Faysal Aziz

Elasticsearch Operations with Kubernetes and Grafana

Operate Elasticsearch in legacy CentOS environments, orchestrate with Kubernetes and monitor with Grafana.

FA-0827 · 2 days

DevOps, Cloud & Infrastructure · Data & Analytics

Build practical Elasticsearch operations skills, then use a prepared Kubernetes environment and Grafana dashboards to explore deployment, monitoring and recovery.

Introduction

Operating Elasticsearch requires an understanding of the services, networks, storage and resources that keep event data available. Engineers need to recognise healthy behaviour, identify early signs of trouble and follow a systematic approach when ingestion slows, nodes disconnect or storage becomes constrained.

This course develops those skills through a combination of practical theory and guided hands-on work. Participants install and administer Elasticsearch on CentOS, investigate common operational problems and apply essential performance and maintenance practices. The focus is on network and infrastructure engineers responsible for reliable event storage and SIEM environments.

The course also introduces Kubernetes and Grafana as tools for deployment and visibility. Participants use a prepared Kubernetes environment to deploy and orchestrate containerised Elasticsearch nodes, explore persistent storage and observe recovery behaviour. They install Grafana on an Ubuntu server, connect monitoring data and build dashboards that support everyday operational decisions.

The training balances explanations with practical work, keeping Kubernetes and Grafana at an introductory operational level. Prepared remote environments support the exercises. CentOS Linux 7 and 8 are end-of-life systems and are used only for controlled legacy training. The Kubernetes, operator and Elasticsearch versions are selected as a compatible stack; CentOS Stream and supported production platforms are considered separately from legacy CentOS Linux.

The training emphasises practical skills shaped by real operational demands, connecting service administration, deployment behaviour and monitoring evidence.

Learning outcomes

By the end of the course, participants will be able to:

- Explain Elasticsearch architecture, node roles, indices, shards and replicas.
- Install and configure Elasticsearch in a version-compatible legacy CentOS training environment with appropriate connectivity and access controls.
- Monitor cluster health, resource utilisation, storage capacity and ingestion.
- Diagnose common failures and perform essential maintenance and recovery procedures.
- Explain Kubernetes concepts relevant to operating Elasticsearch.
- Deploy and manage a small Elasticsearch cluster on Kubernetes using supplied manifests and an operator.
- Install Grafana on Ubuntu and create operational dashboards and basic alerts.

Prerequisites

- Working experience troubleshooting IP connectivity, DNS, routing and firewall rules.
- Ability to use SSH, navigate Linux directories, edit configuration files and run commands with elevated privileges.
- A laptop or workstation with an SSH-capable client and a web browser.
- Unrestricted internet access.
- Ability to establish SSH connections to remote training servers at public IP addresses, with outbound SSH permitted by the device and network.

Training outline

1. 1. Elasticsearch Fundamentals for Infrastructure Operations

1. Event storage and SIEM architecture
2. Master-eligible, data and coordinating node roles
3. Indices, primary shards and replicas
4. Single-node and cluster deployment principles
5. Version compatibility and legacy CentOS considerations

2. 2. Elasticsearch Installation and Cluster Configuration on CentOS

1. Server resources, storage preparation and system limits
2. Package installation, configuration files and service management
3. Cluster discovery, initial cluster formation and node membership
4. Hostnames, network binding, communication ports and firewall rules
5. Authentication, TLS and installation verification

3. 3. Daily Administration, Capacity and Performance

1. Administrative checks using curl and Elasticsearch APIs
2. Cluster health, node statistics and shard allocation
3. Event ingestion, index status and SIEM-managed retention
4. JVM heap, filesystem cache, CPU and disk utilisation
5. Shard sizing, replicas and storage headroom

4. 4. Troubleshooting, Maintenance and Recovery

1. Startup failures, configuration errors and bootstrap checks
2. Connectivity failures, missing nodes and unassigned shards
3. Disk watermarks, write blocks and resource pressure
4. Controlled restarts and recovery validation
5. Snapshot creation and restoration

5. 5. Kubernetes Fundamentals for Elasticsearch

1. Containers, orchestration and desired state
2. Kubernetes control plane, worker nodes and Pods
3. Services, namespaces and cluster networking
4. StatefulSets, persistent volumes and storage classes
5. kubectl commands and supplied YAML manifests

6. 6. Deploying and Operating Elasticsearch on Kubernetes

1. Elastic Cloud on Kubernetes operator installation
2. Elasticsearch cluster deployment and node configuration

3. Persistent storage, credentials and secure service access
4. Resource requests, limits and controlled scaling
5. Pod replacement, recovery behaviour and diagnostic commands

7. 7. Grafana Installation and Administration on Ubuntu

1. Grafana architecture and data source fundamentals
2. Ubuntu preparation and Grafana package installation
3. Service startup, network access and initial configuration
4. Administrator credentials and access permissions
5. Service logs and installation troubleshooting

8. 8. Operational Dashboards, Alerts and Fault Diagnosis

1. Elasticsearch exporter and Prometheus metrics collection
2. Prometheus data source connectivity and authentication in Grafana
3. Cluster health, JVM, storage and ingestion dashboards
4. Dashboard panels, time ranges and basic threshold alerts
5. Correlating Grafana indicators with Elasticsearch and Kubernetes diagnostics

Disclaimer

This outline serves as a guideline. The trainer reserves the right to amend the content, sequence, emphasis or delivery approach without prior notice, as reasonably necessary to accommodate participant needs, technical conditions and the course objectives.

© 2026 Faysal Aziz · Kuala Lumpur