


MARCH 2020



NUTANIX AHV

Security at the Virtualization Layer

1. INTRODUCTION.....	3
1.1. What is AHV?.....	4
2. SECURITY.....	5
3. HARDENING AHV IN 8 STEPS.....	7
3.1. Security-Enhanced Linux (SELinux).....	7
3.2. Address Space Layout Randomization.....	8
3.3. Exec-Shield.....	8
3.4. Linux Advanced Intrusion Detection Environment (AIDE).....	9
3.5. Host Secure Boot.....	9
3.6. Security Technical Implementation Guides.....	9
3.7. Security Configuration, Management and Automation.....	10
3.8. The Security Development Lifecycle (SecDL).....	10
4. CONCLUSION.....	
Nutanix AHV.....	11
LIST OF FIGURES.....	
Figure 1: The kernel Process.....	8



1. Introduction

Virtualization has changed the way we consider security in the data center. Historically the approach may have been what we call platform-centric, in that the platform(s) used in the data-center dictated the security process. Limitations of the hardware such as compatibility with security devices and throughput and core functionality had to be understood in order to architect a solution that worked to protect the environment. Compromise in some areas be it performance or security was always part of the solution and subsequently factored into the organization's risk tolerance.

As applications have become the driving force behind business objectives, security compliance frameworks helped in the process of protecting sensitive data. Some of the techniques used in these frameworks were the restriction of access to this data and isolation/ segmentation of the underlying machines running the workloads. Virtualization introduced new challenges whilst also being a good step forward in securing the workloads. Security now has to be much more application-centric given the freedom with which you can now deploy and move apps.

One of the benefits of virtualization is security; applications running in separate virtual machines are isolated from each other and, ideally, it is very hard for a compromised guest to attack other virtual machines running on the same host. The hypervisor itself is the place where most attacks on a virtualization system will be aimed. To that end, it is essential that insecurities in the hypervisor be addressed swiftly. A good hypervisor should have:

- A small attack surface
- Regular and thorough code audits
- Have Global support through utilization

Given these caveats a Linux kernel with a widely used user space emulator (QEMU) adapted to meet enterprise needs seems to be the most appropriate for a secure hypervisor variant.



1.1 WHAT IS AHV?

AHV is an enterprise-class virtualization solution included with the Nutanix Acropolis Enterprise Cloud OS (AOS), with no additional software components to license, install or manage. AHV as a Type 1 (native) hypervisor, comes preinstalled on Nutanix appliances, is configurable in minutes, and provides all the necessary foundation for your virtual infrastructure. AHV was designed and optimized for the modern Cloud era, built for Hyperconverged Infrastructure (HCI) and on the principles of web-scale engineering with Operational Intelligence, Security, and Automation delivered with 1-click simplicity.

To give you a little history, Nutanix introduced KVM support back in 2013 on AOS 3.5 having already supported ESX, this was a precursor to launching our own native hypervisor and AHV had its formal release with AOS 4.1 in 2015. Under the covers, It is an adaptation of the proven open-source Linux KVM and QEMU. Since release, AHV has become a popular choice for Nutanix customers. AHV adoption has increased year on year and, as of the writing of this paper, we reported a quarterly adoption of more than 47%¹ of nodes sold choosing to deploy AHV over other virtualization solutions.

AHV is more than a branded version of Linux KVM virtualization, Nutanix has invested in creating a virtualization solution that is as easy to manage as our industry-leading HCI offering and provides all the performance and enterprise-grade features needed to run ANY workload.

¹ As of Q2FY20 and based on a trailing four-quarter average see: ir.nutanix.com



2. Security:

In virtual environments, there are unique concerns. These concerns can give pause to the Security Team, whose reticence to these technologies can block the adoption of public and private cloud platforms or at least running disparate workloads on the same shared infrastructure. These concerns can be grouped into two primary categories:

- [VM malware cascading between VMs](#) - i.e. One machine gets compromised or the Hypervisor gets compromised and the entire infrastructure is vulnerable potentially bringing the entire business down.
- [The potential for Guest VMs to escape and control the Hypervisor gaining access to other VMs, and their data.](#)² (Otherwise known as Hyperjacking and VMescape).

These concerns can be mitigated somewhat by implementing a defense-in-depth strategy. One line of that defense is placed upon the vendor and the techniques they use to harden the underlying hypervisor. For an Advanced Persistent Threat (APT) to be adequately deterred in their attempt to compromise virtual infrastructure the following These aren't considerations exclusively for hypervisors, I'm making the point that many considerations for a secure virtual environment have to be made:

- [Network security](#)
As you would in non-virtual environments, make the same considerations for network protection in virtual environments. Isolate the management and hypervisor traffic to its own non-routable VLAN. Ensure public traffic makes no ingress to protected environments.
- [Endpoint security](#)
Monitor for spurious or suspicious activity on all endpoints and inline of traffic flows with IPS/ IDS and/ or packet inspection firewalls.
- [System Hardening](#)
As most attack vectors can be traced to human oversight or failure to properly harden the stack, system security hardening is essential:
 - Patch known CVEs
 - Configure User privileges
 - Remove default accounts
 - Close unused ports and protocols
 - Enforce strong password policies
 - Remove unwanted services

² Successful execution of these hypervisor attacks outside of “proof of concept” has not yet been recorded, this is mainly due to the difficulty of directly accessing hypervisors; however, they are still considered real-world threats.

- [Stringent access management policies](#)

Where possible utilize the capabilities of directory services for Authentication and Authorization to systems. Employ the use, where possible, of Multi-factor authentication. Don't distribute local and privileged accounts beyond the fewest most essential persons.

Log all creations, deletions, elevations of privileges, of privileged accounts and monitor/ log their activity within the environment.

The above list is by no means exhaustive and might not be enough of themselves to protect against hypervisor attacks but they are essential to produce a modicum of due-diligence when attempting to secure a virtualized environment.

In the next section, this white paper will cover how Nutanix secures our AHV hypervisor.



3. Hardening AHV in 8 steps:

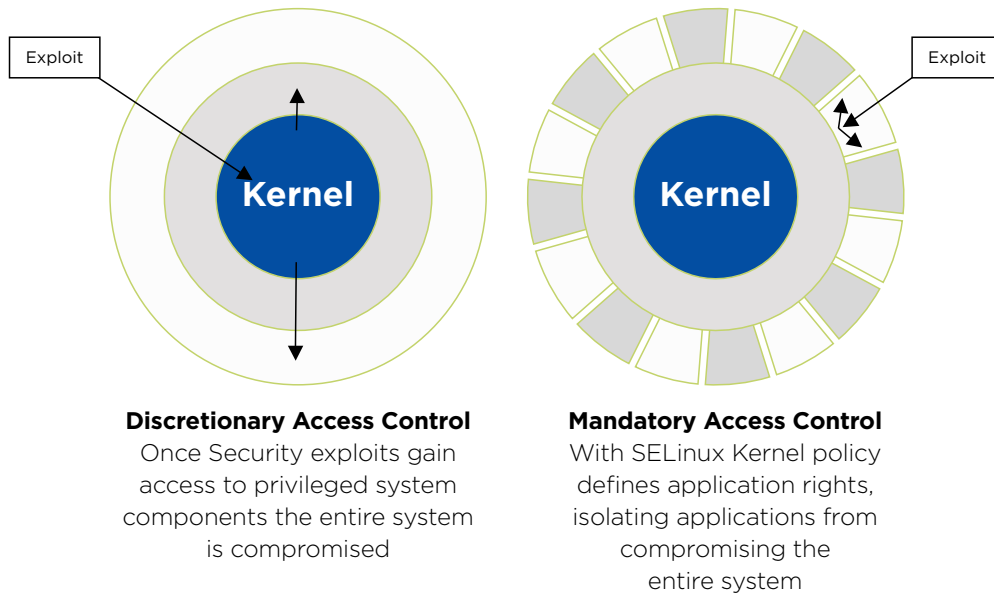
Nutanix uses the open-source Linux KVM as a baseline. A hypervisor that has been validated by scores of Open Source users globally, large enterprise cloud players and the Nutanix Security Engineering and Research Teams (nSERT) who have, through their testing, determining the best security practices for hardening and protection of the User VM environment which is to be supported.

3.1. SECURITY-ENHANCED LINUX (SELINUX)

Given that the KVM is a portion of the Linux kernel, along with additional services that AHV runs, things like libvirt and other pieces of the Linux ecosystem, to ensure security we run Security-Enhanced Linux (SELinux) full context enabled.

When using SELinux processes and files are labeled with an SELinux context, that contains additional information such as an SELinux user, role, type and level. All of this information is used to make access control decisions.

All files, directories, devices, and processes have this security context (or label) associated with them. For files, this context is stored in the extended attributes of the file system. The SELinux context contains additional information such as user, role, type and level allowing access control decisions on processes, Linux users, and files. This allows for policy type-enforcement, by tagging every process and file with a security context to which it belongs, a security enforcement module in the kernel permits or denies access to all objects (such as files and devices). The kernel process decides which subjects can access which objects, this is called Type Enforcement.



So SELinux is a flexible Mandatory Access Control architecture within the standard Linux kernel. What is effectively accomplished is a method to ensure hardware resource isolation. Each VM runs in its own security context, which isolates that CPU or Memory or Cache, based on the security context of that VM.

3.2. ADDRESS SPACE LAYOUT RANDOMIZATION

ASLR aims to introduce randomness into addresses used by a given task. This is a security mechanism that increases control-flow integrity by making it more difficult for an attacker to properly execute a buffer overflow attack even in systems with vulnerable software. Its strength lies in the randomness of the offsets it produces in memory.

3.3. EXEC-SHIELD

SELinux is used in conjunction with Exec-Shield and ASLR, which is a method developed by Red Hat to reduce the risk of worms or buffer overflow attacks or function pointer overflows on Linux systems, in other words, manipulating data in memory for malicious intent.³

³ For further details on exec-shield visit, <https://access.redhat.com/blogs/766093/posts/3534821>

3.4. LINUX ADVANCED INTRUSION DETECTION ENVIRONMENT (AIDE)

The Advanced Intrusion Detection Environment is a utility that creates a database of files on the system and then uses that database to ensure file integrity and detect system intrusions.

AIDE is the freeware branch of Tripwire. An Intrusion Detection system that conducts a checksum verification of all static binaries and libraries in AHV and in the CVM. This functionality has to be enabled by the customer via the Nutanix Command Line Interface (nCLI), directions on how to do this can be found in the Nutanix Security Guide via <https://portal.nutanix.com>. Once enabled AIDE runs weekly and sends its report directly to syslog which should be forwarded to a central log host when the CVM is configured to forward logs.⁴

3.5. HOST SECURE BOOT

Validating the authenticity of the boot sequence is key to defending against bootloader attacks. Secure boot is a technology where the system firmware checks that the systems boot loader is signed with an appropriate cryptographic key authorized by the manufacturer. Nutanix Secure Boot comes in two flavors:

- AHV Secure Boot for Host - which ensures that the AHV binaries are trusted and have not been compromised as part of the boot process of the node. AHV Secure Boot for host only works with Nutanix Supported Hardware that supports both UEFI and Secure Boot.
- AHV Secure Boot for UserVM - which verifies User VM Operating systems running on AHV. Supported for Windows or Linux versions that support UEFI and Secure Boot.

3.6. SECURITY TECHNICAL IMPLEMENTATION GUIDES

As mentioned, AHV is further protected by a bespoke AHV STIG (Security Technical Implementation Guide). A STIG is a hardening guide, a security methodology for standardizing security protocols within networks, servers, computers, and logical designs to enhance overall security. It describes the potential risks and vulnerabilities an attacker might exploit either physically or over a network, and includes a description of the action to take to prevent such a measure from being carried out.

⁴ Details on log forwarding can also be found in the Nutanix Portal <https://portal.nutanix.com>

Details of the contents of the Nutanix AHV STIG can be found in a PDF via the [Nutanix Portal](#). Nutanix STIGs are written in eXtensible Configuration Checklist Description Format (XCCDF) in support of the Security Content Automation Protocol (SCAP) standard. This allows the STIG to be a machine-readable STIG format which automates assessment tools and eliminates time-consuming testing.

Because the STIGs are machine-readable, they are ideal candidates for third-party apps that probe for deficiencies in a system configuration. The Nutanix STIG takes inspiration from the NIST SP800-53 as well as some bespoke protections which didn't exist for Hyper-converged platforms at the time of writing. Effectively the process of system/security hardening is completed by the STIG with several hundred checks to ensure fidelity, conducted periodically and checked for continuity.

3.7. SECURITY CONFIGURATION, MANAGEMENT AND AUTOMATION

A self-healing mechanism or adjustment to the potential "drift" that can happen from a secure configuration state is achieved by our Security Configuration Management Automation daemon (SCMA).

SCMA is a SaltStack daemon that runs periodically and can be adjusted to run more or less frequently in the command line. With SCMA Nutanix AHV and AOS can maintain this adherence to the IA posture that is provided out of the box for Nutanix products, ensuring a secure system on Day 0 and throughout the life of the platform.

3.8. THE SECURITY DEVELOPMENT LIFECYCLE (SECDL)

Finally, a word on our development process. Nutanix AHV engineers are training to think of secure coding as principal to the coding process. Code Analysis techniques, Vulnerability scanning, Threat Modelling, image and artifact scanning, AV scanning and rigorous Penetration testing are all conducted against the hypervisor. The nSERT team work with Nutanix developers to deliver a stripped down hypervisor while still having many enterprise features that customers have come to expect, delivering a smaller attack surface through the removal of superseded or superfluous code.



4. Conclusion

In closing, Nutanix AHV is just one infrastructure component that operates as part of the Nutanix Enterprise Cloud platform, safeguards exist throughout AOS and AHV in combination which have not been fully explored in this paper. Nutanix products and documented best practices are available to protect the data plane such as, Data-at-Rest Encryption, Role-Based Access Control, Micro-segmentation and network visualization in Nutanix Flow and detailed Security compliance auditing and reporting with Nutanix Xi Beam. For details on these capabilities and much more visit <https://www.nutanix.com>.



info@nutanix.com | www.nutanix.com |  @nutanix

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power their business. The Nutanix Enterprise Cloud OS leverages web-scale engineering and consumer-grade design to natively converge compute, virtualization, and storage into a resilient, software-defined solution with rich machine intelligence. The result is predictable performance, cloud-like infrastructure consumption, robust security, and seamless application mobility for a broad range of enterprise applications. Learn more at www.nutanix.com or follow us on [Twitter @nutanix](https://twitter.com/nutanix).