



CARVE SYSTEMS, LLC

June 5, 2020

Container Based Applications

Defense in Depth

Version 1.0

<https://carvesystems.com>

The material contained in this document represents proprietary, confidential information pertaining to Carve Systems, LLC ("Carve") products and services. The client hereby agrees that the information in this document shall not be disclosed outside of the Client, without prior, written approval from Carve. Client will have the right to duplicate, use or disclose the material contained herein to the extent provided for in the contract relating to this work.

Contents

1	Introduction	3
1.1	A Brief History of Run Time	3
1.2	The Three Problems	4
1.3	How to Read This Paper	4
2	Threat Model	5
2.1	Environments	6
2.2	Components	6
2.3	Flows	7
3	Host Operating Systems	8
3.1	Threats	8
3.2	Tools	9
4	Containers	10
4.1	Threats	10
4.2	Tools	12
5	Container Orchestrators	13
5.1	Threats	13
5.2	Tools	14
6	Bonus Problems	15
6.1	Managed Clusters	15
7	Conclusion	17

Introduction

Some people, when confronted with a problem, think “I know, I’ll use a container orchestrator.” Now they have three problems. - Carve¹

1.1 A Brief History of Run Time

Computer software evolves like lots of things in life. Various layers accrue over time. These layers make some problems more manageable, but increase the overall complexity in the system.

One of the most complex pieces of software is the foundational layer of Operating Systems.² There has been a lot of focus and engineering effort put into evolving OSes for decades. This effort has been from a functional point of view and a security point of view.

From a security point of view, lots of the engineering effort has been spent on how best to enforce the Principle of Least privilege.³ Historically, in Unix OSes, much of the model was based on what ordinary users and “root” users are allowed to do.⁴ However, over time, this model proved to not be granular enough. Your choices are essentially to have very limited access or access to everything.

In Linux, which is the Unix operating system we will focus on, there have been various attempts to add this granularity. AppArmor⁵ and SELinux⁶ introduced mandatory access control (MAC)⁷ implementations. Namespaces⁸ allow you to partition kernel resources like mounts, process IDs, networking stack, and user IDs such that each process thinks they have their own set of unique resources. Capabilities⁹ added a way to control what operations a particular thread can perform (this effectively provides the needed granularity for operations so that users don’t need root or suid). Cgroups¹⁰ added a way to throttle resource use of memory, CPU, and I/O. Finally, seccomp¹¹ added the ability to precisely control what system calls are allowed to be invoked by a process.

Many of these features have been used to implement OS-level virtualization¹² in the form of a container. There are several alternative implementations for containers on Linux like Docker,¹³ LXC,¹⁴ and rkt.¹⁵ We will focus on Docker.

¹<http://regex.info/blog/2006-09-15/247>

²https://en.wikipedia.org/wiki/Operating_system

³https://en.wikipedia.org/wiki/Principle_of_least_privilege

⁴<https://en.wikipedia.org/wiki/Superuser>

⁵<https://en.wikipedia.org/wiki/AppArmor>

⁶https://en.wikipedia.org/wiki/Security-Enhanced_Linux

⁷https://en.wikipedia.org/wiki/Mandatory_access_control

⁸https://en.wikipedia.org/wiki/Linux_namespaces

⁹<http://man7.org/linux/man-pages/man7/capabilities.7.html>

¹⁰<https://en.wikipedia.org/wiki/Cgroups>

¹¹<https://en.wikipedia.org/wiki/Seccomp>

¹²https://en.wikipedia.org/wiki/OS-level_virtualization

¹³[https://en.wikipedia.org/wiki/Docker_\(software\)#Components](https://en.wikipedia.org/wiki/Docker_(software)#Components)

¹⁴<https://linuxcontainers.org/lxc/introduction/>

¹⁵<https://coreos.com/rkt/>

Once a development team has decided to containerize their application a new problem arises – how are containers deployed and managed? Container orchestration¹⁶ systems were engineered to solve that problem. The main job of a container orchestrator is to ensure that the desired containers composing a software system are running and healthy.

Now we have our three problems.

1.2 The Three Problems

Modern web applications are increasingly being deployed in containerized platforms. Development teams that build and manage these platforms on their own must manage at least three complex layers of software:

1. The host systems running the containers and cluster nodes.
2. The containers running on the host systems and the tools used to build and configure them.
3. The container orchestration system used to deploy and manage the containers.

The security of all of these layers is extremely important as they build on top of each other. Appropriately securing each layer is essential for defense in depth of the containerized platform.¹⁷

This whitepaper will focus on defense in depth measures for these three layers by giving a set of guidelines to think about for each layer. The focus will be for infrastructure and development teams building container-based applications that need guidance on important threats to think about.

These guidelines will use Docker as the reference implementation for containers and Kubernetes for the reference implementation of container orchestration. Although, these guidelines apply to other container implementations and orchestration systems as well.

1.3 How to Read This Paper

We will start out by providing a very basic threat model for the three aforementioned layers in the next section. The following sections will detail the enumerated threats and guidelines for managing them. We also provide some pointers to helpful tools that can be used for findings specific instances of these threats.

The intended path is for development, security, and infrastructure engineers building their own clusters to read the threat model to get acquainted with the general architecture and then walk through each of the three Host Operating System, Container, and Container Orchestrator sections and think about how the threats detailed there relate to what they are building.

That being said, the three threat sections are also useful in isolation. For example, if you are just working on containerizing your Application, then the Container section is still useful.

¹⁶[https://en.wikipedia.org/wiki/Orchestration_\(computing\)](https://en.wikipedia.org/wiki/Orchestration_(computing))

¹⁷[https://en.wikipedia.org/wiki/Defense_in_depth_\(computing\)](https://en.wikipedia.org/wiki/Defense_in_depth_(computing))

Threat Model

Threat modeling is the art of describing and documenting the security-relevant aspects of a complex system in a simplified form. It often starts with a good diagram, and we recommend using a relatively strict and simplistic type of diagram called a Data Flow Diagram (DFD) for developers and teams relatively new to threat modeling. The data flow diagram identifies the components in a system, which then allows for a per component analysis. For each component the threat model then aims to understand “what can go wrong?”.

In this section we will develop a basic threat model for a Kubernetes deployment. We will use the “STRIDE”¹ threat enumeration technique:

Threat	Desired Property
Spoofing	Authenticity
Tampering	Integrity
Repudiation	Non-repudiability
Information Disclosure	Confidentiality
Denial of Service	Availability
Elevation of Privilege	Authorization
Outdated Software	All of the above

The threats will be enumerated against the following DFD:

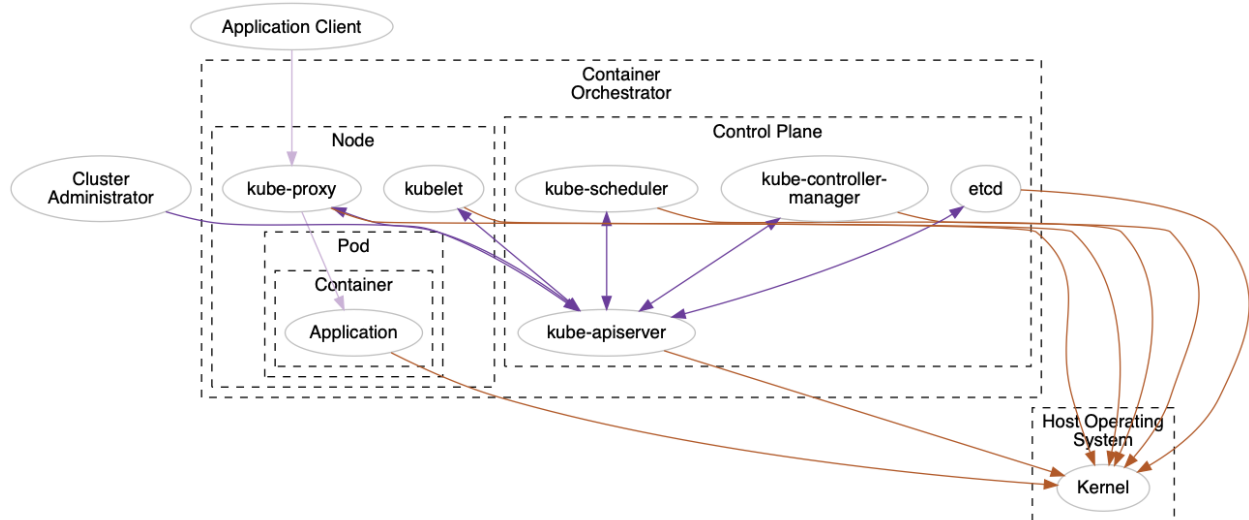


Figure 2.1: Dataflow Diagram

NOTE:

¹[https://en.wikipedia.org/wiki/STRIDE_\(security\)](https://en.wikipedia.org/wiki/STRIDE_(security))

- The Host Operating System in this diagram could be one or several physical or virtual host machines.
- The Pod in this diagram could be one or more instances that may or may not communicate with each other.

The objects of interest in this diagrams are **environments** (the dashed boxes), **components** (the ovals), and **flows** (the arrows):

- **Environment:** A container with one or more components or other sub-environments. An environment may be a physical element such as a data center or a software environment such as a virtual machine. Environments are either trusted or untrusted. Trusted environments require an attacker to breach protections, such as data validation or locked doors, in order to influence the components contained within the system. Environments are also used to indicate ownership boundaries between different organizations.
- **Component:** A data processing element that is evaluated as a non-divisible entity in the threat model. Components may be processes, servers, or a software function implemented by a number of servers in coordination with each other.
- **Flow:** Represents the transfer of data between components. Flows that cross from an untrusted to a trusted environment (and vice versa) are of particular interest when threat modeling.

The specific objects used in this model will be defined in the following subsections.

For more details about threat modeling see Threat Modeling:² Designing for Security (Shostack 2014)

2.1 Environments

The environments in this model define the trust boundaries of the Host Operating System, Container, and Container Orchestrator.

Environment	Description
Container	Contains Operating System objects such as processes, mounts, and networking stacks.
Container Orchestrator	Contains all entities related to the management and deployment of containers. As referred to as a cluster.
Control Plane	Contains all entities related to managing the cluster itself.
Host Operating System	Contains all the kernel and user space components that make up the Operating System.
Node	Contains all entities related to running the containers.
Pod	Contains one or more container instances.

2.2 Components

There are two classes of components we will consider: User Agents and Software Agents. User agents are people that interact with other components through software APIs. Software Agents are the primary software components in the system.

2.2.1 User Agents

Component	Description
Cluster Administrator	A user that administers the Container Orchestrator environment.
Application Client	A user that interacts with the Application through the Application API Call.

2.2.2 Software Agents

Component	Description
Application	An software program running in a Container.
etcd	A distributed key-value store that holds critical data for the cluster.
Kernel	The Host Operating System kernel.
kube-apiserver	The core API server used to manage the cluster. Access is through a REST API.

²Shostack, A. (2014). Threat Modeling: Designing for Security. Wiley.

Component	Description
kube-controller-manager	A daemon made up of control loops that maintain the desired state of the cluster.
kube-proxy	A network proxy responsible for proxying traffic from the kube-apiserver to a Node.
kube-scheduler	A scheduling service that is responsible picking the optimal Node for a Pod to run on.
kubelet	An agent running in a Node that ensures that the desired containers are running and healthy.

2.3 Flows

Flows in this model represent API calls from the various components in the Container Orchestrator environment.

Flows	Description
Application API Call	An API call from the Application Client to the Application.
Kubernetes API Server Call	An API call from one of the Container Orchestrator components to the kube-apiserver.
syscall	A system call from an Application to a Kernel.

Host Operating Systems

Host Operating Systems should be constructed such that it is hard for an attacker to elevate their privileges or pivot to other systems.

Elevation of privileges could allow an attacker to take over the Control Plane and have full control of the cluster. It could also allow an attacker to have complete control of the Applications running on the Nodes.

A Host Operating System may be vulnerable to several forms of attack. The following subsections outline some of the primary threats.

3.1 Threats

3.1.1 Spoofing

- Ensure that strong passwords and hashing algorithms (like SHA512 with 200000 rounds) are used to **authenticate** users.
- If SSH access is used to provide remote access, then ensure that the SSHD configuration provides strong **authentication** options:
 - PermitEmptyPasswords should be set to NO.
 - PermitRootLogin should be set to NO.
 - IgnoreRhosts should be set to YES.
 - HostbasedAuthentication should be set to NO.
 - RhostsRSAAuthentication should be set to NO.

3.1.2 Tampering

- Ensure that the **integrity** of directories and files containing data such as keys, passwords, and configuration settings are **not** world writeable or writeable by an excessive number of users through group settings. Such settings can allow an attacker to tamper with files (under /etc/ for example) which may lead to access or control of critical services.

3.1.3 Repudiation

- Ensure that appropriate system-level logging with syslog¹ is enabled and configured correctly. Accurate logging can help determine the nature of attacks, where they are coming from, and be used to argue **non-repudiability**.

3.1.4 Information Disclosure

- Ensure that directories and files containing **confidential** data such as keys, passwords, and configuration settings are **not** world readable or readable by an excessive number of users through group settings. Such settings can give an attacker important information that can be used to access other parts of the system.

¹<https://linux.die.net/man/5/syslog.conf>

- Ensure that passwords are hashed with a cryptographically secure hashing algorithms with key stretching.² The increase number of rounds will make sure that brute force attacks consume more time and decrease the likely hood that the **confidentiality** of user passwords is violated.
- Limit software name and version information from userspace programs that discloses **confidential** implementation details from appearing in things like banners and HTTP headers. These details can give attackers extra information needed to construct more targeted attacks.
- Limit the number of open TCP and UDP ports. Leaving unnecessary ports open increases the attack surface and can leak **confidential** details about the services in use on the system.

3.1.5 Denial of Service

- Ensure that resources used by critical processes are limited by appropriately using Control Groups³ to prevent DoS attacks and maintain **availability**. Some components like systemd already make use of cgroups and have configuration options for it.⁴

3.1.6 Elevation of Privilege

- Strictly limit the number of users that need sudo access as it is equivalent to **authorization** as root.
- Avoid setuid⁵ binaries. Processes running with the setuid have effective privileges equivalent to an **authorized** root user.
- Limit the permissions of /tmp to non-execute to prevent **unauthorized** users for saving and executing malicious code from /tmp.

3.1.7 Outdated Software

- All host operating systems should be frequently updated. Any kernel exploits can allow an attacker to leverage instances of all threats such as spoofing, tampering, repudiation, information disclosure, and escalation of privileges. Particularly serious vulnerabilities could allow an attacker to escape containers and escalate privileges to the root user.

3.2 Tools

The following scripts are useful for helping identify some of the threats detailed in previous section. They are good for establishing a **baseline** for parts of your system that should be analyzed in more detail.

- The **Linux Basic Security Audit (LBSA)** script⁶ is useful to find weak SSH configurations, poor file permission settings, and other vulnerabilities.
- The **Linux Privilege Escalation Check** script⁷ is useful for findings common privilege escalation paths.
- **Lynis**⁸ searches for common vulnerabilities, poor configurations settings, and poor permission settings.

²https://en.wikipedia.org/wiki/Key_stretching

³<https://en.wikipedia.org/wiki/Cgroups>

⁴<https://www.freedesktop.org/software/systemd/man/systemd.resource-control.html>

⁵<https://en.wikipedia.org/wiki/Setuid>

⁶<http://wiki.metawerx.net/wiki/LBSA>

⁷<https://github.com/sleventyeleven/linuxprivchecker/blob/master/linuxprivchecker.py>

⁸<https://cisofy.com/lynis/>

Containers

Docker containers and their environment should be constructed so that an attacker may not pivot from a compromised container to further exploit other backend services via the network interface, the host operating system, or adjacent containers.

Container compromise occurs when an attacker effectively gains full or partial control of a running container. This can result from flaws in public-facing application logic that allows code injection, the use of vulnerable container images, or injection of malicious images into the CI/CD pipeline.

A docker container may be vulnerable to several forms of attack.¹ The following subsections outline some of the primary threats.

4.1 Threats

4.1.1 Spoofing

- Ensure that the **identity** of a container image is always verified. An image² may be spoofed by replacing an image in a registry or through interception when an image is being transmitted to the registry from a CI/CD pipeline. Consider using something like Docker Content Trust³ to sign your images.

4.1.2 Tampering

- The **integrity** of the image should be verified. An image⁴ may be tampered with by comprising a registry or through interception when it is being transmitted to a container orchestrator. Consider using something like Docker Content Trust⁵ to sign your images.
- Attempt to make the container filesystem read only⁶ to help enforce the **integrity** of the filesystem contents. Otherwise, only make writable the locations that are absolutely necessary for the implementation of the microservice or consider using tmpfs mounts (such as⁷) for writable storage.

4.1.3 Repudiation

- Configure logging and regularly monitor the logs so that suspicious behaviors and attack patterns can be identified by Cluster Administrators.⁸ This will allow **non-repudiation** to be enforced. Docker has several logging drivers⁹ that make it easy to integrate with other services to simplify logging and monitoring.

¹This threat overview is adapted from the [OWASP Docker Top Ten](#) project as of March 20, 2019. Please bear in mind that this documentation project is new and under active development.

²Docker Image. <https://docs.docker.com/v17.09/glossary/?term=image>

³https://docs.docker.com/engine/security/trust/content_trust/

⁴Docker Image. <https://docs.docker.com/v17.09/glossary/?term=image>

⁵https://docs.docker.com/engine/security/trust/content_trust/

⁶Docker read-only option. <https://docs.docker.com/engine/reference/commandline/run/#mount-volume-v--read-only>

⁷Docker tmpfs mounts. <https://docs.docker.com/storage/tmpfs/>

⁸<https://docs.docker.com/config/containers/logging/>

⁹<https://docs.docker.com/config/containers/logging/configure/>

4.1.4 Information Disclosure

- Never store secrets in plain site within a container. This included in environment variables and on the filesystem. Storing secrets in plain site in a container can allow an attacker access to **confidential** information. If the container was compromised in any way an attacker would trivially have access to them. Consider using a secrets manager like Docker secrets¹⁰ or HashiCorp Vault.¹¹
- Only expose that ports that are **absolutely** necessary for for the execution of your application. For example, microservices **should** normally only be using **80** or **443**. Opening unnecessary ports in a container can allow attackers to compromise or glean **confidential** information about a microservice.
- Only mount host filesystems that are **absolutely** necessary for the functioning of the microservice. Mounting host file systems inside a container can allow attackers to modify or extract **confidential** information about the host.

4.1.5 Denial of Service

- Ensure that resources within a container are always limited by appropriately using Control Groups¹² to prevent DoS attacks and maintain **availability**. A container administrator may incorrectly configure the container allowing an attacker to exhaust resources such as CPU cycles and available memory. Docker has several options to enforce these limits.¹³

4.1.6 Elevation of Privilege

- Strictly limit the number of users that need access to the Docker socket. Access to the Docker socket is equivalent to **unauthorized** root access on the **host** machine.¹⁴
- Only mount the Docker socket inside a container when **absolutely** necessary. A compromised container with access to the Docker socket can easily escalate privileges to **unauthorized** root access on the **host** machine.¹⁵
- Always follow the principle of least privilege and **never** grant your containers more capabilities than are absolutely necessary using tools like SELinux or AppArmor profiles. Poorly managed Linux capabilities¹⁶ for Docker containers can greatly increase the attack surface of the host operating system and allow **unauthorized** access to system resources.
- Avoid creating privileged Docker containers with the `--privileged` flag. This flag grants the container almost the same access on the host as processes running outside of the container.¹⁷ This increased level of privilege could allows **unauthorized** agents access to system resources.

4.1.7 Outdated Software

- Frequently rebuild the Docker images to refresh the base Operating System and userspace components to the latest known good versions. Docker images that have not been updated properly allow an attacker to leverage instances of all threats such as spoofing, tampering, repudiation, information disclosure, and escalation of privileges.
- Frequently update the Docker software itself. Docker implementations that have not been updated properly allow an attacker to leverage instances of all threats such as spoofing, tampering, repudiation, information disclosure, and escalation of privileges.

¹⁰<https://docs.docker.com/engine/swarm/secrets/>

¹¹<https://www.vaultproject.io/>

¹²<https://en.wikipedia.org/wiki/Cgroups>

¹³Docker resource constraints. https://docs.docker.com/config/containers/resource_constraints/

¹⁴Docket Socket. <https://docs.docker.com/engine/security/https/>

¹⁵Don't expose the Docker socket (not even to a container). <https://www.lvh.io/posts/dont-expose-the-docker-socket-not-even-to-a-container.html>

¹⁶Linux capabilities. <http://man7.org/linux/man-pages/man7/capabilities.7.html>

¹⁷<https://docs.docker.com/engine/reference/run/#runtime-privilege-and-linux-capabilities>

4.2 Tools

The following scripts are useful for helping identify some of the threats detailed in previous section. They are good for establishing a **baseline** for parts of your system that should be analyzed in more detail.

- **Docker Bench**¹⁸ is a set of auditing scripts that is used to find common flaws like lack of resource limits, privileged containers, missing MAC profiles, and unsafe volume mounts.

¹⁸<https://github.com/docker/docker-bench-security>

Container Orchestrators

Container Orchestrators should be appropriately configured to use proper **authentication**, **authorization**, and use encrypted information channels that preserve **integrity** and **confidentiality**.

Broken **authentication** could allow an attacker to access the kube-apiserver to read sensitive information from the cluster or make damaging modifications. For similar reasons, appropriate **authorization** should be used to limit the scope of resources that one can access. Finally, all Kubernetes API Server Call should happen over encrypted channels to preserve **integrity** and **confidentiality**.

The orchestrator under consideration here is Kubernetes.¹ It is primarily accessed through the `kubectl`² command. The following subsections outline some of the primary threats.

5.1 Threats

5.1.1 Spoofing

- Enforce **authentication** to the kube-apiserver to ensure that only known parties can interact with the APIs. Kubernetes provides several methods of authentication.³ The most common methods are Client Certificates and Service Account Tokens.
- Do **not** enable `--insecure-port` in production.⁴ This is intended for testing only and disables the **authentication** and **authorization** modules.

5.1.2 Tampering

- Use TLS for all API traffic to protect the **integrity** of the data being transmitted.⁵

5.1.3 Repudiation

- Enabled auditing⁶ so that suspicious behaviors and attack patterns can be identified by Cluster Administrators. This will allow **non-repudiation** to be enforced.

5.1.4 Information Disclosure

- Use TLS for all API traffic to protect the **confidentiality** of the data being transmitted.⁷

¹Kubernetes. <https://kubernetes.io/>

²kubectl. <https://kubernetes.io/docs/tasks/tools/install-kubectl/>

³<https://kubernetes.io/docs/reference/access-authn-authz/authentication/>

⁴<https://kubernetes.io/docs/reference/access-authn-authz/controlling-access/#api-server-ports-and-ips>

⁵<https://kubernetes.io/docs/tasks/administer-cluster/securing-a-cluster/#use-transport-layer-security-tls-for-all-api-traffic>

⁶<https://kubernetes.io/docs/tasks/debug-application-cluster/audit/>

⁷<https://kubernetes.io/docs/tasks/administer-cluster/securing-a-cluster/#use-transport-layer-security-tls-for-all-api-traffic>

5.1.5 Denial of Service

- Use resource quotas and limit ranges to control the size and capacity of resources allocated to a namespace.⁸ Putting limits on resources will help ensure the continuous **availability** of the cluster.

5.1.6 Elevation of Privilege

- Use Role-based access control (RBAC) to enforce **authorization** to all APIs. Kubernetes provides an integrated RBAC component.⁹
- Restrict read and write access to etcd to the smallest group possible. Per the Kubernetes documentation, giving write access is equivalent to granting admin access to the cluster and giving read access provides a quick path for **unauthorized** users to escalate their privileges.¹⁰
- Only use user impersonation¹¹ (`--as`) when strictly necessary and audit all existing uses of it carefully. Careless use of this functionality can allow lower privileged users **authorized** access to a wider range of resources in the cluster.
- Use security contexts¹² to establish the minimum permissions that Pods and Containers are **authorized** to have in order to properly function.

5.1.7 Outdated Software

- All host operating systems should be frequently updated. Any kernel exploits can allow an attacker to leverage instances of all threats such as spoofing, tampering, repudiation, information disclosure, and escalation of privileges. Particularly serious vulnerabilities could allow an attacker to take control of the cluster.

5.2 Tools

The following scripts and Kubernetes plugins are useful for helping identify some of the threats detailed in previous section. They are good for establishing a **baseline** for parts of your system that should be analyzed in more detail.

- **kube-hunter**¹³ is useful for mapping the attack surface and discovering known vulnerabilities in Kubernetes clusters.
- **access-matrix**¹⁴ is a krew¹⁵ plugin that discovers and visualizes the RBAC authorization matrix.
- The `kubectl auth can-i`¹⁶ command is a simply and easy way to determine what actions the current user can perform.

⁸<https://kubernetes.io/docs/tasks/administer-cluster/securing-a-cluster/#limiting-resource-usage-on-a-cluster>

⁹<https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

¹⁰<https://kubernetes.io/docs/tasks/administer-cluster/securing-a-cluster/#restrict-access-to-etcd>

¹¹<https://kubernetes.io/docs/reference/access-authn-authz/authentication/#user-impersonation>

¹²<https://kubernetes.io/docs/tasks/configure-pod-container/security-context/>

¹³<https://github.com/aquasecurity/kube-hunter>

¹⁴<https://github.com/corneliusweig/rakness>

¹⁵<https://github.com/kubernetes-sigs/krew>

¹⁶<https://kubernetes.io/docs/reference/access-authn-authz/authorization/#checking-api-access>

Bonus Problems

As we have seen, there are Host Operating Systems, Containers running on Host Operating Systems, and Container Orchestrators managing Containers. Surely that is the end of the line – no. As always, another layer has grown to manage and create Container Orchestrators.

There are several popular platforms used to implement this bonus layer. We will focus on Amazon Elastic Kubernetes Service (EKS), Google Kubernetes Engine (GKE), and Azure Kubernetes Service (AKS). As with all software, these new layers mitigate some threats and create new ones. The following sections will give a brief overview of these platforms and how they can help mitigate some of the threats discussed in previous sections.

A detailed treatment of the security properties of these platforms could easily fill another whitepaper and is not the central focus of this paper. As such, the following sections are strictly meant to be a brief overview for how these platforms relate to the layers and threats that were explored in detail in the preceding sections.

6.1 Managed Clusters

Each of the platforms discussed allows clients to easily create managed Kubernetes Clusters. Instead of having to directly manage your own Control Plane and Nodes the platform provides an API to create them for you. The Control Plane runs in an account fully managed by platform provider. The Nodes run in the creating user's account.

6.1.1 Amazon Web Services EKS

The following table lists AWS services that help mitigate threats that we discussed in the preceding sections:

Threat	Services
Spoofing	IAM
Tampering	KMS
Repudiation	CloudWatch
Information Disclosure	KMS
Denial of Service	AWS Shield
Elevation of Privilege	IAM integration with RBAC and ConfigMap
Outdated Software	All of the above

When deploying Applications in EKS the above service integrations should be used and carefully configured. In addition, review Amazon's documentation on securing your cluster.¹

6.1.2 Google Cloud Platform GKE

The following table lists Google services that help mitigate threats that we discussed in the preceding sections:

Threat	Services
Spoofing	Cloud IAM, GKE Cluster Trust
Tampering	Cloud KMS, GKE Cluster Trust
Repudiation	Cloud Monitoring, Cloud Logging
Information Disclosure	Cloud KMS, GKE Cluster Trust

¹<https://docs.aws.amazon.com/eks/latest/userguide/security.html>

Threat	Services
Denial of Service	Cloud Armor
Elevation of Privilege	Cloud IAM integration with RBAC, GKE Sandbox
Outdated Software	All of the above

When deploying Applications in GKE the above service integrations should be used and carefully configured. In addition, review Google's documentation on hardening your cluster's security.²

6.1.3 Azure AKS

The following table lists Microsoft services that help mitigate threats that we discussed in the preceding sections:

Threat	Services
Spoofing	Azure Active Directory
Tampering	Azure KeyVault
Repudiation	Azure Monitor
Information Disclosure	Azure KeyVault
Denial of Service	Azure Monitor, Azure DDoS Protection
Elevation of Privilege	Azure Active Directory integration with RBAC
Outdated Software	All of the above

When deploying Applications in AKS the above service integrations should be used and carefully configured. In addition, review Microsoft's documentation on hardening your cluster's security.³

²<https://cloud.google.com/kubernetes-engine/docs/how-to/hardening-your-cluster>

³<https://docs.microsoft.com/en-us/azure/aks/concepts-security>

Conclusion

As we have seen there are several complex layers involved when creating modern container-based applications. Each of these layers require a significant commitment to security in order to ensure the system as a whole is secure. A commitment that is **not** a one-time analysis, but an ongoing one.

The threats and tools given in this paper will help with that ongoing analysis. Furthermore, they will hopefully allow the reader to enumerate their own threats and apply proper security engineering practices¹ in their constantly evolving software architectures.

¹<https://carvesystems.com/wp-content/uploads/2020/05/Carve-Systems-White-Paper-Security-Engineering.pdf>