

KSQL & Kafka Connect: Integration and Pipeline Design Training

Duration: 3 Days

Description

This training teaches how to process Kafka data using SQL-like queries with KSQL, perform filtering, transformation, enrichment, windowed aggregations, stream-table joins, and expose materialized views via pull and push queries. It also covers Kafka Connect for integrating databases, file systems, and search engines using ready-made connectors, including CDC with Debezium, DLQ, retry, and monitoring for production-grade pipelines.

Learning Outcomes

- Build real-time materialized views using KSQL stream-table joins and windowing
- Design end-to-end pipelines using Kafka Connect source and sink connectors
- Operate sustainable integrations with schema management, DLQ, retry, and monitoring

Target Audience

- Backend developers
- Data engineers
- Streaming and integration teams
- Teams standardizing KSQL and Kafka Connect

Outline

KSQL Fundamentals

- STREAM vs TABLE concepts
- Filtering, projection, and partitioning
- Joins and window types
- Aggregations and materialized views
- Pull vs push queries

Advanced KSQL & Design Patterns

- Repartitioning strategies

- State and retention tuning
- Schema management with Schema Registry
- Enrichment and deduplication patterns

Kafka Connect Integration

- Connect architecture overview
- JDBC and Debezium source connectors
- JDBC, Elasticsearch, and S3 sink connectors
- SMT patterns
- Error handling and monitoring
- End-to-end demo

Prerequisites

Basic Java programming knowledge is required.