



repmgr Automatic Fail-Over

Cédric Villemain

<cedric@2ndQuadrant.fr>


FRANCE

CHAR(11)

07/11/11, Cambridge

License

- Creative Commons:
 - Attribution-Non-Commercial-Share Alike 2.5
 - You are free:
 - to copy, distribute, display, and perform the work
 - to make derivative works
 - Under the following conditions:
 - Attribution. You must give the original author credit.
 - Non-Commercial. You may not use this work for commercial purposes.
 - Share Alike. If you alter, transform, or build upon this work, you may distribute the resulting work only under a licence identical to this one.



What is repmgr ?

- PostgreSQL replication tool
- Use Streaming Replication and Hot Standby
- Monitoring solution
- NEW: failure detection and action
 - V2.0
 - Development happens months ago
 - Merge to Beta



Again

- Command line tool: repmgr
 - Replication set up
 - Cluster administration
- Daemon tool: repmgrd
 - Monitor replication
 - NEW: Handle automatic tasks
- NEW: Witness repmgrd



Fail-Over Key concepts

- Any number of nodes
- Master dead detection
- Find best Standby to promote
- Automatic and Secure




Automatic Extension

- Repmgrd monitor replication on each node
- New configuration parameters:
 - Failover= automatic|manual
 - Priority (of the node)
 - promote_command
 - follow_command
- Witness server
 - repmgr -f /etc/repmgr/repmgr.conf \
-D /opt/witness -h node1 -U repmgr -d repmgr \
witness create

Detection

- It's a heartbeat check
- Escalation
 - Heartbeat frequency
 - On failure, try several heart massages
 - On dead, speak to others




Voting

- Hardest work is to decide current status
- Postulate : Majority is right
 - Majority needs odd number of nodes in the cluster
 - A Witness helps a group to win the vote



Promotion

- Majority has decided to Fail-Over
- Who will be the new Primary ?
 - Lag decide
 - Priority parameter allows to have preferences

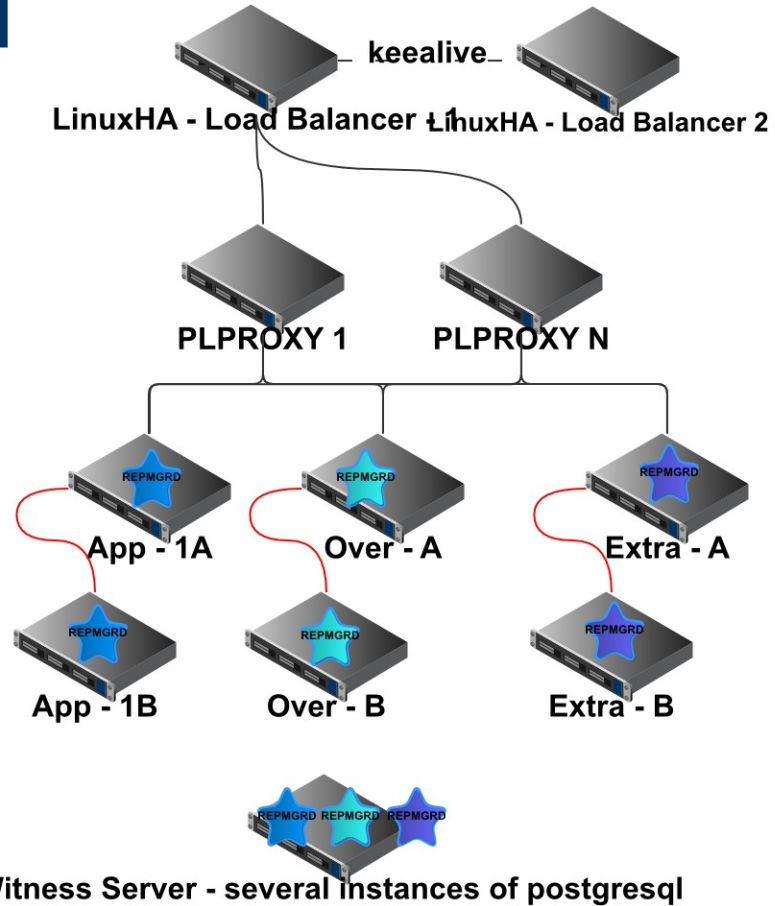


Action

- Apply promote_comand
- Apply follow_command
- **!!! DO NOT HANDLE STONITH !!!**




Example



Ideas & TODO

- pg base backup support
- PgBouncer communication
- Manual fail-over with automatic promotion/action
- Want more ? Ask us ! (or do it, it is **GPL** code)



Thanks

- All contributors of repmgr
- Our great sponsors



More information

- Website
 - <http://www.repmgr.org/>
- Public GIT
 - Github [greg2ndquadrant](#)
- IRC channel
 - freenode [#repmgr](#)
- Mailling-list
 - <http://groups.google.com/group/repmgr>
- Issues
 - <https://github.com/greg2ndQuadrant/repmgr/issues>

