Computer scientists discover gap in the latest security mechanisms used by some chips

April 4 2024, by Daniel Meierhans

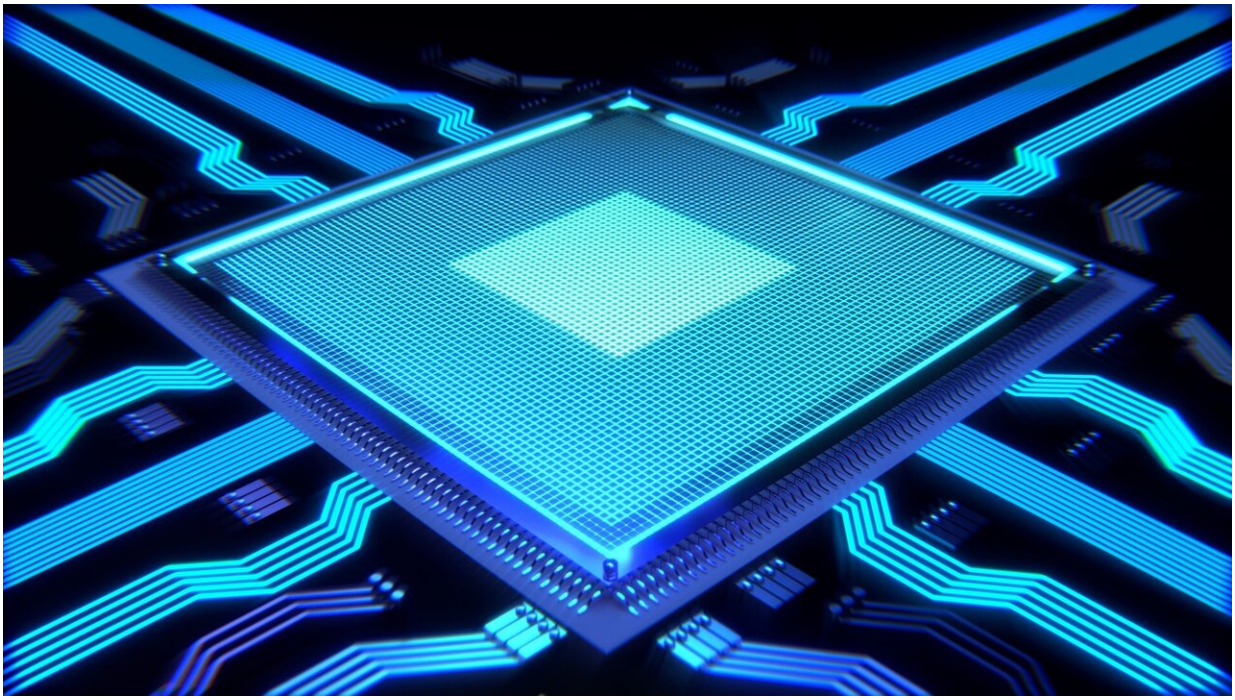

Credit: CC0 Public Domain

Over the past few years, hardware manufacturers have developed technologies that ought to make it possible for companies and governmental organizations to process sensitive data securely using shared cloud computing resources. Known as confidential computing, this approach protects sensitive data while it is being processed by isolating it in an area that is impenetrable to other users and even to the

cloud provider. But computer scientists at ETH Zurich have now proven that it is possible for hackers to gain access to these systems and to the data stored in them.

The researchers ran two attack scenarios, both using what's known as the interrupt mechanism, which temporarily disrupts regular processing—for instance to prioritize a different computing task. There are a total of 256 different interrupts, and each one triggers a specific sequence of programming commands.

"Interrupts are a marginal concern, and it appears that ensuring they have systematic safeguards in place has simply been overlooked," says Shweta Shinde, Professor of Computer Science at ETH Zurich. Together with her Secure & Trustworthy Systems Group, Shinde identified the problematic vulnerabilities in the server hardware used by two leading manufacturers of computer chips, AMD and Intel.

Eavesdrop-proof smartphone project helps find the gaps

Shinde's team uncovered the security gaps while examining the confidential computing technologies used in AMD and Intel processors. The researchers wanted to gain an in-depth understanding of how these processors function because they are working on an eavesdrop-proof smartphone based on confidential computing.

At the core of confidential computing is the trusted execution environment (TEE). The TEE is a hardware-based component that isolates applications while they are being run. Accessing the application memory is then possible only with an authorized code. This means the data is also protected from unauthorized access while it is being stored, unencrypted, in the working memory during processing. In the past, the

only way to ensure such protection was to encrypt data while stored on the hard drive and during transmission.

Instability factor number one: Hypervisors

In the public cloud, applications are isolated using a TEE, specifically from what's known as a hypervisor. Cloud providers use hypervisor software to manage resources ranging from hardware components to their customers' virtual servers. Hypervisors are an important part of cloud services because they create the required flexibility, efficiency and security. In addition to managing and optimizing how the underlying hardware is used, they ensure that different users can work securely in separate areas of the same cloud without disturbing each other.

But the administrative functions hypervisors perform are also an instability factor as they open up a variety of attacks. Under certain conditions, these attacks can make it possible to access data stored in the memories of other active cloud users working with the same hardware. Moreover, cloud providers could also use hypervisors to take a peek at their users' data themselves.

Both these risks are unacceptable to companies and governmental organizations that process sensitive data. Indeed, in an expert report compiled by the Swiss Federal Council, which examined the legal framework for implementing Switzerland's cloud strategy, unauthorized access to what's referred to as "data in use" was rated as the most probable risk associated with using a public cloud.

Fully isolating the hypervisor is impossible

There are, however, fundamental limitations as to how well a user system can be isolated and protected from the hypervisor. After all,

some communication must take place between the two, and as an administrative tool, the hypervisor still has to be able to perform its core tasks. These include allocating cloud resources and managing the virtual server running the secured system in the cloud.

One of the remaining interfaces between the hypervisor and the TEE concerns the management of interrupts. The ETH team launched what are known as Ahoi attacks to exploit the hypervisor as a means of sending coordinated interrupts to the secured system at any time. This exposes the gap in security: instead of blocking the request from the untrustworthy hypervisor, the TEE lets certain interrupts through. Unaware that these interrupts are coming from outside, the system runs its usual programming routines.

Interrupt heckles knock security off its game

By sending coordinated interrupt heckles, the ETH scientists managed to confuse a TEE-secured system so effectively that they were able to gain root access—in other words, take full control.

"Most affected by this problem was AMD's confidential computing, which proved vulnerable to attack from several different interrupts. In the case of Intel, only one interrupt door had been left open," Shinde says in summarizing the results of her "Heckler attack." The researchers also rated AMD's previous means of defense as insufficient. The chip manufacturers have since taken steps to address this.

The second attack scenario, known as WeSee, affects AMD hardware only. It exploits a mechanism that the chip manufacturer introduced to make communication between TEE and hypervisor easier despite isolation. In this case, a special interrupt can cause the secured system to divulge sensitive data and even run external programs.

Byproduct on the path to user control of phones

As important as it is to find gaps in the security for sensitive data stored in the public cloud, for Shinde and her research group this was merely a byproduct on the path to ensuring that users of iPhones and Android smartphones retain full control over their data and applications. A specially designed TEE will do more than make sure user data is protected from eavesdropping by the manufacturer's operating system.

"We also want our TEE to support unmonitored operation of those apps not managed by Apple or Google," Shinde says.

More information: Benedict Schlüter et al, Heckler: Breaking Confidential VMs with Malicious Interrupts (2024). In: 33rd USENIX Security Symposium (USENIX Security), August 14-16, 2024

Benedict Schlüter et al, WeSee: Using Malicious #VC Interrupts to Break AMD SEV-SNP (2024). In: 45th IEEE Symposium on Security and Privacy (IEEE S&P), May 20-23, 2024.

Provided by ETH Zurich