

Databases In the Cloud

Overcoming the challenges

- **Rob Masson – ScaleArc, Manager, Solution Architecture**
- **Jon Tobin – Percona, Director, Solution Engineering**



PERCONA

What's Coming

- Jon
 - Enabling the Organization
 - Advantages of Leveraging “the cloud”
 - Who are the Big Players?
 - I(nfrastructure)aaS vs P(latform)aaS
 - Starting Small
- Rob
 - Cross Region Failover
 - Hybrid Deployments
 - Best Practices & Lessons Learned

“Cloud” – 5 Characteristics

According to National Institute of Standards & Technology

1. On-demand self-service
 2. Broad network access
 3. Resource pooling
 4. Rapid elasticity
 5. Measured service
-
6. Perimeter contained

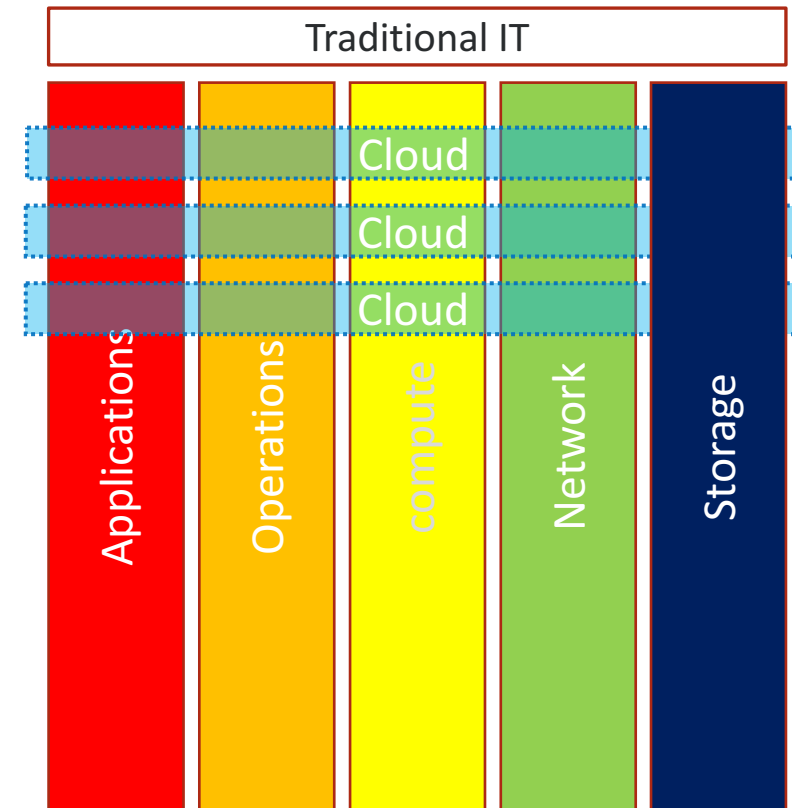
“the cloud” – organizational migration

High Value

- Service oriented
 - Deep understanding of ingredients
- Aligned to business need
 - Not component need
- Linked to business outcomes

Low Value

- Quick provisioning
- Reduced management



Compelling Advantages

- Time to provision
 - Months to minutes
- Geographic distribution
 - Resources available “anywhere”
- Operating vs Capital Expenditure
 - Recurring billing
- Standardization
 - Side effect
- Simplicity

Drawbacks

- Noisy neighbor
- Rigidity
 - [Netflix Outage Christmas Eve 2012](#)
- Resource location
 - Actually, where is that server?
- Regulatory/Security
- Skillset Atrophy
- Lock-in
- Troubleshooting
- Complacency

What are your options?

- **AWS**
 - Market leader
 - Staggering product mix
 - Solid service offerings
- **Azure**
 - Microsoft's reach
 - Enterprise focus
- **Google**
 - Developer focused
 - Fewer, focused options

Azure

- Enterprise focused
- Compliance
 - FIPS 140-2, PCI-DSS (L1), HIPAA/HITECH, etc
- Windows & Linux
- 38 regions
 - Almost every continent
- 99.99% uptime SLAs
- DBaaS: SQL (MSSQL based), NoSQL (Doc, KV), Data Factory (ETL)

AWS

- RDS: MySQL, MariaDB, MSSQL, Postgres, Oracle, Aurora
 - Most feature rich DBaaS
 - *Auto-backups*
 - *Auto-failover*
 - Large traction amongst Percona's user base
- Semi-compliant
- AWS Data Migration Service
 - Homo/heterogenous
- AWS Marketplace
 - Thousands of software solutions

Google Cloud Platform

- Developer focused
 - IaaS, SaaS, PaaS
 - CaaS
- Compliant
 - SSAE 16, ISO 27001, PCI DSS, HIPAA
- Database
 - CloudSQL
 - Cloud Bigtable
 - Cloud Datastore

IaaS

Upside

- Flexible
- Tunable

Downside

- Configuration
- Management @ scale
- Monitoring
- Scaling

AWS Monthly EC2 Reserved Instance Cost
EC2 db.r3large = \$76.75

PaaS

Upside

- Easy to use
- Full featured
- Forgettable
- Great for sweet spot

Downside

- Hard to tune
- Rigid
- Tough to scale
- Forgettable
- Premium cost

AWS Monthly RDS Reserved Instance Cost

RDS db.r3large = \$138.70*

45% more than EC2

Cloud Success

- Service Synergy
 - IaaS is nice, but not valuable
 - Leverage services whenever possible
- Cloud is not a cost reduction strategy
 - Make sure you identify hidden costs
- Know the limits/tradeoffs
 - Test infrastructure & assumptions
- Monitor/alert to those limits
- Avoid integrating app w/infrastructure

Join us at Percona Live

When: April 24-27, 2017

Where: Santa Clara, CA, USA

The Percona Live Open Source Database Conference is a great event for users of any level using open source database technologies.

- Get briefed on the hottest topics
- Learn about building and maintaining high-performing deployments
- Listen to technical experts and top industry leaders

Use promo code “WebinarPL” to save an extra 15% off. Register now and get the early bird rate, but hurry prices go up Jan 31th.

<https://www.percona.com/live/17/register>

Sponsorship opportunities available as well:

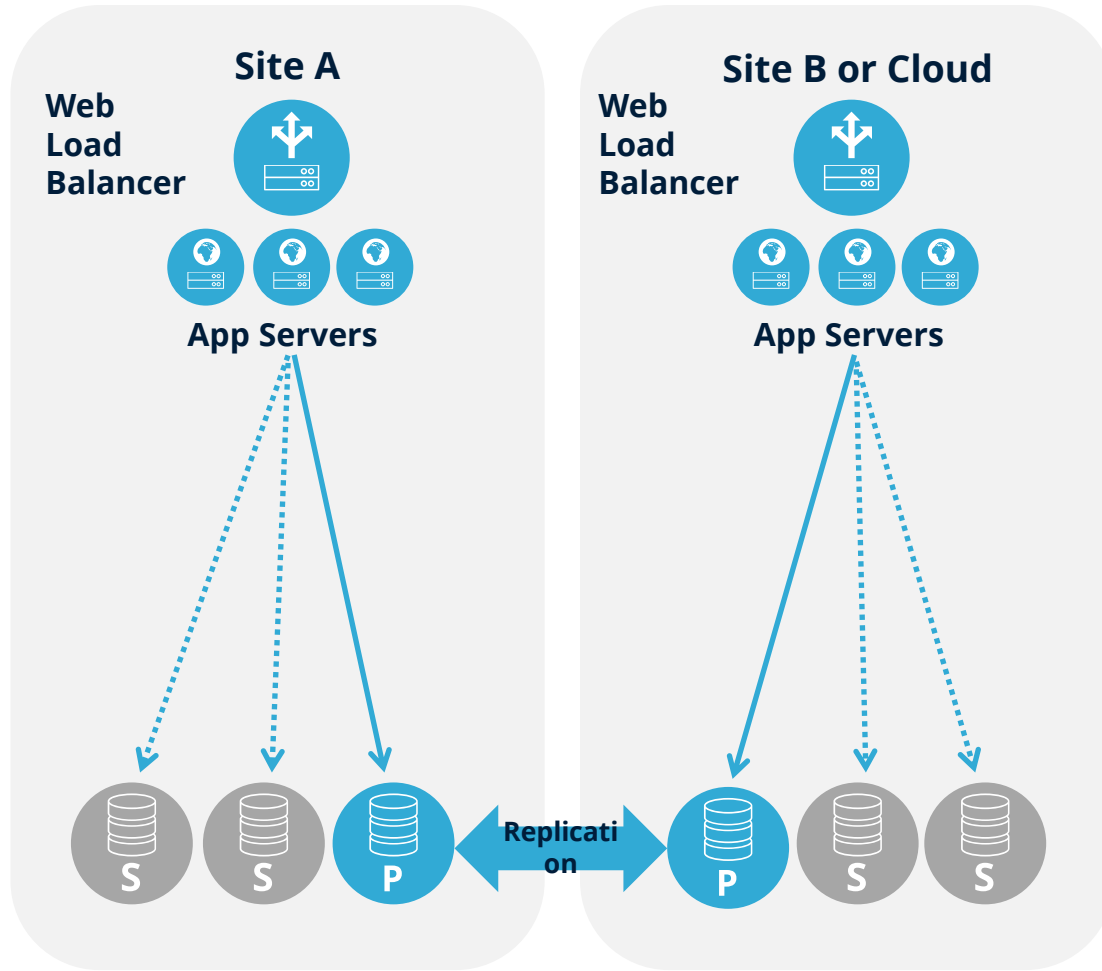
<https://www.percona.com/live/17/be-a-sponsor>



PERCONA
LIVE

Active / Active DCs – Cross Geo, Cloud, Hybrid

Global Traffic Manager



Advantages

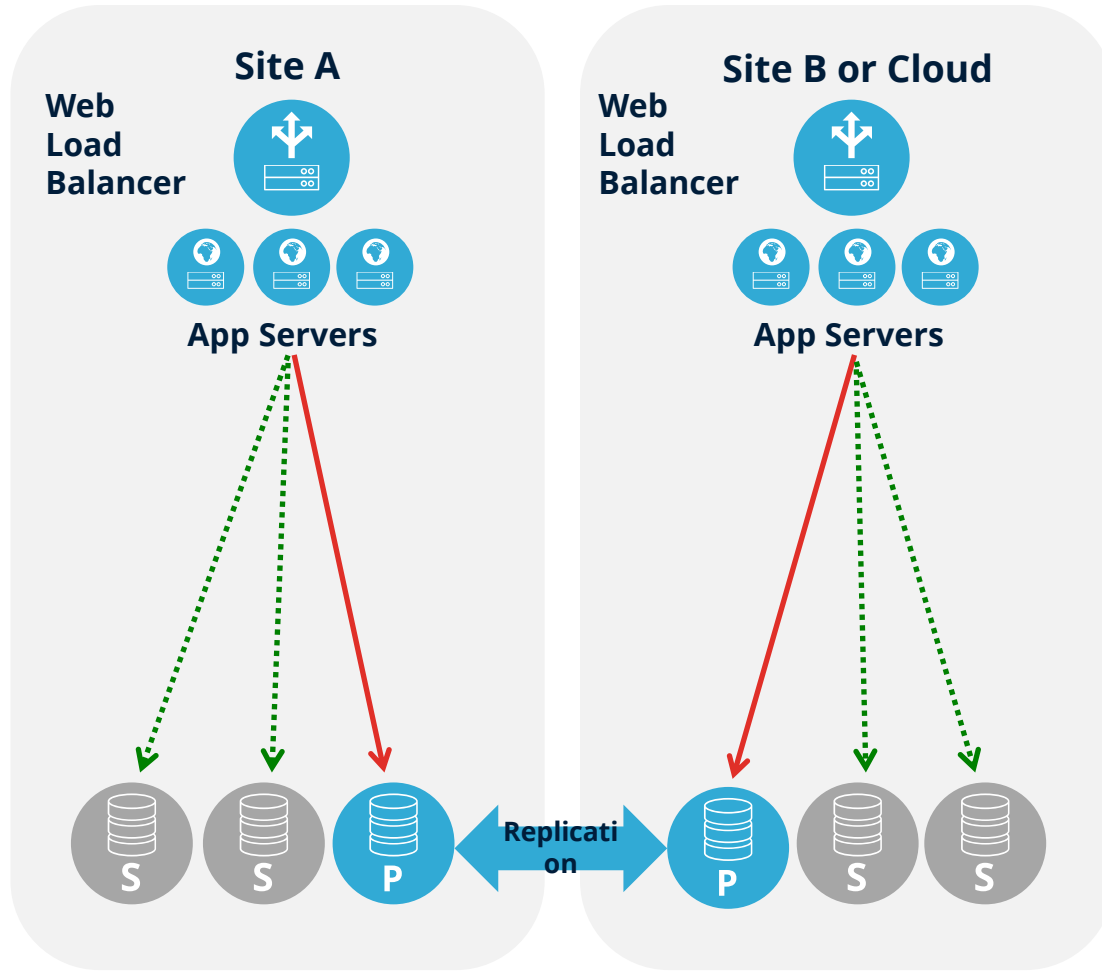
- Geo Load Balancing
- High Availability
- Resilience to Load Spikes

Challenges

- Resolving Writes
- Stale Reads
- Failover and Failback

Active / Active DCs – Strategy 1 : Modify your Apps

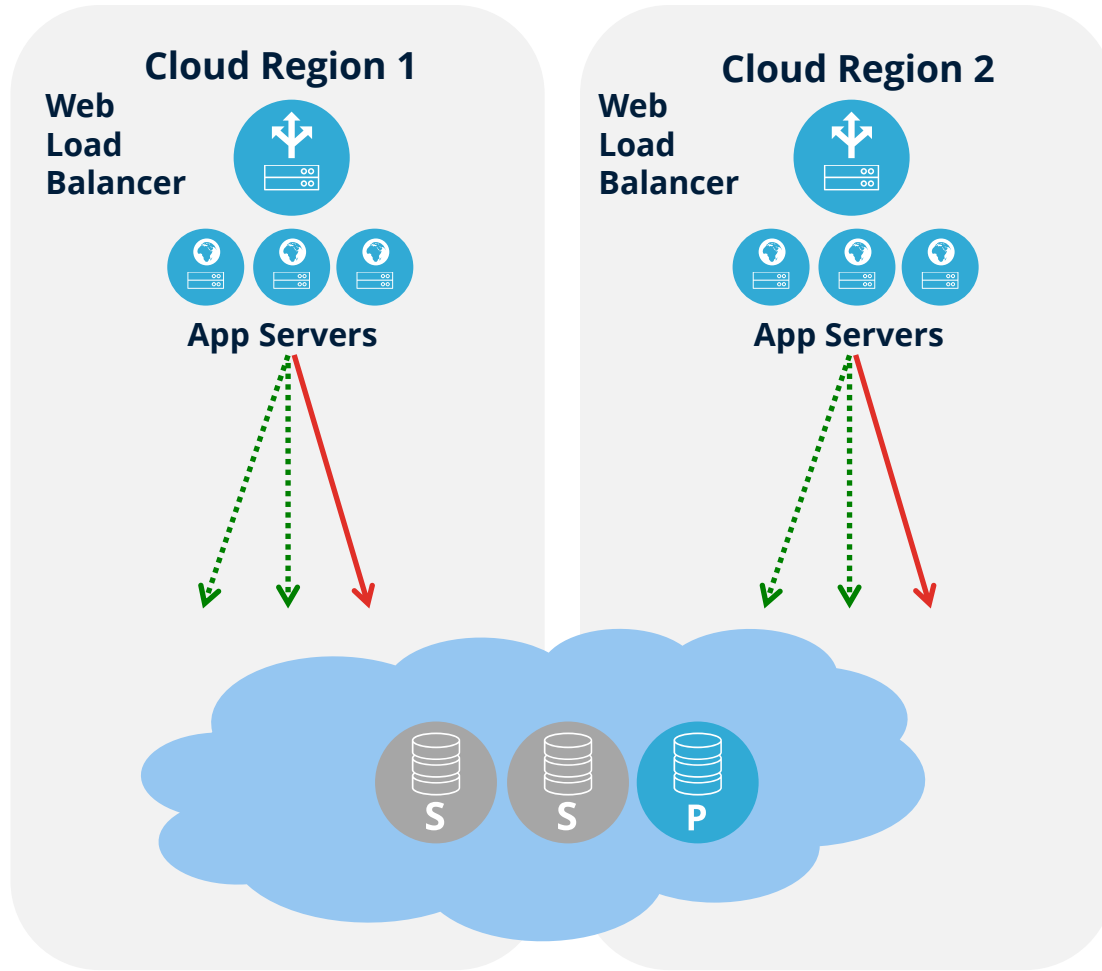
Global Traffic Manager



- 1 Apps will need to understand where to send Writes and where to send Reads
- 2 No inherent Load Balancing
- 3 No inherent Replication awareness
- 4 Does not address Failover or Failback

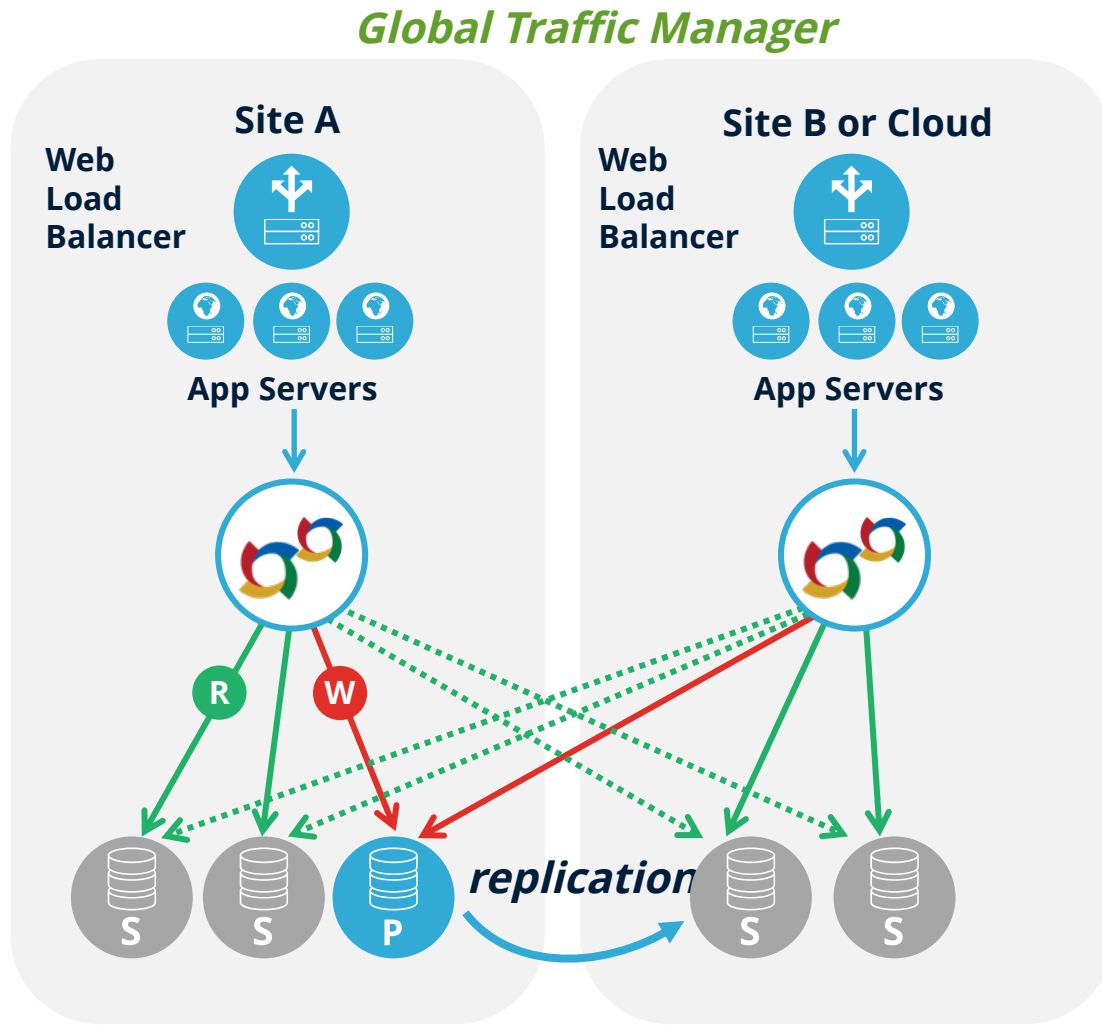
Active / Active DCs – Strategy 2 : PaaS

Global Traffic Manager



- 1 Apps will need to understand where to send Writes and where to send Reads
- 2 No inherent Load Balancing
- 3 No inherent Replication awareness
- 4 Cannot be used for a Hybrid Implementation

Active / Active DCs – Strategy 3: DB Proxy Layer



- 1 ScaleArc performs read/write split and geo-load balance database traffic
- 2 Best performance achieved by Time To First Byte
- 3 Active replication lag monitoring; avoids sending traffic to lagging nodes
- 4 Database auto failover is enabled within and across datacenters

Best Practices



Make your Applications Agnostic to the Implementation Details



Automate the processes as much as possible!!!



Spend Time on the Important Stuff (Security, Migration, Ops)



Be prepared for the Sprawl.....



Don't forget to backup everything!!!!!



Avoid Vendor Lock-In