

MyISAM



InnoDB

# Database Engines



History and comparisons



# Agenda

- What is a Database Engine?
- MyISAM
- Aria
- InnoDB
- XtraDB
- LevelDB
- RocksDB
- SQLite
- Berkeley DB
- Demo (Switching Engines with MySQL)

# Database Engine

- Library that takes care of the on disk storage and CRUD
  - Can be as simple as key-value store
  - Or as rich and complex as full support ACID and transactions with foreign keys
- DBMS can use the database engine and build features on top (server, replication, isolation, stored procedures, etc..)
- Want to write a new database? Don't start from scratch use an engine.
- Sometimes referred as Storage Engine or embedded database

# Database Engines

- Some DBMS gives you the flexibility to switch engines like MySQL & MariaDB
- Some DBMS comes with a built-in engine that you can't change (Postgres)

# MyISAM

- Stands for Indexed sequential access method
- B-tree (Balanced tree) indexes point to the rows directly
- No Transaction support
- Open Source & Owned by Oracle
- Inserts are fast, updates and deletes are problematic (fragments)
- Database Crashes corrupts tables (have to manually repair)
- Table level locking
- MySQL, MariaDB, Percona (MySQL forks) supports MyISAM
- Used to be default engine for MySQL

# Aria

- Created by Michael Widenius
- Very similar to MyISAM
- Crash-safe unlike MyISAM
- Not owned by Oracle
- Designed specifically for MariaDB (mySQL Fork)
- In MariaDB 10.4 all system tables are Aria

# InnoDB

- B+tree - with indexes point to the primary key and the PK points to the row
- Replaces MyISAM
- Default for MySQL & MariaDB
- ACID compliant transactions support
- Foreign keys
- Tablespaces
- Row level locking
- Spatial operations
- Owned by Oracle

# XtraDB

- Fork of InnoDB
- Was the default for MariaDB until 10.1
- In MariaDB 10.2 [switched the default](#) for InnoDB
- “XtraDB couldn’t be kept up to date with the latest features of InnoDB and cannot be used.” ([link](#))
- [System tables](#) in MariaDB starting with 10.4 are all Aria



# LevelDB

- Written by Jeff and Sanjay from Google in 2011
- Log structured merge tree (LSM) (great for high insert and SSD)
- No transactions
- Inspired by Google BigTable
- Levels of files
  - Memtable
  - Level 0 (young level)
  - Level 1 - 6
- As files grow large levels are merged
- Used in bitcoin core blockchain, AutoCad, Minecraft

# RocksDB

- Facebook forked LevelDB in 2012 to become RocksDB
- Transactional
- High Performance, Multi-threaded compaction
- [Many features not in LevelDB](#)
- MyRocks for MySQL, MariaDB and Percona
- MongoRocks for MongoDB
- Many many more use it!

# SQLite

- Designed by D. Richard Hipp in 2000
- Very popular embedded database for local data
- B-Tree (LSM as extension)
- Postgres-like syntax
- Full ACID & table locking
- Concurrent read & writes
- Web SQL in browsers uses it
- Included in many operating systems by default

# Berkeley DB

- Developed by Sleepycat Software in 1994 (owned by Oracle)
- Key-value embedded database
- Supports ACID transactions, locks, replications etc..
- Used to be used in bitcoin core (switched to LevelDB)
- Used MemcacheDB

# DATABASE STORAGE ENGINES

## B-TREE



## LSM TREE



# Demo mySQL

- Spin up a mySQL docker container
- Create a table myISAM
- Create a table innoDB
- Connect to the database with Javascript
- See how transactions behave on myISAM vs innoDB

# Summary

- What is a Database Engine?
- MyISAM
- Aria
- InnoDB
- XtraDB
- LevelDB
- RocksDB
- SQLite
- Berkeley DB
- Demo (Switching Engines with MySQL)