



Mariabackup – too rarely used

FOSDEM 2021, Brussels / remote

Oli Sennhauser

MariaDB Consultant, FromDual GmbH

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



www.fromdual.com

About FromDual LLC

Support



MariaDB



Consulting

remote-DBA



Training





Table of Contents

Mariabackup

- › **Thoughts about Backup**
- › **Steps to do a Mariabackup**
- › **Performance differences**
- › **Point-in-Time-Recovery**
- › **Where has the time gone?**

Why do we need backups?

- **Hardware Crash**
 - **Does NOT happen that often!**
 - HA solution with Failover, is fast (< 10 s)
 - Restore, is not so fast (many minutes to hours!)
- **Setup of Slaves or Testing Systems**
 - **Typically no hard time constraints**
- **Logical Errors!!!**
 - **Most frequent issue**
 - They are spread very fast, no need to tune!

Terminology

- **Since MariaDB 10.4 / 10.5:**
- `mysqldump` → `mariadb-dump`
- `mysql` → `mariadb`
- `mariabackup` → `mariadb-backup`
- `mysqlbinlog` → `mariadb-binlog`
- `etc.`

Types of Backups

- Logical
 - **Dump: mariadb-dump**

```
foreach ( schema )  
  foreach ( table )  
    SELECT * FROM table INTO outfile;
```

- Physical
 - **Mariabackup: mariadb-backup**
 - **Snapshots (LVM, FS, VM)**

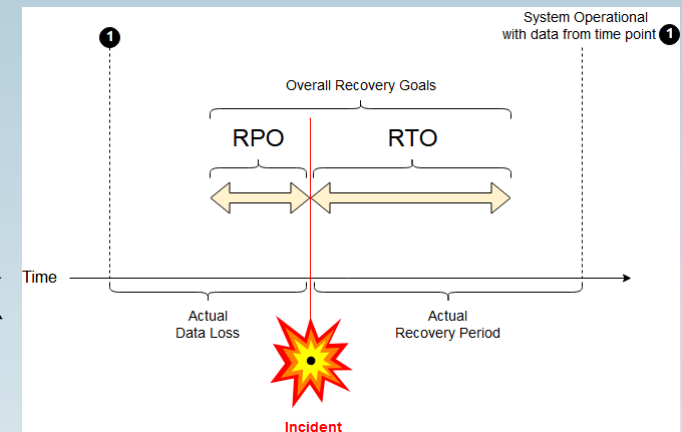
```
cp -r $datadir/* /mnt/tape
```

Restore/Recovery Metrics

- **Restore/Recovery Goal:**
 - **Manager: Immediately and NO loss of data at NO costs!**

- **Realistic Recovery Goals:**

- **Does Restore/Recovery work at all?**



- **Recovery Point Objective (RPO)**

- **How much data can we accept to loose?**

- **Recovery Time Objective (RTO)**

- **How long does a Restore/Recovery take?**

Granularity of Backups

- **Full Instance (`$datadir/*`)**
 - **Full Backup**
- **One or several schemas (`test, world, ...`)**
 - **Partial Backup, Schema Backup**
- **One or several tables (`City, Country, ...`)**
 - **Less than a schema or tables from different schemata**
 - **Partial Backup, Table Backup**
- **What do we need to Restore?**
 - **Full instance? One Schema only? One Table only?**
 - **And is this possible at all? And if yes how?**
 - **When we do the Backup we typically do not know what to Restore!**
 - **→ Full Backup**

Steps to do a Mariabackup

- Backup
 - Copies the files to backup location
- Prepare (old: apply-log)
 - Make backup consistent
 - Can be done at Backup or at Restore time!
- Full Restore
 - Shutdown Instance
 - `rm -rf $datadir/*`
 - Prepare (if not done during Backup yet)
 - Restore
 - Copies the files from backup location to `$datadir`
 - `chown mysql: $datadir/*` → Feature Request!
 - `restorecon -R $datadir` (SELinux)



Mariabackup steps

```
mariadb-backup --user=root --backup \  
--target-dir=/mnt/tape/backup
```

```
mariadb-backup --prepare \  
--target-dir=/mnt/tape/backup
```

```
rm -rf $datadir/*
```

```
mariadb-backup --copy-back \  
--datadir=/var/lib/mysql --target-dir=/mnt/tape/backup
```

```
chown -R mysql: /var/lib/mysql
```

Feature Requests for Mariabackup Developers:

- * `--backup-and-prepare`
- * Preserve ownership of files during `--copy-back/--move-back`

Performance differences

- Our little DWH: < 30 Gibyte
- Logical Backup: `mariadb-dump`:
 - Backup: **229 s**
 - Size (uncompressed): **17'056 Miabyte**
 - Restore: **2671 s**
- Physical Backup: `mariadb-backup`:
 - Backup: **180 s**
 - Size (uncompressed): **28'884 Miabyte**
 - Prepare: **94 s**
 - Restore: **172 s**

Point-in-Time-Recovery

- What for?
 - Reduce data loss (improve RPO)
 - Theoretically up to the last transaction!
- How?
 - Binary Logging (`log_bin = binlog`)
 - Full backup
 - Starting point (`xtrabackup_binlog_info`)

```
mariadb-binlog --disable-log-bin \  
--start-position=47196247 \  
--stop-datetime='2021-02-08 13:59:00' binlog.000023 \  
binlog.000024 binlog.000025 ... | mariadb --user=root
```

Where has the time gone?

- Backup (< 30 Gabyte): 180 s
 - Mostly NOT relevant, we have 24 hours time!
- Prepare: 94 s
 - Can take some time depending on the traffic.
- Restore: 172 s
 - More or less the same time like Backup
- Pre- and Post-Restore tasks: 60 - 300 s?
 - Shutdown, chmod, restorecon, Startup
- PiTR Begin and End point (RPO) evaluation: 15 min?
 - Begin is easy, End is tricky (time or transaction or end)...
- Communication and Synchronization overhead (30 – 120 min ???)
 - Stop application, start application. What to recover, until when?
- PITR itself...
 - Depends on traffic/volume (worst case more than 24 h!?!)
- RTO? 2 h (7200 s)? 4 h (14400 s)?

What is left?

- **Partial Backup/Restore**
 - Partial Table or Schema restore from Mariabackup full backup
 - <https://fromdual.com/partial-table-or-schema-restore-from-mariabackup-full-backup>
 - A bit more complicated!
- **Incremental Backup/Restore**
 - Even more complicated!
 - Consumes more time during restore (at least 2 passes)
 - <https://fromdual.com/mysql-enterprise-backup-incremental-cumulative-and-differential-backup>
- **Schema or Table PiTR**
 - <https://fromdual.com/mysql-table-point-in-time-recovery-from-mysqldump-backup>

Q & A



Questions ?

Discussion?

We have time for some personal talks...

- **FromDual offers neutral and independent:**
 - **Consulting**
 - **remote-DBA services**
 - **Support for MariaDB, MySQL and Galera Cluster**
 - **Trainings**

www.fromdual.com/presentations