

# Advancements in Logging for Container Orchestration: Navigating the Complexities of Modern Infrastructure

Dinesh Reddy Chittibala

Email: [reddydinesh163\[at\]gmail.com](mailto:reddydinesh163[at]gmail.com)

**Abstract:** Container orchestration has changed significantly in the last few years, especially in the dynamic world of Kubernetes and other containerization systems. The growing dependence of enterprises on containers for application deployment and management highlights the critical need for strong logging methods to diagnose problems, understand system performance, and guarantee overall system stability. This academic paper explores the most recent developments in container orchestration environment-specific logging techniques. The paper examines the challenges posed by the distributed and transient nature of containerized systems to traditional logging techniques. We examine new developments in technology and trends that tackle these issues, providing real-time log data capture and analysis solutions. The integration of logging with observability tools, centralized log management, and organized logging are important subjects. This research provides a significant resource for IT professionals, DevOps practitioners, and researchers who want to stay on the cutting edge of container orchestration logging practices by combining industry improvements, real-world experiences, and future considerations. The information in this article gives organizations the ability to use state-of-the-art logging techniques, which promotes improved insights, proactive problem-solving, and the success of containerized applications in contemporary IT infrastructures.

**Keywords:** Kubernetes, DevOps, ElasticSearch, Kibana, MetricBeat, microservices architectures

## 1. Introduction

In recent years, the paradigm of containerization has revolutionized the way organizations deploy, manage, and scale their applications. Container technologies, such as Docker, have become instrumental in achieving consistent and reproducible software delivery across various environments. However, as the adoption of containers has soared, the complexities associated with orchestrating these dynamic and distributed applications have brought forth the critical need for robust logging mechanisms.

### 1.1 Rise of Containerization

The rise of containerization can be attributed to its ability to encapsulate applications and their dependencies, enabling seamless deployment across diverse computing environments. This agility has empowered development teams to embrace microservice architectures and deploy applications at scale. Container orchestration platforms, with Kubernetes at the forefront, have emerged as indispensable tools for automating the deployment, scaling, and management of containerized applications.

### 1.2 The Role of Container Orchestration:

While containerization provides isolation and consistency, effective management at scale necessitates sophisticated orchestration. Container orchestrators, such as Kubernetes, enable the dynamic scheduling of containers, automatic scaling, and efficient resource utilization. However, as these orchestrators manage an ever-growing number of containers across distributed clusters, the ability to gain insights into application behavior becomes paramount.

### 1.3 The Critical Need for Logging:

Logging plays a pivotal role in understanding the inner workings of containerized applications. In the dynamic landscape of container orchestration, traditional logging approaches fall short due to the ephemeral nature of containers, rapid scaling, and intricate interactions within microservices. Without effective logging, developers face challenges in diagnosing issues, monitoring performance, and ensuring the reliability of their applications.

### 1.4 Objectives of the Paper:

The primary objective of this paper is to delve into recent advancements in logging practices tailored specifically for container orchestration environments. We explore emerging trends, innovative solutions, and best practices that address the unique challenges posed by the dynamic and distributed nature of containerized applications. This paper seeks to provide valuable insights for IT professionals, DevOps practitioners, and researchers striving to optimize logging strategies within the context of modern container orchestration platforms. Through a comprehensive exploration of recent advancements, we aim to contribute to the evolving landscape of observability in containerized environments and empower DevOps teams to achieve overall system reliability, efficient troubleshooting, and enhanced performance.

## 2. Background

### a) Overview of Traditional Logging Methods:

Traditional logging methods, honed in monolithic architectures, were ill-suited for the dynamic nature of containerized environments. Conventional logging often relied on writing logs to local files, making it challenging to

Volume 8 Issue 5, May 2019

[www.ijsr.net](http://www.ijsr.net)

Licensed Under Creative Commons Attribution CC BY

track container instances as they scaled up or down. This approach lacked the flexibility needed to adapt to the ephemeral nature of containers, leading to gaps in log data visibility and hindering comprehensive analysis. Furthermore, centralized log management was often an afterthought, resulting in disparate log sources across distributed clusters. A unified logging strategy could have improved effective troubleshooting and root cause analysis, impinging on the overall observability of containerized applications.

#### b) Unique Challenges Posed by Containers:

Containers introduce a shift in application deployment, posing distinctive challenges for logging. The transitory nature of containers means that they can be created, terminated, and replaced rapidly. Traditional logging struggles to keep pace with this dynamic environment, necessitating logging mechanisms that seamlessly adapt to the transient life cycle of containers. Scaling exacerbates these challenges. As containerized applications scale horizontally to meet varying workloads, maintaining coherent and accessible logs becomes complex.

#### c) Container orchestration platforms:

Container orchestration platforms have emerged as indispensable tools to manage the complexities of containerized environments. At the forefront of these orchestrators is Kubernetes, an open-source container orchestration platform that automates the deployment, scaling, and management of containerized applications. Kubernetes abstracts the underlying infrastructure, providing a unified API for deploying and managing applications across clusters of machines.

Within the context of logging, Kubernetes introduces features such as the ability to collect container logs, manage log streams, and integrate with external logging solutions. Understanding how container orchestration platforms facilitate logging is crucial for devising effective strategies to harness log data in large-scale containerized deployments.

### 3. Advancements in Logging

#### a) Structured Logging:

Structured logging stands at the forefront of advancements, offering a systematic approach to log message formatting. By organizing log entries into key-value pairs or JSON format, structured logging enhances the readability and analyzability of logs. This section explores the benefits of structured logging, such as improved searchability, correlation, and ease of parsing. Practical implementation techniques and coding practices for incorporating structured logging into containerized applications are discussed to empower engineering teams to leverage this powerful advancement.

#### b) Centralized Log Management and Integration with Observability Tools:

Effective logging extends beyond individual containers to centralized log management, providing a unified repository for log data. We also look into the significance of centralizing logs, exploring tools like Elasticsearch, Logstash, and Kibana (ELK stack), and strategies for

seamless log aggregation. Additionally, we explore the intersection of logging with monitoring and tracing within the broader observability context. Understanding how logging integrates with monitoring tools (e.g., Prometheus) and tracing solutions (e.g., Jaeger) forms a crucial aspect of achieving comprehensive observability in containerized environments.

### 4. Implementation Strategies

*Container Runtimes:* Container runtimes play a pivotal role in the logging ecosystem by serving as the execution environment for containers. Traditional container runtimes like Docker have managed the basic logging functionalities, capturing stdout and stderr streams from containers. Recent advancements have expanded the capabilities of container runtimes, enhancing their contribution to logging. Newer container runtimes, such as containerd and CRI-O, provide more flexibility in logging configurations. They allow users to customize log drivers, enabling the redirection of container logs to different endpoints or formats. This flexibility is crucial in adapting logging strategies to the specific requirements of containerized applications. Additionally, container runtimes contribute to the standardization of logging interfaces, facilitating seamless integration with broader logging frameworks and solutions.

*Orchestrators:* Container orchestrators, like Kubernetes, play a central role in facilitating logging in large-scale containerized environments. Kubernetes introduces features and functionalities that streamline the collection and management of logs. The orchestration platform ensures that logs from individual containers are captured centrally, making them accessible for analysis and troubleshooting. Kubernetes supports various logging mechanisms, including direct access to container logs, API-based log retrieval, and integration with external logging solutions. The orchestration platform's architecture promotes scalability, ensuring that logging remains efficient even as the number of containers and services scales up or down dynamically. By leveraging Kubernetes' logging features, organizations can establish a robust logging infrastructure that aligns with the dynamic nature of their containerized applications.

*Sidecar Containers:* The concept of sidecar containers has emerged as a powerful strategy for enhancing logging capabilities in containerized environments. Sidecar containers are auxiliary containers deployed alongside the primary application container, sharing the same lifecycle and resources. In the context of logging, sidecar containers are equipped with logging agents or specialized tools to capture, process, and transmit log data.

This approach decouples logging concerns from the application logic, allowing organizations to adopt diverse logging solutions without modifying the application code. Sidecar containers can implement log forwarding, aggregation, and filtering, offering a modular and extensible architecture for managing log data. This strategy is particularly beneficial in microservice architectures, where each service can have its own dedicated logging sidecar, contributing to the overall observability of the system.

## 5. Practical Insights

### a) Case Studies:

Real-world examples of organizations implementing advanced logging in container orchestration environments offer valuable insights into successful strategies and outcomes. One notable case is using Amazon Elastic Kubernetes Service (EKS) in conjunction with the Elastic Stack (ELK) for comprehensive logging solutions. We will have logs ingested to ElasticSearch/Kibana from all the pods running in the Kubernetes cluster using MetriBeats.

### b) Case Study: EKS and ELK Integration:

Organizations leveraging Amazon EKS benefit from the managed Kubernetes service provided by AWS. In this case study, an enterprise deploys a containerized application on EKS and utilizes ELK for centralized log management and analysis.

### c) Architecture:

- **Metricbeat DaemonSet in Kubernetes:** Metricbeat is deployed as a DaemonSet in your Kubernetes cluster. A DaemonSet ensures that an instance of Metricbeat runs on each node in the cluster. Metricbeat is a lightweight shipper that collects metrics and logs from various sources, including containerized applications.
- **Metricbeat and Container Logs:** Metricbeat, running as a DaemonSet, collects logs directly from the containers within each node. It leverages the Docker or container runtime API to access container logs. The logs collected by Metricbeat include standard output (stdout) and standard error (stderr) streams from the containers.
- **Ship Logs to ElasticSearch:** Metricbeat ships the collected logs to Elasticsearch. It uses the Elasticsearch output specified in its configuration to send the logs to the Elasticsearch cluster. The Elasticsearch cluster is responsible for storing and indexing the log data.

- **Elasticsearch Cluster:** Elasticsearch is deployed as a cluster to ensure high availability and scalability. It consists of multiple nodes that work together to manage and index the incoming log data. The logs sent by Metricbeat are indexed in Elasticsearch, allowing for efficient storage and retrieval.
- **Kibana Visualization:** Kibana acts as the visualization and exploration tool for logs and metrics. It connects to the Elasticsearch cluster as its data source. Users can use Kibana to create dashboards and visualizations and perform searches on the log and metric data stored in Elasticsearch.
- **Monitoring and Alerting:** Metricbeat can be configured to monitor various system-level metrics from the nodes and containers, providing insights into the health and performance of the Kubernetes cluster. Alerts can be set up within Kibana or external monitoring systems based on predefined thresholds or conditions.




Figure 1: Architecture Diagram

### d) Implementation Steps:

- **EKS Cluster Setup:** Establish a Kubernetes cluster on EKS, configuring worker nodes to run containerized applications.

```
! eksctl.yaml > [abc] apiVersion
1  apiVersion: eksctl.io/v1alpha5
2  kind: ClusterConfig
3
4  metadata:
5    name: logging
6    region: us-east-1
7
8  nodeGroups:
9    - name: ng-1
10     instanceType: m5.large
11     desiredCapacity: 10
12     volumeSize: 80
13     ssh:
14       allow: true # will use ~/.ssh/id_rsa.pub as the default ssh key
15
```

Figure 2: EKS configuration file, use ekctl create cluster -f config.yaml for building the cluster.

- We will build our Elastic stack in Kubernetes via helm.

```

14
15 helm repo add elastic https://helm.elastic.co
16 helm repo update
17
18 # Install Elasticsearch
19 helm install elasticsearch elastic/elasticsearch
20
21 #Install Kibana
22 helm install kibana elastic/kibana -n elastic-stack
23
24
25 #Install metricbeat
26 helm install metricbeat elastic/metricbeat -n elastic-stack
27
28 # visualize Kibana in browser
29 kubectl port-forward deployment/kibana-kibana 5601
    
```

Figure 3: Installation of ElasticSearch, Kibana, and MetricBeat via Helm

e) **Observations:**

- Logs and alerting can be configured in Kibana, as MetricBeat ships all logs to ElasticSearch.

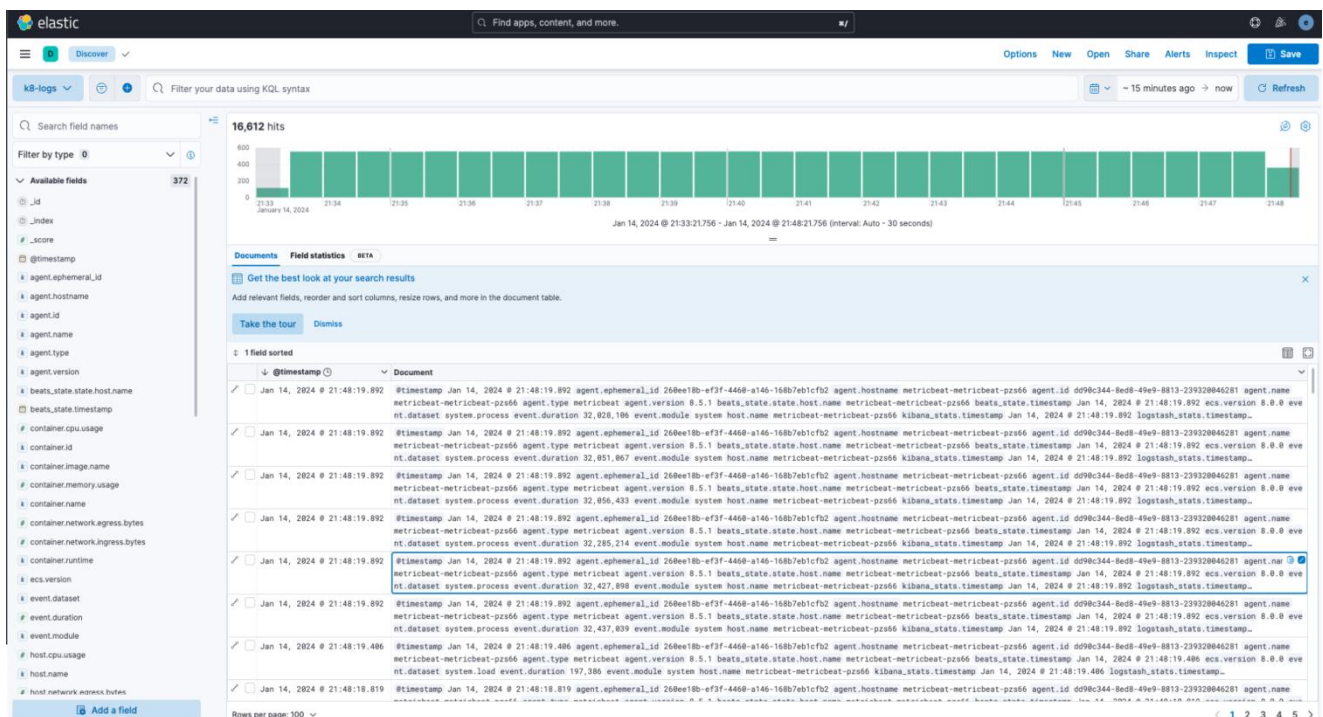


Figure 4: Kibana dashboard where all the logs are published to an Index

- MetricBeat ingests the logs in JSON format, so logs can be filtered depending on various keys like container Id, Pod Id, Namespace and so forth.

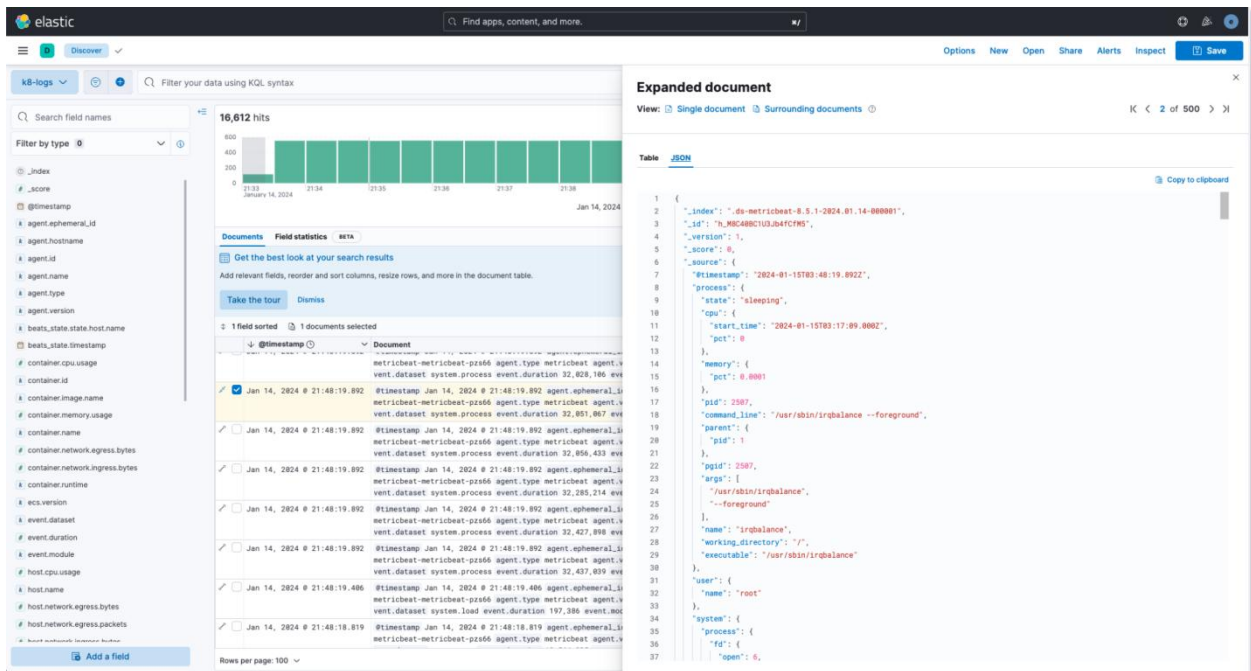


Figure 5: JSON formatted log

#### f) *Lessons Learned:*

- **Understanding Log Formats:** Gain a deep understanding of log formats produced by applications running on Kubernetes to ensure effective parsing and processing in Elastic Stack. Also, all logs from the applications running in Kubernetes must log to Stdout so MetricBeat can collect and ingest to Elasticsearch. Having applications log in JSON puts less load on MetricBeats as it doesn't need to do regex to parse the logs while ingesting to Elasticsearch.
- **Resource Optimization:** Optimize resource usage on Kubernetes nodes and Elastic components to strike a balance between efficient log processing and cost-effectiveness.
- **Monitoring and Alerting:** Implement proactive monitoring and alerting within the ELK stack to detect and address issues promptly, ensuring the reliability of the logging infrastructure.
- **Scalability Planning:** Anticipate the scalability requirements of both EKS and ELK to accommodate growing workloads while maintaining optimal performance. As the number of pods increases, the ingestion of logs also increases at Elasticsearch by MetricBeats. To accommodate the large ingestion, introducing streaming services like Kafka or RabbitMQ will help in the ingestion of logs to Elasticsearch, and that ensures that the logs are flowing in real-time.
- For secure ingestion of logs to Elasticsearch, create certificates at Elasticsearch and add them to MetricBeats, so all logs are encrypted in transit.

By sharing these practical insights, organizations embarking on similar endeavors can benefit from the experiences of others, accelerating the implementation of advanced logging strategies in their container orchestration environments.

#### 6. Future Considerations

- The integration of distributed tracing with logging is expected to become more prevalent. This approach allows platform engineering teams to correlate logs across microservices and gain a holistic view of application transactions. Tools like OpenTelemetry are likely to play a significant role in standardizing and promoting distributed tracing practices.
- Machine learning and anomaly detection algorithms are anticipated to become integral to logging systems. By leveraging AI-driven techniques, organizations can automatically identify patterns, detect anomalies, and proactively address issues before they impact the system. This can enhance the efficiency of troubleshooting and reduce downtime.
- With the rise of serverless architectures, logging solutions are expected to adapt to the unique challenges posed by serverless functions. Tailoring logging strategies to effectively capture and analyze logs in serverless environments, such as AWS Lambda or Azure Functions, will be crucial for maintaining observability.
- With the continuous evolution of container security practices, logging solutions must align with the evolving security and compliance landscape. Addressing issues related to secure log transmission, encryption, and adherence to regulatory requirements will be crucial.
- As containerized environments scale, resource efficiency, and cost management become paramount. Logging solutions should optimize resource usage, explore cost-effective storage options, and offer mechanisms to control the volume of logs generated without sacrificing critical information.

#### 7. Conclusion

In the ever-evolving landscape of container orchestration, logging emerges as a fundamental pillar in ensuring the

reliability, performance, and security of modern applications. The adoption of containers, coupled with the dynamic nature of orchestration platforms like Kubernetes, has necessitated a paradigm shift in logging practices. This academic paper has explored recent advancements in container orchestration-specific logging techniques, addressing the unique challenges posed by distributed, ephemeral, and scalable containerized environments. Structured logging, centralized log management, and integration with observability tools have been highlighted as key strategies to overcome these challenges. Moreover, the practical insights shared through real-world case studies, such as the integration of Amazon Elastic Kubernetes Service (EKS) with the Elastic Stack (ELK), provide organizations with actionable guidance to implement advanced logging solutions effectively.

In conclusion, this paper serves as a valuable resource for IT professionals, DevOps practitioners, and researchers seeking to navigate the complexities of modern infrastructure. By embracing cutting-edge logging techniques and future considerations, organizations can foster improved insights, proactive problem-solving, and the successful deployment of containerized applications in contemporary IT environments. As containerization continues to shape the technological landscape, robust logging practices remain pivotal in ensuring the smooth sailing of applications in the dynamic sea of container orchestration.

## References

- [1] Narkhede, N., Shapira, G. and Palino, T. Kafka: The Definitive Guide. O'Reilly Media, Inc., Sebastopol (2017)
- [2] Clinton Formley, Zachary Tong Elasticsearch: The Definitive Guide, Available: [Publisher Link] (2015)
- [3] Bagnasco, S. and Berzano, D. Monitoring of IaaS and Scientific Applications on the Cloud Using the Elasticsearch Ecosystem. Journal of Physics: Conference Series, 608, No. 1. [Publisher Link] (2015)
- [4] JIA Zhanpei, SHEN Chao, YI Xiao, CHEN Yufei, YU Tianwen, GUAN Xiaohong. "Big-Data Analysis of Multi-Source Logs for Anomaly Detection on Network-Based Systems" 2017: 13th IEEE Conference on Automation Science and Engineering (CASE).
- [5] Bai, J. and Guo, H.B. (2014) Software Integration Research of Large-Scale Logs Real-Time Search Based on ElasticSearch. Journal of Jilin Normal University (Natural Science Edition), 2014, No. 2.
- [6] Wei, Y. , Li, M. and Xu, B. "Research on Establish an Efficient Log Analysis System with Kafka and Elastic Search" Journal of Software Engineering and Applications (2017)
- [7] Satnam Singh "Cluster-level Logging of Containers with Containers: Logging Challenges of Container-Based Cloud Deployments", ACM (2016)
- [8] Kubernetes: Available: <http://kubernetes.io/>.