ELK Stack Workshop

ANI_307 | Expert-Led Live | Automation and Insights | Expert
Course Duration: 2 days

The ELK Stack—Elasticsearch, Logstash, and Kibana—empowers users to collect, search, analyze, and visualize log data in real time. By understanding how each component works together—Logstash for data ingestion, Elasticsearch for indexing and searching, and Kibana for visualization—participants can gain confidence in troubleshooting, monitoring, and gaining insights from their systems that run in a Kubernetes environment. Through many hands-on exercises participants will develop skills ranging from log collection and analysis to building dashboards and alerts that drive operational awareness and decision-making.

Intended Audience

A hands-on, in-depth technical training workshop for personnel involved in engineering and operations and monitoring of applications using the ELK stack on a Kubernetes infrastructure.

Objectives

After completing this course, the learner will be able to:

- Explain Kubernetes fundamentals for the ELK stack
- List and describe the ELK stack components
- Configure the ELK stack on Kubernetes
- Integrate applications with the ELK stack
- Create and manage Logstash pipelines
- Run queries through Kibana and the Elasticsearch API
- Create dashboards to visualize logs with Kibana

Course Prerequisites

- Strong understanding of 4G or 5G Wireless network
- Basic knowledge of Kubernetes

Outline

1. Kubernetes and the ELK Stack
1.1 Kubernetes architecture
1.2 Pods, namespaces, and daemonsets
1.3 ELK stack architecture and components
Exercise: Run and explore minikube and a 3-tier application
Exercise: Verify the ELK stack deployment

2. Kibana
2.1 The Kibana interface and navigation
2.2 Documents and indices
2.3 Discover: Search and query
2.4 Customizing the interface
2.5 KQL: The Kibana Query Language
2.6 Creating Visualizations
2.7 Stack and index management
Exercise: Using Discover
Exercise: Searching with KQL

3. Elasticsearch
3.1 Elasticsearch architecture
3.2 Field types and mapping
3.3 The Elasticsearch API and Query DSL
3.4 Indexing and searching
Exercise: Using the Dev Tools console
Exercise: Working with indices
Exercise: Searches and queries

4. Logstash
4.1 Pipeline architecture: Inputs, filters, and outputs
4.2 Input, filter, and output plugins
4.3 Examples: Beats and grok
Exercise: Simulating a pipeline in Dev Tools and Kibana

5. Advanced Elasticsearch
5.1 Templates
5.2 Aggregations
5.3 Scripting and the Painless language
Exercise: Aggregations and scripting

6. Advanced Logstash
6.1 Using conditionals and advanced filters
Exercise: Parsing and transforming log data
Exercise: Using Elasticsearch ML to ingest logs
Exercise: Building a multi-processor ingest pipeline

7. Advanced Kibana
7.1 Dashboards and visualizations
7.2 Alerting and reporting
7.3 Anomaly Detection
Exercise: Building dashboards
Exercise: Anomaly detection (end-to-end)