the Availability Digest

www.availabilitydigest.com
@availabilitydig

Meltdown and Spectre Security Threats

January 2018

A new class of security threats has just reared its ugly head. These threats, known as Meltdown and Spectre, exploit a common feature of modern microprocessors in order to steal sensitive data from memory. This feature is known as 'speculative execution.'

What is Speculative Execution?

In order to speed up application processing, today's CPUs employ speculative execution. When a program reaches a conditional branch (such as 'If FLAG is true, do Process A; else do Process B'), the CPU decides in advance which branch to take. A 'branch predictor' determines the likely result of the condition and executes the branch that is most likely to run before the test is actually completed.

If the guess is correct, the chip appears to be running faster than had it awaited the results of the test. If the guess is wrong, the chip has to throw away any speculative results. Branch predictions are often 99% accurate.

The Security Risk

The problem is that 'throwing away speculative results' means that these results are left in cache memory. Meltdown and Spectre, and hackers in general, can view what happened in the speculative window by hacking into cache memory.

As a consequence, hackers can manipulate the system. They can steer the behaviour of branch predictors to cause code to run speculatively that shouldn't have run at all. By tricking the code to take a wrong branch, the hackers have access to results that they never should have seen.

Meltdown and Spectre can be especially dangerous since they do not leave any traces in traditional log files.

Meltdown

Meltdown and Spectre use somewhat different techniques for spoofing the CPU to load data into cache memory from which it can be hacked. Meltdown relies upon a common practice that separates loading data into memory from the process of checking permissions. It breaks the mechanism that keeps applications from accessing arbitrary memory.

Meltdown fools the chip into loading privileged data into cache during a speculation window. It then uses a specially constructed code sequence to read this data. Meltdown can read any portion of physical memory, including kernel memory.

Spectre

Spectre tricks applications into accessing arbitrary locations in memory. When speculatively executing code, the chip may load some data that is used later to locate a second piece of data. The chip might then speculatively load the second piece of data into cache. Spectre can access this data by hacking into cache memory.

Spectre can only read memory from the current process. It cannot read kernel memory or other memory.

What is Being Done to Mitigate These Threats?

Intel and others were informed about the flaws back in June, 2017. Intel initially dismissed the importance of the threats. However, in the space of a week, it went from '"should not be significant" to "may be initially higher" to "significant."

Intel and the other companies that had been informed agreed that the vulnerability disclosure be coordinated. If the disclosure were too early, the vulnerabilities could be exploited by hackers. If the disclosure were too late, the exposures might be too big from which to recover. It was agreed to withhold information on Meltdown and Spectre until January 10, 2018, since the remediation and patches had to be available when the threats were announced. This required a coordinated vulnerability disclosure on a scale and scope rarely seen.

The vulnerabilities stayed a relative secret for almost six months. However, the news broke a week early. A message posted by an AMD software engineer revealed that there were software patches for Meltdown and Spectre for Linux kernels. This required that the details of the vulnerabilities and the patches be rushed out.

Fortunately, patches and mitigations had already been released by most major vendors. Unfortunately, vendors are still in the early days of fully understanding the vulnerabilities and issuing patches for them.

All of the patches to correct these vulnerabilities impose performance hits on the chip since they reign in speculative execution to some extent. None of the patches fix the basic CPU flaw – the reading of data into cache as a result of speculative execution. However, they act to prevent practical exploitations of the flaw.

Summary

Until some basic changes are made to CPU chips and the operating systems that run on them, Meltdown and Spectre will continue to pose threats to the stealing of sensitive data from cache. Vendors are hard at work to correct this flaw, but it may remain with us for the foreseeable future.

Acknowledgements

Information for this article was taken from the following sources:

Meltdown and Spectre patches and mitigation released, *Tech Target*; January 4, 2018.
What are Meltdown and Spectre? Here's what you need to know, *Red Hat*; January 5, 2018.
Huge coordinated vulnerability disclosure needed for Meltdown, *Tech Target*; January 5, 2018.
An Explanation of the Meltdown/Spectre Bugs for a Non-Technical Audience, *Cloud Fare*; January 8, 2018.
Intel needs to come clean about Meltdown and Spectre, *The Verge*; January 10, 2018.