

ELASTICSEARCH ENGINEER

This course is designed for both new Elasticsearch users and Elasticsearch professionals. It begins with the basics for getting started with the Elastic Stack, then quickly dives deep into topics for building efficient clusters. You will learn how to write search requests and work with their responses, transform and enrich data, implement strategies to optimize and scale your clusters, create data lifecycle and backup policies, manage multi-cluster operations, and troubleshoot your cluster.

After completing this course, you will be well on your way to becoming an [Elastic Certified Engineer](#).

LESSONS

All lessons include a hands-on lab.

Getting started

Learn how Elasticsearch and the components of the Elastic Stack work together to solve real-world use cases and problems. Identify the different ways to install Elasticsearch within Cloud or self-managed environments. Perform basic operations on data using REST APIs. Understand the type of data Elasticsearch users typically work with. Become familiar with how Kibana can be used to search and analyze data.

Data modeling

Define how Elasticsearch handles strings in your documents. Differentiate between keywords and text fields to understand their capabilities. Use a mapping to configure how Elasticsearch should store and access data. Learn about commonly used field data types and specialized types to optimize mappings. Learn about analyzers, its anatomy, and how to use built-in analyzers. Automate mapping creation using dynamic templates to define a field's mapping.

Continued on the next page

COURSE INFORMATION



Audience

Software Developers
Software Engineers
Data Architects
System Administrators
DevOps



Duration

24 hours



Language

English



Prerequisites

- No prior knowledge of the Elastic Stack required



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

LESSONS

All lessons include a hands-on lab.

Search

Discuss how to define queries using Query DSL, Elasticsearch's search language. Examine the structure of queries and what Elasticsearch returns, particularly how Elasticsearch scores documents. Learn how to find documents based on precise values using term-level queries. Search for numbers, dates, and IPs using range queries. Learn how to compose complex search requests by combining multiple queries using one or more boolean clauses.

Aggregations

Explore metric, bucket, and pipeline aggregations. Learn how aggregations summarize data as metrics, statistics, or other analytics and can be combined to gain deeper insights into your data. Define sub-aggregations to narrow the scope of your search. Perform transforms on Elasticsearch indices to create more efficient aggregations.

Data processing

Implement data processors to transform data by updating fields or enriching documents. Identify the different tools and how they can be used to process data in the Elastic Stack. Implement data enrichments within Elasticsearch. Discuss the benefits of using runtime fields, and how to define them using Painless, an Elasticsearch scripting language.

Distributed datastore

Learn how shards make distributed search possible and how replica shards can be used for high availability to improve search performance. Examine how to choose the number of shards for your cluster, when to scale your cluster up or down, and how to optimize performance. Explore the anatomy of write and search operations to illustrate how shards communicate.

Data management

Determine how to manage and organize data to optimize performance. Explore how to configure an index template to use data streams. Configure nodes for specific data tiers and monitor data as it moves through these tiers. Design index lifecycle management policies to automatically manage indices. Create and manage backups from an Elasticsearch cluster using searchable snapshots.

Cluster management

Learn how to manage cross-cluster operations for distributed Elasticsearch clusters. Replicate and search data across multiple clusters. Identify the health of your cluster, diagnose health issues, and monitor cluster performance using the Elastic stack.