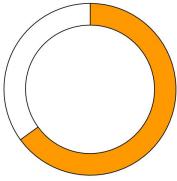


MDS CPU Side-channel Attacks May 2019

CVSSv3 Score - 6.5 Moderate	Advisory ID	nutanix-sa-014-mds	CVE(s)
	Last Updated	17 October 2019	CVE-2018-12126 CVE-2018-12127 CVE-2018-12130 CVE-2019-11091
	Published	16 May 2019	
	Version	13.0	

Final Update

Summary

On May 13th 2019, Nutanix Product Security was made aware of the latest Intel processor vulnerability, dubbed MDS (Microarchitecture Data Sampling). These flaws require local shell access to a system, and if present and exploited could allow data in the CPU's cache to be exposed to unauthorized processes. These vulnerabilities are difficult to execute because of their local access requirements to the host or guest; However a skilled attacker could use these vectors to read memory from virtual or containerized instances, or the underlying host itself.

Issue Description

The Microarchitectural Data Sampling (MDS) family of vulnerabilities refers to a collection of speculative side-channel vulnerabilities, specifically:

- CVE-2018-12126 - Microarchitectural Store Buffer Data Sampling (MSBDS)
- CVE-2018-12127 - Microarchitectural Load Port Data Sampling (MLPDS)
- CVE-2018-12130 - Microarchitectural Fill Buffer Data Sampling (MFBDS)
- CVE-2019-11091 - Microarchitectural Data Sampling Uncacheable Memory (MDSUM)

These vulnerabilities align to functions within modern Intel processors around Load Ports, Store Buffers and Fill Buffers. The Load Port data table is used to store addresses of CPU registers during the loading of data from memory or I/O subsystems. The Store Buffers are a shared buffer scheme that is used during STA (Store Address) and STD (Store Data) speculative operations, while Fill Buffers are used in non-speculative operations.

An attacker may use these mechanisms, under very specific circumstances, to later load and cause a fault in a manner that leaks stale data by way of a side-channel.

The MDSUM vulnerability is a special case to the others in that it pertains to uncacheable memory, which was once believed safe in speculative side-channel attacks.

Impact

An attacker with local access could create a malicious untrusted user process on a trusted guest, or even an untrusted guest, and intercept and sample data and recent operations by way of a side-channel, including recently used memory or I/O port writes.

The types of data included are:

- Previous execution context, including process, guest or hypervisor, at the same privilege level.
- Higher privilege execution context in cases where the attacker's execution was interrupted.

For data to be vulnerable, it must reside on the same core as the attacker. This does include, if Hyper-Threading is enabled, adjacent threads as well.

An attacker is unable to target specific data with these vulnerabilities. Only sampling over a period of time and other methodologies could result in the exposure of meaningful data.

Risks

Attackers exploiting these vulnerabilities within the MDS announcement are unable to target specific data, and would require large inspection and collection periods along with analysis to glean any important data from this exploit. Additionally, each vulnerability has only a local attack vector, meaning access to the local operating system is required along with the execution of malicious code, to perform the attack. Therefore, within a single tenant environment that homes trusted systems, the overall risk is lower. However, in multi-tenant environments, or environments that do not house entirely trusted workloads, the risk stands as Moderate.

Nutanix AHV will receive updated microcode and kernel code first, since it has the higher risk profile of UVM to Hypervisor vectors. Systems such as Nutanix AOS are

considered trusted systems with mechanisms in place to ensure processes are protected and known and will receive updates post hypervisor. If you run a hypervisor other than AHV, please consult with that vendor for updates as they are made available. Links are provided in the sources section for further information.

Affected Products

This document will be updated with the patch release schedules. Please check the [Nutanix Support Portal](#) for the latest update.

Nutanix Products

Updates to processor microcode, and kernel, are required for mitigation.

Product	Fix Release
AHV	<i>Now available via 20170830.279 within 5.5.9.5, 5.10.5 and 5.11 on the Nutanix Support Portal.</i>
Nutanix AOS	<i>Now available via AOS 5.10.7 and 5.11.1 on the Nutanix Support Portal.</i>
Prism Central	<i>Now available via PC 5.11.1 on the Nutanix Support Portal.</i>
Files	<i>Now available via Files 3.6 on the Nutanix Support Portal.</i>
Move	<i>Now available via Move 3.1.0 on the Nutanix Support Portal</i>
X-Ray	<i>Now available via X-Ray 3.5 on the Nutanix Support Portal.</i>
Era	<i>Now available via Era 1.1 on the Nutanix Support Portal.</i>

BIOS Updates (Microcode)

Family	Fix Release
G4 & G5	<i>Now available via FW G4G5T 5.0 and G4G5U within LCM 2.2.3.</i>
G6	<i>Now available via PB41.002 and PU41.002 within LCM 2.2.2.1</i>

3rd Party Products

Product	Fix Release
VMware ESXi	VMSA-2019-0008
Microsoft Hyper-V	ADV-190013

Mitigation Notes

In addition to microcode and operating system updates there are two additional mitigations for concurrent attack vectors.

- a. For ESXi - enable the side-channel-aware scheduler.
- b. For AHV - the AHV scheduler load balancing behaviour is as follows:
 1. If system load < 50%, vCPUs belonging to two different VMs will not be maintained on the same logical core pairing.
 2. If system load > 50%, vCPUs belonging to two different VMs will not be on the same logical core pairing for a significant length of time.

This means it will be highly impractical (but not impossible) to build an inter-VM concurrent attack vector on a AHV with hyperthreading enabled. Particularly if system load is typically below 50%.

To 100% mitigate the concurrent attack vector Hyperthreading can be disabled at the hypervisor level via instruction in [KB6137](#). Note that Nutanix does not recommend you disable hyper-threading.

Sources

Intel:

<https://software.intel.com/security-software-guidance/software-guidance/micro-architectural-data-sampling>

Red Hat: <https://access.redhat.com/security/vulnerabilities/mds>

Red Hat (Video) All about MDS in 3 Minutes: <https://youtu.be/Oeb-O4yKK2c>

XEN: <https://xenbits.xen.org/xsa/advisory-297.txt>

Kernel.org:

<https://www.kernel.org/doc/html/latest/admin-guide/hw-vuln/mds.html>

Fallout: <https://mdsattacks.com/>

ZombieLoad: <https://zombieloadattack.com/>

Support

If you have questions, please open a case with Nutanix Support at <http://portal.nutanix.com>, or by calling Support at the phone number on the website <http://www.nutanix.com/support>.

Revision History

Version	Section	Date
1.0	-	16 May 2019
2.0	Release information	22 May 2019
3.0	Release timelines	29 May 2019
4.0	Updated releases and timelines	12 June 2019
5.0	Updated timelines	25 June 2019
6.0	Updated releases and timelines	10 July 2019
7.0	Updated timelines	25 July 2019
8.0	Updated releases	16 August 2019
9.0	Updated releases	04 September 2019
10	Updated release timelines	18 September 2019
11	Updated releases	03 October 2019
12	Updated releases	08 October 2019
13	Final Update	17 October 2019