WHITEPAPER
Intel Security Technologies
Public Sector

Improving Computer Security in the Public Sector with Intel Technology

Intel security technologies help boost the protection of data, applications, workloads, and the system during its boot sequence.

Table of Contents

Minimizing the Attack Surface

As cyberthreats escalate, it is critical for public sector organizations to adopt advanced security technologies to better defend computing systems. Although no computer system can be absolutely secure, Intel security technologies provide a comprehensive suite of capabilities that can dramatically reduce a computer's attack surface area. These technologies help mitigate attack vulnerabilities across a computing platform, as described in the following and shown in Figure 1:

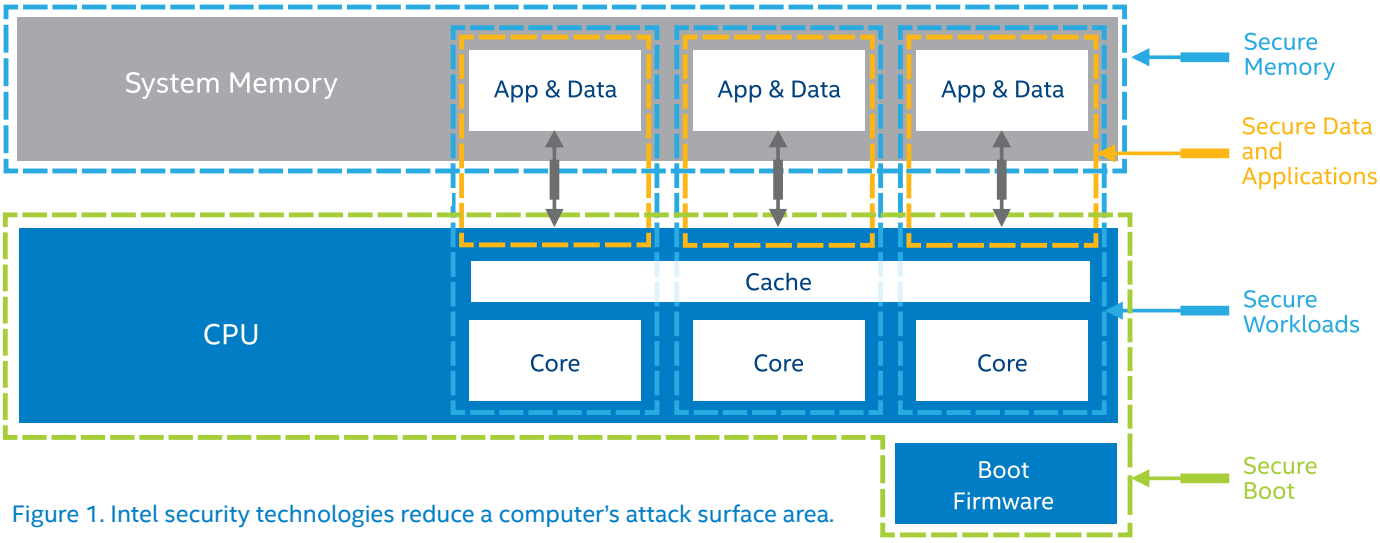

Figure 1. Intel security technologies reduce a computer's attack surface area.

Protection	Threat Example
Secure Boot	Attackers inject malware into or maliciously modify the firmware, which loads before the computer's security defenses are activated.
Secure Memory	Attackers snoop or hijack the external memory bus in order to steal or overwrite data stored in external memory.
Secure Data and Applications	An attacker or malware gains access and control of the computing platform, applications, and data.
Secure Workloads	Attackers, malware, or rogue applications sabotage mission-critical applications by overconsuming or snooping on computing resources (e.g., cache, memory bandwidth).

Intel Security Technologies

How different types of malware operate may vary, but they all seek to corrupt systems, disrupt organizations, steal data, and seize control of platforms.[1] Working to thwart malware in its tracks, Intel is continuously enhancing built-in, hardware-based security capabilities that, when combined with software security products, enable a trusted platform. This white paper describes four Intel Security Technologies that are among many others available on Intel® processor-based computing platforms.

Secure Boot

No single security solution can completely protect a computing system, and that is why a multi-layer security approach is a necessity. But like a sound building that requires a solid foundation, multi-layer protection must be built on a secure foundation, starting with a trusted boot sequence.

Malware and Rootkits

The objective of some malware, like rootkits, is to acquire control of the computer platform by installing itself underneath the operating system or hypervisor during the boot sequence. "Rootkits gain root-level, privileged access to a computer while hiding their existence and actions. Hackers use rootkits to conceal themselves until they decide to execute their malicious malware," according to Gilad David Maayan, who heads Agile SEO, a leading tech marketing agency.[2]

Root of Trust

Providing a foundation for safer computing, Intel combines two powerful boot sequence controls: Intel® Boot Guard and Intel® Trusted Execution Technology (Intel® TXT). This combination, called Intel® Converged BootGuard and Intel TXT, (Intel® CBnT), offers the following protection.

- Intel Boot Guard
  This technology determines whether the firmware booting the platform can be trusted or was inappropriately modified. With Intel Boot Guard, the chipset vendor creates a digital signature for the firmware that must be validated before the boot sequence can be completed.
- Intel TXT
  Offering control beyond the firmware, Intel TXT creates a cryptography-unique identifier for each launch-enabled component and terminates the launch of code that does not match approved code, as shown in Figure 2. The technology establishes a measured launch environment (MLE) that compares all critical launch environment elements against a known good source.

For example, the MLE would detect the presence of the "Blue Pill" rootkit, which installs a renegade hypervisor to control computer resources. An Intel TXT-enabled system would then log the cause of the failure by writing to a non-volatile register before shutting down.

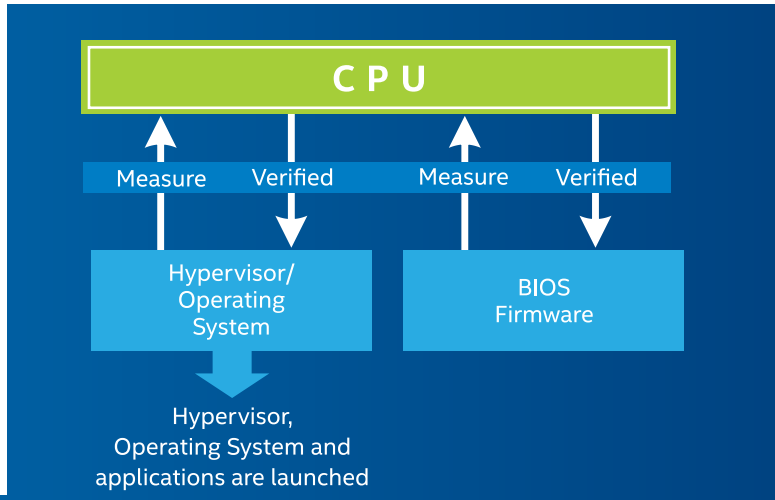

Figure 2. Intel security technologies enable a trusted launch.

Other Benefits

- Following a shutdown due to an unverified launch, residual data is removed, protecting it from memory-snooping software and reset attacks.
- Platform measurement credentials can be reported in support of an organization-wide trust verification process.
- Users can create enforceable lists of "known good" or approved executable code.
- Secure boot is standards-based (e.g., Unified Extensible Firmware Interface (UEFI) and AES).

Secure Memory

Confidential and secret data can be compromised when a computer's memory devices are not secure.

Data Theft

In a computer system, the memory bus transmits data between the CPU and memory devices, like DRAM and FLASH. Attackers can read unprotected memory devices by monitoring the external memory bus (called snooping), or in the case of non-volatile FLASH memory devices, removing and accessing them directly.

Data Encryption

Ensuring data privacy, Intel® Total Memory Encryption (Intel® TME)[3] encrypts the platform's entire physical memory. Enabled in the BIOS, Intel TME helps ensure all memory accessed by the CPU is encrypted, including customer credentials, encryption keys, intellectual property, or personal information transmitted on the external memory bus. This capability is typically enabled during the early stages of the boot sequence, and once configured and locked, will encrypt physical memory using the National Institute of Standards and Technology (NIST) standard AES-XTS algorithm with 128-bit keys, as shown in Figure 3.

The encryption key is created during every boot sequence by a hardened random number generator in the CPU and is never exposed to software. Data stored in memory devices and transmitted on external buses is encrypted and is only unencrypted while inside the CPU, allowing existing software to run unmodified.

Other Benefits

- Full system memory encryption is supported.
- Data is encrypted/decrypted on-the-fly when entering and leaving memory, with minimal performance impact.
- No operating system or application code modifications are required.
- Data is secure against a side channel attack with a memory bus sniffer.

Figure 3. Data in physical memory is encrypted using the standard AES-XTS algorithm.

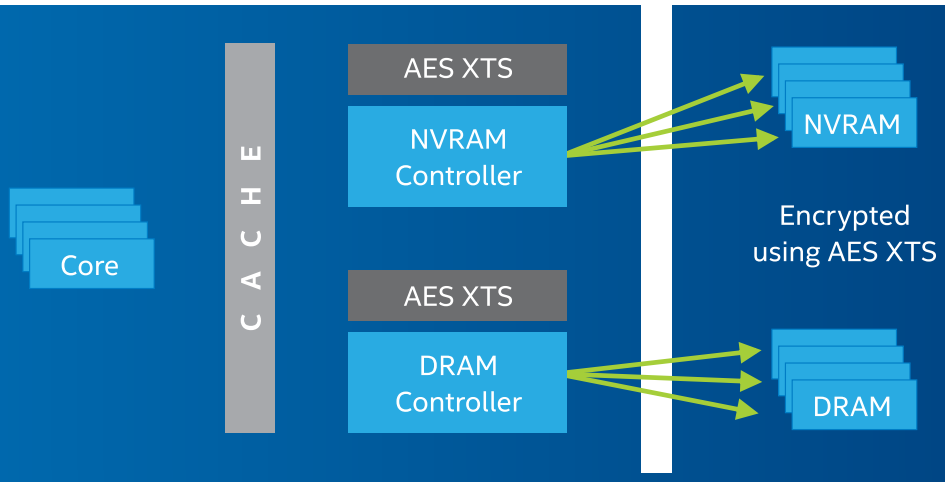

Secure Data and Applications

Government computers may have a combination of classified, confidential, and private data such as personnel deployments, passwords, account numbers, financial information, and health records. A computer's operating system is responsible for enforcing security policies to ensure applications and data are only accessible by authorized individuals, but sometimes malware interferes with this mission.

Malware and Privilege Levels

Typical CPUs support several privilege levels that determine the amount of access and control an application has within the platform. The operating system runs in the highest privileged level; therefore, it has total access to the platform (e.g., CPU and memory). Applications are loaded into less privileged levels so the operating system can fully control them. Not surprisingly, hackers try to find ways to load their malware in the highest privileged level so it will have unrestricted access to the operating system, system resources, and all applications running on the system.

Secure Memory Enclaves

To help protect platforms against malware, even privileged malware, Intel® Software Guard Extensions (Intel® SGX) can be used to partition data and applications into highly protected memory regions, called enclaves. An enclave is a protected area in an application's address space that cannot be read or written by code running outside the enclave environment, regardless of its privilege level. Data within enclaves can only be accessed by code that resides in the enclave.

Intel® SGX can dramatically reduce the attack surface of an application, as shown in Figure 4. This is because malware cannot attack an enclave by compromising the operating system, hypervisor (or virtual machine monitor), or BIOS firmware.

Intel® SGX is a set of special instructions that performs several checks to limit the code segments allowed to enter the address space of an application in an enclave. For example, the enclave cannot be entered through classic function calls, jumps, register manipulation, or stack manipulation. In this way, malware is prevented from accessing code and data in an enclave.

SGX enclave memory is encrypted by the CPU; snooping the memory or connecting the DRAM modules to another system will yield only encrypted data.

Other Benefits

- Software developers can write trusted applications using familiar tools and processes.
- Attestation, the process of demonstrating that a piece of software has been properly instantiated in an enclave on the platform, is supported.
- Enclaves are protected, even when an attacker has physical control of the platform.

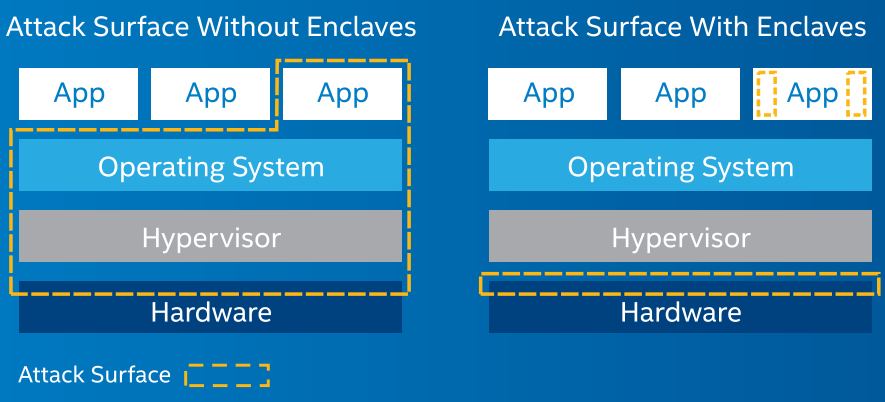

Figure 4. Intel® Software Guard Extensions reduces the attack surface area.

Secure Workloads

National security, safety-critical, and mission applications running on public sector systems need to be protected against malware that attempts to overload computing platform resources, like cache and system memory.

Noisy Neighbor Scenario

When an application overconsumes platform resources, like cache and memory bandwidth, to the performance detriment of other applications, it is called a 'noisy neighbor.' The harm a noisy neighbor causes could be unintentional; or in the case of malware, it could be premeditated, like a DOS attack.

For most CPUs, cache is a resource that is shared among many applications. Application code (and its data) must be loaded into cache before it can run. An application will execute more quickly if all its code remains in cache during execution, but this is less likely when there is a noisy neighbor. A noisy neighbor that consumes large amounts of cache will force the eviction of other applications' code from the cache, resulting in significant performance degradation and potential data leakage.

Shared Resource Monitoring and Allocation

Intel® Resource Director Technology (Intel® RDT) brings new levels of visibility and control over how shared resources cache and memory bandwidth are used by applications, virtual machines (VMs), and containers.

With respect to cache, Intel RDT allows privileged software, like an operating system or hypervisor, to monitor the cache utilization of individual threads, applications, or virtual machines (VMs) and collect information that can be used to:

- characterize workload performance
- enable resource-aware scheduling decisions
- detect noisy neighbors
- aid workload performance debugging

The insights gained by cache monitoring can be used to dedicate regions of cache to critical threads, applications, containers, or VMs (Figure 5), thereby enabling isolation and prioritization of select workloads. In addition, Intel RDT provides separate control over code and data placement in cache.

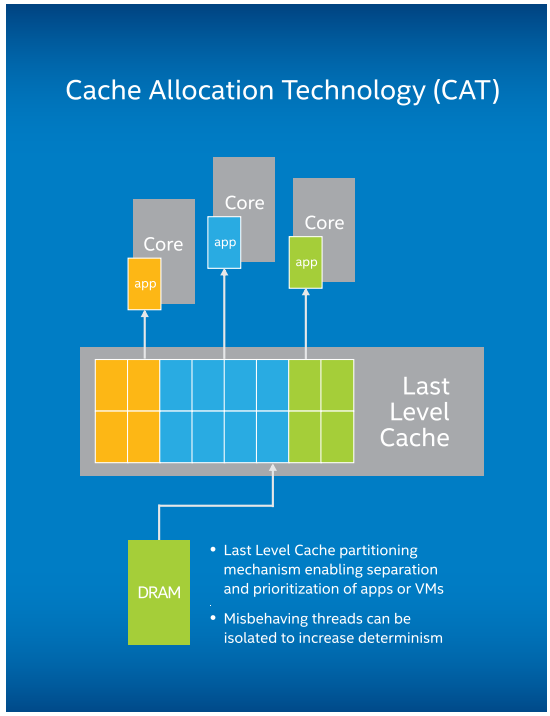

Figure 5. Intel® Resource Director Technology allocates dedicated areas of cache to specific applications[4]

Other Benefits

- Malware cannot commandeer or interfere with the cache and memory bandwidth assigned to priority applications.
- Software developers can increase runtime determinism and workload performance predictability.

- Cache and memory bandwidth contention can be managed.
- Developers can assign platform resources with increased efficiency and flexibility.

Establishing a Root of Trust

With cyberattacks exploiting security vulnerabilities across an entire platform, a software-only security strategy is no longer sufficient. Built-in, silicon-enabled security provides the foundation for a root of trust needed to defend against low-level threats, starting with the CPU and system memory.

Computing systems can establish a root of trust using the Intel security technologies featured in this paper. Intel's products are architected to deliver advanced security to help protect potential attack surfaces and improve computer security in the public sector.

More Information

For more information about the security technologies presented in this white paper, please contact Intel at IOTG-PublicSector@intel.com

1. Bibhu Prasad Biswal, "Intel® TXT," https://www.slideshare.net/BibhuBiswal/intel-trusted-execution-technology-11887933.

2. Gilad David Maayan, "How to Prevent a Rootkit Attack," January 14, 2020, https://blog.malwarebytes.com/how-tos-2/2020/01/how-to-prevent-a-rootkit-attack.

3. Baiju Patel, Intel website, "Intel Releases New Technology Specification for Memory Encryption," Dec 22, 2017, https://software.intel.com/en-us/blogs/2017/12/22/intel-releases-new-technology-specification-for-memory-encryption.

4. Priya Autee, Harpreet Sindhu, "Intel® RDT Hands-on Lab," June 2017, https://www.slideshare.net/MichelleHolley1/intel-rdt-handson-lab.

Intel® technology features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at https://www.intel.com.