



Mitigation Overview for Potential Side-Channel Cache Exploits in Linux*

White Paper

Revision 2.0
May, 2018



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at www.intel.com.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

Intel provides these materials as-is, with no express or implied warranties.

Intel, the Intel logo, Intel Core, Intel Atom, Intel Xeon, Intel Xeon Phi, Intel® C Compiler, Intel Software Guard Extensions, and Intel® Trusted Execution Engine are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2018, Intel Corporation. All rights reserved.



Contents

1.0	Introduction.....	5
1.1	Background.....	5
2.0	Bounds Check Bypass (Spectre variant 1)	6
3.0	Branch Target Injection (Spectre variant 2).....	7
4.0	Rogue Data Cache Load (Meltdown variant 3).....	9
5.0	Rogue System Register Read (Meltdown variant 3a / CVE-2018-3640).....	10
6.0	Speculative Store Bypass (variant 4 / CVE-2018-3639).....	11
7.0	Mitigation Availability	13



Revision History

Date	Revision	Description
January 2018	1.0	Initial release.
May 2018	2.0	Added information on Variant 3a, Rogue System Register Read, and Variant 4, Speculative Store Bypass.



1.0 Introduction

Intel is collaborating with the Linux* kernel community and industry partners to help mitigate potential [side-channel cache exploits](https://www.intel.com/content/www/us/en/architecture-and-technology/facts-about-side-channel-analysis-and-intel-products.html) (<https://www.intel.com/content/www/us/en/architecture-and-technology/facts-about-side-channel-analysis-and-intel-products.html>). This document can help those in the industry make informed decisions regarding various Linux kernel mitigations if they are not already consuming them from other sources like Linux distributions and cloud providers.

To find the latest security updates on these issues, including the latest status on microcode, see the [Security Advisory](https://security-center.intel.com/advisory.aspx?intelid=INTEL-SA-00088&languageid=en-fr) (<https://security-center.intel.com/advisory.aspx?intelid=INTEL-SA-00088&languageid=en-fr>).

1.1 Background

As detailed by [Google* Project Zero](https://googleprojectzero.blogspot.com/) (<https://googleprojectzero.blogspot.com/>) and [security researchers](https://spectreattack.com/spectre.pdf) (<https://spectreattack.com/spectre.pdf>), new side-channel analysis methods were discovered that potentially facilitate access to unauthorized information. These methods rely on a property of modern processors called speculative execution. Despite this commonality, the methods for mitigating these “variants” are very different. As a result, we discuss the mitigations separately.

Anyone consuming security mitigations must be prepared to consume mitigations that will continually be morphed and improved over time as opposed to consuming them at any single point in time. Linux distributions and stable kernel trees both provide these streams of mitigations.

Because Linux is continually improving, the latest upstream kernel is expected to be the most hardened against all security vulnerabilities.



2.0 *Bounds Check Bypass (Spectre variant 1)*

This *bounds check bypass* (Spectre variant 1) exploit targets a processor's *conditional branch instructions*. The operating system uses conditional branches for processing and validating virtually all data, including untrusted data processed by the kernel. Even code appearing "correct" can potentially be exploited to gain unauthorized access to information.

To mitigate this exploit technique in the Linux kernel, we first identify instruction sequences that can be tricked into exploitable behavior. Since a sequence must fit a specific pattern and operate on untrusted data before any potential exploit, not all conditional branches are exploitable. Techniques such as static analysis or manual inspection can identify these sequences.

Linux contains a helper function (`array_index_mask_nospec()`) which can mitigate potentially vulnerable sequences that have been identified.

The Google Project Zero exploit for the bounds check bypass mitigation (Spectre variant 1) targets a very specific kernel feature called eBPF (enhanced Berkeley Packet Filter). When targeting this feature, an attacker actively inserts vulnerable instruction sequences into the kernel which can then be leveraged into a full exploit. The mitigations for eBPF are essential for mitigating bounds check bypass (Spectre variant 1) since this exploit technique is publicly documented.

Pointers supplied by applications are common sources of untrusted kernel data and have always been validated to help ensure security and stability. Linux contains fence instructions in its user access infrastructure to mitigate against bounds check bypass in code paths that handle untrusted data.

As additional manual inspection and static analysis is performed, the list of vulnerable code will be refined and patched if possible.



3.0 Branch Target Injection (Spectre variant 2)

The *branch target injection* (Spectre variant 2) exploit targets a processor's *indirect branch predictor*. Indirect branches are used very differently than the conditional branches that the first exploit may target. In Linux, indirect branches are used relatively rarely compared to conditional branches, but they are used in critical locations. In addition, the compiler may insert indirect branches without the programmer ever being aware.

Since the compiler generates these branches, mitigation against this exploit is the most straightforward when the compiler can simply avoid generating vulnerable branch sequences. A software construct called [retpoline](https://support.google.com/faqs/answer/7625886) (<https://support.google.com/faqs/answer/7625886>) can be deployed to help ensure that a given indirect branch is resistant to exploitation. Retpoline requires updated microcode to make the speculation hardware behavior more predictable on some processor models. The compiler can insert retpolines, or the programmer can insert them manually. Retpoline deliberately steers the processor's branch prediction logic to a trusted location, preventing a potential exploit from steering them elsewhere.

As an example: It's like the gutter on a bowling lane: either the ball hits the pins or it goes to a controlled place (the gutter). It will not end up in the next lane.

Mitigation with retpoline requires obtaining a compiler with retpoline support, as well as a kernel which has been updated to support retpoline. The kernel build then enables special compiler options instructing the compiler to automatically generate retpoline instead of vulnerable indirect branches. In addition to retpoline inserted by the compiler, assembly code making indirect calls has also been manually updated to avoid exploitable indirect calls. This is generally accomplished using retpoline directly from assembly.

Processors based on Broadwell and later microarchitectures may consult the indirect branch predictor when returning from function calls. On these processors, maximum mitigation with retpoline requires following Intel's microcode guidance. See the [Security Advisory](https://security-center.intel.com/advisory.aspx?intelid=INTEL-SA-00088&languageid=en-fr) (<https://security-center.intel.com/advisory.aspx?intelid=INTEL-SA-00088&languageid=en-fr>) for more details. Additional mitigations, enumerated by processors' CPUID instruction and Family/Model signatures, are described in detail in a [whitepaper](https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf) (<https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf>) and include the use of Supervisor Mode Execution Protection (SMEP) and manipulation of the Return Stack Buffer (RSB).

In addition to the software-based retpoline mitigation, Intel is providing a microcode-based solution using indirect branch control mitigations described in the [Speculative Execution Side Channel Mitigations white paper](https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf) (<https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf>). Although some vendors are pursuing



mitigations entirely based on these techniques, the Linux community has chosen to proceed with the retpoline approach.



4.0 Rogue Data Cache Load (Meltdown variant 3)

The rogue data cache load (Meltdown variant 3) exploit targets a processor's speculative data loading mechanisms. Even though the processor's access control may protect a piece of data, it may still be read speculatively before an illegal access is determined to exist. The first two exploits require manipulating the kernel to do something *for* an attacker, while this exploit occurs entirely within code which is under the control of an attacker. This means we cannot modify the kernel to mitigate this exploit—we must fundamentally change where kernel data is available.

The mitigation for this is conceptually very simple: instead of relying on a processor's access-control mechanisms to protect data, simply remove the data instead.

As an example: If you want to protect your laptop from thieves, you can either place it in your car's trunk (rely on access-control) or you can take it with you (remove the data). The latter is more effective.

In Linux, this mitigation is referred to as Kernel Page Table Isolation (KPTI). This mitigation removes the data from the reach of exploits by having the kernel maintain two independent copies of the hardware page tables. One copy contains the minimal set of data and code needed to run an application and enter or exit the kernel, but it does not contain valuable kernel data. This helps ensure no valuable data can be *leaked* by an exploit running outside the kernel. The other set of page tables, active only while the kernel is running, contains everything needed for the kernel to function normally, including its private data.



5.0 Rogue System Register Read (Meltdown variant 3a / CVE-2018-3640)

Rogue System Register Read (Meltdown variant 3a) is mitigated by updated processor microcode. There are no changes to Linux itself to mitigate this issue.



6.0 Speculative Store Bypass (variant 4 / CVE-2018-3639)

The *speculative store bypass* method takes advantage of a performance feature present in many high-performance processors that allows loads to speculatively execute even if the address of preceding potentially overlapping store is unknown. Such a case may allow a load to speculatively read a stale data value.

Intel is providing a microcode-based mitigation referred to as Speculative Store Bypass Disable (SSBD) which is described in the [Speculative Execution Side Channel Mitigations white paper](https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf) (<https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf>). It requires updates to processor microcode. This mitigation may be enabled in applications at the thread level by via a new `prctl()` call: `PR_SET_SPECULATION_CTRL`. This mitigation is inherited by child processes, which makes it possible to deploy from a wrapper program.

For environments where `PR_SET_SPECULATION_CTRL` cannot be deployed, SSBD can be set for the entire system using the kernel boot parameter `spec_store_bypass_disable=on`. This parameter may also be used to mitigate the kernel itself against speculative store bypass.

Managed runtime systems are more likely to be vulnerable because an attacker has more influence on code generation. Linux includes a managed runtime called Berkeley Packet Filter (BPF) which can be mitigated to protect against speculative store bypass. These mitigations may include the use of SSBD while executing BPF programs or detection and mitigation of vulnerable code sequences.

A system's vulnerability to speculative store bypass can be detected by looking for the presence of the `/sys/devices/system/cpu/vulnerabilities/spec_store_bypass` file. The file's contents can be examined to determine if a system is vulnerable and which mitigations are available. For instance, on a vulnerable system, the file contains:

```
Mitigation: Speculative Store Bypass disabled via prctl and seccomp
```

Mitigations for speculation-based side channels which are not specific to Linux may also be used such as address space isolation or the insertion of memory fences inside an application. Some of those approaches are described in the [Speculative Execution Side Channel Mitigations white paper](https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf) (<https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf>).

Intel has provided recommendations in Section 4.2.2.2 "Software Usage Guideline" of the [Speculative Execution Side Channel Mitigations white paper](https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf) (<https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf>).



[Execution-Side-Channel-Mitigations.pdf](#)) to help determine which software needs mitigation.



7.0 Mitigation Availability

- Check our [website \(https://software.intel.com/en-us/side-channel-security-support\)](https://software.intel.com/en-us/side-channel-security-support) for the latest status on both kernel and compiler mitigations.
- See details on kernels with mitigations.
- See also distribution compilers containing retpoline support.