Faysal Aziz

Elasticsearch on CentOS

Deploy, manage and troubleshoot Elasticsearch in legacy CentOS environments.

FA-0826 · 1 day

DevOps, Cloud & Infrastructure · Data & Analytics

Practise Elasticsearch installation, cluster checks, capacity management and fault diagnosis for network and infrastructure teams supporting event storage and SIEM.

Introduction

Keeping Elasticsearch operational takes more than starting a service. Engineers need to understand how nodes communicate, where events are stored, what cluster health indicators mean, and how resource constraints affect ingestion and search. When problems arise, they need a practical way to distinguish connectivity failures from configuration errors, storage pressure and cluster issues.

This course develops those operational skills through focused theory and guided hands-on work on remote CentOS training servers. Participants install Elasticsearch, configure a small cluster, monitor its behaviour and investigate common faults. The emphasis is on the responsibilities of network and infrastructure engineers supporting event storage and SIEM environments, including deployments used with FortiSIEM.

The training connects Elasticsearch concepts directly to everyday administration: checking services, validating network paths, inspecting logs, managing capacity and recovering from failures. CentOS Linux 7 and 8 are end-of-life systems; legacy installations are studied in controlled training environments, not recommended as new production defaults. Participants distinguish CentOS Linux from CentOS Stream and check the operating system, Elasticsearch and SIEM compatibility requirements for the selected deployment.

The training emphasises practical skills shaped by real operational demands, with guided exercises focused on installation, administration and fault diagnosis.

Learning outcomes

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

- Explain Elasticsearch nodes, indices, shards and replicas from an operational perspective.
- Prepare a legacy CentOS training environment and select an Elasticsearch release appropriate to the agreed compatibility requirements.
- Configure a small cluster and verify node communication and access controls.
- Monitor cluster health, resource utilisation and event ingestion.
- Diagnose common service, connectivity, storage and shard allocation problems.
- Apply essential performance adjustments, routine maintenance and snapshot recovery procedures.

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 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

2. 2. CentOS Preparation and Elasticsearch Installation

1. Operating system, Elasticsearch and SIEM compatibility
2. Legacy CentOS repositories and support lifecycle considerations
3. CPU, memory, storage and system resource requirements
4. Package installation, configuration files and directory permissions
5. Service startup, boot configuration and installation verification

3. 3. Cluster Deployment and Secure Connectivity

1. Cluster identity, node discovery and initial cluster formation
2. Hostnames, address binding and node communication
3. HTTP and transport ports, firewall rules and access restrictions
4. Authentication, TLS certificates and operational permissions
5. Cluster membership and connectivity validation

4. 4. Daily Administration and Monitoring

1. Administrative checks using curl and Elasticsearch APIs
2. Green, yellow and red cluster health
3. Node statistics, index status and shard allocation
4. Service status, Elasticsearch logs and system diagnostics
5. Event ingestion checks and SIEM-managed retention boundaries

5. 5. Performance, Capacity and Maintenance

1. Event rates, retention periods and storage headroom
2. JVM heap, filesystem cache and swap considerations
3. Shard and replica sizing fundamentals
4. Controlled restarts and post-maintenance verification
5. Snapshot creation, restoration and recovery verification

6. 6. Troubleshooting and Service Recovery

1. Startup failures, configuration errors and bootstrap checks
2. DNS, firewall, binding and TLS connection failures
3. Missing nodes, unassigned shards and unhealthy clusters
4. Disk watermarks, storage exhaustion and write blocks
5. CPU pressure, JVM memory pressure and rejected requests
6. Corrective actions and end-to-end recovery validation

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