

MySQL in the Hosted Cloud

Colin Charles, MariaDB Corporation Ab

colin@mariadb.org | byte@bytebot.net


<http://mariadb.com/> | <http://mariadb.org/>

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

SCALE13x, Los Angeles, CA, USA

21 February 2015

whoami



- Work on MariaDB at MariaDB Corporation (~~SkySQL Ab~~)
- Merged with Monty Program Ab, makers of MariaDB
- Formerly MySQL AB (exit: Sun Microsystems)
- Past lives include Fedora Project (FESCO), OpenOffice.org

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

MySQL as a service

- Database as a Service (DBaaS)
- MySQL available on-demand, without any installation/configuration of hardware/software
- Pay-per-usage based
- Provider maintains MySQL, 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
```



credit: <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
- HP Helion Public Cloud Relational DB

There are more

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

The new entrants

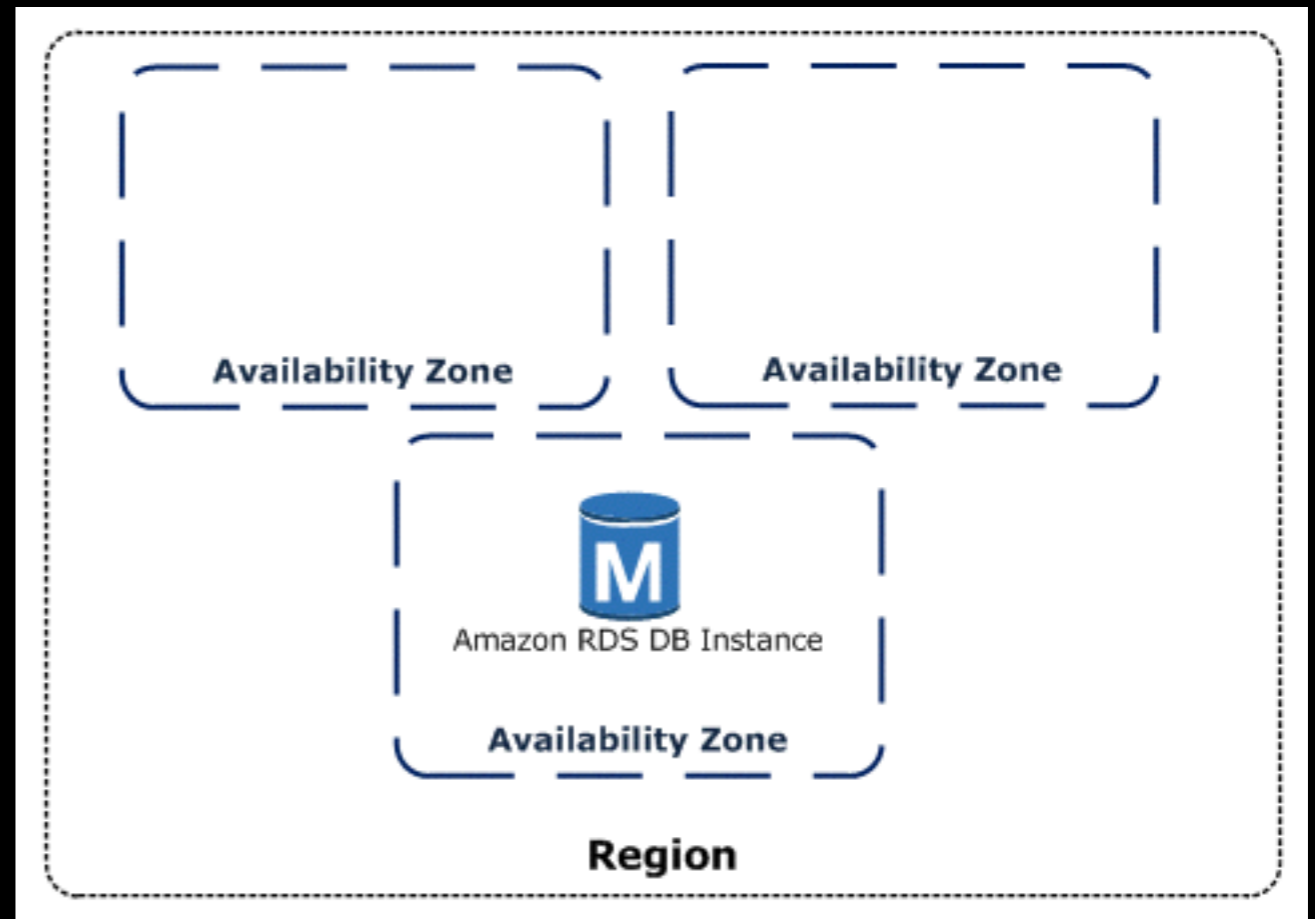
- Google Compute Engine offers Percona XtraDB Cluster as a “click-to-deploy” app
 - comes with Galera 3, Percona Toolkit, XtraBackup as well
- Pivotal CloudFoundry - “MySQL” PaaS which is MariaDB Galera Cluster 10
- Red Hat OpenShift - MySQL 5.1/5.5, MariaDB 5.5

Beware

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

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), US West (Oregon), US West (California), EU (Ireland, Frankfurt), APAC (Singapore, Tokyo, Sydney), South America (São Paulo), GovCloud
- Rackspace: USA (Dallas DFW, Chicago ORD, N. Virginia IAD), APAC (Sydney, Hong Kong), EU (London)*
- Google Cloud SQL: US, EU, Asia
- HP Cloud: US-East (Virginia), US-West

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
- HP Cloud - no specific DB SLA, 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 #
- HP Cloud - phone #, chat, customer forum

Who manages this?

- AWS: self-management, Enterprise (\$15k+)
- Rackspace: \$100 + 0.04 cents/hr over regular pricing
- Google: self-management
- HP Cloud: self-management

MySQL versions

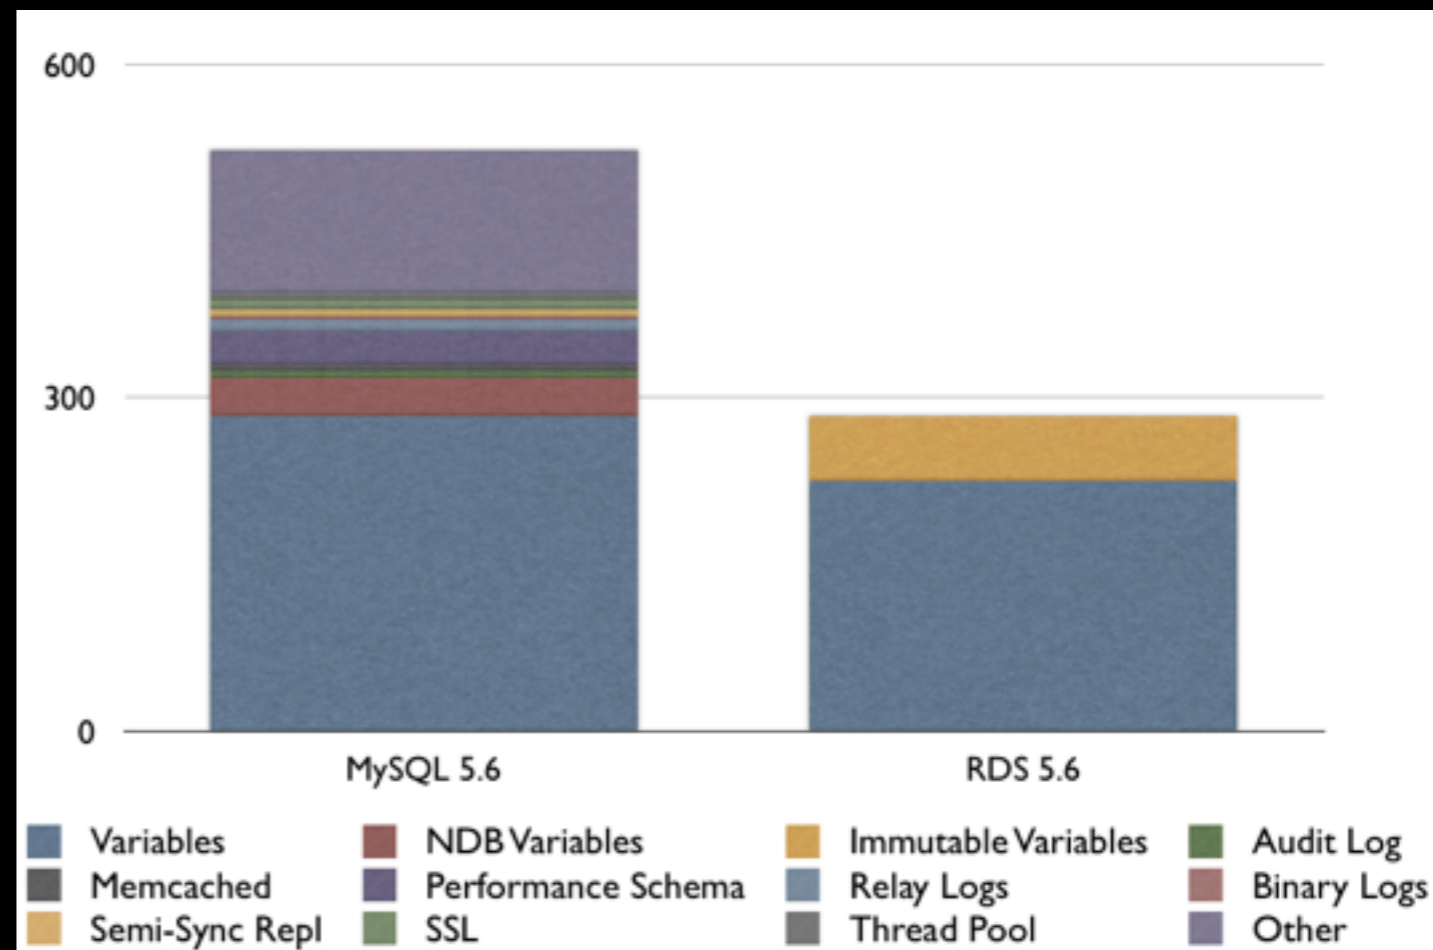
- AWS: MySQL Community 5.1, 5.5, 5.6
- Rackspace: MariaDB 10, MySQL 5.6/5.1, Percona Server 5.6
- Google: MySQL Community 5.5, 5.6 (preview)
- HP Cloud: Percona Server 5.5.28

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
- HP Cloud - within HP Cloud, externally via client (trove-cli, reddwarf), API access, mysql client

Can you configure MySQL?

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



source: <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)
- <https://github.com/ronaldbradford/aws>

Costs: AWS

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

Instance Type	vCPU	ECU	Memory (GiB)	EBS Optimized	Network Performance
db.t1.micro	1	1	0.615	No	Very Low
db.m3.medium	1	3	3.750	No	Moderate
db.m3.large	2	6.5	7.500	No	Moderate
db.m3.xlarge	4	13	15.000	500 Mbit/s	High
db.m3.2xlarge	8	26	30.000	1000 Mbit/s	High
db.m1.small	1	1	1.700	No	Low
db.m1.medium	1	2	3.750	No	Moderate
db.m1.large	2	4	7.500	500 Mbit/s	Moderate
db.m1.xlarge	4	8	15.000	1000 Mbit/s	High
db.m2.xlarge	2	6.5	17.100	No	Moderate
db.m2.2xlarge	4	13	34.200	500 Mbit/s	Moderate
db.m2.4xlarge	8	26	68.400	1000 Mbit/s	High
db.cr1.8xlarge	32	88	244.000	No	10 GiB

Costs: AWS II

- 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])
- m3.2XL (30GB, 8vCPUs) (\$12,964/yr)
 - XL instance (15GB, 8ECUs) (\$9,555/yr)

Costs: Rackspace

- Option to have regular Cloud Database or Managed Instances
- 4GB instance (testing) is \$2,102/yr (vs. \$3,504/yr in 2013)
- 8GB instance (production) is \$4,205/yr (vs \$6,658/yr in 2013)
- Consider looking at I/O priority, and the actual TPS you get

Costs: Google

- You must enable billing before you create Cloud SQL instances
- <https://developers.google.com/cloud-sql/docs/billing>
- Testing (D8 - 4GB RAM) - (\$4,274.15)
- XL equivalent for production (D16 - 8GB RAM) - (\$8,548.30)
- Packages billing plans are cheaper than per-use billing plans

Costs: HP Cloud

- 50% off pricing while in public beta
- 4GB RAM, 60GB storage - \$1,752/yr (usual: \$3,504/yr)
- 8GB RAM, 120GB storage - \$3,504/yr (usual: \$7,008/yr)

Where do you host your application?

- Typically within the compute clusters of the service you're running the DBaaS in
- This also means your language choices are limited based on what the platform offers (eg. AppEngine only offers Java, Python, PHP, Go)

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, MariaDB 10

- MySQL 5.6 in RDS 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
- MariaDB 10 has much of that that, plus multi-source replication, GTIDs that don't require full restarts, threadpool, audit plugin and more

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!
- 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 + HPCloud allow instance backups too

Monitoring

- Options are limited, AWS has the best options currently available
- Today you have CloudWatch
- Google has basic read/write graphs
- Rackspace has started with basic graphs, visuals for MySQL coming soon, have a beta Cloud Intelligence product

Storage Engines

- MySQL (/MariaDB) has many
 - cool ones include TokuDB, SPIDER, CONNECT, CassandraSE
- 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

Unsupported features

- AWS: GTIDs, InnoDB Cache Warming, InnoDB transportable tablespaces, authentication plugins, semi-sync replication
- Google: UDFs, replication, LOAD DATA INFILE, INSTALL PLUGIN, SELECT ... INTO OUTFILE

Provisioned IOPS

- Only available on Amazon
- 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
- 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>

Roadmaps?

- There don't seem to be public roadmaps. You find out when there's a change!
- Presumably HPCloud will get 5.6... and maybe Google will get some MariaDB?

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:

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.

Cancel Launch

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 (RAID 1)
- 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

Some closing thoughts

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

Q&A. Thank You.

colin@mariadb.org | byte@bytebot.net
<http://bytebot.net/blog/> | @bytebot on Twitter

Download MariaDB and give it a try: <http://mariadb.org/>