

Migrating on new servers using repmgr with near zero downtime.

Raoul De Guchteneere

PGConf.EU October 25 2024

University of Louvain – Louvain-la-Neuve – Belgium

- Large university in french-speaking Belgium
- 39,000 students
- 7,000 faculty and staff
- 7 campuses
- 43 bachelor's, 187 master's programmes

Raoul De Guchteneere

- Head of infrastructures (servers and network)
- Former DBA (and still acting)

- Gradually using PostgreSQL since 9.3/2013
- Small projects
- Open source software (OpenNebula, Bacula, Moodle, Drupal, RERO ILS, etc ...)
- Main management tool for the student's academic record (from admission to diploma: entrance exams, courses, schedule, grades, ...). Developed in-house in Python.

- Small 1U DB servers with internal storage
- Debian
- Used in pairs (in two datacenters) with streaming replication
- Haproxy for load balacing
- repmgr

```
postgres@pgsql3:~$ repmgr --help
repmgr: replication management tool for PostgreSQL
```

Usage:

```
repmgr [OPTIONS] primary {register|unregister}
repmgr [OPTIONS] standby {register|unregister|clone|promote|follow|switchover}
repmgr [OPTIONS] node {status|check|rejoin|service}
repmgr [OPTIONS] cluster {showevent|matrix|crosscheck|cleanup}
repmgr [OPTIONS] witness {register|unregister}
repmgr [OPTIONS] service {status|pause|unpause}
repmgr [OPTIONS] daemon {start|stop}
```

Execute "repmgr {primary|standby|node|cluster|witness|service} --help" to see command-specific options

```
postgres@pgsql3:~$ repmgr -f /etc/postgresql/16/main/repmgr.conf cluster show
ID | Name   | Role   | Status   | Upstream | Location | Priority | Timeline |
Connection string
-----+-----+-----+-----+-----+-----+-----+-----+
1 | pgsql3 | primary | * running |          | default  | 100     | 1       |
host=pgsql3 user=repmgr dbname=repmgr port=5465
2 | pgsql4 | standby | running | pgsql3 | default  | 100     | 1       |
host=pgsql4 user=repmgr dbname=repmgr port=5465
```




```
postgres@pgsql5:~$ repmgr -h pgsql3 -p 5465 -U repmgr -d repmgr -f
/etc/postgresql/16/main/repmgr.conf standby clone
NOTICE: destination directory "/var/lib/postgresql/16/main" provided
INFO: connecting to source node
DETAIL: connection string is: host=pgsql3 port=5465 user=repmgr dbname=repmgr
DETAIL: current installation size is 15 GB
NOTICE: checking for available walsenders on the source node (2 required)
NOTICE: checking replication connections can be made to the source server (2 required)
WARNING: data checksums are not enabled and "wal_log_hints" is "off"
DETAIL: pg_rewind requires "wal_log_hints" to be enabled
INFO: checking and correcting permissions on existing directory
"/var/lib/postgresql/16/main"
NOTICE: starting backup (using pg_basebackup)...
HINT: this may take some time; consider using the -c/--fast-checkpoint option
INFO: executing:
  /usr/lib/postgresql/16/bin/pg_basebackup -l "repmgr base backup" -D
/var/lib/postgresql/16/main -h pgsql3 -p 5465 -U repmgr -X stream -S repmgr_slot_4
NOTICE: standby clone (using pg_basebackup) complete
NOTICE: you can now start your PostgreSQL server
HINT: for example: pg_ctl -D /var/lib/postgresql/16/main start
HINT: after starting the server, you need to register this standby with "repmgr standby
register"
```

```
postgres@pgsql5:~$ repmgr -f /etc/postgresql/16/main/repmgr.conf standby register
INFO: connecting to local node "pgsql5" (ID: 3)
INFO: connecting to primary database
WARNING: --upstream-node-id not supplied, assuming upstream node is primary (node ID: 1)
INFO: standby registration complete
NOTICE: standby node "pgsql5" (ID: 3) successfully registered
```




```
postgres@pgsql3:~$ repmgr -f /etc/postgresql/16/main/repmgr.conf cluster show
ID | Name   | Role   | Status   | Upstream | Location | Priority | Timeline |
Connection string
-----+-----+-----+-----+-----+-----+-----+-----+
1 | pgsql3 | primary | * running |          | default  | 100     | 1       |
host=pgsql3 user=repmgr dbname=repmgr port=5465
2 | pgsql4 | standby | running | pgsql3 | default  | 100     | 1       |
host=pgsql4 user=repmgr dbname=repmgr port=5465
3 | pgsql5 | standby | running | pgsql3 | default  | 100     | 1       |
host=pgsql5 user=repmgr dbname=repmgr port=5465
4 | pgsql6 | standby | running | pgsql3 | default  | 100     | 1       |
host=pgsql5 user=repmgr dbname=repmgr port=5465
```



```
postgres@pgsql5:~$ repmgr -f /etc/postgresql/16/main/repmgr.conf standby switchover \  
--siblings-follow  
NOTICE: executing switchover on node "pgsql5" (ID: 3)  
NOTICE: attempting to pause repmgrd on 4 nodes  
NOTICE: local node "pgsql5" (ID: 3) will be promoted to primary; current primary  
"pgsql3"  
(ID: 1) will be demoted to standby  
NOTICE: stopping current primary node "pgsql3" (ID: 1)  
NOTICE: issuing CHECKPOINT on node "pgsql3" (ID: 1)  
DETAIL: executing server command "/usr/lib/postgresql/16/bin/pg_ctl  
-D '/var/lib/postgresql/16/main' -W -m fast stop"  
NOTICE: STANDBY PROMOTE successful  
DETAIL: server "pgsql5" (ID: 3) was successfully promoted to primary  
NOTICE: node "pgsql5" (ID: 3) promoted to primary, node "pgsql3" (ID: 1) demoted to  
standby  
NOTICE: replication slot "repmgr_slot_3" deleted on node 1  
NOTICE: switchover was successful  
DETAIL: node "pgsql5" is now primary and node "pgsql3" is attached as standby
```

```
postgres@pgsql3:~$ repmgr -f /etc/postgresql/16/main/repmgr.conf cluster show
ID | Name   | Role   | Status   | Upstream | Location | Priority | Timeline |
Connection string
-----+-----+-----+-----+-----+-----+-----+-----+
1 | pgsql3 | standby | running | pgsql5   | default  | 100     | 2        |
host=pgsql3 user=repmgr dbname=repmgr port=5465
2 | pgsql4 | standby | running | pgsql5   | default  | 100     | 2        |
host=pgsql4 user=repmgr dbname=repmgr port=5465
3 | pgsql5 | primary | * running |          | default  | 100     | 2        |
host=pgsql5 user=repmgr dbname=repmgr port=5465
4 | pgsql6 | standby | running | pgsql5   | default  | 100     | 2        |
host=pgsql6 user=repmgr dbname=repmgr port=5465
```



```
postgres@pgsql3:~$ repmgr -f /etc/postgresql/16/main/repmgr.conf standby unregister  
INFO: connecting to local standby  
INFO: connecting to primary database  
NOTICE: unregistering node 1  
INFO: standby unregistration complete
```

```
postgres@pgsql3:~$ repmgr -f /etc/postgresql/16/main/repmgr.conf cluster show
ID | Name   | Role   | Status   | Upstream | Location | Priority | Timeline |
Connection string
-----+-----+-----+-----+-----+-----+-----+-----+
3  | psql5  | primary | * running |          | default  | 100     | 2        |
host=pgsql5 user=repmgr dbname=repmgr port=5465
4  | psql6  | standby | running  | psql5    | default  | 100     | 2        |
host=pgsql5 user=repmgr dbname=repmgr port=5465
```



Ευχαριστώ

* No evzones were harmed in the making of this talk.