AssureMOSS

Threat modeling for Kubernetes

Do graphical models helps?

Francesco Minna

19/05/2022

VU

Index

1. Brief introduction about me and the project
2. Kubernetes and STRIDE
3. The experiment
4. Preliminary result and Discussion

whoami

Francesco Minna - 26yo

- PhD Candidate (second year) at Vrije Universiteit Amsterdam
- Part of the Foundational and Experimental Security research group
- Working within the AssureMoss European Project (https://assuremoss.eu/)
- EIT MSc double-degree in cyber security (Trento, IT + Rennes, FR)

AssureMOSS - VU Amsterdam

The Assurance and certification in secure Multi-party Open Software and Services (AssureMOSS) project aims to produce a set of techniques to assess, manage and re-certify the security and privacy of software projects.

We, as the Amsterdam partner, are responsible for the continuous run-time assessment of deployed applications and services.

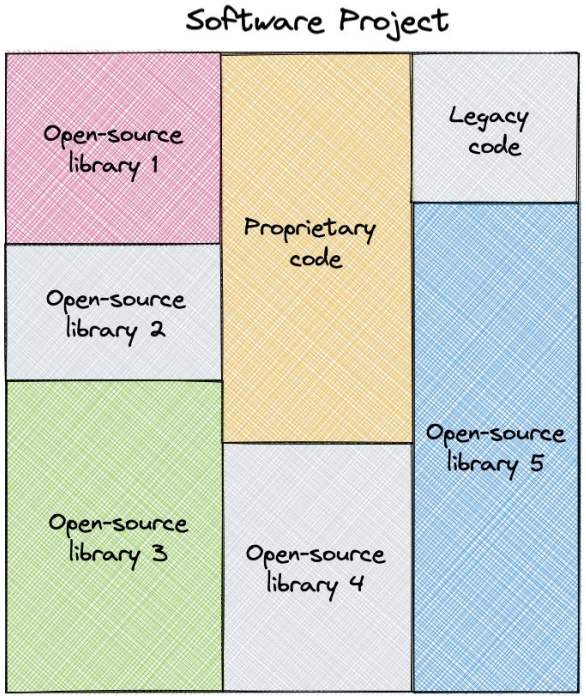

Kubernetes Orchestrator

When you install Kubernetes, you get a Kubernetes cluster, the components of which are the following

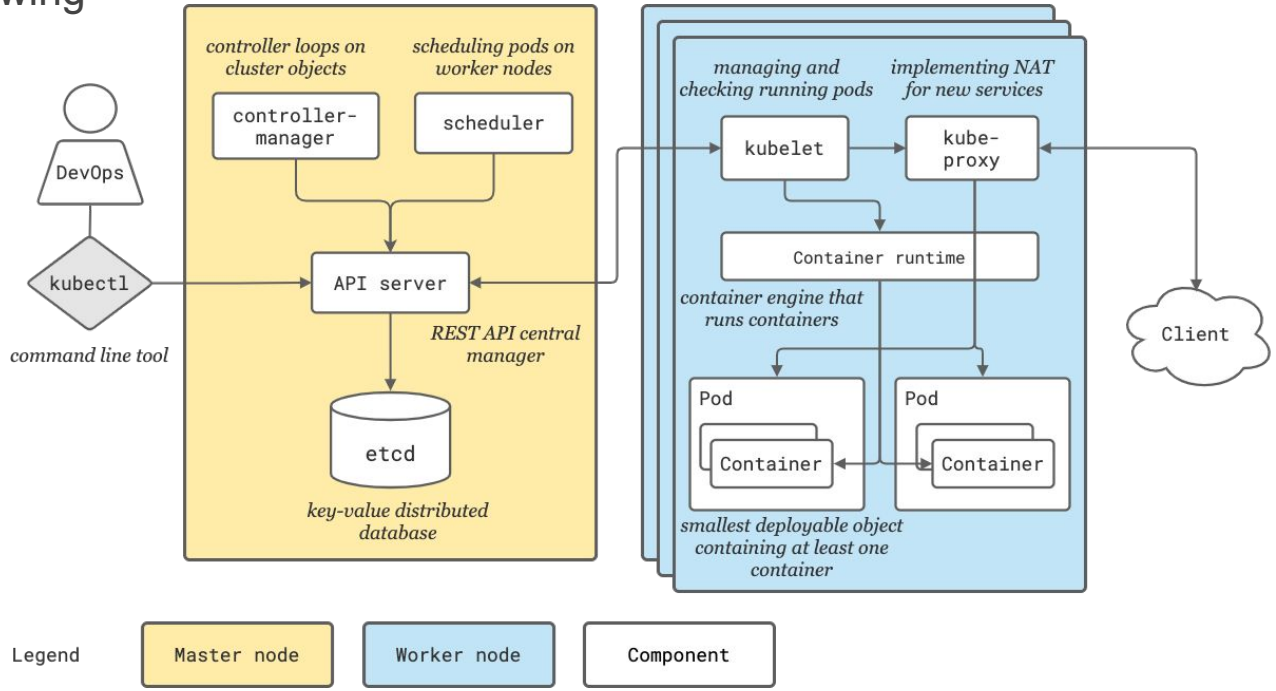

STRIDE Methodology

STRIDE is a Microsoft's threat analysis methodology used to identify security threats software systems.

Property	Threat	Definition
Authentication	**S**poofing	Pretending to be someone else
Integrity	**T**ampering	Modifying data or code
Non-repudiation	**R**epudiation	Claiming to have not performed an action
Confidentiality	**I**nformation Disclosure	Exposing information to someone not authorized to see it
Availability	**D**enial of Service	Deny or degrade service to users
Authorization	**E**levation of Privilege	Gain capabilities without proper authentication

Data Flow Diagram (DFD)

A DFD is a graphical representation of how data enters, leaves, and traverses a system.

- It shows all data sources and destinations;
- It shows all relevant steps that data goes through.

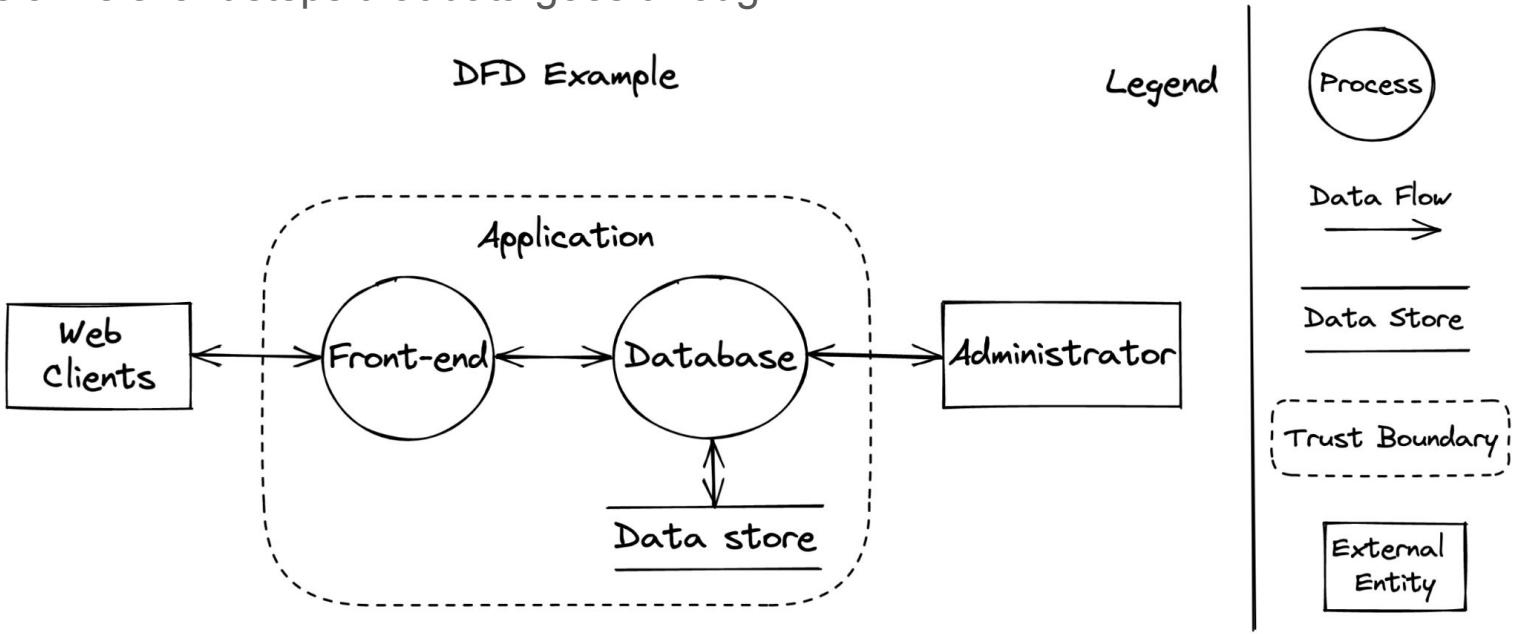

Experiment set-up

Goal: to measure the effectiveness and usefulness of DFDs in identifying and understanding security threats, using the STRIDE methodology.

Scenario: deploying a pod on a Kubernetes cluster.

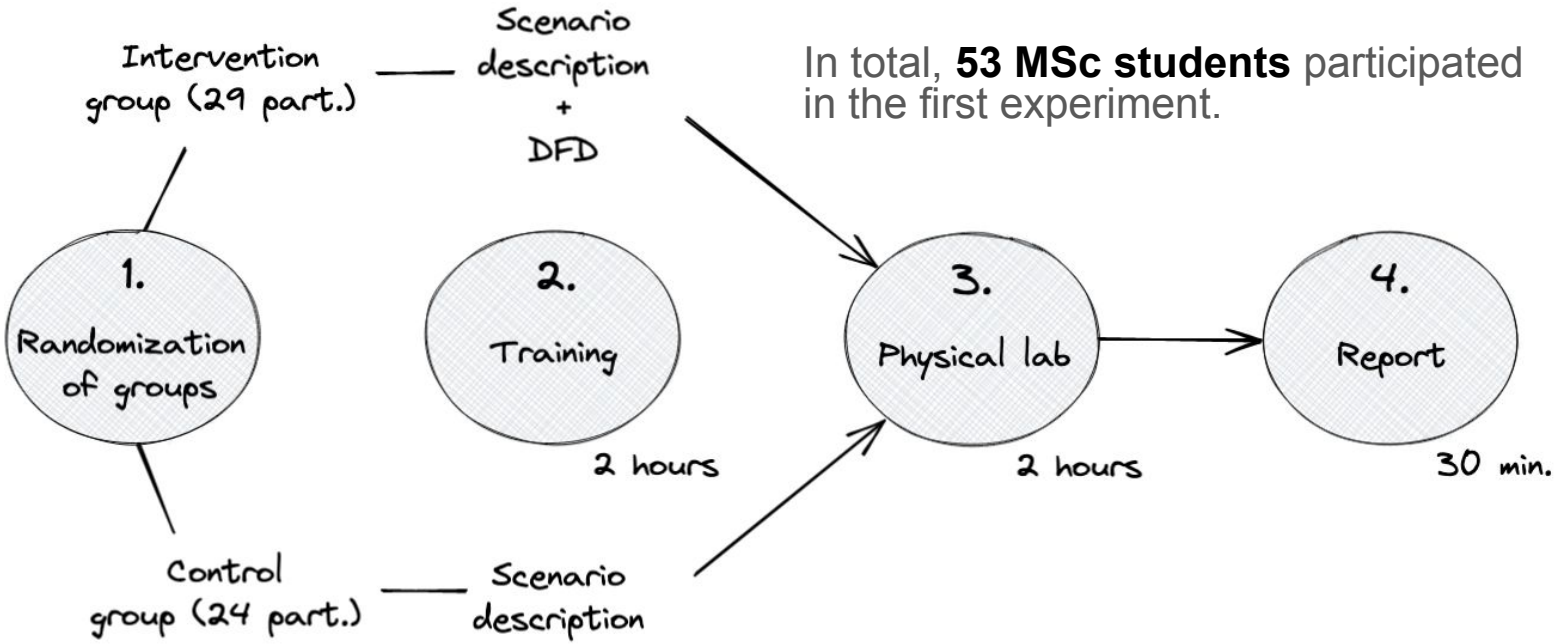

In total, **53 MSc students** participated in the first experiment.

Experiment Training - Kubernetes DFD

The current DFD represents the (simplified) steps of deploying a Pod on Kubernetes.

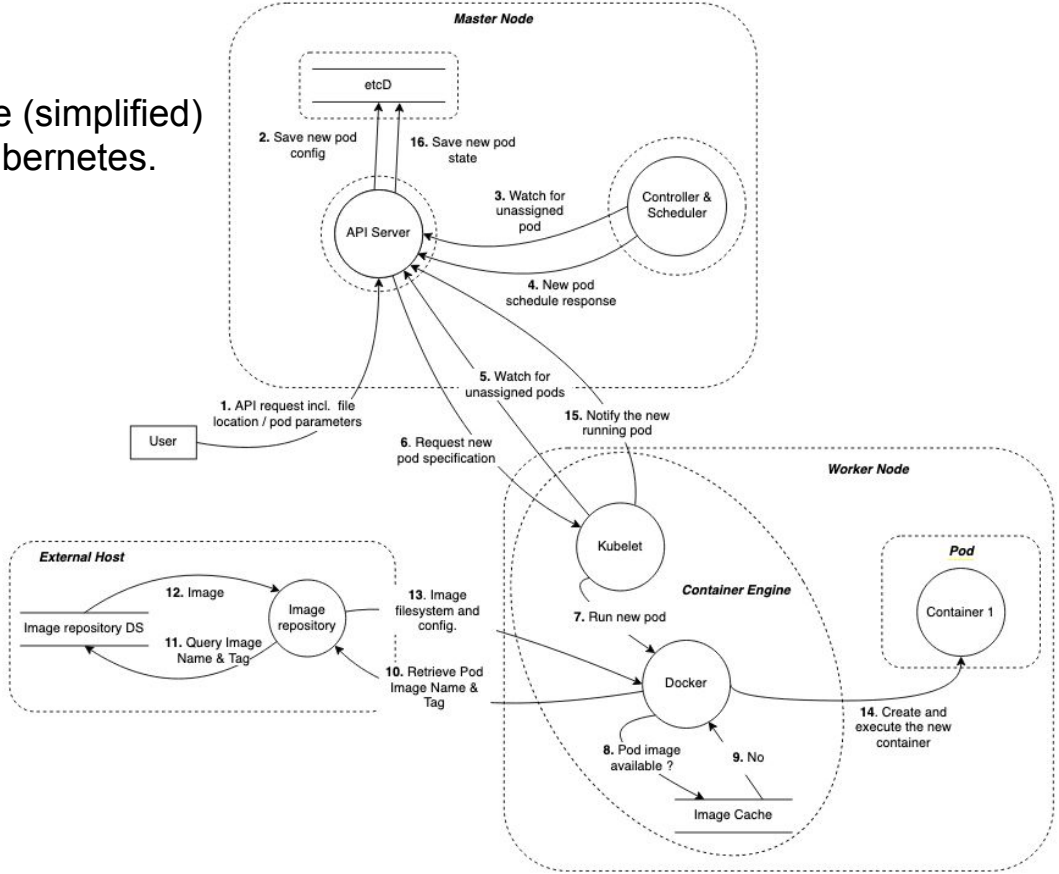

Experiment lab - Threats

We defined a list of (5) real threats and (5) fake threats.

Threat Description	STRIDE Threat	Assumptions	Affected Components
Privileged containers run as root on the host, thus an attacked compromising one of such containers automatically gets root access on the host	Elevation of privilege	• The attacker has access to the container (e.g., exposed on the internet) • The attacker escapes the container	Pod and worker node
…	…	…	
If an attacker compromises a pod in a Kubernetes cluster using a Layer 3 CNI plugin, s/he can steal other pod's identities and laterally move within the cluster using the network bridge	Spoofing	• The attacker exploits one container • The CNI plugin works at layer 3	(Lateral movement between) pods

Metrics of Success

Metric 1.

- **True Positives (TP)** = real threats being selected.
- **False Negatives (FN)** = real threats not being selected.
- **False Positives (FP)** = fake threats being selected.

Precision = TP / (TP + FP)

Recall = TP / (TP + FN)

Metric 2.

Perceived usefulness of the graphical model, i.e., DFD, on a scale from 1 to 5..

Preliminary Results - Metric 1

Summary of the participant's performance.

Metric	(mean) Group A	(mean) Group B	Mann–Whitney U (MWW) test
True Positives (TP)	3.17	2.83	statistic = 277.5 pvalue = 0.09
False Negative (FN)	1.83	2.16	statistic = 277.5 pvalue = 0.09
False Positive (FP)	1.93	2.08	statistic = 331.0 pvalue = 0.37
Precision	0.64	0.60	statistic = 300.5 pvalue = 0.19
Recall	0.63	0.57	statistic = 277.5 pvalue = 0.09

Preliminary Results - Metric 2

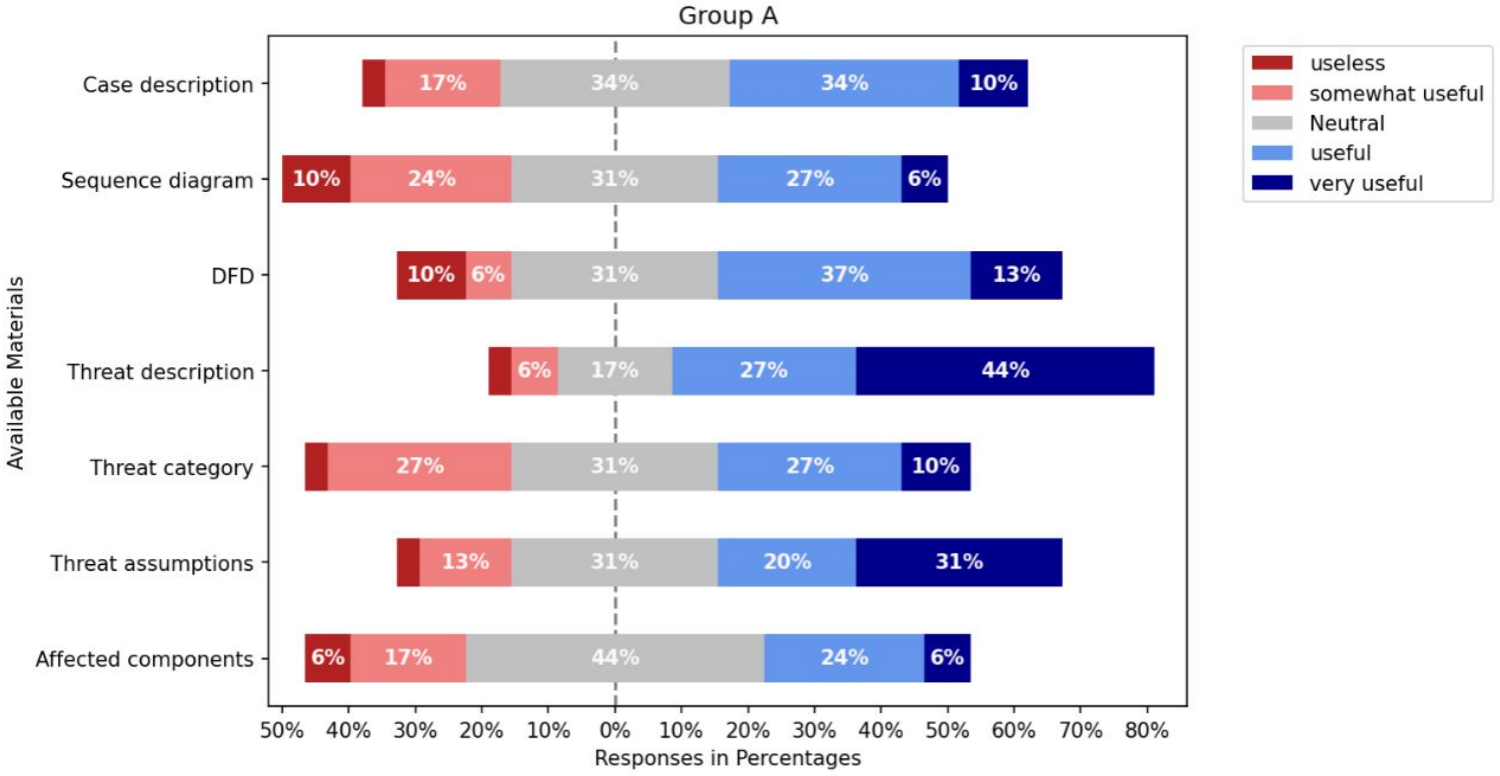

Conclusion and Future Work

At the end of this experiment, we can conclude that DFDs graphical representations do NOT significantly help in identifying threats in Kubernetes clusters.

To expand our contributions, we plan to:

1. Repeat the experiment with experts
   a. Defining new threats
2. Evaluate scalability: are DFDs (more or less) useful when *more complex* scenarios (i.e., DFDs are larger) are considered?
3. Evaluate the effectiveness of other graphical models (e.g., sequence diagrams)

In general, more experiments are required to validate our hypothesis.

List of the REAL threats

Threat Description	STRIDE Threat	Assumptions	Affected Components	Real Threat
An attacker uses leaked (e.g. on Github) cluster configuration file (e.g. /etc/kubernetes/admin.conf) to interact with the cluster, potentially jeopardize running workloads and creating malicious pods.	Elevation of privilege Information Disclosure	- The attacker finds the leaked information; - The attacker can (remotely) interact with the cluster.	- Not present in the DFD.	Yes
Spoofing a cluster admin by stealing the authentication credentials via a social engineering attack.	Spoofing	- The attacker carries out a successful spoofing attack; - The credentials are valid.	- Cluster configuration - Running workloads (pods, container, etc.) - Access control (of users and services) In general, he/she can take over the all cluster.	Yes
An attacker with shell access to only one pod which has no resource limits applied (e.g. max memory and cpus usage) can crash the worker node on which that pod is running, causing a Denial of Service (DoS).	DoS	- The attacker got a remote shell on a pod.	- Compromised pod, worker node, and all the other pods running on the same node.	Yes
An attacker able to upload images to a container registry from which a Kubernetes cluster retrieves pods images, he or she can potentially execute a malicious pod inside the cluster and get a shell from the malicious container.	Elevation of privilege	- The attacker has permissions to upload or modify images in the registry used by the Kubernetes cluster;	- Images Registry - Running pod	Yes
Privileged containers run as root on the host, thus an attacker compromising one of such containers, gets automatically root access on the host.	Elevation of privilege	- The attacker get acccess to a privileged container running in the cluster; - The attacker escapes the container.	- Pod - Worker node	Yes

List of the FAKE threats

If an attacker deploys a malicious pod into a namespace of the cluster with network policies in place, the attacker can use such a pod to send bogus network packages to pods in different namespaces to jam the exposed ports, causing a disruption in the cluster performance.	DoS	- The network policies are implemented correctly to segment the namespaces.	- Pods in different namespaces.	No
If a developer embeds a secret in a Dockerfile A and then builds an image (image A), an attacker with access to image A can reverse the image's layers to not only observe the embedded secret in Dockerfile A but also from other images that were built from different Dockerfiles.	Information Disclosure	- The attacker has access to the image	- Not present in the DFD - Cluster secrets/credentials	No
If an attacker compromises a running container, he/she can only continue exploiting the container by running the software specified in Dockerfile (by the developers) with malicious inputs.	Elevation of privilege Tampering	- The attacker got access to a running container.	- Running pod/container	No
An unauthenticated and non privileged attacker can still upload custom pod configuration into the cluster to prepare the first step of another attack, e.g. turning off logging service or cause a denial of service.	Elevation of privilege Tampering	- The attacker can reach the cluster API server (e.g. through internet)	- New pods	No
If an attacker compromises a pod in a Kubernetes cluster using a Layer 3 network solution, he/she can steal other pod's identities and laterally move within the cluster using the network bridge.	Spoofing	- The attacker exploits one container - The CNI works at layer 3	- Lateral movement between pods	No

Sequence Diagram

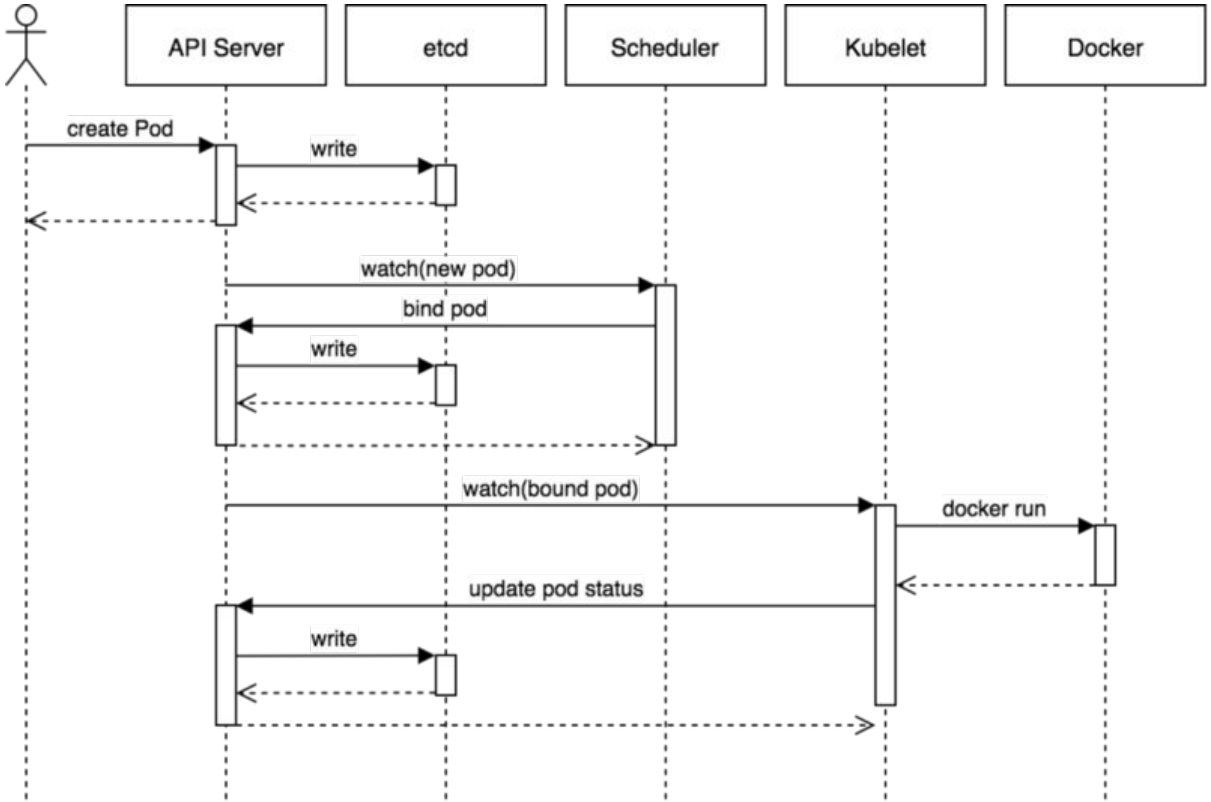

Box Plot of the groups performance

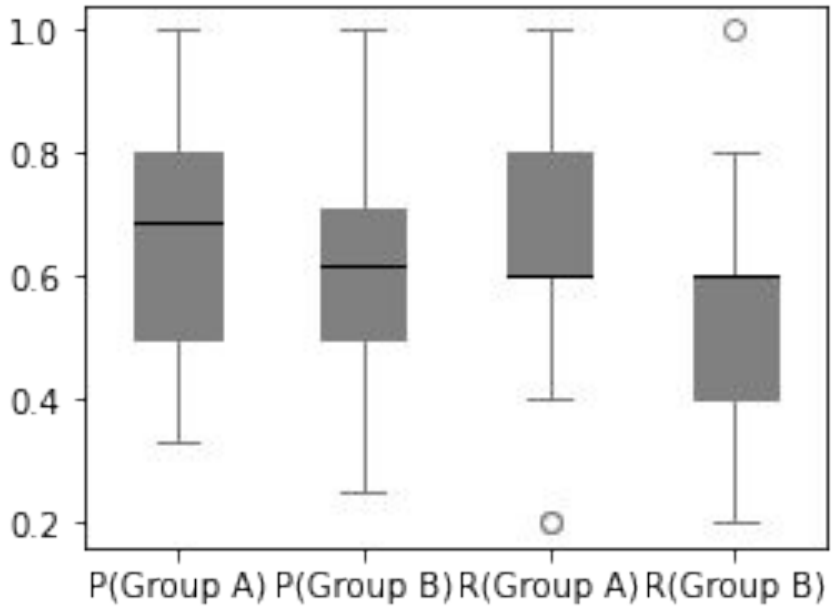

Preliminary Results - Metric 2

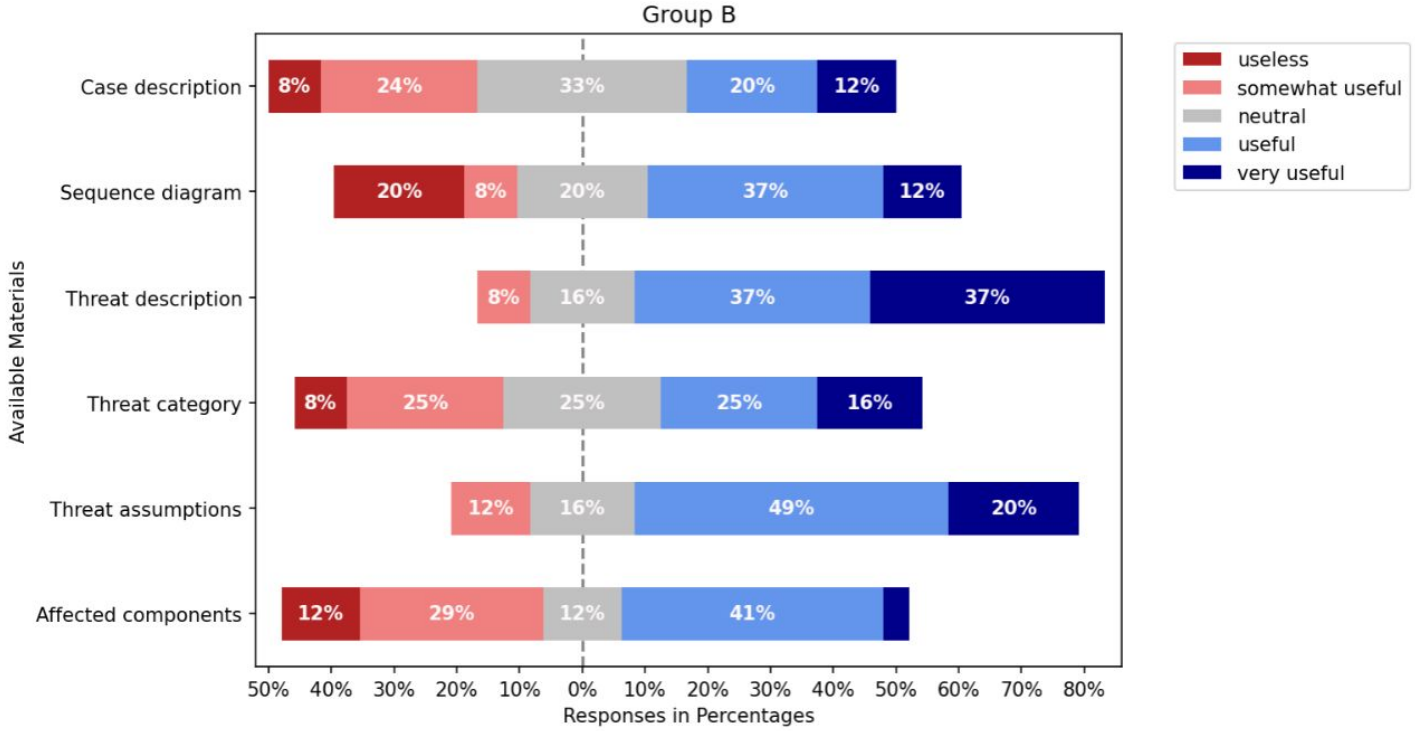

References

[1] AssureMOSS https://assuremoss.eu/en/

[2] Kubernetes Network paper https://www.computer.org/csdl/magazine/sp/2021/05/09497237/1vzYhimJUK4

[3] STRIDE https://www.microsoft.com/security/blog/2007/09/11/stride-chart/

[4] K8s DFD https://github.com/cncf/financial-user-group/tree/main/projects/k8s-threat-model