

Updated on January 28, 2026

Register

ElasticSearch Training

2 days (14 hours)

Overview

[ElasticSearch](#) is a next-generation open source search and indexing engine. Its distributed nature and ability to be resilient and highly available have already won over major players such as Wikipedia, LinkedIn, Netflix, eBay, and WordPress.

ElasticSearch has been specially designed to index very large volumes of data while ensuring high performance scalability and high fault tolerance.

This training course covers the concepts of search engines before detailing the basic features of Elasticsearch.

The training provides all the knowledge necessary to use and operate Elastic Search effectively and resolve the most commonly encountered issues.

Like all our training courses, this one will introduce you to **the latest stable version** of the technology and its new features.

Objectives

- Learn how to use Elasticsearch, from installation to integration and configuration
- Use the Elasticsearch client and Java API to index and search documents
- Monitor Elasticsearch and understand usage statistics
- Determine the challenges and use cases of a search engine
- Situate Elastic Search in a Big Data environment
- Understand how Elasticsearch works
- Know how to index large volumes of data
- Understand system administration and monitoring to ensure availability

Target audience

- Architect
- Developer
- Project manager
- Big Data

Prerequisites

- Knowledge of Java & Linux
- Ideally, you will have completed [our ElasticStack training course](#).
- [Test My Knowledge](#)

ElasticSearch training program

Introduction to search engines

- General
- Overview of open source search solutions
- Integrating a search engine into an application
- Challenges and keys to success
- Presentation of new features in [versions 8 and 9](#)

Introduction to Elasticsearch

- Project history
- Apache Lucene
- What Elasticsearch brings to the table compared to Lucene

The basics of Elasticsearch

- Clustering principles
- Installation
- Configuration
- Concepts of Node, Index, and Type
- Data Partitioning
- Overview of the Rest API

Observability at Elastic

- The pillars of observability
 - Monitoring
 - Logs
 - Metrics
 - Traces API

- Sending data from the Elastic server
- Sending metrics data
- Shipping log data
- Application performance monitoring (APM)
- Observability applications

Document indexing

- Index and document design
- Indexing or deleting documents with the Rest API
- Bulk indexing
- Version
- Weighting
- Other features (routing, consistency, child document, etc.)

Mapping

- Definition and role of mapping
- Field types
- Predefined fields
- Index metadata

Text analysis and extraction

- Basis for text extraction and analysis
 - Analyzers
 - Char filters
 - Tokenizers
 - Token filters
- Use cases
- Configuring and using predefined or customized analyzers
- Multilingual text analysis
- Word removal
- Extracting email addresses and URLs
- Removing HTML tags from text
- Implementation of spell checking
- Indexing of binary files (using Apache Tika)

Searching for documents

- Searching for documents with the Rest API
- Managing results

- Types of queries
 - Match_all query
 - Query_string/simple_query_string query
 - Match query (and derivatives)
 - Term and terms queries
 - Wildcard query
 - Range query
 - Fuzzy query
 - Bool query
 - Exist/missing query
 - And, or, and not queries
 - Type and id queries
- Differentiating between queries and filters
- Types of filters
- Effectively combining filters
- Relevance

Advanced search features

- Managing relevance and scoring
- Suggestions
- Autocomplete
- Highlighting
- Object search / Nested / Parent-Child
- More Like This
- Geospatial Search

Elasticsearch and Java API

- Available Client Types
- Integrating Elasticsearch into a Java application
 - Embedded
 - Client node
 - Client transport
- Using the Java API
- Indexing and searching documents
- Managing indexes
- Managing mappings

Cloud and Clustering

- How an Elasticsearch cluster works
- Preventing split brains
- Configuring an Elasticsearch cluster
- Scalability and data volume
- Backing up and restoring an Elasticsearch cluster

- Monitoring an Elasticsearch cluster:
 - the health API
 - The state API
 - The stats API
 - The pending_task API
 - The _nodes API, also known as node info
 - The hot_threads API
 - The _cat API
 - Monitoring plugins
 - Logs

Advanced features

- Other types of Elasticsearch plugins
- River: definition
- Scripting
- Preheating your cluster
- Percolation
- Node maintenance
- Distributing indexes across different nodes
- Queries and statistics

ADDITIONAL MODULE IN ENGLISH ON REQUEST (+2 DAYS)

- Training language: English
- Course level: Beginner to intermediate

This training course teaches the basic concepts of Elasticsearch and explores all of its features. It provides the knowledge needed to use Elasticsearch effectively, based on real-world use cases. The course covers best practices and common problems.

Theory: 60% Practice: 40% Audience:

- Software Developers
- Data Engineers
- Architects

Prerequisites:

- Knowledge of REST/HTTP, Json, Yaml is appreciated
- No knowledge required

Getting Started

- Elasticsearch Overview
- Key Features
- Basic Concepts
- Install Elasticsearch
- CRUD Operations
- First Steps on Search API

Mapping and Analysis

- Introduction
- Data Types
- Main parameters
- Mapping API
- Analysis and Inverted Index
- Custom Analyzer
- Multi-Fields

Querying

- Search API Overview
- Terms Search
- Full Text Search
- Compound Queries

Aggregations

- Aggregations Overview
- Metrics, Aggregations
- Buckets Aggregations
- Pipeline Aggregations

Data modeling

- Denormalization
- Object relationship
- Nested relationship
- Parent/Child relationship

Dynamic mapping and templates

- Dynamic field mapping
- Dynamic template
- Template API

Nodes and Cluster Management

- Multi-node cluster
- Node Types
- Cluster settings

Ingest API

- Ingest Node
- Pipeline API
- Scripting

Overview

- Script API
- Painless language

More Features

- Alias
- Watcher
- Highlighting
- Suggestion
- Reindex Data
- Update by Query and Delete by Query

Add-on module: Elastic Engineer certification preparation (+2 days)

Companies concerned

This training is intended for both individuals and companies, large or small, wishing to train their teams in new advanced IT technologies or to acquire specific professional knowledge or modern methods.

Positioning at the start of training

The placement test at the start of the training course complies with Qualiopi quality criteria. Once they have finalized their registration, learners receive a self-assessment questionnaire that allows us to gauge their estimated level of proficiency in different types of technologies, as well as their expectations and personal goals for the upcoming training course, within the limits imposed by the selected format. This questionnaire also allows us to anticipate certain connection or internal security issues within the company (intra-company or virtual classroom) that could be problematic for the monitoring and smooth running of the training session.

Teaching methods

Practical training: 60% practical, 40% theory. Training materials distributed in digital format to all participants.

Organization

The course alternates between theoretical input from the trainer, supported by examples and reflection sessions, and group work.

Assessment

At the end of the session, a multiple-choice questionnaire is used to verify that the skills have been correctly acquired.

Certification

A certificate will be issued to each trainee who has completed the entire training course.