Goldman Sachs

@PGCONFNYC 2014 KEYNOTE

INCREASE ENTERPRISE POSTGRESQL ADOPTION AT FIRMS LIKE GOLDMAN SACHS IN FINANCIAL SERVICES

Name	**Bryan Doyle**
Title	**Vice President,**
	Database Technology
Email	**bryan.doyle@gs.com**

[03-Apr-2014]

Goldman Sachs Technology Overview

Goldman Sachs' Largest Division	■ Over 8,000 people in 31 offices, in 15 countries ■ Larger than many pure technology companies
Data Center Statistics	■ Driven by our custom private cloud ■ 300,000 square feet, equivalent to about 140 tennis courts ■ 9.85 petabytes of SAN, 8.5 petabytes of NAS
Production Code	■ 1.2+ billion lines ■ Every major programming language, plus our own
Messages	■ 100 million unique emails per month
Daily Risk Calculations	■ 47,000 CPUs ■ Process 10 billion prices and 2.5 quintillion 64-bit ops

Source: goldmansachs.com

Goldman Sachs Database Organization

Goldman Sachs

Database Team Overview

- Centralized team responsible for:
  - Tooling
  - Platform selection
  - Maintenance
  - Support
- 10+ platforms (RDBMS & NoSQL)
- 1000's of distributed DBs
- Service offering focus
- Developer self-service
- DDL management by AppDev

GS PostgreSQL Positioning

- Supported in-house
- Allowed PostgreSQL use cases are limited
  - 3rd party vendor software
  - Select internally developed non-critical apps
- PostgreSQL tooling
  - Product is missing important hooks
  - Internal offering has less functionality as a result
- We like PostgreSQL!
  - That's why we're here, but…
  - We like commercial DBMS products too

PostgreSQL Economics at Large Enterprises

Goldman Sachs

As costs to deploy another commercial DB instance diminish, PostgreSQL must compete directly

	PostgreSQL	Proprietary Databases
License	■ Free	■ Bulk license and support agreements ■ Minimal marginal cost per unit
Hardware	Commodity servers in a private cloud	

External Support	■ Risk associated to public channel interaction model
	■ Support contract options
	■ Community engagement can be
	  ■ Expensive (time)
	  ■ Met with resistance
		■ Generally fixed and priced into an existing agreement (above)
		■ Greater product influence

Platform Tooling 	■ Onboarding and maintaining new platforms has a significant cost
	  ■ Open source or commercial
	  ■ Similar platforms that don't add value won't be on-boarded
	■ Non-Enterprise ready platforms are more expensive to onboard	

PostgreSQL is a Valuable Database Platform

Goldman Sachs

Key Strengths

- Multi-Version Concurrency Control (MVCC)
- Open Source
- ANSI SQL
- Solid code base
- User community

Value Add

- Many procedural languages
- Extensibility
- Increasing replication capabilities
- Foreign Data Wrappers (FDWs)
- Much More…

PostgreSQL has Challenges in the Enterprise

Goldman Sachs

Features

- Performance considerations
  - Needs better parallelism
  - Better memory and file management
- Needs to be more programmatic
  - Set / Get config settings remotely
  - hba.conf
- Native (tunable) auditing
- Compression
- Service names
  - First class support in server
  - Wildcard LDAP lookup

Externalities

- Lifecycle visibility
  - Great at PostgreSQL level
  - Ecosystem
    - PGXN is a good start
    - Still long term risk of onboarding non-core extensions
- “Contrib” modules
  - Should be “Core Extensions”
  - Currently sounds like a “best effort” by outsiders
- Training modules lacking
- Most developers aren’t demanding yet another RDBMS these days

PostgreSQL Enterprise Engagement and Support

Goldman Sachs

There are challenges engaging the wider PostgreSQL community for support
Introducing a PostgreSQL vendor adds cost, diminishing potential value

Enterprise Engagement with Open Source Projects Often via:

- Commercial entity
- Internal support (only)

Local User Groups

- Great for networking and knowledge sharing
- Not as good for wider influence

Risk Posting Publicly

- Electronic communication retention requirements
- Information Leakage

"Help! My Database is Down!"

- No SLAs, but community is pretty responsive
- Delay and sensitivity concerns

PostgreSQL has Opportunities to Increase Value

Goldman Sachs

Compete Head-To-Head with Other Platforms	■ Developers need to want and demand PostgreSQL over other platforms ■ Need to provide more value than cost of onboarding ■ Value must be found in both features and performance
Rebrand! Perhaps with 10.x Release? **Revisit 2007 Decision**	■ In 1996, “QL” was added to the product to promote SQL support ■ A lot more added to PostgreSQL than SQL but it has lower visibility ■ PostgreSQL deserves more buzz… Postgres?... PostgresDB? … ■ Takes research to find its full capabilities and developers aren’t aware
Industry Focused Advisory Groups?	■ Collect aggregate product feedback from industry verticals ■ Allow consensus to be collected for easier follow-up ■ Contributions attributed to a sector instead of one interest
Better APIs in Addition to SQL Layer	■ Maintain strict SQL conformity layer ■ Some features (e.g.-JSON) not cross compatible with other platforms ■ Great APIs are a reason to adopt NoSQL ■ (Even with less native functionality)


Goldman
Sachs