

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 Webinar

29 August 2018



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. and now PostgreSQL too!
- 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?
- What about PostgreSQL? MongoDB?
- 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)

```
aws ec2 run-instances --image-id ami-xxx -k $  
{EC2_KEYPAIR} -t c3.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
- Alibaba Cloud
- 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

Google Compute Engine

- Google Compute Engine offers Percona XtraDB Cluster as a “click-to-deploy” app
- comes with Galera 3, Percona Toolkit, XtraBackup as well

Pivotal CloudFoundry

- Pivotal CloudFoundry
 - “MySQL” PaaS which is MariaDB Galera Cluster 10 (v1, legacy)
 - MariaDB v10.1.30 and Galera v25.3.20 (latest release, June 12 2018)
 - MySQL for PCF v2
 - Percona Server v5.7.20-21 (latest release, June 22 2018)
- Works with IaaS platforms: AWS, Azure, GCP, OpenStack, and vSphere
- Significance of Galera Cluster versus regular replication

Red Hat OpenShift

- Variants between Online & Enterprise editions (usually you get access to *modern* releases with OpenShift Online)
- MySQL 5.5/5.6/5.7
- MariaDB 10.0/10.1
- PostgreSQL 9.4/9.5
- MongoDB 2.4/2.6/3.2/3.4

Beware

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



HP Comes to Terms With the Cloud

By QUENTIN HARRY
APRIL 13, 2015 1:44 PM

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



HP Helion Strategy to Deliver Hybrid IT Continues Strong

Stephen_Sexton | April 13, 2015

Guest Post by Bill Hill, 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, lower 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 their 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.
- We operate a fast-growing hosted, virtual private cloud – designed specifically for enterprise customers seeking

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+04:00)



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.

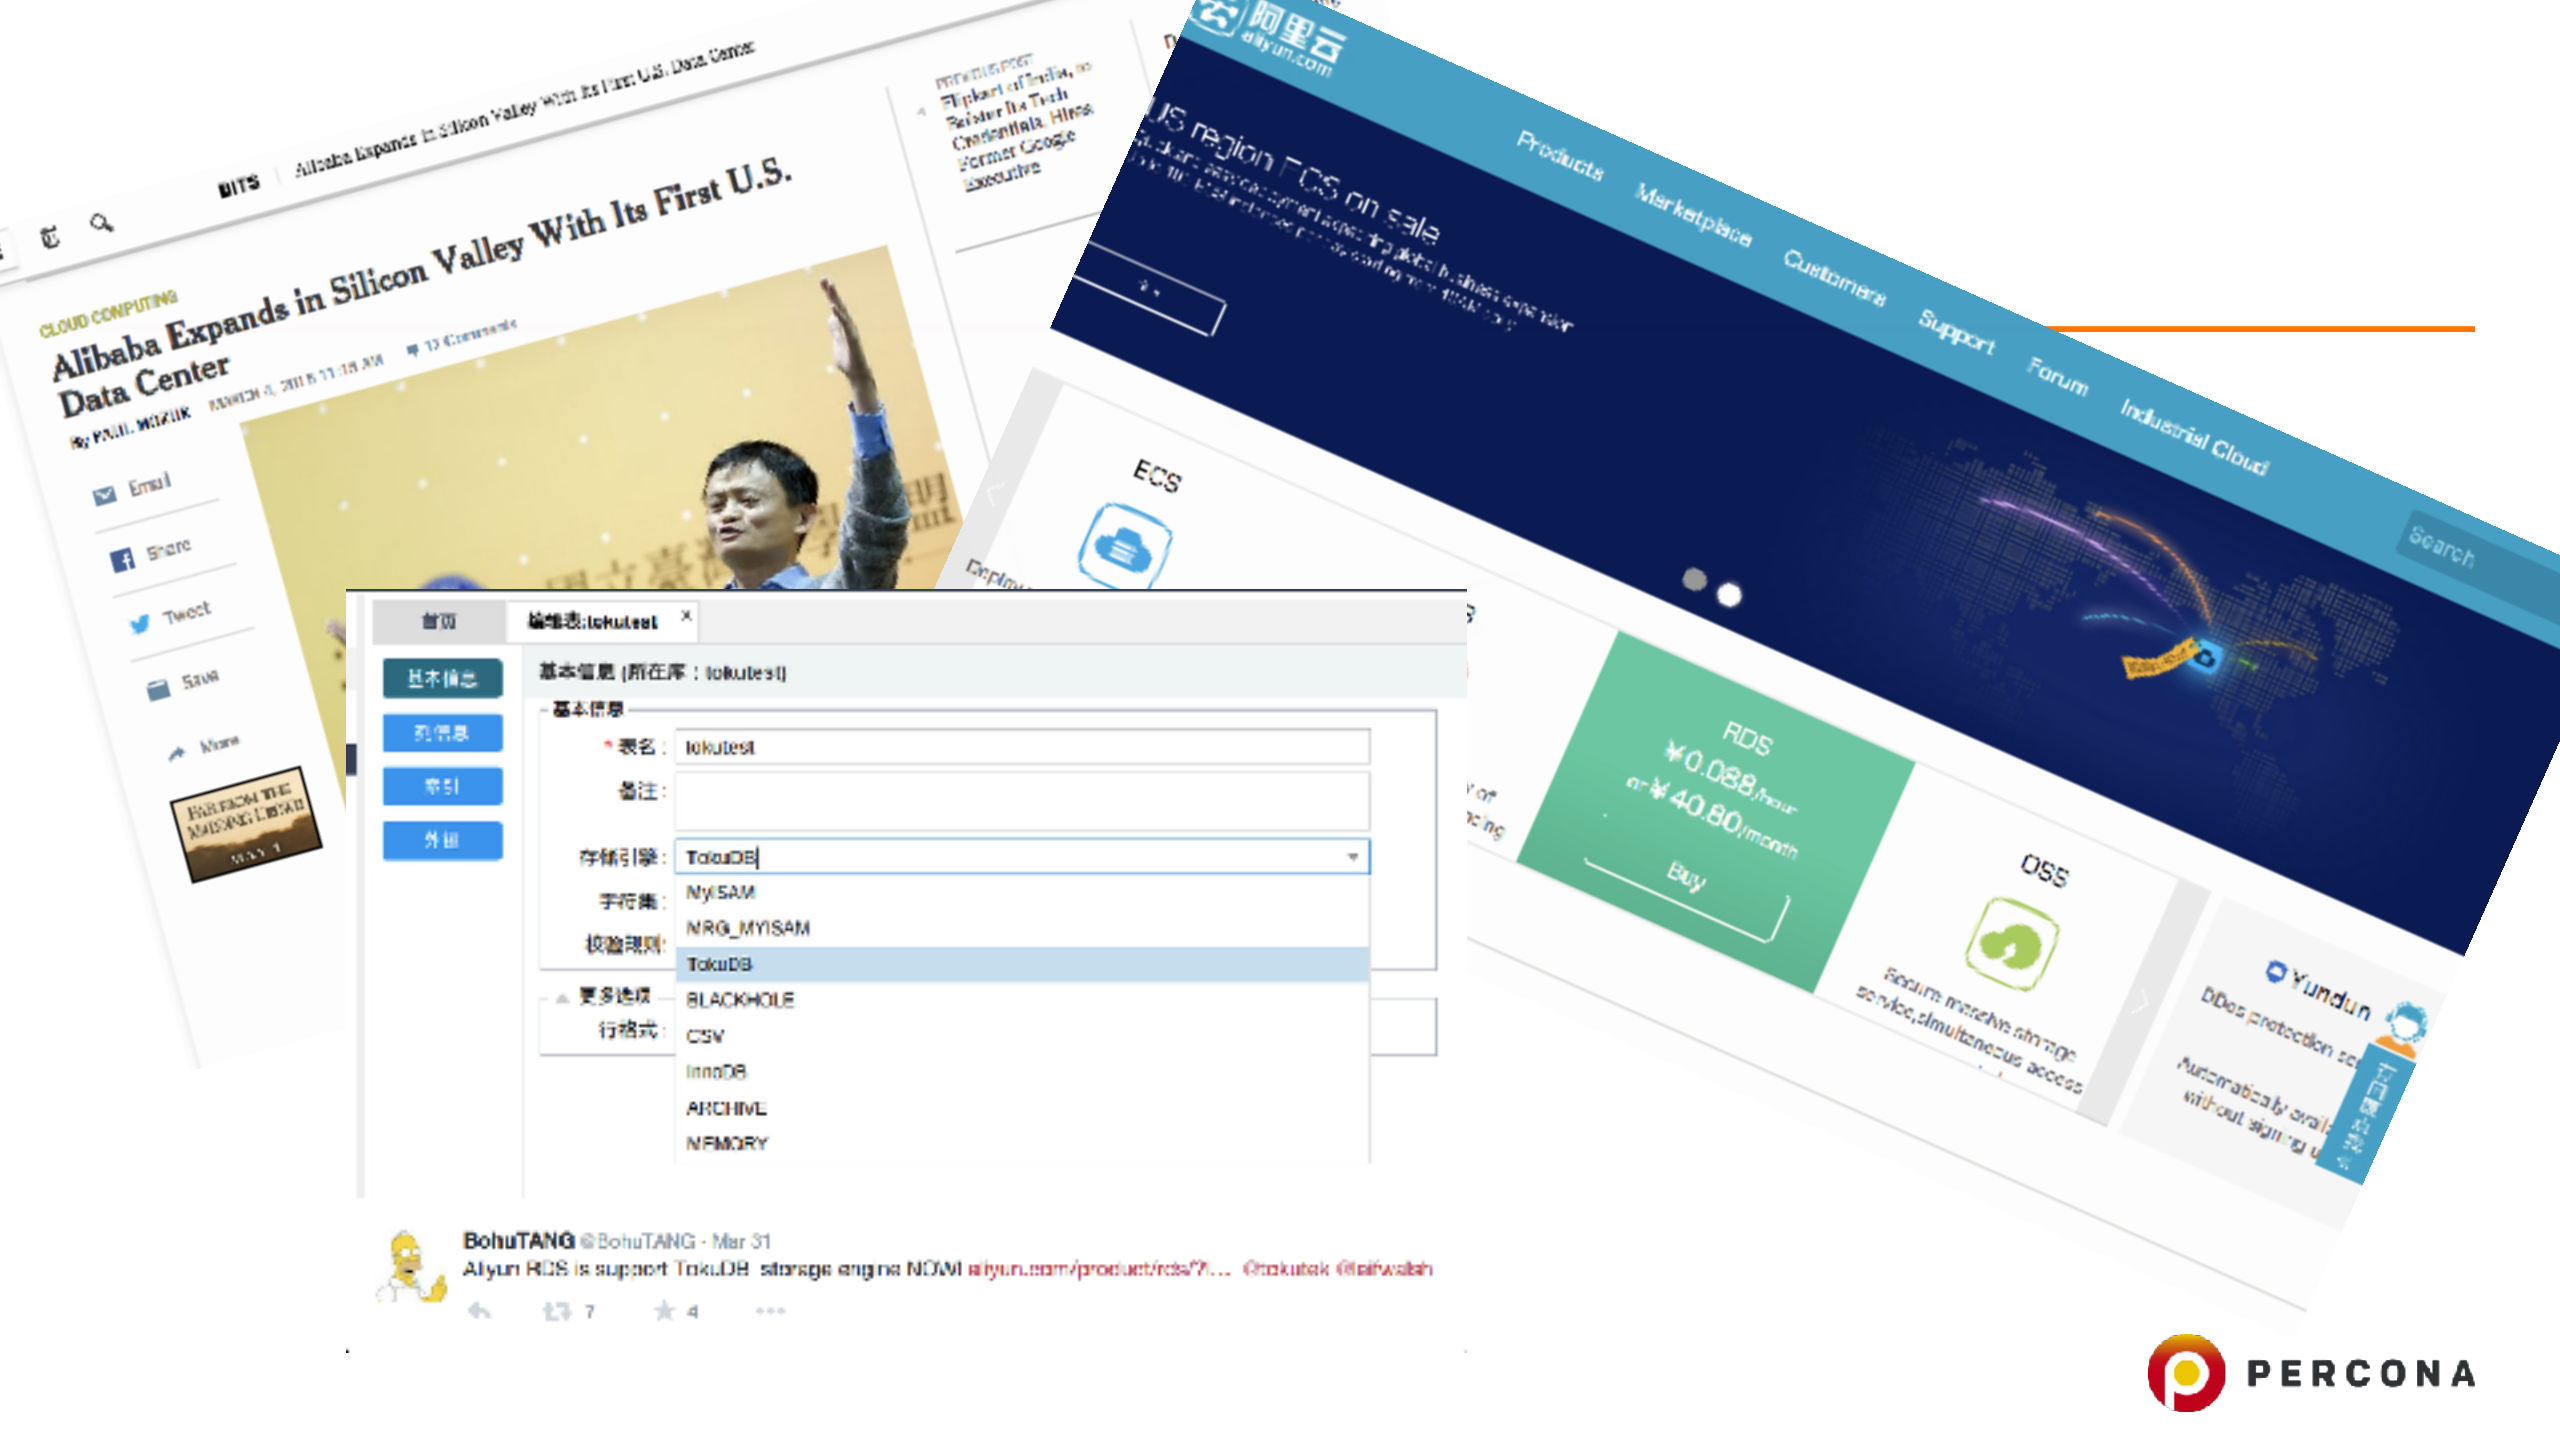
HP's SVP of Helion Product Management Bill Hill has a quote of mine in the media which was interpreted as HP public relations simply exiting the public cloud, which is not the case. In fact, HP public relations simply announced as HP is exiting the public cloud, which is not the case. Our portfolio strategy to deliver on the vision of Hybrid IT continues strong, highlighted by the following:



Bill Hill, SVP of Helion Product Management, announces that HP is not leaving the public cloud market. — 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



CLOUD COMPUTING

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

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



- Email
- Share
- Tweet
- Save
- More



阿里云 aliyun.com

Products Marketplace Customers Support Forum Industrial Cloud

US region ECS on sale

Buy now. Alibaba Cloud is offering global business class year-over-year the best enterprise cloud service at the lowest price.

Search

首页 编辑表: tokutest

基本信息

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: MyISAM

校验规则: MRG_MYISAM

更多选项

行格式: TokuDB

BLACKHOLE

CSV

InnoDB

ARCHIVE

MEMORY

RDS

¥0.088/hour or ¥40.80/month

Buy

OSS

Secure massive storage service, simultaneous access

Yundun

DDoS protection service

Automatically available without signing up

BohuTANG @BohuTANG · Mar 31

Aliyun RDS is support TokuDB storage engine NOW! aliyun.com/product/rds/?... @takutak @alfwalsh

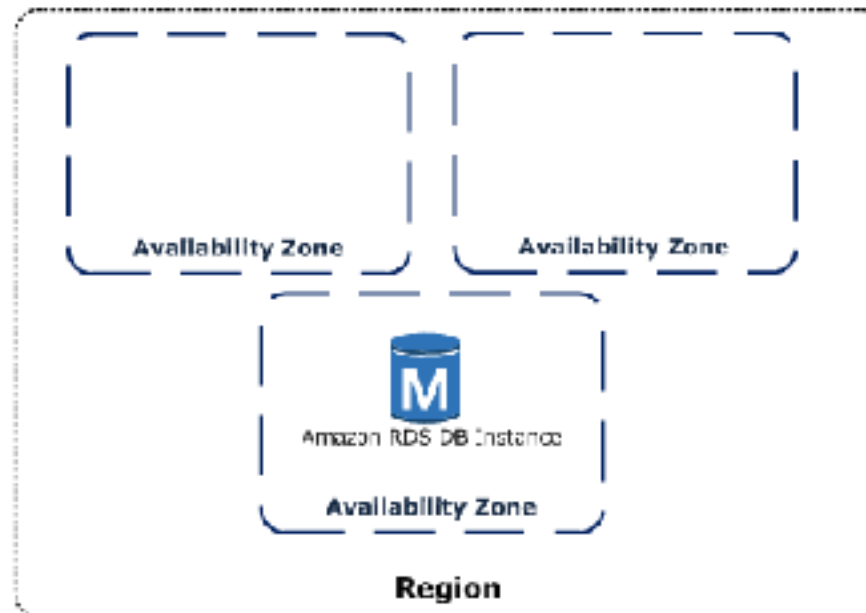
7 4

What else do you get?

- PostgreSQL is getting popular in the cloud: AWS RDS, AWS Aurora, Microsoft Azure, Alibaba Cloud, Google Cloud SQL
- Most MongoDB users run within the cloud — but not hosted, it tends to be “roll your own” or via ObjectRocket (and lately MongoDB Atlas c- reasonable free tier!)
- Compose.io (now IBM) & the like (MongoLab) — hosted MongoDB, Redis, Enhanced PostgreSQL, Elasticsearch, RethinkDB (!?), MySQL with group replication, ScyllaDB, etc.

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**, California), EU (**Ireland, Frankfurt, London, Paris**), 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:
 - 2nd generation instances: Montréal, Iowa, N. Virginia, Oregon, LA, São Paulo, Finland, London, Frankfurt, Netherlands, Tokyo, Mumbai, Singapore, Sydney

Service Level Agreements (SLA)

- AWS - at least 99.95% in a calendar month, less than, 10% service credit
- Rackspace - 99.9% in a calendar month
- Google - 99.95% in a calendar month, less than 10% service credit
- Alibaba Cloud - no less than 99.95% in a calendar month (some services, 99.9%)
- Microsoft Azure - < 99.99% will give you a 10% service credit
- SLAs exclude “scheduled maintenance” which may have impact on storage I/O + elevate latency
 - e.g. AWS is 30 minutes/week

Support

- AWS - active forums; \$100+ (or a % of AWS usage) phone #
- Rackspace - live chat, phone #, forums
- Google - forums; \$150/mo gets support portal; \$400+ (or a % of usage fees) for phone #
- Microsoft - forums; \$100/mo to start
- Alibaba Cloud - \$70 (or a % of usage fees)
- Most have developer support plans at \$19.99-\$29 even; enterprise support around \$8,000-\$15,000

Who manages this?

- AWS: self-management, Enterprise (\$15k+)
- Rackspace Cloud: \$100 + 0.04 cents/hr over regular pricing
- Google: self-management, or via partner, Rackspace Managed Services
- Microsoft: generally self, premier can give you advisory services
- Alibaba Cloud: self-management
- Rackspace Managed Services: AWS, Alibaba Cloud, Google Cloud Platform, Microsoft Azure Cloud, OpenStack Public Cloud

MySQL versions

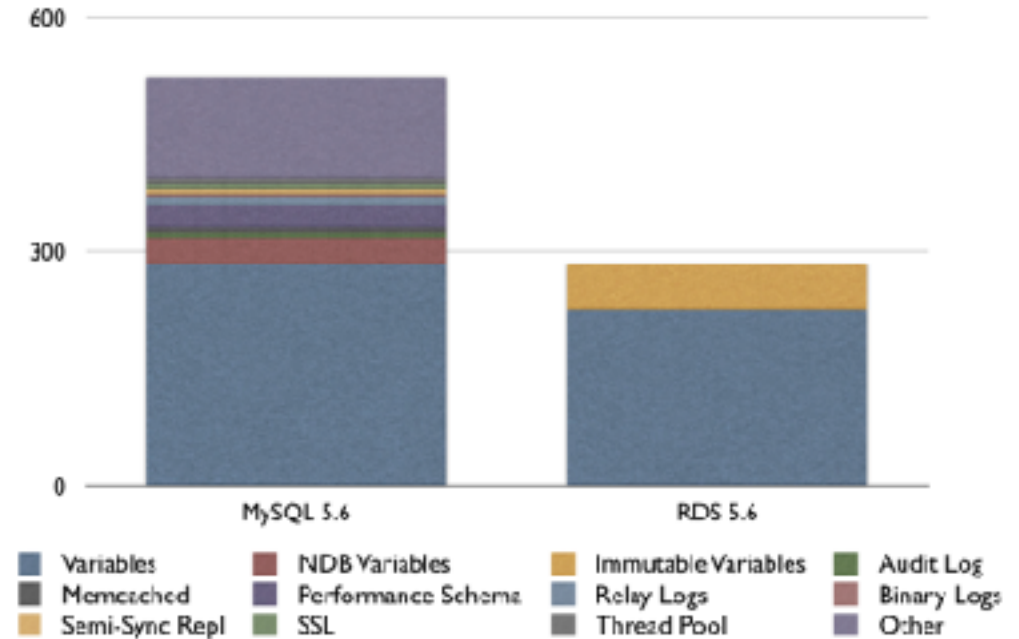
- AWS: MySQL Community ~~5.4~~, 5.5, 5.6, 5.7 / MariaDB Server 10.0, 10.1, 10.2
- Rackspace: MariaDB Server 10, 10.1, 10.1-enc, MySQL 5.7/5.6/5.1, Percona Server 5.6
- Google: MySQL Community ~~5.5~~, 5.6, 5.7
- Microsoft: MySQL Community 5.6, 5.7
- Alibaba Cloud: MySQL Community 5.6, 5.7

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
- Alibaba Cloud - standard 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 (spot pricing — if roll your own)

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 (2vCPU/8GB) 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

Costs: Microsoft Azure

- I/O isn't being charged for yet
- Basic tier with 2 vCores gives you 2GB per vCore (=4GB)
- So... 4 vCores, 8GB RAM, \$3,069/year

Costs: Alibaba Cloud

- rds.mysql.s2.xlarge - 2vCPU/8GB RAM - \$4,537 (\$3,192 on subscription, pay it monthly at \$266)
- In USA, that could be \$3,109 (\$2,112, monthly \$176)
- In China? \$2,487 (\$1,632, monthly \$136)
- SQL Audit, backup, monitoring, all costs more
- Readonly storage instance pricing exist

Where do you host your application?

- Typically within the compute clusters of the service you're running the DBaaS in
- You want a multi-cloud strategy? You will have to have your entire layers of applications replicated on different clouds

RDS: Multi-AZ

- Provides enhanced durability (synchronous data replication)
- Increased availability (automatic failover)
 - Warning: can be slow with large database size
- 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
- This can help you migrate to another platform as well
- Replicate into RDS with 5.5.33 or later
- AWS provides stored procedures like `mysql.rds_set_external_master` nowadays

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 Percona 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).
 - Multi-AZ? Backup taken from the standby!
 - Backup retention default? 1 day. Increase it
 - Aria may not work well with automatic backups, so use InnoDB/XtraDB
- Rackspace allow instance backups too
- Microsoft has automatic backups (with point-in-time recovery). Backup retention is 7 days (up to 35 days), with redundancy options

Monitoring

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

Storage Engines

- MySQL (/MariaDB) have many
 - e.g. MyRocks, TokuDB, SPIDER, CONNECT
- You basically use InnoDB (XtraDB) 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
- RDS? Multi-AZ.
- HA is not equal across options
 - Rackspace has High Availability Database instances built on their ObjectRocket platform
- Google? Deploys semi-synchronous replication, so you do get a usable replica
 - alert for replication lag

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 running in the hosted cloud...

- log access via API
- no SUPER access to skip replication errors easily
- `sync_binlog=0` not available
- no OS access (`sar`, `ps`, `tcpdump`)

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 Benchmark
 - <https://github.com/GoogleCloudPlatform/PerfKitBenchmark>

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 ✕

[Details *](#) **Initialize Databases** [Restore From Backup](#)

Initial Databases ⓘ

Initial Admin User ⓘ
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_retention=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

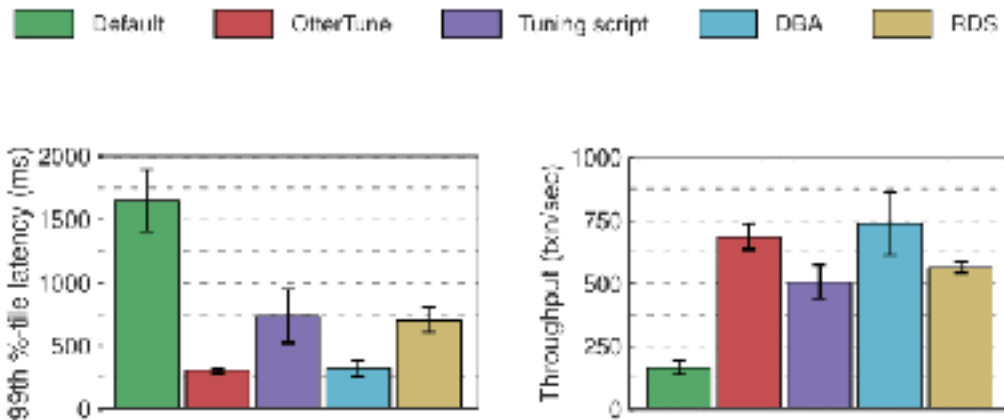
- Can do multiple geographic regions via replication
- Run just one Percona Server/MariaDB server/instance
- Use additional EBS volumes for InnoDB tablespaces
- RAID EBS volumes (RAID1)
- Warm up data partitions, mount partitions with noatime, nodirtime
- Vertical scaling with SSD-backed storage
- Monitoring with Icinga/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 5.6.10 “fork”, no optimiser, not traditional replication (but Aurora <->MySQL works of course)
- MySQL 5.7.12 Aurora launched Feb 2018, with JSON support, spatial indexes, generated columns, etc.
- 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!
- Lab mode (hash joins, scan batching, etc.)

Looking ahead

- OtterTune: automatically find good settings for a database configuration - <https://github.com/cmu-db/ottertune>
- Peloton: self-driving database management system - <http://pelotondb.io/>



60% reduction in latency,
22-35% better throughput

<https://aws.amazon.com/blogs/ai/tuning-your-dbms-automatically-with-machine-learning/>

World's First
"Self-Driving" Database



No Human Labor
Half the Cost
No Human Error
100x More Reliable

ORACLE

Some closing thoughts

- Hardware varies per region
- Sometimes, software manageability varies per region
- Beware costs on your credit card!
- These things change often, in terms of pricing, instance availability — so regularly monitor latest news
- Don't upgrade *immediately* to the latest new releases
- Always read release notes
- If going the EC2 (or equivalent) route, maybe have other management services in place (e.g. Ewhat MongoDB provides)
- Sharding — vitess?
- Proxies — ProxySQL, MySQL Router

Thank you!

Colin Charles

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

<http://bytebot.net/blog> | [@bytebot](#) on twitter

slides: slideshare.net/bytebot