

Updated on 14/02/2024

Sign up

# RocksDB training: The key-value store for fast storage

2 days (14 hours)

## Presentation

[RocksDB](#) training will teach you how to use this open source tool for fast key-value data storage. Use this database efficiently for input/output (I/O) workloads.

During our course, you'll learn how to open and close a database, as well as read and write operations and compaction filters.

You'll be able to store keys and values in various arbitrary byte arrays, sorting data by byte, key by key, or using a custom comparator.

Organize efficiently thanks to RocksDB's ability to adapt to different workloads.

Like all our training courses, this one will introduce you to the latest stable release and its new features ([RocksDB 7.8.3](#)).

## Objectives

- Mastering workload management with RocksDB
- Understanding RocksDB architecture
- Feeding high-performance distributed data systems
- Be able to use the storage engine
- How to use the various RocksDB tools

## Target audience

- Project managers
- Developers
- Architects

## Prerequisites

Knowledge of databases.

## RocksDB training program

### Introduction

- Introducing RocksDB
- Configuration parameters
- Limit the number of buttons to provide good performance
- Use of adaptive algorithms

### Assumptions and objectives

- Fast storage and server workload performance
- Search for effective points and range analyses
- Adjusting trade-offs for different workloads and materials
- Integrated support for tools and utilities to facilitate deployment and debugging in production environments
- Provide certain parameters for performance tuning
- Compatibility

### High-level architecture

- Organize all data in sorted order
- Understand the three basic constructs, memtable, sstfile and logfile
- Log file
- sst file
  - Sort data in an sst file to facilitate key searches

### Features

- Post families
- Updates
- Gets, iterators and snapshots
- Transactions
- Prefix iterators
- Persistence
- Data checksum

- Multi-stage compaction
- Compaction styles
- Metadata storage
- Avoid stalls
- Compacting filter
- Read-only mode
- Database debugging logs
- Data compression
- Table Cache
- I/O control
- ID BD

## Tools

- Empty all key values in an sst file with sst\_dump
- Store, retrieve or scan database contents with ldb
- Modify the number of configured database levels using MANIFEST

## Unit testing

- Use the make check command to run all unit tests
- db\_stress to validate large-scale data accuracy
- Adding fuzzers

## Companies concerned

This course is aimed at both individuals and companies, large or small, wishing to train their teams in a new advanced computer technology, or to acquire specific business knowledge or modern methods.

## Positioning on entry to training

Positioning at the start of training complies with Qualiopi quality criteria. As soon as registration is finalized, the learner receives a self-assessment questionnaire which enables us to assess his or her estimated level of proficiency in different types of technology, as well as his or her expectations and personal objectives for the training to come, within the limits imposed by the selected format. This questionnaire also enables us to anticipate any connection or security difficulties within the company (intra-company or virtual classroom) which could be problematic for the follow-up and smooth running of the training session.

## Teaching methods

Practical course: 60% Practical, 40% Theory. Training material distributed in digital format to all participants.

## Organization

The course alternates theoretical inputs from the trainer supported by examples and

brainstorming sessions and group work.

## Validation

At the end of the session, a multiple-choice questionnaire verifies the correct acquisition of skills.

## Sanction

A certificate will be issued to each trainee who completes the course.