The Aftermath of Spectre and Meltdown

The computing world has witnessed dramatic increases in capacity in terms of CPU, memory, and storage that can be installed in a server. This led industry innovators to develop virtualization and cloud technology that allows for the pooling and scheduling of workloads to optimize the use of compute resources. In response, the microprocessor industry began incorporating hardware features that increasingly allowed for sharing of compute resources including processor cores, cache, and bus access among workloads – all with a focus on stretching the IT dollar by increasing performance and optimizing allocation of resources.

At the same time, security at the OS and application layers was steadily improving. This led threat actors and security researchers to seek out softer targets at other layers of the infrastructure environment including the virtualization stack, management backplanes, firmware, and even hardware as new areas to exploit.

In early 2018, security researchers began disclosing a series of side channel vulnerabilities affecting a range of microprocessors. Two of these, Spectre and Meltdown, became commonly known and widely reported. These vulnerabilities were catastrophic, as they allowed an attacker to read secret information including passwords and encryption keys – potentially *any* sensitive information being handled by the applications running.

The implications were far reaching. In a world that had largely embraced multi-tenant IaaS clouds, public SaaS, and widespread use of virtualization technologies, IT managers were suddenly forced to (re)consider where their data was being processed, and where their workloads were running. Fixes for these vulnerabilities appeared to negatively impact performance, and new variants continued to be discovered.

Those running the most sensitive workloads were facing the harrowing possibility of having to separate those workloads onto dedicated hardware – a drastic step that meant walking away, or at least walking back from IT modernization initiatives that led to significant performance gains, scalability, and reduced costs for their organizations. Spectre and Meltdown were a wakeup call for those responsible for high security and regulated workloads within their organizations. For defenders, it was clear that the whole stack including hardware, firmware, management plane, and virtualization layer could be targeted going forward.

Secure Runtime Environment (SRE)

HPE and Intel have been collaborating to develop a combined hardware and software solution running a new type of virtualization platform known as a microvisor to counter future side-channel variants and other types of attacks. Known as the Secure Runtime Environment (SRE), it was developed to mitigate or eliminate future Spectre and Meltdown variants, future side channel vulnerabilities as well as to provide unprecedented levels of policy-based protections for hardware resources. This SRE would restore confidence in the use of multi-tenant IaaS clouds, public SaaS, and virtualization technologies for workloads with more rigorous security requirements.

HPE ProLiant DL380 Gen10
2x Intel Xeon Gold 6248 CPU (Cascade Lake)
iLO5 Advanced License
HPE and Intel Secure Runtime Environment (SRE)

Figure 1 - Test Subject

InfusionPoints was retained to evaluate the SRE solution via a penetration test in the spring of 2019. We were provided with an HPE ProLiant DL380 Gen10 server configured for secure boot into the SRE. The SRE was preconfigured to provide a multi-tenant setup, where we would attempt to attack the solution to gain access to the

other tenant's virtual machines, secret information, or otherwise compromise the host.

The Results

InfusionPoints first took a forensic bit for bit image of the drive, hoping to disrupt the boot process of the SRE environment or compromise information contained in local domain storage. InfusionPoints analyzed the image to find weaknesses in the implementation of full disk encryption. The disk was encrypted using LUKS, the Linux Unified Key Setup. While inspecting the BIOS of the server, Secure Boot keys were found to be installed and properly enabled. InfusionPoints was unable to find any misconfiguration in the disk encryption or Secure Boot implementations.

InfusionPoints then attempted to launch attacks against the system from our tenant domain against the other tenant and system resources. InfusionPoints was given root access to our domain which was running the popular Kali linux penetration testing distribution. InfusionPoints tried multiple attack vectors from the Kali domain.

The first type of attacks InfusionPoints looked to exploit were the original Spectre and Meltdown vulnerabilities and several variants. InfusionPoints was able to use publicly available exploit code to test for these vulnerabilities, which we could not successfully exploit.

Seeking a more recent exploit, InfusionPoints attempted to use a side channel vulnerability known as ZombieLoad which had just been made public in May of 2019. ZombieLoad is a data-sampling attack that can expose application data by faulting load instructions and exposing data to a program listening on the same CPU thread. The proof of concept code provided by the security researchers failed in stealing memory secrets from other processes. Unlike Spectre and Meltdown, this side-channel attack had just become available and based on public information was believed to affect the Cascade Lake architecture in our test subject[1]. The SRE was able to prevent this exploit in at least two ways. The first mitigation was disabling hyperthreading, a requirement for running workloads with the most rigorous isolation requirements on the SRE. The second mitigation, known as

hyperthreading sibling containment, is used if hyperthreading is enabled. This forces both the physical and virtual cores to be allocated to the same domain, providing the best isolation between domains.

InfusionPoints leveraged a CPU fuzzer called Sandsifter to generate a page fault to see how the SRE would respond. The SRE determined that this action was a security policy violation resulting in the termination and shutdown of our Kali domain to protect the system and other running domains. We also validated that the SRE provides policy-based access to hardware that can be permitted and monitored by verifying the SRE's behavior when writing a string to an LPC IO port.

Conclusion

The HPE and Intel solution running the SRE Microvisor is a unique technology that significantly reduces the attack surface and mitigates against side channel attacks. We were unable to successfully circumvent the protections provided by the SRE. Combined with HPE's Silicon Root of Trust, and Secure Boot technology, the solution provides a trusted execution environment that has integrity validation built in from the moment the server is powered on. Its plug and play architecture allows it to easily integrate into well-known orchestration projects based on libvirt and the ubiquitous OpenStack, providing an opportunity for cloud providers to differentiate their cloud by offering new instance classes with high security profiles. Private cloud operators with regulated workloads can feel more confident that they have strong virtual isolation between instances running in their cloud environment and have an answer to these contemporary threats.

About InfusionPoints

InfusionPoints is your independent trusted partner dedicated to assisting you in building your secure and compliant business solutions, testing your security controls and defending your consumer, employee, and supply chain information.

+1-336-990-0252
info@InfusionPoints.com
https://www.infusionpoints.com/

---

[1] https://techcrunch.com/2019/05/14/zombieload-flaw-intel-processors/