

Break Free From Fake Free—Understanding the Total Cost of Ownership (TCO) of the ELK Stack

Introduction

The ELK Stack—comprising Elasticsearch, Logstash, and Kibana—is a popular solution for log management and data analytics. However, as organizations scale, the complexities and costs of maintaining ELK also grow. This document highlights the key reasons to replace the ELK stack, outlining both business and engineering perspectives.

Business Case for Replacing ELK Stack

1. High Operational Costs

One of the main reasons businesses consider replacing the ELK stack is its growing cost of ownership. ELK can be deceptively inexpensive at the beginning, as it's open-source. However, the complexity of managing and scaling ELK, especially when dealing with high volumes of data, requires substantial investment in both infrastructure and personnel. This escalates operational costs in three key areas:

- **Scaling Costs:** Scaling ELK to handle larger datasets often leads to exponential growth in resource requirements. As the data ingested increases, Elasticsearch clusters must be expanded, leading to more costs in terms of storage and compute resources.
- **People Costs:** ELK requires highly skilled engineers for setup, tuning, and ongoing maintenance. This expertise, often referred to as "DevOps tax," adds significantly to the operational budget.
- **Downtime and Reliability:** Maintaining an ELK stack's uptime can be challenging, resulting in potential business losses during downtime.

2. Poor Cost Transparency

Unlike managed services or alternative platforms that offer clear cost models, ELK stack costs can be unpredictable. Variable infrastructure requirements, especially with scaling, often make it difficult for businesses to budget properly. Additionally, Elasticsearch's resource consumption increases disproportionately with the volume of data, creating opaque and volatile cost projections for businesses.

3. Undifferentiated heavy lifting

The time and energy spent maintaining ELK often distract organizations from focusing on their core business objectives. Teams are frequently overwhelmed with managing Elasticsearch clusters, Logstash pipelines, and Kibana dashboards. This hinders the company's ability to innovate and explore new revenue streams that depend on effective data utilization.

Engineering Case for Replacing ELK Stack

1. Challenges in Scaling

From an engineering perspective, scaling ELK becomes extremely complex as data volumes increase. Elasticsearch clusters, which form the backbone of ELK, struggle under the load of indexing, querying, and storing large datasets. Engineers need to constantly tune the system for performance, which is not a trivial task. As data volumes grow, the inherent architectural limitations of Elasticsearch, such as sharding and replica overhead, become apparent.

2. Performance Issues

ELK is not built for real-time analytics at scale. As query latencies increase and the stack struggles to keep up with data ingestion, engineers spend excessive time troubleshooting. This makes ELK a less-than-ideal solution for teams that need to query data quickly and derive insights in real-time.

- **Indexing Delays:** As log volumes surge, delays in indexing data are inevitable. This makes real-time log analysis inefficient and prone to lag, which is a significant downside for systems requiring rapid insights, such as security operations or infrastructure monitoring.
- **Query Slowdowns:** High data volumes often result in query slowdowns as Elasticsearch struggles to fetch and process logs efficiently.

3. Operational Complexity

Maintaining ELK is resource-intensive. Logstash pipelines require careful monitoring and adjustment to ensure data is ingested correctly. Elasticsearch, on the other hand, needs constant attention to prevent cluster failures, especially as nodes grow and storage becomes a bottleneck. Kibana dashboards also need regular updates to ensure data visualization is effective and accurate. This operational complexity often leads to excessive downtime, negatively impacting both engineering productivity and overall system reliability.

Loved by the Most Innovative Companies

Edgio

reveal



AUDITBOARD



Capital One is focused on building seamless customer experiences that make banking and commerce simpler and easier—whether those customers are, digital or in-person. Like many cloud-first organizations, our data volume continues to expand. Observe provides a centralized and pre-correlated data layer that meaningfully organizes telemetry data from many sources at scale, helping drive faster response times.”

— Mark Cauwels, Managing Vice President, Enterprise Platforms Technology

