

Mastering InfluxDB: Time-Series Data Management and Analytics

Course Description

This intensive 5-day training program provides a deep dive into InfluxDB, the leading time-series database designed for high-performance data ingestion, storage, and real-time analytics. Participants will gain hands-on experience in designing, querying, and managing time-series data, while also exploring integrations with visualization tools, monitoring systems, and enterprise data platforms. The course emphasizes practical labs, real-world use cases, and advanced optimization techniques to prepare participants for production-grade deployments.

Duration : 5 Days

Pre-requisites

- Basic knowledge of databases (SQL or NoSQL concepts)
- Familiarity with Linux command-line operations
- Understanding of data modeling concepts
- Optional: Exposure to monitoring tools (Grafana, Prometheus)

Learning Objectives

By the end of this course, participants will be able to:

- Understand the architecture and core components of InfluxDB
- Design efficient time-series data models
- Perform advanced queries using InfluxQL and Flux
- Optimize ingestion pipelines for high-throughput workloads
- Implement retention policies, continuous queries, and downsampling
- Secure, scale, and monitor InfluxDB in enterprise environments
- Integrate InfluxDB with visualization and analytics tools

Content Coverage

Day 1 – Introduction & Fundamentals

- **Module 1: Time-Series Data Concepts**
 - Characteristics of time-series data
 - Use cases: IoT, monitoring, finance, DevOps

- Comparison with relational and NoSQL databases
- **Module 2: InfluxDB Overview**
 - InfluxDB ecosystem (OSS, Enterprise, Cloud)
 - Core components: storage engine, query engine, APIs
 - Installation and setup (Linux, Docker, Cloud)

Day 2 – Data Modeling & Ingestion

- **Module 3: Data Structures in InfluxDB**
 - Measurements, tags, fields, and time
 - Best practices for schema design
 - Cardinality considerations
- **Module 4: Data Ingestion Techniques**
 - Writing data via CLI, API, and client libraries
 - Batch vs. streaming ingestion
 - Integrating Telegraf for data collection
 - Handling high-throughput ingestion

Day 3 – Querying & Analytics

- **Module 5: Query Languages**
 - InfluxQL basics (SELECT, WHERE, GROUP BY)
 - Flux language introduction
 - Advanced Flux queries (joins, transformations, aggregations)
- **Module 6: Continuous Queries & Retention Policies**
 - Creating continuous queries for automation
 - Downsampling strategies
 - Managing retention policies for lifecycle management

Day 4 – Administration & Optimization

- **Module 7: Performance Tuning**
 - Indexing and query optimization
 - Storage engine internals

- Handling large datasets efficiently
- **Module 8: Security & User Management**
 - Authentication and authorization
 - Role-based access control
 - Encryption and securing endpoints
- **Module 9: Scaling InfluxDB**
 - Clustering and high availability
 - Sharding and replication
 - Backup and restore strategies

Day 5 – Integrations & Real-World Use Cases

- **Module 10: Visualization & Monitoring**
 - Integrating with Grafana
 - Dashboards and alerting
 - Real-time monitoring pipelines
- **Module 11: Enterprise Integrations**
 - Connecting with Kafka, Spark, and cloud-native platforms
 - Using InfluxDB in IoT and DevOps workflows
 - Case studies: predictive maintenance, system monitoring
- **Module 12: Capstone Project**
 - End-to-end lab: ingest, query, visualize, and optimize a real dataset