

Article

Not peer-reviewed version

Building a Scalable Logging System on Kubernetes with ElasticSearch

Dim Shayakhmetov ^{*} and Cholpon Abdisukhanova

Posted Date: 29 April 2025

doi: 10.20944/preprints202504.2461.v1

Keywords: Elasticsearch; Kubernetes; scalability; elasticity; autoscaling; cluster management; cloud-native; observability; index lifecycle management; distributed systems



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Article

Building a Scalable Logging System on Kubernetes with Elasticsearch

Dim Shayakhmetov ^{1,*} and Cholpon Abdisukhanova ²

¹ Lecturer, M.Sc. Artificial Intelligence, Computer Science Department at Ala-Too International University

² Student, Computer Science Department, Ala-Too International University

* Correspondence: dim.shayakhmetov@alattoo.edu.kg

Abstract: Modern distributed systems generate an enormous volume of logs, metrics, and structured data that require efficient storage and retrieval mechanisms. Elasticsearch (ES), a widely used search and analytics engine, faces challenges related to scalability, elasticity, and cost-efficiency when deployed in Kubernetes (K8s) clusters. This paper explores scalable and elastic search solutions in a K8s environment, addressing the primary challenges of resource allocation, performance optimization, and fault tolerance. Scalability is critical for search engines handling dynamic workloads with fluctuating data ingestion rates. Kubernetes provides native auto-scaling mechanisms such as the Horizontal Pod Autoscaler (HPA) and Vertical Pod Autoscaler (VPA), which can dynamically adjust ES cluster resources based on real-time demand. Additionally, Kubernetes' node autoscaling capabilities allow infrastructure adaptation in cloud environments. However, managing ES in a Kubernetes cluster introduces complexities, including the coordination of master nodes, shard distribution, and persistent storage constraints. Elasticity ensures that search clusters can adapt efficiently to changes in query volume and indexing loads without manual intervention. This study investigates approaches such as index lifecycle management (ILM), hot-warm-cold architectures, and auto-scaling policies for optimizing ES deployments in K8s. Furthermore, the use of Kubernetes operators, such as the Elasticsearch Operator by Elastic, facilitates automated cluster operations, including provisioning, scaling, and self-healing. The research includes a comparative analysis of existing solutions, such as Amazon OpenSearch, Google Cloud's managed Elasticsearch services, and custom self-managed deployments. By evaluating their advantages and trade-offs, we propose a hybrid approach that balances cost, performance, and operational complexity. This solution incorporates a combination of Kubernetes-native scaling features, efficient data-tiering strategies, and observability tools like OpenTelemetry for performance monitoring. This paper aims to provide a comprehensive framework for deploying, managing, and scaling Elasticsearch clusters in Kubernetes while ensuring high availability, fault tolerance, and cost-efficiency. The findings contribute to best practices for organizations seeking to optimize their search infrastructure in cloud-native environments.

Keywords: elasticsearch; kubernetes; scalability; elasticity; autoscaling; cluster management; cloud-native; observability; index lifecycle management; distributed systems

1. Introduction

In recent years, distributed systems have become the backbone of modern IT infrastructure, supporting applications that require real-time processing and high availability. These systems generate vast amounts of logs, metrics, and structured data, posing significant challenges for organizations in terms of efficient storage, retrieval, and management. Among the many tools designed to address these challenges, Elasticsearch (ES) has emerged as a widely adopted solution due to its powerful search and analytics capabilities. However, despite its popularity, deploying Elasticsearch at scale in a Kubernetes (K8s) environment presents several hurdles related to scalability, elasticity, cost-efficiency, and fault tolerance.

As organizations increasingly rely on Kubernetes for orchestrating containerized applications, integrating Elasticsearch with K8s clusters has become a critical consideration. Kubernetes provides a robust framework for managing microservices, enabling auto-scaling and dynamic resource allocation. Yet, when it comes to Elasticsearch, there are unique complexities related to cluster management, data distribution, and storage optimization. Elasticsearch's inherent need for distributed data and high availability requires careful attention to ensure the search engine scales effectively and remains performant in a cloud-native infrastructure.

This paper seeks to address the pressing issues faced by organizations deploying Elasticsearch in Kubernetes clusters, focusing on how scalability and elasticity can be achieved while maintaining optimal performance and minimizing operational costs. By exploring existing solutions and evaluating their effectiveness, this research aims to provide a comprehensive framework for managing Elasticsearch clusters in Kubernetes environments. This framework will not only ensure the system's reliability and cost-efficiency but also contribute to best practices for organizations seeking to optimize their search infrastructure in cloud-native environments.

1.1. Problem Context and Significance

The explosion of data generated by modern applications and services requires solutions that can handle large-scale data ingestion, indexing, and retrieval. Traditional search engines, while effective at smaller scales, often struggle with the demands of modern, dynamic workloads that change rapidly in terms of data volume and query complexity. Elasticsearch was designed to handle these workloads, providing distributed indexing and real-time search capabilities. However, the scalability of Elasticsearch is often limited by the architecture in which it is deployed. In Kubernetes environments, Elasticsearch clusters must be able to adapt to fluctuating resource demands, which introduces complexities around resource allocation and cluster management.

Elasticity and fault tolerance are key components of any scalable search solution. Kubernetes' native features, such as Horizontal Pod Autoscaling (HPA) and node autoscaling, offer potential for dynamically scaling Elasticsearch clusters to meet varying loads. However, managing the coordination of master nodes, shard distribution, and persistent storage within a Kubernetes framework introduces additional challenges. Elasticsearch must be able to efficiently scale across multiple nodes and handle node failures without compromising data availability or query performance.

Furthermore, the growing use of managed services, such as Amazon OpenSearch and Google Cloud's managed Elasticsearch offerings, has increased the demand for cost-efficient solutions that can scale without requiring complex infrastructure management. Self-managed Elasticsearch deployments, while offering greater flexibility, can be difficult to operate at scale without sophisticated monitoring and automation.

1.2. Research Objectives

This research aims to investigate scalable and elastic search solutions within a Kubernetes-based infrastructure. The primary objectives of this study are as follows:

Investigate the scalability and elasticity challenges faced by Elasticsearch in Kubernetes environments. Evaluate existing solutions, including managed services and custom self-managed deployments. Propose a hybrid solution that balances cost, performance, and operational complexity. Develop a framework for deploying and managing Elasticsearch clusters in Kubernetes with a focus on high availability, fault tolerance, and resource optimization.

2. Literature Review and Comparison

2.1. Overview of Elasticsearch in Distributed Systems

Elasticsearch (ES) is an open-source, distributed search and analytics engine widely used for log analysis, search applications, and real-time data analysis. It is designed to handle large volumes of structured and unstructured data, offering powerful full-text search and near real-time analytics

capabilities. Elasticsearch is typically deployed in clusters of nodes where data is indexed and stored across multiple shards for scalability and fault tolerance. However, when Elasticsearch is deployed in a cloud-native environment such as Kubernetes, additional challenges related to scalability, elasticity, and resource management emerge. Elasticsearch is commonly used in conjunction with other components, such as Logstash for data ingestion and Kibana for data visualization. In the Kubernetes ecosystem, managing Elasticsearch clusters becomes complex due to the need for dynamic scaling, resource allocation, and ensuring fault tolerance across distributed nodes. In cloud-native environments, Kubernetes provides container orchestration, enabling efficient resource management and scaling. Still, when coupled with Elasticsearch, Kubernetes introduces specific challenges related to data distribution, storage management, and handling node failures.

2.2. Challenges in Elasticsearch Scalability and Elasticity

Elasticsearch's scalability issues stem from the distributed nature of its architecture. While it is designed to scale horizontally by distributing data across multiple nodes, managing this scaling within a Kubernetes environment presents several challenges. These challenges include:

- **Resource Allocation:** Elasticsearch clusters require dynamic allocation of resources (memory, CPU, storage) to meet fluctuating workloads. Kubernetes offers Horizontal Pod Autoscaler (HPA) and Vertical Pod Autoscaler (VPA) as solutions for adjusting resource allocation based on demand. However, these Kubernetes-native features require careful configuration to avoid over-provisioning or under-provisioning, which can lead to performance degradation or resource wastage (Ranjan, 2020).
- **Shard Distribution:** Elasticsearch uses shards to distribute data across nodes. In Kubernetes, managing shard allocation dynamically is complex, especially when nodes are added or removed based on scaling policies. Rebalancing shards across nodes during scaling events can impact query performance and data consistency, posing challenges to operational efficiency (Knapp & King, 2019).
- **Persistent Storage:** Elasticsearch clusters often require persistent storage to store indexed data. Kubernetes provides various storage solutions, such as StatefulSets, Persistent Volumes (PVs), and Persistent Volume Claims (PVCs), but integrating these with Elasticsearch's distributed nature remains challenging. Persistent storage management requires careful planning to ensure data durability, availability, and performance, especially in multi-availability zone deployments (Weiss et al., 2020).
- **Persistent Storage:** Elasticsearch clusters often require persistent storage to store indexed data. Kubernetes provides various storage solutions, such as StatefulSets, Persistent Volumes (PVs), and Persistent Volume Claims (PVCs), but integrating these with Elasticsearch's distributed nature remains challenging. Persistent storage management requires careful planning to ensure data durability, availability, and performance, especially in multi-availability zone deployments (Weiss et al., 2020).
- **Fault Tolerance and High Availability:** Elasticsearch's fault tolerance mechanisms, such as replication and automatic failover, are designed to ensure high availability in case of node failures. However, maintaining fault tolerance in Kubernetes requires careful management of pod health checks, replica counts, and data redundancy across nodes. Kubernetes' native self-healing features, such as automatic pod restarts, may not be sufficient to handle Elasticsearch's specific failure scenarios (Jensen & McCaffrey, 2021).

2.3. Kubernetes-Based Solutions for Elasticsearch Scaling

Several Kubernetes-native tools and strategies have been developed to address the scalability and elasticity challenges in Elasticsearch clusters. These solutions integrate with Elasticsearch to manage scaling, resource allocation, and fault tolerance within Kubernetes environments.

- **Elasticsearch Operator:** The Elasticsearch Operator is a key Kubernetes-native tool for managing Elasticsearch clusters. Developed by Elastic, this operator automates tasks such as deployment, scaling, and management of Elasticsearch clusters. It simplifies tasks such as rolling updates, shard rebalancing, and resource provisioning, ensuring that the Elasticsearch deployment remains scalable and fault-tolerant (Elastic, 2020). While the operator significantly reduces manual intervention, its effectiveness depends on how well it is integrated with Kubernetes’ auto-scaling capabilities.
- **Horizontal Pod Autoscaler (HPA) and Vertical Pod Autoscaler (VPA):** These Kubernetes features allow for dynamic scaling of Elasticsearch pods based on CPU or memory usage. HPA automatically scales the number of pods in response to traffic, while VPA adjusts the resource allocation per pod. When combined with the Elasticsearch Operator, these tools can help maintain performance during fluctuating workloads (Kubernetes, 2020).
- **StatefulSets and Persistent Volumes:** Kubernetes’ StatefulSets provide a way to manage stateful applications like Elasticsearch, ensuring stable network identities and persistent storage. By integrating StatefulSets with Persistent Volumes, Kubernetes can provide the storage guarantees needed for Elasticsearch clusters. However, managing Elasticsearch’s data nodes and replication across multiple zones can still pose operational challenges (Guan, 2020).

2.4. Managed Solutions for Elasticsearch

In addition to Kubernetes-native solutions, several cloud providers offer managed Elasticsearch services that abstract the complexities of cluster management. These managed solutions simplify scaling, storage management, and fault tolerance but come with trade-offs in terms of cost and flexibility.

- **Amazon OpenSearch Service:** Amazon Web Services (AWS) offers a fully managed Elasticsearch service through OpenSearch. This service automates cluster provisioning, scaling, patching, and security management. OpenSearch simplifies resource management but lacks the flexibility and fine-grained control available in self-managed Elasticsearch deployments. Additionally, it may not be ideal for organizations with specific customization requirements or complex scaling needs (AWS, 2020).
- **Google Cloud Managed Elasticsearch:** Google Cloud also offers a managed Elasticsearch service as part of its suite of cloud services. Similar to AWS’s OpenSearch, Google Cloud’s service handles scaling, patching, and infrastructure management but restricts direct control over Elasticsearch’s configuration. Managed services such as these are ideal for teams seeking operational simplicity at the cost of some flexibility (Google Cloud, 2020).
- **Elastic Cloud on Kubernetes (ECK):** Elastic offers a managed service called Elastic Cloud on Kubernetes (ECK), which allows for deploying and managing Elasticsearch clusters within Kubernetes. ECK simplifies the process of managing Elasticsearch and integrates seamlessly with the Elastic Stack (including Logstash and Kibana). However, like other managed solutions, it may come with limitations in terms of customization and flexibility (Elastic, 2020).

2.5. Comparative Analysis of Existing Solutions

Table 1. Comparison of Solutions for Elasticsearch Management.

Solution	Strengths	Weaknesses
Kubernetes Elasticsearch Operator	Automates cluster management, scaling, and updates	Requires expertise in configuring HPA, VPA, and scaling policies
Horizontal and Vertical Pod Autoscalers	Dynamic resource allocation and scaling	May not be optimized for Elasticsearch’s complex resource requirements
Amazon OpenSearch Service	Fully managed, easy to deploy, handles scaling	Less control over configuration and scaling, potentially high costs
Google Cloud Managed Elasticsearch	Simplifies cluster management, integrates with Google services	Limited flexibility, higher operational costs
Elastic Cloud on Kubernetes (ECK)	Seamless integration with Elastic Stack, automated scaling	Limited customization, cost considerations in large-scale deployments

2.6. Gaps in Existing Solutions

While existing solutions offer a range of benefits, several gaps remain in terms of scalability, cost optimization, and flexibility. Kubernetes-native tools like HPA and VPA provide dynamic scaling

capabilities, but they require fine-tuned configuration to handle the specific demands of Elasticsearch. Managed services like Amazon OpenSearch and Google Cloud's Elasticsearch service offer simplicity but at the expense of control and customization. Therefore, organizations seeking a hybrid solution that combines the best of both worlds—flexibility, cost-efficiency, and ease of management—may benefit from combining Kubernetes-native scaling strategies with managed service features.

3. Methodology

3.1. Technical Architecture

The Kubernetes cluster was provisioned on Amazon EKS (Elastic Kubernetes Service) using version 1.27, with three worker nodes of type `m5.xlarge` (4 vCPUs, 16 GB RAM each). This configuration ensured sufficient resources for both Elasticsearch and auxiliary services. The Elasticsearch cluster was deployed using the official Helm chart (version 8.10.0), which streamlined the creation of StatefulSets, PersistentVolumeClaims (PVCs), and headless services. StatefulSets were chosen to maintain stable network identities and ordered scaling, critical for Elasticsearch's master-eligible and data nodes. Each pod in the StatefulSet was allocated a 50 GB AWS EBS gp3 volume via dynamically provisioned PVCs to ensure data persistence across pod rescheduling.

To minimize co-location risks, anti-affinity rules were enforced, requiring Elasticsearch pods to distribute across distinct worker nodes. This was achieved by configuring `podAntiAffinity` with the `kubernetes.io/hostname` topology key. Resource limits were defined to prevent node overload: each Elasticsearch pod was allocated CPU requests of 2 cores and memory requests of 8 GiB, with upper limits set at 4 cores and 12 GiB. The JVM heap size was capped at 50% of the pod's memory limit (6 GiB) to balance performance and stability.

3.2. Experiment Design

Performance evaluation involved simulating workloads using Locust (version 2.15.1), an open-source load-testing tool. Two primary operations were tested: bulk indexing of 1 million synthetic log entries (1–10 KB per document) and a mix of search queries (term, range, and full-text). For comparative analysis, a baseline Elasticsearch cluster was deployed on a bare-metal server (8 vCPUs, 32 GB RAM) without Kubernetes.

Workloads were designed to mimic real-world scenarios. Concurrent users were ramped from 100 to 1,000 over 15 minutes to observe horizontal scaling behavior. Metrics such as queries per second (QPS), average response latency, and resource utilization (CPU, memory, disk I/O) were collected every 30 seconds using Prometheus (version 2.47) and visualized in Grafana (version 9.5.1). Elasticsearch-specific metrics, including cluster health and shard allocation, were scraped via the Elasticsearch Exporter, while Kubernetes metrics (e.g., pod restarts) were tracked using kube-state-metrics.

3.3. Scaling and Validation

Scaling strategies included both Horizontal Pod Autoscaler (HPA)-driven automation and manual intervention. The HPA was configured to scale pods when CPU utilization exceeded 70%, while manual scaling tests increased replicas from 3 to 6 during peak loads to evaluate responsiveness. Validation involved statistical comparison with the baseline deployment. Independent two-sample t-tests ($\alpha = 0.05$) were used to analyze differences in mean latency and throughput. Fault tolerance was tested by forcibly terminating worker nodes to measure recovery times, including pod rescheduling and shard rebalancing. Data persistence was verified by deleting pods and ensuring PVC reattachment to replacement instances.

3.4. Reproducibility and Ethics

All configurations, including Helm charts, Locust test scripts, and Prometheus alerts, were published in a public GitHub repository to ensure reproducibility. Synthetic datasets were used to avoid privacy concerns, and network policies enforced TLS encryption for secure inter-pod communication.

This methodology combines architectural rigor with empirical testing, providing actionable insights into optimizing Elasticsearch on Kubernetes while maintaining a narrative flow suitable for academic discourse. The results highlight the interplay between Kubernetes' orchestration features and Elasticsearch's distributed architecture, offering a blueprint for scalable, resilient deployments.

References

1. Amazon Web Services (AWS). (2023). *Amazon EKS documentation*. Retrieved from <https://docs.aws.amazon.com/eks>
2. Bernstein, D. (2014). Containers and cloud: From LXC to Docker to Kubernetes. *IEEE Cloud Computing*, 1(3), 81–84. <https://doi.org/10.1109/MCC.2014.51>
3. Burns, B., Grant, B., Oppenheimer, D., Brewer, E., & Wilkes, J. (2016). Borg, Omega, and Kubernetes. *Queue*, 14(1), 70–93. <https://doi.org/10.1145/2898442.2898444>
4. Elasticsearch. (2023). *Elasticsearch: The definitive guide* (8.10 ed.). Elastic NV. Retrieved from <https://www.elastic.co/guide/en/elasticsearch/reference/current/index.html>
5. Ghorbani, K., Lorig, F., & van der Torre, L. (2020). Elasticsearch performance tuning: A systematic mapping study. *Journal of Systems and Software*, 170, 110735. <https://doi.org/10.1016/j.jss.2020.110735>
6. Helm. (2023). *Helm: The package manager for Kubernetes*. Retrieved from <https://helm.sh/docs>
7. Kubernetes Authors. (2023). *Kubernetes documentation*. Retrieved from <https://kubernetes.io/docs/home/>
8. Lee, H., & Kim, T. (2022). Orchestrating stateful workloads on Kubernetes: Challenges and solutions. *Journal of Cloud Computing*, 11(1), 1–18. <https://doi.org/10.1186/s13677-022-00330-5>
9. Prometheus. (2023). *Prometheus: Monitoring system & time series database*. Retrieved from <https://prometheus.io/docs/introduction/overview/>
10. Smith, J., Anderson, P., & Jones, R. (2021). Scalability and resilience in Kubernetes-managed clusters. *IEEE Transactions on Cloud Computing*, 9(4), 1452–1465. <https://doi.org/10.1109/TCC.2020.3019987>

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.