PostgreSQL Scale-Out at InMobi

Sivakumar

Agenda

- Introduction
- PostgreSQL Usage Architecture Evolution
- Current challenges

Introduction

- InMobi
  - Founded in 2007
  - World's largest independent mobile advertising network
  - 215 MILLION dollar funding by Softbank, KPCB and Sherpalo.
  - Analyzing 240TB of data/day to enable consumers and businesses make smarter decisions
  - Our ads engage 761 million unique users in our network
  - Technology powerhouse Contributing to various open source projects like Apache Hadoop, Apache Hive, Apache Falcon..

Architecture Evolution V0

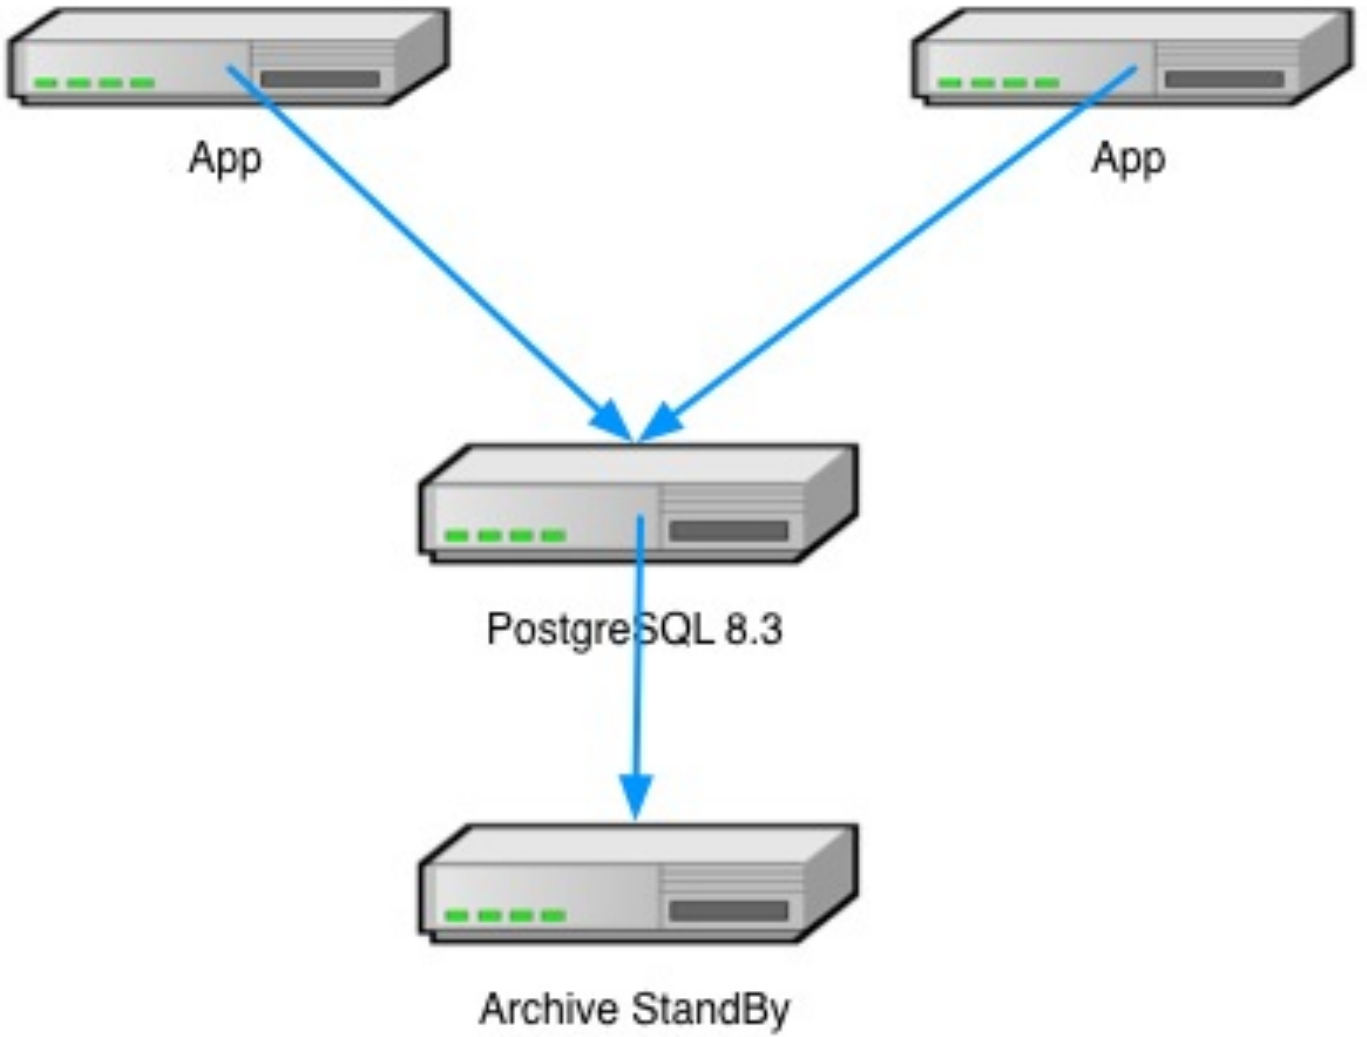

Architecture Today

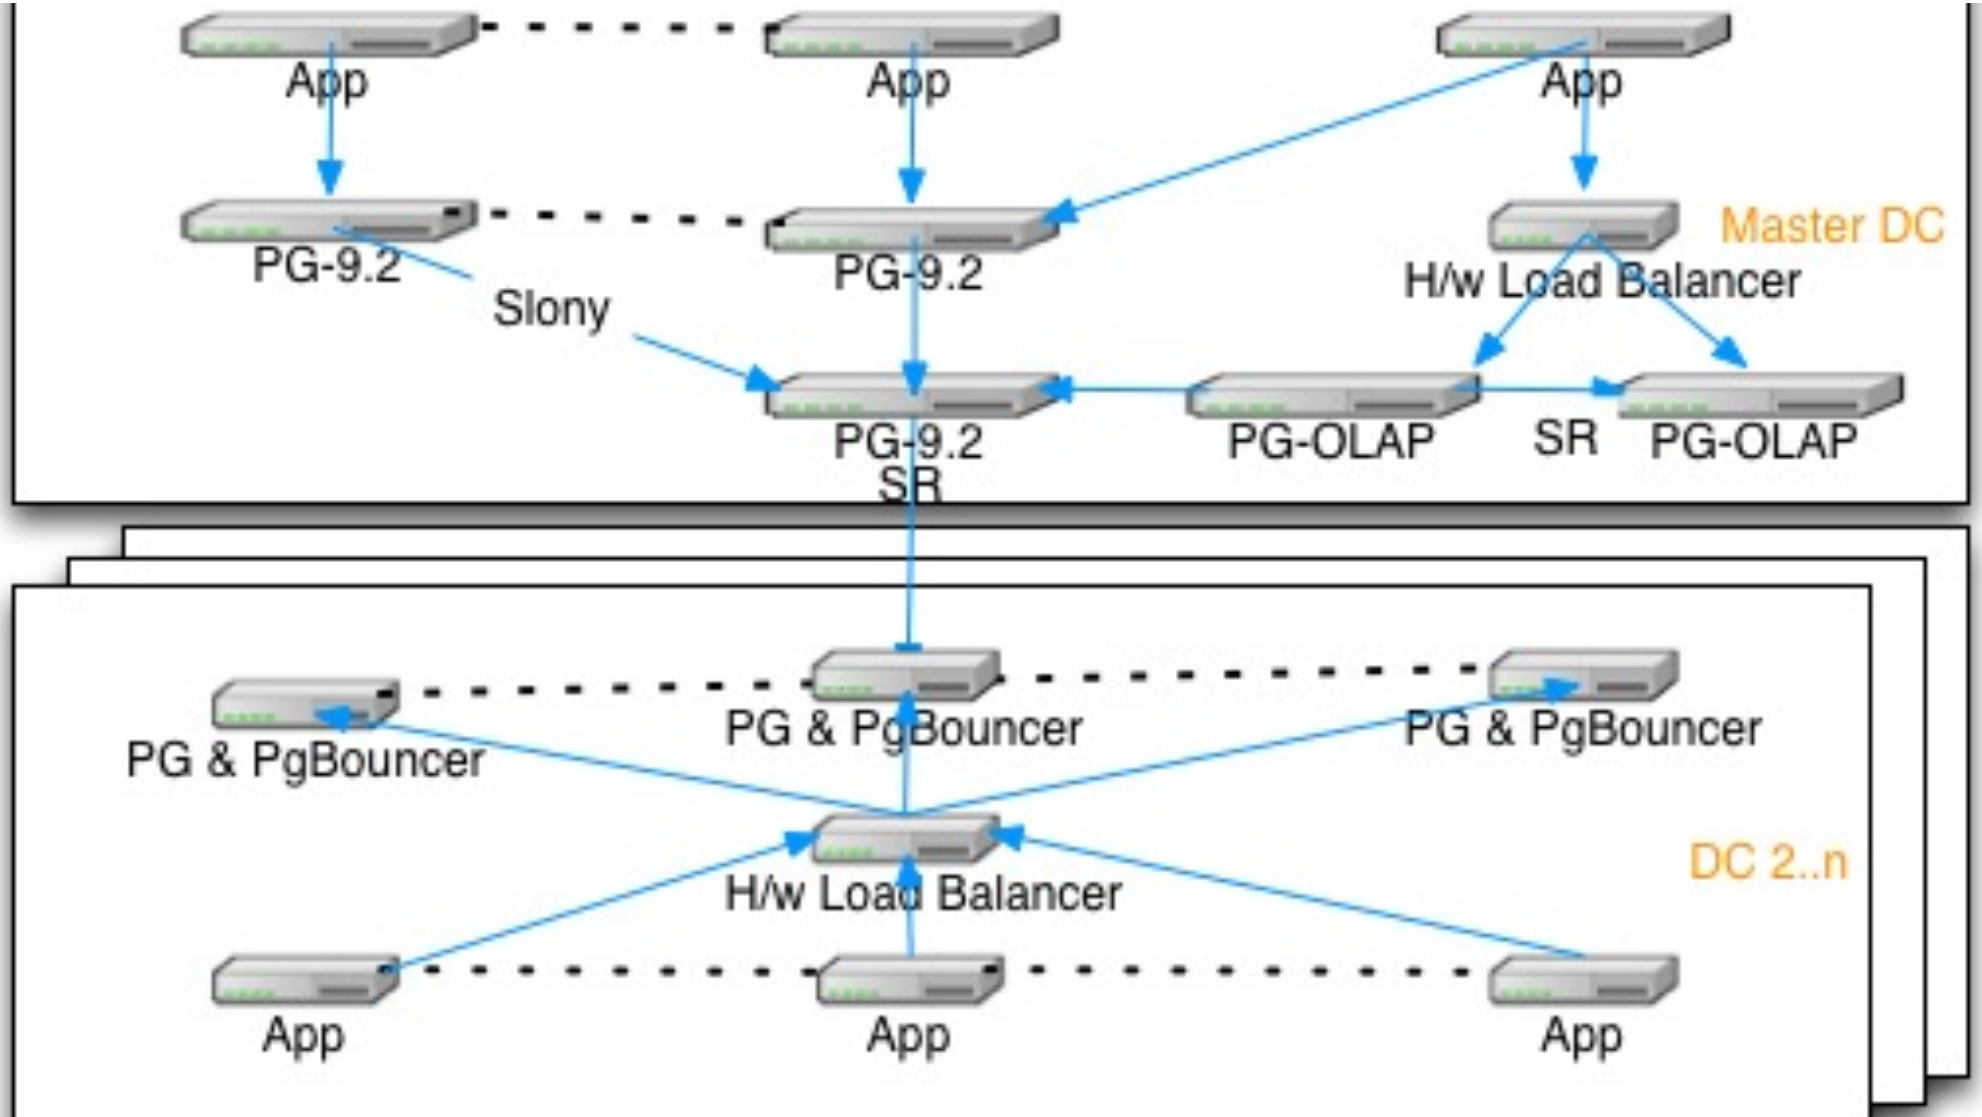

PostgreSQL Architecture Evolution

- Proof Of Concept
- Separation of OLTP & OLAP
- Initial Growth
- New Data Center Expansion
- Service Oriented Architecture
- Operability Issues
- Global reach
- Adapting from other teams
- Cascaded Streaming Replication
- Amazon Web Services
- New Products launch

Proof Of Concept

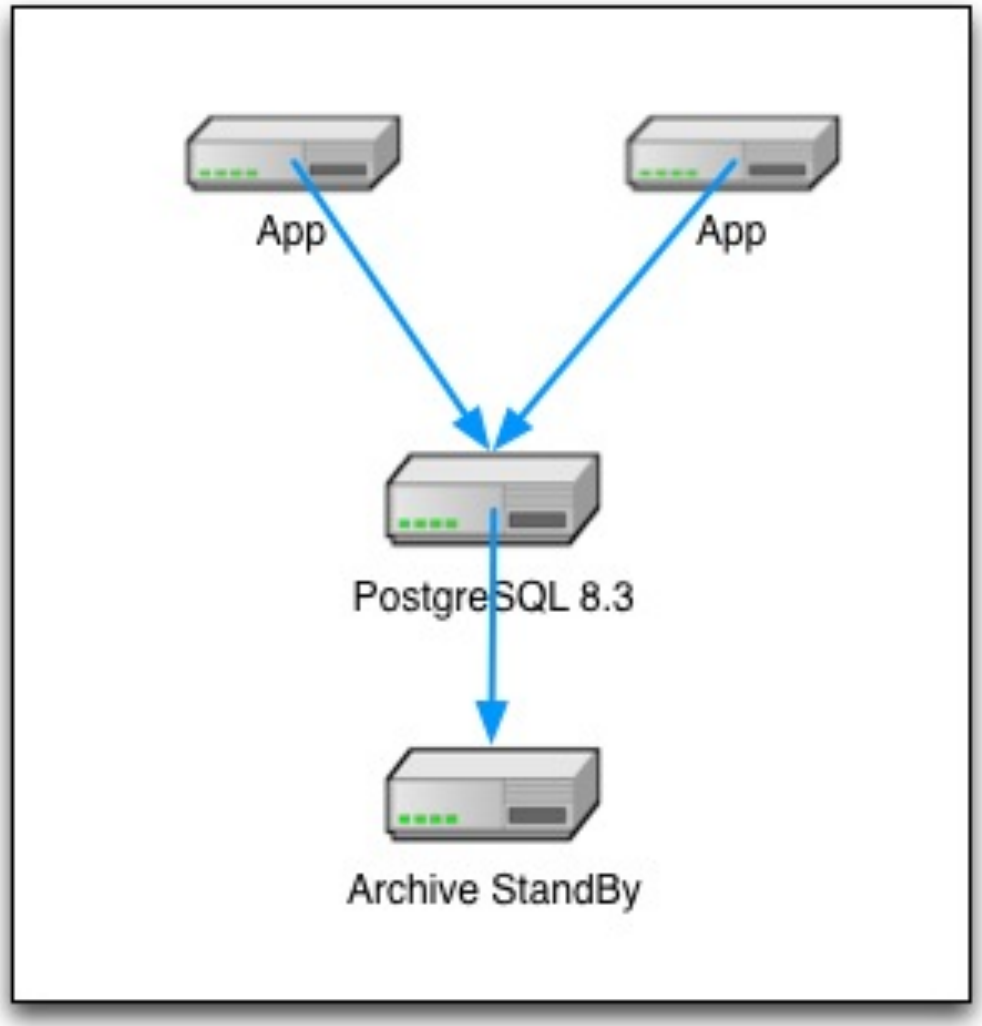

- PostgreSQL Version 8.3
- Typical Startup Architecture
- OLTP Database
- OLAP Database
- Processing Node
- Slave in Archive Mode
- It worked ☺
- Albeit slowness at predictable intervals

Separation of OLTP & OLAP

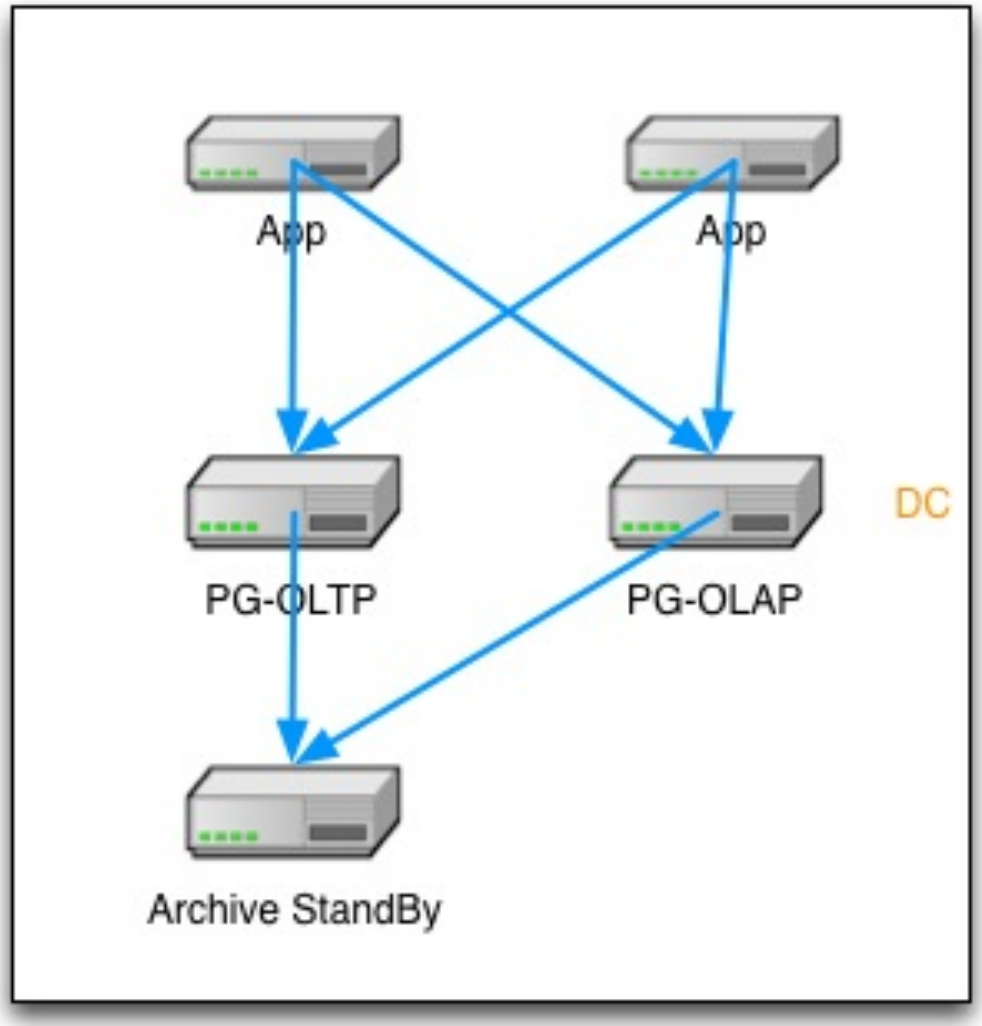

- No more slowness at predictable intervals
- Separation time consuming due to Application changes
- Hourly dump/restore of selected metadata tables
- Made it work ☺ unless
- Issues around missing metadata for some interval
- Wanted to Add more servers for frontend App due to business growth

Initial Growth

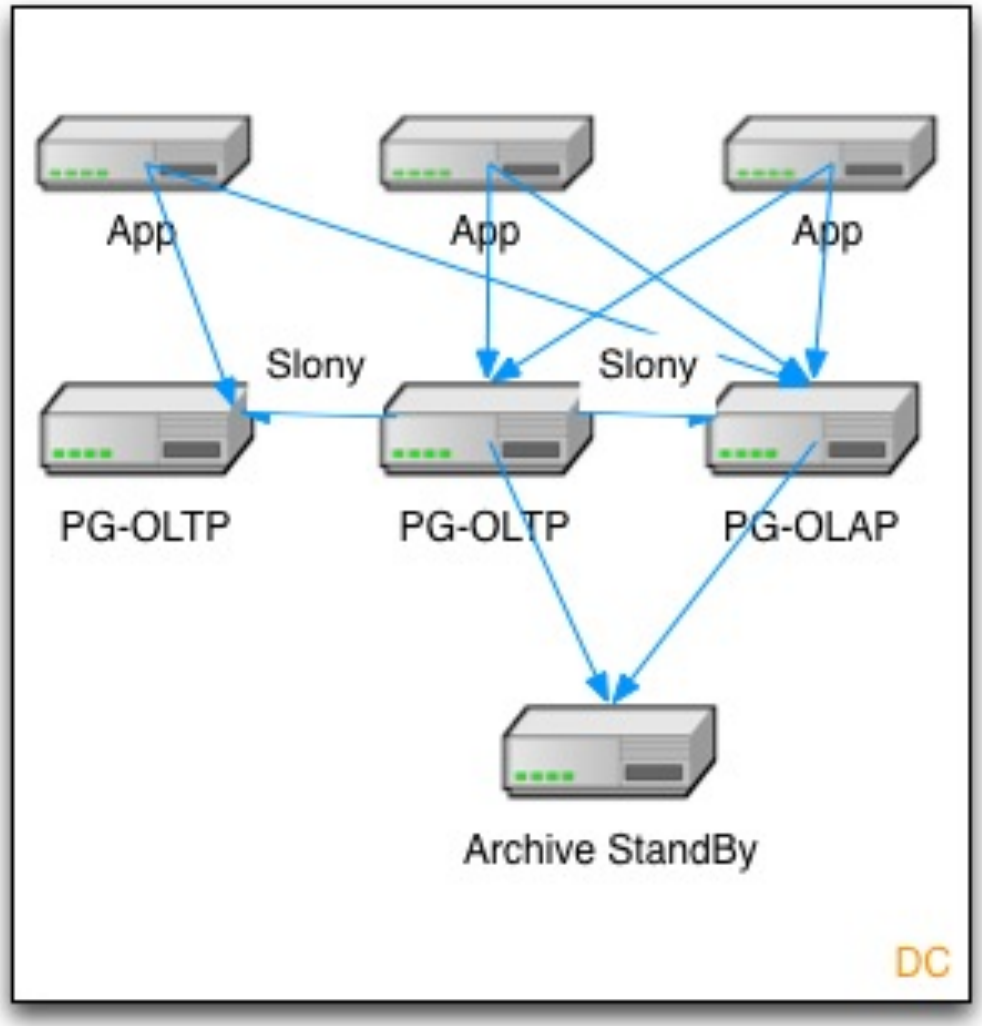

- Slony solved the problem for us
- Training for the team with operational aspects of Slony.
- Introduced database release policy
- Once again made it work ☺ unless
- Business growth plan by entering US market

New Data Center

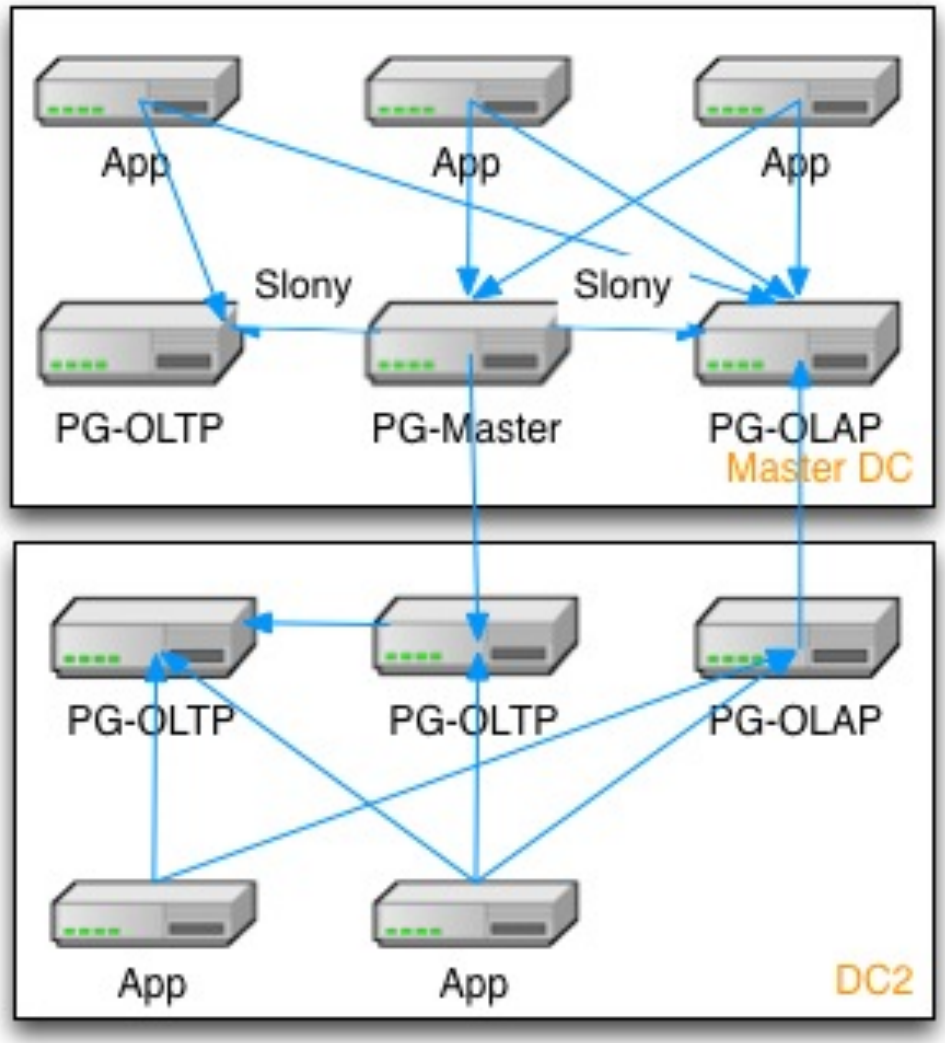

- Easy expansion by adding new nodes in Slony.
- Entered the world of WAN Replication
- Partied Harder ☺ unless
- We experienced issues of WAN Replication
- Maintenance became harder due to Slony & continuous read from apps.
- Releases required more co-ordination among multiple teams

Service Oriented Architecture

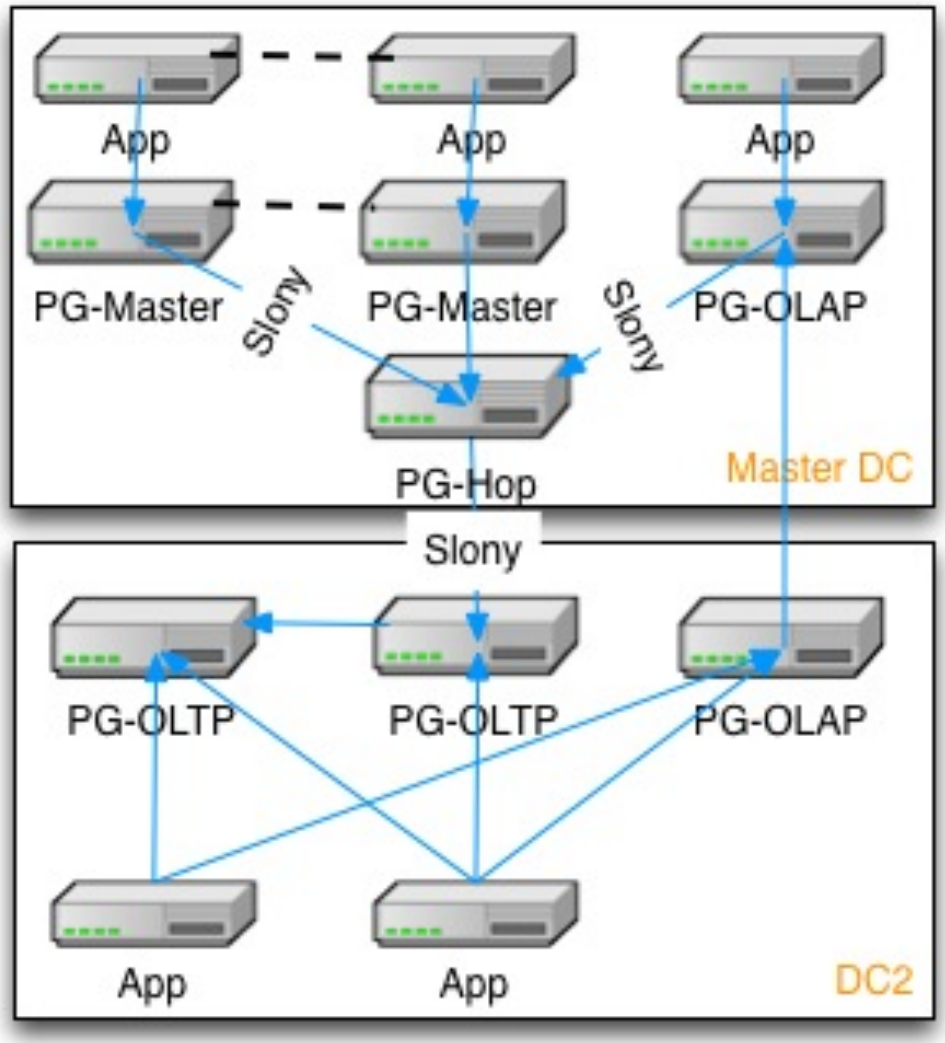

- Monolithic to Micro-lithic database
- Added more masters in Slony cluster
- Increased the number of Slony nodes
- Drastically reduced the release co-ordination
- Tried to minimize network instability by various mechanisms
- In addition wanted to expand for global reach

Operability Issues

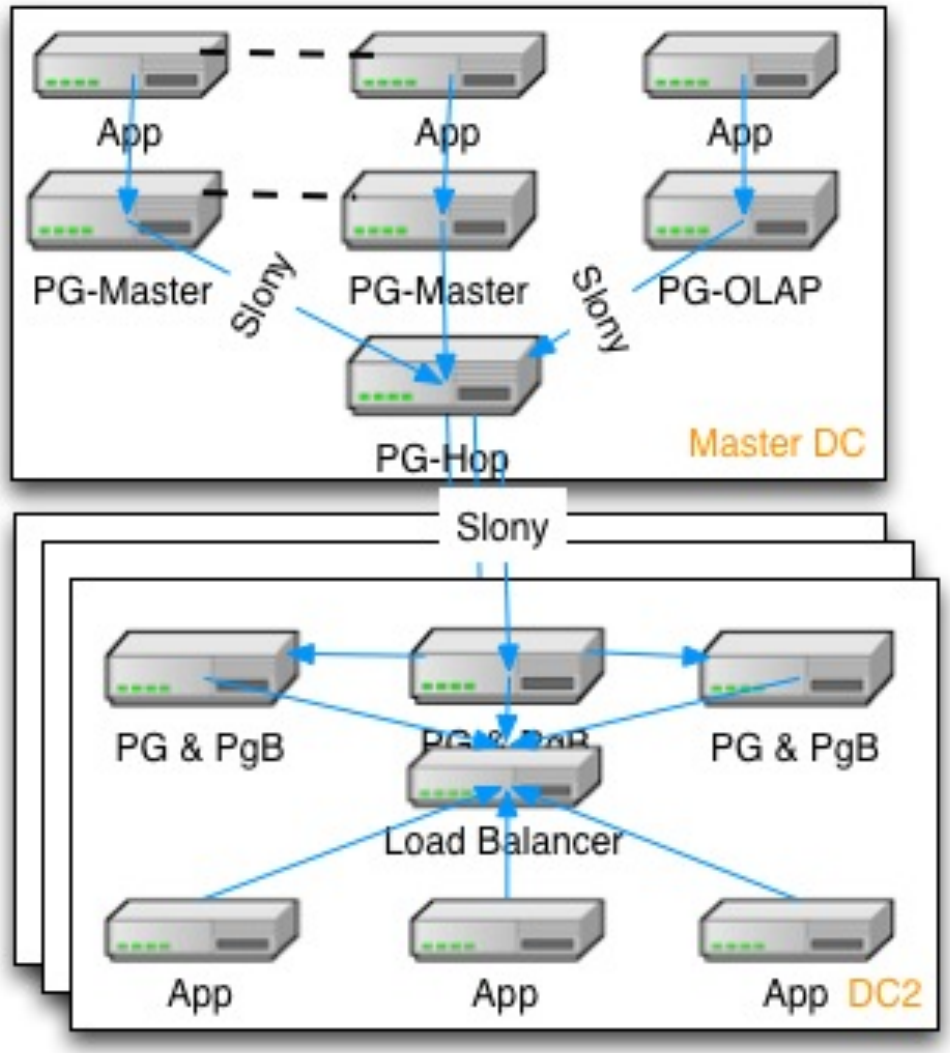

- Addition of even more nodes in Slony
- Got influenced by stateless Apps
- Created VIP in LB to distribute traffic between replicas
- To reduce connection overhead on-boarded pgbouncer
- Blueprint for Read Scalability ☺
- Moved processing from DB nodes to Hadoop

Streaming Replication

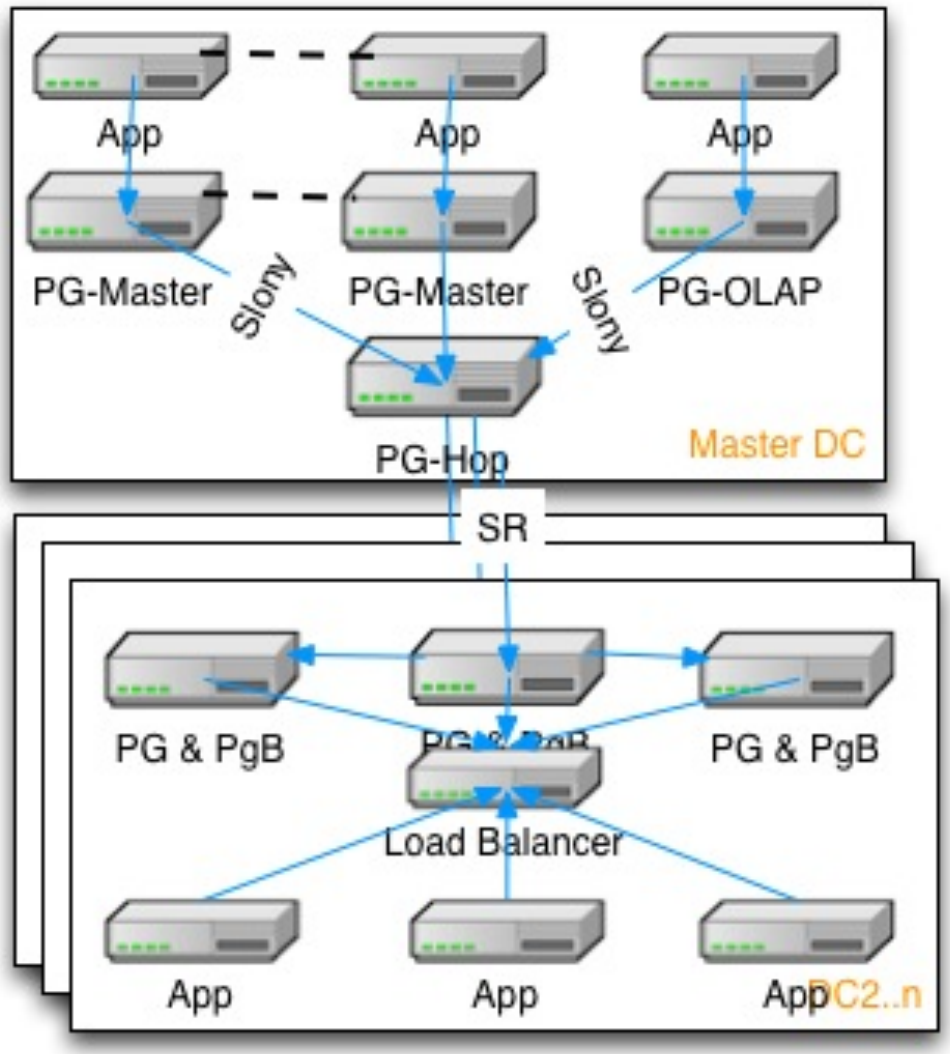

- Upgraded our database to PostgreSQL 9.2
- Migrated from Slony to Streaming Replication for WAN Replication
- Maintenance/Deployments was more efficient
- Delayed Replication for business requirement analysis
- Few quirks with query cancellation

Amazon Web Services

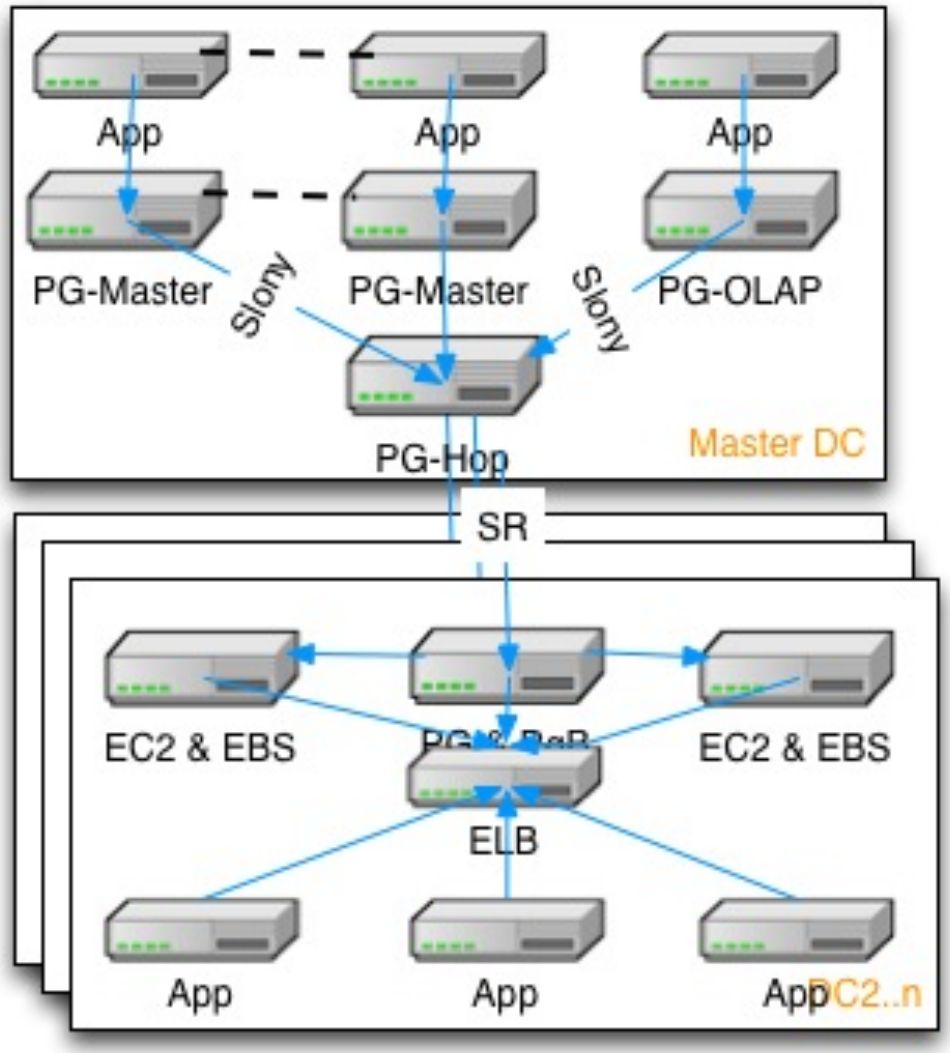

- Embraced AWS for elasticity with respect to computing
- Failure of instances is a norm in AWS
- Automation coupled with Streaming Replication saved us
- ELB distributed the traffic between DB nodes

Current Challenges

- WAL file counts in order of 800/hour
- Write Scalability for Master nodes
- Major Version Upgrades
- FDW for reliable data movement
- Compression and Query Sharding for OLAP systems
- Incremental backup
- Query Visibility in terms of resources consumed
- Priority Queries
- Replication Priority

Thank you.

Sivakumar Krishnamurthy

Database & Production Engineering Manager

sivakumar.krishnamurthy@inmobi.com