

ELASTICSEARCH ENGINEER II

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.

LESSONS

All lessons include a hands-on lab.

Elasticsearch data modeling

Learn the details of how strings are analyzed and indexed in Elasticsearch. Learn how to design and model fields in your documents, including discussions on granular fields and important naming conventions with the Elastic Common Schema. Learn why denormalizing documents is suggested, and how you can do it. Understand in which use cases denormalizing is not enough and how you can leverage nested fields, as well as join fields.

Elasticsearch data processing

Learn how ingest pipelines can modify and enrich your data. Understand the pros and cons of batch processing using the Reindex, Update by Query, and Delete by Query APIs. Learn how to use the Painless scripting language in Elasticsearch, and discuss both the index and search use cases for scripting.

Continued on next page

COURSE INFORMATION



Audience

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



Duration

Classroom - 2 days | 8 hours per day
Virtual - 4 days | 4 hours per day



Language

English



Prerequisites

Complete the [Elasticsearch Engineer I](#) course, or possess equivalent Elasticsearch knowledge



Requirements

- Stable internet connection
- Mac, Linux, or Windows
- Latest version of Chrome or Firefox (other browsers not supported)
- Disable any ad blockers and restart your browser before class

ELASTICSEARCH ENGINEER II

LESSONS (CONTINUED)

All lessons include a hands-on lab.

Elasticsearch from dev to production

Understand the best practices to secure your Elasticsearch cluster. Learn the differences between development and production modes, as well as how caching works. Explore items to consider when moving to production, including index aliases, index templates, search templates, dynamic indexes, and dynamic fields.

Elasticsearch cluster deployment

Learn how to back up and restore a cluster. Understand how to leverage your architecture topology using shard allocation awareness and forced awareness. Learn how to upgrade your system and deal with cluster restarts. Understand the use cases for having multiple clusters and how to leverage both cross-cluster replication and cross-cluster search.

Elasticsearch nodes and index management

Understand what dedicated nodes are and how they can be applied to different use cases. Learn how to build a fully customized architecture that makes the most of each server using shard filtering. Manage your time-series index like a pro with index lifecycle management (ILM). Review the index management UI and understand how rollups can optimize your cluster.

Elasticsearch advanced tips and tricks

Explore the challenges of distributed operations and how Elasticsearch handles them. Learn how to design for scale and how oversharding can crash your cluster. Also learn how to optimize for read or write throughputs. Then explore some common causes of poor search performance and learn how to properly address them. Finally, learn server configuration best practices to consider when moving to production, including network setup, hardware requirements, and JVM settings.