

Research teams confirm “Meltdown” and “Spectre” attacks

Computer systems worldwide exposed to data theft due to CPU design flaws

Mike Ingram
12 January 2018

Research teams have confirmed reports of two attacks, Meltdown and Spectre, that exploit significant flaws in the design of the Central Processing Unit (CPU) contained in all modern computer systems.

Speculation following a January 2 report in the *Register* prompted researchers to go public January 3, ahead of the original January 9 scheduled date coordinated with Intel, AMD and other chip manufacturers as part of the responsible disclosure process which allows time for fixes to be made available before an exploit is publicly exposed.

The flaws were discovered and reported to the chip manufacturers last year by Google’s Project Zero team’s Jann Horn and others when they demonstrated attacks that could take advantage of “speculative execution,” a technique used by most modern CPUs to optimize performance by essentially guessing what executions a given process will require.

Meltdown and Spectre are distinct exploits, but both use side channels to obtain information (including secrets) from an accessed memory location. Side channel attacks are any attack based on information gained from the physical implementation of a computer system, rather than weaknesses in the implementation itself, e.g., software bugs. Both Meltdown and Spectre exploit side effects of the design of computer systems, specifically the CPU.

Meltdown

In the white paper on Meltdown, the researchers point out that memory ensures that the memory assigned to one user’s applications cannot be accessed by another. Memory isolation also prevents user applications from reading or writing kernel memory, which is the space reserved for the operating system. “This isolation is a cornerstone of our computing environments and allows running multiple applications on personal devices or executing processes of multiple users on a single machine in the cloud,” the researchers wrote.

The authors explain that a Meltdown attack “allows overcoming memory isolation completely by providing a simple way for any user process to read the entire kernel memory of the machine it

executes on, including all physical memory mapped in the kernel region.”

Meltdown is not an exploit of a software vulnerability and therefore works on all major operating systems. Meltdown exploits a so-called side-channel, or unintentional flow of information, available on most modern processors. The authors specifically cite modern Intel microarchitectures since 2010, but say it could potentially exploit other CPUs of other vendors.

Meltdown exploits the “out-of-order execution” feature of modern processors. Out-of-order execution is used to overcome latencies, or lags, of busy executions. For example, if a task requires information to be fetched from memory to provide information to other tasks in the execution, the processor will “look ahead” and schedule subsequent operations to idle execution units. An execution unit is a part of the CPU that performs the operations and calculations as instructed by the computer program.

The research team developed a sample application as a proof of concept that was able to successfully access secrets stored in memory when executed against Intel chips. For ARM and AMD CPUs, the team did not manage to successfully leak kernel memory, but the authors caution, “The reasons for this can be manifold. First of all, our implementation might simply be too slow and a more optimized version might succeed.”

Spectre

While Meltdown has only been verified against Intel chips, Spectre is known to affect Intel, Apple, ARM, and AMD processors and works by tricking processors into executing instructions they should not have been able to, granting access to sensitive information in other applications’ memory space. In today’s world of cloud computing this has massive implications as it potentially allows access to other customer data within the cloud provider from compromised systems.

A number of methods exist to achieve the isolation of one application or service from others. These include virtualization, where multiple virtual machines run on the same physical environment, and more recently containerization, where an

application and all its dependencies are run inside a container. A fundamental security assumption underpinning these methods is that the CPU will faithfully execute software, including its safety checks.

The authors note that “Speculative execution unfortunately violates this assumption in ways that allow adversaries to violate the secrecy (but not integrity) of memory and register contents.” As a result, Spectre impacts a broad range of computer systems, whether they are running on dedicated hardware or inside containers or virtual environments. This has massive implications as it impacts public clouds services such as Amazon and Google, where companies can rent virtual machines rather than building their own data centers.

Like out-of-order execution, speculative execution is designed to optimize the CPU for speed. By speculating or guessing what tasks may be needed to complete a job the processor can perform work before it is actually needed. If it turns out the guess was wrong and the work was not needed after all, most changes made by the work are reverted and the results are ignored.

Spectre attacks trick the processor into executing instructions sequences that should not have executed during correct program execution. These are known as transient instructions, as their effects on the given state of the CPU will eventually be reverted. “By carefully choosing which transient instructions are speculatively executed, we are able to leak information from within the victim’s memory address space,” the researchers note.

Intel, by far the biggest chip manufacture with over 80 percent of market share against 15 percent for AMD and 2 percent for all other chip makers, has come under particular scrutiny following reports of the security flaws. CEO Brian Krzanich sold \$24 million worth of stock in November, months after the company was informed of vulnerabilities in its chips, but before it was publicly disclosed. Google informed Intel of the vulnerabilities in June 2017.

A *Business Insider* report of January 3 cited a statement from an Intel representative that “Krzanich’s sale had nothing to do with the newly disclosed chip vulnerability and was done as part of a standard stock-sale plan.”

The article notes: “To avoid charges of trading on insider knowledge, executives often put in place plans that automatically sell a portion of their stock holdings or exercise some of their options on a predetermined schedule, typically referred to as Rule 10b5-1(c) trading plans. According to an SEC filing, the holdings that Krzanich sold in November—245,743 shares of stock he owned outright and 644,135 shares he got from exercising his options—were divested under just such a trading plan.

“But Krzanich put that plan in place only on October 30, according to the filing. The representative said his decision to set up that plan was ‘unrelated’ to information about the security vulnerability. Still, the timeline raises questions.”

The latest generation of Intel’s Core ix processors, Coffee Lake, was made available to desktop computers on October 5, despite Intel being fully aware of the inherent flaws.

While there are software fixes that can be applied at the operating system level to mediate the impact of the security flaws, it is estimated that these could slow devices by anywhere from 5 to

30 percent, dependent on workload.

Three separate class-action lawsuits have been filed by plaintiffs in California, Oregon and Indiana seeking compensation. All three cite the security vulnerability and Intel’s delay in public disclosure and the alleged computer slowdown that will be caused by the fixes needed to address the security concerns. Intel disputed this claim in an earlier statement that flippantly declared, “Contrary to some reports, any performance impacts are workload-dependent, and, for the average computer user, should not be significant and will be mitigated over time.”

Linus Torvalds, the inventor of the open source Linux operating system in 1991 and still in charge of Linux kernel development, posted a sharply worded email to the Linux list on January 3 stating:

“I think somebody inside of Intel needs to really take a long hard look at their CPU’s, and actually admit that they have issues instead of writing PR blurbs that say that everything works as designed.”

Torvalds added, “Or is Intel basically saying ‘we are committed to selling you shit forever and ever, and never fixing anything’?”

Intel and other chip manufacturers, as well as software vendors such as Apple and Microsoft, have emphasized that there is no evidence that the flaws have been exploited by hackers. However, unlike previous software exploits, the nature of these attacks is that they do not leave any fingerprints such as log file entries. The reality is that it is impossible to know if these exploits have been used to date or not.

Due to the critical role played by computer technology—and the vast amounts of information processed, stored and transmitted through computer systems—exploits such as Meltdown and Spectre, as well numerous data breaches reported by companies like Ebay and Equifax, pose a massive threat to millions of people across the planet.

Decisions that affect the security and privacy of the world’s population cannot be left in the hands of a few massive corporations motivated only by profits and share values. The entire information systems infrastructure—from chip manufacturing to cloud hosting providers such as Amazon, Google and Microsoft—must be taken under social ownership and reorganized to satisfy the needs of society and not the private wealth accumulation of corporate billionaires.



To contact the WSWS and the Socialist Equality Party visit:

wsws.org/contact