



ELASTICSEARCH ENGINEER II

Overview

This instructor-led course is designed for Elasticsearch professionals who need to expand their skill set for developing and managing powerful search and analytics solutions with the Elastic Stack. You will learn advanced cluster management techniques, best practices for capacity planning and scaling, tips for monitoring and alerting, considerations for going into production, and more. You will also dig into field and document modeling, fixing data with Painless scripting, cross-cluster search, pipeline aggregations, and more. After completing this course, you will be prepared to attempt the Elastic Certified Engineer exam.

Audience

Experienced Elasticsearch professionals who need to expand their knowledge of Elasticsearch cluster management and application development.

Duration

In-Classroom - 2 Days | 8 hours per day

Virtual Classroom - 4 Days | 4 hours per day

Language

English

Prerequisites

Complete the **Elasticsearch Engineer I** course, or possess equivalent Elasticsearch knowledge

Requirements

- Mac, Linux, or Windows
- Latest version of Chrome or Firefox (Safari is not 100% supported)
- Stable internet connection (virtual classroom)
- Due to virtual classroom JavaScript requirements, we recommend that you disable any ad-blockers and restart your browser before class.

ELASTICSEARCH ENGINEER II

Modules

Field Modeling

- Learn how to design and model the fields in your documents, including discussions on granular fields, range types, dealing with large field cardinality, and designing for proximity matching.

- **Hands-on Lab**

Fixing Data

- Learn how to use the new Painless scripting language in Elasticsearch and discuss use cases for scripting, including the Reindex, Update By Query and Delete By Query APIs.

- **Hands-on Lab**

Document Modeling

- We discuss denormalizing documents, working with nested fields, and using the join field for parent/child relationships.

- **Hands-on Lab**

Advanced Search and Aggregations

- Learn some of the advanced search and aggregation techniques, including cross cluster search and pipeline aggregations.

- **Hands-on Lab**

Cluster Management

- We walk through the details of managing a cluster, including how to configure shard filtering, shard allocation awareness and forced awareness.

- **Hands-on Lab**

Capacity Planning

- Learn about designing for scale, scaling with replicas, scaling with Indices, capacity planning use cases, and working with time-based data.

- **Hands-on Lab**

ELASTICSEARCH ENGINEER II

Elasticsearch Internals

- Take a deep dive into how Elasticsearch works, including the details of Apache Lucene, segments, doc values, and caching.

- **Hands-on Lab**

Monitoring and Alerting

- We discuss of monitoring options, including the Stats API, task monitoring, the cat API, the X-Pack Monitoring component, and guidelines for monitoring a cluster and setting up alerts.

- **Hands-on Lab**

From Dev to Production

- We explore items to consider when moving to production, including network setup, hardware requirements, JVM settings, and also a discussion on some of the common causes of poor query performance and how to fix them.

- **Hands-on Lab**