Postgres Is Different From (*better* than) Your Database

Gurjeet Singh | 2013/09/18

EDB ENTERPRISEDB

POSTGRES IS DIFFERENT FROM (*BETTER* THAN) YOUR RDBMS

Gurjeet Singh
gurjeet.singh.im
EDB - EnterpriseDB.com

FROM THE PERSPECTIVE OF A

- Developer
- Manager

DEVELOPER'S PERSPECTIVE

- License
- Reliability
- Security
- Features
- Extensibility
- Performance
- Diagnosability

DEVELOPER : LICENSE

- BSD-like License
  - Liberal than/Very different from GPL.
  - PostgreSQL License gives you
    - Freedom to read, patch/improve the code.
  - An opportunity to understand the database internals
  - Install, use/develop and deploy full feature-set without a license
  - Upgrade hardware without paying for license cost
  - Unlike commercial RDBMSs where licenses come in the way of development, QA, and hardware upgrades.

DEVELOPER : RELIABILITY

- ACID compliant
- Constraints (Primary Key, Foreign Key, CHECK)
- Data-page checksums to detect corruption

DEVELOPER : SECURITY

- Per-object GRANT/REVOKE permissions
- Per-column GRANT/REVOKE permissions
- SSL connections
- LDAP and RADIUS authentication

DEVELOPER : FEATURES

- Postgres is a Platform
- Vast feature set
- ... and ever-expanding

Postgres Feature Matrix

http://www.postgresql.org/about/featurematrix/

DEVELOPER : FEATURES

- Common Table Expressions (CTEs)
  - Similar to Macro in programming languages
  - Supports recursive evaluation (alternative to Oracle's CONNECT BY)

DEVELOPER : FEATURES : DATA TYPES

- Many advanced data types come builtin

Full List of Builtin Data Types

http://www.postgresql.org/docs/9.3/static/datatype.html

DEVELOPER : FEATURES : DATA TYPES

- Does *not* have/need NVARCHAR data type
  - The CHAR/VARCHAR/TEXT types are capable of storing Unicode data

DEVELOPER : FEATURES : DATA TYPES

- Boolean (a first-class data type)
- Bit (and Bit strings)
- Money
- Bytea (binary data)
- Interval (difference between TIMESTAMPs)
- Enumerated Types
- Geometry Types (Point, Line Segment, Polygon, ...)
- Network Address (inet, cidr, macaddr)

DEVELOPER : FEATURES : DATA TYPES

- Range Types (tsrange, daterange, int4range, ...)
- Composite types
- Builtin multi-dimensional arrays

  Automatically defines array type of builtin and user-defined data types

DEVELOPER : FEATURES : DATA TYPES

- hstore
- JSON
- XML

DEVELOPER : FEATURES

- Streaming Replication
- Hot Standby
- Per-transaction synchronous replication
- True serializable transactions
- LISTEN/NOTIFY
- Triggers
- Exclusion constraints
- Window Functions
- Function overloading
- Operator overloading
- Full-text search
- Large Objects up to 4TB
- Materialized views

DEVELOPER : FEATURES

- Create procedures in various programming languages
  - PL/pgsql
  - PL/perl
  - PL/tcl
  - PL/java
  - PL/v8 - JavaScript

DEVELOPER : FEATURE/PERFORMANCE

- TOAST

  The Oversized Attribute Storage Technique

  Automatic out-of-line storage and automatic compression

DEVELOPER : PERFORMANCE

- Cost-based optimizer
- Index-only scans
- Synchronized seq-scans
- Table partitioning
- Unlogged tables

DEVELOPER : PERFORMANCE

- Partial indexes
- Many types of indexes
  - BTree
  - Hash
  - Gin
  - Gist
  - KNN
  - SP-Gist

DEVELOPER : PERFORMANCE

- Types of join-strategies
  - Nested-loop joins
  - Hash joins
  - Sort-Merge joins
  - Bitmap-index scan; Bitmap-And/Or joins (in-memory)
    - Can use multiple indexes of a table in the same scan
  - Semi joins
  - Anti joins

DEVELOPER : PERFORMANCE

- Faster than NoSQL (at NoSQL use case)
- See slides 32 onwards of Chrisophe Pettus'

PostgreSQL as a Schemaless Database

http://thebuild.com/presentations/pg-as-nosql-pgday-fosdem-2013.pdf

DEVELOPER : DIAGNOSTICS

- Per-function statistics
- Performance Views
  - pg_stat_activity
  - pg_locks
  - pg_stat_*
    - pg_stat_all_tables
    - pg_stat_all_indexes
    - ...

DEVELOPER : DIAGNOSTICS

- Logging Options
  - log_min_duration_statement
  - log_temp_files
- DTrace/SystemTap support

  Scripts to record and diagnose performance issues

DEVELOPER : EXTENSIBILITY

- Highly Extensible
- Create Your Own
  - Data Types
  - Operators
  - Index Types
- `create extension`
- Check out **PGXN.org**
- Foreign Data Wrapper
- Foreign Tables

DEVELOPER : EXTENSIBILITY

- Examples
  - PostGIS

    PostGIS adds support for geographic objects to the PostgreSQL object-relational database.
  - PostgreSQL-HLL

    A PostgreSQL extension adding HyperLogLog data structures as a native data type

MANAGER'S PERSPECTIVE

- Acquisition Cost (License)
- Future Proof
- Vendors/Providers (Companies who can help)
- Maintenance Cost (Support)
- Buy-in From Upper Management

MANAGER : ACQUISITION COST

- Free

MANAGER : ACQUISITION COST

- FREE
- Save Money $$$
- Higher Profit Margins
- A Leg Up On The Competition

MANAGER : ACQUISITION COST

- BSD-like License
  - Liberal than/Very different from GPL.
  - Install, use/develop and deploy full feature-set without a license
  - Upgrade hardware without paying for license cost
  - Unlike commercial RDBMSs where licenses come in the way of development, QA, and hardware upgrades.

MANAGER : FUTURE PROOF

- Community-driven
- No single commercial entity controls it (unlike MySQL)
- Will be always free (cost and project management)

MANAGER : FUTURE PROOF

- Ever-increasing resource-pool; developers and DBAs
- Compare growth in job requirements

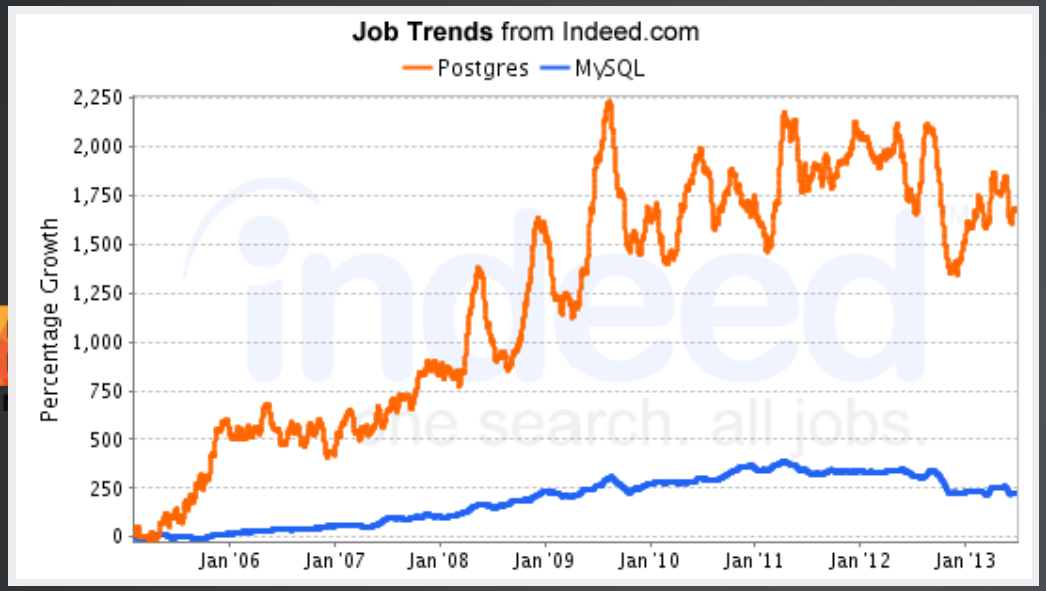

MANAGER : VENDORS/PROVIDERS

- Many top-quality vendors
  - EDB (EnterpriseDB)
  - PG Experts
  - 2nd Quadrant
  - OmniTI
  - Command Prompt
- Many smaller consulting teams/individuals

MANAGER : MAINTENANCE COST

- 24/7 Support available
- Very affordable support contracts

  For e.g. EDB charges per-socket (unlike Oracle, that charges per CPU)

MANAGER : MAINTENANCE COST

- Consulting
- Remote-DBA
- Database Health Checks
- Training
  - For Developers
  - For DBAs

MANAGER : UPPER MANAGEMENT BUY-IN

- Show them the last few slides :)

THANK YOU