



Elephant on Solaris

Zdeněk Kotala

Solaris Sustaining Engineer

Sun Microsystems



Agenda

- About PostgreSQL
- Integration with OpenSolaris
 - > Packages
 - > SMF
 - > RBAC
- Tuning
- DTrace
- Sun cluster

PostgreSQL

- One of the most advanced opensource SQL databases
- Fully ACID compliant (Atomicity, Consistency, Isolation, Durability)
- Full supports for foreign keys, joins, views, triggers, and stored procedures (in multiple languages)
- Support most of SQL92, SQL99
- www.postgresql.org

PostgreSQL history

- Postgres started in 1986
 - > Fork of Ingres
 - > Project led by Michael Stonebraker
 - > Ended in 1993
- Postgres95
 - > Released in 1994 by Andrew Yu and Jolly Chen
 - > New SQL parser
- PostgreSQL
 - > All sources released on the internet in 1996

PostgreSQL on Solaris

- Announced in December 2005
- Released as part of Solaris 10 6/06
- DTrace support from PostgreSQL 8.2
- Niagara T2000 donated to the PostgreSQL community

PostgreSQL 8.1 packages

- First in Solaris 10 6/06
- Advantages
 - > Files are located in standard path `/usr/bin/`, `/usr/lib/`
 - > Datadir is located in `/var/lib/pgsql/`
- Disadvantages
 - > Installation of more PostgreSQL version on one system is not easy
 - > Patch has ~100MB size (contains complete doc and src)

PostgreSQL 8.2 packages

- Planned for Solaris 10 6/07
- In Nevada from build 56
- Advantages
 - > Files are located in /usr/postgres/8.2/ directory and data directory is /var/postgres/8.2
 - > It is integrated with SMF and RBAC
 - > Design of packages allows delivery of changed files only
- Disadvantage
 - > Files are located in non standard path. The library search path must be set.

PostgreSQL compiled with

- Openssl
- Kerberos v5 (not in 8.1 on S10)
- DTrace (new in 8.2)
- NLS – National Language Support
- pg_perl, pg_python, pg_PL, pg_TCL
- Thread safety

Service Management Facilities - SMF

- Two instances for PostgreSQL
 - > svc:/application/database/postgresql:version_81
 - > svc:/application/database/postgresql:version_82
- Three properties
 - > postgresql/bin
 - > postgresql/data
 - > postgresql/log

RBAC – Role Based Access Control

- Similar to sudo, has finer granularity
- New profile - Postgres Administration
 - > Allows user to run commands like initdb, pg_ctl with postgres uid
- New role - postgres
 - > Role is assumed by “su – postgres” command
 - > Necessary for data dir access (e.g. postgresql.conf)

Initial setup

- Setup password for postgres role (it can be blank)
- Add postgres role and “Postgres Administration” profile to database administrator user(s)
- Create database cluster
 - > pfexec /usr/postgres/8.2/bin/initdb /var/postgres/8.2/data
- Setup postgresql.conf
 - > su – postgres
 - > vi /var/postgres/8.2/data/postgresql.conf
 - > exit
- Start PostgreSQL
 - > svcadm enable postgresql:version_82

Tuning

- There are no general rules
- Three tuning places
 - > Database application design
 - > PostgreSQL configuration
 - > Solaris configuration + hardware setup

Tuning – disk subsystem

- UFS

- > Set `wal_sync_method = opendata_sync` in the `postgresql.conf` (open(2) option `O_DSYNC`)
- > Use `forcedirectio mount(1M)` option on filesystem with WAL files
- > Good for OLTP applications

- ZFS


- > Set `wal_sync_method = fsync` in the `postgresql.conf` (all `fdatasync()` at each commit)
- > `full_page_writes = off`
- > Good for Data Warehouse applications

DTrace

- Dynamic tracing framework
- Eleven USDT probes in PostgreSQL 8.2
 - > One group of probes is focused on locks activity
 - > Second is focused on transaction duration
- Example DTrace scripts are on www.pgfoundry.org
- Useful for application bottleneck analysis

Sun Cluster (only for S10)

- PostgreSQL dataservice available since version 3.2
- Only as a failover service





Elephant on Solaris

Zdeněk Kotala

zdenek.kotala@sun.com

