

Apache Kafka: Topic & Schema Design Training

Description

The purpose of this training is to design scalable and manageable Kafka topic architectures aligned with business domains, together with a robust schema strategy. The course covers domain-driven topic modeling, key/partition selection, naming conventions, retention/compaction policies, replication and durability settings, as well as schema design using Avro, Protobuf, and JSON Schema. Schema Registry compatibility modes and schema evolution processes are also addressed to establish production-ready and governed topic and schema standards.

Learning Outcomes

- Design domain-aligned topic topologies with appropriate partitioning, retention, compaction, and replication settings
- Manage schema evolution safely using Schema Registry compatibility modes
- Establish governance standards, templates, and checklists for enterprise Kafka environments

Audience

- Data and platform architects
- Kafka administrators
- Backend developers and data engineers
- Teams standardizing Kafka topic and schema practices

Outline

Topic Design Fundamentals

- Domain-driven topic modeling and naming conventions
- Partitioning strategies and key selection
- Replication, `min.insync.replicas`, and durability
- Retention and log compaction policies
- Producer/consumer behavior impacts
- Workshop: Improving an existing topic map

Schema Design & Evolution

- Schema formats: Avro, Protobuf, JSON Schema
- Field naming, data types, nullability, defaults
- Schema Registry subject strategies
- Compatibility modes
- Versioning and evolution workflows
- Error scenarios and troubleshooting

Prerequisites

Basic Java programming knowledge is required.