

Elasticsearch Engineering

Build, Search and Deploy

FA-0819 · 2 days

Software Development · DevOps, Cloud & Infrastructure · Data & Analytics

An intensive two-day course in application development and deployment with Elasticsearch and AI coding harnesses.

Introduction

Search becomes an engineering problem when it has to work inside a deployed system. Developers need to make deliberate decisions about data modelling, relevance, synchronisation, security and service behaviour, then verify those decisions under real application conditions. This course brings those responsibilities together through intensive development on Ubuntu 24.04 servers.

Participants will use a coding harness to carry out the Elasticsearch installation and configuration, integrate Kibana where it supports the selected requirements, and develop, test and deploy applications that use Elasticsearch. Claude Code and Codex are suitable choices, and participants may use another coding harness that supports the required development and server workflows. Reviewing generated code, checking configuration and verifying actual system behaviour remain part of the engineering process.

The development scope follows a focused set of application requirements. Participants use one agreed framework, such as Node.js, Spring Boot or .NET, and an architecture suited to that application. Modular design, separate frontend and backend components, and microservices are considered where appropriate. Practical development, testing and deployment remain central, with the two days focused on one selected application rather than every framework or architecture.

Learning outcomes

- Use a coding harness to install, configure and verify Elasticsearch on Ubuntu 24.04.
- Configure Kibana for development and diagnostics where applicable.
- Design search indices, mappings and data synchronisation for application requirements.
- Implement search, filtering, relevance and aggregations through application APIs.
- Develop and test applications using a suitable framework and architecture.

- Deploy applications with secure Elasticsearch access and repeatable configuration.
- Diagnose deployed-system issues and validate indexing, performance and recovery.

Prerequisites

- Strong hands-on development experience in Node.js, Java/Spring Boot, C#/.NET or the agreed course stack, including APIs, testing and Git.
- Strong hands-on Linux DevOps experience, including SSH, sudo, service management, networking, TLS and application deployment.
- Deep working understanding of database systems, including data modelling, indexing, transactions and consistency.
- A laptop or workstation capable of running an SSH client, browser, editor and the selected coding harness.
- Access to an Ubuntu 24.04 development server with sudo privileges and resources suitable for the agreed training workload.
- A configured coding harness with authenticated access and sufficient usage allowance for intensive development.
- Unrestricted internet access.
- Ability to establish SSH connections to remote training servers at public IP addresses, with outbound SSH permitted by the device and network.

Training outline

1. 1. Application Requirements and Elasticsearch Architecture

1. Search requirements, data sources and acceptance criteria
2. Documents, indices, shards and replicas
3. Inverted indices, near real-time search and refresh behaviour
4. Framework selection and application architecture

2. 2. Engineering with Coding Harnesses

1. Development context, specifications and implementation plans
2. Harness access to repositories and remote servers
3. Code and configuration review, testing and verification
4. Iterative development, debugging and version control

3. 3. Elasticsearch Installation and Configuration on a Server

1. Server preparation, resource checks and system configuration
2. Harness-driven installation and compatible version selection
3. Service management, startup verification and troubleshooting
4. TLS, API keys, access privileges and network restrictions
5. Kibana installation, secure connection and Dev Tools where applicable

4. 4. Data Modelling and Indexing

1. Explicit mappings, field types and text analysis
2. Document structure, relationships and denormalisation
3. Bulk ingestion, updates and deletions
4. Database synchronisation, retry handling and idempotency

5. 5. Search and Analytics Development

1. Query DSL, full-text queries and structured filters
2. Relevance scoring and boosting
3. Sorting, pagination and highlighting
4. Aggregations and application-facing analytics

6. 6. Application Development and Integration

1. Elasticsearch clients and framework compatibility
2. Backend services, API contracts and frontend integration
3. Modular design, frontend and backend separation, or microservice boundaries
4. Input validation, authorisation, secrets and failure handling
5. Integration testing and application acceptance checks

7. 7. Application Deployment on a Server

1. Build artefacts, dependencies and environment configuration
2. Deployment automation, service startup and reverse proxies
3. HTTPS and secure application-to-Elasticsearch connectivity
4. Deployment verification, restart behaviour and release rollback

8. 8. Elasticsearch in Deployed Systems

1. Cluster health, application logs and Kibana diagnostics
2. Query profiling and indexing performance
3. Shard allocation and replica availability
4. Reindexing, index aliases and application compatibility
5. Snapshot repositories and restore validation

Disclaimer

This outline serves as a guideline for the proposed training. The trainer reserves the right to amend the content, sequence, tools, frameworks and delivery approach, at his discretion and without prior notice, to reflect participant requirements, technical developments and available training time.