+33 (0)1 48 10 09 46
formation@ambient-it.net
140 Avenue Jean Lolive 93500 Pantin - France

Updated on 19/03/2024

Sign up

Redis training

2 days (14 hours)

Presentation

Redis is a key-value NoSQL database, running in memory and providing optional persistence. The Redis data model provides advanced data structures to multiply use cases compared with databases of the same topology. Very close to the system, Redis stands out for its ability to provide a small memory footprint, low latency and very high reliability. Its aim is to deliver the highest possible performance within a highly constrained infrastructure.

Redis is a BSD-licensed, open-source, in-memory data storage framework. Typically used as an in-memory database, cache and/or message broker. It supports data structures such as strings, hashes, lists, sets, sorted sets with range queries, bitmaps, hyperlogs, geospatial indexes with radius queries and streams. Redis features built-in replication, Lua scripting, LRU eviction, transactions and different levels of disk persistence, and offers high availability via Redis Sentinel and automatic partitioning with Redis Cluster.

You can perform atomic operations on these types, such as adding to a string, incrementing the value in a hash, pushing an element to a list, calculating the intersection, union and difference of a set, or getting the member with the highest rank in a sorted set.

In order to achieve its outstanding performance, Redis works with an in-memory dataset. Depending on your use case, you can maintain this either by flushing the dataset to disk from time to time, or by adding each command to a log. Persistence can be disabled as an option, if you just need a feature-rich, networked in-memory cache.

Redis also supports trivial asynchronous master-slave replication, with a very fast, non-blocking first synchronization, automatic reconnection and partial resynchronization over the network. Other features of this DB cache include

including :

- Transactions
- Publish/Subscribe
- Lua scripting
- Limited life keys
- Key deletion by LRU / Maxmemory Configuration (similar to memcached with its cache line replacement algorithms)
- Automatic failover

This course covers the full range of Redis features, including the richness of its API through the Go language, its clustering mode and how to monitor Redis processes. It also covers techniques for optimizing your data model, with particular emphasis on memory footprint.

As with all our training courses, this one will introduce you to the latest stable release, Redis 7.

Objectives

- Know how to implement Redis through its main use cases
- Using the Redis API
- Redis optimization techniques

Target audience

- Developers
- Architects
- Data engineers

Prerequisites

Basic knowledge of a programming language.

Redis training program

Introduction to Redis

- Redis, memory cache
- Positioning in relation to other NoSql engines
- When to use Redis, when not to
- Redis key references
- What's new in version 7

Main data structures and manipulation

- String, List, Set, Hash and Sorted Set
- Main associated commands

Redis architecture

- Communication protocol and data format
- Atomicity of operations
- Start-up process
- Event loop & the different events
- Data durability
- Master-Slave replication

Developing with Redis

- Client access languages
- APIs in detail with Go
- Transaction management
- Server-side scripting with Lua
- Mocking Redis with miniredis

High availability and clustering

- Replicas and the data lifecycle
- Redis clustering
- High availability with Redis Sentinel

Data optimization

- Data expiration
- Pipelining & Multiple Argument commands
- Logical Types vs Physical Types
- Patterns applied to data design

Advanced data structures and organization

- Publish/Subscribe
- HyperLogLog
- BitMap
- Complex query problems

Traffic monitoring

- The "monitor" command

- Event analysis and History
- The different tools on the market

Going further (add-on module +0.5 day in-house)

- Security & Encryption
- Key design and configuration recommendations
- Extending Redis with Redis Modules
- How to debug Python or Lua scripts
- The main deployment topologies used by leading Redis references

Beyond Redis (add-on module +0.5 day in-house)

- What's not in Redis
- The Redis Labs enterprise solution
- Alternatives & potential successors

Companies concerned

This training 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 input from the trainer, supported by examples, with 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.