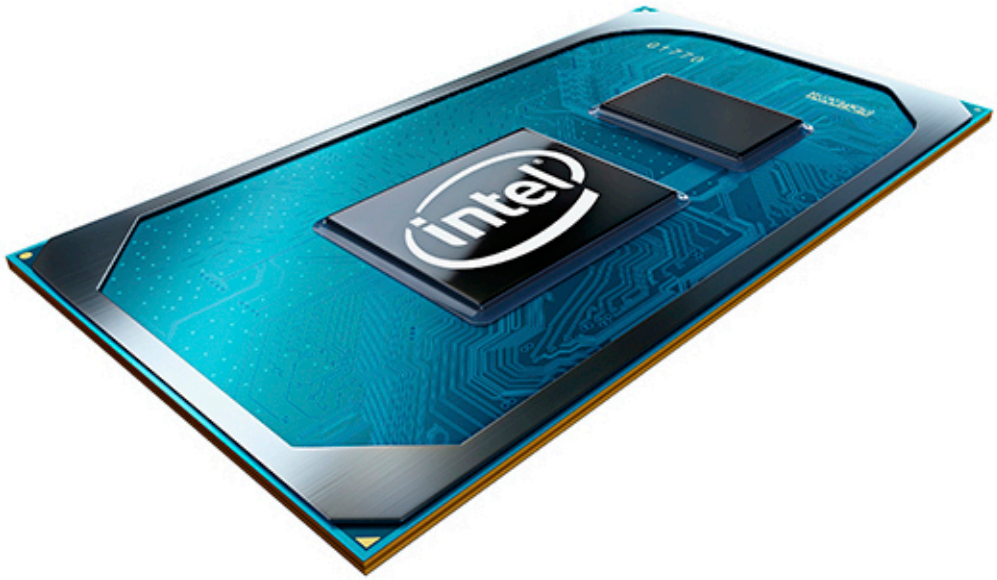
Intel will soon bake anti-malware defenses directly into its CPUs



Enlarge / A mobile PC processor code-named Tiger Lake. It will be the first CPU to offer a security capability known as Control-Flow Enforcement Technology.

reader comments

71 with 49 posters participating

Share this story

The history of hacking has largely been a back-and-forth game, with attackers devising a technique to breach a system, defenders constructing a countermeasure that prevents the technique, and hackers devising a new way to bypass system security. On Monday, Intel is announcing its plans to bake a new parry directly into its CPUs that’s designed to thwart software exploits that execute malicious code on vulnerable

computers.

Control-Flow Enforcement Technology, or CET, represents a fundamental change in the way processors execute instructions from applications such as Web browsers, email clients, or PDF readers. Jointly developed by Intel and Microsoft, CET is designed to thwart a technique known as return-oriented programming, which hackers use to bypass anti-exploit measures software developers introduced about a decade ago. While Intel first published its implementation of CET in 2016, the company on Monday is saying that its Tiger Lake CPU microarchitecture will be the first to include it.

ROP, as return-oriented programming is usually called, was software exploiters' response to protections such as Executable Space Protection and address space layout randomization, which made their way into Windows, macOS, and Linux a little less than two decades ago. These defenses were designed to significantly lessen the damage software exploits could inflict by introducing changes to system memory that prevented the execution of malicious code. Even when successfully targeting a buffer overflow or other vulnerability, the exploit resulted only in a system or application crash, rather than a fatal system compromise.

ROP allowed attackers to regain the high ground. Rather than using malicious code written by the attacker, ROP attacks repurpose functions that benign applications or OS routines have already placed into a region of memory known as the stack. The "return" in ROP refers to use of the RET instruction that's central to reordering the code flow.

Very effective

Alex Ionescu, a veteran Windows security expert and VP of engineering at security firm CrowdStrike, likes to say that if a benign program is like a building made of Lego bricks that were built in a specific sequence, ROP uses the same Lego

pieces but in a different order. In so doing, ROP converts the building into a spaceship. The technique is able to bypass the anti-malware defenses because it uses memory-resident code that's already permitted to be executed.

CET introduces changes in the CPU that create a new stack called the control stack. This stack can't be modified by attackers and doesn't store any data. It stores the return addresses of the Lego bricks that are already in the stack. Because of this, even if an attacker has corrupted a return address in the data stack, the control stack retains the correct return address. The processor can detect this and halt execution.

"Because there is no effective software mitigation against ROP, CET will be very effective at detecting and stopping this class of vulnerability," Ionescu told me. "Previously, operating systems and security solutions had to guess or infer that ROP had happened, or perform forensic analysis, or detect the second stage payloads/effect of the exploit."

Not that CET is limited to defenses against ROP. CET provides a host of additional protections, some of which thwart exploitation techniques known as jump-oriented programming and call-oriented programming, to name just two. ROP, however, is among the most interesting aspects of CET.

Those who do not remember the past

Intel has built other security functions into its CPUs with less-than-stellar results. One is Intel's SGX, short for Software Guard eXtension, which is supposed to carve out impenetrable chunks of protected memory for security-sensitive functions such as the creation of cryptographic keys. Another security add-on from Intel is known as the Converged Security and Management Engine, or simply the Management Engine. It's a subsystem inside Intel CPUs and chipsets that implements a host of sensitive functions, among them the firmware-based

Trusted Platform Module used for silicon-based encryption, authentication of UEFI BIOS firmware, and the Microsoft System Guard and BitLocker.

A steady stream of security flaws discovered in both CPU-resident features, however, has made them vulnerable to a variety of attacks over the years. The most recent SGX vulnerabilities were disclosed just last week.

It's tempting to think that CET will be similarly easy to defeat, or worse, will expose users to hacks that wouldn't be possible if the protection hadn't been added. But Joseph Fitzpatrick, a hardware hacker and a researcher at SecuringHardware.com, says he's optimistic CET will perform better. He explained:

> *One distinct difference that makes me less skeptical of this type of feature versus something like SGX or ME is that both of those are "adding on" security features, as opposed to hardening existing features. ME basically added a management layer outside the operating system. SGX adds operating modes that theoretically shouldn't be able to be manipulated by a malicious or compromised operating system. CET merely adds mechanisms to prevent normal operation—returning to addresses off the stack and jumping in and out of the wrong places in code—from completing successfully. Failure of CET to do its job only allows normal operation. It doesn't grant the attacker access to more capabilities.*

Once CET-capable CPUs are available, the protection will work only when the processor is running an operating system with the necessary support. Windows 10 Version 2004 released last month provides that support. Intel still isn't saying when Tiger Lake CPUs will be released. While the protection could give defenders an important new tool, Ionescu and fellow researcher Yarden Shafir have already devised bypasses for it. Expect them to end up in real-world attacks within the decade.

https://arstechnica.com/?p=1683807