ELASTIC STACK DATA ADMINISTRATION I

Overview

This instructor-led course focuses on how to use the Elastic Stack to turn raw data into valuable insights, a top priority for data administrators. Looking at how each component of the Elastic Stack is used during the data lifecycle, you will first learn how to reliably collect and normalize data using Beats and Logstash. The data is then ingested into Elasticsearch for enrichment and analysis before being visualized and explored within Kibana. In addition to walking through data administration with lectures and hands-on labs, we will also discuss common pitfalls and architecture recommendations for various use cases.

Audience

Data Architects, Data Administrators, System Administrators, DevOps

Duration

In-Classroom - 2 Days | 8 hours per day

Virtual Classroom - 4 Days | 4 hours per day

Language

English

Prerequisites

- No prior knowledge of the Elastic Stack required
- Comfort using the terminal or command line recommended

Requirements

- Laptop with Wi-Fi connectivity (In-Classroom Only)
- Mac, Linux, or Windows
- Latest version of Chrome or Firefox (Safari is not 100% supported)
- Due to virtual classroom JavaScript requirements, we recommend that you disable any ad-blockers and restart your browser before class.

ELASTIC STACK DATA ADMINISTRATION I

Modules

Elastic Stack Data Administration Concepts

- Discuss the different problems and types of data that the Elastic Stack components can solve and the different architectures you can build.
- **Hands-on Lab** (30 minutes)

Systems Metrics

- Collect, store and analyze CPU, RAM, disk usage and operating system processes information.
- **Hands-on Lab** (30 minutes)

Services Metrics

- Understand how to collect and analyze data from services including Apache HTTPD, Docker and PostgresSQL.
- **Hands-on Lab** (30 minutes)

Ingesting File Data

- Learn how Logstash and Beats reliably read single and multi-line events handling complex problems like file rotation and failover.
- **Hands-on Lab** (30 minutes)

Data Processing

- Understand the difference between unstructured and structured data and what a document store is. Learn how to parse, ingest and analyze CSV data.
- **Hands-on Lab** (30 minutes)

Data Enrichment

- Learn how to grok (parse) and enrich log data using different lookup options and to determine if Logstash or Ingest Node is better for your ingestion pipeline.
- **Hands-on Lab** (30 minutes)

Data Store Integration

- Integrating the Elastic Stack with other databases and Hadoop.
- **Hands-on Lab** (30 minutes)

ELASTIC STACK DATA ADMINISTRATION I

Network Monitoring

- Configuring and deploying Packetbeat to capture and filter packet data. Learn different transaction protocols and monitored processes.
- **Hands-on Lab** (30 minutes):

Data Ingestion Architectures

- Discuss recommended architectures to scale the Elastic Stack and when to use persistent or distributed queues.
- **Hands-on Lab** (30 minutes)

Triage and Maintenance

- Monitoring data pipelines and fault identification and recovery. Upgrading Logstash & Beats.
- **Hands-on Lab** (30 minutes)