



Building the slowest MySQL application

Stéphane Combaudon
Froscon
August 2012

Agenda

- Architecture
- Configuration
- Schema/Indexes
- Queries
- Hardware
- Backup/Recovery
- Instrumentation


Architecture

- Does everybody need sharding?
 - Definitely no!
 - Only if a single server can't handle the write load
 - Functional partitioning is safer and easier to implement
- Keep it simple
 - Not every application is Facebook size!
 - Operations are cheaper if the architecture is simple

Typical replication topologies

- Master-master
 - Great to improve HA, not to improve write capacity
 - Never write on both masters!
- Master-slaves
 - The “standard” setup, very good to scale a read-mostly application
 - Promoting a slave if the master crashes is not as easy as it seems

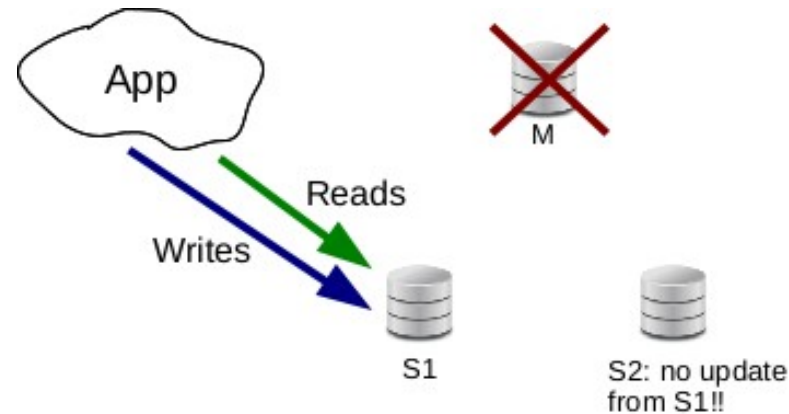
Master-slave: master crash



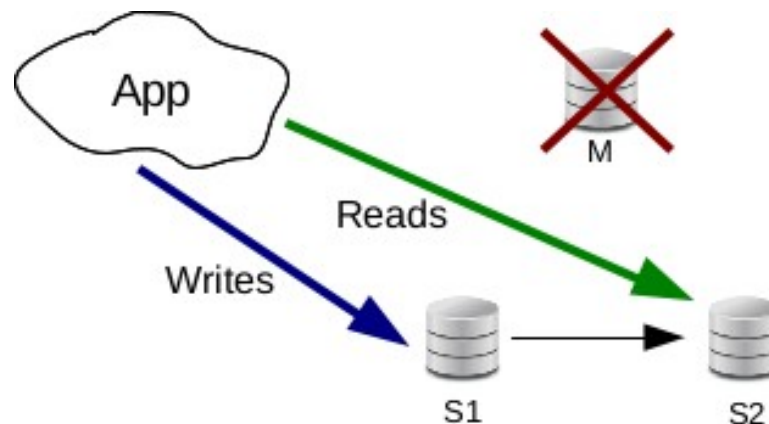
- If master crashes
 - No writes available
 - You have to promote one of the slaves

Promoting a slave

- Promoting S1 is not enough



- Need to set up replication between S1 and S2



Configuration

- Most common traps
 - Keep the default my.cnf
 - Spend weeks to fine-tune every setting
 - New HW 2x powerful does not mean you can simply set settings 2x bigger
 - Anyway bigger is not always better
- These can lead to performance problems, instability and frustration

Essential settings

- InnoDB
 - Buffer pool: `innodb_buffer_pool_size`
 - Redo logs: `innodb_log_file_size`
- No query cache (almost always correct)
 - `query_cache_type = 0`
 - `query_cache_size = 0`
- You're done!
 - Was it easier than you thought?

Easy to configure and helpful

- InnoDB durability: `innodb_flush_log_at_trx_commit`
- Sync binlog contents to disk at each commit:
`sync_binlog = 1`
- Easy to configure by reading the doc
 - Size of the `table_cache`: `table_definition_cache`,
`table_open_cache`
 - Size of thread cache: `thread_cache`
 - Upper limit of concurrent `cnx`: `max_connections`
 - Disabling DNS lookups: `skip_name_resolve`

What you should not configure

- Specialized buffers
 - `sort_buffer_size`, `join_buffer_size`, ...
- Esoteric settings
 - `innodb_concurrency_tickets`, `back_log`, ...
- Look at these settings only if you do know what you're doing

Conf for master and slaves

- If same HW, conf should be approx. the same
- You can take shortcuts on slaves for perf.
 - No binary logging
 - Relaxed InnoDB durability
 - `read_only` parameter to avoid most accidental writes
- If you promote a slave, don't forget to change its configuration!

Schema

- For many: “normalization kills performance”
 - Because it increases the number of joins
 - And joins increase the number of random ops
- But
 - Only true for specific hot spots in high load apps
- Other performance killer designs
 - Having everything in a BLOB column
 - Entity-Attribute-Value

(De)Normalization

- Normalize and index correctly first
 - Be sure to understand all the benefits of indexing: filtering, sorting, covering
 - SSDs mitigates cost of random ops for normalized schemas
- Denormalize if some queries become slow
 - Some combinations of filtering/sorting can't be solved with indexes
 - Some queries will have to read lots of data
 - Often the case with COUNT(*) or GROUP BY queries

Other useful tips

- Don't be too generous when sizing
 - The whole size is used for implicit temp tables
 - Use tinyint instead of int if possible
 - Use varchar(30) instead of varchar(255) if possible
- InnoDB tables and primary keys
 - The PK holds the data (clustering index)
 - It is implicitly included in all secondary keys
 - So set a PK explicitly, as short as possible

A common problem

- Let's design the user table of a social network

```
CREATE TABLE user(  
  user_id INT AUTO_INCREMENT,  
  login VARCHAR(30),  
  password VARCHAR(30),  
  PRIMARY KEY(user_id)  
)ENGINE = InnoDB
```

- Over time, you will add contact information(phone, address, ...), preferences, etc
 - The table will quickly grow very big, queries will be slow
 - Fragmentation and slow ALTER TABLEs will be the norm
- What can you do?

A solution

- Let's create 2 new tables:

```
CREATE TABLE user_info(  
  user_info_id INT AUTO_INC.,  
  value VARCHAR(30),  
  PRIMARY KEY(user_info_id)  
)ENGINE = InnoDB;
```

```
CREATE TABLE user_has_info(  
  user_id INT,  
  user_info_id INT,  
  PRIMARY KEY(user_id,user_info_id)  
)ENGINE = InnoDB;
```

- Now adding a property is easy and generic:
 - Insert a line in user_info to register the property
 - Insert a line in user_has_info for each user_id having the property

Pros and cons

- Benefits
 - No need to alter the user table anymore
 - Development is easier
- Drawbacks
 - The user_has_info table will not scale well
 - Every user will have 10s of rows in the table
 - Some queries are difficult to write efficiently
 - List of the users not having property xxx

Final thoughts

- In this particular situation:
 - Some properties can go to the user table
 - Some properties can go to the user_info table
 - Some properties will need dedicated tables
- As your application grows, you will have to deal with this kind of problem
 - Be creative!
 - But try not to over-engineer

Queries

- If my queries are slow, what can I do?
 - “No pb, just upgrade your HW”
- Not always the solution
 - Cost
 - Physical limit of your HW
 - Contentions in MySQL
 - What if your queries are just waiting to sth external to the database?

Improving queries

- Means improving response time
 - Low and stable response time is your goal
 - Stability is often overlooked
- How to make a query run quicker?
 - Remove unnecessary work
 - Run necessary work as efficiently as possible

Remove unnecessary work

- Select only columns you really need
 - Exception: `SELECT *` can be useful for caching purposes
- Select only rows you really need
 - Use a `LIMIT N` clause if you want the top N results
- Use caching to offload the DB
 - The fastest query is the query you don't run

Optimize necessary work

- All access types are not equal
 - type column in EXPLAIN output: `index` (index scan) and `ALL` (full scan) are the worst ones
 - An index scan can be order of magnitudes slower than a full scan
- Rewriting queries
 - Subqueries may perform very badly (much better in 5.6 and in new versions of MariaDB)
 - Use `INNER JOIN` instead of `LEFT JOIN` when possible

Indexing

- Correct indexing is key to good performance
 - Not as easy as it seems
 - Too few idx kill perf, too many idx kill perf too...
 - Tools can help you find improvements
 - pt-duplicate-key-checker will find duplicate idx
 - pt-query-digest and pt-index-usage will help you find slow queries
 - user_statistics feature in MariaDB and Percona Server is useful to identify useless indexes

Hardware

- In short
 - Use commodity: it doesn't mean junk!
 - 24 cores, 128GB RAM, 640 GB SSD is still commodity
- CPU
 - No parallelization of query execution, so fast CPUs are needed for good response times
 - MySQL scalability has greatly improved

RAM

- Memory
 - MySQL uses memory to cache index/data and for buffers
 - If possible you want your working set to be cached in memory
- Your working set is the fraction of your data that is accessed frequently
 - Not easy to assess it (1% to 100% of your data)

Disks

- The world is moving to flash storage
 - Much more IOPS and lower latency than HDDs
 - Especially good at random IOPS
 - SLC (performance) vs MLC (cost, capacity)
- You may need to update your my.cnf to take advantage of Flash
 - Percona Server and MariaDB offer the most flexibility
 - MySQL 5.6 is catching up

Flash technologies

- SSD
 - SATA interface: drop-in replacement for HDDs
 - You need RAID like for HDDs
- PCIe devices
 - Needs special drivers
 - Better performance than SSD
 - No need for RAID

SSD/HDD usage patterns

- Mixing HDDs and Flash
 - Flash for hot data / HDDs for archives
 - Flash for data files / HDDs for InnoDB redo logs
- More RAM is often the best, but not if
 - The amount of RAM is limited
 - You have a high-throughput write workload

HW for master and slaves

- Things to keep in mind
 - Slaves must be able to keep up with the master's write load
 - If you promote a slave, it should be as powerful as the master
- Common choices
 - Same HW for master and slaves
 - New HW for master, master's old HW for slave
 - Flash for slaves, HDDs for master

Backup/Recovery

- Typical mistake: focus on backup
 - A backup you can't restore is useless
 - So focus on restoring instead of backing up!
- Different needs
 - For backups: low-impact required, quick if possible
 - For restores: quick required, low-impact if possible

Defining the right strategy

- First you should know your RPO and RTO
 - RPO: Recovery Point Objective, ie how much data can you lose?
 - RTO: Recovery Time Objective, ie how much downtime can you afford?
- Relaxed RPO and RTO means you will have more options to choose a tool

Different kinds of backup

- Logical backups
 - Text files, easily readable, editable (grep, sed, awk)
 - Flexible – see list of options for mysqldump
 - Restoring is VERY slow
- Raw backups
 - Binary files
 - Restoring the whole backup is fast
 - Often not obvious to restore a single table/db

mysqldump

- A must for small databases (max ~ 10GB)
 - Very flexible
 - Backups are fast
 - Restores are not too slow
- But totally unusable for larger databases
 - Restore time is a showstopper

XtraBackup

- Online raw backups with low-impact
- Full / incremental backup
- Moving single tables from server to server
- Parallel backups
- Streaming backups
- In active development
- And more...

Instrumentation/monitoring

- Monitoring/alerting will warn you when sth is wrong
 - Nagios, Zabbix...
 - Spend some time designing meaningful checks
- Graphing/trending
 - Cacti, Munin...
 - Will help you identify why things have broken
 - Graph everything you can from CPU usage to size of the InnoDB buffer pool

Instrumentation inside MySQL

- EXPLAIN
- SHOW PROFILE
- SHOW GLOBAL STATUS
- SHOW ENGINE INNODB STATUS
- SHOW MUTEX STATUS
- INFORMATION_SCHEMA
- PERFORMANCE_SCHEMA

Instrumentation outside MySQL

- vmstat, iostat, mpstat...
- top, free...
- innotop
- Percona Toolkit
- Learn at least how to use some of the tools
 - Sometimes you need realtime diagnostics
 - Monitoring tools always have a lag

-
- Thanks for attending!
 - Q & A



stephane.combaudon@percona.com

We're Hiring! www.percona.com/about-us/careers/