2ndQuadrant®
PostgreSQL

PostgreSQL: Hyperconverged DBMS

PGConf India 14th Feb 2019

Simon Riggs, CTO, 2ndQuadrant
PostgreSQL Major Developer

Historical Perspective

Professor Michael Stonebraker

- Leader of the original Postgres project - ***Thanks!***
- Leader of the first commercialised version of Postgres

- Later criticised Postgres
  - **One-size-does-not-fit-all**
  - Postgres .... is an Elephant!

- Then went on to launch Vertica, H-Store etc..

Recent History

Proliferation of DBMS

- DB-Engines lists 343 DBMS currently available
- Why so many?
- For any ***specific*** task, there is always a best tool
  - No single tool is best at every task
- Divide any workload into enough small pieces and you'll be able to pick the right tool for every piece

- But will the whole architecture work smoothly?

Elephants

Long life, Wisdom and Persistence

- Arthur C. Clarke's First Law:
  When a distinguished but elderly scientist states that something is possible, he is almost certainly right. When he states that something is impossible, he is very probably wrong.
- Can an Elephant succeed?
- Can one DBMS do it all? If so, how?

PostgreSQL Plaudits

PostgreSQL public recognition

- DBEngines Database of the Year 2017 and 2018
- Cloud Adoption by just about everyone
- Stack Overflow #1 Best Loved **R**DBMS 2018
- Other databases following PostgreSQL features
- Many forks, but open source version is strong

- Yes, an Elephant can succeed

PostgreSQL Use Cases

Many ideas, many use cases

- Relational
- Post-Relational
- GIS
- Full Text Search
- XML/JSON document
- Graph
- Time-Series
- Business Intelligence

PostgreSQL Relational

The features most people think of

- Tables, Views
- Constraints
- Foreign Keys
- Triggers

- Full and complete implementation of SQL:2016
- Excellent Cost-based Optimizer
- Robust and easy to use

PostgreSQL *Post*-Relational

- Abstract types
- Domains
- Arrays

- Functions
  - Polygamous
  - Variadic
- Procedures
  - Transactional

PostgreSQL GIS via PostGIS

- Implements full OpenGIS standard
- Points
- Lines
- Geometries
- Many specialized indexes for these data types
- Fully integrated with many mapping applications

PostgreSQL Full Text Search

- Many specialized indexes for these data types (again!)
- Phrase search
- Weights, ranking and headlining
- Parser supports 23 modes, including hyphenation
- Dictionary support for synonyms, stop word lists, thesaurus, spelling and stemmer functions
- Full collation support

PostgreSQL XML/JSON Documents

- XML support
- JSON and JSONB
- Benchmarks show better compression than MongoDB
- Accessor functions
- Serialize/Deserialize
- Iterator functions
- In-place update functions

PostgreSQL Graph Query

- Recursive query support in core PostgreSQL
  - Efficient use of indexes
  - Works with relational and XML/JSON
- Proven that all SPARQL queries can be written in SQL
- LDBC Benchmarks are all *expressable* in SQL
- Benchmarks show PostgreSQL **beating Neo4j**

PostgreSQL Time Series

- Full SQL:2016 time and date handling
- Complete timezone library
- Window Functions with Sliding Window queries
- Range data types
- BRIN indexes
- Partitioning for Time Series
- Optimizations for monotonic keys (timestamps)

PostgreSQL Business Intelligence

- Excellent Cost-based Optimizer
- Extended optimizer statistics
- Plan-types suited to VLDB
- Many different optimizations benefit TPC-H workload
- Significant sort improvements
- Aggregation improvements
- Just in Time Compilation
- Parallel Query, Parallel Aggregation, Parallel Join
- Hot Standby for operational database offload

PostgreSQL Extensions

Variety of additional modules to extend PostgreSQL

- PostGIS
- Tsearch2
- Pglogical
- Postgres-BDR
- Citus
- ...many, many others from many contributors

PostgreSQL Extensibility

Features pioneered within PostgreSQL

- User defined functions
- User defined datatypes
- Pluggable indexes
- User-defined ANALYZE
- Many plugin interfaces
- EXTENSION packaging
- Generic WAL (transaction log)
- Table Access Methods (Ok, bit late on that one)

PostgreSQL

- Can one DBMS be best at everything? **No**
- Can a DBMS with huge numbers of contributors collect together to produce something no one mind could contemplate, covering multiple use cases? **Yes, but slowly**
- Can one DBMS provide the facilities for multiple additional features via extensibility? **Definitely**

PostgreSQL Solutions

- PostgreSQL Core can do **more** than people admit
  - Developing rapidly with releases every 1 year
- PostgreSQL System Extensions add even more value
  - Developing even more rapidly every ~3 months
- PostgreSQL Custom Server Applications
  - Developing custom solutions in ~1 month

PostgreSQL Ecosystem

- Consider the whole ecosystem as a unit, PostgreSQL is a hyperconverged system
- Agile, extensibility not monolithic development allow 343 ideas to be brought together into one whole
- **Multiple use cases** in one *flexible* system of interlocking components
- This is why Postgres succeeds

How 2ndQuadrant follows this thinking

- Heavy contributor to open source PostgreSQL
- Provides own distro, 2QPG, **if** you need it
- Provides many tools and value-added extensions
- Rapid development of practical solutions, now

2ndQuadrant Solutions

- 24/7 Bug-fix Support and RemoteDBA
- Consulting and Training
- Product Stack focused on enhanced
  - Performance
  - High Availability
  - Security

  for your PostgreSQL apps

2ndQuadrant PostgreSQL Solutions

Website https://www.2ndquadrant.com/
Blog https://blog.2ndquadrant.com/
Email info@2ndQuadrant.com