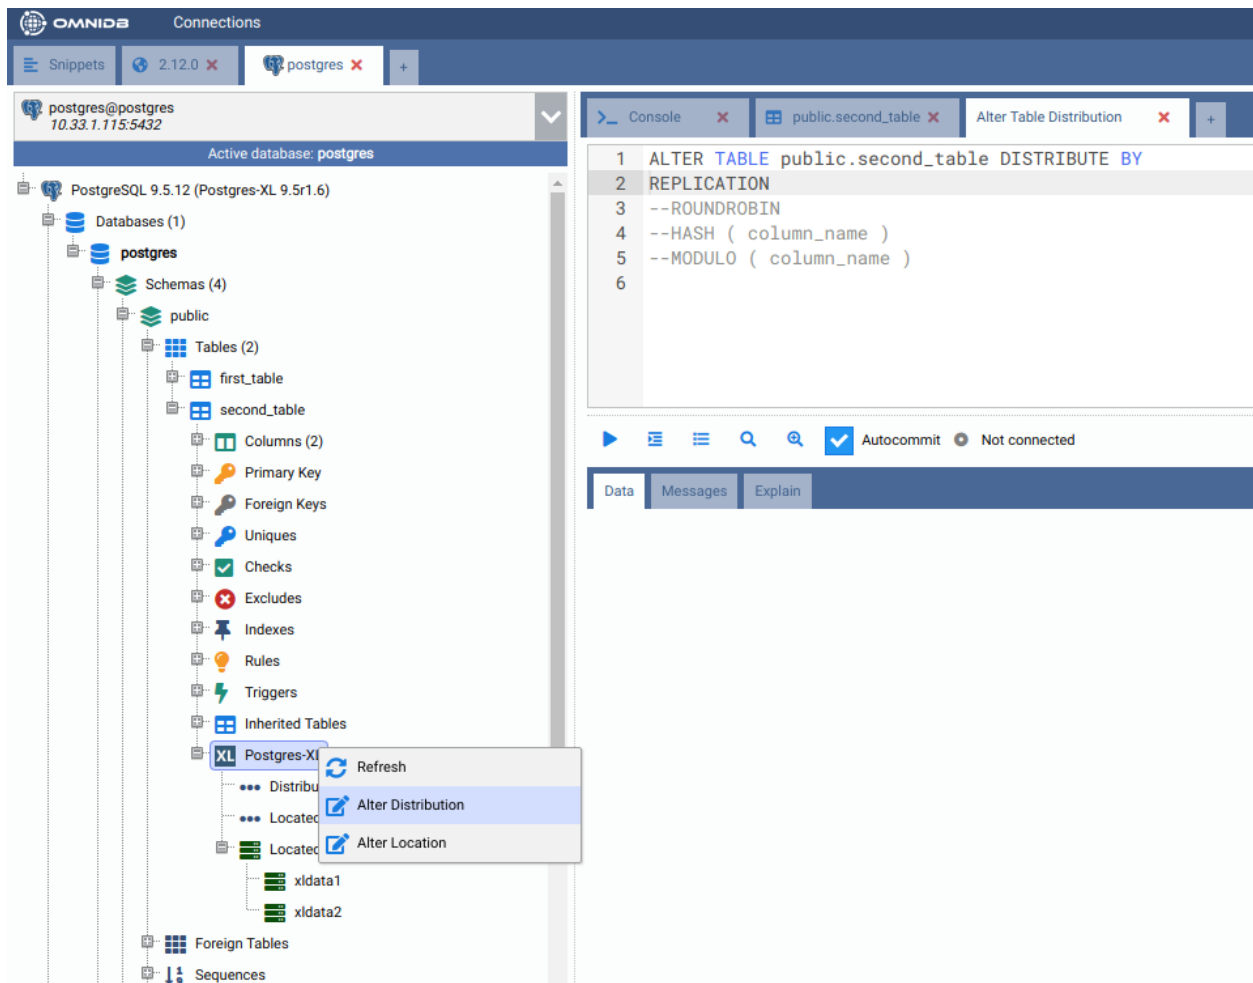
	Column Name	Data Type	Nullable	
1	id	integer	NO	✗
2	name	text	YES	✗
3				

Constraint Name	Type	Columns	Referenced Table	Referenced Columns	Delete Rule	Update Rule	
pk_second_table	Primary Key	name					✗

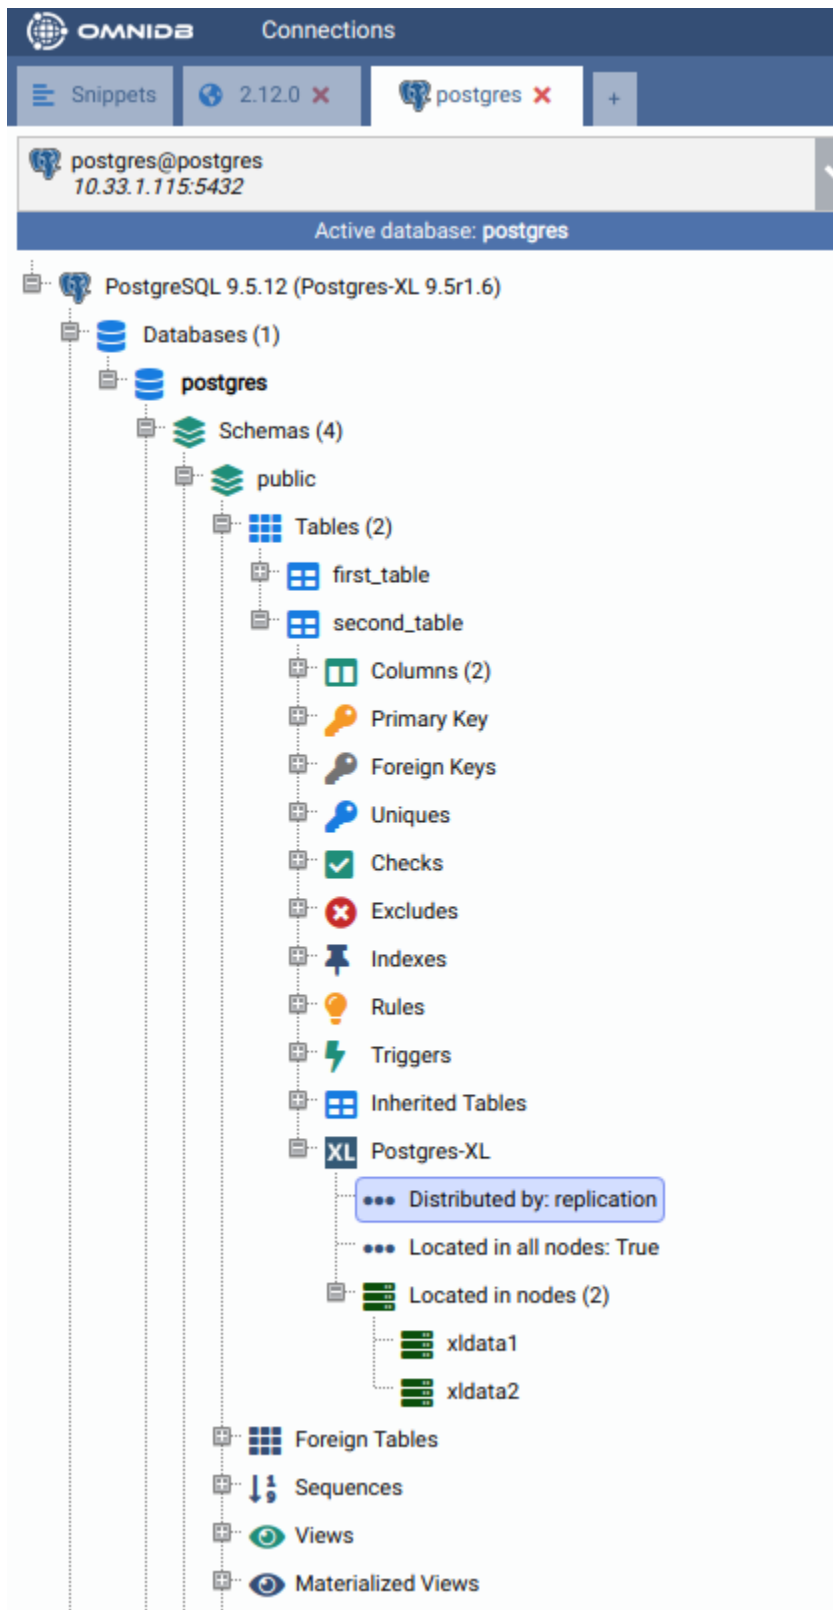
Note how by default it was created as a HASH table:



Let us change the distribution type of the table by right-clicking on the *Postgres-XL* node inside the table, and then clicking on *Alter Distribution*. Uncomment the “REPLICATION” line and execute the command:



You can check the distribution was successfully changed by right-clicking on the *Postgres-XL* node and clicking on *Refresh*. The properties will now show *Distributed by: replication*.



Now add some data to the table:

The screenshot shows the OmniDB interface with a PostgreSQL 9.5.12 connection. The left sidebar displays the database structure: Databases (1) - postgres - Schemas (4) - public - Tables (2) - first_table, second_table. A context menu is open for 'second_table', showing options like 'Query Data', 'Edit Data', 'Insert Record', 'Update Records', 'Count Records', 'Delete Records', and 'Truncate Table'. The SQL console on the right contains the query: `select * from public.second_table t order by t.name`. Below the console, a results table shows 5 records with columns 'id (integer)' and 'name (text)'. The records are: (1, John), (2, Paul), (3, Ringo), (4, George), and (5, Yoko). The bottom status bar shows 'Query 10 rows', 'Number of records: 0', 'Response time: 0.047 seconds', and a 'Save Changes' button.

		id (integer)	name (text)
1	✗	1	John
2	✗	2	Paul
3	✗	3	Ringo
4	✗	4	George
5	✗	5	Yoko
6	+		

And then check that all data exist on all data nodes:

The screenshot displays the OmniDB interface for a PostgreSQL-XL database. The top bar shows the connection 'postgres@postgres 10.33.1.115:5432' and the active database 'postgres'. The left sidebar shows a tree view of the database structure, including Schemas (public), Tables, Foreign Tables, Sequences, Views, Materialized Views, Functions, Trigger Functions, pg_catalog, information_schema, storm_catalog, Extensions, Foreign Data Wrappers, Tablespaces, Roles, Replication Slots, and PostgreSQL-XL. The PostgreSQL-XL section shows Nodes (3) with xlcoord and xldata nodes, and a Groups section. A context menu is open over the xldata node, showing options: Refresh, Execute Direct, Pool Reload, Alter Node, and Drop Node. The right pane shows the SQL console with the query: `EXECUTE DIRECT ON (xldata1) 'SELECT * FROM second_table'`. Below the console, the execution status is shown: Autocommit, Idle, Number of records: 5, Start time: 11/14/2018 15:32:00, Duration: 58.984 ms. The bottom pane shows the 'Data' tab with a table of results:

	id	name
1	1	John
2	2	Paul
3	3	Ringo
4	4	George
5	5	Yoko

OMNiDB

Connections

Snippets2.12.0postgres+

postgres@postgres10.33.1.115:5432

Active database: postgres

PostgreSQL 9.5.12 (Postgres-XL 9.5r1.6)

Databases (1)

postgres

Schemas (4)

public

Tables

Foreign Tables

Sequences

Views

Materialized Views

Functions

Trigger Functions

pg_catalog

information_schema

storm_catalog

Extensions

Foreign Data Wrappers

Tablespaces

Roles

Replication Slots

Postgres-XL

Nodes (3)

xlcoord

xldata1

xldata

Groups

Refresh

Execute Direct

Pool Reload

Alter Node

Drop Node

PropertiesDDL

PropertyValue

Consolepublic.second_tableExecute DirectExecute Direct+

1EXECUTE DIRECT ON (xldata2)

2'SELECT * FROM second_table'

3

AutocommitIdleNumber of records: 5Start time: 11/14/2018 15:32:50 Duration: 52.696 ms

DataMessagesExplain

	id	name
1	1	John
2	2	Paul
3	3	Ringo
4	4	George
5	5	Yoko

19. Deploying omnidb-server

Whenever deploying omnidb-server the user must be aware of how OmniDB works in terms of ports so the environment can be properly configured taking the infrastructure into account.

OmniDB uses 2 servers to answer user requests, one is the default webserver serving the application itself and the other is a websocket server used by several parts of OmniDB, such as Query, Console and Debugging Tab, allowing a bi-directional communication between the client and the server which enhances performance and user experience. This means that 2 ports need to be properly configured:

- OmniDB server:
 - **Technology:** CherryPy
 - **Default port:** 8000
- Websocket server:
 - **Technology:** Tornado
 - **Default port:** 25482

Both servers support SSL so OmniDB can run by itself securely without the need of a load balancer or reverse proxy, such as Nginx.

The configuration of ports and certificates can be done via command options or configuration file.

19.1 Command options

```
Usage: omnidb-server [options]
```

Options:

```
--version          show program's version number and exit
-h, --help         show this help message and exit
-H HOST, --host=HOST listening address
-p PORT, --port=PORT listening port
-w WSPORT, --wspport=WSPORT
```

(continues on next page)

(continued from previous page)

```

                                websocket port
-e EWSPORT, --ewsport=EWSPORT
                                external websocket port
-d HOMEDIR, --homedir=HOMEDIR
                                home directory containing local databases config and
                                log files
-c CONF, --configfile=CONF
                                configuration file

```

- `-H` specifies in what addresses the servers will listen, the default value is `0.0.0.0` meaning that all addresses bound to the machine will be used (127.0.0.1, 192.168.0.100, 162.154.12.35, for example).
- `-p` specifies in what port OmniDB server will listen, this is the port used in the browser's URL if OmniDB is being accessed directly. The default value is

1.

- `-w` specifies in what port the websocket server will listen. If OmniDB is being accessed directly the websocket client will connect to this port. The default value is 25482.
- `-e` specifies in what port the websocket client (the page opened in your browser) will connect. This option is used when OmniDB is behind a load balancer and the tornado server isn't being accessed directly, in this case we must tell websocket client what port to use. If not specified the client will use the port specified in `-w`.
- `-d` This parameter let's the user choose what folder will store the persistent files, such as `omnidb.conf`, `omnidb.log`, `db.sqlite3` (sessions database) and `omnidb.db` (application database). With this option is possible to have several instances of `omnidb-server` running, each one pointing to a specific directory. It also facilitates the deployment with Docker as it enables to point OmniDB to a mounted volume.
- `-c` Points OmniDB to a specific configuration file, can be used along with `-d` to specify a storage folder but choosing a specific config file.

19.2 Configuration File

The configuration file, `omnidb.conf` by default, can be used to set all the parameters specified in the previous category and a few additional parameter related to SSL and some about the query server itself.

This file is created when OmniDB is started for the first time or when a new folder is specified with the option `-d`. If no folder is specified the default location for files is:

- Linux: `~/.omnidb/omnidb-server/`
- Windows: `User Folder/.omnidb/omnidb-server/`

Here is the default configuration file:

```

# OmniDB Server configuration file

[webserver]

# What address the webserver listens to, 0.0.0.0 listens to all addresses bound to_
↪the machine
listening_address    = 127.0.0.1

# Webserver port, if port is in use another random port will be selected
listening_port      = 8000

```

(continues on next page)

(continued from previous page)

```
# Websocket port, if port is in use another random port will be selected
websocket_port = 25482

# External Websocket port, use this parameter if OmniDB isn't directly visible by the
# client
external_websocket_port = 25482

# Security parameters
# is_ssl = True requires ssl_certificate_file and ssl_key_file parameters
# This is highly recommended to protect information
is_ssl = False
ssl_certificate_file = /path/to/cert_file
ssl_key_file = /path/to/key_file

# Trusted origins, use this parameter if OmniDB is configured with SSL and is being
# accessed by another domain
csrf_trusted_origins = origin1,origin2,origin3

[queryserver]

#Max number of threads that can used by each advanced object search request
thread_pool_max_workers = 2

#Number of seconds between each prompt password request. Default: 30 minutes
pwd_timeout_total = 1800
```

- `is_ssl`: specifies whether to run securely or not.
- `ssl_certificate_file`: path to the certificate file.
- `ssl_key_file`: path to the key file.
- `csrf_trusted_origins`: list of trusted origins. When OmniDB is started with SSL and the browser is accessing it through another domain this parameter must specifies the domain in order to properly establish communication.
- `thread_pool_max_workers`: defines the max number of threads that can be used in advanced object search requests. That feature uses such mechanism to perform searches in parallel. This requires a tuning. Too much workers can be even worse than less of them.
- `pwd_timeout_total`: defines the timeout of typed password in the interface, that is, the time before the last typed password being considered as expired. The value is set in seconds. Defaults to 30 minutes.

Let's take a look on how to deploy `omnodb-server` in different scenarios:

19.3 Deploying OmniDB directly

In this case no load balancers or reverse proxies are used, OmniDB is accessed directly and is **extremely** recommended to start it with SSL enabled if it will be visible to the outside world.

For this scenario the user needs to specify the following parameters:

- `-H` or `listening_address`: Specify the address visible to the clients, can be a domain.
- `-p` or `listening_port`: Specify a port that will be used in the browser url: `https://mydomain.com:PORT`
- `-w` or `websocket_port`: Specify a port that will be used by javascript to connect to Tornado server directly.

- `is_ssl`: True
- `ssl_certificate_file`: /path/to/file
- `ssl_key_file`: /path/to/file
- `-e` or `external_websocket_port`: external websocket port isn't needed as `-w` will be used directly.

It is important to mention here that both ports need to be visible to every client trying to access OmniDB.

19.4 Deploying OmniDB behind a reverse proxy

In this case OmniDB won't be accessed directly but through a properly configured load balancer or reverse proxy.

For this scenario a possible approach is to run `omnidb-server` listening to the local address `127.0.0.1` and without SSL, given that the balancer will handle the security part.

The following parameters are required:

- `-H` or `listening_address`: `127.0.0.1`.
- `-p` or `listening_port`: Specify a port to which the load balancer will redirect all the OmniDB server requests.
- `-w` or `websocket_port`: Specify a port to which the load balancer will redirect all the Websocket server requests.
- `-e` or `external_websocket_port`: Specify a port that will be used by JavaScript to connect to Tornado server. Since OmniDB is behind a load balancer, a port being listened by the load balancer should be specified here and the balancer will redirect all requests to the port specified with `-w`. It is possible to specify the same port used to access OmniDB but then the load balancer needs to proxy requests to the specific server according to the URL pattern.

Consider this example of OmniDB being hosted behind Nginx:

- Starting `omnidb-server`:

```
omnidb-server -H 127.0.0.1 -p 9000 -w 26500 -e 443
```

In this case OmniDB can only be accessed locally and the browser will try to connect to the websocket server with the default https port (443).

- Nginx configuration file:

```
server {
    listen 443 ssl;
    listen [::]:443 ssl;
    include snippets/ssl-domain.conf;
    include snippets/ssl-params.conf;
    server_name domain.org;
    client_max_body_size 75M;

    location /wss {
        proxy_pass http://127.0.0.1:26500;
        proxy_set_header    X-Real-IP $remote_addr;
        proxy_set_header    X-Forwarded-For $proxy_add_x_forwarded_for;
        proxy_set_header    X-Forwarded-Ssl https;
        proxy_set_header    X-Forwarded-Proto https;
        proxy_set_header    X-Forwarded-Port 443;
        proxy_set_header    Host $host;
```

(continues on next page)

(continued from previous page)

```
    proxy_http_version 1.1;
    proxy_set_header Upgrade $http_upgrade;
    proxy_set_header Connection "upgrade";
}

location / {
    proxy_pass http://127.0.0.1:9000;
    proxy_set_header    X-Real-IP $remote_addr;
    proxy_set_header    X-Forwarded-For $proxy_add_x_forwarded_for;
    proxy_set_header    X-Forwarded-SSL https;
    proxy_set_header    X-Forwarded-Port 443;
    proxy_set_header    Host $host;
    proxy_http_version 1.1;
    proxy_set_header Upgrade $http_upgrade;
    proxy_set_header Connection "upgrade";
}
```

As can be seen, Nginx is listening for requests to `domain.org` in port 443. Since we also specified the external websocket port to 443, websocket requests will be dealt here too.

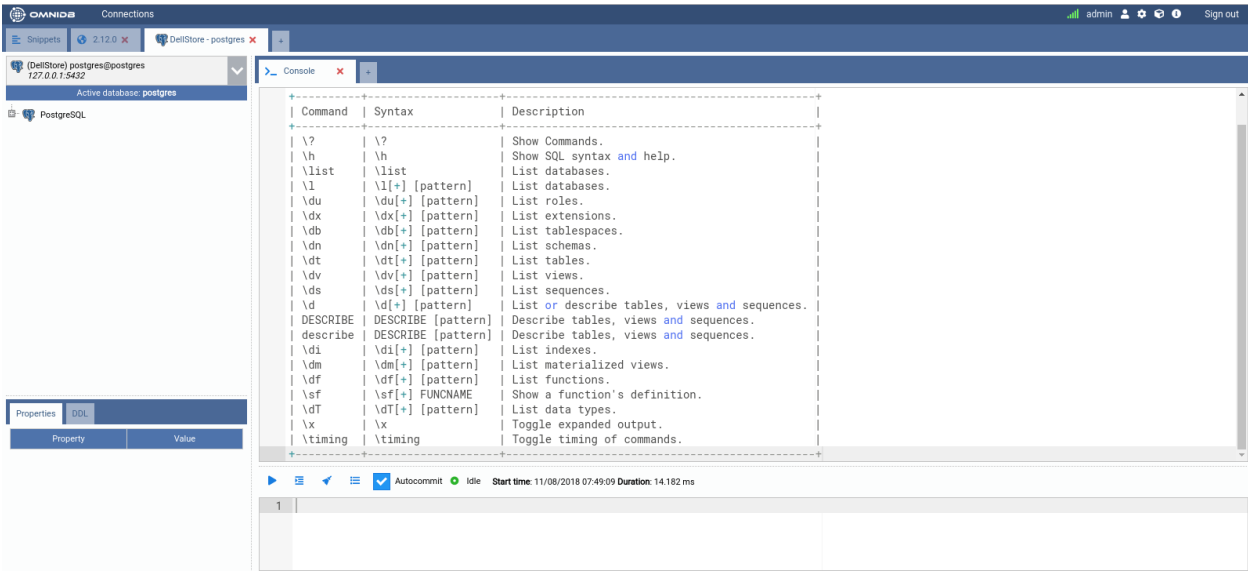
Websocket requests are always directed to the pattern `/wss` so we use a specific location configuration to redirect all requests to the port specified with `-w`, 26500 in this case.

Other requests that are not to `/wss` should all be redirected to OmniDB server, 9000 in this case.

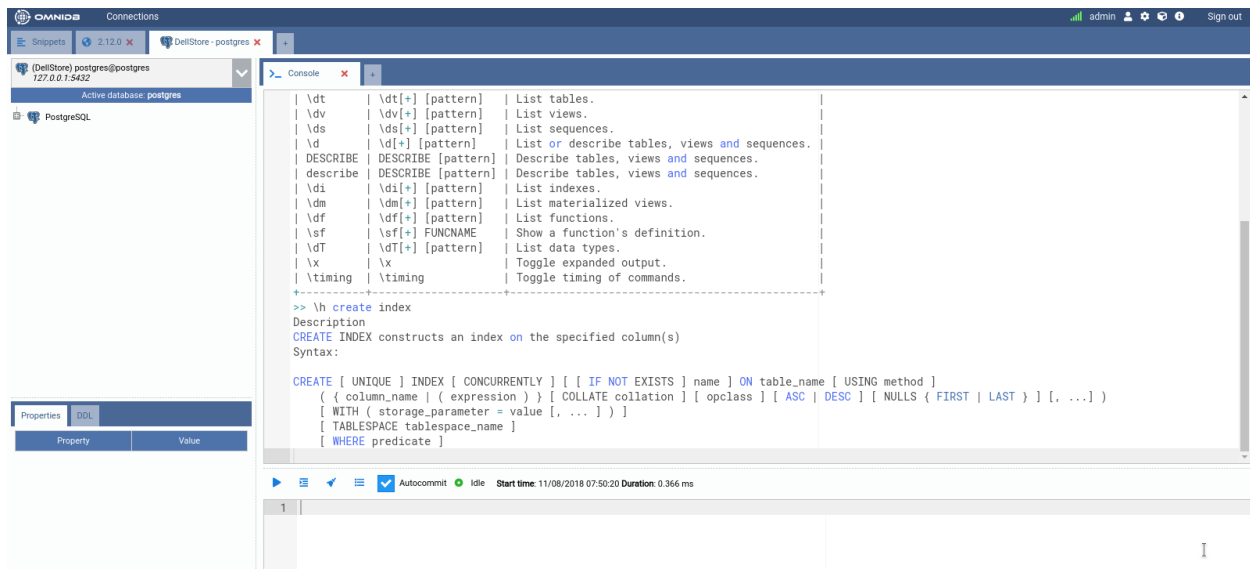
CHAPTER 20

20. Console Tab

Introduced in OmniDB 2.6.0, the new OmniDB Console Tab provides an easy and comfortable way to interact with your databases. Users familiar with the `psql` command line utility will find that Console Tab behaves very similarly. In fact, many of the backslash commands Console Tab provides are present in `psql`.



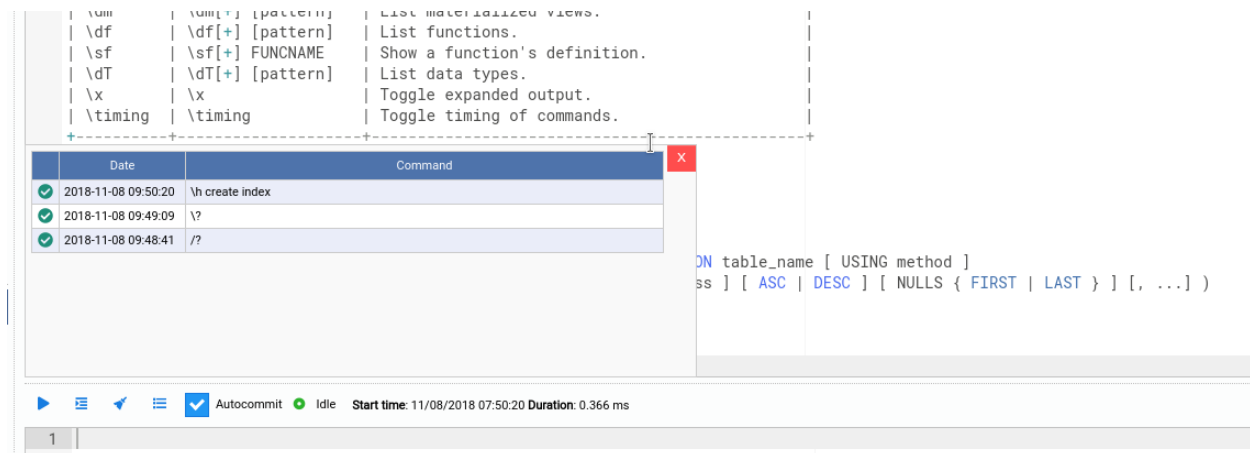
For example, `\?` shows a list with all commands available, its syntax and description. The command `\h` can be another friend of yours, because it shows a comprehensive help about any PostgreSQL SQL command.



The editor on the bottom of the tab area is full-featured just like the *Query Tab* editor (it provides syntax highlight and autocomplete with *Ctrl-Space*). To execute a command, just type it in this editor. If the command is a backslash (starts with \), just type *Enter* and it will be executed. If it is a regular SQL command, then it can be multi-lined, and you will need to type *Alt-Q* to execute it.

All commands and its output will be logged into the display area, which is textual and read-only, so you can copy and paste its contents somewhere else. You can clear the display area by clicking on the *Clear Console* button.

All commands also are logged in the connection query history, and also in a local console history, which you can by clicking in the *Command History* button.



By clicking in the green check, you can borrow the command and put it into the editor, so you can amend it and execute it. Another comfortable way to navigate through the history is using the shortcuts *Ctrl-Up* and *Ctrl-Down*, to quickly paste in the editor the previous and next commands, respectively.

Backslash commands such as `\dt`, `\d+`, `\x` and `\timing` are very useful when dealing with databases every day. The console tab will also show any errors and the results of any SQL command you type in a pretty way. Try it out!

Console tab. Type the commands in the editor below this box. \? to view command list.

```
>> \dt
SELECT 0
>> \d+ categories
```

Column	Type	Modifiers	Storage	Stats target	Description
category	integer	not null default nextval('categories_category_seq'::regclass)	plain	None	None
categoryname	character varying(50)	not null	extended	None	None

Indexes:
 "categories_pkey" PRIMARY KEY, btree (category)
 Has OIDs: no

```
>> \timing
Timing is on.
>> select *
from categories
```

category	categoryname
1	Action
2	Animation
3	Children
4	Classics
5	Comedy

Autocommit Idle Start time: 11/08/2018 07:54:44 Duration: 8.663 ms

```
>> \timing
Timing is on.
>> select *
from categories
```

category	categoryname
1	Action
2	Animation
3	Children
4	Classics
5	Comedy
6	Documentary
7	Drama
8	Family
9	Foreign
10	Games
11	Horror
12	Music
13	New
14	Sci-Fi
15	Sports
16	Travel

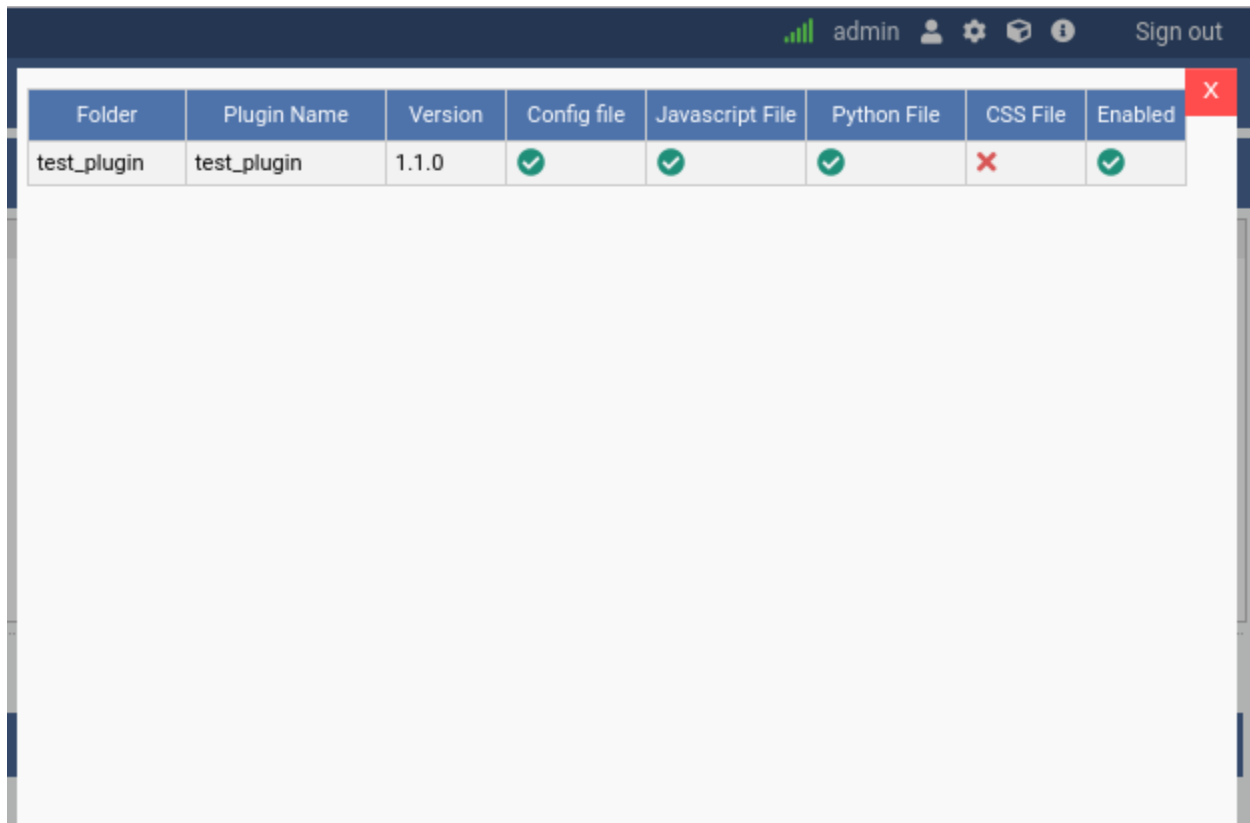
CLOSE CURSOR
 Time: 0:00:00.007430

Autocommit Idle Start time: 11/08/2018 07:54:44 Duration: 8.663 ms

CHAPTER 21

21. Plugin System

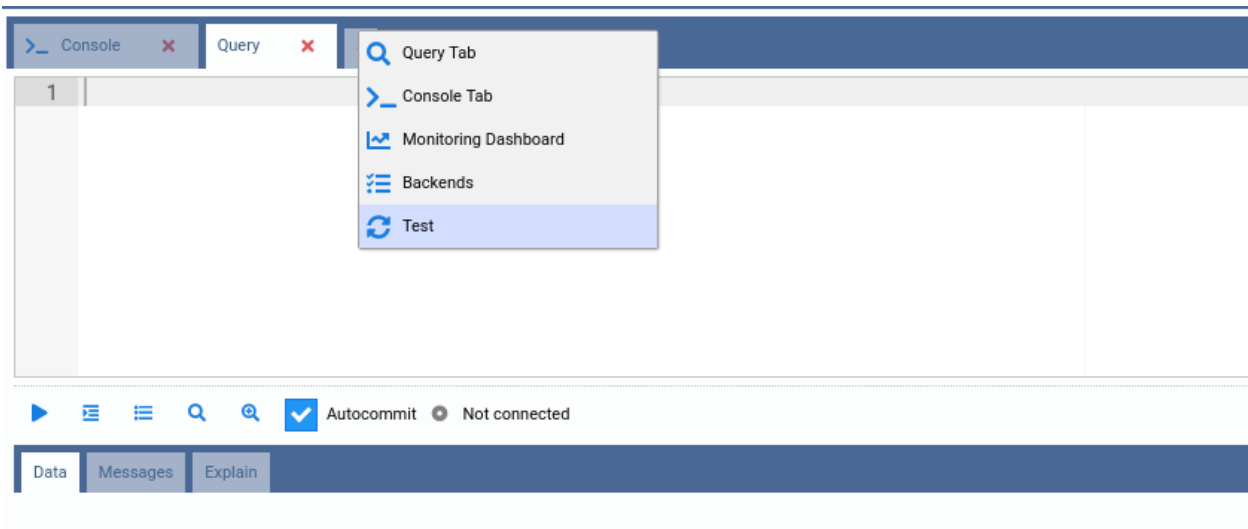
OmniDB 2.9 introduces the plugin system, a feature that allows users to develop and share their own features that can be plugged into OmniDB without having to deploy the whole application again.



Folder	Plugin Name	Version	Config file	Javascript File	Python File	CSS File	Enabled
test_plugin	test_plugin	1.1.0	✓	✓	✓	✗	✓

The plugin system is based on hooks that are located in different parts of the interface. Each plugin can subscribe to any hook and have a collection of API functions to perform different tasks, such as creating inner/outer tabs, creating tree nodes and calling python functions in the plugin's python code.

Here is an example of a plugin that adds the Test action into the inner tab + context menu:



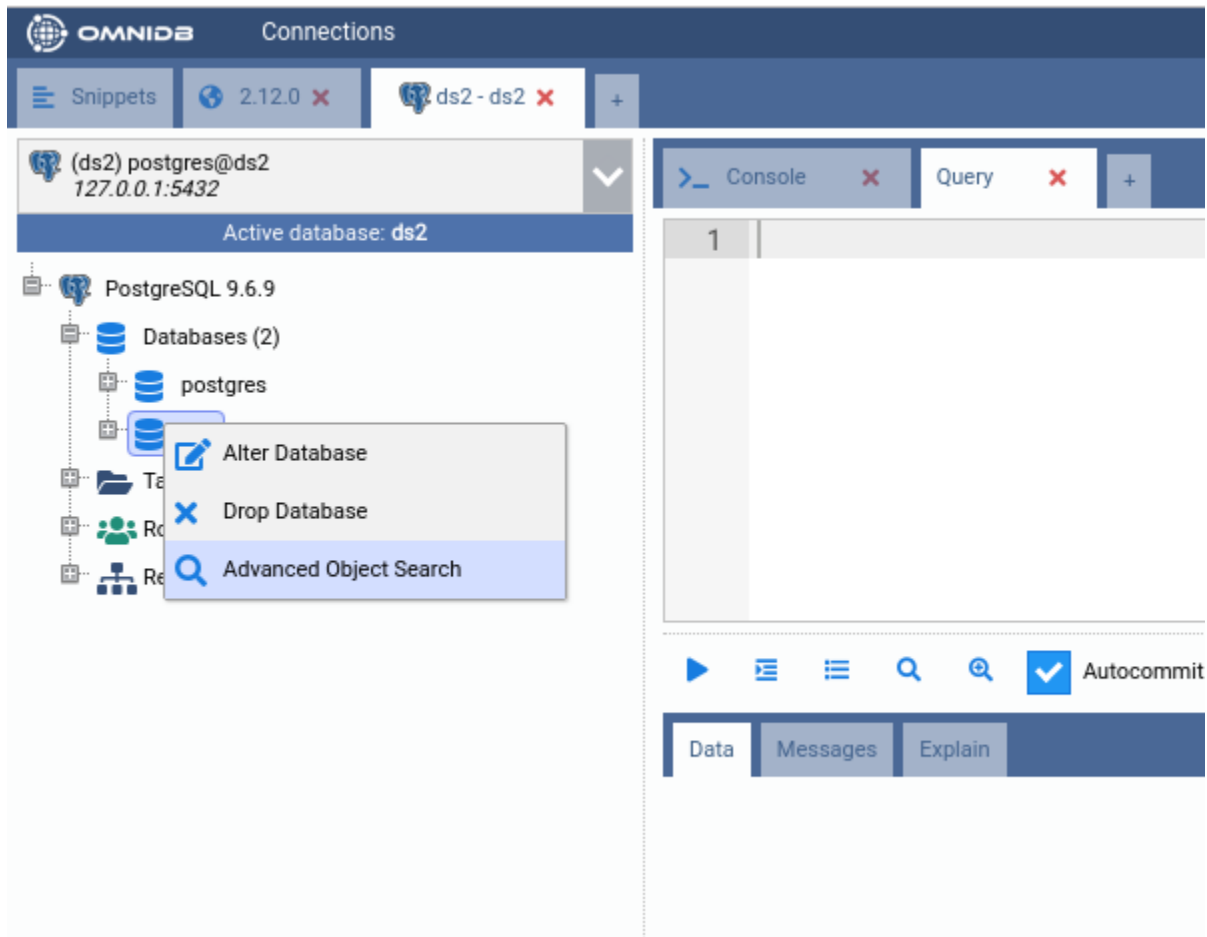
For more details about the Plugin system, instructions on how to install and also to develop plugins, please refer to the [github page](#):

[Plugin System](#)

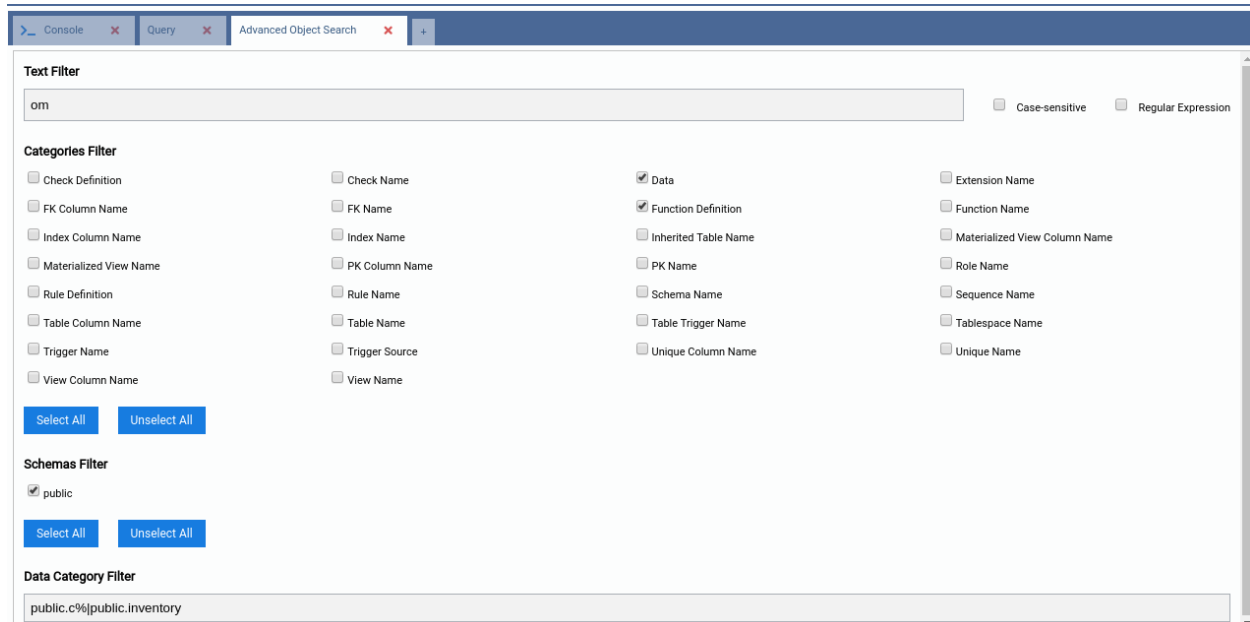
22. Advanced Object Search

OmniDB 2.9 introduces the a Advanced Object Search feature, allowing users to use an advanced pattern matching to search database objects and tables data. The feature allows to use the default SQL LIKE operator and also complex regular expressions.

You can access the Advanced Object Search feature by right clicking in a specific database node in the tree:




The interface allows you to filter categories of objects, schemas where searches will be executed and also to limit the search space when the Data category is selected, so you search for a pattern in a subset of tables:



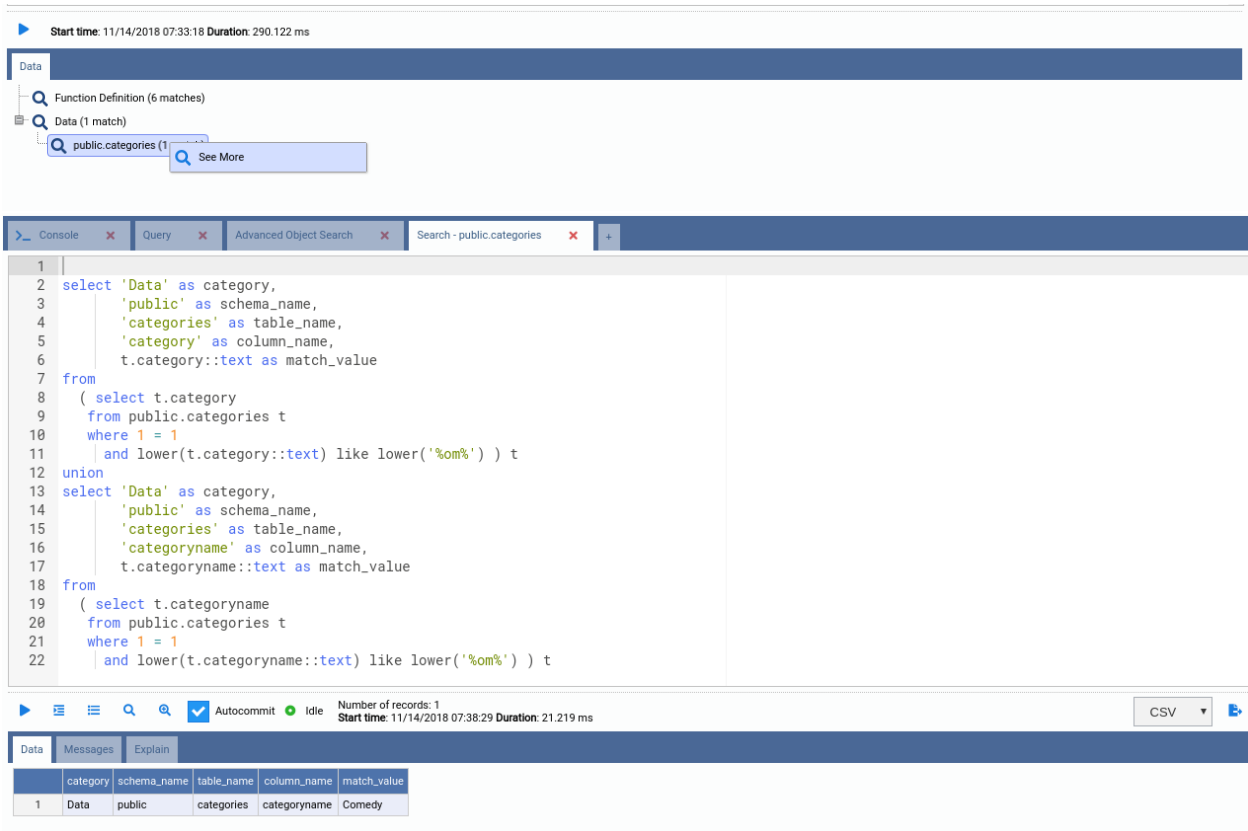
After filling the fields and running OmniDB will perform the search using several threads that will speed up the process

by running in parallel (It is customizable. For more info, see chapter 19 - Deploying OmniDB).

When the search is finished OmniDB will display the result in a tree:



For more details about the search in each category, right click the desired node and select ‘See More’. OmniDB will open a query tab with the SQL command used to perform that specific search. Just run the command to get the results:



23. Debugger Plugin Installation

- *1- Linux Installation*
- *2- Windows Installation*
- *3- FreeBSD Installation*
- *4- MacOSX Installation*
- *5- Post-installation steps* ** **REQUIRED** **

23.1. Linux Installation

You can install from Debian PGDG repository **or** from standalone packages **or** compile from source.

- 1.1. Installing from Debian PGDG repository (**recommended**)
- 1.2. Installing from DEB/RPM packages
- 1.3. Compiling the extension from source

24.1 23.1.1. Installing from Debian PGDG repository

On Debian and Ubuntu systems, this is the recommended way of installing the OmniDB debugger for PostgreSQL PL/pgSQL functions and procedures.

24.1.1 23.1.1.1. Install Debian PGDG repository (if not already)

```
sudo echo "deb http://apt.postgresql.org/pub/repos/apt/ $(lsb_release -cs)-pgdg main" > /etc/apt/sources.list.d/pgdg.list
sudo wget --quiet -O - https://apt.postgresql.org/pub/repos/apt/ACCC4CF8.asc | apt-key add -
```

24.1.2 23.1.1.2. Install omnldb_plugin for your PostgreSQL version X.Y

```
sudo apt install postgresql-X.Y-omnldb
```

24.1.3 23.1.1.3. Set shared_preload_libraries

```
nano /etc/postgresql/X.Y/main/postgresql.conf
    shared_preload_libraries = 'omnidb_plugin'

sudo systemctl restart postgresql
```

24.1.4 23.1.1.3. Post-installation steps

23.1.1.3.1. Create omnidb_plugin extension (should be done by a superuser)

```
psql -d <database> -c 'CREATE EXTENSION omnidb_plugin'
```

23.1.1.3.2. Create sample functions (optional)

```
psql -d <database> -f sample_functions.sql
```

23.1.1.3.3. Next steps

Follow *Post-installation steps* in section 5.

24.2 23.1.2. Installing from DEB/RPM packages

24.2.1 23.1.2.1. Install the package

```
# For example, Debian-like 64 bits:
sudo dpkg -i omnidb-plugin_2.16.0-debian-amd64.deb

# For example, for CentOS-like 64 bits:
sudo rpm -ivU omnidb-plugin_2.16.0-centos-amd64.rpm
```

24.2.2 23.1.2.2. Create a symlink

```
# Find the PostgreSQL version and path for $libdir and create a link to the specific_
↳library. For example:
sudo ln -s /opt/omnidb-plugin/omnidb_plugin_96.so /usr/lib/postgresql/9.6/lib/omnidb_
↳plugin.so
```

24.2.3 23.1.2.3. Set shared_preload_libraries

```
nano /etc/postgresql/X.Y/main/postgresql.conf
    shared_preload_libraries = 'omnidb_plugin'

sudo systemctl restart postgresql
```


24.2.4 23.1.2.4. Post-installation steps

23.1.2.4.1. Create omnidb schema in your database (should be done by a superuser)

```
psql -d <database> -f debugger_schema.sql
```

23.1.2.4.2. Create sample functions (optional)

```
psql -d <database> -f sample_functions.sql
```

23.1.2.4.3. Next steps

Follow *Post-installation steps* in section 5.

24.3 23.1.3. Compiling the extension from source

24.3.1 23.1.3.1. Install headers for PostgreSQL and libpq

```
sudo apt install postgresql-server-dev-X.Y libpq-dev
```

24.3.2 23.1.3.2. Compile omnidb_plugin

```
make
```

24.3.3 23.1.3.3. Install omnidb_plugin

```
sudo make install
```

24.3.4 23.1.3.4. Set shared_preload_libraries

```
nano /etc/postgresql/X.Y/main/postgresql.conf
    shared_preload_libraries = 'omnidb_plugin'

sudo systemctl restart postgresql
```

24.3.5 23.1.3.5. Post-installation steps

23.1.3.5.1. Create omnidb_plugin extension (should be done by a superuser)

```
psql -d <database> -c 'CREATE EXTENSION omnidb_plugin'
```

23.1.3.5.2. Create sample functions (optional)

```
psql -d <database> -f sample_functions.sql
```

23.1.3.5.3. Next steps

Follow *Post-installation steps* in section 5.

23.2. Windows Installation

25.1 23.2.1. Downloading the plugin

Download the zip corresponding to your architecture from the website.

25.2 23.2.2. Installing the plugin

Move the `omnidb_plugin.dll` corresponding to your PostgreSQL version to the folder *lib*, which is inside the folder where PostgreSQL was installed.

25.3 23.2.3. Set `shared_preload_libraries`

Change the file *PostgreSQL_directory/data/postgresql.conf*, including the following line:

```
shared_preload_libraries = 'omnidb_plugin'
```

Then restart PostgreSQL.

25.4 23.2.4. Post-installation steps

25.4.1 23.2.4.1. Create `omnidb` schema in your database (should be done by a superuser)

```
psql -d <database> -f debugger_schema.sql
```

25.4.2 23.2.4.2. Create sample functions (optional)

```
psql -d <database> -f sample_functions.sql
```

25.4.3 23.2.4.3. Next steps

Follow *Post-installation steps* in section 5.

23.3. FreeBSD Installation

26.1 23.3.1. Downloading the plugin

Download the tar.gz corresponding to your architecture from the website.

```
wget --no-check-certificate https://omnidb.org/dist/2.16.0/omnidb-plugin_2.16.0-  
↪freebsd.tar.gz
```

26.2 23.3.1. Installing the plugin

Move the omnidb_plugin.so corresponding to your PostgreSQL version to the folder *lib*, which is inside the folder where PostgreSQL was installed.

```
tar -xzf omnidb-plugin_2.16.0-freebsd.tar.gz  
cp omnidb-plugin_2.16.0-freebsd/omnidb_plugin_10.so /usr/local/lib/postgresql/omnidb_  
↪plugin.so
```

26.3 23.3.3. Set shared_preload_libraries

Change the file *PostgreSQL_directory/data/postgresql.conf*, including the following line:

```
shared_preload_libraries = 'omnidb_plugin'
```

Then restart PostgreSQL.

26.4 23.3.4. Post-installation steps

26.4.1 23.3.4.1. Create omnidb schema in your database (should be done by a superuser)

```
psql -d <database> -f debugger_schema.sql
```

26.4.2 23.3.4.2. Create sample functions (optional)

```
psql -d <database> -f sample_functions.sql
```

26.4.3 23.3.4.3. Next steps

Follow *Post-installation steps* in section 5.

23.4. MacOSX Installation

27.1 23.4.1. Limitations

If you have PostgreSQL installed in your Mac and want to also install OmniDB debugger, please be aware that currently we don't offer any packages for the debugger for Mac OS X. Your only option is to compile and install from sources. It is not that hard, as you can see below.

27.2 23.4.2. Compiling the extension from source

27.2.1 23.4.2.1. Install SDK headers for Mac OS

```
sudo installer -pkg /Library/Developer/CommandLineTools/Packages/macOS_SDK_headers_for_macOS_10.14.pkg -target /
```

27.2.2 23.4.2.2. If not installed, install PostgreSQL from Homebrew

This will also install PostgreSQL headers and libpq.

If brew is not installed yet, you can install it like this:

```
/usr/bin/ruby -e "$ (curl -fsSL https://raw.githubusercontent.com/Homebrew/install/master/install) "
```

Then:

```
brew install postgresql
```

27.2.3 23.4.2.3. Compile omnidb_plugin

```
make
```

27.2.4 23.4.2.4. Install omnidb_plugin

```
sudo make install
```

27.2.5 23.4.2.5. Set shared_preload_libraries

```
vim /usr/local/var/postgres/postgresql.conf  
    shared_preload_libraries = 'omnidb_plugin'  
  
brew services restart postgresql
```

27.2.6 23.4.2.6. Post-installation steps

23.4.2.6.1. Create omnidb_plugin extension (should be done by a superuser)

```
psql -d <database> -c 'CREATE EXTENSION omnidb_plugin'
```

23.4.2.6.2. Create sample functions (optional)

```
psql -d <database> -f sample_functions.sql
```

23.4.2.6.3. Next steps

Follow *Post-installation steps* in section 5.

23.5. Post-installation steps **** REQUIRED ****

28.1 23.5.1. Grant privileges to each database user that will debug functions (should be done by a superuser)

Every database user that uses the debugger needs access to the debugger control tables.

```
psql -d <database> -c 'GRANT ALL ON SCHEMA omnidb TO <user>; GRANT ALL ON ALL TABLES_
↪IN SCHEMA omnidb TO <user>;'
```

28.2 23.5.2. Enable passwordless access to each database user that will debug functions

Every database user that uses the debugger needs local passwordless access to the target database. This is because the database will create an additional local connection to perform debugging operations.

We need to add a rule to *pg_hba.conf* of type `host`, matching the PostgreSQL user and database OmniDB is connected to. The method can be either `trust`, which is insecure and not recommended, or `md5`.

28.2.1 trust

- Add a rule similar to:

#	TYPE	DATABASE	USER	ADDRESS	METHOD
host		<database>	<user>	127.0.0.1/32	trust
host		<database>	<user>	:::1/128	trust

28.2.2 md5

- Add rules similar to:

#	TYPE	DATABASE	USER	ADDRESS	METHOD
host		<database>	<user>	127.0.0.1/32	md5
host		<database>	<user>	:::1/128	md5

- Create a `.pgpass` file with a similar content:

```
localhost:<port>:<database>:<username>:<password>
```

More information about how `.pgpass` works can be found here: <https://www.postgresql.org/docs/11/static/libpq-pgpass.html>

CHAPTER 29

Indices and tables

- `genindex`
- `modindex`
- `search`