

Chapter 01: First Steps

General	Connection
Host name/address	<input type="text"/>
Port	5432
Maintenance database	postgres
User name	postgres
Password	<input type="password"/>
Save password?	<input type="checkbox"/>
Role	<input type="text"/>
SSL mode	Prefer

pgAdmin 4 File Object Tools Help

Browser

- Servers (5)
 - articmonkeys
 - Databases (1)
 - postgres
 - Casts
 - Catalogs
 - Event Triggers
 - Extensions
 - Foreign Data Wrappers
 - Languages
 - Schemas (2)
 - example
 - public
 - Collations
 - Foreign Tables
 - Functions
 - Materialized Views
 - Sequences
 - Tables (4)
 - customer
 - foo
 - foo2

Dashboard Properties SQL Statistics

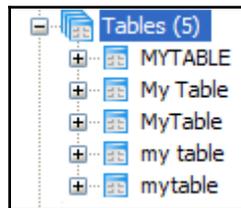
Statistics	
Backends	5
Xact committed	10382
Xact rolled back	18
Blocks read	452
Blocks hit	46632
Tuples returned	246096
Tuples fetched	24778
Tuples inserted	310
Tuples updated	6
Tuples deleted	172
Last statistics reset	2017-01-17 12:35:41.5
Tablespace conflicts	0
Lock conflicts	0
Snapshot conflicts	0
Bufferpin conflicts	0
Deadlock conflicts	0
Temporary files	0

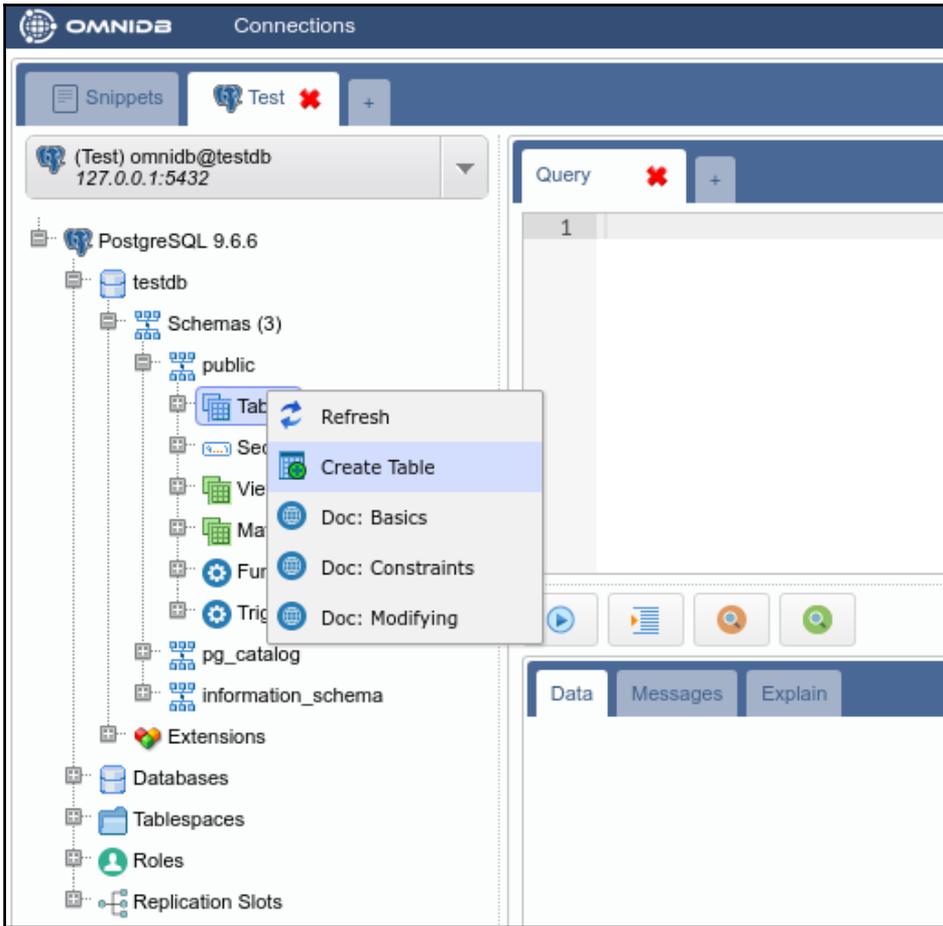
postgres on sriggs@articmonkeys

```
1 SELECT count(*)
2 FROM demo d1 JOIN demo d2 ON d1.col1 = d2.col1
3 WHERE d1.col1 = 'DEMO';
```

Data Output Explain Messages History

The diagram illustrates the execution plan for the provided SQL query. It starts with two 'demo' tables. The first 'demo' table is processed, and its output is materialized. This materialized result is then joined with the second 'demo' table using a 'Nested Loop Inner Join' operation. Finally, the result of the join is passed to an 'Aggregate' operation, represented by a sigma symbol (Σ).





Query ✖ +

```

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

```

▶ 📄 🔍 🔄 **Start time: 12/05/2017 10:03:26 Duration: 15.073 ms**

Data Messages Explain

☰ ⬆ ⬇ ⬆

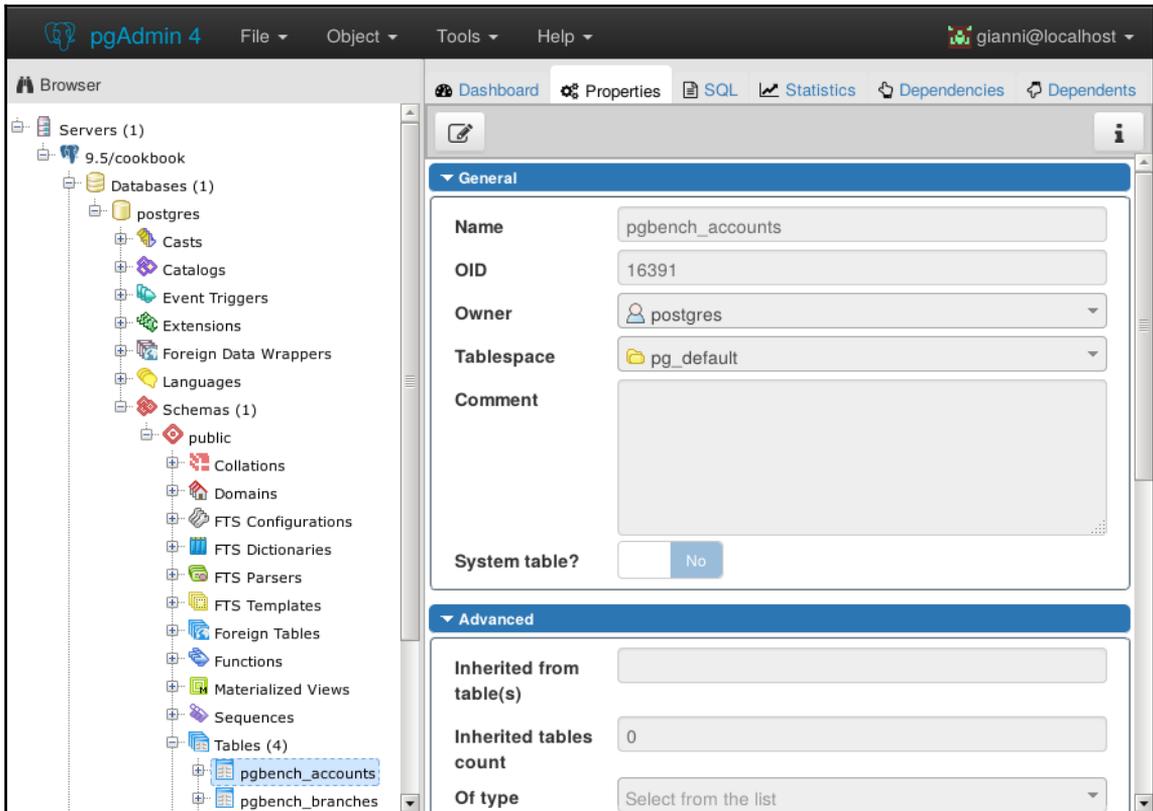
QUERY PLAN

```

#
1  Seq Scan on customer cust (cost=0.00..38.27 rows=2 width=226) (actual time=0.029..0.029 rows=1 loops=1)
2    Filter: ((cust_name)::text = 'Rafael'::text)
3    Rows Removed by Filter: 8
4  SubPlan 1
5  ⦿ Aggregate (cost=12.13..12.14 rows=1 width=8) (actual time=0.017..0.017 rows=1 loops=1)
6  ⦿ Seq Scan on address addr (cost=0.00..12.12 rows=1 width=0) (actual time=0.013..0.014 rows=1 loops=1)
7    Filter: (cust_id = cust.cust_id)
8    Rows Removed by Filter: 8
9  Planning time: 0.323 ms
10 Execution time: 0.097 ms

```

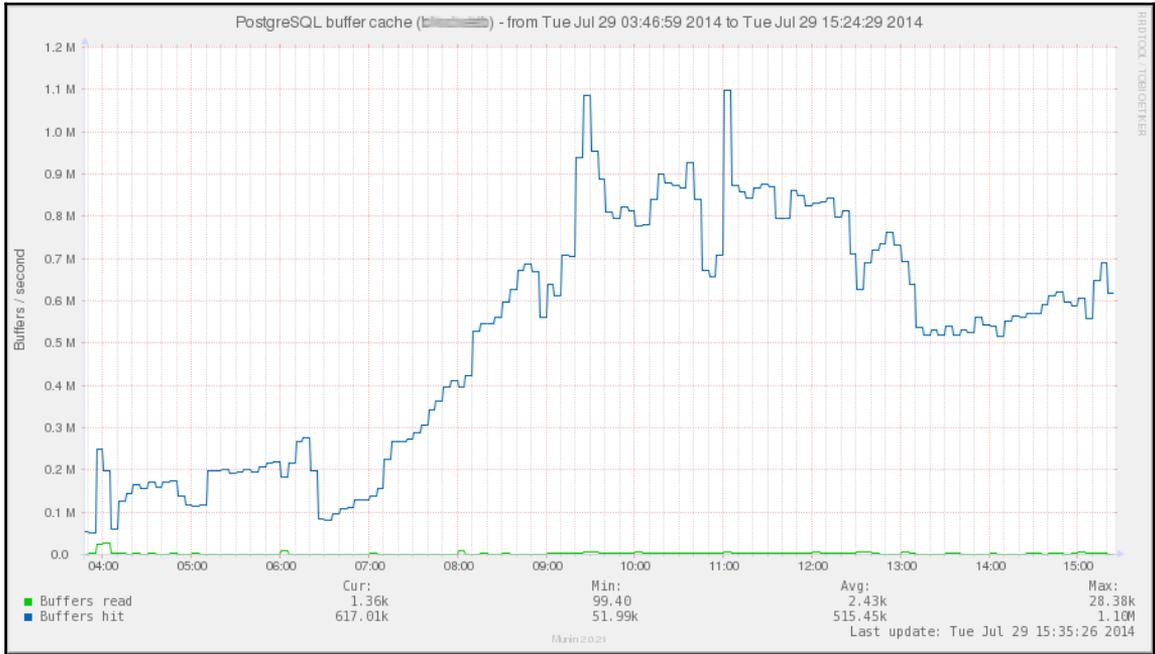
Chapter 02: Exploring the Database



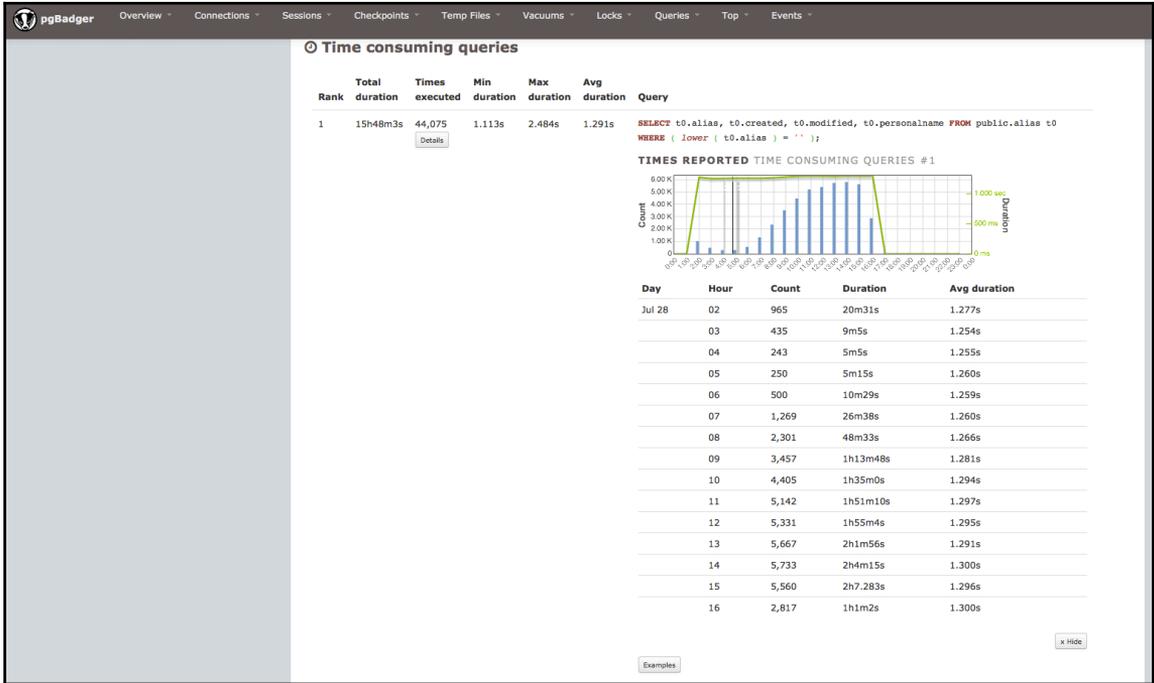
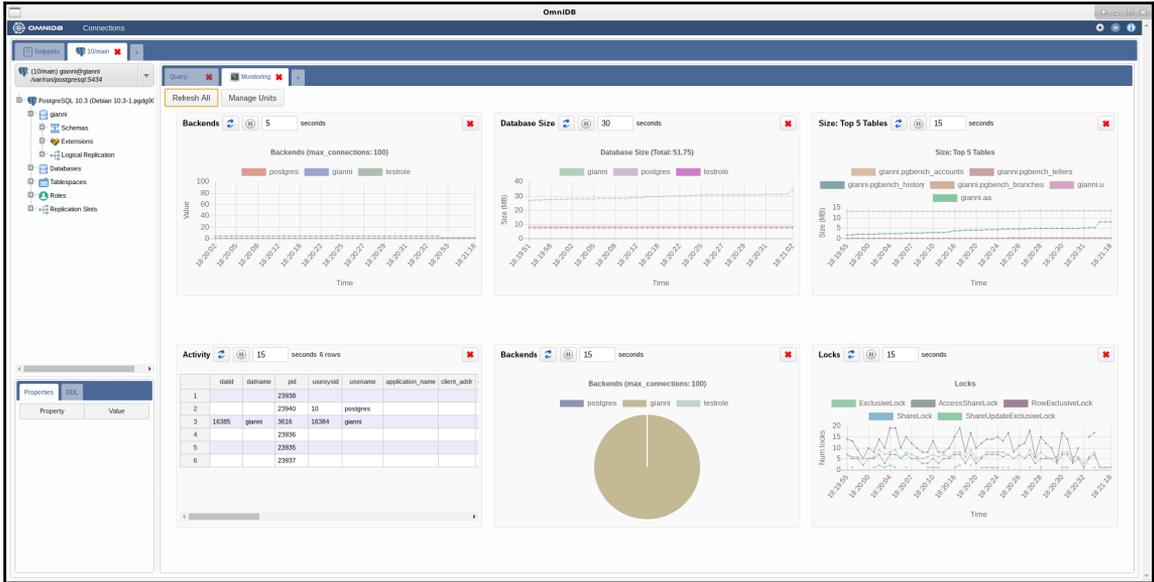
Chapter 05: Tables and Data

	A	B	C	D
1	Key	Value		
2		1c		
3		2d		
4				
5				
6				

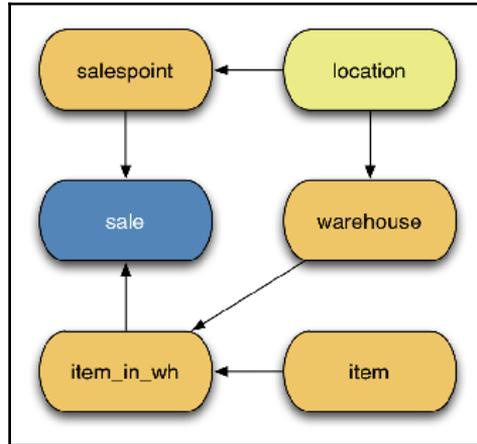
Chapter 08: Monitoring and Diagnosis



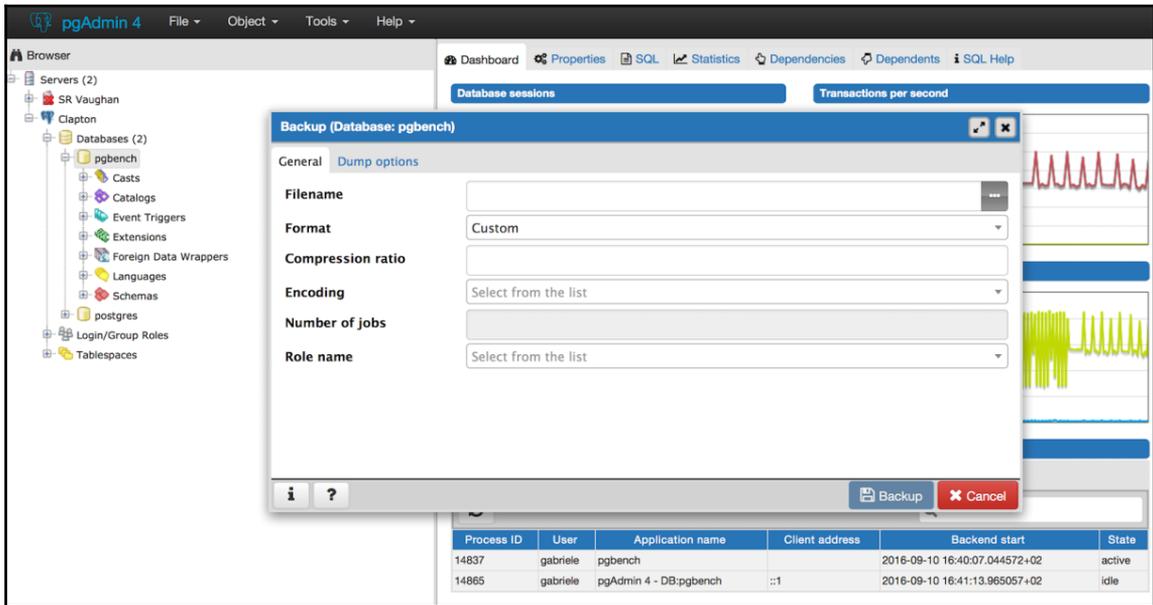




Chapter 10: Performance and Concurrency



Chapter 11: Backup and Recovery



The screenshot shows the pgAdmin 4 interface. On the left is a tree view of the database structure. The main area displays the 'Database sessions' view for the 'pgbench' database. A 'Backup (Database: pgbench)' dialog box is open, showing the 'General' tab with fields for 'Filename', 'Format' (set to 'Custom'), 'Compression ratio', 'Encoding', 'Number of jobs', and 'Role name'. Below the dialog, a table shows the current database sessions:

Process ID	User	Application name	Client address	Backend start	State
14837	gabriele	pgbench		2016-09-10 16:40:07.044572+02	active
14865	gabriele	pgAdmin 4 - DB:pgbench	:::1	2016-09-10 16:41:13.965057+02	idle