# The PlanetScale Binlog Server

Shlomi Noach
**PlanetScale**

MySQL Belgian Days 2026

# About me

Engineer at **PlanetScale**

Maintainer for **Vitess**, building **Neki**

Author for **planetscale/binlogsrv**

**github.com/shlomi-noach**

github.com/planetscale/binlogsrv

# The lede

Released as OSS, Apache 2

Not used in production (yet)

Maintenance considerations

# Agenda

What is a binlog server?

The PlanetScale binlog server

Vitess/PlanetScale notes

Community notes

# What is a binlog server?

# Binlog Server

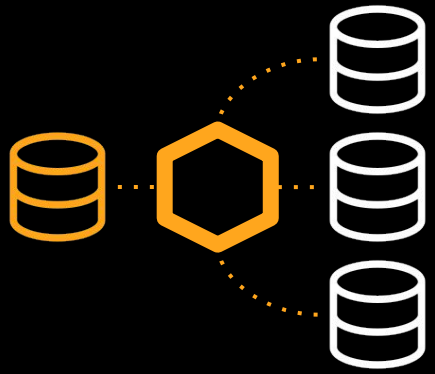

Conceived by Jean-François Gagné

A “repeater” for binary logs

Looks like a replica, looks like a server, but is neither

# MySQL Server: writes binlog

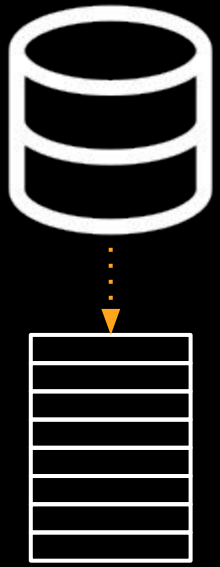

# Binlog Server: connect as replica

# Binlog Server: duplicate(*) binlog

# Binlog Server, bonus: archive

# Binlog Server use case #1

Duplicate binlogs

Possibly offload/archive for safekeeping

Advantage: does not contain data, has more space for binlogs

# Binlog Server use case #2

Act as semi-sync replica

# binlogsrv

Replica usage

```
binlogsrv
  --data_dir "${BINLOGSRV_DATADIR}"
  --port 3000
  --mysql_source_host "${MYSQL_SOURCE_HOST}"
  --mysql_source_port "${MYSQL_SOURCE_PORT}"
  --mysql_user "${MYSQL_SOURCE_USER}"
  --mysql_pass "${MYSQL_SOURCE_PASS}"
  --source_retry_interval 10s
  --source_heartbeat_interval 1s
  --auto_start_replica
  --semi_sync_replica_enabled
```

# binlogsrv

Replica usage, offload to S3

```
binlogsrv
  --s3
  --s3_bucket "my_bucket"
  --s3_prefix "my_path/my_keyspace/0/"
```

# binlogsrv

Replica notes

Near-exact binlog copy

Rotate/Format/heartbeat/...

Data agnostic

# binlogsrv

Directory layout

```
/pathto/binlogsrv
/pathto/binlogsrv/00000000000000000001.bin
...
/pathto/binlogsrv/00000000000000000007.bin
/pathto/binlogsrv/0000000001
/pathto/binlogsrv/0000000001/binlog_checksum
/pathto/binlogsrv/0000000001/previous-gtids
/pathto/binlogsrv/0000000001/replicating-binlog
/pathto/binlogsrv/0000000001/source-gtid-purged
/pathto/binlogsrv/0000000001/source-server-id
/pathto/binlogsrv/0000000001/source-uuid
/pathto/binlogsrv/0000000001/source-version
/pathto/binlogsrv/binlogsrv-server-id
/pathto/binlogsrv/binlogsrv-term
/pathto/binlogsrv/binlogsrv-uuid
/pathto/binlogsrv/mysql-bin.index
/pathto/binlogsrv/next-dump-gtids
/pathto/binlogsrv/primary-position
```

# Binlog Server: standalone

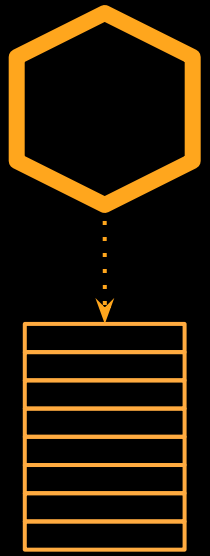

# Binlog Server: accept replica connections

# Binlog Server: send binlogs

# Binlog Server, bonus: fetch binlogs

# Binlog Server use case #3

Feed binlogs to MySQL servers

Common use case: catchup after restoring from backup

Advantage: does not overload production servers

Advantage: manages binlog storage

# binlogsrv

Server usage

```
binlogsrv
  --data_dir "${BINLOGSRV_DATADIR}"
  --port 3000
  --mysql_source_host "${MYSQL_SOURCE_HOST}"
  --mysql_source_port "${MYSQL_SOURCE_PORT}"
  --mysql_user "${MYSQL_SOURCE_USER}"
  --mysql_pass "${MYSQL_SOURCE_PASS}"
  --mysql_serve_port 33060
  --auto_start_server
```

# Binlog Server use case #1 + #3

Point in time recovery

# Binlog Server: relay

# Binlog Server: relay

# Binlog Server: relay

# Why?

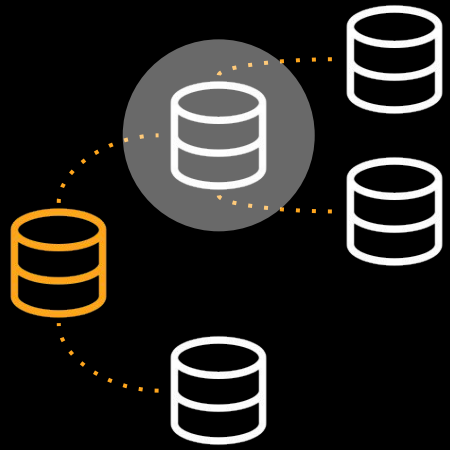

# MySQL relay / intermediate replica

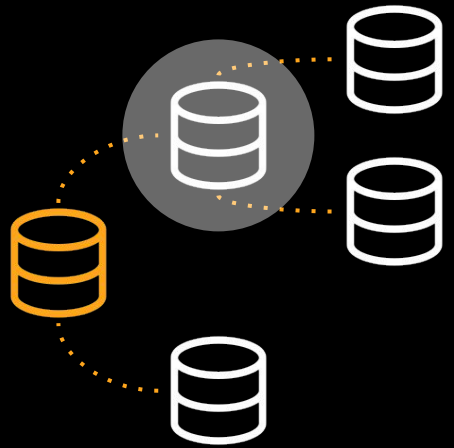

# MySQL relay / intermediate replica

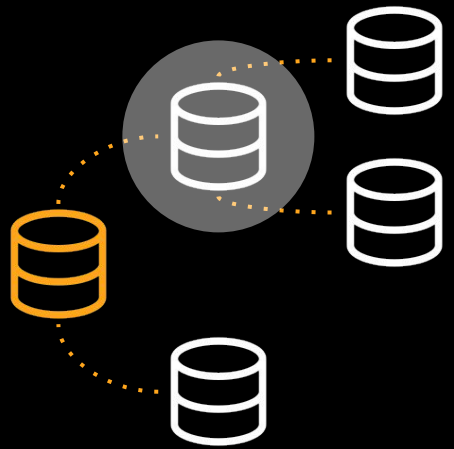

- Disk space

# MySQL relay / intermediate replica

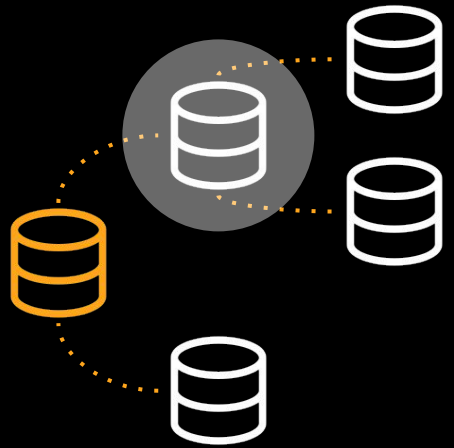

- Disk space
- Lag

# MySQL relay / intermediate replica

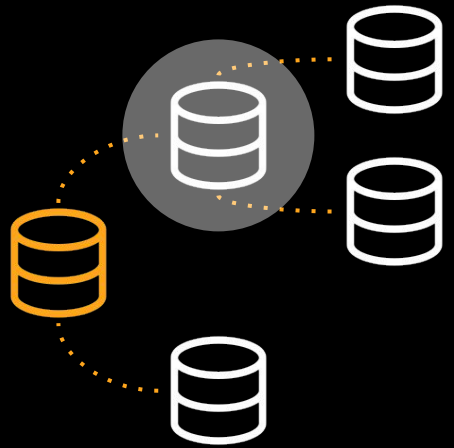

- Disk space
- Lag
- Rewrite

# MySQL relay / intermediate replica

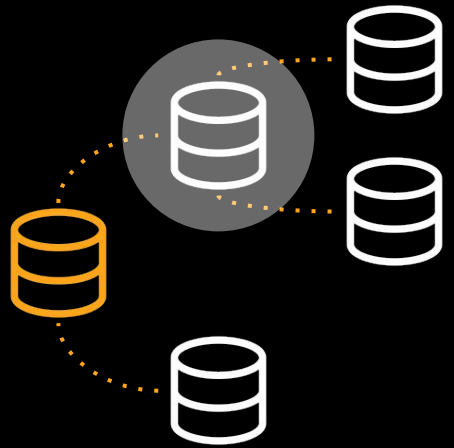

- Disk space
- Lag
- Rewrite
  - Different version

# MySQL relay / intermediate replica

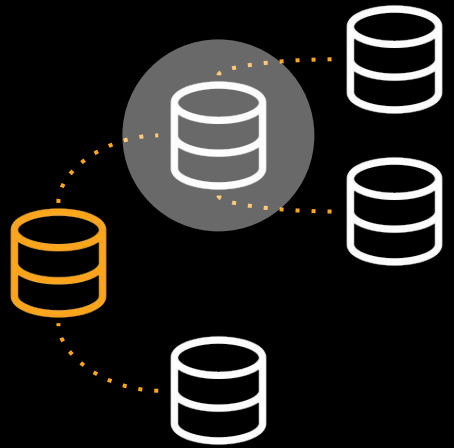

- Disk space
- Lag
- Rewrite
  - Different version
  - Risk of different content

# MySQL relay / intermediate replica

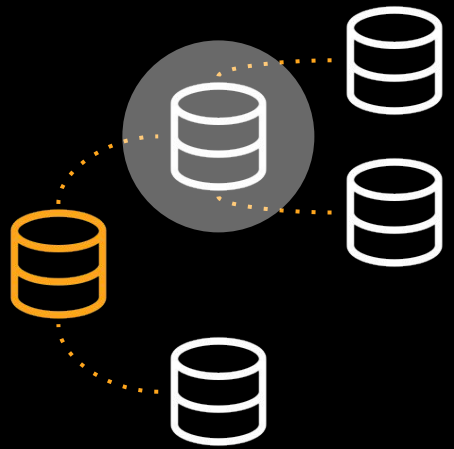

- Disk space
- Lag
- Rewrite
  - Different version
  - Risk of different content
  - Different log names/rotation(*)

# MySQL failover

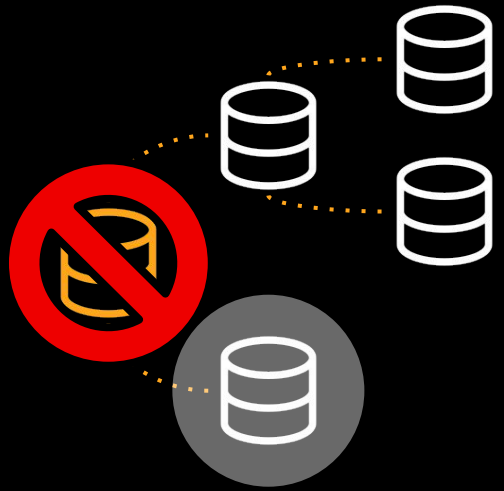

Possibly most up-to-date replica not a good candidate

- Higher version
- Errant GTIDs
- Lagging

# binlogsrv failover

# binlogsrv failover

# binlogsrv failover

# binlogsrv failover

# Binlog Server use case #4

Play a role in failover scenario

Requires external failover orchestration

# binlogsrv failover

Pick binlogsrv

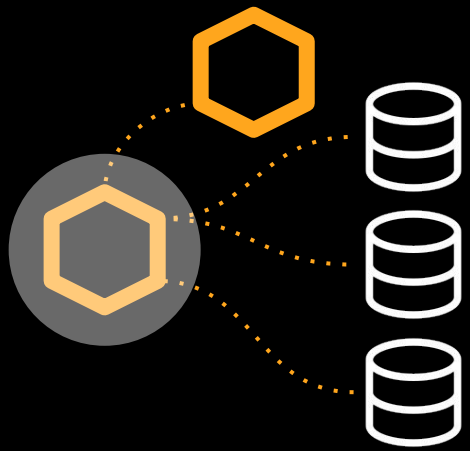

# binlogsrv failover

Pick replica to promote

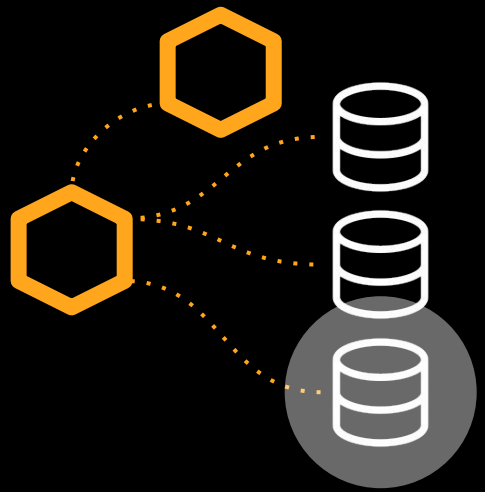

# binlogsrv failover

Promote

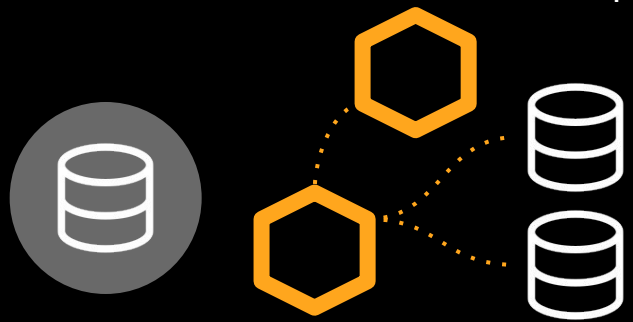

# binlogsrv failover

Then what?

# binlogsrv

- Attach to new replication source
- New “term”
- Potential duplication of binlog data
- Consecutive binlogs

# binlogsrv

Directory layout

```
/pathto/binlogsrv/0000000001
/pathto/binlogsrv/0000000001/binlog_checksum
/pathto/binlogsrv/0000000001/previous-gtids
/pathto/binlogsrv/0000000001/replicating-binlog
/pathto/binlogsrv/0000000001/source-gtid-purged
/pathto/binlogsrv/0000000001/source-server-id
/pathto/binlogsrv/0000000001/source-uuid
/pathto/binlogsrv/0000000001/source-version
/pathto/binlogsrv/0000000002
/pathto/binlogsrv/0000000002/binlog_checksum
/pathto/binlogsrv/0000000002/previous-gtids
/pathto/binlogsrv/0000000002/replicating-binlog
/pathto/binlogsrv/0000000002/source-gtid-purged
/pathto/binlogsrv/0000000002/source-server-id
/pathto/binlogsrv/0000000002/source-uuid
/pathto/binlogsrv/0000000002/source-version
/pathto/binlogsrv/0000000003
/pathto/binlogsrv/0000000003/binlog_checksum
```

# binlogsrv

Directory layout

```
/pathto/binlogsrv/00000000000000000001.bin
/pathto/binlogsrv/00000000000000000002.bin
/pathto/binlogsrv/00000000000000000003.bin
/pathto/binlogsrv/00000000000000000004.bin
/pathto/binlogsrv/00000000000000000005.bin
/pathto/binlogsrv/00000000000000000006.bin
/pathto/binlogsrv/00000000000000000007.bin
/pathto/binlogsrv/00000000000000000008.bin
/pathto/binlogsrv/00000000000000000009.bin
/pathto/binlogsrv/00000000000000000010.bin
/pathto/binlogsrv/00000000000000000011.bin
/pathto/binlogsrv/00000000000000000012.bin
/pathto/binlogsrv/00000000000000000013.bin
```

# binlogsrv

- GTID based source switchover
- Requires same GTID “family”
- Good for a single replication tree

# Vitess

MySQL-compatible, horizontally scalable, cloud-native database clustering system.

- CNCF graduated project
- Open source, Apache 2.0 licence
- Contributors from around the community

# Vitess

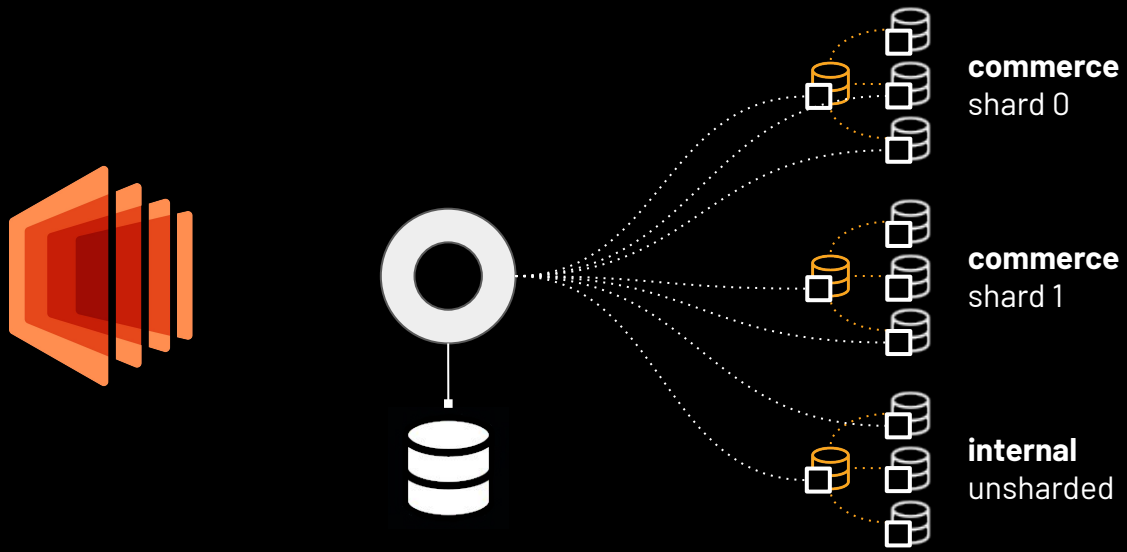

# binlogsrv

Vitess config

```
binlogsrv
  --data_dir "${BINLOGSRV_DATADIR}"
  --port 3000
  --grpc-port 3001
  --service-map grpc-tabletmanager
  --topo_implementation etcd2
  --topo_global_server_address localhost:2379
  --topo_global_root /vitess/global
  --keyspace commerce
  --shard 0
  --tablet_path zone1-0000000041
```

# Vitess/PlanetScale notes

# Binlog Server use case #5

Filtering

# Community

PERCONA

MySQL Open Source

Subscribe to RSS Feed

# Let's Rebuild the MySQL Community Together

November 26, 2025

Vadim Tkachenko

PERCONA

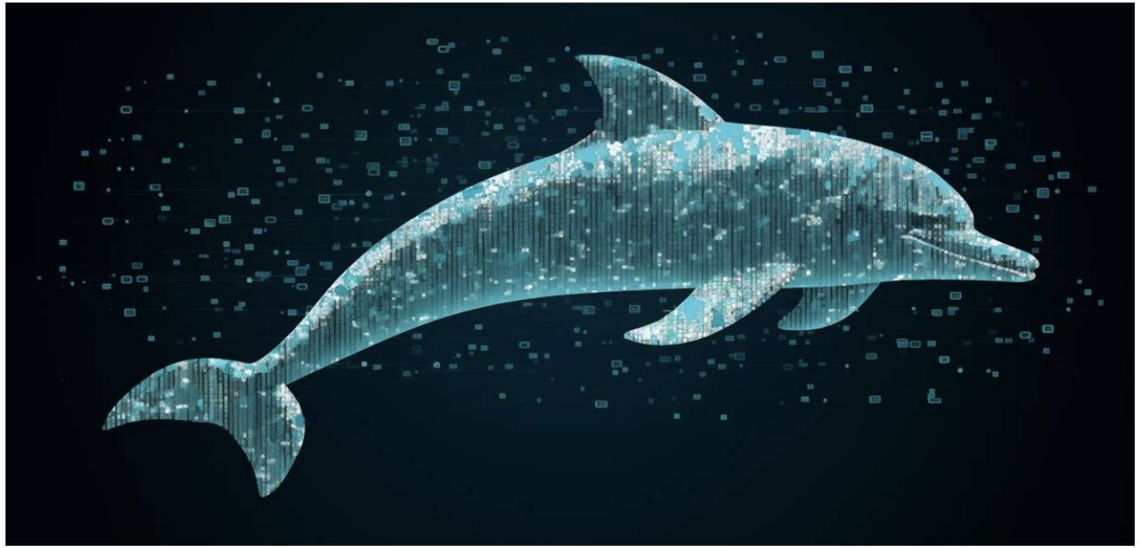

Insight for DBAs | MySQL | Percona Software

Subscribe to RSS Feed

# Building the Future of MySQL: Announcing Plans for MySQL Vector Support and a MySQL Binlog Server

November 25, 2025

Dennis Kittrell

# Closing thoughts

# Thank you!

Reach out on the Vitess Slack workspace

https://vitess.io/slack