# Why your backup strategy is probably wrong?

# Hello!

## I am Pep Pla

Father of three kids and three cats. In my spare time I'm a Consultant at Percona.

You can find me at @peppla

*Strategy: the skill of making or carrying out plans to achieve a goal.*

# All backup strategies are wrong

- Backups are not a goal.
- A goal is to protect your data.
- A goal is to ensure your business survival.
- A goal is legal compliance.

# Recovery Strategy

- Your should focus on recovery.
- You should have a Recovery Strategy.

**Recover**: To become successful or normal again after being damaged or having problems.

**Restore**: Return something or someone to an earlier good condition.

# Infinite Possible Incidents

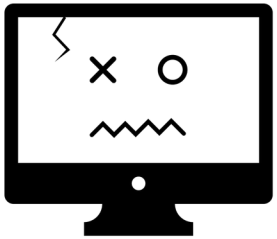

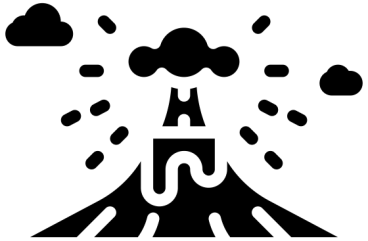





PERCONA

# Absolute protection is impossible





PERCONA

# Recovery Strategy

- Protection from a limited set of incidents
  - Higher probability.
  - Higher risk.
  - Lower protection cost vs risk.

- Design Backup-Restore/Availability Methodologies to ensure recovery from these incidents.

# RPO & RTO

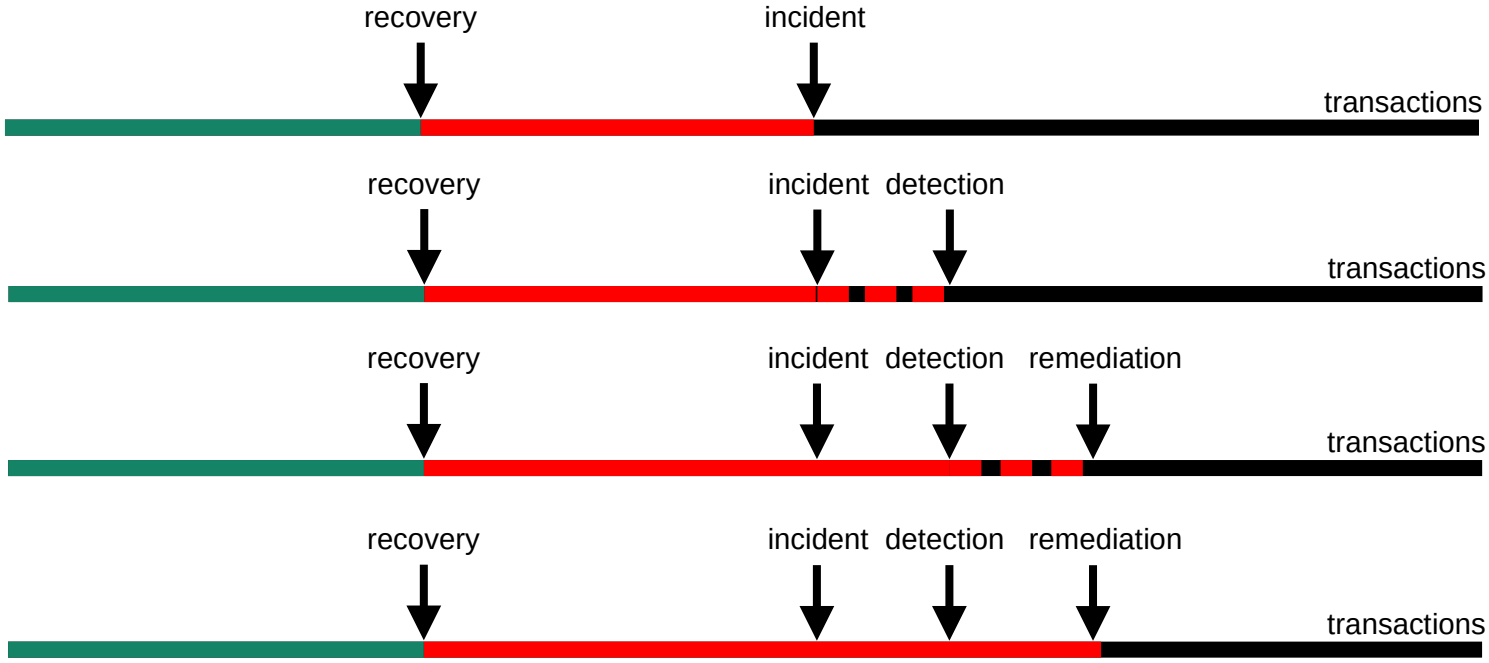

# RPO & RTO

## What does “Business” believes?

- RPO & RTO are worst case scenario
- If everything goes wrong we will be down X hours at most.
- If everything goes wrong we will lose a maximum of Y transactions.

## What SRE and DBA know?

- RPO & RTO are best case scenario
- If everything goes fine we will be down only X hours.
- If everything goes fine we will lose only Y transactions.

# The Schröendinger Backup

A backup can suffer of quantum superposition:

- It was completed successfully.
- You can’t restore it (in a relevant time).

Why a successful backup can be irrecoverable?

- Too long to restore.
- Bugs in backup/storage.
- Wrong backup.
- Infrastructure not available.
- ...

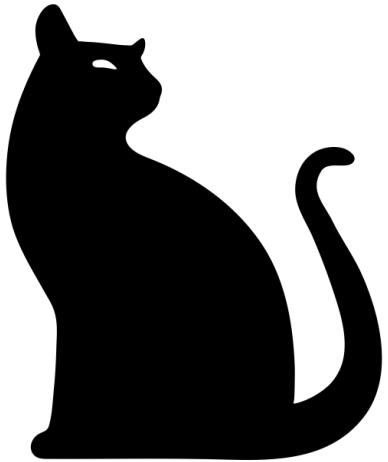

# How to know if the cat is dead or alive?

Open the box!

Perform restore and recovery tests.

As frequent as possible.

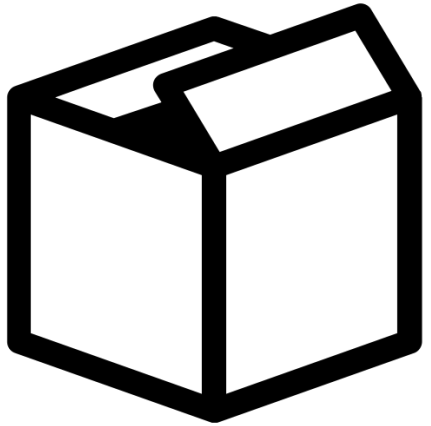

# Backups: Cold, Hot and Warm

The database consists of two types of files:

- dbfiles: contain all data. random access. read write. long term.
- logfiles: contain some data. sequential access. write. short term.

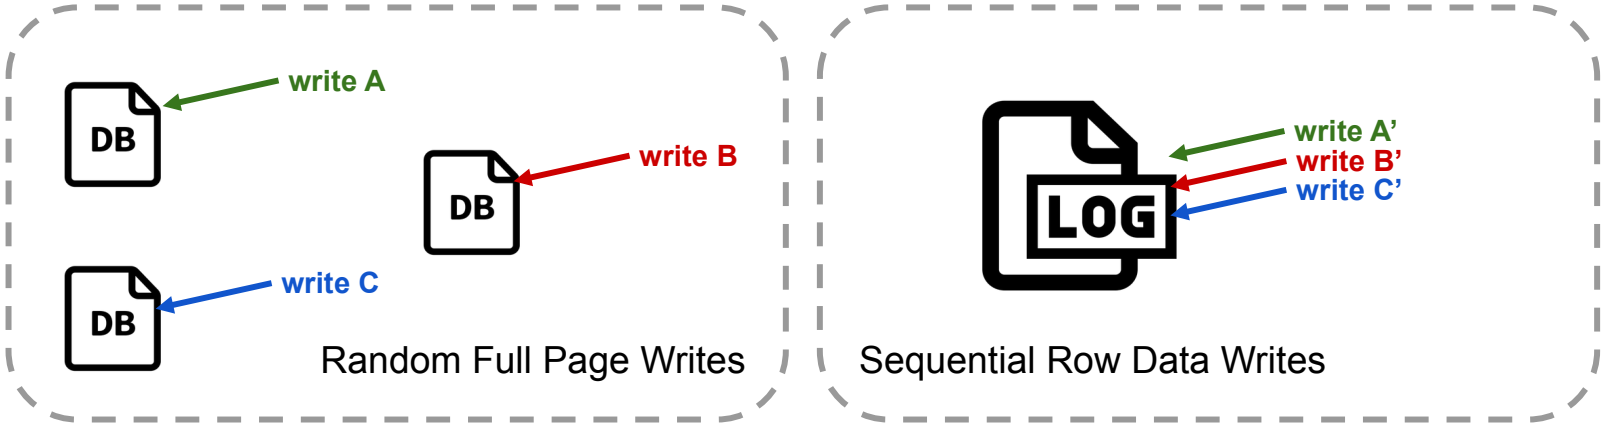

# Cold Backup

- Writes frozen
- We copy everything

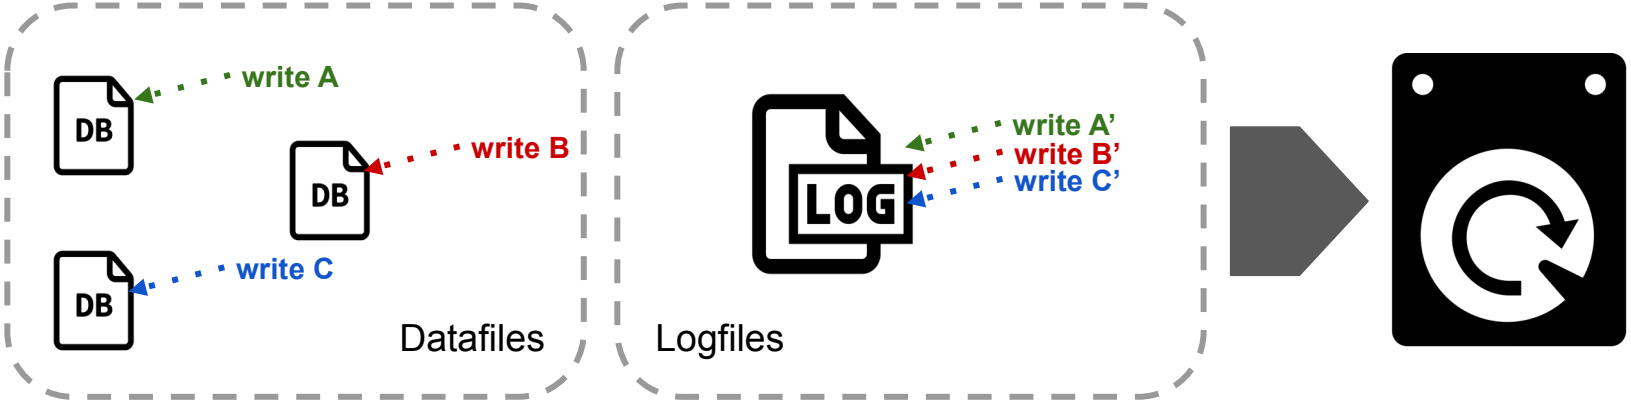

# Cold Backup

- Copy of files: cp, tar, rsync.
- Filesystem snapshot (LVM)
- Storage snapshot.

# Hot Backup

- Copy datafiles ignoring changes
- Copy the required redo log entries
- We have all the data required

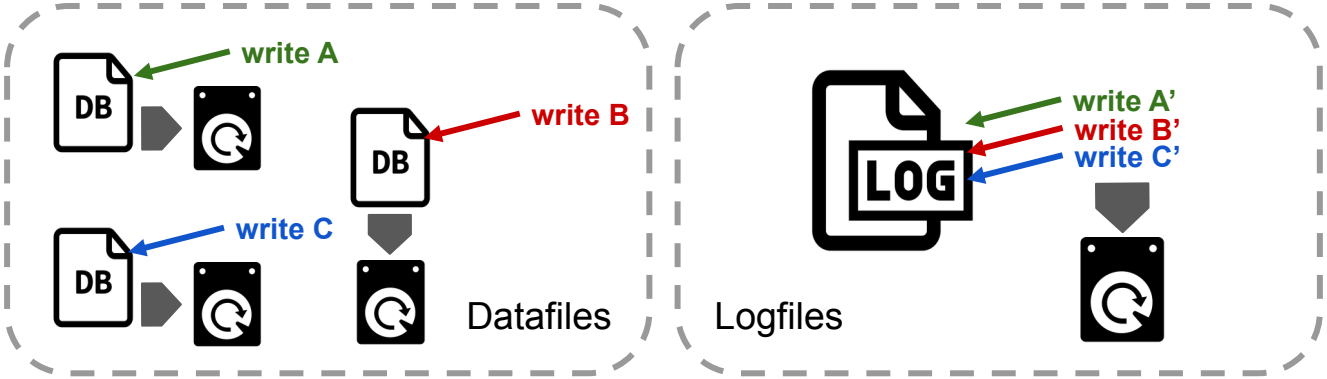

# Hot Backup

- Percona Xtrabackup
- MySQL Enterprise backup
- Clone Plugin

# Warm backup

- Cold backup without closing the database…
- … or closing it for a shorter period of time.
  - Snapshots without shutdown
    - Recovery needed and make sure snapshot is consistent.
  - Physical copy of unchanged while open. Sync modified while closed.
    - Multiple rsync. Last one with the database closed.

# The logical song

Generate SQL code to rebuild the database.

- mysqldump
- mysqlpump
- mydumper
- MySQL Shell dump utilities

# Logical backups

- Rebuilding the database can take a long time.
- Good for partial recovery.
- You have a human readable copy of data.

# Point in time recovery

- We use binary logs to recover until a certain point of time.
- We need 3 things:
  - Valid backup.
  - Position at the time of the backup.
  - Binary logs.
- Reduces data loss.
- Time consuming to apply the binary logs.

# Replication is not the solution

- Replication is not a backup.
- Protects from a reduced set of incidents.
- False sense of protection:
  - Delayed replicas… how much is enough?

# Don’t try this at prod!

- Test and document your recover procedures.
- Disasters always happen in the worst moment:
  - Time to decide.
  - No time to do research/learn.
  - “Business” needs to be involved in the decision making process.

# Thanks!

## Any questions?

You can find me at:

- @peppla
- pep.pla@percona.com