



Real-World Impact Measurement of Spectre and Meltdown Patches

Executive Summary

Spectre and Meltdown are architectural flaws in many modern microprocessors which can, under some circumstances, allow malicious software to steal data from other running processes or applications. The discovery of these flaws, announced in early 2018, set off a scramble to understand the impact of those vulnerabilities, and create remedies for future and existing hardware.

Within a few months, the initial patches for Spectre and Meltdown emerged, but there were two problems still to overcome.

The first is that the patches quickly gained the reputation for affecting the systems' performance, though the precise nature of the performance impact was a matter of anecdote, not rigorous measurement.

The second challenge is that Spectre and Meltdown, as announced in January 2018, were the tip of the iceberg; using those vulnerabilities as a guide, researchers have discovered other, similar architectural flaws. This process – and the process of creating patches – is still ongoing.

This paper summarizes real-world impact measurements of the Spectre and Meltdown patches, based on benchmarks run at a Spirent customer site on real-world servers and real-world workloads. Empirical data on the impact of the patches is presented – eliminating guesswork.

Armed with this information, organizations can assess their own plans for patching against Spectre and Meltdown, while being prepared with mitigating the significant performance impact through additional hardware deployments, workload shifts, or other methods.

Real-World Impact Measurement of Spectre and Meltdown Patches

A Big Processor Architectural Flaw

There are two competing forces at work within many Enterprises and Cloud Service Providers (CSP). On one hand, Security managers, such as the Chief Risk Officer (CRO) and Chief Information Security Officer (CISO), want to install patches to address the [Spectre and Meltdown vulnerabilities](#) found in many microprocessors. On the other hand, IT operations staff are also concerned about the performance implications of those patches and wish to delay patching until all side-effects are fully understood.

Why is this an issue? There is a serious concern being voiced that the patches may cause a significant performance impact on servers. This paper discusses the implications of these two new threats while providing further insight into authoritative test data based on real-world tests, conducted by Spirent, at a major customer data center.

Spectre and Meltdown burst onto the scene in January 2018, as two separate teams of researchers discovered fundamental flaws in many microprocessors, including those from AMD, ARM and Intel. Due to these flaws, malicious applications could access private memory used by other processes by exploiting side channels to access the microprocessors' caches.

Spectre and Meltdown are two separate flaws, although both exploit caches through side channels, which is why they are often discussed together.

Introducing Meltdown

Meltdown was independently discovered and reported by three teams: Jann Horn (Google Project Zero); Werner Haas and Thomas Prescher (Cyberus Technology); and Daniel Gruss, Moritz Lipp, Stefan Mangard and Michael Schwarz (Graz University of Technology).

The Meltdown flaw breaks down the isolation between an application and the operating system, such as Linux, Unix, Windows, Android, Mac OS or iOS. This vulnerability basically melts security boundaries which are normally enforced by the hardware, hence the Meltdown name.

Leveraging this vulnerability, a malicious application could read or change protected operating system memory that is normally off-limits to applications, potentially letting the rogue application corrupt the operating system, or steal information. Operating system and firmware patches can remediate against Meltdown.

Introducing Spectre

Spectre was independently discovered and reported by two groups of people: Jann Horn (Google Project Zero) and Paul Kocher in collaboration with Daniel Genkin (University of Pennsylvania and University of Maryland), Mike Hamburg (Rambus), Moritz Lipp (Graz University of Technology) and Yuval Yarom (University of Adelaide and Data61).

Spectre allows a malicious application to access, or overwrite, the memory of another application, potentially allowing it to steal information or corrupt that application's operation. For example, malware could attack an application, accessing passwords, financial account information or encryption keys. The Spectre name is based on the root cause, speculative execution, which we'll cover in more detail further on.

The Spectre vulnerability is extremely difficult for malware to exploit, due to the number of steps the malware must execute to target the microprocessor, and because the memory access is indirect - in other words, the microprocessor is being tricked in subtle ways to reveal what is stored in memory. Unlike the Meltdown situation, there are no general means of patching against every possible attack that could exploit the Spectre vulnerability. However, there are patches that protect against known exploits that leverage the Spectre vulnerability.

Fixing the Vulnerabilities

The long-term, permanent fixes for both Meltdown and Spectre will require changes to microprocessor design that change how internal caches operate and eliminate the side channel flaws. Future processors, still being designed, will be immune to some or all those flaws.

These updated chips will soon be available in the commercial market. [Intel has said](#) that it will release versions of its Xeon and Core processors, which will address some of the Spectre and Meltdown vulnerabilities, in the second half of 2018. The other vulnerabilities will continue to be addressed through firmware and operating system patches, at least for now.

However, even with new chips coming, the reality is that servers, desktops, notebooks, mobile devices and embedded systems are still vulnerable today. It may take several years before most systems with vulnerable microprocessors are replaced. Therefore, it is necessary to install software patches that will protect systems by bypassing the Spectre and Meltdown flaws, even at the cost of reduced system performance.

Patches against Spectre and Meltdown come from two main types of organizations.

- **Hardware manufacturers** are offering firmware updates for specific computer models, including mobile devices, desktops/notebooks and servers. For example, Dell and Hewlett-Packard Enterprise offer firmware (BIOS) updates for some of their enterprise servers.
- **Operating system makers** are providing more general patches. Many Linux providers, Microsoft (for Windows), Google (for Android and Chrome OS), Apple (for Mac OS and iOS) and others have provided patches or operating-system upgrades for current and other supported versions of their platforms.

Even applications makers are offering Spectre and Meltdown patches. Google has updated its Chrome 64 web browser to detect if a malicious website is trying to exploit the Meltdown or Spectre vulnerabilities via embedded JavaScript code.

There are performance implications to these software patches and firmware updates, because they don't allow the operating system and applications access to vulnerable operations - including the chip's cache. The upside being improved security, the downside resulting in reduced memory read or write speeds, as well as reduced CPU throughput and in some cases, I/O throughput.

Real-World Impact Measurement of Spectre and Meltdown Patches

A New Round of Vulnerabilities

The original vulnerabilities were revealed in January 2018. In May 2018, [researchers found a new class of side channel vulnerabilities](#). One type of attack works a processor's registers, another leverages the chip's speculative store bypass (SSB) functions. At the time of writing, software patches are not yet available to address these newest vulnerabilities, and as such, they are not covered by this paper's test results.

Think of the Spectre and Meltdown situations like a game of whack-a-mole. As soon as one set of side-channel vulnerabilities are understood, and patches made available, new side-channel vulnerabilities appear.

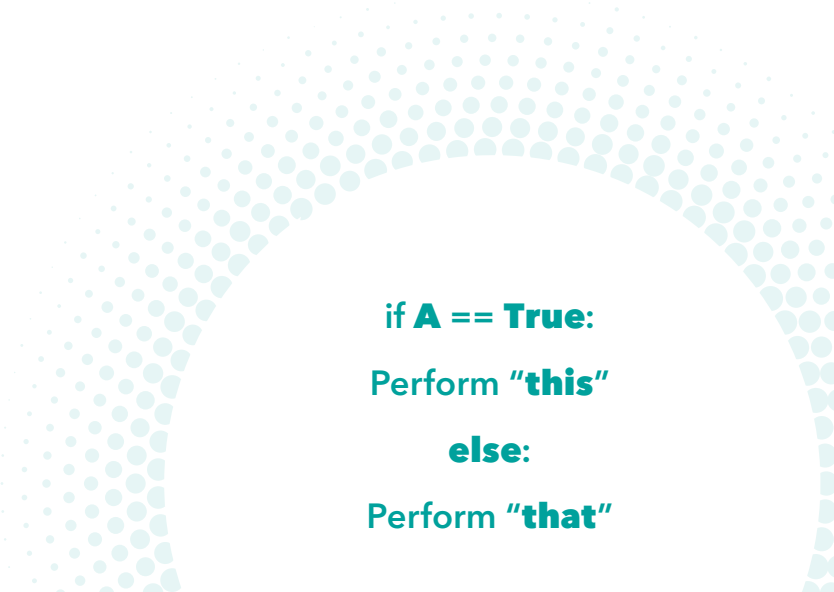
We are now about one year into Spectre and Meltdown. The research teams behind these vulnerabilities notified microprocessor makers about the flaws on June 1, 2017; this information was kept confidential, so that the chip makers could work on understanding and mitigating the flaws. The public learned about these vulnerabilities on January 3, 2018. The first patches were released in late January.

Then came the whack-a-mole scenarios, with eight new flaws revealed in early May, affecting microprocessors from AMD, ARM and Intel. Since then, the chipmakers have been working hard to understand these flaws and determine ways to mitigate them. As of June 13, technical details on three of the eight new flaws have been published by various sources, providing information that can help processor makers and software vendors develop patches, but it also helped attackers create further exploits.

A significant challenge is the makeup of the Spectre vulnerabilities, which affect a processor feature called speculative execution. In **speculative execution**, the processor looks ahead to future branches in software code - think "If A is true do this, if A is false, do that." Instead of waiting to determine if A is true or false, advanced microprocessors begin calculating the actions for both "this" and "that." When the processor gets to the branch condition, the actions based on that branch are already prepared. So, if A is true, the "this" results are used, and the "that" results are discarded. Speculative execution can offer significant performance benefits.

The weakness exploited by Spectre points to the unintended consequences of speculative execution. In theory, the results for "that" are discarded and have no lingering effect. However, the Spectre flaws show that malicious software can exploit speculative execution to access normally inaccessible data or manipulate the processor into revealing information about other running programs - which is the side-channel flaw.

Speculative execution is an essential feature of modern processors. It is not a specific instruction or function that can be ignored or disabled. That is why it is impossible to block all possible Spectre exploits. However, researchers continue to discover new methods of speculative execution attacks, which helps with writing software patches or microprocessor fixes to mitigate those specific attacks.



if A == True:
Perform "this"
else:
Perform "that"

How Spirent Tested Spectre and Meltdown Patches

Spirent worked with one of its customers, a large cloud provider, to test the effects of Spectre and Meltdown patches on servers in a test facility. This facility is used to test new applications, as well as other operating system and application updates, before they are deployed in the customer's production environments.

The testbed consisted of two identical servers – one the control, the other patched to mitigate the effects of Spectre and Meltdown. The tests were part of an engagement designed to help the carrier determine the effect of the patches on actual workloads. In other words, this was real-life – not a laboratory simulation.

The servers were built by the carrier, based on [Super Micro X9DR3-F motherboards](#):

- Dual socket R (LGA 2011) with two 2.5GHz ES-2670 v2 processors, 20 cores, 64GB ECC DDR3 RAM
- Intel C606 chipset
- Intel i350 Dual port GbE LAN
- 4x SATA2 and 2x SATA3 ports
- 8x SAS ports from the C606 chipset
- Integrated IPMI 2.0 and KVM with Dedicated LAN

The servers were running [CentOS 7.4](#), a version of Linux derived from Red Hat Enterprise Linux, with kernel 3.10.0-327.13.1.el7.x64_64.

Before patching, both servers were tested with the [Spectre and Meltdown mitigation detection tool v0.37+](#), which showed that the servers were vulnerable to all five of its speculative-execution vulnerabilities:

- CVE-2017-5753 aka Spectre Variant 1
- CVE-2017-5715 aka Spectre Variant 2
- CVE-2017-5754 aka Meltdown or Variant 3
- CVE-2018-3640 aka Variant 3a
- CVE-2018-3639 aka Variant 4

The test server was patched, and then mitigation test tool was re-run. As expected, the tool showed that the first three Spectre variants were no longer a threat, but the test server was still vulnerable to variants 3a and 4. At the time of this project, no broad patches were available to protect against those newly discovered vulnerabilities:

- In the case of Spectre Variant 3a, up-to-date CPU microcode is needed to mitigate this vulnerability, and that microcode was not available.
- In the case of Spectre Variant 4, the Intel Xeon ES-2670 v2 processors don't support the necessary SSBD (Speculative Store Bypass Disable) bit needed for remediation. A future firmware or microcode update may address that situation.

Real-World Impact Measurement of Spectre and Meltdown Patches

Benchmarking Infrastructure

Two separate tools were used to set up the workload testbed (Spirent CloudStress) and run the benchmarks (Spirent CloudScore).

- [Spirent CloudStress](#) is a workload generator that can simulate the physical footprint of any virtual machine. It allows for application workloads to be run, recorded, and then played back. This not only simplifies the execution of multiple types of workloads during a test but minimizes potential errors due to unintended changes in environment or configuration. These workloads stress the CPU, memory, storage and network I/O.
- [Spirent CloudScore](#) is the industry's first benchmarking system able to assess and compare the performance of virtualized or cloud infrastructure. It provides a scorecard, backed-up by industry standard benchmarking tests, to help grade various elements of workload performance, including CPU, memory access, storage I/O and network I/O. CloudScore can be used test the impact of any change to an environment, from hardware changes to operating system upgrades to applications patches. Though not specifically designed for relatively simple software and firmware patches (such as Spectre and Meltdown), CloudScore is perfectly suited to this task.

The tests using CloudStress and CloudScore were configured and run in the customer's test-lab environment by Spirent engineers, working on-site with the customer's own IT staff and software architects.

A Dramatic Slowdown



Figure 1. Cloud Infrastructure Performance Summary.

Testing of the patch-and-control performance of the servers was focused on three areas: Compute, network I/O, and storage I/O. In two of the three cases, the impact of the patches was severe.

The compute tests used a common methodology, that of compressing files using [gzip](#) and [7-Zip](#). The gzip test showed a dramatic 25% performance hit; while the 7-Zip test shows a minimal 3% performance degradation. The CPU tests also tested kernel compilation time, which dropped by 2% after applying the patch. The industry-standard [UnixBench](#) index showed a 1% decrease.




Figure 2. Compute Performance Summary.




Figure 3. Network Performance Summary.




Figure 4. Storage Performance Summary.

The network tested measured throughput and latency for different network scenarios, such as transfers within same availability zone (intra-AZ) or across different availability zones (inter-AZ). The results were uniformly lower. For example, the Intra-AZ latency in microseconds test showed that the patched server had much less performance (by approximately -24%) than the control server. Similarly, the measurement of jitter in milliseconds test shows that the patched server had much less performance (by approximately -138%) than the control server.

The storage tests showed considerable impact. The tests analyze I/O performance using different tools and different test scenarios and provide deep analysis of cloud performance related to data storage tasks. The test of write throughput in operations-per-second showed a drop from 3,000 op/sec to 2,700 op/sec with 8M 4K blocks. With 16M 4K blocks, the throughput drop was even larger, from 4,700 op/sec to 2,800 op/sec.

Real-World Impact Measurement of Spectre and Meltdown Patches

About Spirent Communications

Spirent Communications (LSE: SPT) is a global leader with deep expertise and decades of experience in testing, assurance, analytics and security, serving developers, service providers, and enterprise networks.

We help bring clarity to increasingly complex technological and business challenges.

Spirent's customers have made a promise to their customers to deliver superior performance. Spirent assures that those promises are fulfilled.

For more information, visit: www.spirent.com

Analysis and Recommendations

Performance suffers after the application of the Spectre and Meltdown patches. These and other tests have shown that the degree of performance degradation varies widely, depending on the nature of the server, the model of the specific processor and operating system kernel, and the nature of the workload. For example, a server engaged in displaying web pages - that is, doing a lot of database reads - will degrade differently than a server engaged in storing e-commerce transaction with a workload executing many database writes.

The degree of virtualization also matters in ways that may not be expected. Communication between guests and kernel hosts on the machine to process network I/O can be significantly affected by the Spectre and Meltdown patches. For example, the tests showed significant performance impact in an OpenStack environment with traffic flowing to/from Open vSwitch (OVS) bridges on the local system.

The variability in results based on hardware, software, and workload, the constantly evolving nature of Spectre and Meltdown - and their patches - makes it impossible to offer general guidance. Will datacenters need additional compute resources to offset the performance impacts of the patches? Almost certainly, but there is no fixed rule.

In some cases, there may be no need for additional hardware. Simply rebalancing workloads will be sufficient. In other scenarios, it may be advisable to start investing in more hardware or begin looking at shifting workloads to the cloud. If you are concerned about the performance impact of Spectre and Meltdown patches on your environment or workloads, the best solution is to test for their impact. Infrastructure benchmarking and simulation tools such as Spirent CloudStress and Spirent CloudScore, as well as [Spirent Professional Services](#) can assist in making precise predictions of the impact of these patches to help mitigate risks and maximize productivity. [Contact Spirent today to learn more.](#)



Contact Us

For more information, call your Spirent sales representative or visit us on the web at www.spirent.com/ContactSpirent.

www.spirent.com

© 2018 Spirent Communications, Inc. All of the company names and/or brand names and/or product names and/or logos referred to in this document, in particular the name "Spirent" and its logo device, are either registered trademarks or trademarks pending registration in accordance with relevant national laws. All rights reserved. Specifications subject to change without notice.

Americas 1-800-SPIRENT

+1-800-774-7368 | sales@spirent.com

US Government & Defense

info@spirentfederal.com | spirentfederal.com

Europe and the Middle East

+44 (0) 1293 767979 | emeainfo@spirent.com

Asia and the Pacific

+86-10-8518-2539 | salesasia@spirent.com