Booking.com

Copying your databases

Percona Live 2021

Nicolai Plum – Booking.com Database Engineering

Topics

When and why?

Where from?

MySQL-side methods

Storage-side methods

Comparisons and Recommendations

MySQL at Booking.com

- In use 15+ years
- Thousands of instances
- Hundreds of replication chains
- Data storage across our business

Database copies – Why?

- New instances
- Replacements
- Upgrades
- **Backups**

Database copies – Where From?

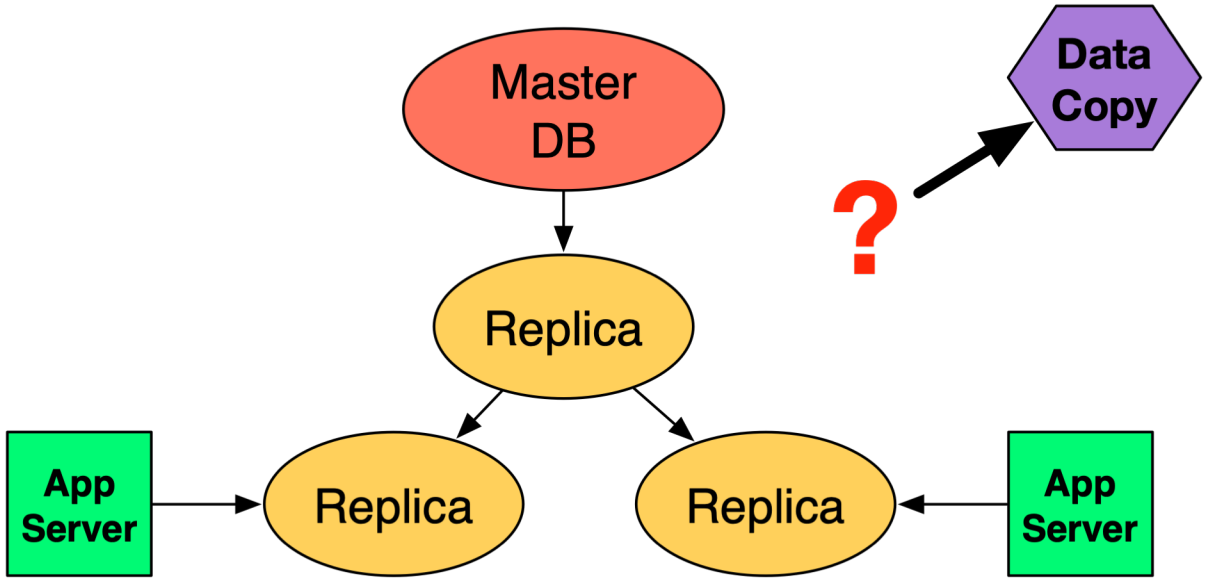

Primary master

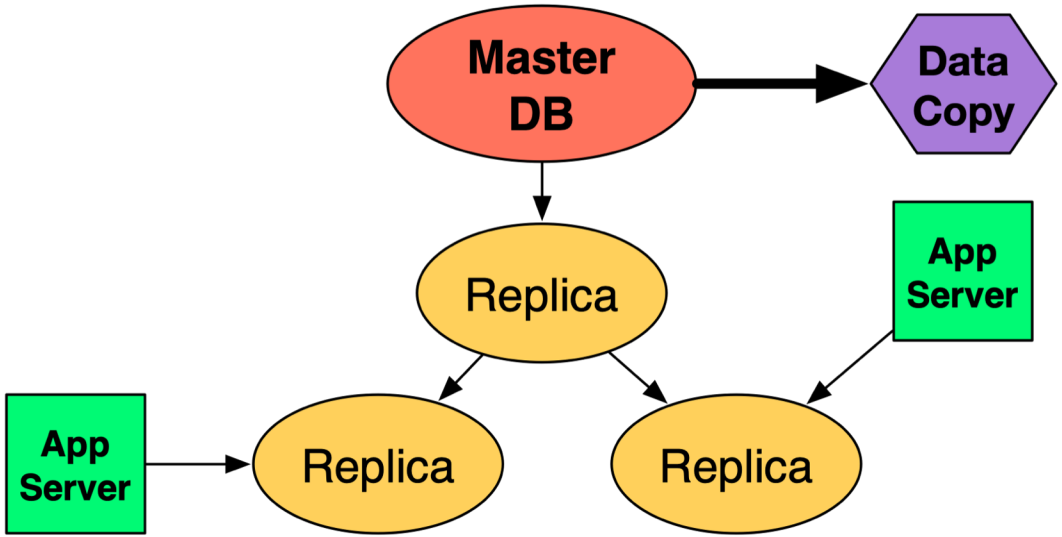

Intermediate/standby replica

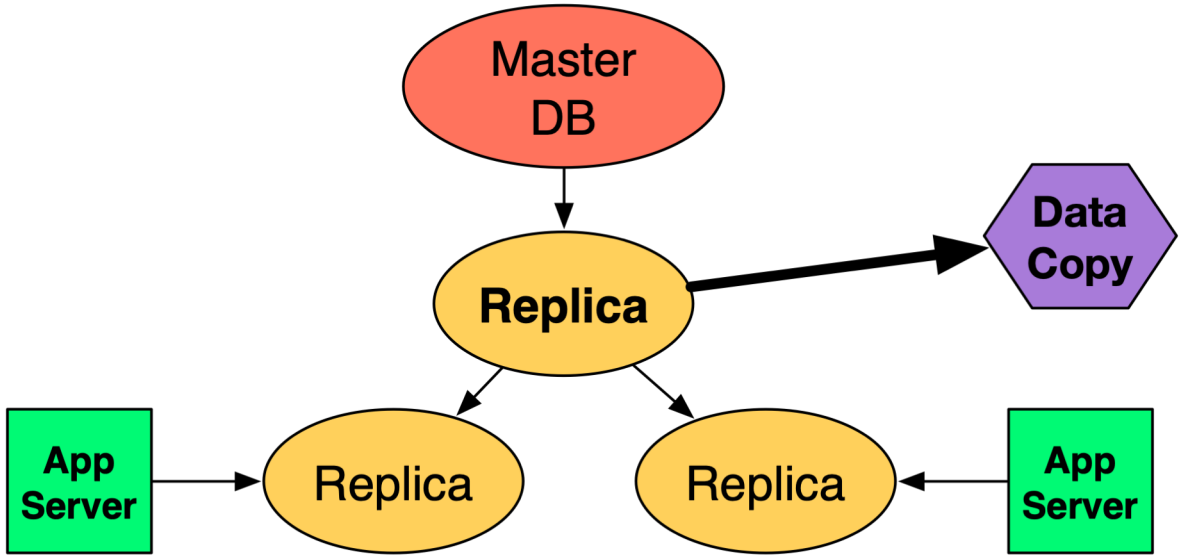

Shared replica serving users

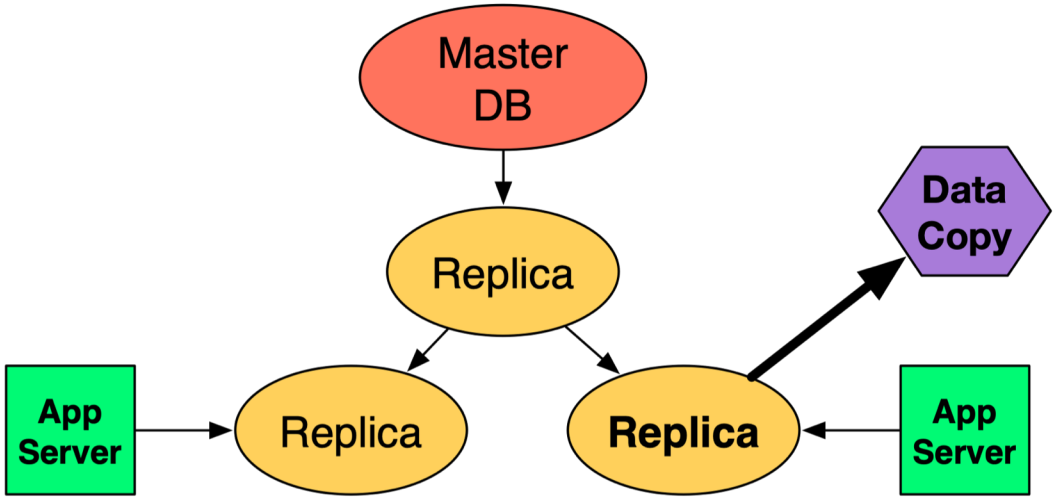

Dedicated replica

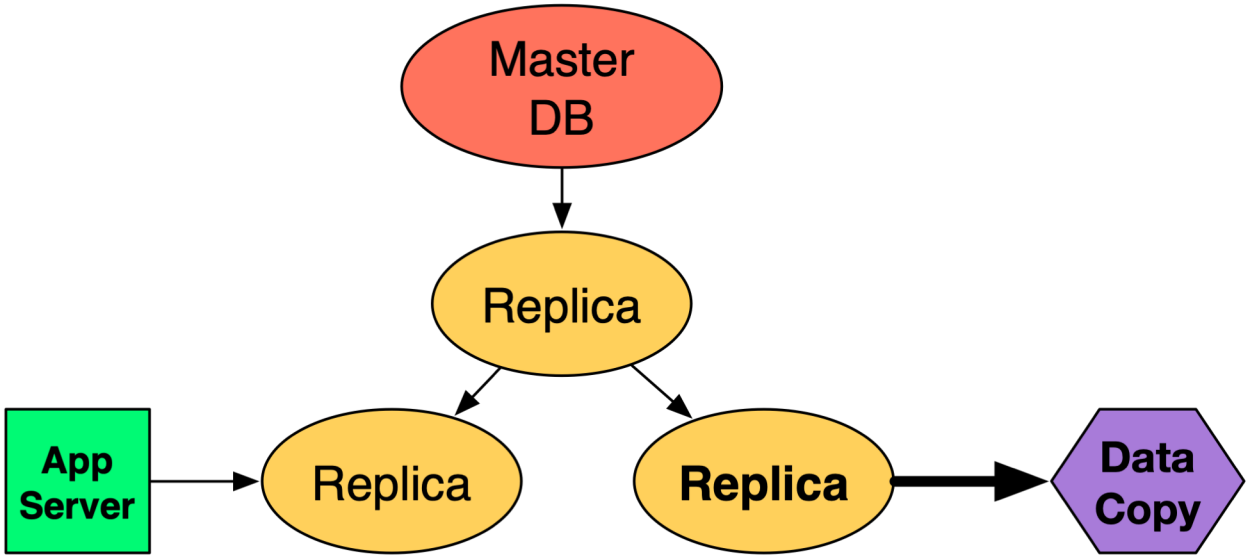

Comparison

	Efficiency (cost saving)	Speed	User Impact
Primary Master	High	Medium	High
Intermediate / standby	High	High	Medium
Existing replica	High	Medium	High
Dedicated replica	Low	High	Low

Recommendation

- First choice: Dedicated replica
- Second choice: Intermediate/standby

- For creating new chains: Primary master

MySQL-side methods

Online	Offline
mysqldump style	file rsync
xtrabackup style	ZFS send/receive
Native Cloning	

Backup & restore

mysqldump / mariadump

- SQL: DDL + DML + GTID
- Very flexible
  - Including version downgrade and load to other vendors
- Very slow
- The last resort; don't even bother automating

xtrabackup / mariabackup

- Tablespaces + log
- "prepare" after copy
- No downgrade in MySQL 8, MariaDB 10.5
- Fast enough for many uses

Native Cloning

- Oracle MySQL 8 only
- Exact version match only
- Simple command via SQL interface
- 2-3 times faster than xtrabackup

File copy

- Create consistent copy of database files
  - Shutdown for snapshot or copy
- Copy using rsync (daemon), pigz, …
  - or filesystem tools – ZFS send/receive
- Transport efficiency is a factor here

Storage-side methods

- Command storage to make snapshot of data
- Copy/clone/restore snapshot
- Start new instance
- Profit!

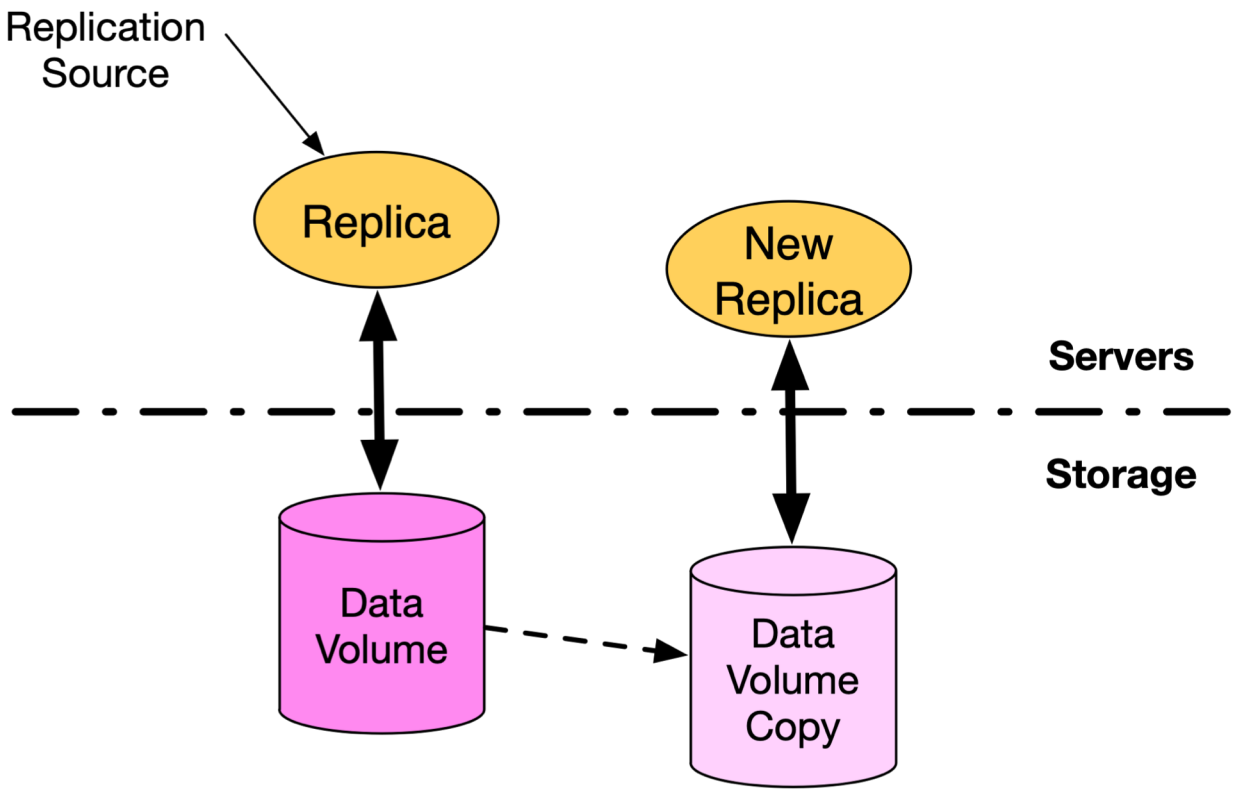
Replication
Source
Replica
New
Replica
Servers
Storage
Data
Volume
Data
Volume
Copy

Inconsistent snapshot

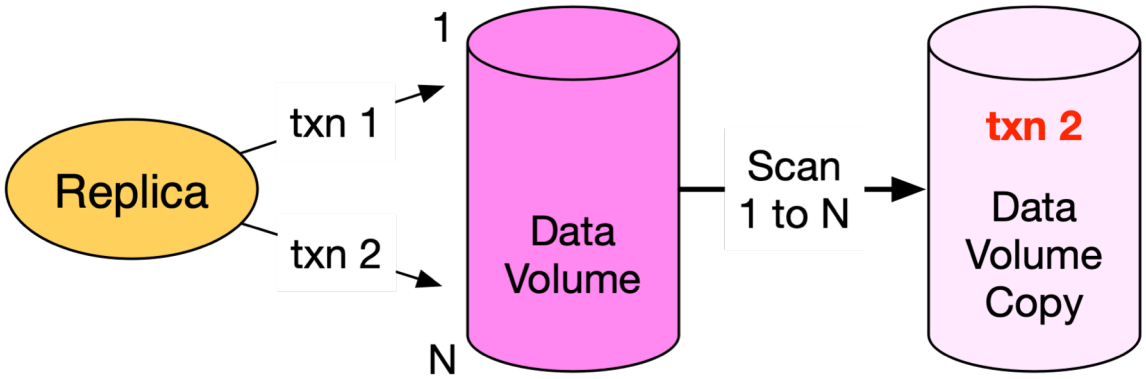

Consistent snapshot

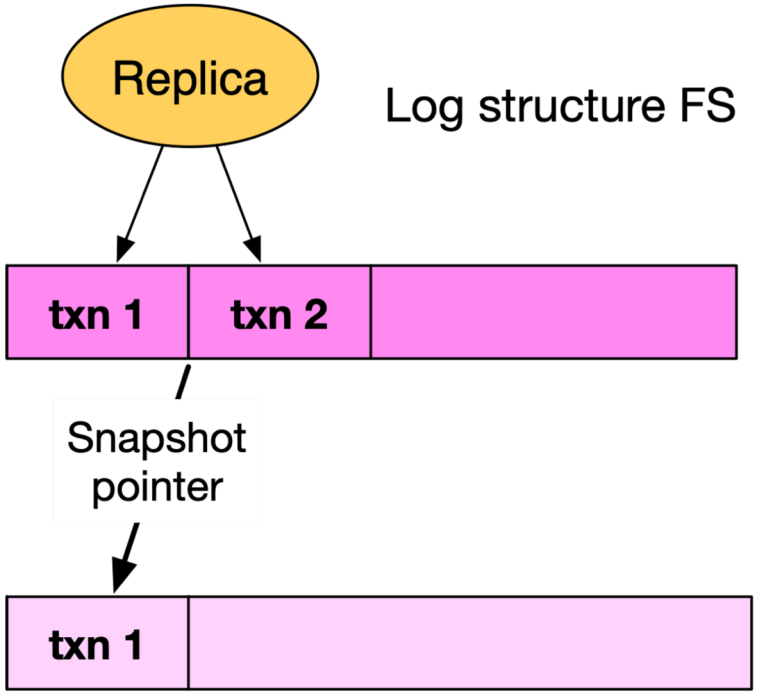

Shutdown for copy?

- Have you analysed the storage system in great detail?
- Do you feel lucky?
- No?
- Then shutdown MySQL for snapshot.

Advantages of storage copy

- Fast (usually)
- Features are already available (with most modern storage)
- Can be simpler than host copy

Limitations?

- Available immediately?
- Performance while copying?
- Ancestor volume must continue to exist?
- Limited number of descendents?
- Concurrency limits?
- Copy between storage clusters?
- Management complexity?

Backup and restore

- You have backup, and also restore.
- Restore is…
  - Not scalable (restore a backup to one target only)
  - Often not very fast
  - Often not made for intensive use 24/7
- Useful secondary copying method

Concurrency

- Hardware limits
  - 6Gb/s SATA, 10Gb/s Ethernet
- Mutually incompatible methods
  - Online vs offline
- Snapshot freshness
- Replicating during snap copying works fine

After copying data

- Attach the volume to server instance, if needed
- Change server_id and server_uuid
- Set up replication as needed
  - Use GTID and AUTO_POSITION
  - Or join to a group or cluster
- System automation setup
- Register with monitoring & metrics

Comparisons

	Flexible (version)	Speed	Online (source)?	Recommendation
MySQLdump	Very	Very slow	Yes	Last resort
xtrabackup	Yes *	Medium	Yes	First choice
Native	No *	Fast	Yes	Handy extra
File copy	Yes *	Fast	No	First choice
Storage copy	Yes *	Very fast (or maybe slow)	No	If available, first choice
Restore	Yes	Slow	Varies	Second choice

* Remember no downgrades on Oracle MySQL 8, Percona 8, and MariaDB 10.5

Recommendation

- Online copy will sometimes be necessary
- Implement automation for several options:
- First choices: xtrabackup, File copy, Storage copy (if available)
- Second choices: Native, Restore
- Do not:
  - Assume storage cloning is always possible
  - Rely on Native alone (makes upgrades painful)

?

nicolai.plum@booking.com