

A Security Perspective on Unikernels

Joshua Talbot*, Przemek Pikula*, Craig Sweetmore*, Samuel Rowe*,
Hanan Hindy*, Christos Tachtatzis†, Robert Atkinson† and Xavier Bellekens*†

*Division of Cyber-Security, Abertay University, Dundee, Scotland

†EEE Department, University of Strathclyde, Glasgow, Scotland

Abstract—Cloud-based infrastructures have grown in popularity over the last decade leveraging virtualisation, server, storage, compute power and networking components to develop flexible applications. The requirements for instantaneous deployment and reduced costs have led the shift from virtual machine deployment to containerisation, increasing the overall flexibility of applications and increasing performances. However, containers require a fully fleshed operating system to execute, increasing the attack surface of an application. Unikernels, on the other hand, provide a lightweight memory footprint, ease of application packaging and reduced start-up times. Moreover, Unikernels reduce the attack surface due to the self-contained environment only enabling low-level features. In this work, we provide an exhaustive description of the unikernel ecosystem; we demonstrate unikernel vulnerabilities and further discuss the security implications of Unikernel-enabled environments through different use-cases.

Index Terms—Unikernel, Docker, Container, Security

I. INTRODUCTION

Cloud computing is comprised of various virtualisation architectural models enabling users to build heterogeneous services comprised of multiple resources such as network devices, software components, serverless components, and containers. However, a new paradigm focusing on transient microservices based on Unikernels has emerged and is becoming progressively popular. Unikernel’s on-demand properties, low running costs and elasticity make it a perfect candidate for transient services. With the rise of multi-tenancy, multi-cloud infrastructure and the heterogeneity of the services proposed, the complexity of the ecosystem is constantly increasing, leading to a tremendous attack surface to cover and protect. In addition, the supply chain is often composed of various third party libraries, Operating Systems (OS) and re-implemented operating system functions or legacy code enabled, to ensure retro-compatibility between systems. Existing work on Unikernels focuses mainly on its applications across a broad range of technologies, as well as, its integration with the host platform. However, while Unikernels claim a reduced attack surface, to the best knowledge of the authors, the security of Unikernels and their attack surface has not yet been explored in depth. In this manuscript, we review and explore Unikernel security ecosystems as the rise of transient microservices will make Unikernels prevalent in cloud infrastructure. The remainder of this manuscript is structured as follows;

II. VIRTUALISATION, CONTAINERS AND UNIKERNELS

A. Virtualisation

Virtualisation is the process of emulating a system, or multiple systems, using the resources of a host machine [1]. This can be used to re-create networks, or completely isolated machines, increasing the versatility of hardware. Virtualisation improves security through isolation as individual virtual machines cannot communicate with each other without explicitly specifying a connection. This isolation means that if the virtual machine administrator account is compromised, the attacker will not be able to access the host or other virtual machines running on it. This is facilitated by an additional level of user-privilege on the host that controls the guest [2].

1) *Virtualisation Types*: There are a variety of virtualisation types, each with their own advantages and disadvantages.

a) *Full Virtualisation*: This type of virtualisation virtualises the hardware the guest machines runs on. This can either be hardware-assisted, with the hardware itself supporting the virtualisation or software-assisted, where the operating system interfaces with the hardware [3]. The former type of virtualisation’s main appeal is its ability to emulate hardware, allowing for consistent performance, improved reliability, and isolation. If an attacker gained control over the machine he would have no knowledge of the real hardware of running on the host. While the malicious user might not be able to interact with the host environment, he might, however, be able to discover that he is interacting with a virtualised environment [4]. This, in turn, will allow the attacker to fine-tune his attacks to target the virtual machine itself, reducing its security.

b) *Para-Virtualisation*: Though largely antiquated (having support removed from the Linux kernel in 2009 [5]), this type of virtualisation ensures that the virtual machine interacts with a software interface instead of the hardware directly. This allows the virtual machine to use Application Programming Interfaces (APIs) to make system calls that would otherwise be hard to virtualise. This improves the overall efficiency since, the most complex system calls are abstracted through the API [6]. However, with the improving efficiency of hardware virtualisation, para-virtualisation no longer provides tangible performance benefits [5].

c) *OS-Virtualisation*: OS-Virtualisation is where a single kernel can run multiple occurrences of the operating system as containers, each of which acting as an isolated machine. These containers place less emphasis on recreating an entire machine, but rather focus on the user space, allowing users

to run multiple operating systems and associated software on a single machine for convenience. The containers do not have access to the hardware of the physical machine, and will typically use the same OS as the host, which in turn can limit the application of the machine [7]. This also implies that if the kernel is ever compromised, all associated containers will be compromised [8].

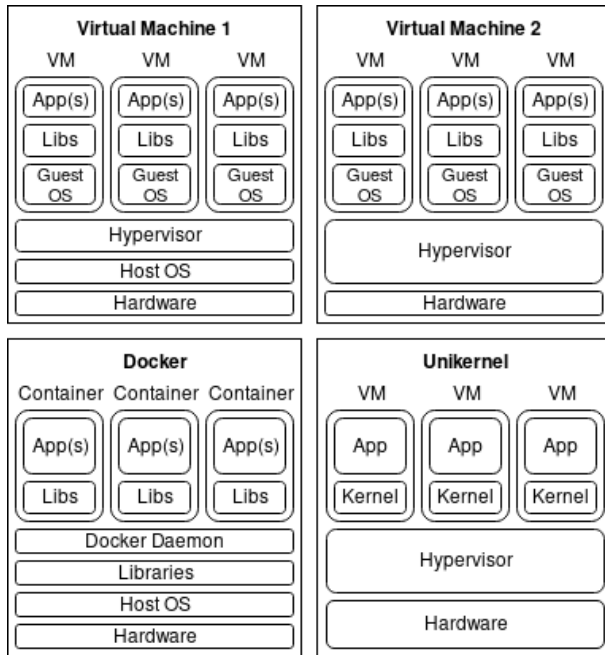


Fig. 1: Virtualisation Types

B. Docker

Docker is software enabling OS-Virtualisation through the use of the ‘docker engine’ which manages the containers on the host. Docker, like most OS-Virtualisation, opts for process-level isolation over full isolation. Whilst this makes it more efficient for running isolated applications, Docker systems demonstrate numerous vulnerabilities, as highlighted in [9].

C. Unikernel

1) *Unikernel Types*: There are two major kinds of Unikernels, whose security profiles differ slightly: Clean Slate and Legacy [10].

a) *Clean Slate*: Clean slate Unikernels do not try to emulate classical OS in any regards. They are written in a single programming language and provide interfaces for external communications (i.e. networking) in the same language. Examples include MirageOS (OCaml), IncludeOS (C++), HalVM (Haskell), LING (Erlang) and runtime.js (JavaScript). Clean slate Unikernels allow language specific virtual machines, like the Java Virtual Machine (JVM), to function as actual virtual machines. C libraries, while present, are transparent to userspace code. The most straightforward implementation can

be found in runtime.js, which wraps Chromium’s V8 Javascript engine inside a lightweight kernel [11].

b) *Legacy*: Legacy Unikernels, on the other hand, implement a subset of POSIX to ensure unmodified software to run, while some only require minor configuration changes. They don’t support timesharing (the ability to simultaneously run multiple independent programs), instead, they delegate this role to the virtualisation layer. Unikernels such as OSv and Graphene focus on ensuring Linux compatibility and software interoperability, re-implementing system call interfaces, while the Rumprun unikernel implements a subset of FreeBSD’s syscalls [12] [13].

Figure 1 summarises the primary differences between the different types of virtualisation including unikernels. Figure 1 (top-left) shows the layout of software assisted virtualisation, and OS-virtualisation. The virtualisation is run on top of the OS by a hypervisor, which can allow for the creation of virtual hardware. Figure 1 (top-right) shows hardware-assisted virtualisation, and para-virtualisation. The hardware itself (sometimes assisted by the hypervisor) runs the virtual machines itself, while Figure 1 (bottom-left) demonstrates how containers are implemented and finally Figure 1 (bottom-right) provides an overview of Unikernels virtualisation architecture. As demonstrated, Unikernels do not require an operating system to function correctly.

2) *Isolation*: Software running on a Unikernel is less isolated from the hypervisor than software running on a virtual machine, but more isolated than software running in a container. While Unikernels rely on the hypervisor for their isolation, they also bring their own kernel with a reduced attack surface. This bespoke kernel makes Unikernels more isolated than containers. However, similar to any software, it is up to the Unikernel developers to supply Unikernels with intrinsic security.

D. Ecosystem

Figure 2 shows how the development and production environments differ, along with the Unikernel’s distribution. In development, an OS would be required in setting up the Unikernel and its features, while in production it is optimal to deploy the Unikernel without a host OS for increased optimisation and security. This can be managed through an online repository to host and modify the Unikernel image. Moreover, the environment/functionalities vary depending on the Unikernel’s distribution, which has been generalised in the example above. Features include a package manager (e.g. Conan) for building and downloading Unikernel compatible applications, a toolstack/domain manager (e.g. Cosmos) to manage unprivileged domains, and an API (e.g. Rest) to manage Unikernels remotely through issuing commands.

E. Unikernel Usage

Unikernels are primarily used in cloud computing, however, they also show potential in Internet of Things (IoT) and networking devices. These platforms benefit from Unikernels flexible, lightweight and scalable framework. For cloud

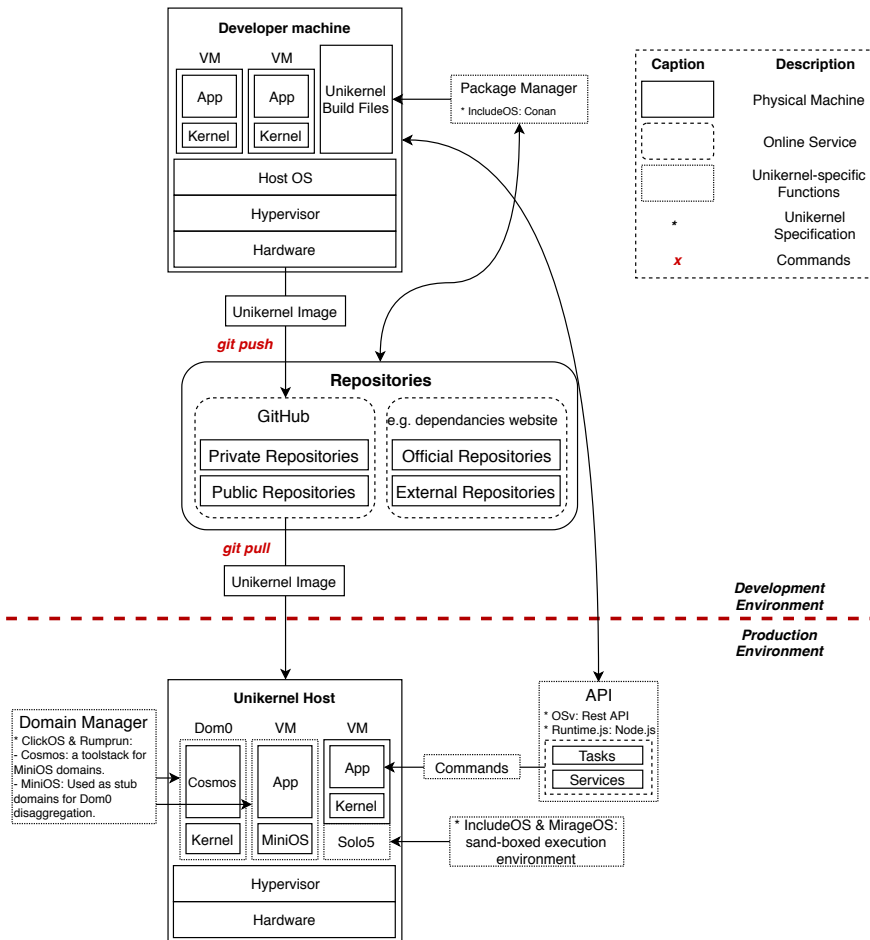


Fig. 2: Unikernel Ecosystem

computing and networking devices, Unikernels prove more effective at utilising the available hardware, allowing for increased scalability in a more isolated, heterogeneous computing environment with a more optimised code base. Due to their efficiency, Unikernels also come with the benefit of much faster boot times, allowing for downed services or network nodes to be quickly restored with minimal overhead. The most compelling reason to use Unikernels is their potential for security which takes advantage of their reduced attack surface, isolation and, depending on the distribution, a robust set of security features [14]. IoT devices also largely benefit from Unikernels lightweight code-base and scalability, allowing for a more complex and diverse set of services to be deployed on-demand with very low overhead, despite the hardware limitations.

III. UNIKERNELS SECURITY OVERVIEW

By limiting the code base of deployed applications, Unikernels inherently have a small and unique attack surface making them relatively secure. This is further achieved in some im-

plementations by evaluating and modifying the implemented code [15], where developers are able to focus on hardening security to effectively mitigate existing attack vectors.

1) *Shell*: Numerous Unikernels do not implement a shell natively, making most types of payloads, that typically rely on bash, ineffective. Further preventing automated attacks or less experienced attackers from successfully exploiting a vulnerability by increasing the complexity of the payloads.

2) *System Calls*: System calls are often removed, or are not supported by Unikernels, hence, malicious users are required to know the exact memory layout in order to invoke a function call such as `open()` or `write()`. The attack surface is further reduced through implementing randomised memory layouts at every build [15].

3) *Hardware Emulation*: By not emulating hardware interfaces such as floppy drives, Peripheral Component Interconnect (PCI) bus or Graphics Processing Unit (GPU), possible breakouts such as the 2015 Venom attack that affected QEMU can be prevented [16]. Unikernels such as IncludeOS and Mirage-OS are also able to mitigate this through running




Fig. 3: Pivoting Attack

on solo5; a sandboxing interface between the Unikernels and the hypervisor. Solo5 allows for a minimal code base, removing code that would otherwise be unnecessary to cloud computing. This is achieved during compilation by determining the required dependencies from the imported libraries, application code and configuration files, further reducing the Virtual Machine Monitor (VMM) overall attack surface.

Many Unikernel implementations rely solely on their reduced attack surface from the previously mentioned features or lack thereof, not taking into consideration known vulnerabilities relating to their existing attack vectors, such as minimal security features in Address Space Layout Randomisation (ASLR) [17]. This has led to criticism and debate discussing their design; The lack of separation between the user and kernel-space comes with great security concerns. For example, a successful buffer overflow attack on the Unikernels limited functionality could give an attacker a foothold into kernel space, making privilege escalation unnecessary, as well as, code execution and pivoting to another target a potential threat. Furthermore, Figure 3) shows an attacker blocked by the target’s firewall can leverage a web-application running on a unikernel to exploit the network. By exploiting the web-application, the Unikernel API or through compromising the Unikernel itself, the attack could use kernel functions to forge malicious packets, for denial of service, or to target other devices on the network enabling him to pivot [17]. Hence, while demonstrating potential for scalability and improved security by reducing their attack surface, some Unikernels are yet to implement key security features.

A. Immutable Infrastructure

In traditional infrastructure, patches can be applied to upgrade packages, change the servers configuration and modify or upload code. This, however, poses the threat of malicious modifications being made by an attacker. Having an immutable infrastructure mitigates this by employing the “destroy and provision” approach of needing to rebuild the Unikernel to make any changes. This not only prevents malicious modifications but also reduces the overall code complexity of having outdated configurations that could lead to new bugs and vulnerabilities.

a) *Para-Virtualisation*: In some Unikernel implementations, para-virtualisation is implemented to restrict privileged operations to the hypervisor through an API. This essentially allows applications to run in ring 3 rather than ring 0, isolating the Unikernels application from the hardware level. Therefore, the hypervisor is able to enforce Write xor Execute (W^X) to its page tables by making executable pages immutable.

b) *Heterogeneous Networking*: Increased heterogeneity can be achieved through either enabling multiple instances of a certain application to run with varying configurations and libraries or by using a separate Unikernel instance for each function in a network. For example, this can be implemented through using a set of Unikernels for hosting databases, each with their own protections and privileges depending on the sensitivity and nature of their data, rather than having a singular database that stores all of this data. Similarly, by allocating specific Unikernels to run the server, web service, etc. . . , the network’s architecture can be isolated to its critical functions, preventing possible knock-on effects of certain attacks, such as DDoS, from compromising the entire network and allowing for the affected elements of the network to be isolated and easily identifiable.

B. Entropy

Unikernels often have a low entropy as the hardware is virtualised, hence, randomly generated values persist across reboots, meaning, that if an attacker were to crash a Unikernel, they may be able to determine randomized values even after rebooting. This lack of entropy can lead to security features such as ASLR, stack cookies, TCP sequence numbers and access tokens, etc. . . becoming ineffective [17]. In Rumprun and MirageOS, this has been effectively mitigated through the implementation of RDRAND, however, since this is a common issue among virtualised platforms, this vulnerability requires to be checked on a platform basis before use [18]. Enforcing entropy persistence amongst Unikernels may affect incohesive generated Unikernels by having duplicate values. To mitigate this, generated seeds should be validated as non-duplicates otherwise, they may be vulnerable to nonce reuse attacks.

```

IncludeOS 0.15.1-5 (x86_64 / 64-bit)
++>> Running [ Hello world - OS included ]
### .TEXT ###
printf:0x23c28c fn1:0x201000 fn2:0x201010 fn3:0x201020
### .DATA ###
str1:0x305620 str2:0x305628 str3:0x305630
### STACK ###
var1:0x1ffcf8 var2:0x1ffcec var3:0x1ffcf0
### HEAP ###
Round 1 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 2 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 3 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 4 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 5 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 6 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 7 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 8 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 9 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 10 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0

IncludeOS 0.15.1-5 (x86_64 / 64-bit)
++>> Running [ Hello world - OS included ]
### .TEXT ###
printf:0x23c28c fn1:0x201000 fn2:0x201010 fn3:0x201020
### .DATA ###
str1:0x305620 str2:0x305628 str3:0x305630
### STACK ###
var1:0x1ffcf8 var2:0x1ffcec var3:0x1ffcf0
### HEAP ###
Round 1 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 2 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 3 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 4 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 5 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 6 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 7 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 8 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 9 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 10 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0

IncludeOS 0.15.1-5 (x86_64 / 64-bit)
++>> Running [ Hello world - OS included ]
### .TEXT ###
printf:0x23c28c fn1:0x201000 fn2:0x201010 fn3:0x201020
### .DATA ###
str1:0x305620 str2:0x305628 str3:0x305630
### STACK ###
var1:0x1ffcf8 var2:0x1ffcec var3:0x1ffcf0
### HEAP ###
Round 1 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 2 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 3 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 4 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 5 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 6 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 7 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 8 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 9 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
Round 10 --> ptr1:0x3d59a0 ptr2:0x3d59c0 ptr3:0x3d5ae0
    
```

Fig. 4: ASLR Test Results

IV. UNIKERNEL WEAK LINKS

A. ASLR Vulnerability

IncludeOS mentions its use of ASLR, however, it has been demonstrated that their implementation is either flawed or a bogus claim [17]. In this section we replicate a successful attack against ASLR and the latest IncludeOS version. The ASLR attack is performed against latest stable release of Ubuntu (18.04.2 LTS) and IncludeOS (v0.15.0). The cloud Hello World demo from IncludeOS’s GitHub repository was used as a baseline to host the ASLR testing service. This service prints the memory addresses of strings, the output of functions and variables on the stack, as well as, pointers to newly added blocks of memory onto the heap. To verify the assertion of ASLR being implemented after each build the tests were repeated after deleting and then rebuilding the Unikernels files.

B. Results

Figure 4 illustrates the results showing that ASLR was not implemented in the latest version of IncludeOS, which is evidenced by the unchanging memory addresses of the stored values. Furthermore, Rumpun, IncludeOS and MirageOS versions of Unikernels have been tested and confirmed to not implement ASLR, page protections or stack canaries, and set their memory to RWX and hence are vulnerable against multiple attacks.

C. Unikernel Limitations

a) *Protection Rings*: In traditional operating systems, protection rings are used to set increasing levels of access to the operating system. The issue with Unikernels is that they run their applications in Ring 0 or the “kernel ring”. Specifically, one of the core ideas behind the structure of a Unikernel is that the kernel, operating system and application is contained in a single system [19]. One of the factors of this is that the Unikernel does not have the capability to provide additional protection rings, and so must go without.

b) *Guard Pages*: unmapped pages between memory allocations used to cause segmentation faults are not implemented on these Unikernels either [17].

c) *Debugging Tools*: Unikernel deployment is limited due to the increased difficulty in debugging . This is due to the removal of components of the operating system; standard commands for debugging such as netstat, tcpdump and ping are not present on the Unikernel [20], [21]. Making it difficult for developer maintaining and updating Unikernels as there is no easy maner to determine an issue in the code and would likely be required to perform trial-and-error testing consuming time and resources. An additional issue is that Unikernels cannot be updated while running, and require to be shut down, updated, rebuilt and run again for changes to take effect.

The combination of these factors creates the potential for the exploitation of a buffer overflow vulnerability to directly overwrite the instructions of the program. In turn, this allows for remote code execution, effectively granting the attacker remote access to the system, which will be met with administrator privileges due to the lack of protection rings.

D. Mitigations

- **Entropy**: Enforcing entropy persistence amongst Unikernels may affect groups of frequently generated Unikernels by sharing duplicate values. To mitigate this generated seeds should be validated as non duplicates otherwise they may be vulnerable to nonce reuse attacks.
- **Hardening**: A method of mitigating the chance of the Unikernel being compromised is to implement hardening. Hardening is where efforts are made to restrict what vulnerable parts of the system can be accessed by an attacker. Common techniques for this are consistent patching of the operating system and application, closing of unused ports, enforcement of password complexity and removal of default accounts [22].
 - 1) *Host Hardening*: Host hardening makes it more difficult to exploit the applications running on the host. This can be accomplished through:
 - reducing the attack surface
 - protecting against stack overflows
 - randomising memory layout (ASLR)
 - protecting against buffer overflows
 - encrypting data wherever possible

Unikernels, at least in theory, provide the ultimate attack surface reduction. How much the surface is reduced

varies by Unikernel. However, in general, clean slate Unikernels reduce the attack surface more than legacy Unikernels do. This is because compiling the application and the kernel it will run on into a single program allows the compiler to verify all methods data is passed around. Interfacing directly within a typesafe, high level programming language allows clean slate Unikernels a vastly reduced attack surface compared to passing data through pipes and C library functions. This reduces the risk of buffer overflows.

2) *Library Hardening*: For Unikernels, there are two very different forms of library hardening: - C standard library hardening, which affects legacy Unikernels and, to a lesser degree, clean slate Unikernels, - Native library hardening, which affects clean slate Unikernels. However, unlike traditional OS where C is the native language, a clean slate Unikernel the language the Unikernel was implemented in is the native.

Unikernels use many different C standard libraries. OSv uses musl, IncludeOS uses newlib, rumprun uses libc from NetBSD and UKL uses glibc [23] [17] [24]. NCC testing has revealed that some, like newlib, lack support for the `_FORTIFY_SOURCE` macro that can be used to automatically detect some buffer overflows in common C functions. The lack of `_FORTIFY_SOURCE` forces developers to manually verify bounds checks in all the programs. Other parts of the Standard C library, that are often exploitable and should be hardened include format specifiers.

Native library hardening depends on both how robust the Unikernel's native programming language standard library is, and how securely are the Unikernel specific wrappers over external interfaces implemented.

3) *Networking*: Network hardening usually involves disabling unnecessary services. Unikernels come with few, if any, making them relatively tricky to crack out of the box. However, the network services they come with can be quite vulnerable if left exposed. For example, in OSv the REST API used to control the Unikernel can replace the command line, read and write files and directories. Exposing it to attackers gives them command execution, Local and Remote File Inclusion.

4) *Security Modules and their comparison with Traditional Linux Host Hardening*: There are several potential security modules that could be installed and implemented by an administrator. All provided examples making use of a system for Mandatory Access Controls, where the OS will control the ability for individual users to grant or deny access to resource objects on a file system [25] These are: SELinux (Security-Enhanced Linux), which provides a system for Mandatory Access Controls and defines the access and transition rights of all users, applications, processes and files on the system. [26], AppArmor, which binds the access control attributes to the programs instead of the users, which will either report or outright prevent access from chosen profiles [27]

V. CONCLUSION

Cloud environments are constantly evolving to reduce deployment costs and decrease the complexity of virtualised solutions. Throughout this paper we presented an overview of the strength and weaknesses of Unikernels. We demonstrated that some Unikernels are vulnerable to known attacks such as buffers overflows and did not yet integrated best practices to alleviate common vulnerabilities.

REFERENCES

- [1] Microsoft Azure, "What is a virtual machine?" 2019. [Online]. Available: <https://azure.microsoft.com/en-us/overview/what-is-a-virtual-machine/>
- [2] VMware, "Vmware vcloud suite 6.0 documentation," 2015. [Online]. Available: <https://pubs.vmware.com/vcloudsuite-60/index.jsp?topic=/%2Fcom.vmware.vcloudsuite.doc/%2FGUID-E2D5949F-33A1-49D3-B20F-4ACA9EA14C1.html>
- [3] MicE, "What is the difference between hardware and software virtualization?" March 2010. [Online]. Available: <https://superuser.com/questions/117774/what-is-the-difference-between-hardware-and-software-virtualization>
- [4] P. Ferrie, "Attacks on virtual machine emulators." [Online]. Available: https://www.symantec.com/avcenter/reference/Virtual_Machine_Threats.pdf
- [5] A. Katria, "commit d0153ca35d344d9b640dc305031b0703ba3f30f0," September 2009. [Online]. Available: <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=d0153ca35d344d9b640dc305031b0703ba3f30f0>
- [6] C. Horne, "Vmware paravirtualization," October 2007. [Online]. Available: https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/VMware_paravirtualization.pdf
- [7] R. Lingeswaran, "Para virtualization vs full virtualization vs hardware assisted virtualization," December 2017. [Online]. Available: <https://www.unixarena.com/2017/12/para-virtualization-full-virtualization-hardware-assisted-virtualization.html/>
- [8] V. S. Pethuru Raj, Jeeva S. Chelladurai, "Containerization vs virtualization an introduction to docker," September 2015. [Online]. Available: <https://jaxenter.com/containerization-vs-virtualization-docker-introduction-120562.html>
- [9] J. Sanders, "Docker containers are filled with vulnerabilities: Here's how the top 1,000 fared." [Online]. Available: <https://www.techrepublic.com/article/docker-containers-are-filled-with-vulnerabilities-heres-how-the-top-1000-fared/>
- [10] M. Bright, "Unikernels in action," presented at the DevConf.cz, Brno, January 2018. [Online]. Available: https://mjbright.github.io/Talks/2018-Jan-28_Devconf.cz_Unikernels/2018-Jan-28_Devconf.cz_Unikernels.pdf
- [11] J. O. system written in JavaScript, "Jsos." [Online]. Available: <https://github.com/JsOS-Team/JsOS>
- [12] O. Lab, "Graphene library os." [Online]. Available: <https://github.com/oscarlab/graphene>
- [13] C.-C. Tsai, K. S. Arora, N. Bandi, B. Jain, W. Jannen, J. John, H. A. Kalodner, V. Kulkarni, D. Oliveira, and D. E. Porter, "Cooperation and security isolation of library oses for multi-process applications," in *Proceedings of the Ninth European Conference on Computer Systems*. ACM, 2014, p. 9.
- [14] I. Eyberg, "Assessing unikernel security," May 2019. [Online]. Available: <https://nanovms.com/dev/tutorials/assessing-unikernel-security>
- [15] P. Buer, 2017. [Online]. Available: <http://unikernel.org/blog/2017/unikernels-are-secure>
- [16] M. Corporation, April 2015. [Online]. Available: <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2015-3456>
- [17] S. Michaels and J. Dileo, "Assessing unikernel security," 2019. [Online]. Available: https://www.nccgroup.trust/globalassets/our-research/us/whitepapers/2019/ncc_group-assessing_unikernel_security.pdf
- [18] T. Leonard, "Mirageos security advisory 02," March 2019. [Online]. Available: <https://mirage.io/blog/>
- [19] R. Bias, "Unikernels will create more security problems than they solve," June 2016. [Online]. Available: <https://thenewstack.io/unikernels-will-create-security-problems-solve/>

- [20] B. Cantrill, "Unikernels are unfit for production," January 2016. [Online]. Available: <https://www.joyent.com/blog/unikernels-are-unfit-for-production>
- [21] G. Rushmore, "A discussion of the operational challenges with unikernels," August 2015. [Online]. Available: <https://www.morethanseven.net/2015/08/21/operating-unikernel-challenges/>
- [22] N. Kapoor, "Host hardening - achieve or avoid," 2016. [Online]. Available: https://www.owasp.org/images/8/83/Host_review_owasp_2016.pdf
- [23] V. Dimitri, 2015. [Online]. Available: <https://github.com/cloudius-systems/osv/blob/a3cd022fcd02c88eae89476aa6c29e3c4be04926/libc/manifest.txt>
- [24] R. Ali, November 2018. [Online]. Available: <https://next.redhat.com/2018/11/14/ukl-a-unikernel-based-on-linux/>
- [25] M. Rouse, "mandatory access control (mac)," December 2013. [Online]. Available: <https://searchsecurity.techtarget.com/definition/mandatory-access-control-MAC>
- [26] Red Hat, "49.2 introduction to selinux." [Online]. Available: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/5/html/deployment_guide/ch-selinux
- [27] AppArmor, "sbeattie," April 2019. [Online]. Available: <https://wiki.ubuntu.com/AppArmor>