

The Pennsylvania State University
The Graduate School

CO-RESIDENCY ATTACKS ON CONTAINERS ARE REAL

A Thesis in
Computer Science and Engineering
by
Sushrut D. Shringarputale

© 2019 Sushrut D. Shringarputale

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

December 2019

The thesis of Sushrut D. Shringarputale was reviewed and approved* by the following:

Patrick McDaniel
William L. Weiss Professor of Information and Communications Technology
Thesis Advisor

Thomas La Porta
William E. Leonhard Professor, Computer Science and Engineering and
Electrical Engineering

Chitranjan R. Das
Distinguished Professor of Computer Science and Engineering
Head of the Department of Computer Science and Engineering

*Signatures are on file in the Graduate School.

Abstract

Public clouds are inherently multi-tenant: applications deployed by different parties (including malicious ones) may reside on the same physical machines and share various hardware resources. With the introduction of newer hypervisors, containerization frameworks like Docker, and managed/orchestrated clusters using systems like Kubernetes, cloud providers downplay the feasibility of co-tenant attacks by marketing a belief that applications share no part of the stack. In this paper, we demonstrate that attackers can confirm co-residency with a victim application from inside state-of-the-art containers running on virtual machines. We analyze the degree of vulnerability present in containers running on various systems including within a broad range of commercially utilized orchestrators. Our results show that on realistic cloud environments, we can obtain 90% success rates for co-residency detection. Our investigation confirms that co-residency attacks are a significant concern on containers running on modern orchestration systems.

Table of Contents

List of Figures	vi
List of Tables	vii
List of Symbols	viii
Acknowledgments	ix
Chapter 1	
Introduction	1
Chapter 2	
Background	4
2.1 Containers	5
2.2 Orchestrator implementation	5
2.3 Virtual machines and isolation	6
2.4 Co-residency Attacks	7
Chapter 3	
Co-resident Watermarking Attack	10
3.1 Attack Phase	11
3.2 Co-residency Detection Phase	13
Chapter 4	
Evaluation	16
4.1 Experimental setup	17
4.1.1 Simple setup	17
4.1.1.1 <i>Local</i>	17

4.1.2	Realistic setup	18
4.1.2.1	<i>Local Separated</i>	18
4.1.2.2	<i>Quiet Cloud</i>	18
4.1.3	Non-quiescent environment	18
4.1.3.1	<i>Noisy Cloud</i>	18
4.2	Co-residency Detection	20
4.2.1	Local and Local Separated	20
4.2.2	Cloud Environment	22
4.2.3	Non-quiescent setup	23
4.2.3.1	Realistic noise setup	23
4.2.4	Adversarial deception	24
4.3	Orchestrator Comparison	25
4.3.1	<i>Local</i>	25
4.3.2	<i>Local Separated</i>	26
Chapter 5		
	Discussion and Future Work	29
5.1	Effect of noise	29
5.2	Why does this attack generalize?	29
5.3	Defenses	30
Chapter 6		
	Conclusion	32
	Bibliography	33

List of Figures

2.1	Components of orchestration platform	5
3.1	Attack setup diagram	11
3.2	Expected trace of RTTs	13
3.3	Expected histogram of RTTs	13
3.4	Expected histogram of RTTs	13
4.1	RTT trace for one sample run	19
4.2	RTT Histograms for one sample run	20
4.3	ROC curves for controlled experiment across different system configurations	21
4.4	Success rates by varying experimental variables: l and f	24
4.5	System performance variation based on orchestrator	25
4.6	RTT trace for various orchestrators	27

List of Tables

- 4.1 Experimental Variables 17
- 4.2 Successful detection rates across orchestrators and configurations . . 21
- 4.3 T-test statistic values for each orchestrator 28

List of Symbols

- $aMSE$ Mean Square Error value after noise is factored out
 $aMSE_0$ Threshold $aMSE$ value for detection

Acknowledgments

This research was sponsored by the U.S. Army Combat Capabilities Development Command Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Combat Capabilities Development Command Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

Dedication

To my parents, Dhawal and Chitra, and my brother Suhrud. Your unconditional love and support have made this possible.

Chapter 1

Introduction

Cloud computing adoption has grown exponentially in the last decade, with revenue of \$175 billion in 2018 [1] and \$206.2 billion projected in 2019. Several diverse industries have begun migrating workloads to the cloud including healthcare, financial services, technology, and insurance [1]. Dependence on cloud services continues to increase as all companies, from small startups to large multi-national corporations and governments, store and process their data on the cloud. The data handled by many of these organizations can be highly sensitive. In most cases, customers run their code on the same physical machine as others (a state called *co-residency*). Those other customers could include potential adversaries, and this co-residency exposes a major attack surface. Despite the isolation between processes and virtual machines, shared access to hardware resources introduces security and privacy concerns for tenants running on the same hardware.

Some reported attacks against data-centers include NordVPN's recent data-center breach, which took place on an undisclosed third-party provider, leading to leakage of TLS keys [2]. Misconfiguration by Tesla for a Kubernetes cluster led to a cryptojacking attack and exposure of secret access keys to the AWS account, leading to attackers launching crypto-currency mining software on Tesla's compute resources on AWS [3]. In cases where misconfiguration was not the issue, a rare disclosed datacenter breach showed a targeted attack where the adversaries employed Trojan malware to compromise websites hosted on a Central-Asian datacenter [4]. Cases like these show that attackers are employing sophisticated mechanism to attack data-centers with specific targets. As providers patch misconfiguration vulnerabilities,

the adversaries are bound to turn towards attacks on the virtualization layer. Zero-day vulnerabilities have been publicly exploited in virtualization software in controlled environments [5], which are bound to make their way into actual exploits. In order to locate targets then, we suspect that attackers will implement co-residency detection on clouds environments as a major attack vector.

So far, Cloud Service Providers (CSPs) have adopted hypervisors, e.g. *KVM*, container runtimes, e.g. *Docker*, and orchestration systems, e.g. *Kubernetes* to streamline access to cloud resources. CSPs depend on low-layer abstractions such as hypervisors and container sandboxes [6] to provide security from co-residency attacks. The processes themselves are run using abstractions that elide a shared-hardware view: providing an illusion that the virtual machines and processes share no part of the hardware with others. A study by Eder [7] claims that the combination of container-based isolation and hypervisor-based virtualization may provide a much more improved security model. The direction of engineering and research in the industry indicates that this belief is widespread. Development of orchestration systems regularly downplay the potential impact of co-residency. CSPs are introducing multi-tenant, managed Software-as-a-Service implementations of container orchestrators that that make it easier for customers to run their applications [8, 9]. However, these implementations share the cluster-level abstractions between tenants, making co-residency more likely.

The growth of such orchestration systems has fundamentally changed the way cloud platforms manage virtualized execution of computation. In these systems, containers become the fundamental unit of computation, rather than customers managing discrete VMs per application. Containers are rapidly poised to replace VMs, as made evident by the recent development of Kata containers [10]. These containers claim to have the security of VMs, while still remaining lightweight and easy to run.

Research by Atya et al. [11, 12], Bazm et al. [13] have confirmed that virtual machines are vulnerable to co-residency attacks. The structure of containers induced speculation that similar vulnerabilities also exist for containers, even when run within virtual machines and orchestration systems. For example, Gao et al. [14] showed that the way containers mount the `procfs` pseudo file system exposes many exploitable side-channels.

In this paper, we are the first to confirm the feasibility of co-residency attacks on containers in modern orchestration systems. We adapt the co-residency watermarking attack [15] which uses hardware side-channels to establish whether an adversary is co-resident with its target. We show that the attack is successful when deployed against containers running on virtual machines, even within commercially utilized, modern orchestration systems. We analyze the effectiveness of the attack on a large variety of system configurations (*Local*, *Local Separated*, *Quiet Cloud*, and *Noisy Cloud*) and measure the sensitivity of the attack to architecture, load, orchestrator selection, and adversarial activity. Using modern orchestrator implementations (*Docker Swarm*, *Kubernetes*, *Helios*), we show that variations in these choices have little effect on the success of such an attack. Our evaluation yielded results with 90% success rates on cloud-like environments.

Finally, we analyze how inconspicuous an adversary can be during the attack and show that an adversary can detect co-residency with little work, making it much harder to detect.

We make the following observations in this paper:

- We empirically confirm that systems running on container runtimes within virtual machines are vulnerable to co-residency attacks, with a very high (>90%) detection rate on cloud-like systems.
- We demonstrate that co-residency detection can tolerate background noise at a 70% success rate, as long as it does not exceed hardware capacity.
- We show that changes to the architecture and choice of orchestrator reduce the fidelity of detection by at most 10%, demonstrating little effect on adversarial success.
- We analyze the amount of adversarial work necessary for the adversary to be successful in detecting a target and confirm that an adversary can be successful by generating interference without reaching the hardware capacity, around 70% of the capacity in one of our experiments.

Background

Orchestration platforms have ushered in a new way of virtualizing execution on shared cloud systems. They provide a way for administrators to manage the deployment of their applications at a cluster-level abstraction. The primary goal of orchestration platforms is to help orchestrate the placement, scaling, and lifecycle of applications in a distributed manner, providing high availability and capabilities to respond to large spikes in traffic and demand quickly. This removes the effort needed to plan how the applications run. As a result, large scale applications are quickly adopting these strategies because fast scaling, cost effectiveness, and rolling updates leading to lower downtimes help make their systems much more reliable. Applications recently have been growing in size and complexity, requiring layers of authentication, web-capabilities, integration with other platforms, caching, and storage. To make these applications easier to maintain, engineers are developing and deploying them as smaller units, or *micro-services*, that can be scaled and managed independently.

Google's Borg project led to the development of the Kubernetes orchestrator [16] which helps automate the process of assigning containers to physical machines. In addition, the platform performs various additional virtualization at the compute, device, and network levels to simplify application development. There are other developers that have invested significant effort into such systems, such as Docker Swarm [17], Spotify's Helios [18], and Apache Mesos [19].

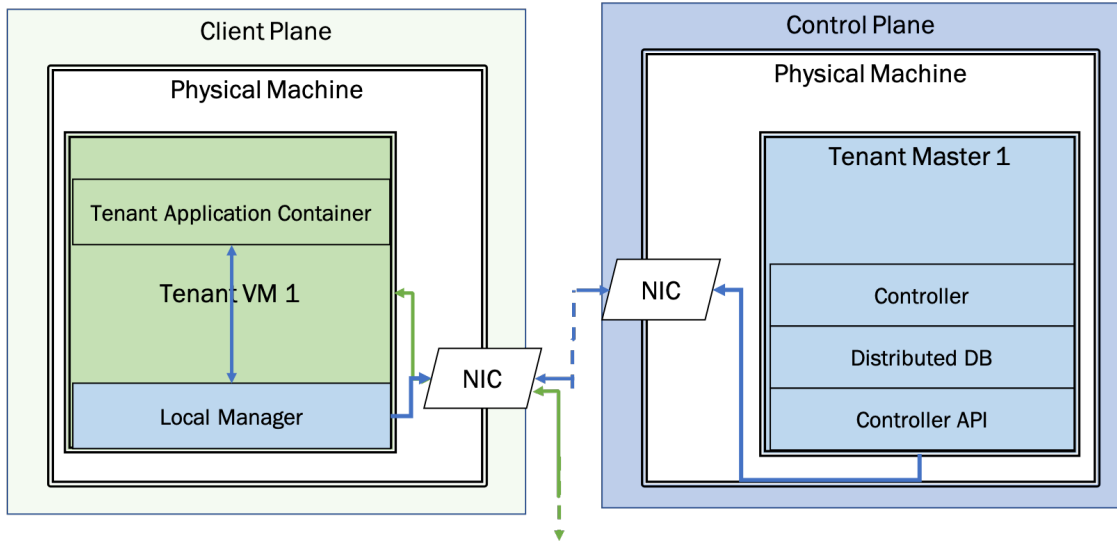


Figure 2.1: Components of orchestration platform

2.1 Containers

Micro-services demand that each running service be able to start and stop quickly (scale), have a small footprint, and automate easily. This makes virtual machines a very poor choice. Creating a VM per application would require a very long setup time, requiring a new initialization of the operating system per service instance, leading to much higher resource usage. To solve this problem, the industry is rapidly switching over to a new type of process-isolated runtime for applications. This runtime uses Linux namespaces, control groups, and the Union File system (OverlayFS [20]) to make the application system-agnostic as well as easy to spin-up and isolate. This runtime defines the process as a *container*, which contains all the libraries necessary to run the application. By building the blueprint for this ahead of time, the spin-up time and resource usage for containers tends to be very low [21]. Due to these reasons, containers are becoming the standard unit of computation within micro-services.

2.2 Orchestrator implementation

Most orchestrators are quite similar in structure and differ in architectural structures (See Figure 2.1). The core concept includes a coordinator application (Master) that

communicates in a distributed way with all the physical nodes (VMs) running the client’s applications. The coordinator uses a distributed database that stores the latest state of the cluster of client containers. The state includes which containers need to run, how many instances of a container are needed active, which node they should run on, and other metadata. All these administrating pieces together constitute the *control-plane* of the cluster.

The control-plane is responsible for continuously monitoring the cluster, so it implements health checks that verify the states of all container instances using periodic HTTP calls. The control-plane is also responsible for implementing a cluster network to allow all the containers to communicate with each other using private IP addresses. Each client node runs a local manager that the coordinator communicates with. The local manager communicates with the container runtime and operating system to keep track of the running containers and their health. Finally, it provides the firewall rules and port mappings necessary to implement the cluster network subnet needed by the containers running on that node. Utilizing the health checks, if a container is in a non-ideal state, the control-plane takes steps to fix that. In most cases, the control-plane is managed by the cloud-provider and is physically separated from client applications.

We looked into the source of three orchestrators listed above: Docker Swarm, Kubernetes, and Helios. Docker Swarm and Kubernetes expose HTTP-based APIs on each node in the cluster that the coordinators make requests to. Thus, the majority of computation is performed by the coordinators, independent from the nodes. On the other hand, Helios implements a poll-based paradigm, where each node queries the state database directly to compute the action to take for the cluster to reach the “goal state”. Since the orchestrators is critical to the operation of the cluster, best practices suggest that they be separated from the actual compute nodes by administrators [22].

2.3 Virtual machines and isolation

CSPs such as Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform, etc., provide services to their customers with a set of guarantees (i.e., service level agreements). While it is evident to customers that they are sharing

the physical hardware with other customers, hereby referred to as *tenants*, the cloud platform also provides certain isolation guarantees for any code running on these systems. More specifically, the operating system provides isolation guarantees between processes that ensure that the memory in use is completely isolated. Access Control mechanisms in the operating system provide isolation at file and device levels. Similarly, hypervisors provide isolation to virtual machines running complete operating systems. These isolation guarantees are supposed to ensure that tenants' code will not be adversely affected by others from a security, performance, and privacy perspective. Since each layer of isolation describes a new attack surface, it is important to that each layer is properly secured.

2.4 Co-residency Attacks

A co-residency attack considers an active, malicious adversary that has no affiliation to the cloud provider [23]. The adversary appears as an innocuous tenant to the CSPs with no elevated privileges on the actual machine they are logged onto. They access the standard interface for launching instances, same as the victim, are free to launch as many VM instances on the cloud service as they want (barring some limits for specific CSP), and can choose the hardware configuration the instances will have. This hardware configuration is applied at the VM level, however, certain CSPs tend to restrict the type of VM instance that can be placed on a specific type of machine [24]. This means that the adversary has the capability to choose from a subset of physical machines in some cases. For example, if an adversary decides to use *m2.large* instances on AWS, they will most likely be placed on machines specified for *m2.large* instances [25]. As a result, co-resident placement becomes much simpler. Zhang et al. [26] shows that influencing co-tenancy with a 90% probability takes only 20 instances on a chosen instance type.

The victim is a legitimate cloud tenant performing some security sensitive operations using hardware that is shared with the adversary. The victim has the same capabilities as the adversary when it comes to launching instances. Both members in this model trust the cloud provider and have little to no insight into the placement policies of the datacenter. Additionally, the adversary gains little to no information from metadata provided to the instance such as the IP address.

While this information was useful in older attacks for cloud cartography [23], this is no longer viable [25] on current cloud architectures due to the introduction of Virtual Private Clouds, where all user VMs are launched in a logically isolated section of the cloud, giving the user the ability to pick their IP addresses [27].

Instead, the adversary aims to use one of the available side channels in order to gain unauthorized access to some information belonging to the victim. This can range from knowledge of existence [15, 23], to application data [28], to cryptographic private keys [29]. Gao et al. [14] show the extent of channels that can be used by containers on cloud providers, including seemingly benign information such as per-core power usage.

Most of these studies looked at the isolation provided by virtual machines on clouds, assuming that each application runs within its own virtual machine. However, containers and orchestrators seem to provide an added abstraction layer that may improve the isolation guarantees. Indeed Eder [7] claims that the addition of containerization-based isolation within hypervisor-based virtualization may be an improved security model. The direction CSPs have taken with their development of orchestration services seem to certainly imply that this is the case [8, 9]. For example, managed orchestration clusters support multiple tenants on an identical set of virtual machines, vastly increasing the possibility of co-residency. Honig and Porter [30] claim that the KVM hypervisor is secure enough to protect from side-channel attacks since the vulnerabilities in libraries are patched by the developers quickly. Allclair and kaczorowski [31] from Google suggest that the layers of isolation in Kubernetes generate enough security for containers on VMs and orchestrators. Recent research by Atya et al. [11] shows that this is certainly not the case for VMs.

Sources intuit that containerized isolation is not secure, even running within virtual machines and orchestrators. Edwards et al. [22] recently completed a security review of all systems within the Kubernetes architecture and outlined some of the major attack vectors and best practices relating to clusters running on Kubernetes. They note that multi-tenant clusters are a major attack surface for orchestrators today. However, their definition of multi-tenancy differs from ours in the following way: they define a multi-tenant cluster as an orchestrated collection of containers deployed on compute nodes that have containers deployed by multiple tenants. This may induce co-resident multi-tenancy, but the attack surface explored by the

paper looks at access control and over-privilege as the main issues to protect from. Our definition deals with applications that are assigned the same physical node to containers deployed by two unrelated tenants on *different clusters*. As a result, the attack surface we investigate deals with adversaries circumventing the combined isolation guarantees provided by the hypervisor and container runtime using some side-channels.

Co-resident Watermarking Attack

The goal of this study is to ascertain that containers running on virtual machines are also vulnerable to co-residency attacks, regardless of system configuration. In order to achieve this goal, we picked a representative attack that requires reasonably low setup cost. The attack attempts to determine if the adversary's container is running on a VM that is co-resident to their victim's container. We then analyze the various sensitivities of the attack to the architecture, system load, orchestrator choice, and adversarial activity.

For measuring architectural sensitivity, we run the attack on a quiescent setup where only the adversary and the victim are active. With this setup, we are able to show that the attack is possible and ensure that the measurements obtained are not due to external interference. After this, the attack is re-run on a cloud-like environment to analyze the effect of other tenants on the success of the attack.

Similar to realistic cloud environments, we analyze the effect of other load in the system to the attack. In addition to the victim and adversary, we analyze the attack's success when other processes and tenants are demanding a high amount of CPU and network bandwidth.

The third sensitivity we measure is the choice of orchestrator. We show that the presence and choice of an orchestrator plays a minimal role in increasing resiliency against the attack. Finally, we vary the amount of activity the adversary performs and measure how deceptive an adversary can be while still succeeding at detecting the co-resident victim.

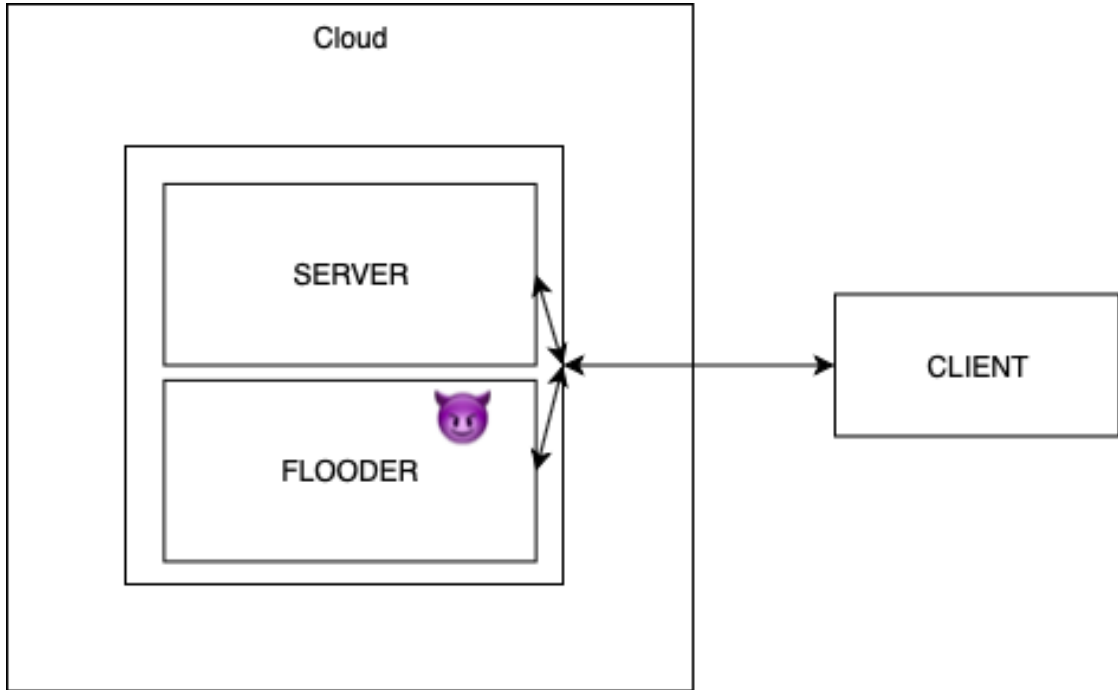


Figure 3.1: Attack setup: FLOODER and SERVER are co-resident while the CLIENT communicates with them externally

3.1 Attack Phase

We use the co-resident watermarking attack [15], which works against public-facing web servers to detect when a malicious VM is running co-resident to the server. For the purposes of this paper, we found that the watermarked network-flow attack for detecting co-resident victims was appropriate as it can run very quickly (<2 minutes for each run) and has a much higher success rate than other attacks. In our setup, we have three applications running inside containers that play the following roles: the SERVER, the CLIENT, and the FLOODER as shown in Figure 3.1.

The SERVER is the modeled victim in this attack and has no knowledge of the adversary. It listens for web requests, which for simplicity were restricted to one type, where the SERVER sends a 10 MB file in chunks of 8 KB.

The CLIENT is controlled by the adversary and is aware about the co-residency detection attack in progress. It maintains communication with the adversary that is potentially co-resident with the SERVER, and it measures round trip time (RTT)

```

1 set p as period
2 set d as duration
3 notify_client()
4 loop for d seconds // On sequence
5     notify_client()
6     loop p seconds: // Watermarking Sequence (On-WM)
7         send_burst_packet()
8     end loop
9     notify_client() // Notify end of watermarking sequence
10    sleep(10-p) // Idling Sequence (On-Idle)
11    notify_client() // Notify end of idling sequence
12 end loop
13 notify_client() // Notify Start of off sequence
14 loop while client.active()
15     sleep(p) // Watermarking Sequence (Off-WM)
16     notify_client() // Notify end of watermarking sequence
17     sleep(10-p) // Idling Sequence (Off-Idle)
18     notify_client() // Notify end of idling sequence
19 end loop

```

Algorithm 3.1: Flooder implementation

for requests made to the **SERVER**. The **CLIENT** continues to make requests to the **SERVER** until the adversarial container i.e. **FLOODER** notifies completion.

Finally, the **FLOODER** is a container that generates interference in the network interface card (NIC) of the physical machine by flooding it with large amounts of data. In order to minimize the interference from the hypervisor, the **FLOODER** signals the **CLIENT**, runs for period p , signals the **CLIENT**, and idles for the next $10 - p$ seconds. The code for flooder is provided in Algorithm 3.1.

The flooder runs in two main sequences: *On* and *Off*. During the *On* sequence, the **FLOODER** runs in a loop for d seconds. In this loop, the first p seconds are the watermarking sequence (*On-WM*), while the idling sequence (*On-Idle*) is made up of the $10 - p$ second window that follows. During *On-WM*, the **FLOODER** generates large bursts of traffic sent out of the machine, approaching close to the capacity of the NIC. Due to the generated contention in accessing the NIC, we expect that the Round Trip Time (RTT) for the **CLIENT**'s requests should increase here.

The **FLOODER** signals the **CLIENT** the end of the *On* sequence and goes into the *Off* sequence. This sequence is useful to obtain a model of the environmental noise (load l), which is factored out during the co-residency detection phase. During this sequence, the **CLIENT** makes the same RTT observations with signals from the **FLOODER**. The **FLOODER** sleeps for p seconds (instead of actually sending packets), which is the *Off-WM* sequence. After notifying the client, the **FLOODER** sleeps again




Figure 3.2: Expected trace of RTTs: The trace shows a square wave (left) when the flooder causes interference (*On* sequence), while a near constant RTT when the flooder is not causing interference (right) (*Off* sequence)




Figure 3.3: Expected histogram of RTTs: The histograms show clear separation when the flooder causes interference (*On* sequence)




Figure 3.4: Expected histogram of RTTs: The histograms look identical when the flooder is not causing interference (*Off* sequence)

for $10 - p$ seconds, which is the *Off-Idle* sequence. We expect that the RTTs should be nearly identical during the entire *Off* sequence. The expected RTT trace for both the *Off* and *On* sequences is shown in Figure 3.2. Using the collected RTT trace, the *CLIENT* analyzes whether the *FLOODER* is actually co-resident with the *SERVER*.

3.2 Co-residency Detection Phase

Varadarajan et al. [25] describe the performance covert-channel based detection mechanisms for co-residency detection: If an adversary is co-resident with the victim and is generating contention that leads to a measurable performance degradation

for the victim, it can be measured by the adversary to detect co-residency. Prior work have used various methods to analyze this drop in performance on various different attack surfaces. When generating contention in the memory bus [25], the authors compared the means and medians of the distributions, taking into account the noise of the system when the adversary was not active. Building on this methodology, when detecting co-resident VMs in VPC enabled clouds, Xu et al. [27] thresholded the latency of a trace-route to a possible candidate VM in order to determine co-residency. Finally, Bates et al. [15] generated contention in the NIC and measured the difference in distributions of the packets that were watermarked and ones that were not. They measured the difference in these distributions using the Kolmogorov-Smirnov statistical test, confirming co-residency if the test rejects the null hypothesis.

The closest prior methodology for detection to our work is the network watermarking attack [15], however, we found during our evaluation that applying the KS-test to round trip times does not work. Due to the high variance in the RTTs, we found that unless the measured values generated the ideal distributions described in Figure 3.4, the KS-test generated extremely high false-positive rates. For our case, it was sufficient to use a simple distribution distance metric such as the Mean Square Error (MSE). We calculate the MSE value between the histograms for the *On-WM* and *On-Idle*, since histograms where the FLOODER has generated enough interference should be separate, as demonstrated in Figure 3.3. In contrast, in scenarios where the FLOODER is not actually co-resident, the distributions will be nearly identical, as shown in Figure 3.4. Thus, the MSE value should be high when the FLOODER is co-resident, while close to zero when the FLOODER is not.

$$MSE = \frac{1}{n} \sum_{i=1}^n (Y_i - \hat{Y}_i)^2 \quad (3.1)$$

The final piece to detection is generating the correct threshold MSE value. We need to determine the threshold to separate the configurations that are co-resident from those that are not. In the case of the ideal clean square wave, that is trivial. Any MSE value greater than zero should be a successful detection, since there will be no effect generated when the adversary is not co-resident. However, during our evaluation we found that as the system gets noisier, the distributions for the *On-*

WM and *On-Idle* sequences get closer together while the distributions for *Off-WM* and *Off-Idle* show pseudo-separation. This makes it difficult to determine a ground-truth threshold value for the MSE.

In order to determine and eliminate the effect of this background noise, we use a model of the background noise by obtaining an RTT trace when the FLOODER is inactive (*Off* sequence). This technique is similar to the technique used when determining co-residency based on memory bus access contention [25]. Using the *Off* sequence, we generate a baseline MSE value for the system without the activity of the FLOODER. Since the *Off* sequence is obtained immediately after, it should provide a reasonable measurement of the background noise. We subtract this value (MSE_{off}) from the (MSE_{on}) and use the result to generate the correct threshold.

$$aMSE = MSE_{on} - MSE_{off} \quad (3.2)$$

Since the model of the background noise may change on each system we evaluate, we run a controlled experiment where we assume that the adversary has access to the SERVER and can force co-residency. Using the $aMSE$ value from these runs, we generate the Receiver Operating Characteristic (ROC) curve, find the Pareto point on the curve given by Equation 3.3, and obtain the optimal threshold value ($aMSE_0$) for that system configuration.

$$aMSE_0 = \operatorname{argmax}(TruePositive - FalseNegative) \quad (3.3)$$

Chapter 4

Evaluation

As part of the evaluation, we try to answer the following questions:

- Is co-residency detection possible when running containers on virtual machines?
- Is co-residency detection affected when containers run on varying system configurations and how?
- Do orchestrators play a role in the success of co-residency detection? Which components of orchestration software play the largest role?
- What is the effect of system load and architecture on success metrics?
- How deceptive can an adversary be before co-residency detection is no longer possible? What is the least amount of data transmission needed by the adversary for the attack to succeed?
- Is there a difference in how much network-watermarking co-residency detection reacts to the vendor of orchestrator used? If yes, why?

For each measured sensitivity, we establish the experimental variables listed in Table 4.1.

Variable Name	Variable Symbol	Variable Description
Period of sequence	p	How long the flooding sequence (<i>On-WM</i>) lasts
Detection threshold	$aMSE_0$	The threshold at which an adversary detects co-residency for the given environment
Load	l	The amount of load in the system around the shared hardware
Orchestrator choice	o	The chosen orchestrator for the system
Adversarial Activity	f	Amount of adversarial activity (flooding)

Table 4.1: Experimental Variables

4.1 Experimental setup

We performed our experiments with two different setups: the first emulated a small, simple cloud environment and the second mimicked a realistic cloud provider.

4.1.1 Simple setup

4.1.1.1 *Local*

Our simple setup involved just two physical nodes capable of running workloads. We setup two orchestration clusters: one for the target container (**SERVER**) and the other controlled by the adversary. Each cluster had its coordinating application, i.e. Master, running on the same node, meaning that the Master node VM for each cluster was allowed to run client containers directly. While not being very realistic, this allowed us to generate a footprint for the amount of computation and IO interference generated by the Master, which is a significant piece in any orchestration system.

Since the role of the Master is to simply designate a specific node (including itself) to run the container, perform health checks, and virtualize access to other portions of the system, we expected that once an application has been successfully assigned to a node, i.e. scheduled, the orchestration system should generate minimal interference in the actual containers' performance.

4.1.2 Realistic setup

4.1.2.1 *Local Separated*

Running client containers on the Master node does not reflect realistic cloud system setups. Cloud providers generally handle the setup for the Master node(s) and keep them separate from the user’s applications [32]. To mimic this, we separated out the Master nodes and CLIENT onto different physical machines. With this, we were able to isolate the victim and adversary with no other system-level interference (barring the OS and hypervisor).

In such a setup, the only piece of software from the orchestration platform that runs on the same machine as the application is the local node manager. This software only responds to health-check HTTP requests and ensures that the application container is healthy. We expect much lesser interference from the orchestration here and the attack should be much more robust.


4.1.2.2 *Quiet Cloud*

We created the same setup as above on a realistic, cloud-like environment. The Master node ran on separate physical hardware, with the local manager running along side the client applications. On this setup, other tenants did not share the physical hardware with the SERVER and FLOODER, however, other tenants were active in the local network, generating more noise than the *Local and Local Separated* setups.

4.1.3 Non-quiescent environment

4.1.3.1 *Noisy Cloud*

We also ran the attack in a setup that includes a lot more system noise on the victim. To emulate this, we used the popular benchmarking tool Yahoo! Cloud Serving Benchmark (YCSB) [33] to continuously benchmark Memcached, a fast, persistent key-value store. This generates a high number of disk IO and network calls on the system. During this environment we kept the Master node separate to ensure that the load in the system is predictable.



(a) Measured RTT during a flooding sequence: The flooding sequence generates a square wave watermark that confirms co-residency

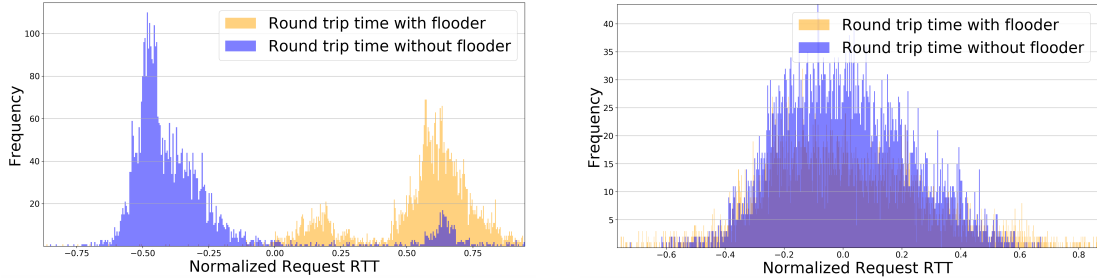
(b) Measured RTT when adversary is not co-resident: The flooding sequence does not generate a square wave watermark

Figure 4.1: RTT trace for one sample run

In addition to the YCSB benchmark, we use a realistic server trace of traffic as background noise to evaluate how realistic load in the system will affect the detection rates. We use a recent trace (April 2019) from the MAWI dataset [34] to mimic realistic background traffic. We adjust the multiplier on the trace replay to increase the contention in the NIC due to noise.

The local setups (*Local* and *Local Separated*) for experiments used two Dell Precision T7600 machines with identical internal configurations. They both had 32 GB RAM, six-core Intel Xeon E5-2630 processor, Ubuntu 18.04, and were connected to a local network over ethernet with a Cisco Linksys E1550 router. The NIC for target machine had a 1 Gbps ethernet connection. A third machine, a 2013 Mac Pro was used as a packet sink and existed on the local network, but took no computational part in the actual attack. The KVM hypervisor was used as the virtualization layer, which is the standard in industry [35].

The experiments were repeated on a cloud-like cluster (*Quiet Cloud* and *Noisy Cloud*) which contains sixteen compute nodes. Each node is a Dell PowerEdge M620 machine running eight-core Intel Xeon E5-2609 processors with 64 GB RAM. All nodes were running Ubuntu 18.04 Server and were connected to the cluster network. Each node communicated through a central NIC for the cluster, which had a bandwidth of 1 Gbps. Any node that required virtualization used the KVM hypervisor like above. The cluster is currently in use by other researchers to run various workloads and applications, making it similar to an Infrastructure-as-a-Service (IaaS) cloud service. Other tenants of the cluster were accessing some nodes on the same shared network, leading to slightly higher background noise than *Local Separated* and *Local*.



(a) Normalized RTT histogram for a flooding sequence: The dark histogram has clear separation from the light, determining co-residency

(b) Normalized RTT histogram for a control sequence: The dark histogram has no separation from the light, showing no adversarial interference

Figure 4.2: RTT Histograms for one sample run

4.2 Co-residency Detection

Recall that the job of the FLOODER was to generate interference in the network interface (NIC) of the physical machine by flooding it with large amounts of data. In this experiment, we fixed the value of the period p to 5 seconds. The setup was on the *Local Separate* setup. The result of this can be seen in Figure 4.1a. Every time the flooder ran, the measured RTT from the SERVER increased as expected. This sequence continued for a minute, after which it signaled the CLIENT and went into the *Off* sequence, where it continued the above cycle without actually sending any data. During this period, the RTTs show little to no pattern as seen in Figure 4.1b.

During the co-residency detection, we compared the histograms of the two sequences. Figure 4.2a shows a clear separation in the mean RTTs. The histogram during the control sequence, as seen in Figure 4.2b, has nearly no separation and shows that the adversary is not co-resident. For each setup, we calculated the $aMSE_0$ threshold value from the ROC curves using Equation 3.3, seen in Figure 4.3.

4.2.1 Local and Local Separated

We ran this experiment 20 times on each setup and created the confusion matrices for all setups and orchestrators. This section presents and compares the success rates for detection on all of these.

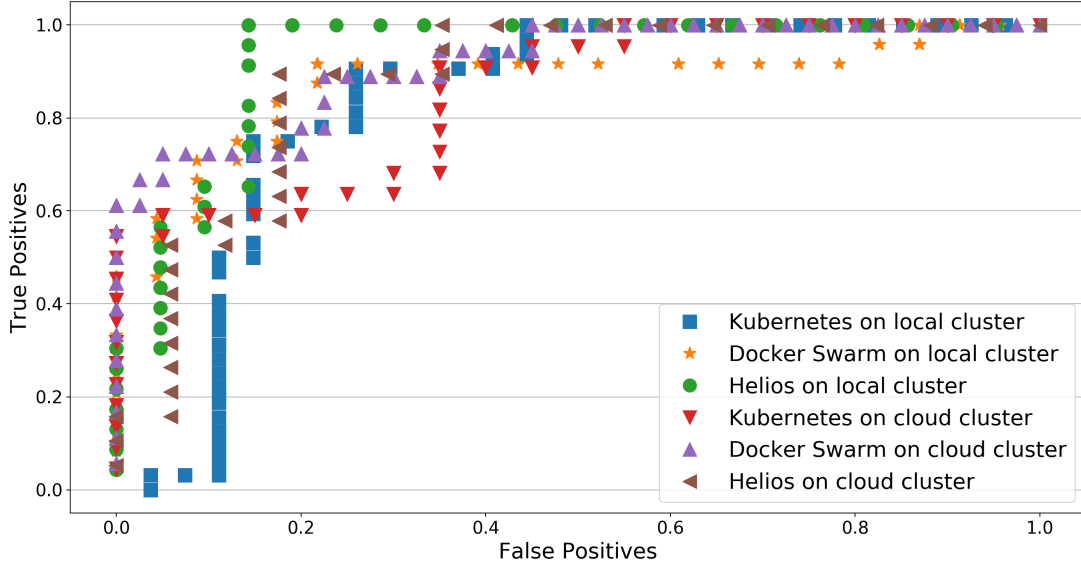


Figure 4.3: ROC curves for controlled experiment across different system configurations

		Configuration							
		<i>Local</i>		<i>Local Separated</i>		<i>Quiet Cloud</i>		<i>Noisy Cloud</i>	
		TP	FP	TP	FP	TP	FP	TP	FP
Orchestrator	Kubernetes	100	38	94	15	91	35	35	36
	Docker Swarm	100	22	92	9	72	5	58	29
	Helios	100	21	100	10	88	18	45	47

Table 4.2: Successful detection rates across orchestrators and configurations

Initially, the local setup was made to be extremely simple. The orchestrator’s coordinator shared the physical node with the application containers, the effect of which is described in much more detail in section 4.3. Confusion matrices for this run are presented in Table 4.2. The setup here is fairly unrealistic as the orchestrator’s coordinator is generally managed by the cloud provider. As such, the interference generated by the coordinator seemed to be reacting to external traffic in the system. As a result, when the FLOODER was active during *On-WM*, the RTTs took much longer, generating very high success values. While this shows that the attack is possible, the results seem unrealistically high.

Once we separated the coordinators onto different machines for the realistic setup, the noise in the system became extremely low, making it possible to get

realistic success rates for detection. Table 4.2 describes the confusion matrices generated for each of the orchestrators.

Due to the Master containers being co-resident, we saw that the ideal MSE threshold value determined for this setup was able to get great separation between true positives and true negatives. As a result, the prediction rate here was extremely high. While running the containers co-resident, the response to the noise generated by the Master was extremely high on Helios and Docker, creating very high MSE values for most runs. The same effect was not observed when the Master containers were moved to an independent node. The MSE values for co-residency detection were in a much tighter range, leading to a higher number of false positives and lower success rates. However, even in this state the success rates were reasonably high - close to 95% for all orchestrators. This detection was carried out with an $aMSE_0$ value of 25.

Summary: Co-residency detection is possible with high accuracy and low false-positives on containers within virtual machines, even when deployed within modern container orchestration systems.

4.2.2 Cloud Environment

The first experiment ran on a semi-quiescent network (*Quiet Cloud*) where the only other traffic belonged to some users in the cluster, however, no VM or user shared the CPU and RAM allocated for this experiment. The success rates for this run are presented in the confusion matrices in Table 4.2. Success rates on all three orchestrators on this cloud-like environment are high and show that co-residency attacks are significant.

Noisy Cloud: The second run included the YCSB benchmark running in the `SERVER` container in the background. This generates load that makes the system non-quiescent. The results of this run are described in the confusion matrices in Table 4.2.

The addition of the noise on this setup generates an expected drop in the success rates. Because this drop varies, it shows that the orchestrators are variably sensitive to noise in the environment.

Finally, we also investigated the reasons behind the failures during detection. Some platforms seemed to have a much higher failure rate than others. In general,

we found that during the cases when the detector failed, the FLOODER was throttled by the KVM. In such cases, the CPU time that the FLOODER VM ran for was about 20x higher than the SERVER VM.

In our initial solution to this problem on the local cluster, we added a cooldown to the FLOODER so that it may gain back scheduling priority, but since the SERVER was also inactive during those times, it had little effect. This was not an entirely realistic model of a server, since a real server would not stop computation when the FLOODER was idle. So we finally chose to add more CPU usage to the SERVER’s VM by making it read from `/dev/urandom`. As a result, the FLOODER VM’s priority with the scheduler increased, giving much better results.


Summary: Co-residency detection within virtual machines is possible on cloud-like environments. Background noise in the system has a significant effect on the success of co-residency detection.

4.2.3 Non-quiet setup


To verify that IO and network interference from virtualization or the master plays a role, we also ran the setup on the *Noisy cloud* while a separate application issued a lot of IO and network calls on the SERVER container. We used the Memcached benchmark tool from the YCSB benchmarking tool [33] for this purpose, which added a high level of load on the system. The result here seems to build on the above postulation. The additional network and disk usage make Kubernetes and Docker Swarm react more than Helios. Kubernetes seems especially sensitive to such load in the system. The success values shown in Table 4.2 are very low. This shows that the attack is fairly sensitive to system noise.

4.2.3.1 Realistic noise setup

The YCSB benchmark generates very high noise in the NIC, since it makes requests as fast as possible. The successful detections in this scenario are low since the amount of traffic generated in the NIC is over the capacity of the NIC, leading to thrashing. We confirm that this is the case by running the same experiment with a recent (April 2019) traffic trace from the MAWI dataset [34]. The results of this experiment are displayed in Figure 4.4a. When the overall load l in the system is within the



(a) Success rates with variable background noise: Detection has high success rates while system load l is within hardware network capacity



(b) Success rates at varying rates of flooding f

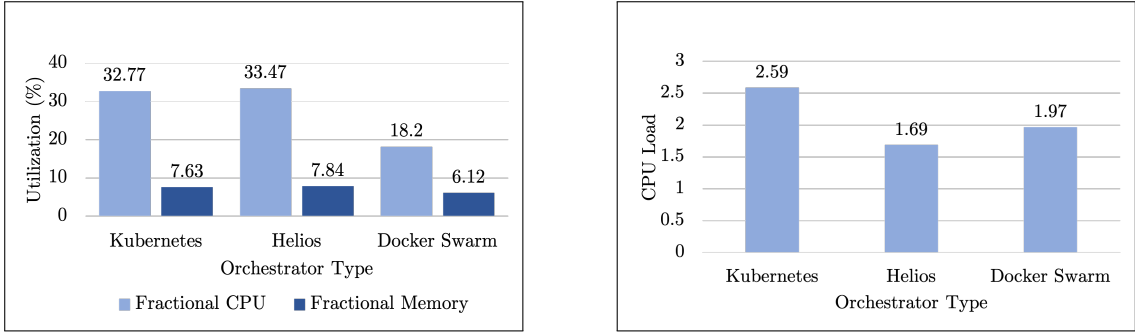
Figure 4.4: Success rates by varying experimental variables: l and f

capacity of the NIC, we find that detection is quite successful. As l passes this value, we find that successful detection gets much harder. However, we find that when the background noise is generated by one regular server, detection is quite successful.

Summary: Background noise in the system significantly affects the success of the attack. However, this is only the case when l exceeds the capacity provided by the hardware. Any background noise under that limit is tolerated and does not reduce the detection success rates significantly.

4.2.4 Adversarial deception

On the *Quiet Cloud*, we controlled the amount of interference generated by the adversary for a number of runs of the attack and evaluated the number of times the attack succeeded for each. Figure 4.4b shows the successes the adversary obtained when flooding at varying rates. This result presents an interesting finding to us. As flooding reaches 1 Gbps (network capacity), the success rates seem to flatten out. As the FLOODER's activity approaches network capacity, the NIC starts to thrash, making it difficult for the system to keep up. At this point, the adversary generates the highest contention possible. This result concurs with the previous section: as the background noise generates enough contention to generate thrashing in the NIC, the success of the attack reduces. However, if the adversary can prevent



(a) Orchestrators have varying CPU usage yielding varying resiliency due to hypervisor scheduling

(b) CPU Load: A lower load shows that the orchestrator performs more CPU bound computation than IO

Figure 4.5: System performance variation based on orchestrator

the system to reach such a state, the amount of activity needed by the adversary can be quite reasonable. Based on this result, we can conclude that the adversary can have significant success at detection while only using around 70% of network capacity for the node in a quiescent environment.

Summary: An adversary can be fairly deception when generating contention in the shared resource, requiring close to 70% of network capacity for high successes on our particular setup.

4.3 Orchestrator Comparison

4.3.1 *Local*

When running the same experiment on all three orchestration systems, we see that the effect of the FLOODER is very clearly visible in Figure 4.6. Inspecting the RTT graph shows a clear square wave that coincides with the FLOODER’s interference. We can conclude then that orchestrators do not prevent such attacks.

However, the successes for each orchestrator seem to differ on each of the systems. To understand why this is the case, we compared the system performance while running the applications containers with the Master on the same node (*Local* setup). We found that all systems had differing CPU and memory usage on average. Figure 4.5a shows a comparison for each of these metrics. The architecture of these systems provides a better understanding for why this is the case.


Kubernetes and Docker Swarm both use a centralized database similar to Etcd to store the state of the cluster. All possible Master nodes maintain a consistent state of the cluster in the Etcd instances they run. Each Etcd instance communicates with others to maintain this consistency. Whenever the Master realizes that the state of the cluster is less than ideal, for example, when a new application is deployed on the cluster, an application crashes, etc., the Master node picks a client node to run the application container on and enforces this by making an HTTP request to the node’s manager. Additionally, in order to understand the state of the cluster, the Master node sends a heartbeat, health-check request to all client nodes (including itself). The Master itself also runs a number of virtualization for IO, DNS, etc. While Docker Swarm and Kubernetes are similar in architecture, their differences arise from the actual implementation of the softwares used. By exploring the source code for Kubernetes [36], Helios [18], and Docker Swarm [17], we were able to draw the following conclusions: Docker Swarm has a much simpler connection to talk to the Docker runtime while Kubernetes implements many more abstractions to ensure container-runtime agnosticity. As a result, Docker Swarm has much lower CPU, memory, and IO usage than Kubernetes.

Helios on the other hand uses a completely different architecture. It maintains an Apache Zookeeper cluster instance that both the Master and client nodes connect to. The Master stores whatever changes are necessary to the cluster and the clients continuously poll Zookeeper on a KeepAlive TCP connection. Whenever they notice a change, they implement these into their state and communicate it back to Zookeeper. Due to this reason, Helios uses more CPU and memory while it queries the Zookeeper instances for data, but spends less time waiting on IO. Thus, the CPU load when Helios runs is much lower than its competitors.


Summary: The inclusion of orchestrators does not have any major impact on the success of co-residency detection.

4.3.2 *Local Separated*


By separating out the Master and the client nodes, we were able to isolate the interference generated in the system by the Master node. While the data was useful, it did not represent a realistic setup. In such a setup, we can now compare the



(a) RTT trace on Kubernetes



(b) RTT trace on Helios



(c) RTT trace on Docker Swarm

Figure 4.6: RTT trace for a flooding sequence across Kubernetes, Helios, and Docker Swarm: Square wave is visible on all three platforms

realistic effect an orchestrator will have on the success of co-residency detection. Any effect by the orchestrator will exist whether or not the attack is currently in progress. Hence, we used the RTTs generated during the *Off* sequence for each orchestrator as well as a baseline measurement where no orchestrator was present. We compare the average RTT for each run on the orchestrator and compare the average for each run during the baseline. We can compare the averages using a simple statistical *T*-test. The *T*-statistic and *p*-value values for each are presented in Table 4.3. For each orchestrator, we reject the null hypothesis suggesting that the orchestrator does have some effect on the RTTs. This means that the interference generated by the orchestrator will raise the baseline RTTs, possibly generating

higher number of false positives. However, all orchestrators only raise the baseline by at most $0.55ms$. Compared to the RTT increase generated by the FLOODER at maximum flooding is around 4 ms, the increase by the orchestrator is quite minimal. Finally, it is interesting to note that rise in baseline is lower in Helios than the other orchestrators, which, due to the poll-based implementation, sees comparatively more CPU than IO usage.

	Kubernetes	Helios	Docker Swarm
P-value	0.00013	$3.9 * 10^9$	$5.29 * 10^{-6}$
T statistic	4.198	7.922	5.395
Mean RTT increase (ms)	0.51	0.36	0.55

Table 4.3: T-test statistic values for each orchestrator

Summary: The orchestrators have some effect on the success of co-residency detection, but the effect is minimal and not enough to be claimed secure against co-residency attacks.

From our experiments, we highlight a few key takeaways:

- Co-residency detection attacks are feasible on containers running within virtual machines. Additionally, orchestration platforms play a minimal role in the success of attacks. These results were verified with a simple setup as well as an architecture that mimics a real cloud environment.
- Adding load to the attack surface reduces the success of co-residency attacks, as one can expect. However, this is only the case when the load exceeds the hardware capacity of the system.
- The attacker can be fairly deceptive in their attack and only needs about 70% of the bandwidth every five seconds to detect co-residency.

Discussion and Future Work

5.1 Effect of noise

While we found that added noise to the system reduces the success of the attacks, that is not necessarily the end of the road for co-residency detection. The detection success rates we generated were for one-shot detection: this means that the adversary runs the attack once with duration d seconds. Based on the data obtained in that run, the adversary detects whether or not they are co-resident. However, modifying the detection mechanism can help improve the accuracy of detection. If the adversary were to increase d , the background noise would be more likely to normalize and improve detection. Alternately, if the adversary used a combination of runs to detect co-residency, they could define a successful detection when more than 50% of the runs return positive results. This could be a significant boost in the detection success.

5.2 Why does this attack generalize?

Singh and Somani [37], Bazm et al. [13], and Zhang et al.[26] performed comprehensive studies on the various types of side-channel attacks. They demonstrated that VMs and containers are both individually vulnerable to the entire class of co-residency attacks. These include last-level cache attacks [38], network topography and instance placement [25], memory bus attacks [39], and network interface cards [15]. Each of these attacks follows a very specific pattern:

- The attack surface is some shared hardware resource that may or may not be virtualized. Despite the status of virtualization, there is no isolation guarantee provided by the software for such shared hardware.

For example, in case of the attack on the shared cache during encryption/decryption, the attack relies on access to shared memory pages in the encryption libraries.

During the Ristenpart et al. [23] placement attack as well as the knowledge of existence [15] attack used in this paper, the NIC is the shared hardware resource that comes under contention. Despite the kernel scheduler’s best efforts, when the attacker demands access to the NIC for transmitting data, it takes away bandwidth from the victim, delaying its transmissions.

- Once the attacker has some form of shared hardware, the attacks attempt to generate some sort of contention while accessing it. As a result, the original execution of the victim takes a latency hit, which the adversary can measure and glean information from.
- Since the software that enables and performs this access does not play a significant role in isolating access to the hardware, each of these attacks is made possible.

The addition of another abstraction layer (orchestrator) to the applications has no major impact on how the hardware is isolated between tenants. Since noise in the system seems to reduce detection success, this may be implemented as a defense. The results in this paper show that while it is possible for the orchestrators to add noise to the system during the attack and make it slightly less effective, they don’t have any effect on whether the hardware is completely isolated to prevent co-residency attacks altogether.

5.3 Defenses

While the attack analysis here is a good representation of co-residency attacks on containers, there are some caveats when the correlations may not apply. Since this attacks depends on water-marking network traffic, it is necessary that the CLIENT

be able to communicate with the same server constantly. In cases when the web-server sits behind a load balancer, the `CLIENT`'s successive requests may get routed to other server instances and the attack may be unsuccessful.

For other attacks, specific defenses should also work. Blocking access to the `clflush` instruction for VMs seems to solve last-level cache attacks [40]. Duplicating pages instead of sharing them across VMs will be an effective defense against cache based side-channel attacks [41]. Zhang et al. [26] suggest periodically flushing cache lines to prevent all manners of `FLUSH+RELOAD` attack variants. Finally, adding noise to the specific attack surface may be a possible defense. However, since the success of the attack is only affected when hardware capacity is reached, such a defense is not encouraged since it will also adversely affect the performance of client applications as well.

To implement defenses that encompass a wider range of attacks, moving target defenses [42, 11, 43] have been shown to be successful in mitigating co-residency attacks. Since orchestrators often manage a large number of nodes for a cluster, performing container migration on these nodes will be easier than migrating the entire VM as suggested by Atya et al. [11]. There are multiple benefits to implementing migrations of applications at the container level. The migration cost for a virtual machine is very high, since a new operating system instance needs to be instantiated at each migration. As a result, migrations cannot be made frequently. Additionally, migrating VMs generate enough detectable traffic for an adversary to potentially follow the target to the new location [44]. Due to the lightweight nature of containers, if the persistent data used by an application is handled properly, there may potentially be no trace of a migrated container. The orchestration platform may also implement a hybrid migration strategy, where VMs are also migrated, but less frequently.

Since the orchestration components are closely linked to the running kernel, we also suggest implementing migration triggered through VM-monitoring as described by [45, 46]. Such a solution would further reduce the cost incurred by the migration while providing stronger security against co-residency attacks.

Chapter 6

Conclusion

In this paper, we showed that containers running on virtual machines are susceptible to co-residency attacks. We analyzed the attack on systems varying in architecture, load, and orchestrators and showed that this plays a minimal role in how they affect resiliency to co-residency based attacks. Our analysis showed that while orchestrators may add some amounts of noise to the system, there is no evidence that enough interference is provided by the orchestrators to make the system significantly and provably secure to network-based co-residency detection attacks. Any differential addition to the resiliency of the system comes at a performance cost that may not be worth paying. Cloud customers that have secure computing requirements should not depend on the orchestration platforms to provide enough protection from co-residency attacks.

Bibliography

- [1] COLUMBUS, L. (2016), “Roundup Of Cloud Computing Forecasts And Market Estimates, 2016 - Forbes,” .
URL <http://www.forbes.com/sites/louiscolombus/2016/03/13/roundup-of-cloud-computing-forecasts-and-market-estimates-2016/>
- [2] MARKUSON, D. (2019), “Why the NordVPN network is safe after a third-party provider breach,” .
URL <https://nordvpn.com/blog/official-response-datacenter-breach/>
- [3] (2018), “Lessons from the Cryptojacking Attack at Tesla,” .
URL <https://redlock.io/blog/cryptojacking-tesla>
- [4] LEGEZO, D. (2018), “LuckyMouse hits national data center to organize country-level waterholing campaign,” .
URL <https://securelist.com/luckymouse-hits-national-data-center/86083/>
- [5] MORTENOIR1 (2018), “VirtualBox Zero-Day exploit,” .
URL https://github.com/MorteNoir1/virtualbox_e1000_0day
- [6] “GKE Sandbox — Kubernetes Engine — Google Cloud,” .
URL <https://cloud.google.com/kubernetes-engine/sandbox/>
- [7] EDER, M. (2016) “Hypervisor-vs. Container-based Virtualization,” *Future Internet (FI) and Innovative Internet Technologies and Mobile Communications (IITM)*, 1.
- [8] (2019), “Cluster multi-tenancy — Kubernetes Engine Documentation — Google Cloud,” .
URL <https://cloud.google.com/kubernetes-engine/docs/concepts/multitenancy-overview>

- [9] TRUYEN, E., D. VAN LANDUYT, V. RENIERS, A. RAFIQUE, B. LAGASSE, and W. JOOSEN (2016) “Towards a container-based architecture for multi-tenant SaaS applications,” in *ARM 2016 - 15th Workshop on Adaptive and Reflective Middleware, colocated with ACM/IFIP/USENIX Middleware 2016*, pp. 1–6.
- [10] “About Kata Containers: Kata Containers,” .
URL <https://katacontainers.io/>
- [11] ATYA, A. O. F., Z. QIAN, S. V. KRISHNAMURTHY, T. LA PORTA, P. MCDANIEL, and L. M. MARVEL (2019) “Catch Me if You Can: A Closer Look at Malicious Co-Residency on the Cloud,” *IEEE/ACM Transactions on Networking*, **27**(2), pp. 560–576.
- [12] ATYA, A. O. F., Z. QIAN, S. V. KRISHNAMURTHY, T. L. PORTA, P. MCDANIEL, and L. MARVEL (2017) “Malicious co-residency on the cloud: Attacks and defense,” in *Proceedings - IEEE INFOCOM*.
- [13] BAZM, M. M., M. LACOSTE, M. SÜDHOLT, and J. M. MENAUD (2017) “Side-channels beyond the cloud edge: New isolation threats and solutions,” in *2017 1st Cyber Security in Networking Conference, CSNet 2017*, vol. 2017-Janua, IEEE, pp. 1–8.
URL <http://ieeexplore.ieee.org/document/8241986/>
- [14] GAO, X., Z. GU, M. KAYAALP, D. PENDARAKIS, and H. WANG (2017) “ContainerLeaks: Emerging Security Threats of Information Leakages in Container Clouds,” *Proceedings - 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2017*, pp. 237–248.
URL <http://ieeexplore.ieee.org/document/8023126/>
- [15] BATES, A., B. MOOD, J. PLETCHER, H. PRUSE, M. VALAFAR, and K. BUTLER (2012) “Detecting Co-residency with Active Traffic Analysis Techniques,” in *Proceedings of the 2012 ACM Workshop on Cloud Computing Security Workshop, {CCSW} ’12*, ACM, New York, NY, USA, pp. 1–12.
URL <http://doi.acm.org/10.1145/2381913.2381915>
- [16] VERMA, A., L. PEDROSA, M. KORUPOLU, D. OPPENHEIMER, E. TUNE, and J. WILKES (2015) “Large-scale cluster management at Google with Borg,” in *Proceedings of the 10th European Conference on Computer Systems, EuroSys 2015*, Bordeaux, France.
- [17] DOCKER (2019), “Docker Swarm,” <https://github.com/docker/swarm>.
URL <https://github.com/docker/swarm>
- [18] SPOTIFY (2019), “Helios,” <https://github.com/spotify/helios>.
URL <https://github.com/spotify/helios>

- [19] APACHE (2015), “Apache Mesos,” <https://github.com/apache/mesos>.
URL <https://github.com/apache/mesos>
- [20] DOCUMENTATION, D. (2017), “Overlayfs storage driver docker documentation,” .
- [21] KANG, H., M. LE, and S. TAO (2016) “Container and Microservice Driven Design for Cloud Infrastructure DevOps,” in *2016 IEEE International Conference on Cloud Engineering (IC2E)*, pp. 202–211.
- [22] EDWARDS, S., D. CZARNOTA, and R. TONIC (2019) *Kubernetes Security Whitepaper, Tech. rep.*, Trail of Bits.
URL <https://github.com/trailofbits/audit-kubernetes/blob/master/reports/KubernetesWhitePaper.pdf>
- [23] RISTENPART, T., E. TROMER, H. SHACHAM, and S. SAVAGE (2009) “Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds,” in *Proceedings of the ACM Conference on Computer and Communications Security*, ACM, pp. 199–212.
URL <http://dl.acm.org/citation.cfm?id=1653662.1653687>
- [24] (1987), “EC2: Dedicated Hosts,” .
URL <https://aws.amazon.com/ec2/dedicated-hosts/getting-started/>
- [25] VARADARAJAN, V., Y. ZHANG, T. RISTENPART, and M. SWIFT “A Placement Vulnerability Study in Multi-Tenant Public Clouds,” .
- [26] ZHANG, W., X. JIA, C. WANG, S. ZHANG, Q. HUANG, M. WANG, and P. LIU (2016) “A comprehensive study of co-residence threat in multi-tenant public PaaS clouds,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 9977 LNCS, pp. 361–375.
URL http://link.springer.com/10.1007/978-3-319-50011-9_28
- [27] XU, Z., H. WANG, and Z. WU (2015) “A Measurement Study on Co-residence Threat inside the Cloud,” *24th USENIX Security Symposium (USENIX Security 15)*, pp. 929–944.
URL <https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/xu>
- [28] HORNBY, T. (2016) “Side-Channel Attacks on Everyday Applications: Distinguishing Inputs with FLUSH+RELOAD,” in *BlackHat USA 2016*.
- [29] ZHANG, Y., A. JUELS, M. K. REITER, and T. RISTENPART (2014) “Cross-tenant side-channel attacks in PaaS clouds,” in *Proceedings of the ACM*

- Conference on Computer and Communications Security*, pp. 990–1003.
URL <http://dl.acm.org/citation.cfm?doid=2660267.2660356>
- [30] HONIG, A. and N. PORTER (2017), “Google Cloud Platform Blog,” .
URL <https://cloud.google.com/blog/products/gcp/7-ways-we-harden-our-kvm-hypervisor-at-google-cloud-security-in-plaintext>
- [31] ALLCLAIR, T. and M. KACZOROWSKI (2018), “Google Cloud Platform Blog,” .
URL <https://cloud.google.com/blog/products/gcp/exploring-container-security-isolation-at-different-layers-of-the-kubernetes-stack>
- [32] “Amazon EKS - Managed Kubernetes Service,” .
URL <https://aws.amazon.com/eks/>
- [33] COOPER, B. F. (2010), “Yahoo ! Cloud Serving Benchmark,” .
- [34] SONY, C. and K. CHO (2000) “Traffic data repository at the WIDE project,” in *Proceedings of USENIX 2000 Annual Technical Conference: FREENIX Track*, pp. 263–270.
- [35] “Amazon EC2 FAQs - Amazon Web Services,” .
URL <https://aws.amazon.com/ec2/faqs/>
- [36] FOUNDATION, C. N. (2019), “Kubernetes,” .
URL <https://kubernetes.io>
- [37] SINGH, G. K. and G. SOMANI (2018) *Cross-VM Attacks: Attack Taxonomy, Defense Mechanisms, and New Directions*, Springer International Publishing, Cham, pp. 257–286.
URL https://doi.org/10.1007/978-3-319-97643-3_8
- [38] GRUSS, D., C. MAURICE, and K. WAGNER (2015) “Flush + Flush : A Stealthier Last-Level Cache Attack,” in *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*, pp. 279–299.
- [39] WU, Z., Z. XU, and H. WANG (2015) “Whispers in the hyper-space: High-bandwidth and reliable covert channel attacks inside the cloud,” *IEEE/ACM Transactions on Networking*, **23**(2), pp. 603–615.
- [40] YAROM, Y. and K. FALKNER (2014) “Flush + Reload : a High Resolution , Low Noise , L3 Cache Side-Channel Attack,” in *USENIX Security 2014*, vol. 1, pp. 1–14.
URL <http://eprint.iacr.org/>

- [41] IRAZOQUI, G., M. S. INCI, T. EISENBARTH, and B. SUNAR (2014) “Wait a minute! A fast, Cross-VM attack on AES,” in *International Workshop on Recent Advances in Intrusion Detection*, Springer, pp. 299–319.
- [42] AZAB, M. and M. ELTOWEISSY (2016) “MIGRATE: Towards a Lightweight Moving-Target Defense Against Cloud Side-Channels,” *Proceedings - 2016 IEEE Symposium on Security and Privacy Workshops, SPW 2016*, pp. 96–103.
- [43] MOON, S.-J., V. SEKAR, and M. K. REITER (2015) “Nomad: Mitigating arbitrary cloud side channels via provider-assisted migration,” in *Proceedings of the 22nd acm sigsac conference on computer and communications security*, ACM, pp. 1595–1606.
- [44] ACHLEITNER, S., T. L. PORTA, P. MCDANIEL, S. V. KRISHNAMURTHY, A. POYLISHER, and C. SERBAN (2017) “Stealth migration: Hiding virtual machines on the network,” in *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*, pp. 1–9.
- [45] MISHRA, P., E. S. PILLI, V. VARADHARAJAN, and U. TUPAKULA (2017) “Out-vm monitoring for malicious network packet detection in cloud,” in *2017 ISEA Asia Security and Privacy (ISEASP)*, IEEE, pp. 1–10.
- [46] WAHAB, O. A., J. BENTAHAR, H. OTROK, and A. MOURAD (2017) “Optimal load distribution for the detection of VM-based DDoS attacks in the cloud,” *IEEE Transactions on Services Computing*.