EnterpriseDB®
The Enterprise PostgreSQL Company

PostgreSQL Clustering with Red Hat Cluster Suite

Devrim GÜNDÜZ
Principal Systems Engineer
EnterpriseDB
devrim.gunduz@EnterpriseDB.com


Use Red Hat Cluster Suite
for PostgreSQL Clustering

Goals

- Clustering goals
  - Active/passive clustering
  - Having a redundant system
    - Data redundancy
    - Network redundancy
    - Server and power redundancy
  - Maximum uptime
  - Service failover (=PostgreSQL)
  - Data integrity

Limitations

- What are the limitations?
  - No Active/active clustering in core.
  - No more than 16 nodes (uh? See next slide)

Simple Active/Passive Setup

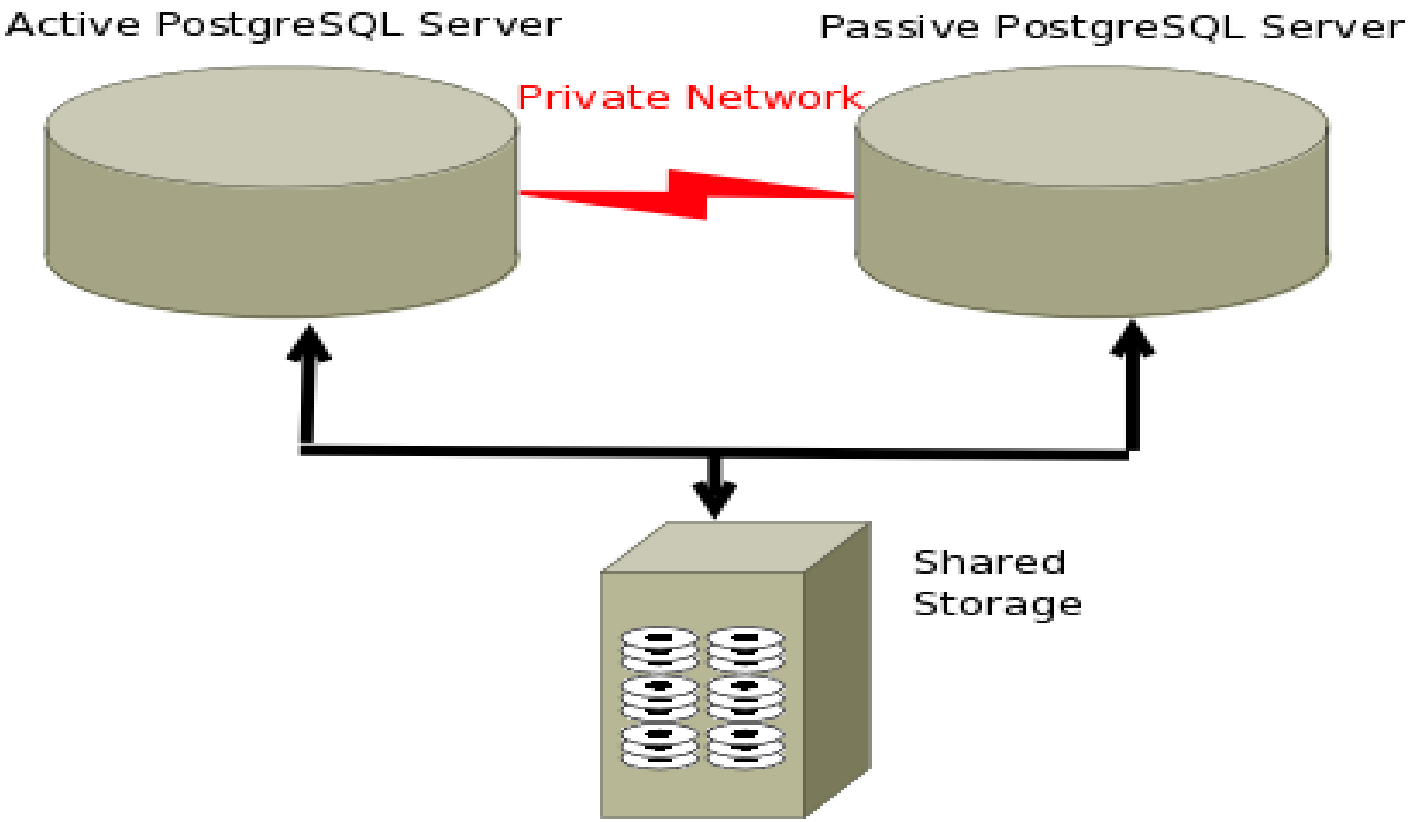

Sharding, multipe nodes

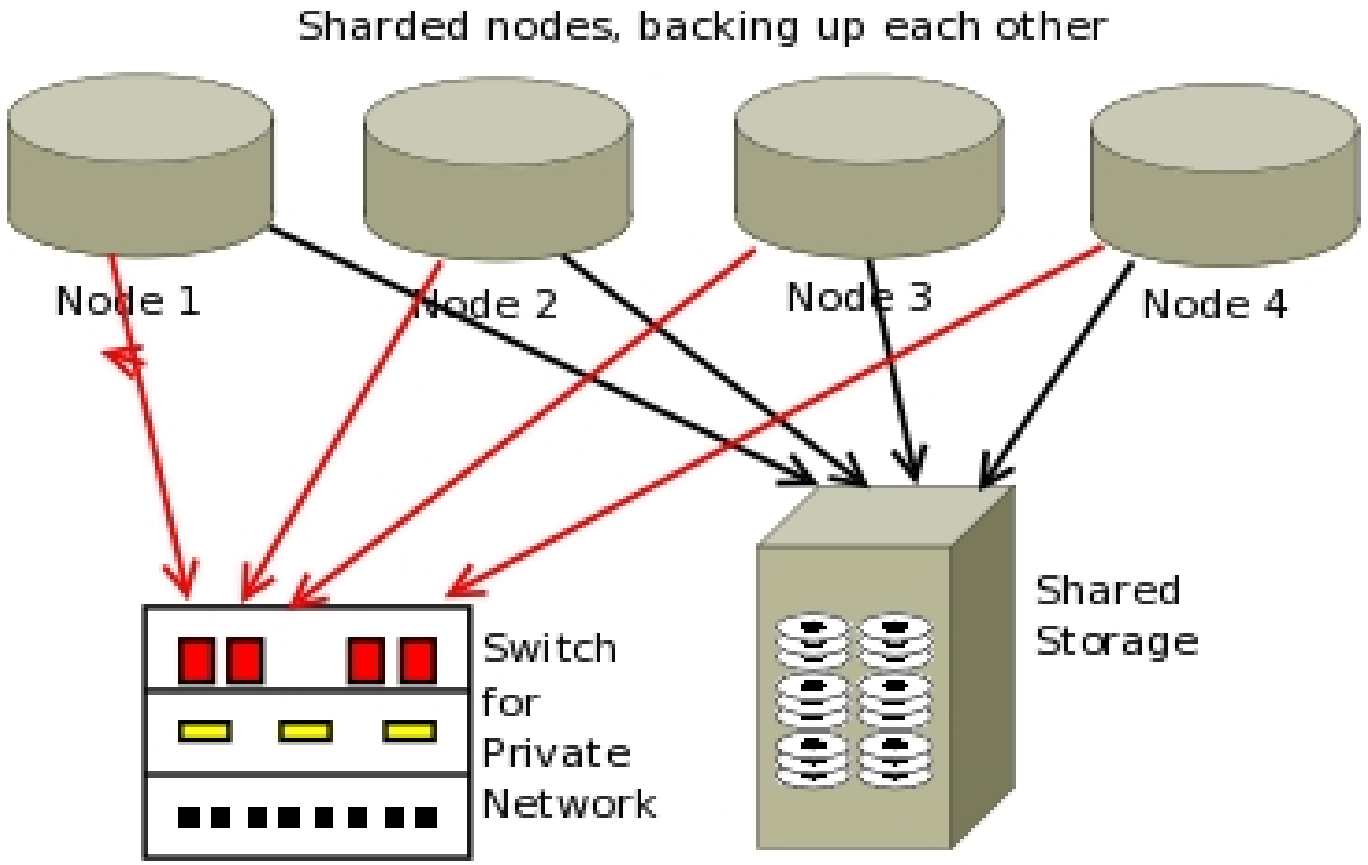

Hardware and software requirements

- Hardware
  - Minimum hardware: An hardware that Red Hat Enterprise Linux can run.
  - Typical hardware : Depends on your needs.
  - SAN : Storage is the most important part – Use Raid arrays.
  - Please read Greg's book.
- Software
  - RHCS is built on GFS.
  - GFS is built on LVM.
  - PostgreSQL :-)
  - Use RHEL 5.5+, or better, RHEL 6.1 (which was released today)

Setup and design

- We need two servers that has been setup identically.
  - Only OS and PostgreSQL will run
  - Same PostgreSQL versions.

- Using GFS, all data will be mounted from the storage. GFS is not a requirement, but we would better be safe.

- When node1 goes down, node2 will act as “active” server by announcing specified cluster ip. When node1 comes back, the process is reverted.

EnterpriseDB®
The Enterprise PostgreSQL Company

PostgreSQL Clustering with Red Hat Cluster Suite

Devrim GÜNDÜZ
Principal Systems Engineer
EnterpriseDB
devrim.gunduz@EnterpriseDB.com