



MariaDB Point-in-Time-Recovery (PiTR)

FOSDEM 2022, virtual

Oli Sennhauser

CTO, FromDual GmbH


<https://www.fromdual.com/presentations>

How are you doing backups?

- **Most of us do backups? Why?**
 - **Hardware gets broken**
 - **You are too slow to fix hardware! This can take hours to days to recover! → HA**
 - **Logical errors**
 - **Most common error type!**
e.g. DROP TABLE . . .
- **How are we doing backup?**
 - **Logical: mariadb-dump**
 - **Physical: mariadb-backup**

Point-in-Time-Recovery (PITR)

- Some of us do several backups a day (pulsed backup):



- How many hours between the backups?
- But what is in between?
- This is where PiTR comes into play!
 - Recover ALL changes since LAST full backup!
By replaying the binary logs...

Binary Logs

- We must record all the DB changes after the last backup
 - Or better: All the changes at all
- This is done with enabling the "Binary Logs"
= Change Log, Journal of the DB

```
# /etc/mysql/mariadb.conf.d/99-fromdual.cnf

[mariadb]

log_bin          = /var/log/mariadb/binary-log
binlog_format = ROW    # Recommended, see later
```

- Enabling Binary Logs requires DB restart!



Logical Backup for PiTR

- We need a consistent Full Backup of the DB
- Our recommendation for Logical Backup:

```
# for only InnoDB tables!!!  
mariadb-dump --user=root --all-databases --flush-logs \  
--single-transaction --master-data=1 --flush-privileges \  
--quick --triggers --routines --events \  
--hex-blob > /var/backup/full_dump.sql
```



Physical Backup for PiTR

- We need a consistent Full Backup of the DB
- Our recommendation for Physical Backup:

```
mariadb-backup --user=root --backup --target-dir=/var/backup  
mariadb-backup --user=root --prepare --target-dir=/var/backup
```

- You need physical access to the DB files!

* Wish to the developers: --backup-and-prepare Option!!!

Restore backup

- Restore logical Backup

```
cat /var/backup/full_dump.sql | mariadb --user=root
```


- Restore physical Backup

```
rm -rf /var/lib/mysql/*  
mariadb-backup --copy-back \  
--datadir=/var/lib/mysql --target-dir=/var/backup  
chown -R mysql: /var/lib/mysql
```

* Wish to the developers: remember ownership of the files?

Point-in-Time-Recovery

- We need exact point in time of the backup in the Binary Log stream
 - Binary Log File and Position



Starting point for PiTR?

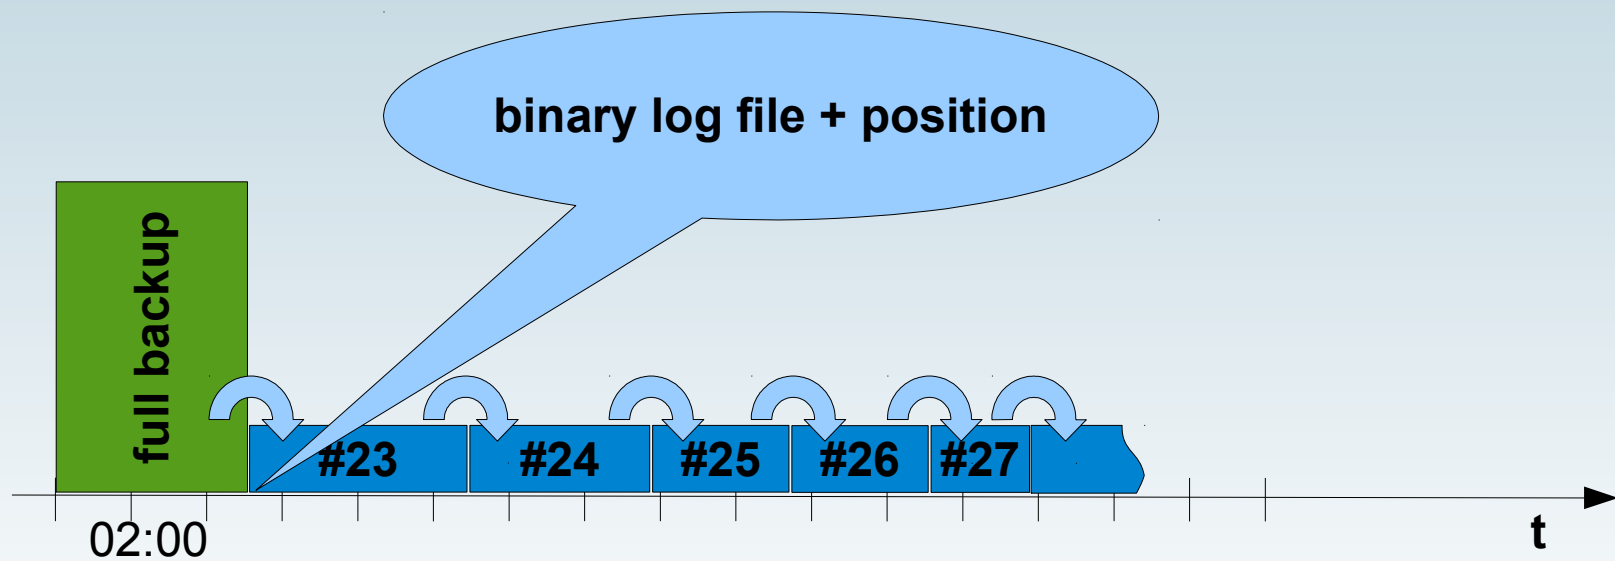
- `mariadb-dump:`

```
head -n 25 /var/backup/full_dump.sql
18 --
19 -- Position to start replication or point-in-time recovery from
20 --
21
22 CHANGE MASTER TO MASTER_LOG_FILE='binary-log.000023', MASTER_LOG_POS=429;
```

- `mariadb-backup:`

```
cat /var/backup/xtrabackup_binlog_info
binary-log.000023 429 0-3324-563
```

Point-in-Time-Recovery



Point-in-Time-Recovery


- How to do a PiTR:

```
mariadb-binlog --disable-log-bin \  
--start-position=429 \  
--stop-datetime='2022-02-01 13:59:00' \  
binary-log.000023 binary-log.000024 \  
binary-log.000025 ... | mariadb --user=root
```

- To be more precise: `--stop-position` can be used!
 - Needs `mariadb-binlog` utility to find exact stop position.
- First try against `less` to check if you got it right!

Advanced: Flashback

- Instead of doing all this:



- What about just fixing the logical error?
- Binary Log Flashback → Reverse DML operation (needs `binlog_format = ROW`)

```
mariadb-binlog --flashback --start-position=348622898 \  
--stop-position=349828120 binary-log.000080 \  
| mariadb --user=root test --force
```

Advanced: GTID PiTR

- New feature in MariaDB 10.8!

```
mariadb-dump --gtid --master-data=1 ...  
  
25 -- GTID to start replication from  
28 CHANGE MASTER TO MASTER_USE_GTID=slave_pos;  
29 SET GLOBAL gtid_slave_pos='0-3322-4';  
  
mariadb-binlog --gtid-strict-mode --start-position=0-3322-4 \  
--stop-position=0-3322-7 chef_mariadb-108_binlog.* | mariadb
```

- Stop GTID is inclusive!!! (--stop-position)

- Literature:

<https://mariadb.com/kb/en/mariadb-1080-release-notes/>

<https://jira.mariadb.org/browse/MDEV-4989>

Test your restore!!!

- **Personal advice: Test your restore!!!**
- **Why should you test your restore?**
- **Your restore might not work at all!?!**
 - **And you find it out much, much later...**
- **Restore can take a very long time.....**
 - **Verify your MTTR or RTO/RPO!**
 - **MTTR - Mean Time to Repair**
 - **RTO - Recovery Time Objective**
 - **RPO - Recovery Point Objective**
- **Getting experience!**

Thank you!



Questions ?

Discussion?

We have some time for a personal talk...

FromDual provides neutral and independent:

- Consulting
- remote-DBA
- Support for MariaDB and Galera Cluster
- Training

www.fromdual.com/presentations