



Analysing and Alerting on Application Logs Within Kubernetes Infrastructure

Danila Koryugin

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

June 28, 2023

Analysing and alerting on application logs within Kubernetes infrastructure

Danila Koryugin
University of Twente
Drienerlolaan 5, 7522 NB Enschede
The Netherlands
d.koryugin@student.utwente.nl

Abstract

Software applications designed as a set of micro-services constantly require supervision. Logging of applications can help development teams explore and discover new insights in the workflow within micro-services, in addition to errors, exceptions, mistakes, and inefficiencies. Kubernetes is the most popular technology to maintain, manage and work with micro-services. Nevertheless, Kubernetes technology lacks in ability to effectively and efficiently provide logging information to the end-users. It is worth mentioning that Kubernetes has its logging tools, however, they are inefficient and not representative to obtain useful and insightful information. To cope with this, there exist various solutions to provide powerful logging capabilities. However, in this paper main focus will be on two of the most popular stacks of logging solutions that can also be integrated into Kubernetes infrastructure, namely Elastic stack and PLG stack with Prometheus. In addition to logging analysis, it is also necessary to have an opportunity to alert on these logs based on specified rules, e.g. alert on errors. The goal of this research paper is to understand which stack is the most suitable for providing application logging within Kubernetes infrastructure based on chosen metrics for a Java logger application, as well as define which stack is more relevant in terms of implementation based on a literature review.

Keywords

Kubernetes, Logging, ELK stack, Elastic search, PLG stack, Loki, Grafana, K8s, Elastic Stack, Docker, Promtail, Logstash, Kibana, Prometheus

1. Introduction

Software applications play a huge role in people's lives. Since the beginning of software development, developers were focusing on building monolithic applications, which [1] have a single code base that includes multiple services. These services can communicate with external systems or consumers via different interfaces like Web services, TCP connection, and REST API, as well as communicating with each other, for instance, web application service most of the time requires database service to store data. However, with the growth of companies providing more and more software solutions, the need for faster and more efficient development became extremely important. In addition to switching from a more conservative development method (Waterfall) to iterative and [2] with emphasis on responding to change over following a plan (Agile), and focusing on continuous integration and continuous development (CI/CD), companies tend to switch from monolithic software to micro-services. That is where Kubernetes technology appears as one of the best for managing and scaling micro-services applications.

Kubernetes is [3] an open-source system for automating the deployment, scaling, and management of containerised applications. Nevertheless, especially with Agile iterative development, it is necessary to supervise the state of the services, such as discovering potential errors, exceptions, the overall behaviour of a micro-service application, and other important information, that would help development teams to understand what should be done to improve or fix the services. This process is called logging. Kubernetes is a very popular technology, but there is no unified solution for providing an opportunity to effectively, efficiently, and easily apply logging technologies in its infrastructure to analyse and store application logs. Although Kubernetes has a built-in solution for gathering such data, e.g. [10] writing to standard output and standard error streams, it does not provide the possibility to process these logs, filter, cluster, and visualise them. To cope with this, there exist various solutions, and the most popular are [4] Elastic stack (Elastic search, Logstash, Kibana, Metricbeat) and the [5] PLG stack (Promtail, Loki, Grafana) with Prometheus. These approaches can help enhance working with logging, help to find useful and insightful patterns, and explore application bottlenecks, errors, exceptions, and inefficiencies. Moreover, these stacks provide opportunities to alert on application logs to the end-user, the best examples would be notifying users in Slack, Jira, and Email.

However, the application of the aforementioned logging tools and technologies is popular among DevOps developers and in general among programmers, it lacks scientific research. The goal of this research paper is to analyse and compare these stacks. The main focus of this research will be on the implementation of both stacks within Kubernetes infrastructure for a simple Java logger application on a physical device MacBook Air 18. Moreover, the literature review analysis will be conducted to discover which stack is more relevant in terms of implementation for different use cases.

2. Background

2.1 Terminology

Before continuing with the problem statement section, it is necessary to establish the terminology that will be used further in this paper, so the reader not familiar with DevOps has an idea of what was used during the research. Therefore, looking ahead, in this section technologies used will be listed with their tools and supported with respective documentation from official sources. It is worth mentioning that in this section only the description of technologies and tools will be observed, the steps taken to implement them will be described in more detail later in this paper.

The list of used technologies and tools:

- Docker
- Kubernetes
- ELK
- PLG & Prometheus

- Helm
- Java logger application

2.1.1 Docker

Docker [11] is an open platform for developing, shipping, and running applications on a physical/cloud machine (here and after physical/cloud machine will be referred to as the node). It provides the ability to package the application called containerisation. Docker consists of different objects, each with its objective to make Docker work.

- *Image* is a read-only template with instructions for creating a Docker container.
- *Container* is a runnable instance of an image. It can be created, started, stopped, moved, or deleted using the Docker API or CLI.
- *Docker-compose* is a tool for defining and running multi-container Docker applications. In Docker-compose, a YAML format file is used to configure the application's services.
- *Docker hub* is a hosted repository service provided by Docker for finding and sharing container images, as well as pushing and pulling images.

2.1.2 Kubernetes

Kubernetes [3] or just K8s (here and after Kubernetes will be referred to as K8s, official short name) was already defined in the Introduction section. Here only the objects [12] used by K8s will be described.

- *Pod* is a group of one or more containers, with shared storage and network resources, and a specification for how to run the containers. A Pod's contents are always co-located and co-scheduled, and run in a shared context.
- *Kube-proxy* is a K8s network proxy that runs on each node. This reflects services as defined in the K8s API on each node and can do simple TCP, UDP, etc.
- *Minikube* is a lightweight K8s implementation that creates a VM on a local machine and deploys a simple cluster containing only one node.

2.1.3 ELK

ELK or *Elastic stack* [13] is a stack of logging solutions made by Elasticsearch company. There exist different products, however, in this paper, only 4 of them will be used.

- *Elasticsearch* is a distributed, RESTful search and analytics engine capable of addressing a growing number of use cases. Moreover, it centrally stores the data for lightning fast search, fine-tuned relevancy, and powerful analytics that scale with ease.
- *Kibana* is a free and open frontend application that provides search and data visualisation capabilities for data indexed in Elasticsearch and acts as the user interface for monitoring, managing, and securing an Elastic Stack cluster.
- *Logstash* is a free and open server-side data processing pipeline that ingests data from a multitude of sources, transforms it, and then sends it to chosen stash.
- *Metricbeat* is a data shipper for collecting and shipping various system and service metrics to a specified output destination.

2.1.4 PLG & Prometheus

PLG stack [14] is a stack of logging solutions provided by Grafana labs company. The company has various products, however, in this paper only 3 will be used. In addition to that,

for purpose of retrieving the system metrics, Prometheus technology [15] provided by SoundCloud company will be used.

- Grafana is an open-source frontend application to query, visualise, alert on, and understand data pipelined to it.
- *Loki* is a horizontally scalable, highly available, multi-tenant log aggregation system inspired by Prometheus. It is designed to be very cost effective and easy to operate.
- *Promtail* is an agent which ships the contents of local logs to a private Grafana Loki instance. It primarily is able to discover targets, attach labels to log streams, and push them to the Loki instance.
- *Prometheus* is a free software application used for event monitoring and alerting. It records metrics in a time series database built using an HTTP pull model, with flexible queries and real-time alerting

2.1.5 Helm

Helm [16] is a Kubernetes deployment tool for automating creation, packaging, configuration, and deployment of applications and services to Kubernetes clusters.

2.1.6 Java Logger application

Java logger application will be used for testing purposes and represent the logs' spamming application with different log levels, namely info, warn, debug, and error, which is the most popular and widely used. The program was created to spam these logs with different chances of these logs appearing, thus the probability of the log with info level is 50%, both warn and debug levels are with a probability of 20%, and the error logs are with a probability of 10%. Moreover, each level has its number of times to be printed into the console log, namely info - 100, debug - 50, warn - 30, and error - 10. The process iterates infinitely until manually stopped with a delay of 5 seconds between each log's print. There exist two applications, with the same functionality for logging, but with different configurations to connect to each stack.

3. Problem Statement

Software application development highly depends on logging capabilities integrated into micro-services architecture. Finding the right solution appears to be challenging and complex, as existing solutions require different resources of a node, have limitations on where they can be deployed, and require substantive knowledge of the development team for proper implementation.

The problem statement will lead to the following research question:

What is the comparative performance for a Java logger application and implementation relevance based on a literature review for ELK and PLG stacks?

To answer this question, it is necessary to subdivide it into smaller sub-questions. Each subquestion will explore the problem statement in depth observed from different perspectives, namely:

1. Is there a significant difference between stacks in terms of performance for a Java logger application?
2. What is the prevalence and usage of the ELK stack and PLG stack in various research papers, and how do they

compare in terms of their implementations across different domains and use cases?

4. Research Methods

This research will be divided into 2 stages according to the subquestions mentioned above. The results of each subquestion will help to answer the research question of this paper. For the 1st subquestion, it is necessary to analyse the performance of each stack deployed on a shared single physical node when integrated into K8s infrastructure for a Java logger application. To conduct such analysis, specific metrics of a physical node are chosen, including a hypothesis to be tested with these metrics with the aforementioned logger application, providing a practical foundation for the research. For the 2nd subquestion, a systematic literature review will be conducted to understand the relevance of each stack based on different implementation scenarios. A literature review will help to discover which stack is more popular for various implementations. Thus, the focus will be made on the number of appearances in chosen databases, as well as providing an overview of the most popular IT-related sectors where the combination of K8s and either of the stack is used and deployed. It is necessary to mention, that only research papers will be used in the literature review and not the documentation of mentioned technologies and tools.

4.1 Performance Analysis

The performance analysis will consist of 3 hypotheses per each chosen metric to be tested. These metrics are essential for understanding the differences between stacks in terms of performance. The chosen metrics are:

- **CPU Usage:** measure the percentage of the total CPU capacity that is being used when the stacks are deployed. Lower CPU capacity indicates that the stack requires fewer node resources.
- **Indexing Latency:** measure the time taken by each stack to process and index log events. Lower latency indicates faster processing and indexing capabilities.
- **System Load:** system load is a measure of the amount of computational work that a computer system performs. In this case, the processes are running K8s and stacks.

All three metrics are available in Kibana for the ELK stack and in Grafana for PLG. Each hypothesis will be tested with a suitable statistical method. For research, it is intended to conduct hypothesis testing with two data sets. There exist two methods, i.e. two-sample *t*-test and paired *t*-test. According to Xu M., et. al. (2017), the two-sample *t*-test is used when the data of two samples are statistically independent, while the paired *t*-test is used when data is in the form of matched pairs. The data sets that will be gathered are independent of each other, as they will be gathered separately concerning the setup that will be running at that time, i.e. setup for ELK or setup for PLG. Thus, the method is chosen to be a 2-sample *t*-test. Moreover, the advantage of this method is that it will help to understand if there is a significant difference between each stack in terms of each metric. All tests will be conducted using the MacBook Air 18 as the node where stacks will be integrated. The data samples will be gathered in a timeframe of 10 minutes.

4.1.1 Setup limitations

Before continuing with the setup section, it is necessary to mention the limitations with which I coped during the integration of stacks on my node. The initial plan was to integrate each stack into the K8s minikube cluster and run the Java application at the same time, as stacks are better in performance and simplicity of integration when deployed within the same environment. However, during the integration of the ELK stack into the minikube, it was found that there are not enough resources on my node, as it requires higher CPUs and RAM for deployment. Therefore, it was decided to run the stack as the Docker-compose file within Docker and connect it to the Java application within the minikube. To keep the stacks within the same environment, i.e. in Docker, it was decided to also run PLG in Docker as a Docker-compose file. However, during the implementation of Promtail and the lack of documentation on it, as well as the time limitation given on research, I could not perform the connection between the Java application inside of the minikube to the Promtail running inside Docker. To make PLG work and save time, it was decided to run PLG inside the minikube, as was intended initially.

4.1.2 Setup

In this section, the setup, made for testing the hypotheses, will be shortly described to give an overview of how Docker, K8s, and both stacks operate with each other.

First of all, the Docker desktop application was installed on the node. After that, the K8s' minikube Docker image was pulled into the Docker using the CLI in the node's terminal. As it was mentioned in the limitations section, the ELK stack was deployed within Docker via the Docker-compose file, containing the configurations for Elasticsearch, Kibana, Logstash, and Metricbeat. After that, using the Helm package manager the configuration file containing configurations for Grafana, Loki, Promtail, and Prometheus was deployed as pods within minikube using the CLI in the terminal. After that two different Java applications were dockerised (containerised using Docker), pushed to the Docker hub repository created in my name, and then pulled with the help of a configuration file as an image to run a Java application pod within minikube. As it was mentioned earlier, both applications have the same functionality, but different configurations according to Promtail and Logstash documentation. The connection between Logstash and Java application was established using the TCP protocol and managed by kube-proxy for service discovery. The connection between Promtail and Java application was established automatically, as Promtail and Java application are running on the same cluster, i.e. minikube. The deployment of each stack was successful, so Logstash and Promtail can gather Java logs and pipeline them to respective services.

4.1.3 Hypothesis Description

In this section the description of the hypothesis will be shown and the steps to perform the testing procedure will be listed. The testing procedure will be the same for each metric with respective data sets.

1. *Model:* two data samples will be gathered while the application is running within the given timeframe with respect to each stack per each metric
2. *Hypothesis:*
 - 2.1 H_0 : the mean (μ_{elk}) metric's value of ELK is equal to the mean (μ_{plg}) metric's value of PLG ($\mu_{\text{elk}} = \mu_{\text{plg}}$)

2.2 H₁: the mean (μ_{elk}) metric's value of ELK is significantly different from the mean (μ_{plg}) metric's value of PLG ($\mu_{\text{elk}} \neq \mu_{\text{plg}}$)

3. *Significance level*: $\alpha = 5\%$

4. *Notation*:

- Sample size - n_x
- Sample mean - \bar{x}_x
- Sample standard deviation - s_x

5. *Test statistic*:

Pooled standard deviation is a measure of the variation, spread or dispersion of the data around the mean (see Fig. 1), s_p :

$$s_p = \sqrt{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2 / (n_1 + n_2 - 2)}$$

Fig. 1 Pooled standard deviation formula

Test statistic is a statistical test that is used to compare the means of two groups (see Fig. 2), t :

$$t = (\bar{x}_1 - \bar{x}_2) / s_p \sqrt{(1/n_1 + 1/n_2)}$$

Fig. 2 T test formula

6. *Compute p-value*: reject H_0 = if the p-value $\leq \alpha = 5\%$

7. *Draw statistical conclusion*.

Each hypothesis per each metric will be calculated and the results evaluated in the section 5.

4.2 Literature review

In this section literature review analysis will be conducted. The purpose of the analysis is to explore which stacks are more relevant in terms of practical implementation on various cases and to provide an overview of these cases. This will help to provide a theoretical foundation for the research. As it is mentioned before, the method is chosen to be a literature review, as it will help provide a summary and synthesis of the existing knowledge.

4.2.1 Search terms

To conduct a literature review it is necessary to state the search terms. In this research, the focus is mainly on two aspects: K8s and both stacks. Since the objective is to analyse existing literature, a state which stack is more relevant in terms of implementation, and in what are these cases, the search terms would contain "Kubernetes or K8s and" any of the following terms:

- ELK
- PLG
- Elasticsearch
- Logstash
- Kibana
- Promtail
- Grafana
- Loki
- Prometheus

4.2.2 Database Search

In order to locate the relevant research papers, the search was conducted in the following databases: IEEE, ScienceDirect, Wiley, and JSTOR. Moreover, the filters were limited to find ing papers starting from 2014, the year when Kubernetes was officially published as open-source software, and to research papers written in English.

4.2.3 Inclusion & Exclusion criteria

The inclusion and exclusion criteria are defined to find papers focused on the implementation of ELK or PLG within Kubernetes infrastructure in IT-related sector (see Table 1).

Inclusion criteria	Exclusion criteria
Research must contain only research articles	Research must exclude conference proceedings, news, documentations, and book chapters
Research must contain an example of either of stack implementation within K8s	
Research must be done in English	
Research must be done after 2014	
Full-text is available	

Table 1.

4.2.4 Screening process

The screening process was conducted with an emphasis on inclusion and exclusion criteria and used the aforementioned search terms. Moreover, to filter the gathered research papers PRISMA [17] four-phase flow diagram will be used and a research paper by Carrera-Rivera A. (2022) was used as an additional guidance. All articles were imported into the Mendeley software for screening.

Following the four-phase flow diagram, the following steps were conducted (see Fig. 3). The search resulted in 177 articles from the aforementioned search terms. After that, 69 papers were deleted as duplicates. Next, the papers were assessed and checked for eligibility based on inclusion and exclusion criteria. Out of 108 papers that were left after removing the duplicates, 76 of the papers were removed as assessed as not eligible for this study. Therefore, only 32 papers will be analysed in the literature review.

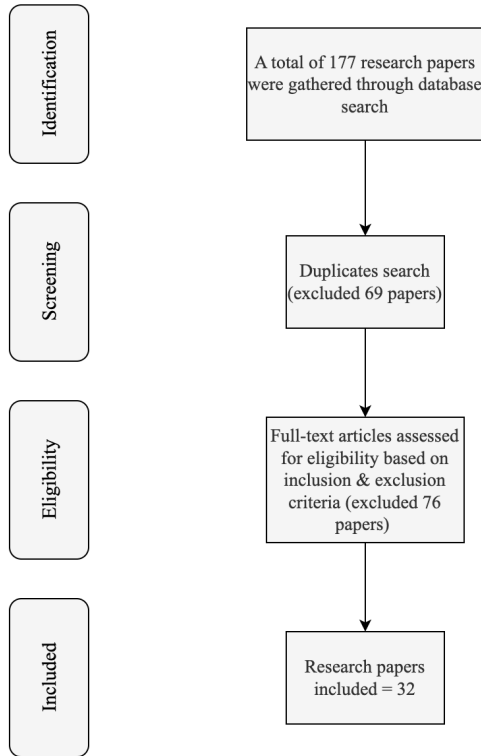


Fig. 3 Four-phase flow diagram

5. Results

5.1 Hypothesis analysis results

In the next 3 subsections, the hypothesis will be tested for each metric. The data gathered is shown in Appendix A at the end of the paper. All calculations are made according to formulas stated in the testing procedure. The result of hypothesis testing is then discussed according to the results provided in each subsection.

5.1.1 CPU Usage

First, it is necessary to state the variables.

For ELK we have:

- $n_{elk} = 10$
- $\bar{x}_{elk} = 40.72$
- $s_{elk} = 15.659$

For PLG we have:

- $n_{plg} = 10$
- $\bar{x}_{plg} = 43.92$
- $s_{plg} = 11.701$

Second, compute s_p : $s_p = 13.822$

Third, compute t : $t = -0.517682$

Lastly, as we intend to find if there is difference between stacks, we find two-tailed p-value = 0.610984

According to the testing procedure, we defined the significance level to be equal to 5% or 0.05. Since p-value = 0.610984 is not less than the significance level, we fail to reject the null hypothesis H_0 . Thus, we can say that there is no significant

difference between the ELK and PLG stacks in terms of CPU Usage.

5.1.2 Latency Indexing

First, it is necessary to state the variables.

For ELK we have:

- $n_{elk} = 41$
- $\bar{x}_{elk} = 9.3956$
- $s_{elk} = 12.529$

For PLG we have:

- $n_{plg} = 41$
- $\bar{x}_{plg} = 7.8802$
- $s_{plg} = 11.431$

Second, compute s_p : $s_p = 11.9925$

Third, compute t : $t = 0.572097$

Lastly, as we intend to find if there is difference between stacks, we find two-tailed p-value = 0.568860

According to the testing procedure, we defined the significance level to be equal to 5% or 0.05. Since p-value = 0.568860 is not less than the significance level, we fail to reject the null hypothesis H_0 . Thus, we can say that there is no significant difference between the ELK and PLG stacks in terms of Latency Indexing.

5.1.3 System Load

First, it is necessary to state the variables.

For ELK we have:

- $n_{elk} = 17$
- $\bar{x}_{elk} = 23.0317$
- $s_{elk} = 6.3111$

For PLG we have:

- $n_{plg} = 17$
- $\bar{x}_{plg} = 23.5705$
- $s_{plg} = 7.5003$

Second, compute s_p : $s_p = 6.9312$

Third, compute t : $t = -0.226644$

Lastly, as we intend to find if there is difference between stacks, we find two-tailed p-value = 0.822142

According to the testing procedure, we defined the significance level to be equal to 5% or 0.05. Since p-value = 0.822142 is not less than the significance level, we fail to reject the null hypothesis H_0 . Thus, we can say that there is no significant difference between the ELK and PLG stacks in terms of System Load.

5.1.4 Overall results

After the conduction of the hypothesis testing with the stated testing procedure using a two-sample t-test and from the results of testing each metric, we can conclude that concerning the chosen metrics, there is no significant difference between stacks deployed on the shared physical node. Both stacks were deployed independently and could not affect the performance of each other during the data gathering. All the data were gathered as mentioned above in a timeframe of 10 minutes. The results were expected, as they were deployed on a shared physical node with limited resources and a single cluster. In addition, the stacks are similar in terms of approaches to building a logging solution, i.e. search engines on non-

relational databases, using JSON format files for faster and more efficient parsing, etc. Two of the chosen metrics are related to the performance of the physical node when stacks are deployed, and the latency metric is mostly related to the processing powers of the stack itself. From the results, it is seen that even for different kinds of metrics there is no sufficient evidence that there is a significant difference between stacks. Nevertheless, section 7 (Limitations) will describe how various aspects of deployment could have affected the results.

5.2 Literature review results

5.2.1 Usage overview

The analysis of the included research papers shows that all of the papers after 2017, with the biggest rate of 30.03% were published in 2021. All of these papers contain information on the implementation of Kubernetes and either stacks or any technology from these stacks. In the analysis, the results were categorised into either the full stack or some of the technologies from these stacks, or the mix of technologies from both stacks. The analysis concludes that in 53.1% ELK stack or any technology included in the stack was used, in 31.3% PLG stack or any technology included in the stack was used including Prometheus, and in 15.6% the mix of technologies included in both stacks were used (see Fig. 4), which mostly was the mix of Elasticsearch and Prometheus. It can be seen that in over a half of use cases, the ELK stack was implemented for particular solutions.

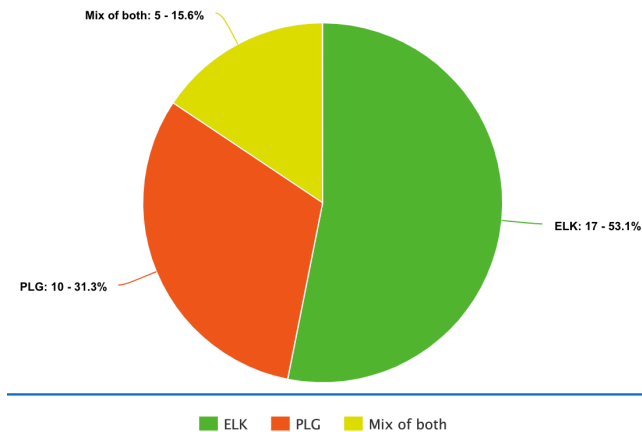


Fig. 4 Pie chart of technologies used in papers

5.2.2 Stacks' Implementation

In this section, the included research papers were analysed in terms of the use cases in which the stacks were implemented. The analysis showed that the use cases' topics are extremely diverse and capture different domains. The largest number of papers were related to the Monitoring and Cloud-edge Environment (3 papers per each). The next were Security Architecture, Data Analysis, Networking, Blockchain, and VMi Management (2 papers per each). Lastly, the use cases' topics that were mentioned only once in the research papers' list and couldn't be clustered further are Trigger System (Physics related), Data Visualisation, Data Management, Automatic DFD Extraction, Machine Learning, Log analysis, Malware Exposure, Cloud Scheduling, Cloudification Middleware, Workload Simulation, Integration with Raspberry Pi, Auto Scaling System, and Science Platform

(Astronomy) (1 paper per each). From such a distribution, we can derive possible patterns if there is any trend for a particular stack usage for particular use case (see Table 2), use cases mentioned only once will not be considered.

	Stack usage per 1 paper		
Monitoring	Prometheus & Grafana	Mix of both full stacks	Prometheus & Grafana
Cloud-edge Environment	Elasticsearch & Kibana	Elasticsearch & Kibana	Grafana
Security Architecture	Full ELK stack	Elasticsearch	
Data Analysis	Full ELK stack	Full ELK stack	
Networking	Grafana & Prometheus	Elasticsearch & Kibana	
Blockchain	Elasticsearch & Kibana	Elasticsearch & Logstash	
VMi Management	Full ELK stack	Prometheus	

Table 2.

It can be observed from the table that the prevailing stack is ELK. For instance, in Data Analysis-, Security Architecture-, and Blockchain-related use cases the authors of the research prefer using either a full ELK stack or a combination of ELK products. In the case of PLG, it can be seen that this stack and its products separately prevail in Monitoring-related use cases. For the other use cases, it can be seen that there is no specific preference over the stack, so both ELK and PLG products were used to implement particular use cases.

To summarise, the literature review was conducted, and out of 177 located research articles, 32 were eligible to be analysed concerning the inclusion and exclusion criteria. Based on the analysis, it can be concluded that ELK stack technologies, either altogether or separately, prevail in terms of implementation for various use cases. Moreover, although the use cases are very different and could not be grouped any further to gather similar use case topics, the prevalence of ELK products can still be seen.

5.2.3 Overall Results

From the results of previous sections, the 2nd subquestion can be answered as there is a large prevalence and usage of ELK over PLG in various research papers (over a half = 53.1%). Moreover, when compared to specific use cases, the ELK still shows a higher rate of usability and implementation in similar and different use case domains.

6. Recommendations

In this section, the recommendations concerning the stacks will be discussed. ELK and PLG are both logging solutions, which are capable of monitoring application logs, system metrics, networking, etc. Both stacks have different implementations in terms of the languages they were written on, thus they provide different functionality as well as compatibility with other technologies. Therefore, to implement a particular solution for a specific use case, it is necessary to understand how each

product within the stack works, which functionality it provides, and which resources are needed for deployment. It is needless to say, that the substantive knowledge of Docker and Kubernetes is required from the development team to be able to manage the containers and pods and configure them. From a personal perspective, the ELK stack seems more complete as all solutions are part of a single stack and have their necessary functionality, which helps to reach a goal without separate products. Other products can be deployed as well if needed, but usually, it is not the case. However, the PLG stack provides a larger flexibility in terms of different implementations, as a great number of products can be connected to Grafana, and these products can be mixed to achieve a required objective. Another advantage of Grafana against Kibana is the dashboard flexibility since it provides the opportunity to create dashboards with required metrics and parameters. Nevertheless, to be able to create such a dashboard it is required to know Grafana QL language, which is very time-consuming and was a limitation during this research (see section 7). It is certain that stacks serve the common solution and are extremely popular among developers, to be able to efficiently and correctly deploy them, it is required to refer to its documentation.

7. Evaluation

This section elaborates on the limitations faced during the research.

The main limitation is undoubtedly related to the given timeframe of the research. Within the 8 weeks provided to conduct the research, 4 were spent on getting acquainted with the aforementioned stacks of technologies, namely Docker, Kubernetes, ELK, and PLG products, and for the deployment of it on my node. Docker and Kubernetes are the backbones of DevOps development, and it is crucial to obtain enough knowledge to be able to deploy applications as well as configure them. Apart from the particular technologies, DevOps requires substantive knowledge of networking, as all the Docker and Kubernetes objects, especially minikube, have their networking parameters, such as IP addresses, ports, etc., to be able to connect them and manage them. For instance, one of the huge problems was that Loki could not listen to the application logs deployed within minikube until I found that it is necessary to port-forward the minikube pod to expose it to the local network.

(Regarding the 1st subquestion)

- Another huge limitation is concerning the resources of my MacBook Air 18, as it is not new and not powerful enough for quick and efficient manipulations with the Docker and Kubernetes objects, such as starting pod/container, stopping it, restarting, deleting, etc. Sometimes the process of stopping the set of containers where ELK was deployed could take up to 15-20 minutes or even stuck in an infinite loop of being stopped, which could only be resolved by restarting the Docker engine. After realising that node's resources are not enough for ELK deployment within Kubernetes, it was decided to switch to AWS infrastructure. However, the nodes with needed resources are expensive, so it was decided to switch to the setup mentioned in section 4.1.2 concerning the limitation mentioned in 4.1.1.
- Another limitation regards the ELK and PLG capability to provide needed metrics. In this sense, Grafana and Kibana have different implementations on it, and while Kibana has a default set of dashboards with important metrics to explore the performance of either CPU usage or Networking, Grafana does not have such default dashboards. Such

dashboards can only be either manually created (within a given timeframe it would be impossible) or found suitable in the dashboard marketplace (dashboards are free). Nevertheless, searching for dashboards that would match the ones Kibana provides was hard, as most of them require their own tools setup, such as InfluxDB and Telegraf, which are not included in this research.

(Regarding the 2nd subquestion)

- The only limitation faced was the sample size for the literature review due to the given timeframe. During the research conduct, only specific libraries were chosen according to chosen methodology, and during the screening (without duplicates) 108 research papers that are related to Kubernetes and both stacks. However, there undeniably exist a greater amount of research papers mentioning needed technologies, for instance, the search made via Google Scholar scientific browser gave thousands of results. Needless to say, such an amount could not be processed by a single person within the given timeframe. Without this limitation, the research might have been statistically more accurate due to the larger sample size, thus, be more representative not only in terms of the stacks' usage and implementation, but also in which domains these stacks were deployed.

8. Conclusion

In this paper, the two most popular stacks of logging solutions were considered and analysed from two different perspectives. First of all, the necessary information, such as introduction and background, was presented to the reader to get familiar with the purpose of this research and to get familiar with the objects and technologies that were required for the research.

Second of all, the research was conducted to analyse these stacks from both practical and theoretical foundations.

For the practical foundation, the statistical analysis was conducted based on the chosen metrics to establish if there is any significant difference between the stacks in terms of performance. These metrics were chosen to be CPU usage, System Load, and Latency Indexing. The statistical analysis was conducted with the help of hypothesis testing concerning the testing procedure. The statistical method chosen was a two-sample t-test, as it helps to discover the significant differences between the two samples. The results of the hypothesis testing showed, that for each metric there is no significant difference between the stacks.

For the theoretical foundation, the literature review was conducted to establish which stack prevails in terms of usability in various research papers, and compare them in terms of their implementations across different domains and use cases. From the literature review based on PRISMA's our-phase flow diagram and inclusion and exclusion criteria, 32 research papers were found where Kubernetes and any of the stacks, or a separate product of these stacks were used. The analysis showed that the ELK stack prevails in terms of usability across different research papers showing 53.1%, whereas the PLG showed only 31.3%, and the mix of the products implemented showed 15.6%. From the clustering of the use cases where the stacks were implemented, it was discovered that the ELK also prevails in terms of usability, especially in such cases as Data Analysis, Security Architecture, and Blockchain, whereas PLG products were used with the same amount of appearance as ELK (for instance, in Networking it is 50/50: Grafana and Prometheus and Elasticsearch and Kibana).

Thus, answering the main research question of this paper, the results showed that in terms of practice, there is no significant difference between the stacks, and at the same time there is a

large prevalence and usage of ELK over PLG in various research papers, as well as when compared to specific use cases, the ELK still shows a higher rate of usability and implementation in similar and different use case domains.

8.1 Future work

In the process of writing the Limitations and Recommendations sections, 3 points were formulated to enhance the research and go deeper into the analysis.

1. It is needless to say that the timeframe should be extended, as it would extensively help for proper research conduct.
2. The researchers should either be professionals in the DevOps field or have a team of such developers to be able to ask for a piece of advice and recommendation on how to properly setup the technologies and deploy them.
3. A larger group of people should conduct the literature review to be able to capture a larger sample size of research papers to further enhance the theoretical foundation of the analysis.

I believe that with these suggestions, the analysis would be highly improved and show useful and valuable results. Moreover, I believe that in the future, such analysis would help to formulate a methodology for developers, such as guidance on how to deploy and in which domains which stack is more suitable.

References

1. O. Al-Debagy & P. Martinek, (2018). "A Comparative Review of Microservices and Monolithic Architectures," *2018 IEEE 18th International Symposium on Computational Intelligence and Informatics (CINTI)*, pp. 000149-000154, doi: [10.1109/CINTI.2018.8928192](https://doi.org/10.1109/CINTI.2018.8928192).
2. Shore, J., & Warden, S. (2021). "The art of agile development", *O'Reilly Media, Inc.*, link: <https://books.google.nl/books?id=kXZIEAAAQBAJ&lpg=PP1&ots=M1TeAuVz4b&dq=agile%20development&lr&pg=PP1#v=onepage&q=agile%20development&f=false>
3. Official Kubernetes Website. Link: <https://kubernetes.io/>
4. Gupta, Y., & Gupta, R. K. (2017). "Mastering Elastic Stack". *Packt Publishing Ltd.*, link: <https://books.google.nl/books?id=EVQoDwAAQBAJ&lpg=PP1&ots=-0K4beEdeo&dq=Mastering%20Elastic%20Stack%20By%20Yuvraj%20Gupta%2C%20Ravi%20Kumar%20Gupta&lr&pg=PP1#v=onepage&q=Mastering%20Elastic%20Stack%20By%20Yuvraj%20Gupta,%20Ravi%20Kumar%20Gupta&f=false>
5. Holopainen, M. (2021). "Monitoring Container Environment with Prometheus and Grafana", *Metropolia University of Applied Sciences*, link: https://www.theseus.fi/bitstream/handle/10024/497467/Holopainen_Matti.pdf?sequence=2
6. Horalek, J., Urbanik, P., Sobeslav, V., Svoboda, T. (2023). "Proposed Solution for Log Collection and Analysis in Kubernetes Environment", *Nature of Computation and Communication*. ICTCC 2022.
7. M. Imran, et. al. (2023). "Evaluation and Implementation of Various Persistent Storage Options for CMSWEB Services in Kubernetes Infrastructure at CERN", *Journal of Physics: Conference Series, Ser.* 2438 012035, doi: [10.1088/1742-6596/2438/1/012035](https://doi.org/10.1088/1742-6596/2438/1/012035)
8. N. Sukhija & E. Bautista, (2019). "Towards a Framework for Monitoring and Analyzing High Performance Computing Environments Using Kubernetes and Prometheus", *IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation*, pp. 257-262, doi: [10.1109/SmartWorld-UIC-ATC-SCALCOM-IOP-SCI.2019.00087](https://doi.org/10.1109/SmartWorld-UIC-ATC-SCALCOM-IOP-SCI.2019.00087).
9. M. Dávideková & M. Gregu Ml, (2016). "Software Application Logging: Aspects to Consider by Implementing Knowledge Management," *2nd International Conference on Open and Big Data (OBD)*, pp. 102-107, doi: [10.1109/OBD.2016.22](https://doi.org/10.1109/OBD.2016.22).
10. Official Kubernetes documentation on built-in logging solutions, url: <https://kubernetes.io/docs/concepts/cluster-administration/logging/>
11. Official Docker documentation, url: <https://docs.docker.com/>
12. Official Kubernetes documentation, url: <https://kubernetes.io/docs/home/>
13. Official Elastic stack documentation, url: <https://www.elastic.co/guide/index.html>
14. Official Grafana labs documentation, url: <https://grafana.com/docs/>
15. Official Prometheus documentation, url: <https://prometheus.io/docs/introduction/overview/>
16. <https://helm.sh/docs/> helm
17. Moher, D., Liberati, A., Tetzlaff, J., Altman, D. G., Prisma Group. (2009). Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement. *PLoS medicine*, 6(7), e1000097. <https://journals.plos.org/plosmedicine/article?id=10.1371/journal.pmed.1000097>
18. Carrera-Rivera A., Ochoa W., Larrinaga F., and Laso G., (2022). "How-to conduct a systematic literature review: A quick guide for computer science research", *MethodsX*. doi: [10.1016/j.mex.2022.101895](https://doi.org/10.1016/j.mex.2022.101895)

Appendix A

	CPU Usage (%)	Latency Indexing (ms)	System load (1m)
ELK	72.8, 61.1, 39.3, 47, 38.3, 35.9, 33.4, 28.5, 28.8, 22.1	0.71, 0.73, 1.32, 14.74, 0.27, 3.74, 2.99, 6.57, 0.68, 0.41, 10.06, 1.19, 0.81, 0.32, 3.78, 4.55, 5.93, 37.58, 45.69, 1.47, 0.33, 3.46, 13.95, 32.41, 26.97, 18.04, 4.21, 7.62, 1.32, 7.32, 6.62, 4.24, 23, 45.1, 18.63, 2.18, 0.81, 0.45, 1.31, 1.02, 22.69	19.02, 20.66, 22.98, 15.95, 25.13, 27.17, 35.97, 30.22, 24.95, 14.82, 19.58, 29.09, 31.7, 22.81, 17.75, 13.26, 20.48
PLG	48.4, 42.9, 38.5, 65.2, 48.1, 38.8, 41.9, 58.3, 29.5, 27.6	0.48, 4.91, 0.674, 7.7, 0.77, 0.674, 0.77, 9.63, 0.963, 0.674, 0.77, 5.78, 0.578, 40.1, 0.309, 37.9, 0.285, 0.285, 1.93, 7.92, 2.12, 7.44, 32, 34.1, 11.3, 2.1, 12.41, 0.55, 2.31, 1.01, 0.69, 1.89, 24.76, 1.05, 4.23, 10.59, 0.78, 3.46, 29.92, 15.8, 1.48	37.9, 22.1, 13.5, 20.7, 17.3, 37.4, 18.5, 22.3, 32.8, 16.7, 21.7, 20.2, 33.0, 20.5, 17.7, 19.5, 28.9