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

Updated on 25/09/2025

Sign up

Materia KV training

3 days (21 hours)

Presentation

Materia KV is a serverless key-value database designed for simplicity, high availability and scalability. Compatible with the Redis protocol, this transactional and distributed solution replicates your data synchronously across multiple datacenters, while offering pay-per-use billing and infrastructure-free operation.

Our Materia KV training course will guide you through integration with your applications, automation via CI/CD and Infrastructure as Code, secure access, observability and performance optimization. You'll learn how to manipulate JSON commands, set up efficient TTLs and define robust architecture patterns (cache-aside, write-through, anti-dogpile protections).

At the end of the course, you'll know how to deploy, monitor and industrialize Materia KV in a Data/DevOps/SRE context, from POC to production, with a resolutely practical and operational approach.

As with all our training courses, this one will introduce you to the latest documented edition and its recent evolutions.

Objectives

- Master the architecture (serverless, FoundationDB, replication)
- Provision and integrate Materia KV (CLI, Redis clients, CI/CD)
- Implement security, observability and SRE runbook
- Exploit TTL, JSON and cache patterns
- Optimize performance and costs

Target audience

- Data Engineers
- DevOps / SRE

Prerequisites

- Notions of NoSQL/key-value (e.g. Redis)
- Basic CI/CD / IaC experience
- Familiarity with a programming language (Go / Node / Python / PHP)

Materia KV training program

[Day 1 - Morning]

Architecture, uses and getting started

- Understanding Materia KV: serverless, transactional and distributed key-value base
- Overview of use cases (cache, sessions, feature flags, lightweight state, real-time)
- API/protocol overview (Redis compatibility, DynamoDB / GraphQL roadmap)
- Add-on creation, environment variables, TLS and tokens
- Connection via redis-cli, official clients and Clever Tools
- Hands-on workshop: provisioning an add-on, running GET/SET/DEL, testing TTL and persistence

[Day 1 - Afternoon]

Internals, sustainability and consistency

- FoundationDB foundations: ACID, replication, simulation-driven testing
- Synchronous multi-DC replication: RTO/RPO, consistency, latencies
- Data model and limits (sizes, supported commands, JSON)
- TTL/EXPIRE: expiration strategies, idempotent patterns
- Structuring key spaces (naming, prefixes, hashing)
- Practical workshop: consistency/resilience tests (concurrent writes, latency, TTL)

[Day 2 - Morning]

Security, costs and operations

- Authentication, TLS, token management, isolation
- Observability: metrics, logs, dashboards, alerting
- Scalability and pay-per-use: anti-patterns & best practices
- Comparisons: managed Redis vs. Materia KV
- Production safeguards: limits, retries, backoff, circuit-breakers
- Practical workshop: incident scenarios, instrumentation and alerts

[Day 2 - Afternoon]

CI/CD, IaC and application integration

- Terraform / OpenTofu provisioning, Clever Tools automation
- 12-factor models: config, secrets, blue/green, canary
- Node/Go/Python/PHP integration with Redis clients
- Patterns: cache-aside, write-through, lease, dogpile protection
- Tests/QA: datasets, contract testing, load
- Practical workshop: CI pipeline deploying app + KV add-on, load testing

[Day 3 - Morning]

JSON data and advanced models

- JSON commands: targeted reading, partial updating
- Serialized storage (JSON, MsgPack, Protobuf)
- Transactions: atomic sequences, contention and logical locks
- Simulated structures, batching and pipelining
- Optimization: key length, compression, differentiated TTLs
- Practical workshop: mini-API REST Materia KV JSON for a business use-case

[Day 3 - Afternoon]

Migration, governance and runbook

- Mapping existing systems and progressive migration strategies
- Governance: prefixes, conventions, quotas
- SRE runbook: incident SOPs, SLO/SLI, error budgets
- FinOps: estimating/optimizing cost-per-use
- Go-live checklist: security, observability, limits & testing
- Practical workshop: deployment, rollback and incident recovery

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 forthcoming course, 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 training: 60% hands-on, 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 is used to check that skills have been correctly acquired.

Certification

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