CASE STUDY

SQUARESHIFT

Upgrade of ELK Stack From RHEL 7 to RHEL 8 With Consideration of Amazon Linux for One of the World's Largest Investment Company

About the Client

The client is one of the world's largest investment companies offering a broad selection of investments, advice, retirement services, and insights to individual investors, institutions, and financial professionals.

Project Context

The current ELK stack production environment consists of:

- 10 Elasticsearch nodes
- 2 Kibana nodes
- 2 Fleet servers
- 2 Logstash servers

All servers were running on RHEL 7 and needed to be upgraded to RHEL 8 across both production and disaster recovery (DR) clusters. As part of the project, testing was conducted to explore the possibility of migrating to Amazon Linux OS, but the decision was made to proceed with a RHEL to RHEL upgrade due to the need for rolling upgrades and to avoid the complexities introduced by a heterogeneous OS environment.

Migration Strategy

The RHEL 7 to RHEL 8 migration followed a methodical, tiered approach to minimize downtime and ensure the stability of critical services:

Tiered Upgrade:
The upgrade process began with the Frozen, Cold, Warm, and finally the Hot data tiers.

Data Nodes First:
Upgrades started with the data nodes, followed by the master-eligible nodes.

Master Node Last:
The currently elected master node was upgraded last, identified using the Kibana UI (GET _cat/master).

DR Cluster Before Production:
To reduce risk, the DR cluster was upgraded before the primary production cluster.

Cross-Cluster Replication (CCR):
CCR was paused in the follower cluster before upgrading the leader cluster. After the upgrade, CCR was resumed.

CASE STUDY

SQUARESHIFT

Upgrade of ELK Stack From RHEL 7 to RHEL 8 With Consideration of Amazon Linux for One of the World's Largest Investment Company

Consideration of Amazon Linux

A PoC was conducted to test the migration of servers from RHEL 7 to Amazon Linux 3 (ALX), where 3 master and 3 data nodes were evaluated. However, this approach was not pursued for the production environment due to:

- **Rolling Upgrades:** Amazon Linux would have introduced complexities in performing rolling upgrades, as a homogeneous OS environment is preferred to ensure seamless upgrades across nodes.
- **Heterogeneous OS Challenges:** Mixing RHEL and Amazon Linux within the same cluster posed potential compatibility and maintenance challenges.

Thus, the decision was made to proceed with the RHEL to RHEL migration for both production and DR clusters to maintain a uniform environment.

Solutions and Recommendations

- **Health Check Report:**
  - Identification of unassigned shards, space optimization, and CCR setup verification.
  - Diagnostic analysis revealed bottlenecks in DR clusters, recommending downsizing by removing one warm node.
  - Capacity planning and scaling recommendations were provided to align with workload needs.

- **Licensing Guide:**
  - Elastic licensing is based on the total number of nodes and features in use. By monitoring usage across hot/warm tiers, the report provided a breakdown of licensing costs and optimization tips.
  - Recommendations included right-sizing the infrastructure to avoid over-allocation and reducing costs.

- **Observability Expansion:**
  - The scope was expanded to include Real User Monitoring (RUM), synthetic monitoring, and bi-directional CCR for higher resilience and monitoring of data flows.
  - Implementation of ILM policies for disaster recovery to handle data lifecycle management efficiently.

CASE STUDY

SQUARESHIFT

Upgrade of ELK Stack From RHEL 7 to RHEL 8 With Consideration of Amazon Linux for One of the World's Largest Investment Company

Proposed Architecture

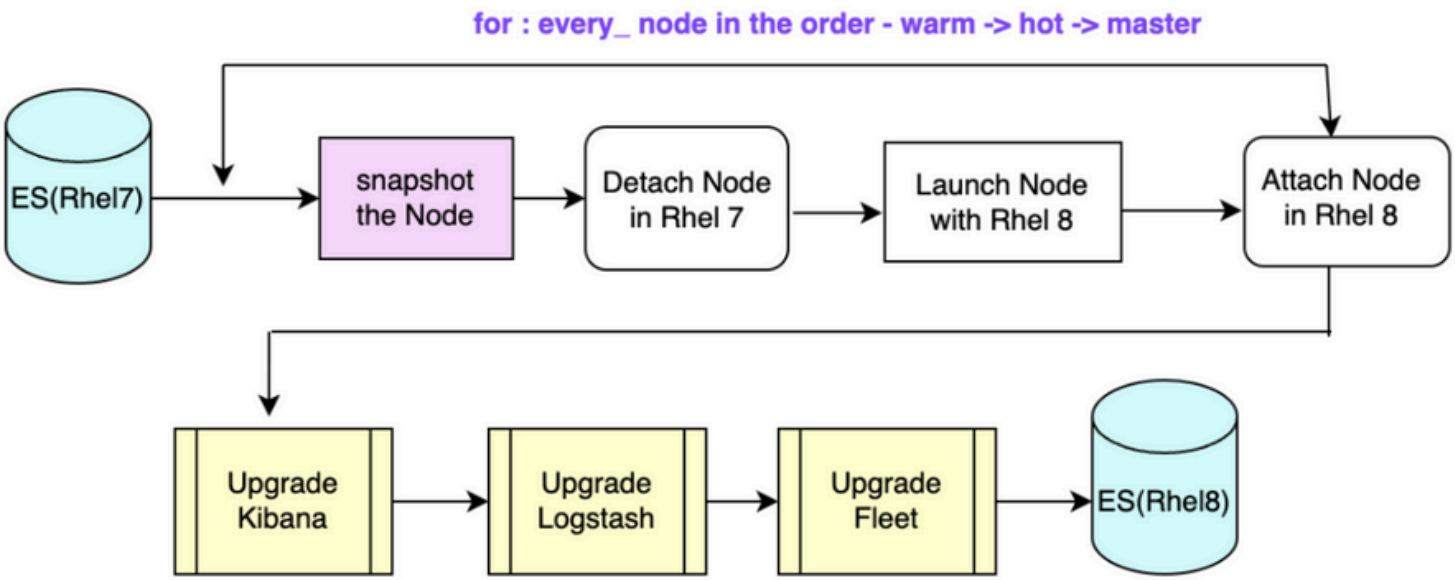

Upgrade Outcome

- **Seamless OS migration:** Successful migration of nodes from RHEL 7 to RHEL 8 using rolling upgrades across both production and DR clusters.

Resolved issues:

- Optimized free space and removed unassigned shards from hot and warm tiers.
- Corrected LDAP configurations for smoother authentication flows.
- Improved alerting mechanisms and connectors for Kibana dashboards.

Future Considerations

- **Synthetic Monitoring and Behavioral Analytics:** Expansion to monitor user interaction with the system in real-time.
- **Elastic Agent vs. Beats:** Recommendations to adopt Elastic Agent for simplified management over legacy Beats.

This upgrade and migration project enhanced the stability and performance of the ELK stack, optimized licensing costs, and expanded the observability capabilities for the organization's infrastructure.