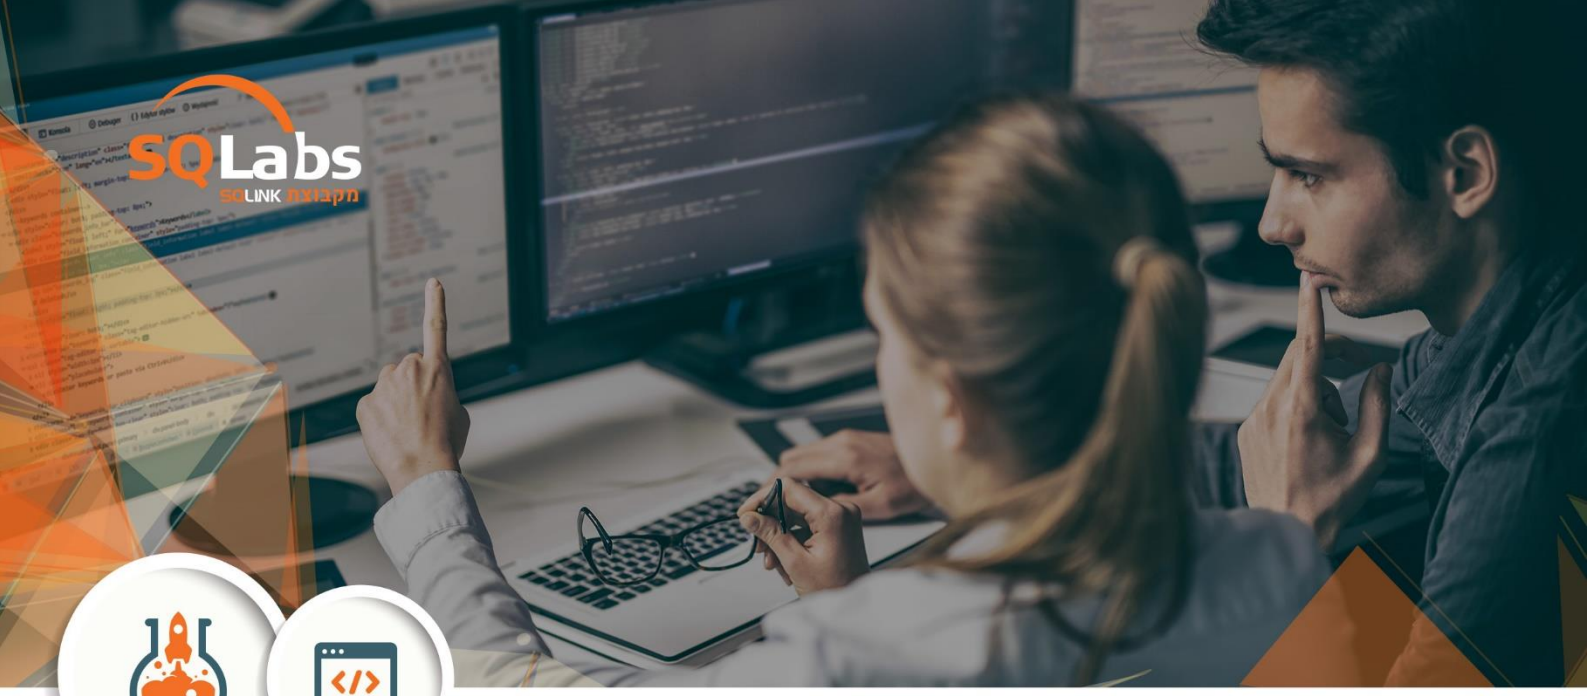

Elasticsearch

Duration – 32 hours

Description:

This ELK training intends to provide a solid foundation in search and information retrieval. It starts with fundamental concepts and follows with internals, best-practices and key features. Each topic is followed by a hands-on lab. At the end of the training, attendee will understand how Elasticsearch works, how to load data into Elasticsearch using Logstash, will be able to reliably analyze data and will be ready to build search applications and present visualization using Kibana

How does Elasticsearch work?

You can send data in the form of JSON documents to Elasticsearch using the API or ingestion tools such as Logstash and Amazon Kinesis Firehose. Elasticsearch automatically stores the original document and adds a searchable reference to the document in the cluster's index. You can then search and retrieve the document using the Elasticsearch API. You can also use Kibana, an open-source visualization tool, with Elasticsearch to visualize your data and build interactive dashboards.

Who should take is course?

Developers who want to learn Elasticsearch. The course is intended for developers and not system administrators.

Main syllabus subjects:

Module 1 - Introduction to Elasticsearch

- The Story of Elasticsearch
- Documents
- Indexes
- Indexing Data
- Searching Data
- The Bulk API

Module 2 - The Search API

- Introduction to the Search API
- URI Searches
- Request Body Searches
- The match Query
- The match phrase Query
- The range Query
- The bool Query
- Source Filtering

Module 3 - Text Analysis

- What is Analysis?
- Building an Inverted Index
- Analyzers
- Custom Analyzers
- Character Filters
- Tokenizers
- Token Filters
- Defining Analyzers
- Synonyms
- How to Choose an Analyzer
- Segments

Module 4 - Mappings

- What is a Mapping?
- Dynamic Mappings
- Defining Explicit Mappings

- Adding Fields
- Dive Deeper into Mappings
- Specifying Analyzers
- Dynamic Templates
- Index Templates

Module 5 - More Search Features

- Filters
- Term Filters
- The match_phrase_prefix Query
- The multi_match_Query
- Fuzziness
- Highlighting
- More Like This

Module 6 - The Distributed Model

- Starting a Node
- Creating an Index
- Starting a Second Node
- Shards: Distribution of an Index
- Distributing Documents
- Replication
- Split Brain
- Other Node Types
- Development vs. Production Mode

Module 7 - Working with Search Results

- The Anatomy of a Search
- Relevance
- Boosting Relevance
- DFS Query-then-fetch
- Sorting Results
- Doc Values and Fielddata
- Pagination
- Scroll Searches
- Choosing a Search Type

Module 8 - Aggregations

- What are Aggregations
- Types of Aggregations
- Buckets and Metrics
- Common Metrics Aggregations
- The range Aggregation
- The date_range Aggregation
- The terms Aggregation
- Nesting Buckets

Module 9 - More Aggregations

- Global Aggregation
- The missing Aggregation
- Histograms
- Date Histograms
- Percentiles
- Top Hits
- Significant Terms
- Sorting Buckets

Module 10 – Handling Relationships

- The Need for Data Modeling
- Denormalization
- The Need for Nested Types
- Nested Types
- Querying a Nested Type
- Sorting on a Nested Type
- The Nested Aggregation
- Parent/Child Types
- The has_child Query

The has_parent Query