What is wrong with PostgreSQL?
OR
What does Oracle have that PostgreSQL should?

Richard Stephan

PostgreSQL is an Enterprise 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
- Secure TCP/IP with SSL
- Enterprise Accounts (i.e., LDAP Authentication)

PostgreSQL is an Enterprise RDBMS

- PostgreSQL's stability, flexibility, security, manageability, reliability, scalability, performance, adherence to standards, and supportability (EnterpriseDB, CommandPrompt, 2ndQuadrant, …) set it amongst its proprietary commercial competitors (Oracle, DB/2, Sybase, SQL Server).
- However, there are some features or capabilities that are missing or diminished when compared to the number one RDBMS - Oracle.

No Parallel Query

- All queries run within a single process.
- This is acceptable for most OLTP applications; however, not for warehousing and business intelligence.
- Solution … fork it! into an MPP (Massively Parallel Processing) environment.
  - Netezza
  - Greenplum
  - Aster Data
  - Stado

No Horizontal Scalability or Clustering

- Currently PostgreSQL scales up to 32 CPUs on a single host. The community is working on getting it to scale up to 64 CPUs for version 9.2.
- There are two projects currently underway to create a clustered solution.
  - Postgres-XC
  - Postgres-R
- Those are both forks and they are designed to be a shared nothing architecture versus Oracle's RAC which is a shared disk solution.

No In Place Upgrades!

- When changing the third version position (8.4.7 → 8.4.8; 9.0.3 → 9.0.4) there is no need to upgrade. These usually are software patches.
  - Stop the database and switch out the software.
- However, moving from the first or second version position (8.3 → 8.4; 8.4 → 9.0) requires an export and import of the data.
- This is fine for small data sets, but an outage for a rather large database would be time consuming.
- There is a new contributed program called pg_upgrade.
  - Unfortunately that still is not an in place upgrade. It is just a fast migration tool. It takes the current database and makes a new one in the new version.

No Multiplexing of Reliability Components

- Oracle multiplexes the control file and the redo logs.
- There is a program called pg_resetxlog. Its intention is to clear out WAL files and reset them and it updates the control file as needed. This is the current solution for "fixing" those files if they get corrupted.
- I have inquired whether there were any plans to multiplex or some other method to increase the reliability of these very important components.
  - I was told not yet. And the likelyhood was very low. Just make sure you have good backups like always.

RMAN

- Oracle RMAN.
  - Backup and Recovery management.
  - Incremental backups.
  - Single datafile restore.
  - Datafile corruption detection.
- Luckily backup and recovery is very easy in PostgreSQL. Of course you might argue that it is not very flexible. It is an all or nothing solution. There is no ability to restore and recover a single data file.
- There are third parties developing an RMAN like solution.

Data Guard

- Oracle has three modes:
  - Maximum Performance
  - Maximum Availability
  - Maximum Protection
- PostgreSQL Streaming Replication and Hot Standby
  - 9.0 Asynchronous
  - 9.1 Synchronous

Now for a Rebuttal – Cool things PostgreSQL has that Oracle does not.

- Two Phase Commit
  - Oracle says it has two phase commit. And it sort of does through the proprietary implementation of database links.
- ANSI-SQL Standard Data types (Integer, Boolean, Date, Time, …)
- A zero length string is different than a null string.
- Transactional DDL (Including table truncates.)
- Table Inheritance (ORDBMS)
- Proprietary: Notify – Listen
- SQL/MED (Management of External Data)