Elastic Stack

Duration 3 day(s) (ELASTIC-STACK)

Centralize, explore highlight the logs

Description

Elastic stack is mainly built around Elasticsearch, Kibana and Logstash (ELK) even if now, the Beats tools come to supplement them advantageously. Elasticsearch is the fundamental brick of the stack as a data index, storage base and finally as a search engine. All projects must transfer and transform the data before sending it to them in order to make them usable and centralized. While the Beats or Logstash tools can transfer the data, the transformation can be done by the ingestion nodes of the Elasticsearch cluster or by the Logstash filter plugins. It is also desirable to have a facilitator tool allowing us to query / analyze or represent our data with a fluid and comfortable graphical interface. This is where Kibana gives its power to create queries, build thematic visualizations or dashboards just from your favorite web browser.

Goals

- Discover Elasticsearch and its key concepts
- Discover how to bring and transform raw data via different technological bricks (brokers, sockets ...) before sending them to the Elasticsearch search engine
- Discover how to produce efficient dashboards with many diagrams, maps or times series compositions

Public

Architect, Developer, Ops

Prerequisites

Minimal knowledge of Web and REST architectures, message brokers, development and linux is required.

Structure

60% Theory, 40% Practice

Program

Overview

- New uses of journaling
- The ecosystem around Elasticsearch
- The role of Elasticsearch, Logstash and Kibana
- Architectural examples

Introduction to Elasticsearch

- From indexing to research
- Textual analysis
- Mappings and configuration of the analysis
- Querying the possibilities of Elasticsearch
- Queries and Filters
- Aggregations
- Replication and partitioning
- Use cases around Logstash (index templates, dynamic names with date resolution, ...)
- Installation and configuration

Logstash

- The key concepts: Input, Output, Filter ...
- The Inputs: File, Redis, RabbitMQ ...
- The Filters: Grok, Date, Mutate ...
- Outputs: File, Elasticsearch, Redis ...
- Threading and high availability

Kibana

- Data Discovery and Queries Construction
- Aggregations and construction of Visualizations
- Compositions of 'time series' curves with Timelion
- Assemble views into a dashboard and administer it
- Installation and configuration