Hands on Lab
Part 2 – PgPool-II Cluster

DBTech's Virtual Workshop on Distributed, Replicated, Embedded and Mobile Databases (Deliverable #28).

Antonio César Gómez Lora

Dpto. Lenguajes y Ciencias de la Computación

Universidad de Málaga

Introduction

We are going to create a PgPool-II Cluster using three machines. The first one (192.168.192.1) will contain PgPool-II. The second and third machines (192.168.192.2 and 3) will contain the PostgreSQL replicas.

We start now using the previously created and cloned virtual machines. Those machines have Debian (5.0.5/5.0.6), MySQL, PostgreSQL and PgPool-II installed.

Creating the cluster

Start the three machines and login as root on third machine.

We must edit postgresql.conf (Fig. 1).

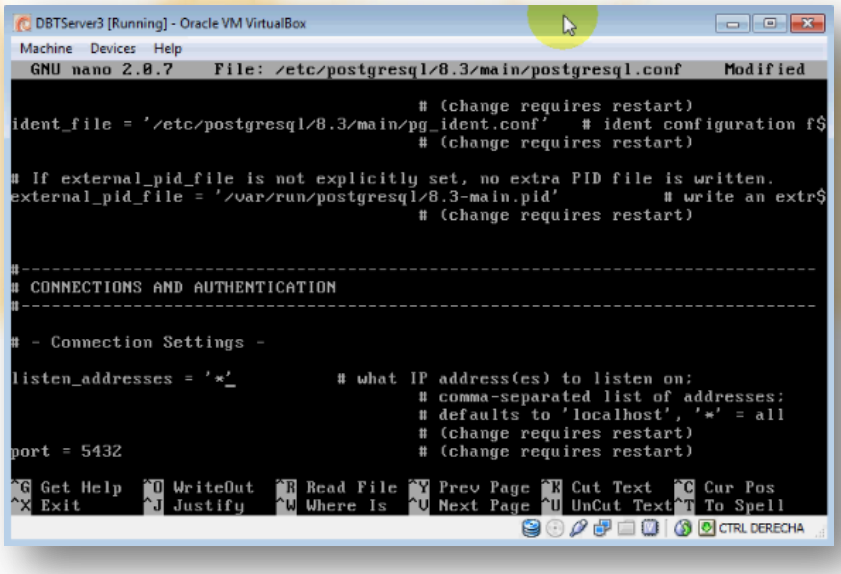

Fig. 1. postgresql.conf

Version to be presented on
DBTech's Helsinki Meeting, Nov'10
Please feedback any comment to: cesar@lcc.uma.es

So execute:

nano /etc/postgresql/8.3/main/postgresql.conf

uncomment and change listen_addresses to:

listen_addresses = '*'

Now edit pg_hba.conf (Fig. 2), execute:

nano /etc/postgresql/8.3/main/pg_hba.conf

Add the folowing line at the end of the file:

host all postgres 192.168.192.0/24 trust

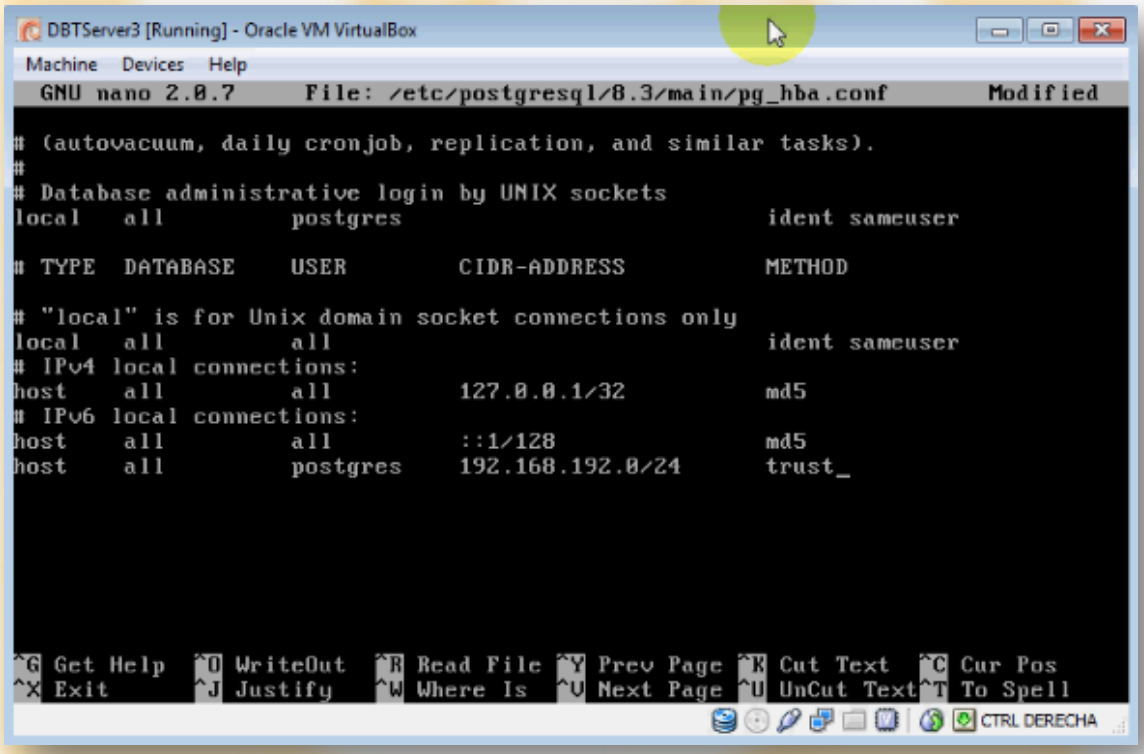

Fig. 2. pg_hba.conf

Repeat the same process on the second machine, editing postgresql.conf and pg_hba.conf.

On first machine edit pgpool.conf (Fig. 3), execute:

nano /etc/pgpool.conf

Change listen_addresses, replication_mode and load_balance_mode to:

listen_adresses = '*'

replication_mode = true

load_balance_mode = true

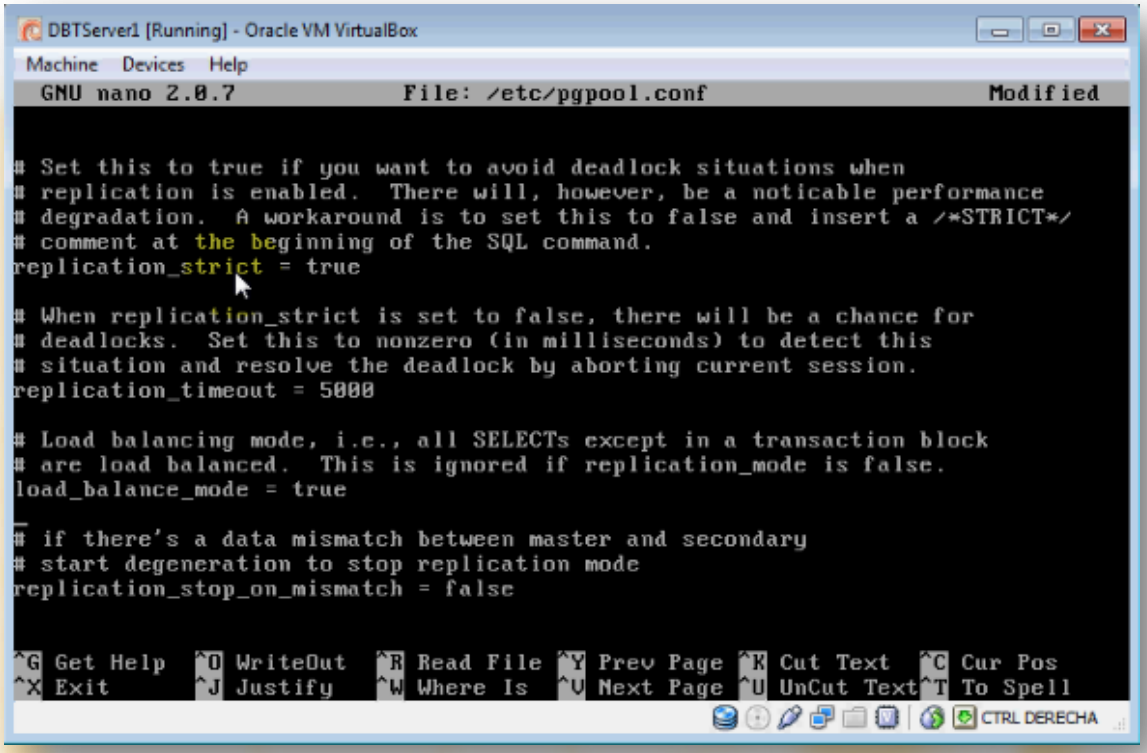

Fig. 3. pgpool.conf

Also include near the end of the file (Fig. 4):

backend_hostname0 = '192.168.192.2'

backend_port0 = 5432

backend_weight0 = 1

backend_hostname1 = '192.168.192.3'

backend_port1 = 5432

backend_weight1 = 1

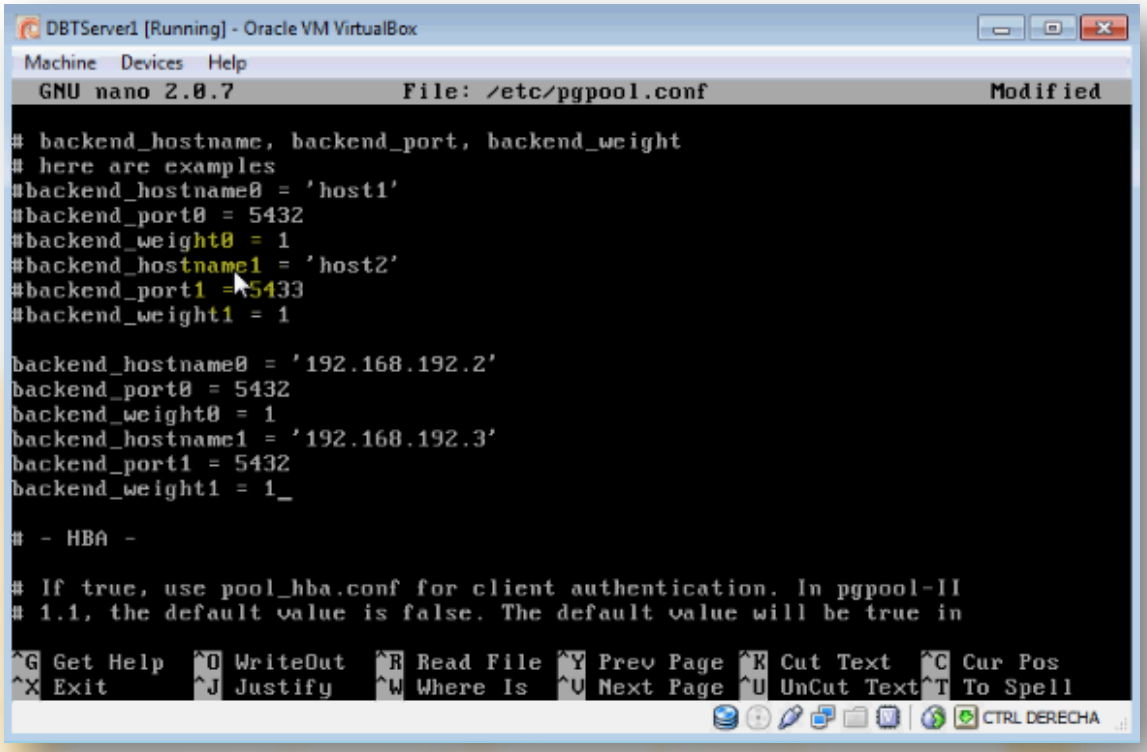

Fig. 4. backends on pgpool.conf

Now its time to restart postgresql service on all machines (Fig. 5), execute on each machine:

/etc/init.d/postgresql-8.3 restart

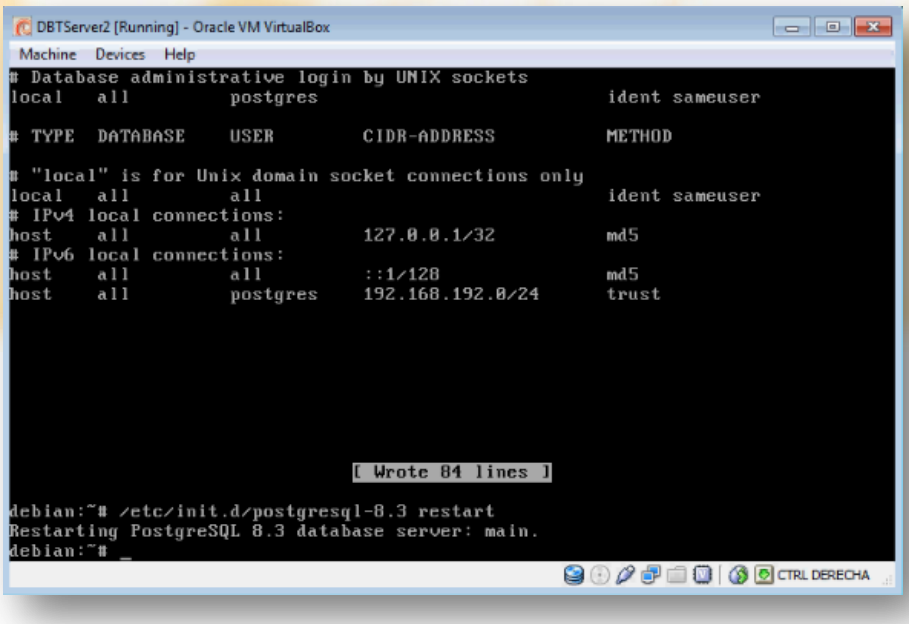

Fig. 5. Restarting postgreSQL

Version to be presented on
DBTech's Helsinki Meeting, Nov'10
Please feedback any comment to: cesar@lcc.uma.es

On first machine also restart pgpool service (Fig. 6), execute:

/etc/init.d/pgpool2 restart

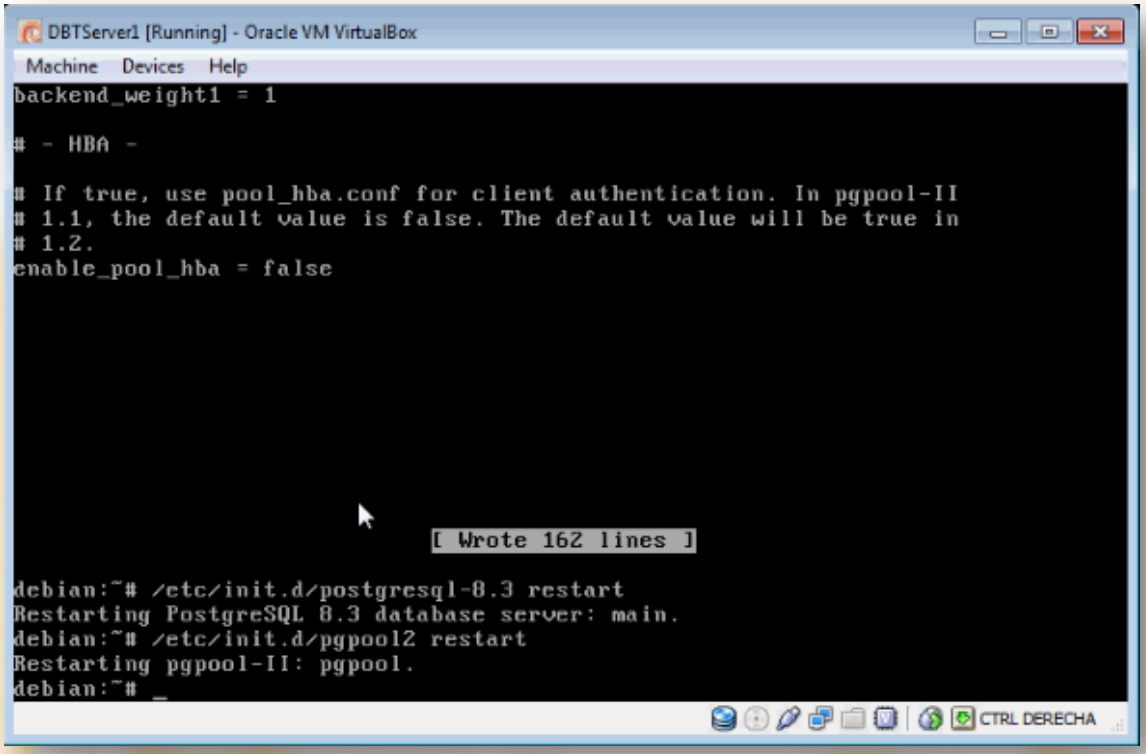

Fig. 6. Restarting pgpool-II

Testing the cluster

Verify that PostgreSQL is clean (Fig. 7), execute on second and third machine:

su postgres

psql -l

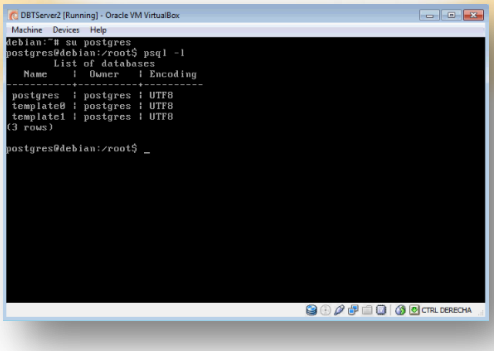

Fig. 7. Clean postgreSQL

Version to be presented on
DBTech's Helsinki Meeting, Nov'10
Please feedback any comment to: cesar@lcc.uma.es

On first machine execute:

su postgres

psql -l -p 5433

create a database using pgpool2 (Fig. 8),

execute:

createdb -p 5433 pgptest

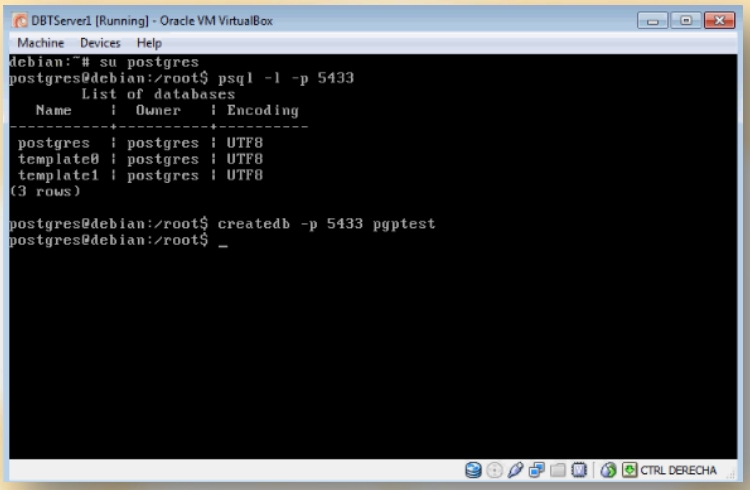

Fig. 8. Create pgptest

Let's see if pgptest exists on storage nodes (Fig. 9), second and third machines, on them execute:

psql -l

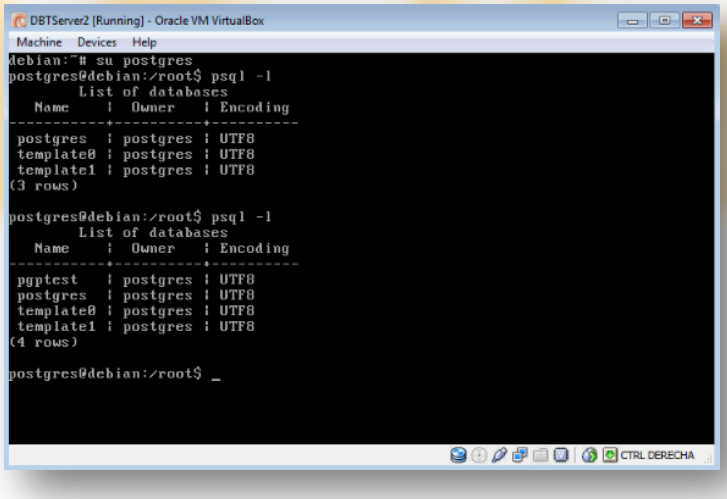

Fig. 9. Database replicated

Now let's monitorize second and third machine activity using htop (Fig. 10) or top.

If htop is not installed you can install it by executing

```
sudo apt-get install htop
```

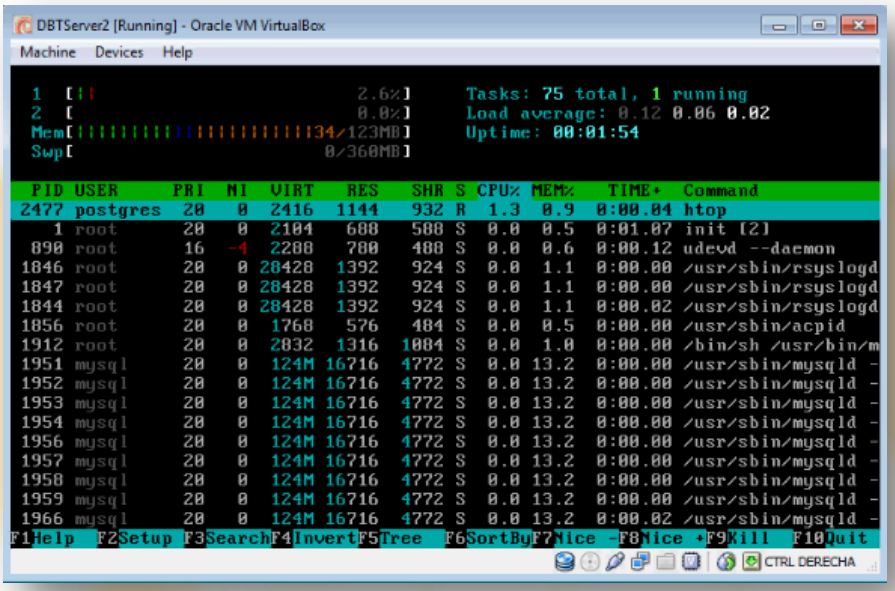

Fig. 10. Htop

Let's start psql interpreter connecting to pgpool on first machine, execute:

```
psql -p 5433 pgptest
```

Now let's create a simple table, insert a value and duplicate the table contents, so we get a little big table (Fig. 11).

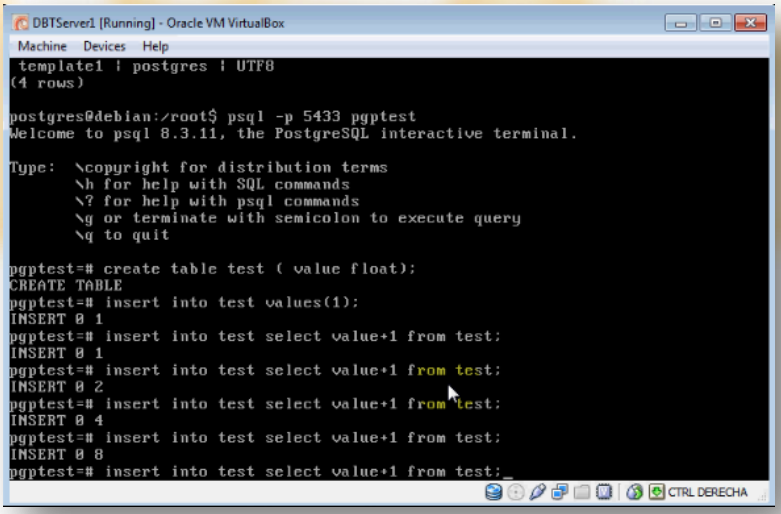

Fig. 11. Test table

Now we create two "heavy" views (Fig. 12), with no meaning.

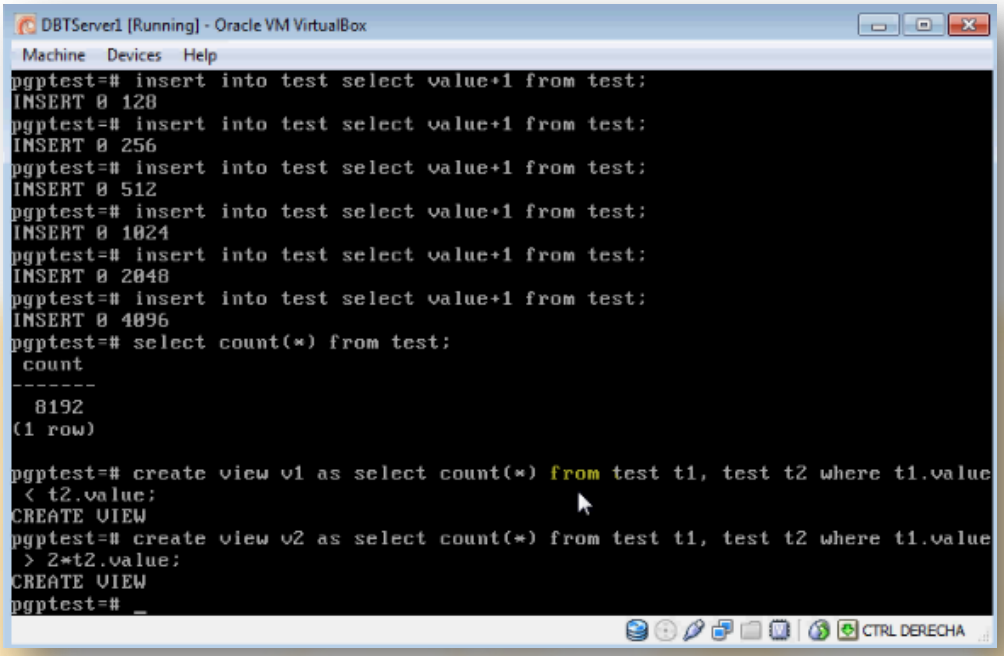

Fig. 12. Little heavy views

Now open a second terminal session on first machine (by pressing ALT+F2), login, change to user postgres and start the psql interpreter (Fig. 13), execute:

su postgres

psql -p 5433 pgptest

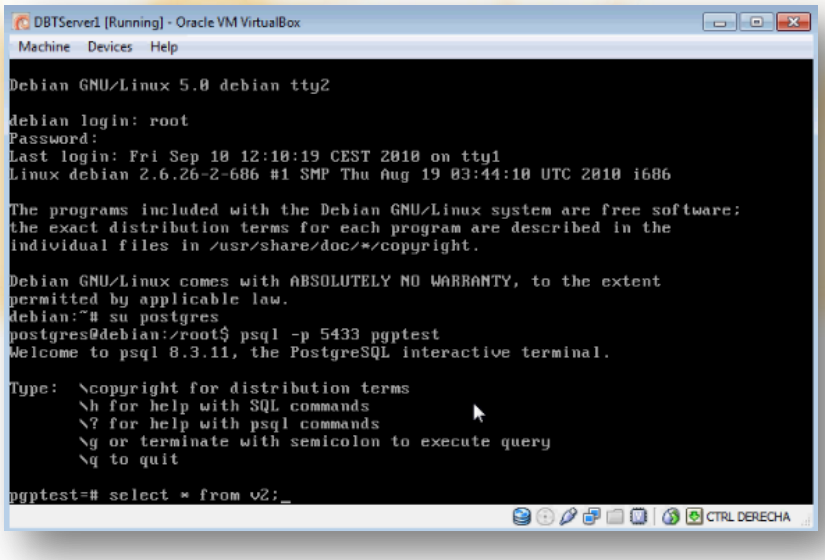

Fig. 13. Second session

Version to be presented on
DBTech's Helsinki Meeting, Nov'10
Please feedback any comment to: cesar@lcc.uma.es

Using ALT+F1 and ALT+F2 alternate between sessions (Fig. 14), and prepare but do not execute on first session:

```
select * from v1;
```

And on second one:

```
select * from v2;
```

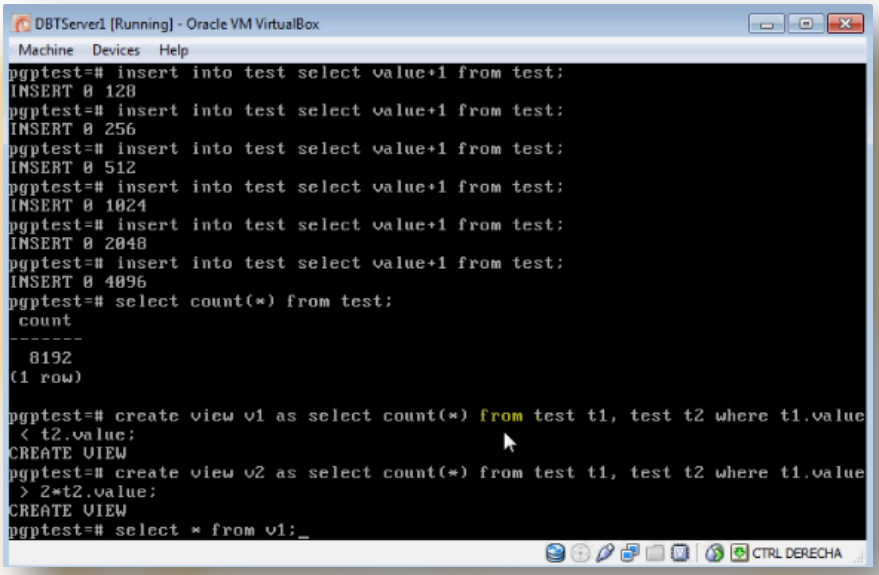

Fig. 14. First session

When ready press ENTER and quickly alternate to other session and press ENTER again. You will see activity on both machines (Fig. 15). In this case second session's query was launched first. Pgpool sends the query to the first free database (the one on the left), and when first session launch its query Pgpool sends it to the second datatabase (the one on the right), beacuse the first one was working. If you launch first a query and when it finishs the second one you will see that they are executed by the first database.

Fig. 15. Processing

Version to be presented on
DBTech's Helsinki Meeting, Nov'10
Please feedback any comment to: cesar@lcc.uma.es

Now let's see that both databases contains the same table created through pgpool2 (Fig. 16).

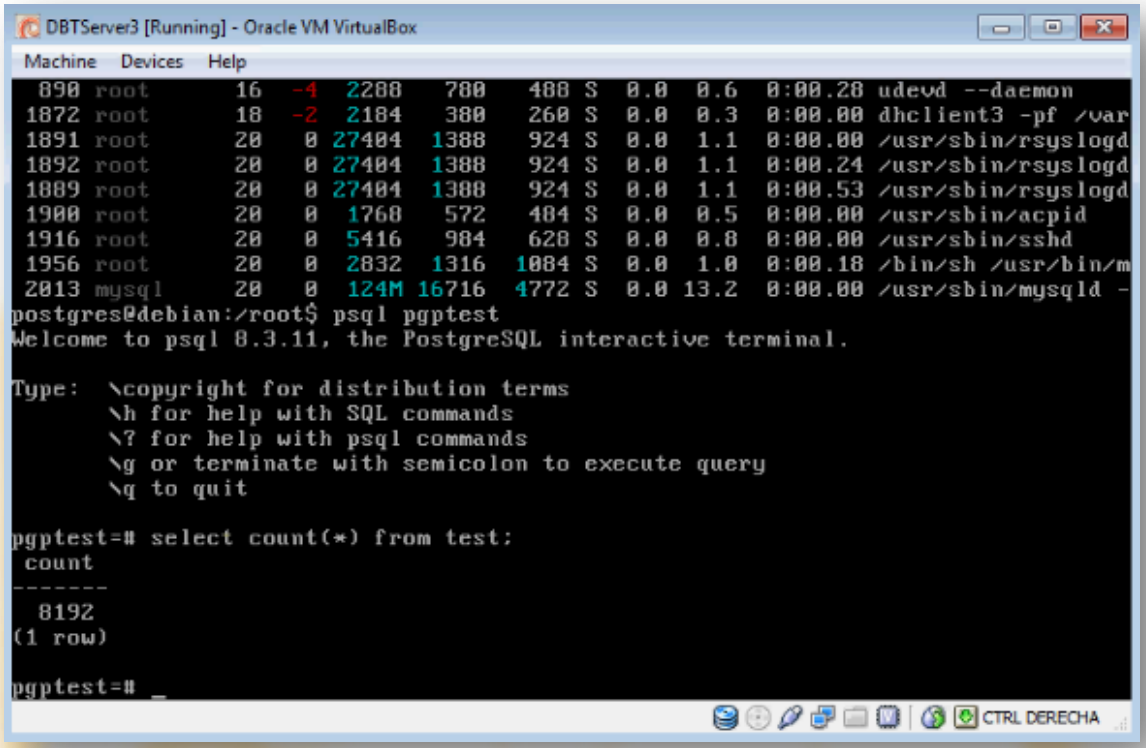

Fig. 16. Replicated

Proposed exercises

Repeat the proposed exercises 1 and 2 for MySQL Cluster, using now PgPool-II.

1. If you click on VirtualBox network icon (on the bottom of each virtual machine status bar) you can view network interfaces and disconnect and reconnect whenever you want. Disconnect a virtual machine and try to determine what happens with replicated data if both replicas continue modifying the replicated data. First disconnect cluster manager, and after testing it, disconnect the first storage node and finally the second storage node. If cluster lost its integrity regenerate the replicated table. If a data node cannot reconnect then restart the services properly (as it has been done in this document).
2. Start transactions on storage nodes and try to analyze what kind of transactional model it follows.

And now:

3. Consulting the PgPool manuals try to configure PgPool to perform parallel queries.

Figure Index

Fig. 1. postgresql.conf 1
Fig. 2. pg_hba.conf 2
Fig. 3. pgpool.conf 3
Fig. 4. backends on pgpool.conf 4
Fig. 5. Restarting postgreSQL 4
Fig. 6. Restarting pgpool-II 5
Fig. 7. Clean postgreSQL 5
Fig. 8. Create pgptest 6
Fig. 9. Database replicated 6
Fig. 10. Htop 7
Fig. 11. Test table 7
Fig. 12. Little heavy views 8
Fig. 13. Second session 8
Fig. 14. First session 9
Fig. 15. Processing 9
Fig. 16. Replicated 10