Invisible Things Lab presents another attack on Intel® Trusted Execution Technology.

December 21, 2009 — Invisible Things Lab publishes today a detailed research paper describing a full circumvention attack against Intel® Trusted Execution Technology (TXT). This is a different attack against Intel TXT, than the one presented by the same researchers in February at Black Hat DC.

What is Intel Trusted Execution Technology?

Intel Trusted Execution Technology (TXT), formerly code named LaGrande, is currently part of the Intel® vPro™ brand and is a key component of the Intel's Safer Computing Initiative. Intel TXT comprises a set of extensions to the CPU and to the chipset, and also makes extensive use of the Trusted Platform Module 1.2 (TPM). TXT goal is to provide a mechanism for safe loading of system software, e.g. a hypervisor or a kernel.

Our new attack

We again showed that an attacker can compromise the integrity of a software loaded via an Intel® TXT-based loader in a generic way, fully circumventing any protection TXT is supposed to provide. As usual we have created a proof-of-concept code that demonstrates the successful attack.

This time our attack exploits an implementation error in the so called SINIT Authenticated Code modules. Intel distributes the SINIT modules for each of its TXT-capable chipset. Currently those modules are made available as part of the tboot project -- Intel's reference implementation of TXT-based loader for Xen and Linux. In the future those modules are expected to be distributed by OEMs as part of the BIOS firmware.

SINIT modules are digitally signed and cannot be modified. The SINIT modules are executed by the SENTER instruction (the central CPU instruction for TXT operation). Thus, they are somewhat similar to microcode, although they are written in a regular x86 assembly.

We should stress that our attack is software-only and does not require physical presence.

Relation to our previous TXT attack

Earlier this year our team has presented an attack against Intel TXT that exploited a design problem with SMM mode being over-privileged on PC platforms and being able to interfere with the SENTER instruction.

There is no relation between the two attacks. They are totally independent, each exploiting a different problem in TXT infrastructure. Both attacks, however, allow to achieve the same goal.

Working with the vendor

We have informed Intel about the SINIT implementation error, together with description how it could be exploited by an attacker to circumvent TXT secure launch, on September 30, 2009. Intel has confirmed the vulnerability shortly afterwards, and we agreed to withhold the publication of the paper describing the attack until Intel fixes the problem and publishes updated SINIT modules and a security advisory.

Affected products

Intel TXT is a very new technology. The first desktop machines supporting TXT started appearing only at the end of 2007, while laptops with TXT support have not been available before 2009. Consequently very few products use TXT. One example is the open source Xen hypervisor, that uses the earlier mentioned tboot for implementing TXT-supported boot. Recently also the Linux kernel got support to be loaded via tboot.

There is also at least one upcoming commercial product, the Citrix's XenClient, that is said to make use of Intel VT-d and TXT technologies.

Seriousness of the attack

We should stress, that if the virtualization system is well designed, a buggy TXT doesn't automatically render the system vulnerable to an attack. Rather, in normal circumstances, the attacker would still need to bypass some system-imposed protections first (e.g. exploit a potential bug in the hypervisor or a backend driver). The difference that TXT makes is that in the event of such an attack, if the attacker tried to introduce any permanent changes to the system, e.g. subvert the system binaries on disk, TXT would be able to prevent such a subverted system from booting the next time.

However, there exist scenarios where TXT attack can be fatal and no other bug in the system software is required to undermine security of the platform.

One such case is when the system's goal is to contain a potentially malicious user, e.g. by forbidding the user to connect a USB device to the "corporate" virtual machine in order to prevent (intentional or not) information leaks. TXT provides mechanisms to allow the IT department to implement such protections very effectively with the help of Remote Attestation feature of TPM.

For the above scenario TXT is absolutely necessary, and if the attacker can bypass TXT secure launch, e.g. using our exploit, the attacker can circumvent such restrictions at will. And this all using an extremely cheap software-only attack[1].

Another scenario where TXT circumvention might be fatal is full disk encryption, especially in case of laptops. It's widely known that a full disk encryption scheme that is not based on a trusted boot scheme is subject to simple attack where the attacker can easily subvert the boot loader in order to capture the user passphrase. Later the attacker can steal the encrypted laptop and will know the passphrase needed to decrypt it. Such attacks have been demonstrated in practice, incidentally also by our team (see the Evil Maid Attack).

FAQ

Q: Intel TXT seems to be quite buggy. Is there any real benefit for the industry to adopt it as a building block for future systems?

Yes, definitely. TXT provides unique features that should allow to create more secure systems in the future, exploiting advanced architectures that could allow to e.g. minimize system attack surface. Our team believes that the research we do can enable safer systems in the future, by eliminating implementation errors from TXT before its wide adoption.

Q: What about other vPro™ technologies?

Intel vPro™ brand comprises several technologies: Intel VT-x, VT-d, TXT, and AMT. The research described in this press release affects only the TXT technology. We have presented a separate research earlier this year, at Black Hat USA in Las Vegas in July, about attacks on Intel AMT technology (titled "Ring -3 Rootkits").

Q: Instead of looking for bugs in low-level system components, why don't you focus on finding bugs in Web browsers or other applications that people use? Wouldn't that be more important for protecting ordinary users?

We believe that the process of finding, and then patching, bugs in applications like Web browsers, is an endless arm-race and doesn't lead to more secure systems in the long run. We believe that really secure systems need to be build system on solid foundations, exploiting the Security by Isolation principle, and that's why we look at those low-level technologies currently.

Q: Which Intel chipsets are affected?

We test our proof-of-concept attack on Intel Q35- and Q45-based chipsets, but it's very likely that all other TXT-capable chipsets available currently on the market are also affected. Please see the Intel official advisory for the details on affected platforms.

---

[1] Of course, the person that is in a physical possession of the machine, e.g. laptop, can theoretically, always gain full control over the software executing on this machine. In particular, such an attacker can e.g. replace the processor with a malicious processor that would allow for certain backdoors (e.g. ring3 to ring0 escalation), or can retrieve the secrets stored in the TPM using electron microscope, or can perform active attacks on the LPC bug in order to reset the PCR17 and PCR18 registers without executing the SENTER instruction, or can replace the DRAM chips with ones that would record the contents of the memory onto external device. Such physical attacks are however considered very expensive to perform, often much more expensive than the data that are supposed to be protected by such systems.

About the research authors

Rafal Wojtczuk, Principle Researcher, specializes primarily in kernel and virtualization security. Over the years he has disclosed many security vulnerabilities in popular operating system kernels (Linux®, SELinux, *BSD, Windows™), virtualization software (Xen®, VMWare® and Microsoft® virtualization products), and low-level system technologies (Intel TXT, Intel AMT, Intel BIOS). He is also well known for his articles on advanced exploitation techniques, including novel methods for exploiting buffer overflows in partially randomized address space environments. Rafal holds a Master's Degree in Computer Science from University of Warsaw. He is based in Warsaw, Poland.

Joanna Rutkowska, founder and CEO of Invisible Things Lab, leads a team of researchers who focus on system-level security. This includes kernel, hypervisor, chipset and CPU security issues. The recent achievements of the team include: bypassing Intel TXT, attacks on SMM, Intel AMT and BIOS, and demonstration of practical Xen hypervisor compromises. She is also known for writing Blue Pill -- the first virtualization-based rootkit with nested hypervisors support, and also for her work on various kernel mode malware for Windows and Linux. Joanna holds a Master's Degree in Computer Science from Warsaw University of Technology. She is based in Warsaw, Poland.

Alexander Tereshkin, Principal Researcher, is an experienced reverse engineer and an expert into Windows® kernel and hardware virtualization, specializing in rootkit technology and kernel exploitation. He is known for his research on sophisticated ideas for novel rootkit creation and personal firewall bypassing in the past years. Recently he has done significant work in the field of virtualization based malware and Microsoft® Vista™ kernel security. He is a co-author of "Understanding Stealth Malware" course. Alex holds the Russian equivalent of a Master's Degree in Applied Mathematics, and also the Russian equivalent of a PhD degree in Information Security from Taganrog State University Of Radioengineering (Southern Federal University).

About Invisible Things Lab

Invisible Things Lab focuses on cutting-edge research in computer security, specializing in system-level security. We are well known for our pioneering research in the areas of kernel security, virtualization security and system/firmware-level security. In particular we were the first to disclose vulnerabilities in such low-level system elements as Intel chipsets, BIOS, and TXT. Our work has been widely quoted by international press and the members of our team often speak at industry conferences around the world.

Contact

For press inquiries, Invisible Things Lab can be contacted via email:
`contact@invisiblethingslab.com`

Relevant Links

- http://www.intel.com/technology/security/
- http://www.intel.com/technology/vpro/index.htm
- http://security-center.intel.com/advisory.aspx?intelid=INTEL-SA-00021&languageid=en-fr
- http://community.citrix.com/download/attachments/100303689/CitrixXenClient_SolutionBrief.pdf?version=1
- http://invisiblethingslab.com/