PostgreSQL

Richard Stephan
Senior Database Administrator
New York Independent System Operator

What is PostgreSQL?

Brief History of PostgreSQL

PostgreSQL History

- 8-May-2000 — 7.0.0 – Foreign keys, JOIN syntax
- 13-Apr-2001 — 7.1.0 – WAL, TOAST, Outer Joins
- 4-Feb-2002 — 7.2.0 – Vacuum nolonger locks tables, MD5 passwords.
- 27-Nov-2002 — 7.3.0 – Schemas, selects on functions, alter table drop column, prepared queries.
- 17-Nov-2003 — 7.4.0 – Multikey hash join
- 19-Jan-2005 — 8.0.0 – Savepoints, Tablespaces, P-I-T-R (hot backups)
- 8-Nov-2005 — 8.1.0 – Two phase commit, Roles replaced groups, Table partitioning
- 5-Dec-2006 — 8.2.0 – Better warm standbys
- 4-Feb-2008 — 8.3.0 – XML
- 1-Jul-2009 — 8.4.0 – Windowing functions, Common table expressions, Column permissions

PostgreSQL – Open Source RDBMS

- Open source has the high potential to create quality software. Everyone collaborates, the best software wins. Not just within one company, but among an Internet-connected, worldwide community.
- As a result, the open source model often builds higher quality, more secure, more easily integrated software. And it does it at a vastly accelerated pace and often at a lower cost.
- However, there exists support and licensing Issues.

PostgreSQL

- Database cluster.
  - *Accounts and roles*
  - *Tablespaces*
  - *Write Ahead Logs (WAL)*
  - *Server configuration file*
  - *Server log file*
  - *Databases*
    - Schemas
      - Database objects (tables, views, indexes, triggers, functions, …)

PostgreSQL – Storage

- Tablespaces – OS directories
- Tables, indexes, sequences – data files
  - *Multiples of 1 GB.*
  - *TOAST – The Oversized Attribute Storage Technique*
- Write Ahead Logging (WAL)
  - *Transaction logging*
  - *Roll forward recovery (i.e., REDO)*
  - *"Recycle" and archival*
- Temporary space
  - *New with 8.3 – temp_tablespaces*

PostgreSQL – Procedural Languages

- User-defined functions (Stored Procedures and Triggers).
- Four procedural languages come standard:
  - *PL/pgSQL*
  - *PL/Tcl*
  - *PL/Perl*
  - *PL/Python*
- PL/pgSQL
  - *Easy to use and similar to Oracle's PL/SQL.*

PostgreSQL – Export/Import

- pg_dump - Export
  - *Output to script or archive*
- pg_dumpall – Cluster level export
- pg_restore – Import for a pg_dump archive

PostgreSQL – Enterprise Class RDBMS?

- Schemas, Roles, Accounts
- Tablespace Management
- Table Partitioning
- Write-Ahead Logging (WAL)
  - *Point-In-Time-Recovery (Archive logs)*
  - *On-Line/Hot Backups*
- Multi-Version Concurrency Control (MVCC)
- Nested Transactions (savepoints)
- Query Planner/Optimizer
- Network Encryption
  - *Secure TCP/IP with SSL*
  - *MD5 Password Encryption in Database and on Network*
- Enterprise Accounts (i.e., LDAP Authentication)

PostgreSQL at the NYISO

PostgreSQL on the Veritas Cluster for Intranet Applications

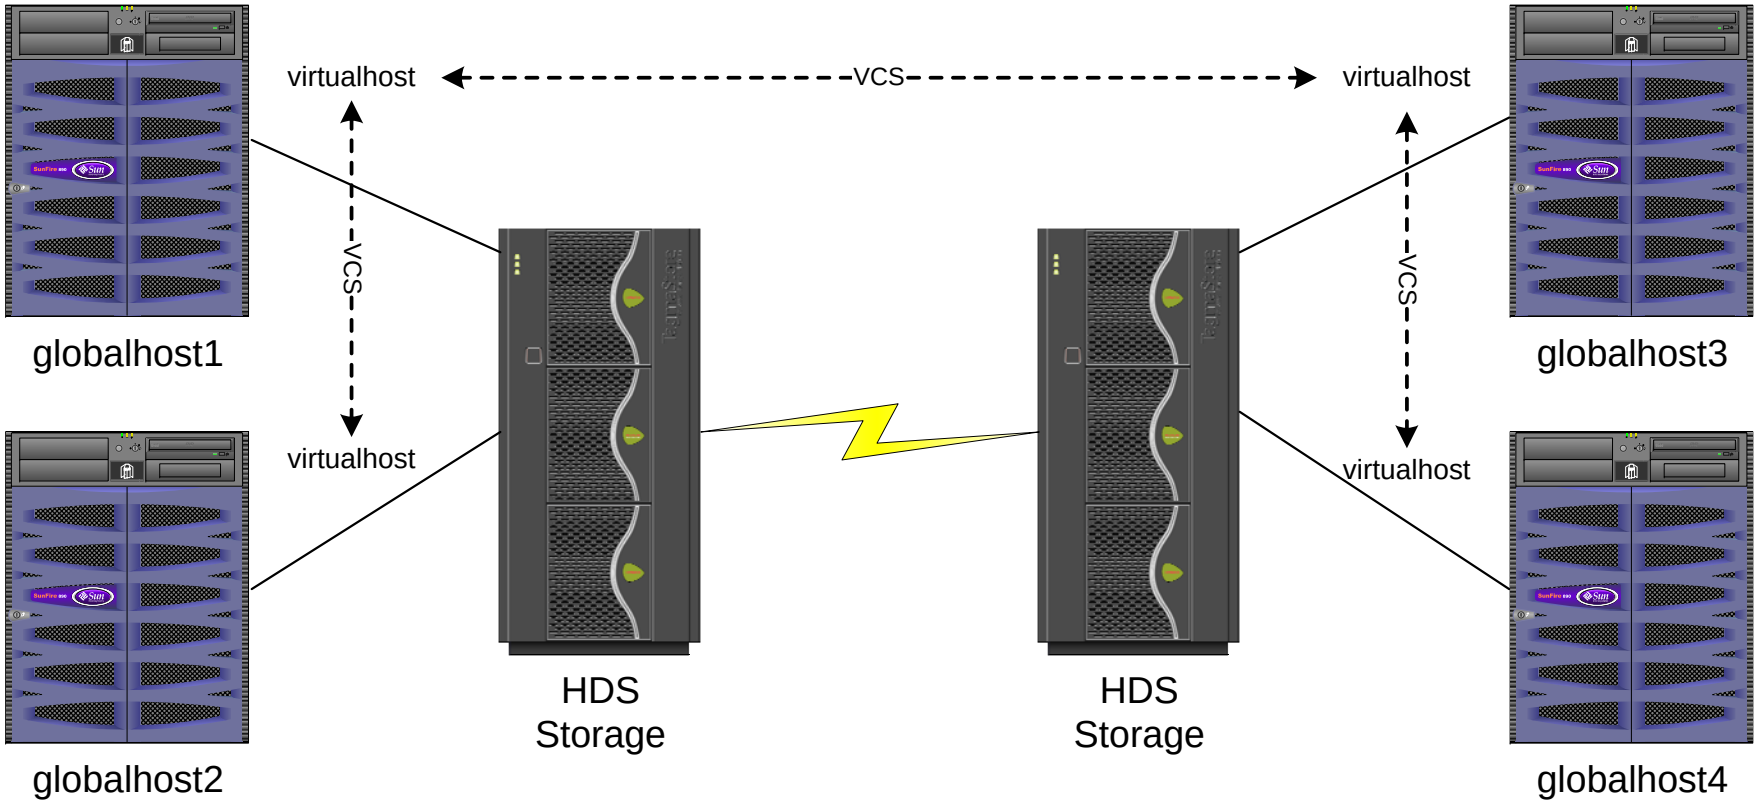

PostgreSQL at the NYISO

PostgreSQL on the Veritas Cluster for Intranet Applications

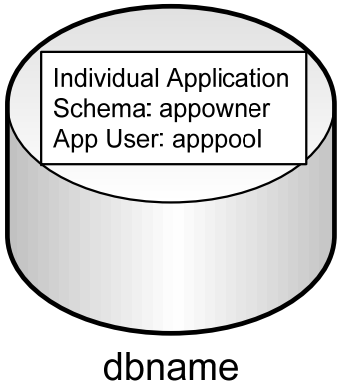

/opt/postgres/8.3
/archive1/dbname
/backup1/dbname
/pgcluster/admin/data
/pgcluster/admin/log
/pgdata1/dbname/system
/pgdata2/dbname/data
/pgdata3/dbname/index

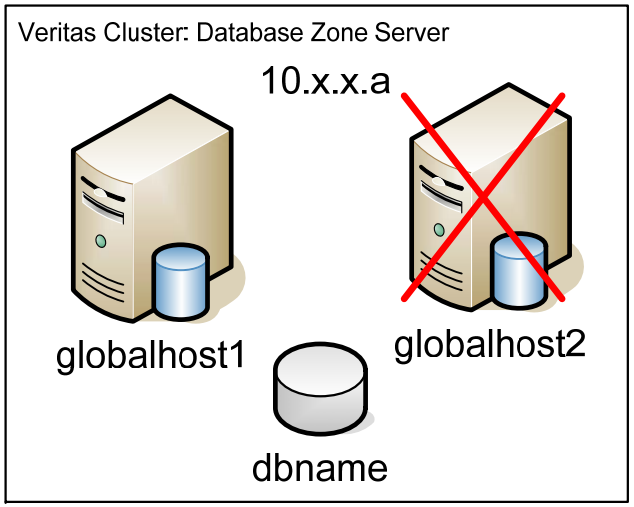

virtualhost

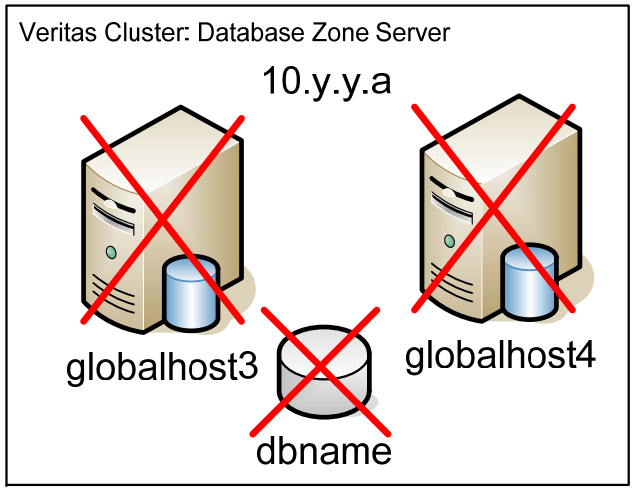

Oracle's Advantages over PostgreSQL

- Oracle Real Application Cluster (RAC)
  - *$23K/CPU & $5060/yr/CPU*
- Oracle Data Guard
  - *PostgreSQL's warm standby with archive log shipping is similar to Oracle DG in Maximum Performance mode.*
- Parallel query.
- RMAN
- Lack of tools and wide spread adoption.

NEW YORK INDEPENDENT SYSTEM OPERATOR

The New York Independent System Operator (NYISO) is a not-for-profit corporation that began operations in 1999. The NYISO operates New York's bulk electricity grid, administers the state's wholesale electricity markets, and provides comprehensive reliability planning for the state's bulk electricity system.

www.nyiso.com