Databases in the Hosted Cloud

Colin Charles, Chief Evangelist, Percona Inc.

colin.charles@percona.com / byte@bytebot.net

http://bytebot.net/blog/ | @bytebot on Twitter

PERCONA

whoami

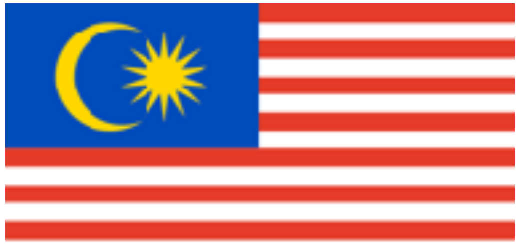

- Chief Evangelist, Percona Inc
    - Focusing on the MySQL ecosystem (MySQL, Percona Server, MariaDB Server), as well as the MongoDB ecosystem (Percona Server for MongoDB) + **100% open source** tools from Percona like Percona Monitoring & Management, Percona xtrabackup, Percona Toolkit, etc.
- Founding team of MariaDB Server (2009-2016), previously at Monty Program Ab, merged with SkySQL Ab, now MariaDB Corporation
- Formerly MySQL AB (exit: Sun Microsystems)
- Past lives include Fedora Project (FESCO), OpenOffice.org
- MySQL Community Contributor of the Year Award winner 2014

Agenda

- MySQL as a service offering (DBaaS)
- Choices
- Considerations
- MySQL versions & access
- Costs
- Deeper into RDS
- Should you run this on EC2 or an equivalent?
- Conclusion

Database as a service

- Database as a Service (DBaaS)
- MySQL/PostgreSQL available on-demand, without any installation/configuration of hardware/software
- Pay-per-usage based
- Provider maintains database, you don’t maintain, upgrade, or administer the database

New way of deployment

- Enter a credit card number
- call API (or use the GUI)

```
ec2-run-instances ami-xxx -k ${EC2_KEYPAIR} -t m1.large
nova boot --image centos6-x86_64 --flavor m1.large db1
```

http://www.flickr.com/photos/68751915@N05/6280507539/

Why DBaaS?

- “Couldn’t we just have a few more servers to handle the traffic spike during the elections?”
- Don’t have a lot of DBAs, optimise for operational ease
- Rapid deployment & scale-out

Your choices today

- Amazon Web Services Relational Database Service (RDS)
- Rackspace Cloud Databases
- Google Cloud SQL
- Oracle Cloud MySQL Service
- Microsoft Azure Database

There are more

- Jelastic - PaaS offering MySQL, MariaDB Server
- ClearDB - MySQL partnered with Heroku, Azure clouds
- Joyent - Image offers Percona MySQL and a Percona SmartMachine

Some new entrants

- Google Compute Engine offers Percona XtraDB Cluster as a “click-to-deploy” app
  - comes with Galera 3, Percona Toolkit, XtraBackup as well
- Pivotal CloudFoundry - “MySQL” PaaS which is MariaDB Galera Cluster 10
- Red Hat OpenShift
  - MySQL 5.5/5.6/5.7, MariaDB 10.0/10.1
  - PostgreSQL 9.2/9.4/9.5
  - MongoDB 2.4/2.6/3.2

Beware

- GenieDB - globally distributed MySQL as a service, master-master replication, works on EC2, Rackspace, Google Compute Engine, HP Cloud
- Xeround - 2 weeks notice...

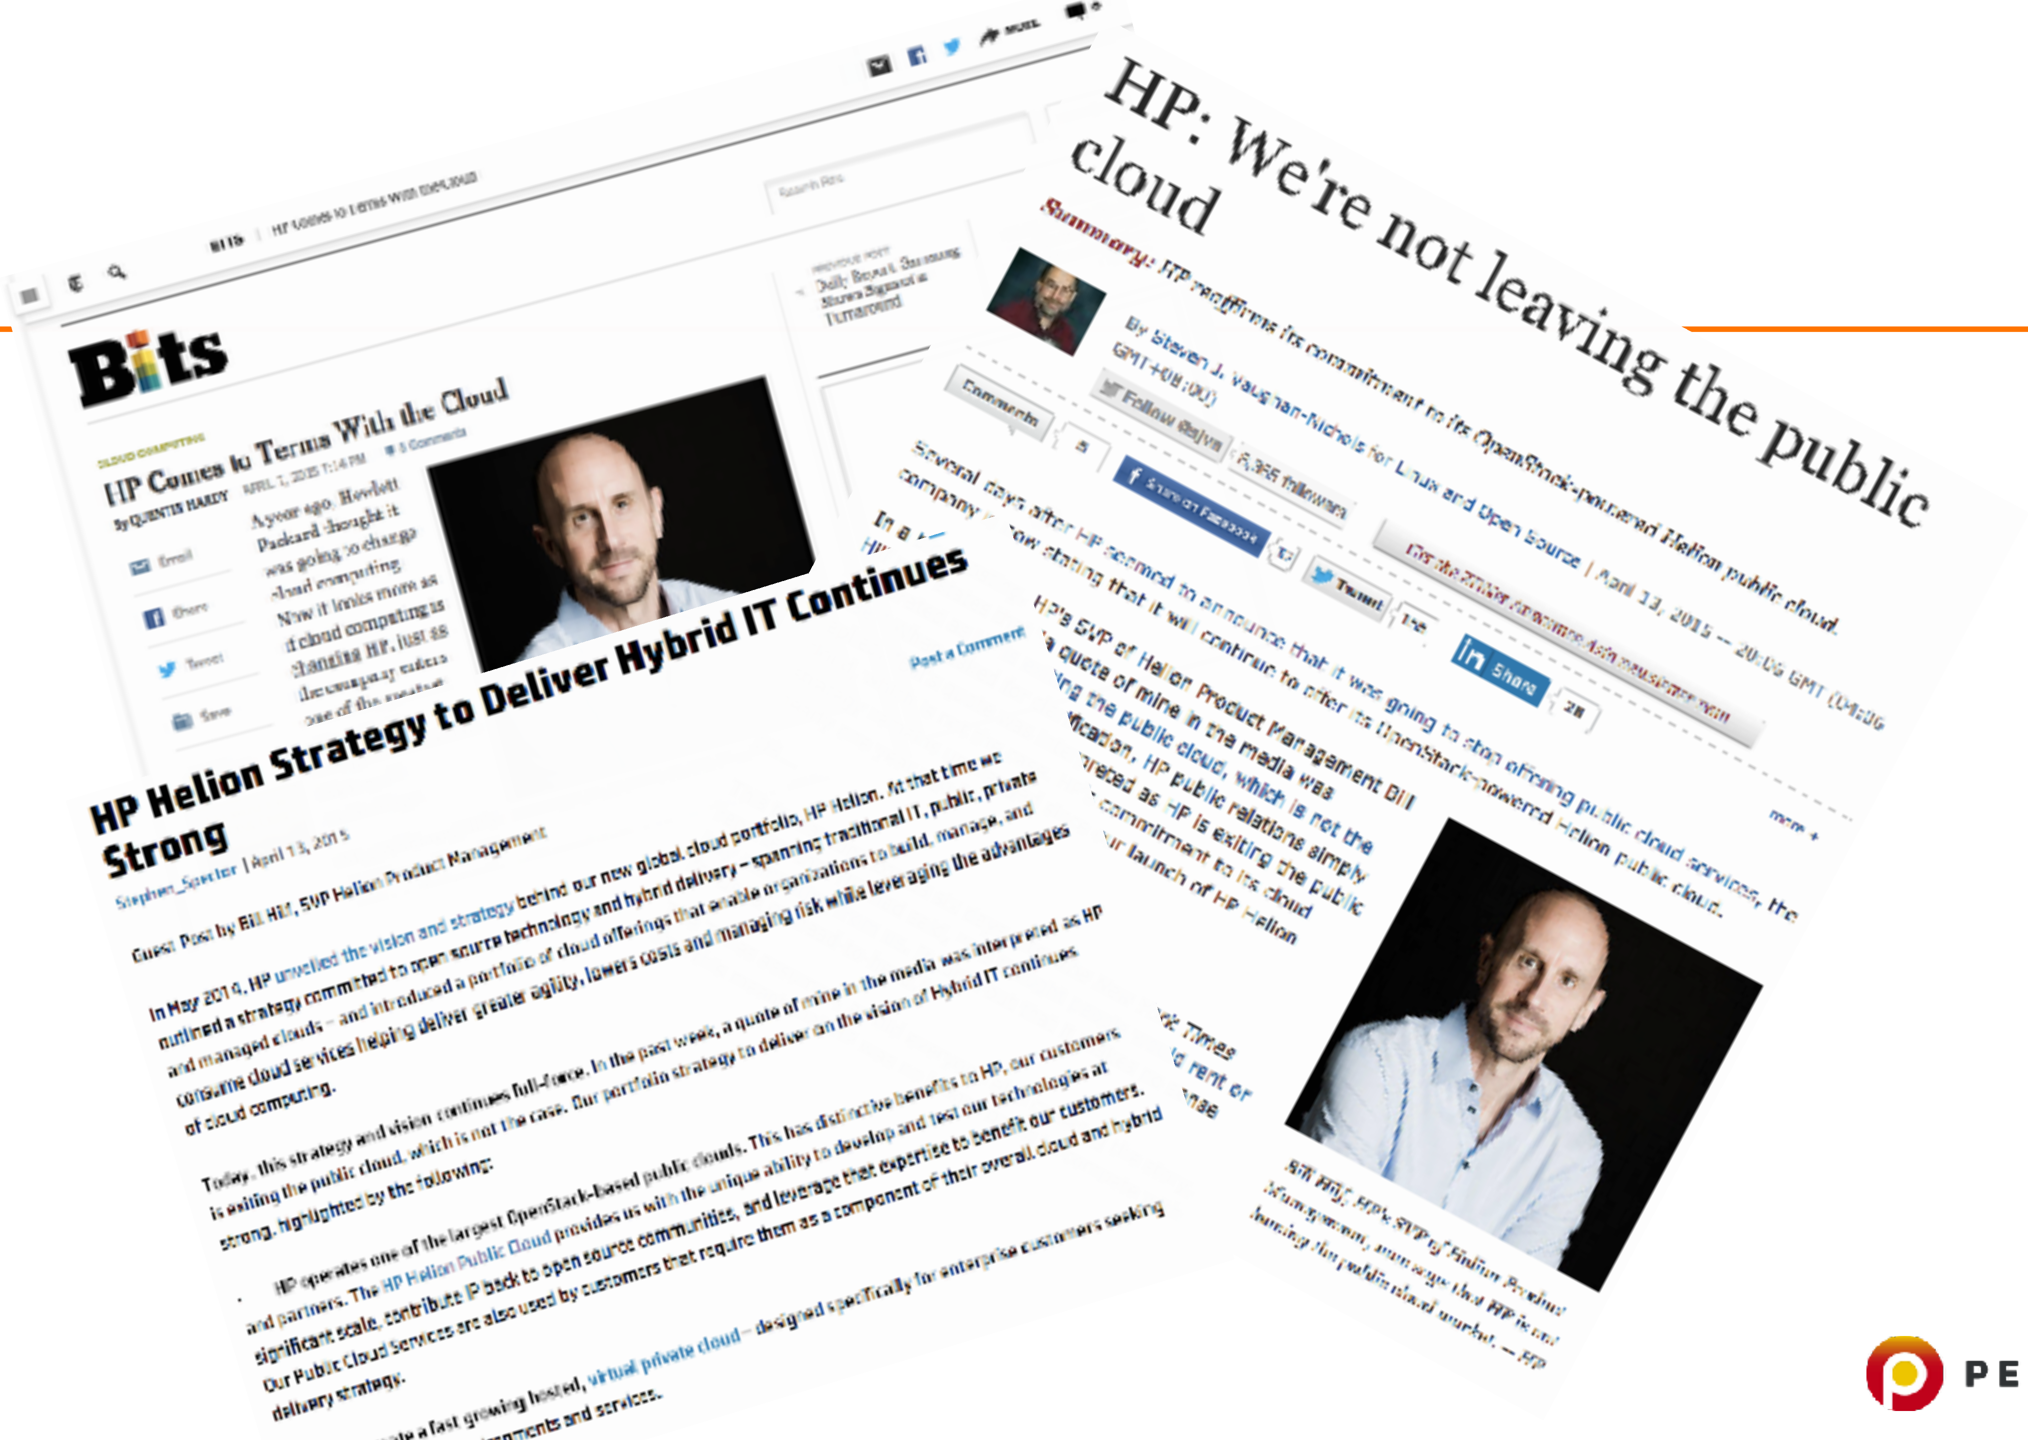

Bits | HP Comes to Terms With the Cloud

Search Bits

PREVIOUS POST
Daily Report: Samsung Shows Signs of a Turnaround

Bits

CLOUD COMPUTING

HP Comes to Terms With the Cloud

By QUENTIN HARDY APRIL 7, 2015 7:14 PM 5 Comments

A year ago, Hewlett Packard thought it was going to change cloud computing. Now it looks more as if cloud computing is changing HP, just as the company enters one of the greatest

HP Helion Strategy to Deliver Hybrid IT Continues Strong

Post a Comment

Stephen_Spector | April 13, 2015

Guest Post by Bill Hilf, SVP Helion Product Management

In May 2014, HP unveiled the vision and strategy behind our new global cloud portfolio, HP Helion. At that time we outlined a strategy committed to open source technology and hybrid delivery – spanning traditional IT, public, private and managed clouds – and introduced a portfolio of cloud offerings that enable organizations to build, manage, and consume cloud services helping deliver greater agility, lowers costs and managing risk while leveraging the advantages of cloud computing.

Today, this strategy and vision continues full-force. In the past week, a quote of mine in the media was interpreted as HP is exiting the public cloud, which is not the case. Our portfolio strategy to deliver on the vision of Hybrid IT continues strong, highlighted by the following:

- HP operates one of the largest OpenStack-based public clouds. This has distinctive benefits to HP, our customers and partners. The HP Helion Public Cloud provides us with the unique ability to develop and test our technologies at significant scale, contribute IP back to open source communities, and leverage that expertise to benefit our customers. Our Public Cloud Services are also used by customers that require them as a component of their overall cloud and hybrid delivery strategy.
- ... a fast growing hosted, virtual private cloud – designed specifically for enterprise customers seeking ... environments and services.

HP: We're not leaving the public cloud

Summary: HP reaffirms its commitment to its OpenStack-powered Helion public cloud.

By Steven J. Vaughan-Nichols for Linux and Open Source | April 13, 2015 -- 20:06 GMT (04:06 GMT+08:00)

Follow @sjvn 6,366 followers

Comments 5 Share on Facebook 10 Tweet 136 Share 28

Create ZDNet Asia newsletter now

Several days after HP seemed to announce that it was going to stop offering public cloud services, the company ... now stating that it will continue to offer its OpenStack-powered Helion public cloud.

more +

In a ... Hi... HP's SVP of Helion Product Management Bill ... a quote of mine in the media was ... ing the public cloud, which is not the ...fication, HP public relations simply ...preted as HP is exiting the public ... commitment to its cloud ...ur launch of HP Helion

...k Times ...d rent or ...nage

Bill Hilf, HP's SVP of Helion Product Management, now says that HP is not leaving the public cloud world. — HP

Thank you HPCloud

- http://www.bytebot.net/blog/archives/2015/10/27/sunsetting-hpcloud-whom-contributed-to-making-mysql-better
- Utility user
- enforce_storage_engine
- prevent LOAD DATA INFILE/SELECT INTO OUTFILE
- restrict # of binlog files

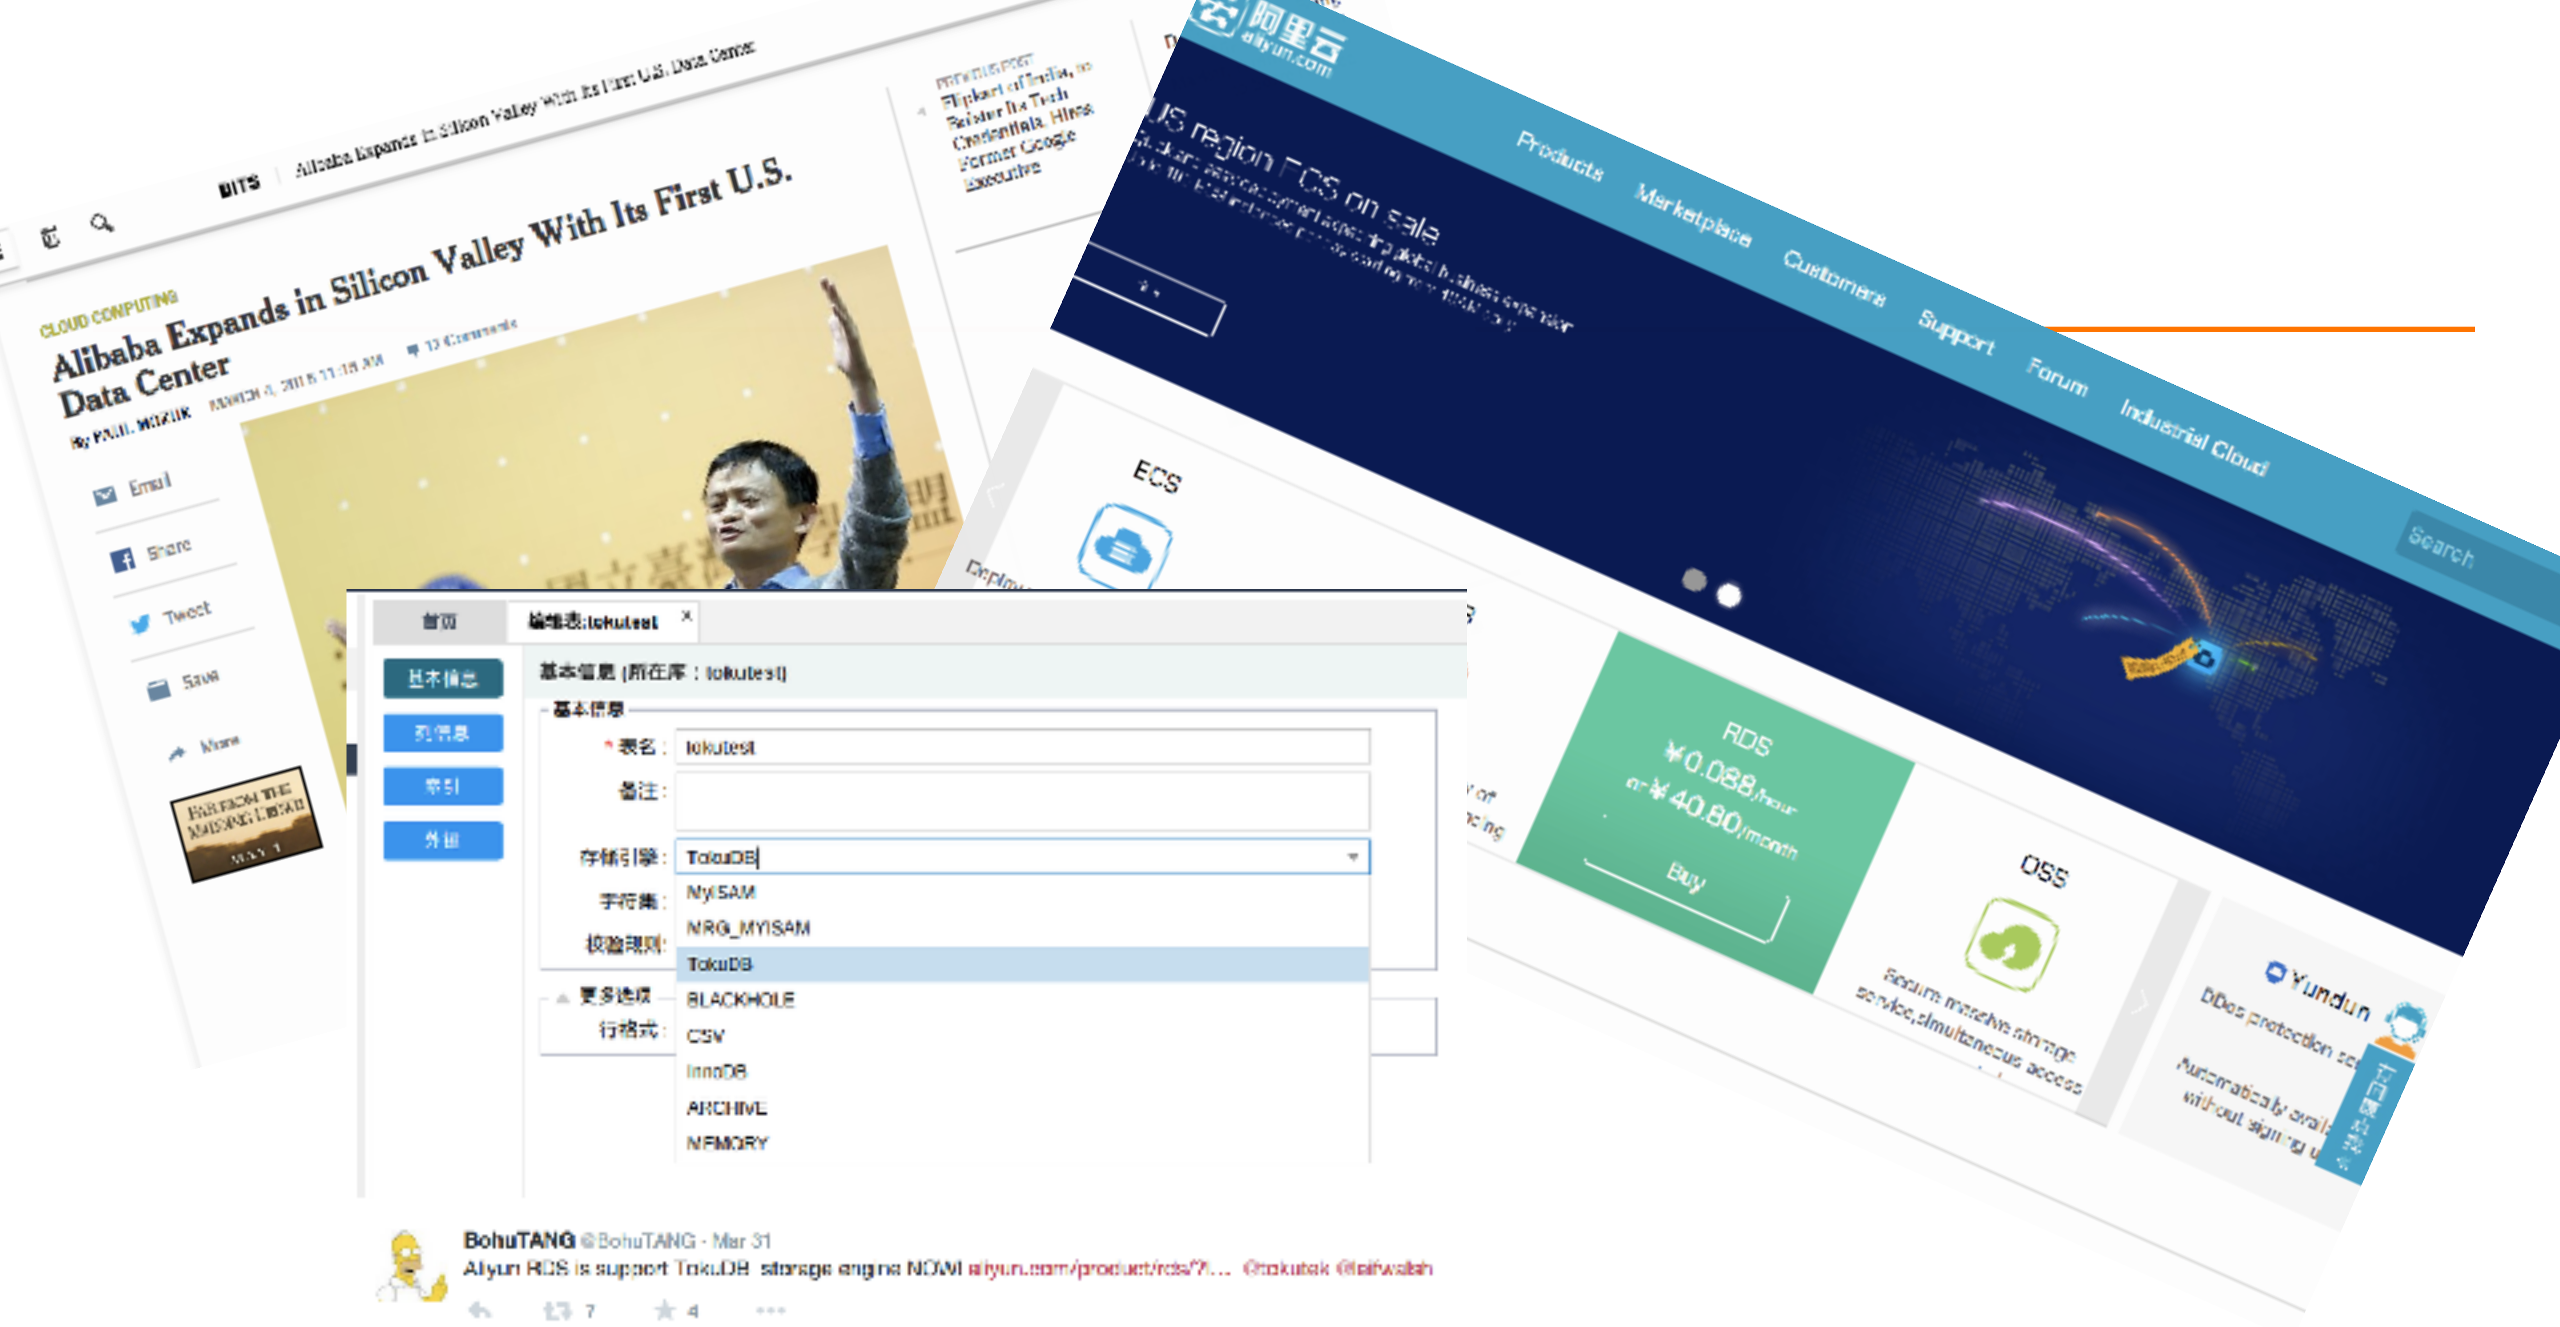

BITS | Alibaba Expands in Silicon Valley With Its First U.S. Data Center

CLOUD COMPUTING

Alibaba Expands in Silicon Valley With Its First U.S. Data Center

By PAUL MOZUR   MARCH 4, 2015 11:15 AM   13 Comments

PREVIOUS POST
Flipkart of India, to Bolster Its Tech Credentials, Hires Former Google Executive

US region ECS on sale

Secure massive storage service, simultaneous access

Yundun
DDos protection se
Automatically avail
without signing u

基本信息 [所在库：tokutest]

基本信息

* 表名：	tokutest
备注：	
存储引擎：	TokuDB
字符集：	
校验规则：	

- MyISAM
- MRG_MYISAM
- TokuDB
- BLACKHOLE
- CSV
- InnoDB
- ARCHIVE
- MEMORY

更多选项

行格式：

BohuTANG @BohuTANG · Mar 31
Aliyun RDS is support TokuDB storage engine NOW! aliyun.com/product/rds/?I… @tokutek @leifwalsh

What else do you get?

- AWS RDS has PostgreSQL
- Most MongoDB users run within the cloud — but not hosted, it tends to be “roll your own” (and lately MongoDB Atlas based - reasonable free tier!)
- Compose.io (now IBM) & the like (MongoLab) — hosted MongoDB, Redis, Enhanced PostgreSQL, ElasticSearch, ~~RethinkDB~~, MySQL with group replication, ScyllaDB, etc.

MariaDB Server on AWS

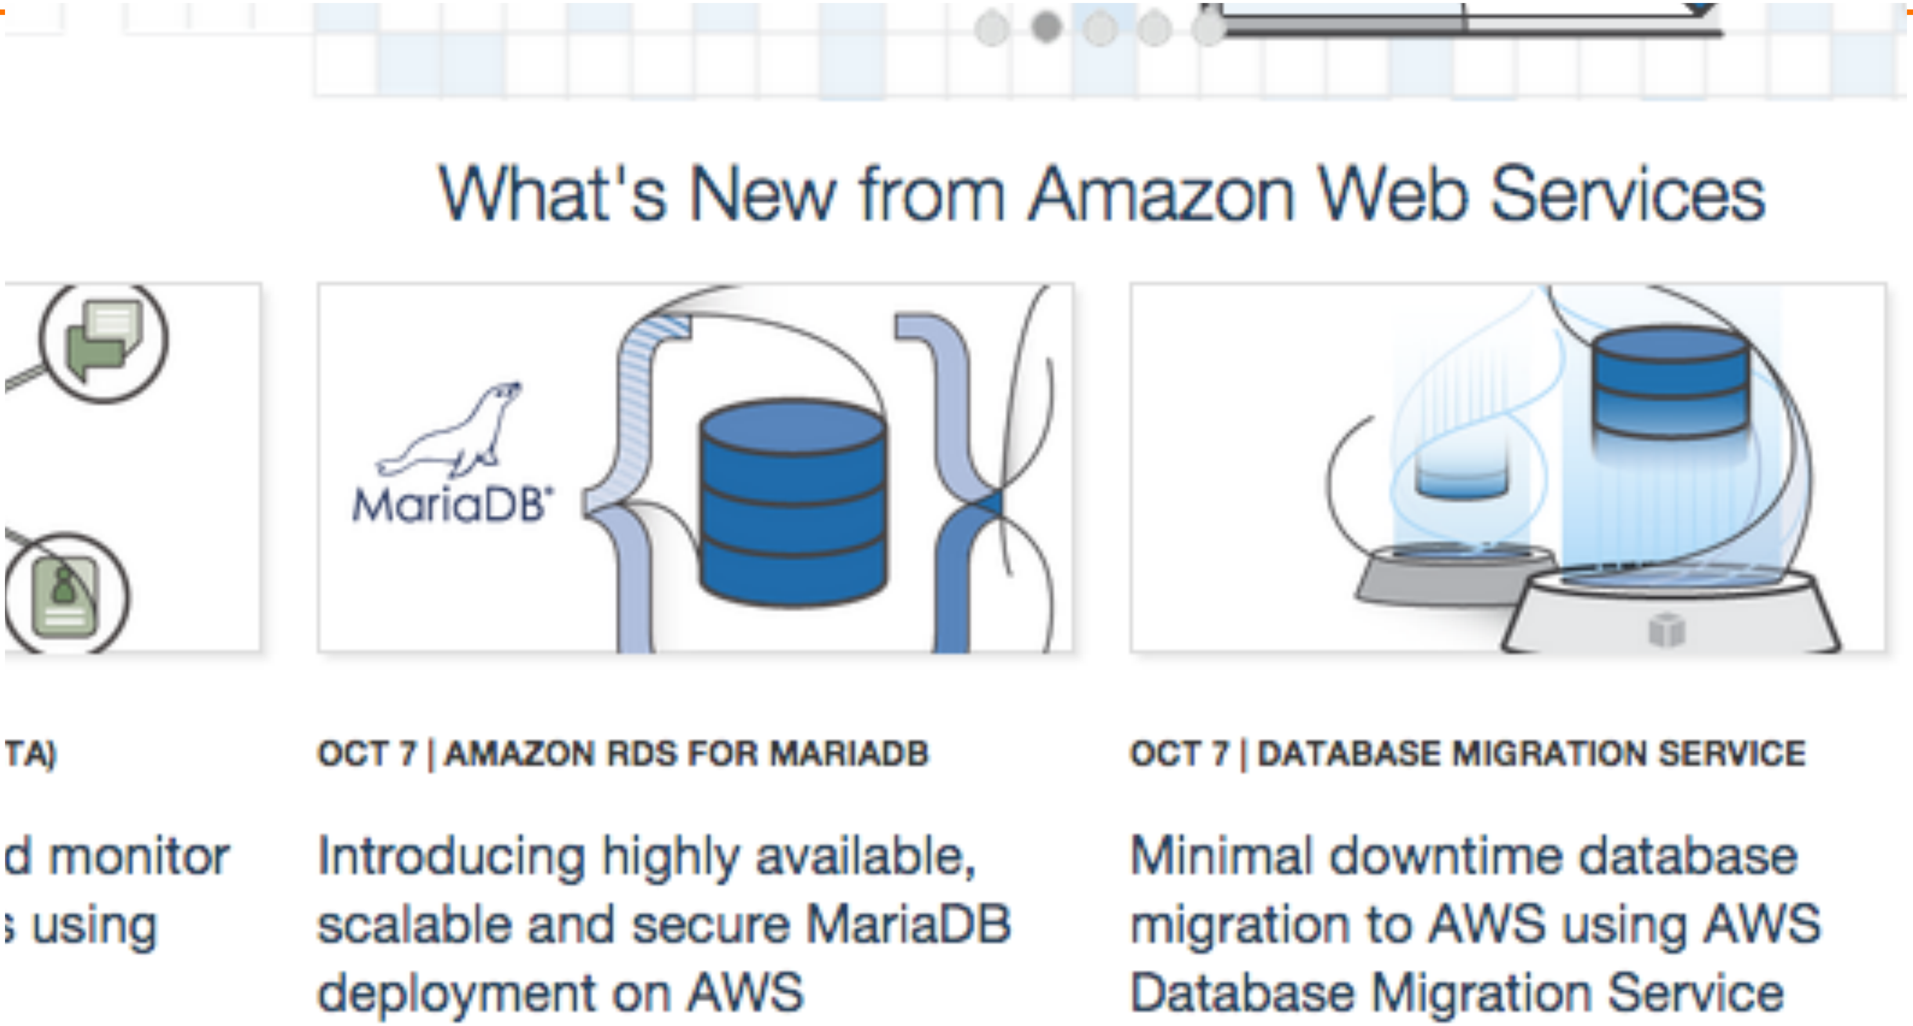

What's New from Amazon Web Services

TA)

d monitor
s using

OCT 7 | AMAZON RDS FOR MARIADB

Introducing highly available, scalable and secure MariaDB deployment on AWS

OCT 7 | DATABASE MIGRATION SERVICE

Minimal downtime database migration to AWS using AWS Database Migration Service

Regions & Availability Zones

- Region: a data centre location, containing multiple Availability Zones
- Availability Zone (AZ): isolated from failures from other AZs + low-latency network connectivity to other zones in same region

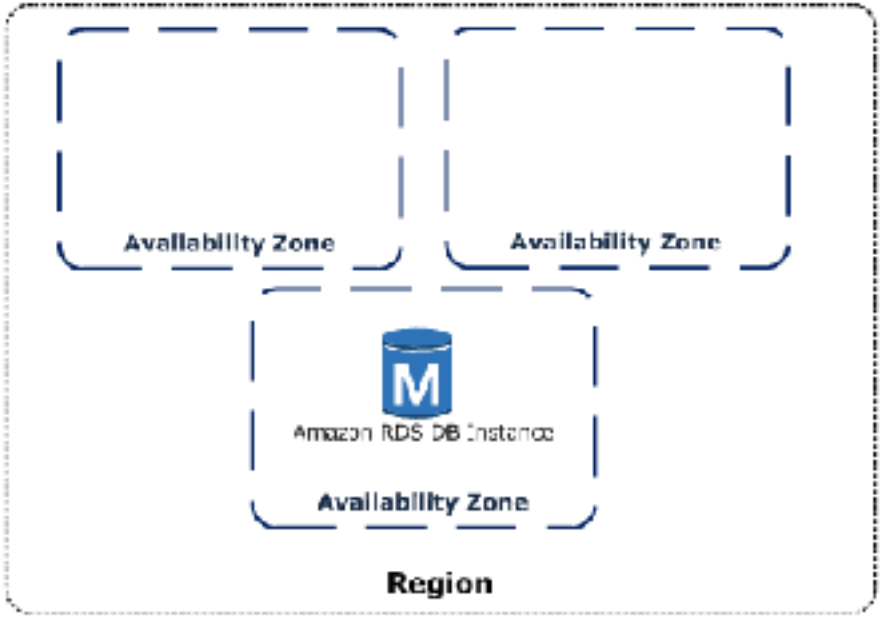

Location, location, location

- AWS RDS: US East (N. Virginia, **Ohio**), US West (Oregon, **Northern California**), US West (California), EU (Ireland, Frankfurt, **London**), APAC (Singapore, Tokyo, Sydney, **Seoul, Mumbai**), South America (São Paulo), GovCloud, Canada (**Central**), China (**Beijing**)
- Rackspace: USA (Dallas DFW, Chicago ORD, N. Virginia IAD), APAC (Sydney, Hong Kong), EU (London)*
- Google Cloud SQL: US (Berkeley County, SC, Council Bluffs, IA), EU (Belgium), Asia (Taiwan, Japan)

Service Level Agreements (SLA)

- AWS - 99.95% in a calendar month
- Rackspace - 99.9% in a calendar month
- Google - 99.95% in a calendar month
- SLAs exclude “scheduled maintenance” which may storage I/O + elevate latency
  - AWS is 30 minutes/week, so really 99.65%

Support

- AWS - forums; $49/mo gets email; $100+ phone #
- Rackspace - live chat, phone #, forums
- Google - forums; $150/mo gets support portal; $400+ for phone #
- Microsoft - forums; $29/mo to start

Who manages this?

- AWS: self-management, Enterprise ($15k+)
- Rackspace: $100 + 0.04 cents/hr over regular pricing
- Google: self-management
- Microsoft: generally self, premier can give you advisory services

MySQL versions

- AWS: MySQL Community ~~5.1~~, 5.5, 5.6, 5.7 / MariaDB Server 10.0, 10.1
- Rackspace: MariaDB Server 10, 10.1, 10.1-enc, MySQL 5.7/5.6/5.1, Percona Server 5.6
- Google: MySQL Community ~~5.5~~, 5.6, 5.7
- Microsoft: MySQL Community 5.6, 5.7

Access methods

- AWS - within Amazon, externally via mysql client, API access.
- Rackspace - private hostname within Rackspace network, API access.
- Google - within AppEngine, a command line Java tool (gcutil), standard mysql client
- Microsoft - within Azure, externally via mysql client

Can you configure MySQL?

- You don’t access my.cnf naturally
- In AWS you have parameter groups which allow configuration of MySQL

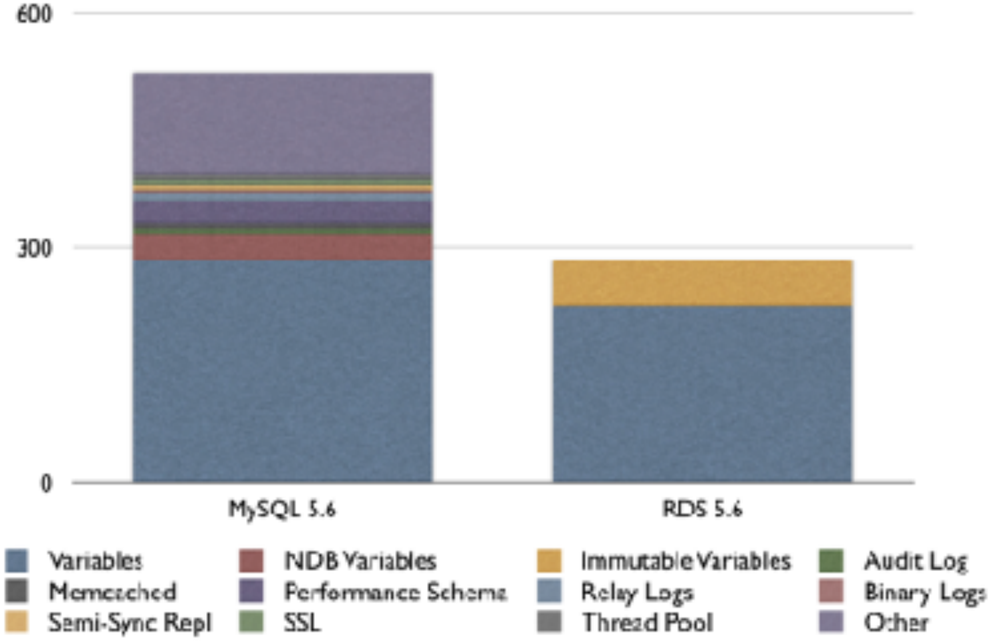

http://www.mysqlperformanceblog.com/2013/08/21/amazon-rds-with-mysql-5-6-configuration-variables/

Cost

- Subscribe to relevant newsletters of your services
- Cost changes rapidly, plus you get new instance types and new features (IOPS)
- Don’t forget network access costs
- Monitor your costs daily, hourly if possible (EC2 instances can have spot pricing)

Costs: AWS

- AWS prices vary between regions
- http://aws.amazon.com/rds/pricing/

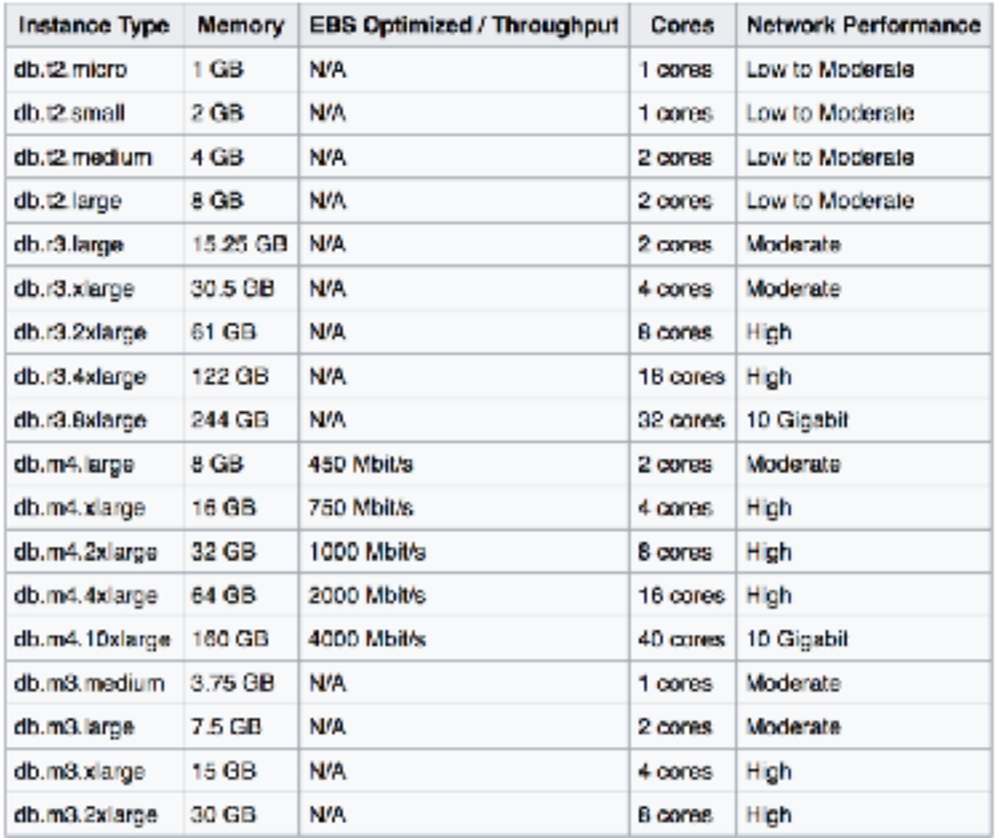

Instance Type	Memory	EBS Optimized / Throughput	Cores	Network Performance
db.t2.micro	1 GB	N/A	1 cores	Low to Moderate
db.t2.small	2 GB	N/A	1 cores	Low to Moderate
db.t2.medium	4 GB	N/A	2 cores	Low to Moderate
db.t2.large	8 GB	N/A	2 cores	Low to Moderate
db.r3.large	15.25 GB	N/A	2 cores	Moderate
db.r3.xlarge	30.5 GB	N/A	4 cores	Moderate
db.r3.2xlarge	61 GB	N/A	8 cores	High
db.r3.4xlarge	122 GB	N/A	16 cores	High
db.r3.8xlarge	244 GB	N/A	32 cores	10 Gigabit
db.m4.large	8 GB	450 Mbit/s	2 cores	Moderate
db.m4.xlarge	16 GB	750 Mbit/s	4 cores	High
db.m4.2xlarge	32 GB	1000 Mbit/s	8 cores	High
db.m4.4xlarge	64 GB	2000 Mbit/s	16 cores	High
db.m4.10xlarge	160 GB	4000 Mbit/s	40 cores	10 Gigabit
db.m3.medium	3.75 GB	N/A	1 cores	Moderate
db.m3.large	7.5 GB	N/A	2 cores	Moderate
db.m3.xlarge	15 GB	N/A	4 cores	High
db.m3.2xlarge	30 GB	N/A	8 cores	High

Costs: AWS II

- Standard latest generation: db.m4.large (8 GB) at $1,533/yr [2017]
  - Previous generation: Medium instances (3.75GB) useful for testing ($1,577/yr [2014] vs $2,411/yr [2013])
  - Large instance (7.5GB) production ready ($3,241/yr vs $4,777/yr [2013])
- Multi-AZ production ready: db.m4.2xlarge (32GB, 8vCPUs) at $12,264/yr [2017]
  - Previous generation: m3.2XL (30GB, 8vCPUs) ($12,964/yr)

Costs: Rackspace

- Option to have regular Cloud Database or Managed Instances
- 4GB instance is $2,111.4 (vs. $2,102/yr in 2015 and $3,504/yr in 2013)
- 8GB instance is $4,213.8 (vs. $4,205/yr in 2015 and $6,658/yr in 2013)
- Consider looking at I/O priority, and the actual TPS you get
- High Available instances with automatic failover available too

Costs: Google

- You must enable billing before you create Cloud SQL instances
- https://developers.google.com/cloud-sql/docs/billing
- Previously (first generation):
  - Testing (D8 - 4GB RAM) - $3,204.7 vs $4,274.15 (in 2015)
  - XL equivalent (D16 - 8GB RAM) - $6413.05 vs $8,548.30 (in 2015)
- Today:
  - db-n1-standard-8 (8 vCPUs, 30GB RAM) - $6,762.72
- Packages billing plans are cheaper than per-use billing plans

Where do you host your application?

- Typically within the compute clusters of the service you’re running the DBaaS in
  - Even Google now offers Compute Engine so there is no such language limitation, ala the early AppEngine days

RDS: Multi-AZ

- Provides enhanced durability (synchronous data replication)
- Increased availability (automatic failover)
    - Warning: can be slow (1-10 mins+)
- Easy GUI administration
- Doesn’t give you another usable “read-replica” though

External replication

- MySQL 5.6 you can do RDS -> Non-RDS
  - enable backup retention, you now have binlog access
  - target: exporting data out of RDS
- Replicate into RDS with 5.5.33 or later
  - AWS provides stored procedures like mysql.rds_set_external_master nowadays

MySQL 5.6/5.7, MariaDB 10/10.1

- MySQL 5.6 provides crash-safe slaves, the InnoDB memcached interface, online schema changes, full-text InnoDB indexes, optimizer improvements, INFORMATION_SCHEMA enhancements, scalability/replication improvements, PERFORMANCE_SCHEMA enhancements
  - 5.7 - encryption, JSON
- MariaDB 10 has some of that that, plus multi-source replication, GTIDs that don’t require full restarts, threadpool, audit plugin and more
  - 10.1 - encryption!

Getting started

- Importing data into the cloud?
  - mysqldump is a good choice today
- Upgrading from RDS 5.5 to RDS 5.6?
  - mysqldump before, but nowadays you can do this via Read Replicas

Handling backups

- You don’t get to use xtrabackup! (or mydumper)
- Google Cloud SQL automates backups (has a backup window - 4h)
- Amazon has automated backups (with point-in-time recovery), with full daily snapshots (has a backup window).
- Rackspace allow instance backups too

Monitoring

- Options are limited, AWS has the best options currently available
  - Today you have CloudWatch
- Google has improved on their read/write graphs
- Rackspace has started with basic graphs, visuals for MySQL, have a Cloud Intelligence product
- Percona Monitoring & Management - http://pmmdemo.percona.com/
- DataDog, etc.

Storage Engines

- MySQL (/MariaDB) have many
  - cool ones include TokuDB, SPIDER, CONNECT
- You basically use InnoDB and MyISAM with cloud solutions
  - MyISAM on RDS won’t guarantee point-in-time recovery, snapshot restore

High Availability

- Plan for node failures
- Don’t assume node provisioning is quick
- Backup, backup, backup!
- “Bad” nodes exist
- HA is not equal across options - RDS wins so far
  - Rackspace has High Availability Database instances built on their ObjectRocket platform

Unsupported features

- AWS MySQL: GTIDs (but MariaDB Server GTIDs work!), InnoDB Cache Warming (intra-schema parallel replication in 5.7 works - this was an XtraDB 5.6 feature), InnoDB transportable tablespaces, authentication plugins, password strength plugin, replication filters, semi-sync replication
- AWS MariaDB: Data at Rest Encryption, MariaDB Galera Cluster, HandlerSocket, Multi-source Replication, Password validation plugin, simple_password_check, and cracklib_password_check, Replication Filters, Storage engine-specific object attributes, table and Tablespace Encryption
- Google: UDFs, PERFORMANCE_SCHEMA, LOAD DATA INFILE, INSTALL PLUGIN, SELECT ... INTO OUTFILE
- mysqlsh?

Option groups

- InnoDB memcached
- MariaDB Server Audit Plugin

Provisioned IOPS

- Available on Amazon, and recently Google
- Faster, predictable, consistent I/O performance with low latencies
- Good throughput, RAID on backed
- EBS is more reliable

More on RDS

- log access via API
- no SUPER access to skip replication errors easily
- sync_binlog=0 not available
- no OS access (sar, ps, tcpdump)
- https://github.com/boto/boto

Warning: automatic upgrades

- Regressions happen even with a minor version upgrade in the MySQL world (though this is happening a lot less as quality has increased tremendously)
- InnoDB update that modifies rows PK triggers recursive behaviour until all disk space is exceeded? 5.5.24->5.5.25 (fixed: 5.5.25a)
- Using query cache for partitioned tables? Disabled since 5.5.22->5.5.23!

Benchmarking for use

- sysbench
  - OLTP test, use tables with 20M rows and 20M transactions, check 1-128 threads/run (run this on RDS, Rackspace)
  - June 2013, tps, performance per dollar, Rackspace delivers more performance across all flavours except 512MB instance
- Yahoo! Cloud Serving Benchmark
  - https://github.com/brianfrankcooper/YCSB
- Google's PerfKit Benchmarker
  - https://github.com/GoogleCloudPlatform/PerfKitBenchmarker

Roadmaps?

- There don’t seem to be public roadmaps. You find out when there’s a change!
- Watch out for those events…

Usability

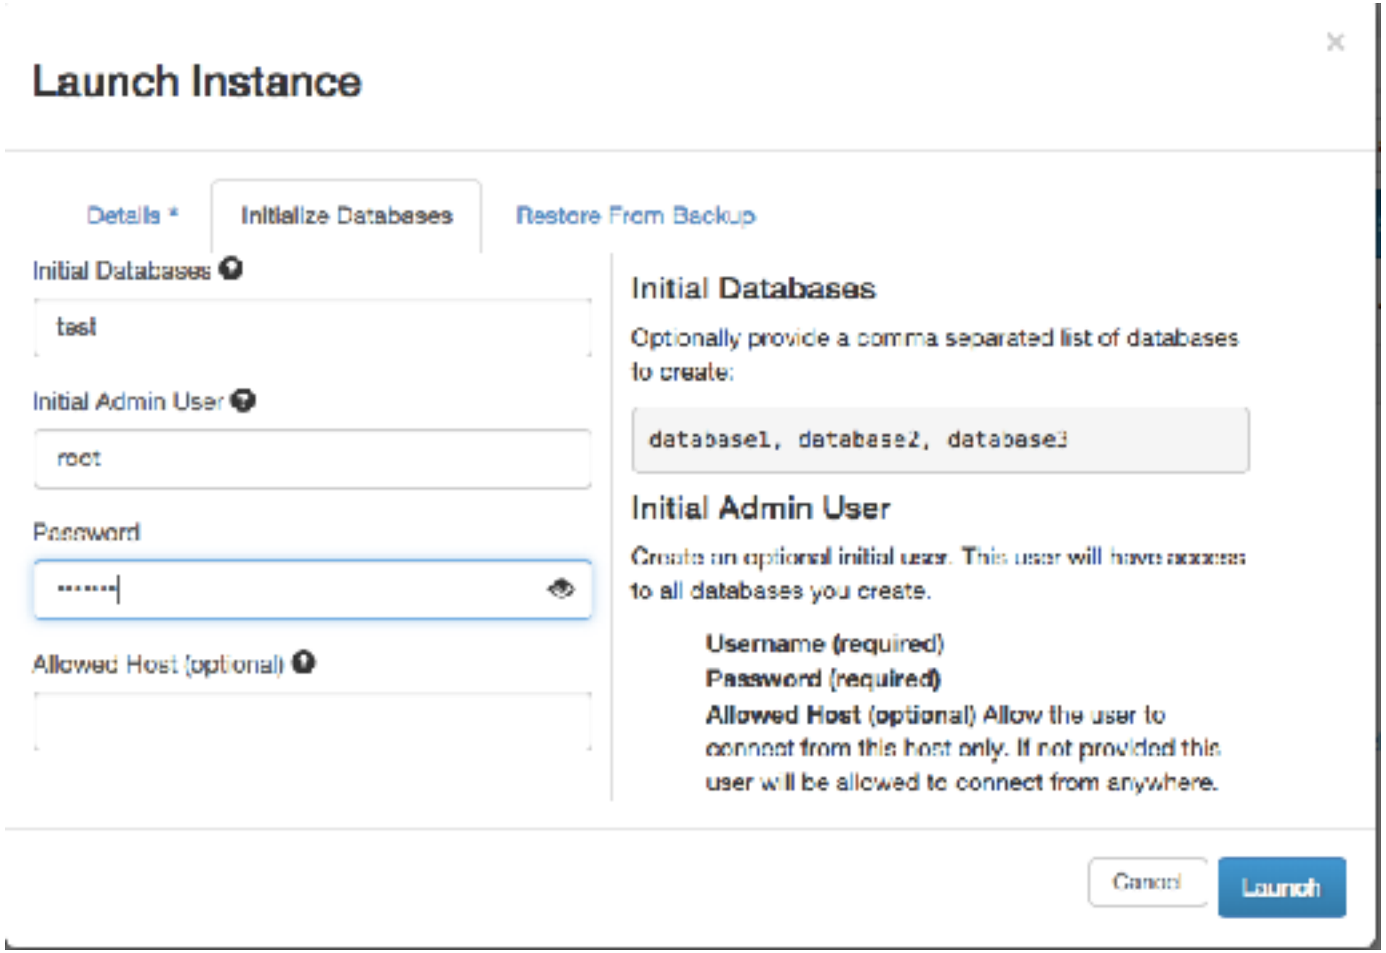

Launch Instance

Initial Databases

test

Initial Admin User

root

Password

•••••••

Allowed Host (optional)

Initial Databases

Optionally provide a comma separated list of databases to create:

```
database1, database2, database3
```

Initial Admin User

Create an optional initial user. This user will have access to all databases you create.

Username (required)
Password (required)
Allowed Host (optional) Allow the user to connect from this host only. If not provided this user will be allowed to connect from anywhere.

PostgreSQL in RDS

- loading data? backup_retntion=0
- disable multi-AZ when loading
- disable autovacuum
- dump compressed, restore in parallel
- don’t disable fsync (really!)
- sync replication using multi-AZ
- you can control the upgrade time though — this is a bonus
- Use PIOPS
- SSL should be on

Running MySQL in EC2

- Can do multiple geographic regions via replication
- Run just one Percona Server/MariaDB server/instance
- Use additional EBS volumes for InnoDB tablespaces
- RAID EBS volumes (RAID1)
- Warm up data partitions, mount partitions with noatime, nodirtime
- Vertical scaling with SSD-backed storage
- Monitoring with nagios
- Snapshot backups and save to S3
- Can use Elastic Load Balancer
- Can use spot instances
- Can use tools like MHA to provide automatic failover
- Can use MariaDB Galera Cluster/Percona XtraDB Cluster

AWS Aurora

- Bigger instances work better
- Zero-downtime migration from RDS
- Metrics via CloudWatch, Connectors via MariaDB
- 99.99% uptime
- MySQL 5.6 “fork”, no optimiser, not traditional replication (but Aurora <->MySQL works of course)
- Auto scaling - compute, memory, storage
- Replicas (15) for reads
- Automated backups in S3, DB snapshots
- Encryption with key server being Amazon KMS
- Spatial data support - like InnoDB 5.7!

Some closing thoughts

- Hardware varies per region
- Sometimes, software manageability varies per region
- Beware cost on your credit card!

Thank you!

Colin Charles

colin.charles@percona.com / byte@bytebot.net

http://bytebot.net/blog | @bytebot on twitter

slides: slideshare.net/bytebot