

# Time-Series Analytics with InfluxDB 2 and Grafana

## Flux queries, monitoring dashboards and version-aware lifecycle management

FA-0702 · 3 days

Data & Analytics · DevOps, Cloud & Infrastructure

---

## Introduction

This three-day course builds a bounded time-series monitoring example using InfluxDB 2.x, Flux and Grafana. Participants ingest simulated telemetry, explore schema choices, query time windows and build dashboards and alert rules.

The explicit 2.x scope preserves the source's Flux and task-based learning path. Flux is in maintenance mode and is not supported in InfluxDB 3; current SQL/InfluxQL alternatives and migration implications are discussed. The sample is a learning system, not a production-readiness guarantee.

## Learning outcomes

- Explain time-series use cases and measurements, tags, fields and timestamps.
- Operate a prepared InfluxDB 2.x Linux lab and apply selected access/TLS controls.
- Design and evaluate a small telemetry schema, including cardinality trade-offs.
- Use Flux filtering, aggregation, windows and selected joins in the 2.x environment.
- Configure a matching Grafana data source, dashboard and alert example.
- Apply bucket retention and a selected Flux downsampling task.
- Explain version-dependent backup, scaling, permission and migration considerations.

## Prerequisites

- Basic understanding of programming concepts.

- Familiarity with SQL and database principles.
- Comfort working with Linux command-line environments.
- No prior knowledge of InfluxDB or Grafana is required.

## Training outline

### 1. Day 1 — Time-series foundations and platform choices

- Definition and characteristics of time-series data.
- Comparison with relational and transactional databases.
- Common use cases: IoT, DevOps monitoring, industrial analytics, and financial telemetry.

### InfluxDB concepts

- Core concepts: measurements, tags, fields, and timestamps.
- Architecture and storage model overview.
- InfluxDB product/version differences: the lab uses 2.x; newer 3.x offerings use SQL/InfluxQL rather than Flux
- Data flow within InfluxDB: ingestion, storage, querying, and visualization.

Compare the 2.x lab with InfluxDB 3 SQL/InfluxQL products; do not transfer query, administration or task assumptions without checking compatibility.

### 2. Day 1 — Linux operation and schema design

- Supported Linux prerequisites for the chosen 2.x release
- Install using the supported package or prepared deployment for the chosen release
- Configuring systemd services for automatic startup.
- Managing InfluxDB directories, logs, and service health.
- Version-matched upgrade, backup and restore procedures in the lab
- Securing installations (firewall, authentication, and TLS setup).

### Measurements, tags and lifecycle

- Mapping SQL tables to InfluxDB measurement models.
- Choosing tags vs. fields effectively.
- Managing cardinality for scalability.
- InfluxDB 2.x bucket retention and schema/cardinality decisions

## 3. Day 2 — Flux and Grafana integration

- Overview of the Flux language and syntax.
- Comparison with SQL query structures.
- Data filtering, aggregation, and transformation techniques.
- Working with time windows and joins.

### Data source and dashboard

- Adding InfluxDB as a data source in Grafana.
- InfluxDB API token, organisation and bucket configuration; distinguish these from Grafana API credentials
- Data source configuration and validation.
- Designing visualizations: time-series panels, tables, and gauges.
- Using dashboard variables for dynamic filtering.
- Thresholds, annotations and appropriate refresh intervals; dashboards are not a hard real-time guarantee

## 4. Day 3 — Retention, performance and alerts

- Understanding retention and downsampling strategies.
- InfluxDB 2.x Flux tasks for downsampling; continuous queries are a different version-specific mechanism
- Automating cleanup and archival policies.

### Measure before tuning

- Indexing and query efficiency.
- Handling high cardinality datasets.
- Hardware and scaling considerations for large deployments.

### Grafana alerting

- Configuring alert rules in Grafana.
- Integrating with email, Slack, and webhook notifications.
- Implementing proactive monitoring dashboards.

## 5. Day 3 — Permissions and end-to-end review

- InfluxDB 2.x organisations and tokens; Grafana organisations, teams and available roles
- Implementing least-privilege access and API tokens.
- Security best practices for production systems.

Review the simulated ingestion-to-dashboard path, validate selected alerts and document limitations and next operational steps.