PostgreSQL Clustering with Red Hat Cluster Suite

Devrim Gündüz

PostgreSQL Project

devrim@PostgreSQL.org

Goals

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

Limitations

- No Active/active clustering (PostgreSQL limitation)
  - PgCluster? (Not that suitable, header (“Location: PgCluster talk”);
- No more than 8 nodes (who would need it)(Red Hat limitation for Cluster Manager that we will use)

Hardware and software requirements

- Minimum hardware: An hardware that Red Hat Enterprise Linux can run.
- Typical hardware : Depends on your needs. An Opteron is preferred.
- SAN : Storage is the most important part – Use Raid arrays.
- RHCS is built on GFS.
- GFS is built on LVM.
- PostgreSQL :-)

Design and howtos

- 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.

PostgreSQL Clustering with Red Hat Cluster Suite

Devrim Gündüz

PostgreSQL Project

devrim@PostgreSQL.org