

intel security

2026 Intel Platform Security Report

Introduction

From confidential AI to post-quantum cryptography, Intel is not just securing today's workloads—we are shaping the standards for tomorrow's trusted computing.

For seven consecutive years, Intel's annual Product Security Report has set the industry benchmark for transparency and technical leadership in product security assurance. In 2025, Intel's competitive advantage is clearer than ever: no other silicon vendor matches our relentless pace of innovation, our depth of investment in proactive security, or our influence in shaping the future of trusted computing.

Intel's approach is fundamentally different. While competitors are still working to secure the basics, Intel is architecting the future—embedding security at the hardware level, driving global standards, and delivering technologies that anticipate and neutralize tomorrow's threats. Our platforms are not just secure by design—they are engineered to be the foundation of trust for enterprises navigating the complexities of cloud, edge, and AI.

This year's report offers an unfiltered look into Intel's security-first culture and the tangible results of our investments:

- Unmatched innovation in confidential computing and confidential AI, enabling customers to protect data and models throughout the AI lifecycle.
- Leadership in post-quantum cryptography, ensuring resilience against the next generation of cryptographic threats.
- Industry-defining advancements in software robustness, with silicon-level features that eliminate entire classes of vulnerabilities.
- A proven, holistic approach to architectural hardening that integrates threat modeling, formal verification, and negative space testing.
- Open collaboration and standards leadership, ensuring that Intel's breakthroughs benefit the entire ecosystem—not just our customers.

As the only silicon provider consistently delivering proactive security assurance, Intel empowers organizations to innovate with confidence. This seventh annual report reaffirms our commitment to lead, be transparent, and secure the digital future for everyone.

Highlights:

150
MITRE ATT&CK Techniques
Mitigated in Hardware¹

**SOFTWARE
ROBUSTNESS**
Intel platforms deliver
industry-leading
software execution
robustness through
advanced hardware
protection
technologies

PQC
Intel is the first CPU
vendor to implement
Post-Quantum Ready
Cryptography in its
platforms

7x
From 2023 to 2025, AMD
reported 7x more
vulnerabilities in its
hardware Root-of-Trust
firmware than Intel

**CONFIDENTIAL
AI**
Intel sets the benchmark
for AI Security by
integrating Confidential
Computing and Intel®
TDX Connect to deliver
end-to-end protection
for sensitive
AI workloads

77%
From 2023 to 2025,
AMD reported 77%
more vulnerabilities in
their Confidential
Computing firmware
than Intel


x86
Intel is advancing
memory-safety across
the x86 ecosystem by
contributing its
innovative Memory
Tagging Technology

#1
Intel's Product Security
Assurance ranked #1 in the
silicon industry²

10 to 1
In new processors
released since 2023,
external researchers
have found 10
vulnerabilities in AMD
platforms vs 1 in Intel
platforms

Contents

Platform Security	4
Platform Spotlight: Intel vPro®	6
Confidential AI	14
Post-Quantum Cryptography (PQC) Readiness & Compliance	21
Software Robustness	25
Product Security Assurance	31
Advancing Silicon Security Through Security Hardening, Formal Verification, & Negative Space Testing	34
Industry Collaboration & Leadership	40
References	42
















Platform Security


Powered by industry-leading product security assurance, Intel platforms integrate advanced silicon-level defenses and comprehensive security features—enabling enterprises to mitigate current threats and proactively address emerging attack vectors across cloud, edge, and AI environments.

Introduction to Intel Platform Security

A look at some of the security features, technologies, and processes built into Intel platforms

-  **Intel® Trusted Execution Technology (Intel® TXT)**
Enabled by Microsoft/OEM, Intel® TXT uses a Measured Launch Environment (MLE) to boot within a verified trust boundary
-  **Intel® Boot Guard**
Hardware-based boot integrity that meets the Microsoft Windows requirements for UEFI Secure Boot
-  **Intel® BIOS Guard**
Trust boundary to help protect BIOS image updates
-  **Intel® Platform Trust Technology (Intel® PTT)**
Delivers the capabilities of a Trusted Platform Module (TPM) 2.0 without the possibility of attacks on the TPM Low Pin Count (LPC) bus
-  **Intel® Runtime BIOS Resilience (IRBR)**
-  **Intel® System Resources Defense (ISRD)**
-  **Intel® Tunable Replica Circuit – Fault Injection Detection**
-  **Intel® Partner Security Engine (Intel PSE)**
-  **Intel® Total Memory Encryption (TME)**
-  **Intel® Secure Key**
Generates high-quality keys for security and cryptography
-  **Post-Quantum Ready Memory Encryption**
Intel was the first silicon vendor to implement quantum-safe memory encryption in client and server platforms

-  **Confidential Computing Intel® TDX Connect**
-  **Confidential Computing**
Application Isolation with Intel® SGX | Virtual Machine Isolation with Intel® TDX
-  **Intel Converged Boot Guard and TXT (CnBT)**
For trusted BIOS, OS, and VMM
-  **Runtime Data Protection**
Helps ensure data is isolated and encrypted even while in memory
-  **Intel® Platform Firmware Resilience (Intel® PFR)**
Platform root-of-trust
-  **Intel® Threat Detection Technology (Intel® TDT)**
Detects fileless memory attacks
-  **FUTURE: Memory Tagging Technology (MTT)**
Introduces new x86 CPU instructions to detect memory safety violations such as buffer overflows and use-after-free errors
-  **Linear Address Space Separation (LASS)**
Prevents unauthorized reads/writes to SPI Flash
-  **CET (Control Flow Enforcement Technology)**
A hardware-based security feature developed by Intel to prevent control flow hijacking attacks
-  **VT-rp (Virtualization Technology for Redirected Protection)**
An Intel virtualization-based security feature that helps protect against memory corruption exploits
-  **Post-Quantum Ready Microcode Authentication**
Intel was the first silicon vendor to implement quantum-safe microcode encryption in client and server platforms



With industry-leading product security assurance as the foundation, Intel platforms integrate unrivaled silicon-level defenses and comprehensive security features—enabling customers to mitigate current threats and proactively address emerging attack vectors across cloud, edge, and AI environments.

Intel® technologies may require enabled hardware, software, or service activation. Not all features are available on all platforms. Please refer to the summary tables for feature availability details. No product or component can be secure. Your costs and results may vary.

Built on Industry-Leading Product Security Assurance

- Comprehensive Platform Security Requirements*
- Silicon Industry Leading Security Development Lifecycle (SDL)*
- Silicon Industry Leading Product Security Incident Response Team (PSIRT)*
- Predictable servicing through the quarterly Intel Platform Update (Intel® IPU)*
- Dedicated Offensive Security Research Team*

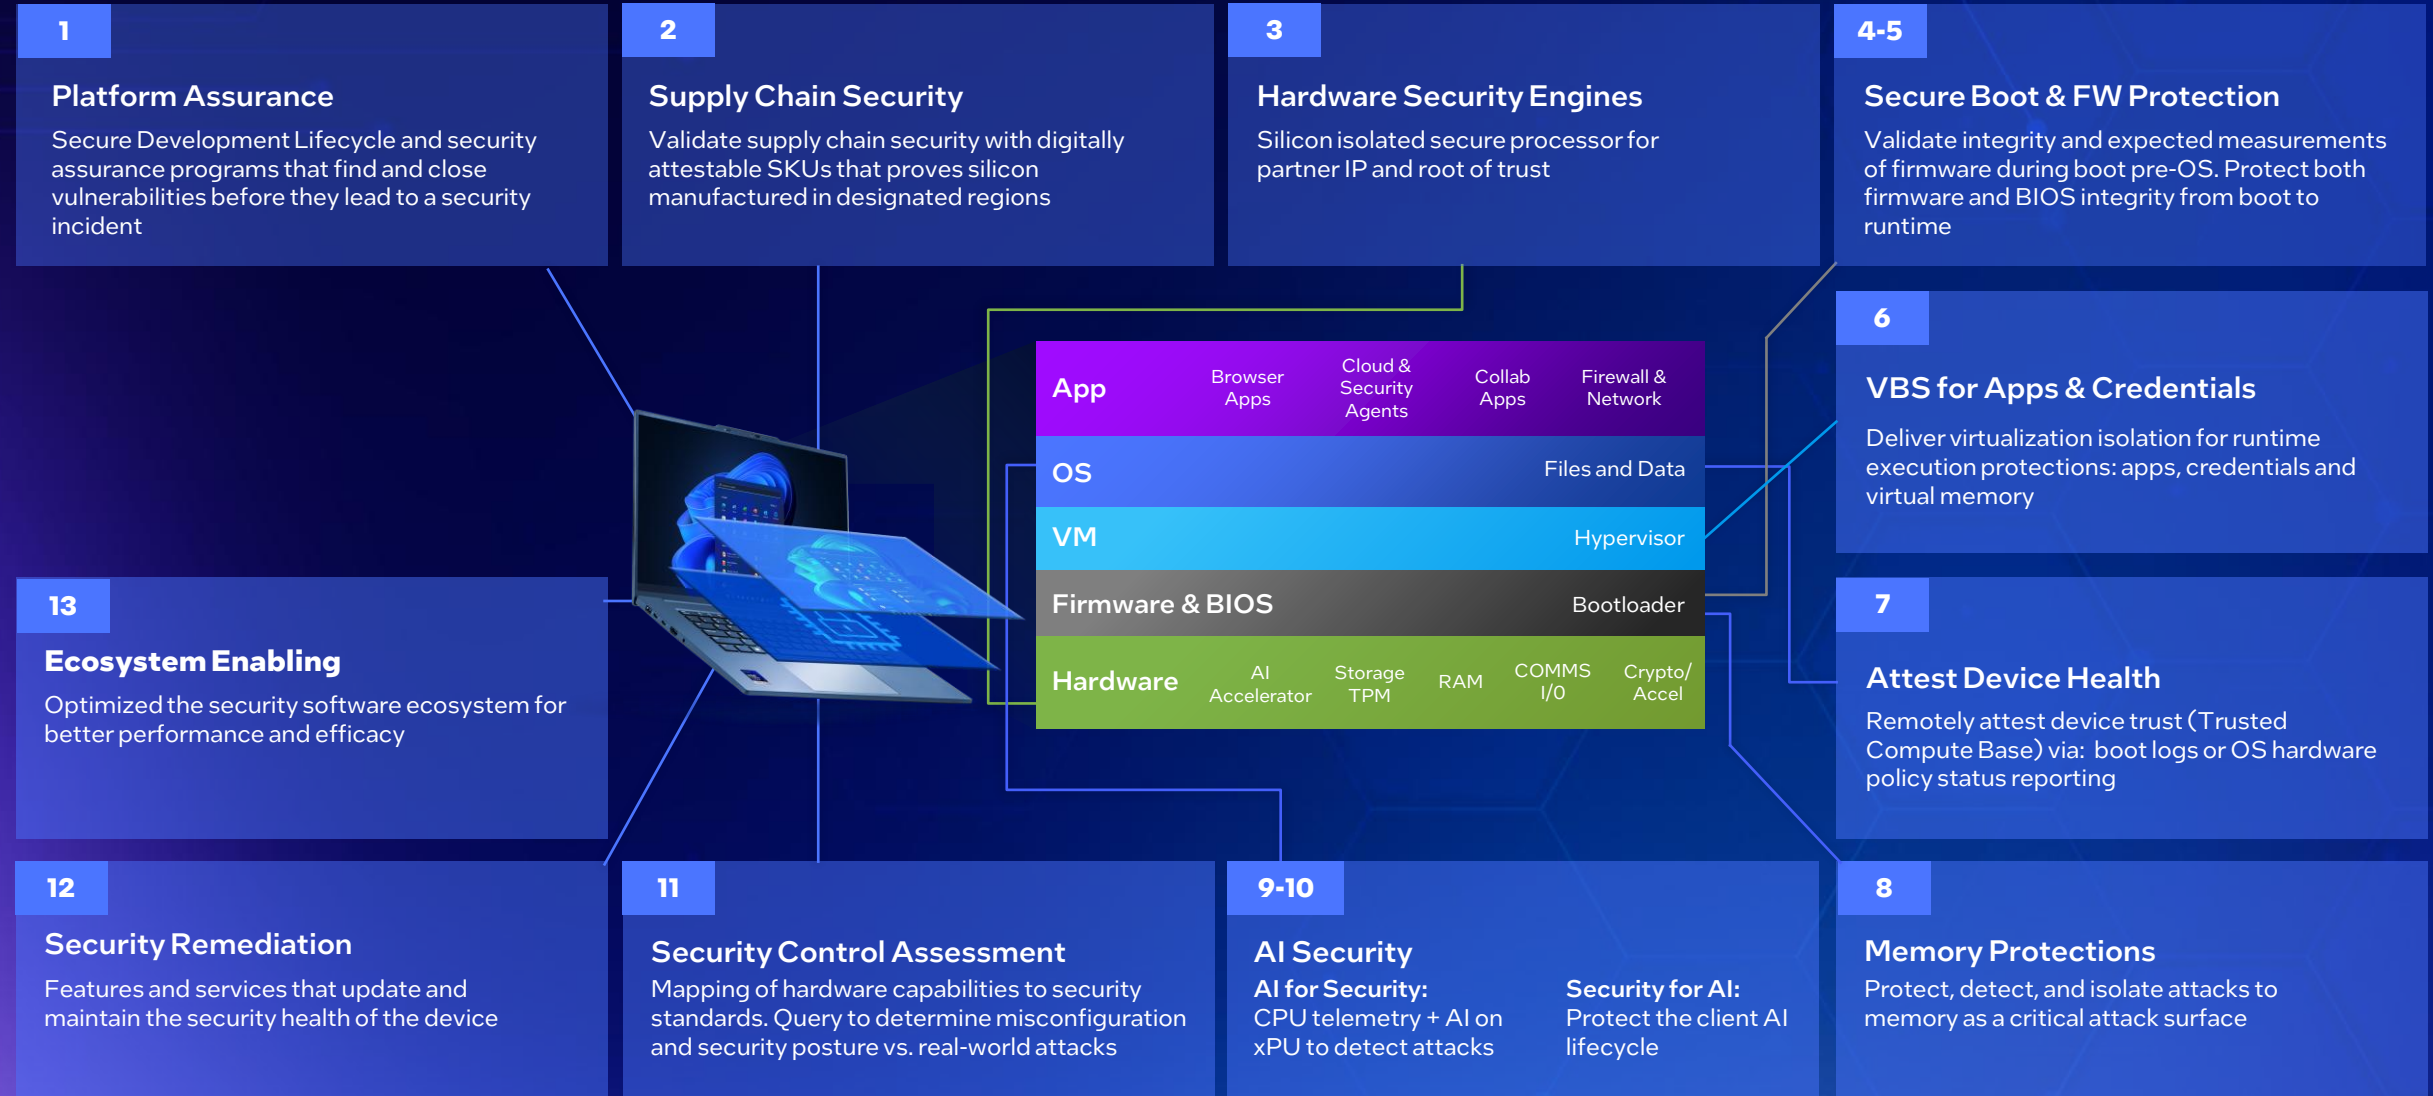
-  Root of Trust Security Features
-  Post-Quantum Cryptography
-  Confidential Computing
-  Server Security Features
-  Hardware-Assisted Protection Technologies



Platform Spotlight: Intel vPro[®]

Enterprise-grade, hardware-rooted security capabilities that surpass competitive offerings.

13 Categories that Define Modern Enterprise Security




Intel vPro[®] powered by Intel[®] Core[™] Ultra Series 3 Processors

VS AMD

Intel evaluated 13 essential security use cases that define a modern enterprise security posture for PC fleets. Intel scored the features & security programs that are needed to achieve each use case and applied a business scoring criteria that measured: feature deployment reach, impact for IT, market demand, and implementation difficulty. The percentage scores shown in the side-by-side dials reflect abstracted comparison results across five higher-level categories based on normalized Intel and AMD scores grouped by relevant Usages Compare.¹

Usages Compare: Product Security Assurance



Assurance Programs: Proactive secure development practices, vulnerability management & bug bounty programs help find and eliminate silicon platform vulnerabilities.


Intel Differentiator:

#1 Ranked
#1 in product security assurance²

Market Impact:

96%
of disclosed vulnerabilities closed³

Usages Compare: Supply Chain



Attest Validity of Hardware Components: Service and cryptographic capabilities that enable attestation across all supply chain phases: manufacture, transport, deployment, operation, and decommissioning.


Intel[®] Assured Supply Chain: Digitally attestable SKUs that prove silicon manufactured in designated regions

Intel Differentiator:


5 Supply chain phases supported

Market Impact:

Millions
of devices are more secured from espionage attacks⁴



Usages Compare: AI Security



AI for Security:

Intel silicon has evolved to provide a significant threat detection assist by using: AI models that profile malware using CPU telemetry with GPU/NPU accelerators

Security for AI:

Intel silicon can help secure key attack surfaces for client AI inferencing

Intel Differentiator:


+24%

of ransomware detection assist.⁵

Market Impact:

60%
of AV/EDR market leaders integrated.⁶

Usages Compare: Below-the-OS Security



Security Engines, Protected Boot/Firmware:

Modern SOC's all manage root of trust, digital rights, critical systems, and Microsoft Pluton usages but differ on attack protection/hardening capabilities

Virtualization Isolation:

Intel silicon technologies enable virtualization for a secure runtime to protect code, credentials, secrets, and keys but differ on VM & hypervisor attack protections

Memory Protections:

Intel Silicon can accelerate memory scanning, encrypt data, & block Living Off the Land attacks that are the most common first attack entry points


Intel Differentiator:

1st silicon CPU vendor to turn CSNA 2.0 post-quantum crypto standards into deployable hardware-rooted crypto solutions for business PCs.⁷

Market Impact:

2/3 of organizations surveyed concerned about harvest now, decrypt later attacks⁸

Usages Compare: Attest & Remediate Device Health



Assess Security Posture:

Third party security assessments, functional query APIs, and MITRE ATT&CK[®] mappings can deliver insights into fleet hardware misconfigurations and threat posture.

Security Remediation:

Out of band, chip-level KVM is a primary recovery tool for BSOD incidents

Intel Differentiator:

150

MITRE hardware feature mappings & config status that can be queried across fleets.⁹

Market Impact:

55%
of manageability software leaders integrated for Intel vPro recovery¹⁰

1) Among x86-based PCs, based on Intel analysis of available information as of March 2026 comparing unique architecture and unrivaled security capabilities of Intel vPro powered by Intel Core Ultra series 3 versus AMD Ryzen[™] Pro AI 300 Series processors as measured by primacy and ubiquity of relevant use case or threat to IT professionals and users, ease of solution implementation, and breadth of implementation across deployed systems. Details at intel.com/performance-vpro. No product can be absolutely secure.

2) See slide 13 for footnotes 2-10

Intel® Threat Detection Technology – DTECT

The only AI-enhanced silicon that detects the most advanced threats in real time.¹

1st ISV integrations expected 2H'26

Leverage PC hardware to discover the toughest threats

Major threat techniques evading security software	Intel® TDT-DTECT	Competitors' silicon ¹
Data stealers	✓	✗
File-less malware	✓	✗
Malware obfuscation	✓	✗
Evasion techniques	✓	✗
Trojans/backdoors	✓	✗

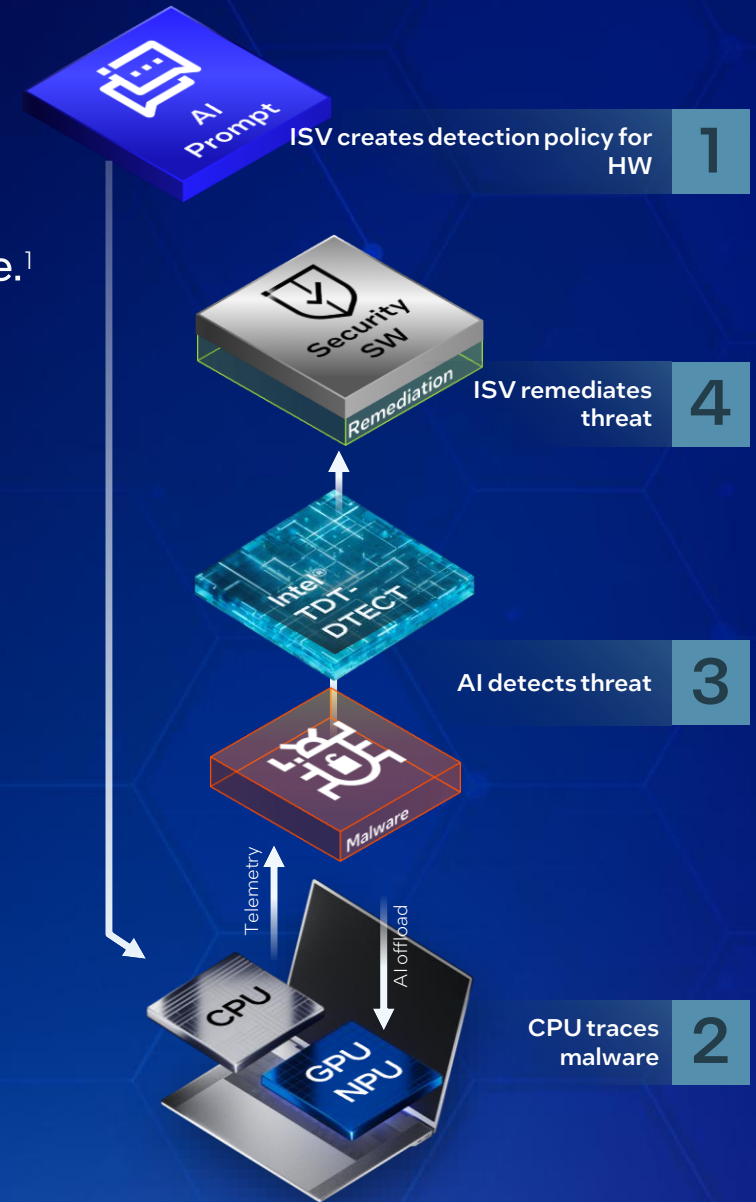
Delivering what others can't¹

Discover more threats, earlier

All stages of malware execution chain

Runs local on AI PC

No significant disruption to performance



Intel® Total Storage Encryption

New



Enhances data-at-rest security while exceeding modern enterprise security requirements.

Intel® Total Storage Encryption (Intel® TSE)

- Uses AES-256 encryption, incorporating post-quantum-resistant design principles to help prevent 'harvest now, decrypt later' attacks¹
- Protects encryption keys in a dedicated hardware engine to reduce attack surface
- Offloads encryption/decryption and frees up the CPU for productivity
- Intel® TSE is an Integrated solution offered all Intel® Core™ Ultra Series 3 processors with Intel vPro® platforms

70% average CPU cycle savings enabled by hardware-accelerated BitLocker compared with software-based BitLocker.²

Validated Defenses to Real World Attacks

PC Security Stack Mappings – Hardware Enabled Defense: MITRE ATT&CK® mapping shows how modern hardware-based security can significantly help enterprises counter threats and protect systems.¹ [Read more](#)

Intel is the only silicon vendor with industry-validated hardware security capabilities against the MITRE ATT&CK® framework, as demonstrated through a collaborative project with Microsoft, CrowdStrike, and ATTACK IQ that mapped Intel vPro® and Intel® Core™ Ultra AI PC mitigations to real-world threats, with 150 total mitigations validated — including 90 with Windows 11, Secured-core PC, and Microsoft Defender, and 80 with CrowdStrike Falcon, all proven using Attack IQ. Additionally, Intel mapped over 30 hardware security mitigations to MITRE ATLAS, further distinguishing vPro security in the industry.

intel
vPRO



MITRE
ATT&CK™

150

Overall Intel vPro
mitigations

30

Mitigations for
security for AI

90

Mitigations for Win11,
Secured-core PC,
Defender

93

Mitigations for


MSFT & MITRE Testimonial | [click image for video](#)




¹ Project in collaboration with Intel, Microsoft, CrowdStrike, and MITRE. Read an [Intel blog](#) to learn more.

Intel® Assured Supply Chain (Intel® ASC)

Trace and verify the source of silicon components and manufacturing locations

Intel Provides

The industry's only commercial PC silicon manufacturing corridor purpose-built for enhanced transparency and reduced geopolitical supply-chain risk.



Pre-determined end-to-end silicon manufacturing corridor spanning 5 locations:
US, Ireland, Vietnam, Taiwan, Malaysia

Digital supply chain verification:
Intel ASC identifier embedded at the silicon level
Customer verifiable using OEM and industry tools



Available on Intel® Core™ Ultra Series 3 processors with Intel vPro®
eligible SKUs in commercial systems from leading PC manufacturers¹

¹As of March 2026, based on Intel's unique silicon-enabled PC supply-chain authentication service among x86-based systems, which can validate hardware component authenticity and attest to manufacture in designated geographies. See intel.com/vpro for details.

Additional Resources

[AI PC Security: Next-Gen Threat Defense with MITRE and Intel](#)

[Intel AI PCs Deliver an Industry Validated Defense vs Real World Attacks](#)

[Stacked Defense from the Hardware Up – MITRE](#)

[Zero Trust Endpoints: Enabling Local AI with Privacy and Security Advantage](#)

[Windows 11 Upgrade – The Hardware Security Focused Refresh](#)

[The Intel® Core™ Ultra 200V Series with Intel vPro®: Leading Commercial System Manageability and Security](#)

Security Comparison Supporting Details		
2	Intel is #1 in product security assurance	Ranked #1 in 2024 ABI Report (independent research commissioned by Intel): Embracing Security as a Core Component of the Tech You Buy .
3	In the latest Product Security Report, Intel internally discovered 96% of disclosed vulnerabilities.	Source: 2026 Intel® Platform Security Report
4	Millions of Intel PCs devices are more secured from espionage attacks	Intel provides the only commercial PC silicon manufacturing corridor purpose-built for enhanced transparency and reduced geopolitical supply-chain risk (as of March 2026). Intel's unique silicon-enabled PC supply-chain authentication service among x86-based systems can validate hardware component authenticity and attest to manufacture in designated geographies.
5	Intel® Threat Detection Technology has been shown to increase overall EDR ransomware protection efficacy by 24% over software alone.	Based on SE Labs Enterprise Advanced Security (Ransomware – Intel) , Threat Detection Technology study published March 2023 (commissioned by Intel) analyzing ransomware detection rates with Intel TDT versus non-hardware based solutions.
6	60% of AV/EDR market leaders have integrated Intel® Threat Detection Technology	Based on Intel analysis as of March 2026.
7	Below the OS Security, two-thirds of organizations surveyed are concerned by harvest now, decrypt later attacks	Source: Capgemini, " Nearly two-thirds of organizations consider quantum computing as the most critical cybersecurity threat in 3–5 years " published July 2025
8	Intel mapped Intel vPro® Security across 150 MITRE hardware features that can be queried across fleets.	Refers to Intel being the first and only CPU silicon vendor to implement CSNA 2.0-aligned post-quantum cryptography into deployable, hardware-rooted platform security capabilities for commercial client systems among x86-based business PCs. Based on analysis of publicly available information as of March 2026.
9	Intel delivers industry-leading boot-time protection to establish trusted business PCs.	As of March 2026, among x86-based PCs, including in comparison with competition and prior generation products; refers to Intel vPro® systems powered by Intel® Core™ Ultra Series 3 processors, based on unique architecture and unrivaled security capabilities engineered to help safeguard organizations of all sizes. Intel provides the broadest suite of x86-based Windows PC capabilities to expose accurate device-security status and to help the Windows OS enforce boot-time policies that establish device trust.
10	55% of manageability software leaders have integrated recovery enabled by Intel vPro®.	Based on Intel analysis as of March 2026.

Confidential AI

Secure, confidential AI isn't optional—it's the future. Intel is making that future possible.



Confidential AI: Securing Innovation at the Speed of Transformation

By 2029, Gartner predicts more than 75% of operations processed in untrusted infrastructure will be secured in-use by confidential computing.⁴

Generative AI (GenAI) and Agentic AI are reshaping industries—from healthcare and finance to manufacturing and government—by unlocking new efficiencies, insights, and capabilities. Enterprises are approaching AI deployments with caution and considering solutions that protect enterprise data and models throughout the AI workflow and ensure security and privacy from the edge to the cloud. Solutions must integrate with legacy and emerging IT systems and be able to leverage the latest and greatest AI hardware.

To address these needs, Intel is advancing a silicon-rooted strategy for **Confidential AI, built on over a decade of leadership as the inventor of the general-use Confidential Computing paradigm**. This approach enables enterprises to protect proprietary data, models, and algorithms throughout the AI lifecycle—during training, fine-tuning, and inference—without compromising performance or scalability.

Intel's **Confidential Computing** portfolio, including Intel® SGX and Intel® TDX, provide end-to-end solutions to secure all AI workflows, all built on the industries most trusted platform.



Addresses
GenAI
threats



Accelerates
enterprise AI
adoption



Unlocks
proprietary
data



Powered
by open
software



Protects
agentic AI
workflows

With Confidential AI, Intel is working through industry groups to shape a future where enterprises can confidently deploy AI in high-stakes environments, knowing their data and models are protected by design.



Intel Platforms Provide the Key Capabilities of Confidential AI

Data-in-Use Protection

Intel's Confidential Computing technologies, including Intel® Trust Domain Extensions (TDX) and Intel® Software Guard Extensions (SGX), create hardware-based Trusted Execution Environments (TEEs) that encrypt and isolate sensitive data while it is actively processed.

Model Integrity and IP Safeguards

Proprietary models are shielded from reverse engineering and unauthorized access, helping enterprises preserve competitive advantage and reduce the risk of intellectual property theft.

Compliance-Ready Execution

Attested environments provide cryptographic proof of software integrity and data handling practices, supporting regulatory frameworks such as GDPR, HIPAA, and the EU AI Act.

Scalable Ecosystem Integration

Intel collaborates with OEMs, ISVs, and cloud providers to embed Confidential AI into real-world deployments—from PrivateGPT appliances to federated learning platforms—ensuring secure AI adoption across sectors.

Agentic AI

Protect Agentic AI workflows, orchestration, prompts, and context with Confidential Computing.

Transparency

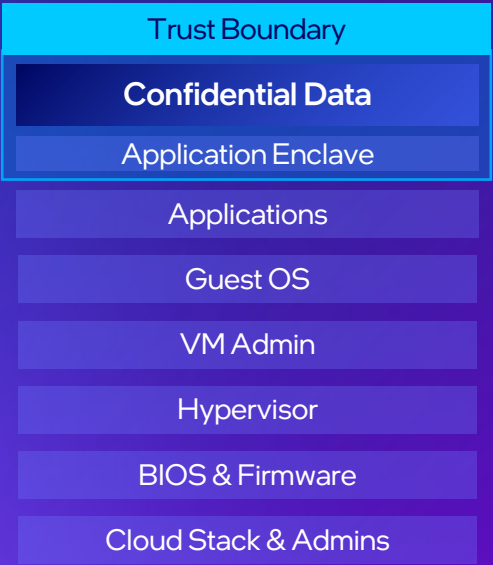
The Intel® TDX module source code is publicly available giving customers full visibility and confidence in security, while accelerating industry-wide innovation and trust.

Intel Offers the Most Comprehensive Confidential Computing Portfolio in the Industry

App Isolation

Intel® SGX


Smallest trust boundary for greatest data protection & code integrity



VM Isolation

Intel® TDX

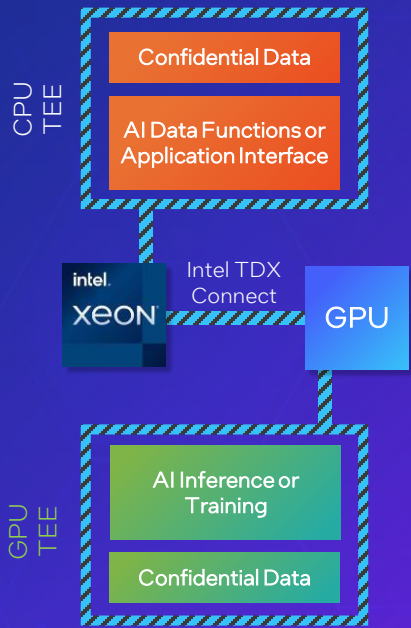
Most straightforward path to greater security and compliance



Encrypted Connection

Intel® TDX Connect

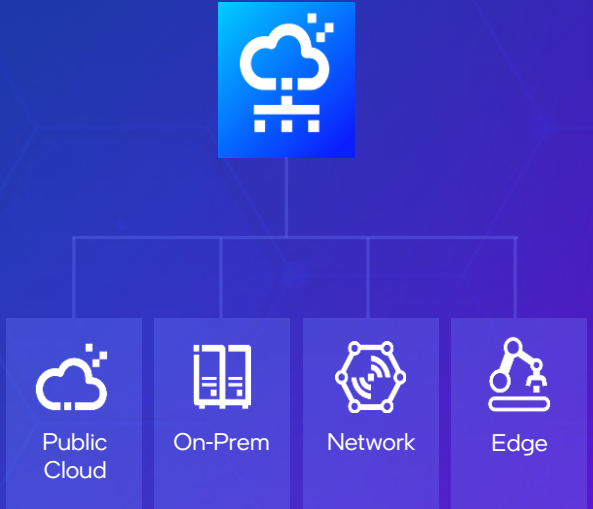
High-performance encrypted connection between CPU and PCIe devices



Trust Services

Intel® Trust Authority

Uniform, independent attestation of trustworthy environments with DCAP libraries and the Intel Trust Authority Service



The Next Milestone in Confidential AI with Intel® TDX Connect

Intel® TDX Connect is an extension of Intel® Trust Domain Extensions (Intel® TDX), designed to secure I/O communications between Confidential Virtual Machines (VMs) and PCIe devices such as GPUs, SmartNICs, and storage devices. It is especially valuable for Confidential AI and other sensitive workloads, providing a high-performance, encrypted connection between the CPU and PCIe devices with direct memory access and lower overhead.

*Activating TDX Connect will require Intel® Xeon 6 with P-cores, TDX Module updates, an enabled OS, and a TDISP enabled generation 6 PCIe device.

“

Microsoft is excited to productize Intel TDX Connect into future generations of Azure confidential VMs, which is in early development between our hardware and software developers. Intel TDX Connect represents a major milestone in our journey to improve confidential computing performance and extensibility. It enables existing and net-new workloads to benefit from the privacy assurances of confidential computing without compromising on price or performance.”

— **Vikas Bhatia**, Head of Product for Azure Confidential Computing at Microsoft



Intel® TDX Connect Key Benefits

End-to-End Confidentiality

Establishes a high-performance encrypted channel between CPU and PCIe devices, enabling data confidentiality during transit.

Secure Direct Memory Access (DMA)

Enables secure DMA between CPU and accelerators, minimizing risks of data leakage or tampering.

Enhanced I/O Virtualization

Improves I/O virtualization performance with reduced latency, ideal for high-throughput AI and analytics workloads.

Multi-Vendor Device Support

Compatible with PCIe-compliant accelerators from various vendors, including support for platforms like Nvidia Blackwell.

Cloud and Edge Deployment

Supports secure deployment across cloud-native, hybrid, and edge environments.

Zero Trust Compliance

Facilitates cryptographic attestation and verifiable execution environments, aligning with Zero Trust principles and regulatory standards.

Seamless Infrastructure Integration

Designed to work with Intel Xeon 6 processors and future confidential VMs, enabling easy adoption without major architectural changes.

Security Research on Intel Confidential Computing

We encourage offensive research on our products to continuously strengthen our security



Proactive Security Validation

At Intel, we're dedicated to making our Confidential Computing technology as secure as possible. This kind of collaborative research extends our internal threat models and helps uncover and address security vulnerabilities that can emerge in these complex environments before malicious actors can take advantage of them.

“

Intel TDX is an instrumental technology helping to achieve our confidential computing goals. Now that we are finished, it's even more secure, and I'm very confident, after this hackathon, with this technology.”

— **Yair Netzer**, Principal Security Research Manager, Microsoft

“

Our deep collaboration with Intel allows us to battle-test and strengthen the security of foundational technologies that power Confidential Computing. By proactively identifying vulnerabilities in critical features like Live Migration and TD Partitioning using advanced AI tools like Gemini, we are helping to raise the security bar for the entire ecosystem.”


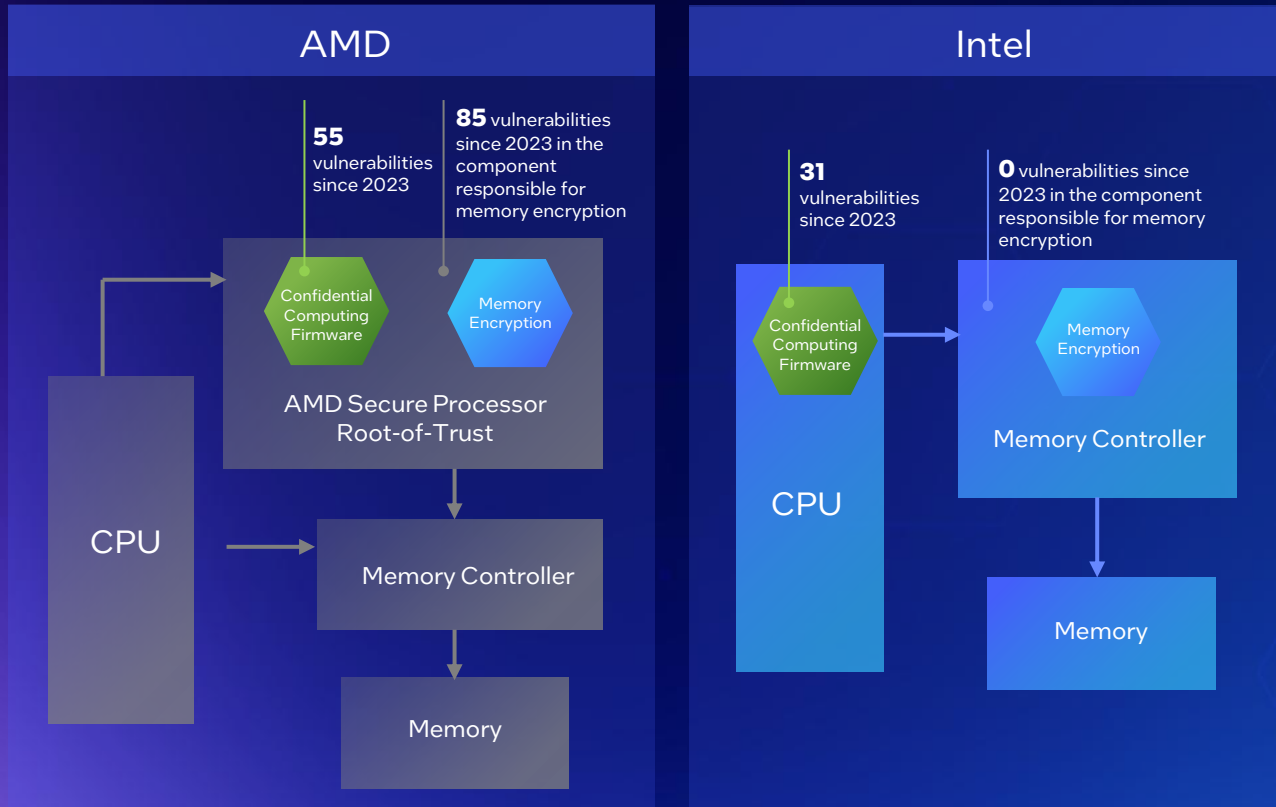
— **Andrés Lagar-Cavilla**, Distinguished Engineer, Google Cloud

Hardware, the Foundation of Trust

A Trusted Execution Environment (TEE) is only as secure as the underlying hardware and firmware. With industry leading product security assurance as the foundation, Intel reduces the attack surface and total cost of ownership (TCO) over competing products.

AMD reported **100% more vulnerabilities** than Intel 2023 - 2025 in the component responsible for **memory encryption**.

AMD reported **77% more vulnerabilities** than Intel 2023 - 2025 in their **Confidential Computing firmware**.



VULNERABILITY SPOTLIGHT: RMPocalypse

What: A medium severity (CVSS 6.0) race condition in the AMD Secure Processor. Intel platforms are not affected.

Impact: Potential to allow a malicious hypervisor to corrupt the Reverse Map Table (RMP) during initialization, affecting VM guest memory integrity.

Engineering Decision: The vulnerability centers on a component called the Reverse Map Table (RMP). This table is the "protector" of the system's memory. Its job is to ensure that only authorized programs can access specific data. Because the RMP is very large, it has to be stored in the main memory. To protect that memory, the system uses the RMP itself.

Secure Design Gap: This creates a "chicken-and-egg" or Catch-22 scenario: How do you protect the RMP before it has been set up? Researchers found that during the initialization process (when the computer is starting up or setting up a virtual machine), there is a tiny window of time where the RMP is not fully protected.

CWE-284 (Improper Access Control): The core of RMPocalypse is that the system fails to restrict "write" access to the Reverse Map Table (RMP) during its sensitive initialization phase. An unauthorized actor (the hypervisor) can modify data it should not be able to touch.

Intel View: Intel takes "secure by design" very seriously. At different parts of our product development process, our designs and implementations are carefully evaluated against a variety of common security weaknesses, including proper access control for sensitive content. Our assurance efforts focus on system interactions under normal operations as well as corner-case scenarios. This attention to detail ensures access control works as intended amid race conditions as well as power state transitions.

Who: Published October 2025 in a paper titled "RMPocalypse: How a Catch-22 Breaks AMD SEV-SNP" by researchers Benedict Schlüter and Shweta Shinde from ETH Zurich.

Additional Confidential AI Resources

[Intel Confidential Computing Solutions](#)

[Intel® TDX on Dell PowerEdge: Fraud Detection](#)

[Confidential AI – Protecting Data and Models with Intel Confidential Computing](#)

[Securing AI Workloads with Intel® TDX, NVIDIA Confidential Computing and Supermicro Servers with NVIDIA HGX™ B200 GPUs: A Foundation for Confidential AI at Scale](#)

[Confidential Computing: Powering the Next Generation of Trusted AI](#)

[Intel® Architecture Memory Protections for Confidential Computing](#)

[Safeguarding Foundational Technologies: How Intel and Google Collaborate to Strengthen Intel® TDX](#)

[Strengthening the Foundation: A Joint Security Review of Intel TDX 1.5](#)

Post-Quantum Cryptography (PQC) Readiness & Compliance

Intel is not just preparing for the quantum era—we're shaping it by co-authoring global PQC standards and accelerating hybrid cryptography across client and server ecosystems to secure the digital backbone of tomorrow.

Securing the Future: Intel's Post-Quantum Cryptography Strategy for Enterprise Resilience

As quantum computing advances toward practical viability, enterprises face a critical inflection point: the cryptographic foundations securing today's digital infrastructure are at risk of obsolescence. The looming threat of "Q-Day"—when quantum computers can break widely used public-key algorithms—has catalyzed a global shift toward Post-Quantum Cryptography (PQC). For technical decision-makers, this transition is not merely a compliance exercise but a strategic imperative to safeguard long-lived data, maintain trust, and ensure business continuity.

Addressing PQC Threats

Complete

Resilience to Data Harvesting

Intel has adopted larger key sizes for symmetric crypto to protect against Harvest Now Decrypt Later (HNDL) attacks.

Nearing Completion

Code Signing & Authentication of Firmware

Intel has deployed CNSA 2.0 compliant PQC algorithms for CPU and SoC firmware.

In-progress

Secure Internet with new Digital Signature & Key Establishment Standards

Intel is helping in the ongoing effort to develop NIST PQC standards and implementing them in our products.

Intel's PQC Roadmap

PQC Algorithm Standards

Intel is deeply engaged in the development of PQC algorithms across the industry. Standards development is expected to continue for the next several years.

PQC Technology Implementation

Intel began offering PQC capabilities in its platforms in 2023.

Full PQC Compliance

All new Intel platforms will incorporate PQC resistant algorithms across the full stack by 2030.

Intel PQC Leadership

2016

Intel begins work on PQC standards with NIST and other regulators worldwide.

2017

Intel begins increasing key sizes for symmetric crypto to protect against HNDL.

2023

Intel first to implement quantum-resistant CPU microcode authentication and memory encryption in client platforms (Intel® Core™ Ultra Series 1).

2024

Intel first to implement quantum-resistant CPU microcode authentication and memory encryption in server platforms (Intel® Xeon® 6 Processor with P-cores) and security engine firmware (Intel® Core™ Ultra 200v Series Processor).

2024

Intel co-authored the SLH-DSA (FIPS-205) standard for PQC digital signatures.

Be Ready for Q-Day with Intel Platforms!

Platform Security Compliance Leadership

Intel leads in [Product Security Certifications](#) through its FIPS approach, providing a structured process for certifying products and ensuring product readiness to meet diverse customer requirements. Through mandates, terminology, dedicated labs, and certified solutions, Intel is uniquely positioned to support government, regulatory, and defense industries—reinforcing its role as a semiconductor leader.



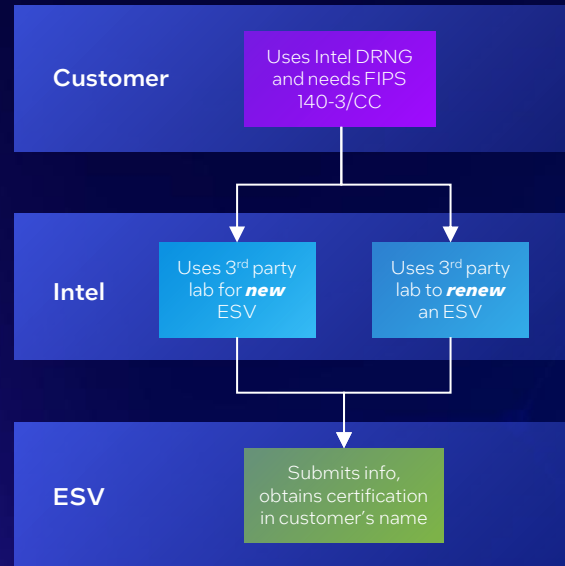
Intel manages entire process



Efficient path to NIST cert



For new or updated ESV



Highlights:

- Intel is the first to provide **Entropy Source Validation (ESV)** as a service to its customers under its **Intel Secure Key validation service** enabling customers to seamlessly acquire the ESV validation in their name.
- Intel has **certified modules** such as its security engines and hardware accelerators across different client and server products and maintains a certified Digital Random Number Generator (DRNG).
- Intel's commitment to product readiness—referred to as **FIPS Certifiable**—has enabled the company to deliver solutions like the Cryptography Primitives Library (CPL), ensuring compliance with rigorous security standards.
- The Intel® Secure Key Validation Service supports validation of NIST approved PQC algorithms implemented in Intel products.

*Intel has defined the term "Certifiable" as the identification of cryptographic module, designed and implemented to meet applicable FIPS PUB 140-3 requirements, documented, pre-tested, that includes one or more FIPS Approved Algorithms (CAVP), entropy source validation (ESV), with the capability to perform required self-tests and other operations, and eventually enables acquisition of FIPS 140-3 validation from CMVP by customers (OEMs...) or Intel if desired.

Intel Product Mandate: All products incorporating cryptography that reach production must be FIPS-certifiable*.



Intel CST Laboratory

Intel's Cryptography Security Testing Lab was established to support validation of all implemented crypto algorithms on its products.

FIPS Certified and Certifiable Products

Intel MEV-TS	IPP Crypto	QAT
Intel DRNG	CSME	IPSEC-MB



VULNERABILITY SPOTLIGHT: RDSEED

What: A high-severity (CVSS 7.2) vulnerability in the RDSEED instruction on AMD platforms compromising the integrity of hardware-generated random numbers critical for cryptographic functions. Intel platforms are not affected.

Impact: Systems relying on the affected RDSEED instruction may produce predictable or weak values, potentially allowing an attacker to infer sensitive information, compromise encryption keys, or bypass authentication mechanisms.

Engineering Decision: This vulnerability is a microarchitectural defect in Zen 5 processors that incorrectly signals "success" (Carry Flag = 1) even when the internal entropy pool is exhausted. Specifically, under heavy system load, the 16-bit and 32-bit RDSEED instructions fail to detect entropy depletion and silently return non-random zeros as if they were valid cryptographic seeds.

Secure Design Gap: Failure to enforce the hardware-software contract: the processor is expected to monitor the entropy pool and return an error when it no longer has enough entropy to generate a random number. This "silent failure" prevents software from detecting hardware depletion, leading to the use of predictable values in security-critical operations. The escape also highlights a lack of robust security testing of the product. Error handling needs to be carefully implemented AND verified.

CWE-333 (Improper Handling of Insufficient Entropy in TRNG): The core of the RDSEED vulnerability is a failed "Failure Signal" leading to predictable output.

Intel View: Intel works to continuously harden its Secure Key technology. Intel's recent Entropy Security Certification is a testament to its robust design conforming to NIST 800-90A/B/C, earning FIPS 140-3 and Common Criteria certifications.

Who: Discovered in 2025 by Gregory Price, a Linux kernel engineer with Meta.

Additional PQC and Compliance Resources

[Be Ready for Post-Quantum Security with Intel® Cryptography Primitives Library](#)

[Who's at Fault? A Look at Post-Quantum Cryptography and Fault Injection Attacks](#)

[Accelerate Post-Quantum Cryptography with Intel Crypto Technologies](#)

[Intel Co-Develops One of Three New Post-Quantum Crypto Standards Released by NIST](#)

[Post-Quantum Cryptography: Defending Against Future Adversaries with NIST Standards](#)

Software Robustness

Intel platforms integrate silicon-level security primitives that enforce memory safety, control flow integrity, and runtime isolation—making them among the most resilient environments for executing software securely and reliably.

Safe Platforms for Software Execution

Intel platforms stand out for their robust hardware-assisted protection technologies, which directly address the most common and dangerous software vulnerabilities. **Control-flow Enforcement Technology (CET)** provides hardware-enforced integrity for application and OS execution paths, using shadow stacks and indirect branch tracking to block Return-Oriented Programming (ROP), Jump-Oriented Programming (JOP), and Call-Oriented Programming (COP) attacks. This makes code-reuse exploits significantly harder to execute compared to software-only defenses. Similarly, **Virtualization Technology for Redirect Protection (VT-rp)** leverages Intel VT-x extensions to create hypervisor-enforced security boundaries, preventing malicious remapping of memory and isolating critical system components from kernel-level exploits and rootkits. These features help ensure runtime resilience on clients, servers, and in complex virtualized environments.

Intel also leads in x86 memory safety innovations with **Linear Address Space Separation (LASS)** and **Memory Tagging Technology (MTT, code-named ChkTag)**. LASS hardens the platform by enforcing strict hardware-level separation between user and kernel memory, making it significantly harder for attackers to exploit privilege-escalation paths or extract kernel data. MTT, once available, will go further by adding new x86 instructions for hardware-accelerated detection of buffer overflows, use-after-free errors, and other memory safety violations—eliminating entire classes of vulnerabilities that plague modern software. These technologies help to harden applications, OS kernels, hypervisors, and even UEFI firmware, creating a silicon-level defense that scales from client PCs to cloud infrastructure. By embedding these protections directly into the CPU, Intel platforms deliver production-grade security that software-only solutions cannot match, making them the safe choice for executing critical workloads.



“

CET protected software from ROP exploits (which was the most significant technique for real-world exploits on the OS kernel after CFG, and the other mitigations were put in place), and after the deployment of CET, attackers were looking at data-only attacks or to modify page tables to achieve code execution, simply because it was the only attack vector that was still available in the kernel.”


— **Andrea Allievi**, Senior Windows Core OS Developer, Microsoft

Software Robustness: A Defense-in-Depth Stack

Software Robustness is the ability of a system to **tolerate software bugs (defects)** such that they do not result in exploitable security vulnerabilities.

KEY THREATS

These threats **still account** for the majority of high-severity real-world exploits



Memory Safety (70%)

Buffer Overflow, Use-After-Free, OOB



Control-Flow Attacks

ROP, JOP, COP



Isolation Failures

Sandbox Escape, In-Process Breach



Boundary Violations

User/Kernel, VM/Host Escapes

KEY TECHNOLOGIES TO COMBAT THREATS



HOW TECHNOLOGY COMBATS THE THREAT

Memory-safe languages prevent whole bug classes from existing

Structural correctness & checks

Damage containment at Runtime

Hardware protections help prevent inevitable bugs from being exploitable

Robust systems use layers of protection:
programming language, software, and hardware

Hardware-Assisted Protection Technologies

Hardware Layer Protection

LASS (Linear Address Space Separation) enforces a strict "no user access to kernel memory/no kernel access to user memory" policy at the hardware level, making it much harder for attackers to exploit the traditional user/kernel boundary violations that many privilege escalation exploits depend on.

LASS



Software Layer Attacks

Attack Type: Meltdown and some Spectre variants that rely on cache-based covert channels.

What it does: Enforces strict separation between user and kernel address spaces without relying on paging accesses that may reveal kernel memory layouts.

Why it matters: It hides the kernel memory layout to make it harder for attackers to compromise the operating system. It also mitigates attempts to exfiltrate kernel data.

CET (Control-flow Enforcement) provides hardware-enforced control flow integrity that's much harder to bypass than software-only solutions, making traditional code-reuse attacks significantly more difficult to execute successfully.

CET



Attack Type: Return-Oriented Programming (ROP), Jump-Oriented Programming (JOP), Call-Oriented Programming (COP).

What it does: Protects against the above attacks using shadow stacks and indirect branch tracking.

Why it matters: It hardens the control flow of applications and the OS, making exploitation of memory corruption bugs significantly harder.

VT-rp (Virtualization Technology for Redirect Protection) leverages Intel VT-x virtualization extensions to create hypervisor-enforced security boundaries to isolate critical system components from potential exploits.

VT-rp



Attack Type: Rootkits

What it does: Secures virtual memory mapping mechanisms in virtualized environments.

Why it matters: It allows hypervisors to ensure that VMs can't be tricked into remapping memory in unsafe ways.

FUTURE: MTT (Memory Tagging Technology) is an x86 CPU instruction set enhancement developed jointly by Intel and the x86 Ecosystem Advisory Group (EAG) to address memory safety violations, which are a leading cause of software vulnerabilities in production usages.

MTT



Attack Type: Memory safety violations (buffer overflows, use-after-free) in production environments.


What it does: MTT introduces new x86 instructions that detect memory safety violations like buffer overflows and use-after-free errors.

Why it matters: It brings hardware-accelerated protection to the x86 ecosystem, addressing the most prevalent class of software vulnerabilities and enabling safer deployment of production workloads.

These technologies require an Intel equipped processor, an equipped OS and software designed to use it.

Introducing Memory Tagging Technology: Intel's Latest Contribution to the x86 Ecosystem

Intel is redefining the future of secure computing by making its x86 platforms the most memory-safe in the world—uniting hardware precision with developer flexibility through Memory Tagging Technology (MTT). This silicon-level innovation positions Intel to lead in protecting AI and mission-critical workloads from the ground up.



Memory tagging creates a hardware-enforced validation layer by pairing pointers with specific memory locations through matching metadata tags. By storing a "key" within the pointer's bits and a "lock" within the physical memory's metadata, the CPU can instantly detect unauthorized access. If a pointer attempts to access a memory object with a different tag—whether due to a buffer overflow or a use-after-free error—the hardware identifies the mismatch and halts the operation, effectively neutralizing most common memory corruption vulnerabilities.

Key Components of MTT

(MTT was originally introduced as "ChkTag" in 2025⁵)

Instruction Set Enhancement

MTT introduces new x86 CPU instructions to detect memory safety violations such as buffer overflows and use-after-free errors. These operate at instruction-level granularity enabling targeted tag checks where needed.

Hardware-Accelerated Detection

Unlike software-only approaches (e.g., address sanitizers), MTT uses hardware acceleration to enable production-grade memory safety checking, making it viable for real-world deployment.

Flexible Developer Control

Developers can tune its behavior via compiler optimizations, language features, or intrinsics to balance security with performance and operational needs.

Broad Applicability

It is designed to harden a wide range of environments, including applications, OS kernels, hypervisors, and UEFI firmware.

Compatibility and Portability

MTT software remains functional on processors without MTT support, eliminating the need for multiple binaries and easing deployment.

Ecosystem Collaboration

Invented by Intel and developed in collaboration with the x86 Ecosystem Advisory Group (EAG), MTT aims to unify memory tagging across the x86 ecosystem.

Complementary to Existing Security Features

MTT works alongside existing x86 security mechanisms like shadow stacks and confidential computing, enhancing overall platform robustness.

Additional Software Robustness Resources

[Protecting linear address translations with Hypervisor-enforced Paging Translation \(HVPT\)](#)

[ChkTag: x86 Memory Safety](#)

[A Technical Look at Intel® Control-Flow Enforcement Technology](#)

[Modern AI PCs Need Advanced Security](#)



Product Security Assurance

The Architecture of Trust Starts Here

Intel Continues to Lead the Silicon Industry in Product Security Assurance

Intel technology is designed to accelerate a Zero Trust strategy, enabling hardware as the root of trust, with industry-leading security assurance as the foundation everything is built on.

An independent 2024 study by ABI Research offers a comparative assessment of the Security Assurance Practices of top silicon vendors, ranking Intel number one in the industry.

Product security assurance at Intel is an investment in people, processes, and tools extending from development and manufacturing to the end of the product lifecycle. It means that customers can be confident in Intel's Security-First Pledge: we design with security in mind, continually look for ways to strengthen our products, and disclose vulnerabilities we find.

Intel Ranked #1 vs Key Competitors for Product Security Assurance¹



SPOTLIGHT:

Intel Secure Development Lifecycle (SDL)

Intel introduced its Secure Development Lifecycle (SDL) in 2008 and has continuously refined the framework to strengthen platform security. In 2019, SDL underwent a significant transformation, mandating component-level integration across hardware, firmware, and software. This granular approach embeds security requirements at the lowest design layers, amplifying SDL's effectiveness and establishing Intel's leadership in silicon security engineering.

Areas of Intel Leadership:

- Security Development Lifecycle
- Proactive Security Practices
- Threat Discovery and Response
- Offensive Security Research
- Security Training
- Community Engagement

To learn more about Intel's differentiating product security assurance investments, visit: <https://www.intel.com/content/www/us/en/security/product-security-assurance.html>

SDL Impact:


80%

Reduction in security escapes for products developed since 2019

The Product Security Assurance Advantage

Intel's proactive approach to product security differentiates it from competitors. Here, we analyze two firmware categories that are critical to the platform's security and to the data processed on the platform.

Intel and AMD platforms both have security processors (or engines) acting as the **root-of-trust**. Comparing Intel® Converged Security and Management Engine (Intel® CSME) combined with Intel® Server Platform Services (Intel® SPS) to the AMD Secure Processor (ASP), we see differences in implementation. Intel's security engines are purpose-built to authenticate firmware during boot up, while AMD's security processor holds responsibility for many other tasks, such as cryptographic functions and random number generation for instructions like RDRAND and RDSEED. For this report, we include vulnerabilities disclosed in the known functions of each engine.



From 2023 to 2025, AMD reported

7x

more vulnerabilities in its hardware Root-of-Trust than Intel.



VULNERABILITY SPOTLIGHT: SinkClose

What: A high-severity (CVSS 7.2) SMM vulnerability (CVE-2023-31315) affecting AMD platforms dating back to 2006, including Ryzen, EPYC, and Threadripper processors. SinkClose could allow a privileged attacker to escalate privileges into SMM. Intel platforms are not affected.

Impact: Potential installation of persistent malware such as bootkits and rootkits that survive reboots, OS reinstalls, or even firmware updates. Infections in SMM are nearly undetectable and could compromise system integrity below the OS level.

Engineering Decision: The vulnerability stems from improper validation of values written to a specific Model-Specific Register (MSR), combined with a legacy backward-compatibility feature, Tclose, in AMD's address translation mechanism.


Secure Design Gap: In a secure design, the SMI Lock bit should render critical SMM configurations immutable post-boot. However, AMD's hardware did not enforce this lock on the TClose feature, creating an unintended bypass that persisted for nearly two decades due to overlooked legacy paths.

CWE-20 (Improper Input Validation): This is the primary classification, as AMD describes the root cause as "improper validation in a model-specific register (MSR)," which enables chained weaknesses like CWE-94 (Improper Control of Generation of Code) for arbitrary code execution in SMM.

Intel View: Intel recognized this space as a potential target for researchers and attackers over a decade ago and has implemented proactive product security assurance and security hardening efforts to safeguard our products. As a result, Intel is not affected by SinkClose.

Who: SinkClose was discovered by IOActive security researchers Enrique Nissim and Krzysztof Okupski, who publicly disclosed it at DEF CON 32 (2024) in a talk titled "AMD SinkClose: Universal Ring-2 Privilege Escalation." This follows a pattern of AMD SMM flaws, akin to the "Ghosts in the Machine Check" exploit (DEF CON 33, 2025), using Machine Check Exceptions for similar escalations.

Intel pioneered the general-use **confidential computing** paradigm in 2015 with Intel® Software Guard Extensions (Intel® SGX). Our confidential computing technologies undergo continuous validation through rigorous internal hardening and collaborative security assessments with leading cloud service providers.



From 2023 to 2025, AMD reported

77%

more vulnerabilities in its Confidential Computing firmware than Intel.



Advancing Silicon Security Through Security Hardening, Formal Verification, & Negative Space Testing

A closer look at advanced methods of security validation



INT31 is Intel's world-class security research team that performs cutting-edge offensive security research and analysis across all stages of the product engineering lifecycle, fortifying those products against emerging threats.

1000 **10**

Combined Years
of Experience

Countries

Areas of Expertise

- Attacks on Privileged FW
- Cryptography
- Fault Injection
- Formal Methods
- Fuzzing
- Industry Research
- Memory Research
- Microcode and Micro-architecture
- Networking Technologies
- Novel SW and HW Mitigations
- Physical Attacks
- Security of AI and AI to Break Security
- Side Channels (HW and SW)
- Supply Chain Security
- Telemetry-based Attacks

Recent publications by the INT31 team, including groundbreaking work on IODyne, Intel.BIN, and Negative Space Verification, make it clear that INT31 is not just pushing boundaries in hardware security, they're redefining the entire security assurance and verification landscape with precision, creativity, and unmatched technical depth.

New Accordion Mode Improves Protection of Cloud-Scale Data

What is it:

A new cryptographic mode called the *accordion*, implemented as *ddd-aes*, designed to replace AES-GCM for large-scale data encryption with better security and scalability.

Why it matters:

- Solves nonce reuse and birthday-bound limitations.
- Enables secure encryption for cloud-scale workloads.
- Supports future NIST standardization efforts.

Intel.BIN: The BIOS Binary Instrumentation Framework

What is it:

Intel.BIN is a dynamic instrumentation framework for BIOS binaries that enables runtime analysis of firmware vulnerabilities without modifying the BIOS.

Why it matters:

- Detects buffer overruns and privilege escalation.
- Operates in pre-OS environments with high fidelity.
- Fills a critical tooling gap in firmware security research.

Precise Dataflow Tracking in Hardware: Improving Hardware Security Verification with IODyne

What is it:

IODyne is a dataflow tracking tool that improves pre-silicon security validation by eliminating false positives and automating vulnerability case reduction.

Why it matters:

- Accurately detects threats like Meltdown and MFBDS.
- Reduces manual debugging and verification effort.
- Scales across complex SoC designs.

Hardening Security of Hardware IPs by Verifying Negative Space Formally

What is it:

A methodology for formally verifying how hardware IPs respond to malformed or malicious inputs, uncovering vulnerabilities missed by traditional simulation.

Why it matters:

- Strengthens trust in hardware designs.
- Enables early detection of hidden bugs.
- Promotes adoption of formal methods in security workflows.

Who's at Fault? A Look at Post-Quantum Cryptography and Fault Injection Attacks

What is it:

An analysis of fault injection attacks on ML-DSA, a post-quantum cryptographic algorithm, with evaluation of countermeasures to protect key integrity.

Why it matters:

- Highlights risks in quantum-era cryptography.
- Guides secure PQC implementation strategies.
- Encourages hardware-based protections like TRC.

Static RTL Analysis for Pre-Silicon (Security) Validation

What is it:

Cobra 3 is Intel's next-generation static RTL analysis tool designed to validate integration and security requirements in large-scale SoC designs before fabrication. It inspects RTL code without simulation, enabling early detection of structural flaws, misconfigurations, and security vulnerabilities.

Why it matters:

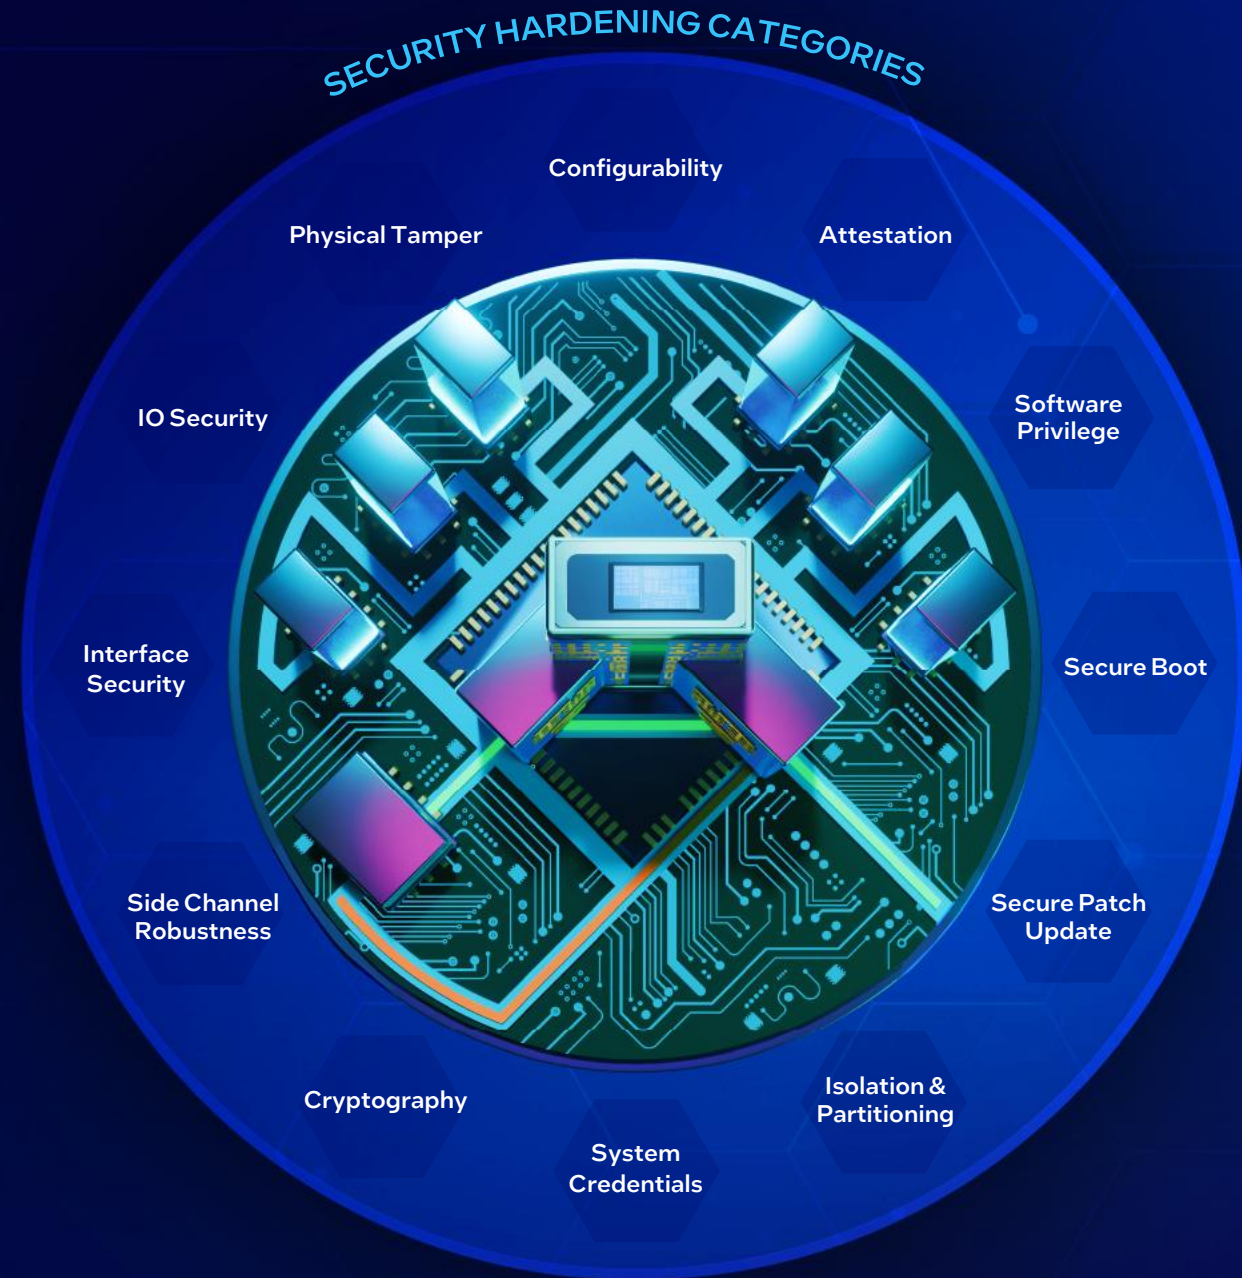
- Detects bugs before simulation or silicon, reducing costly late-stage fixes.
- Validates critical security features like Intel® Debug Protection Technology (DPT).
- Scales to millions of gates and thousands of modules in minutes.
- Supports CI/CD workflows and reusable validation templates.
- Identifies hardware CWEs such as improper access control and uninitialized security registers.
- Enables shift-left validation and fosters "security by design" culture.

Security Hardening Framework

Intel's security hardening framework defines the foundational security controls, design principles, and implementation standards that must be evaluated across all Intel reference platforms and their associated peripherals. It is designed to guide platform architects, silicon designers, and software engineers in building secure systems that meet enterprise-grade security expectations.

The framework spans three domains—platform, silicon, and software—and establishes a hardware-rooted chain of trust that extends from firmware through the software stack. It outlines mandatory and recommended requirements across critical security categories, including secure boot, attestation, cryptography, patch management, privilege isolation, and physical tamper resistance.

This holistic approach helps ensure that Intel platforms are resilient against known and emerging threats, enabling secure deployment across diverse environments, including cloud, edge, and enterprise infrastructure. By aligning with industry standards like NIST and IEEE, and incorporating best practices from Intel's SDL and security design principles, the framework supports compliance, auditability, and long-term platform integrity.



Formal Verification: Provable Security at Silicon Scale

What Is Formal Verification in Silicon?

Formal verification (FV) is a mathematical approach to proving the correctness of hardware designs by exhaustively exploring all possible input states and behaviors. Unlike simulation, which tests specific scenarios, FV uses formal specifications (assertions) to validate whether a design meets its intended properties across the entire state space.

In silicon security, FV is especially valuable for **negative space validation**—ensuring that hardware behaves securely even under **undefined, unsupported, or malicious inputs**.

Intel has mastered the extreme complexity of Formal Verification in silicon and is sharing its methodologies to strengthen ecosystem security through collaborative innovation.

Benefits of Formal Verification for Silicon

- 1. Exhaustive Coverage**
FV performs a complete search over the design's state space, making it possible to detect bugs that simulation might miss—especially in deeply embedded modules that are hard to trigger via test vectors.
- 2. Early Bug Detection**
FV enables early identification of critical bugs, such as:
 - Missing error detection logic
 - Malformed packets not dropped
 - Incorrect assertion of error signals
- 3. Ease of Use for Embedded Modules**
Once formal properties are defined, modules can be verified independently without needing full system integration. This accelerates validation for hard-to-reach logic blocks.
- 4. Scalable Specification**
FV allows concise expression of complex behaviors. For example, a single assertion like illegal, unsecure flow can capture an entire negative space, improving maintainability and auditability.
- 5. Proof of Fixes**
FV can mathematically prove that a bug fix works across all inputs and does not introduce new issues—something simulation cannot guarantee.
- 6. Security Hardening**
FV is increasingly used to verify architectural hardening against security threats, including malformed packet generation and protocol violations.


Challenges of Formal Verification in Silicon

- 1. Scalability**
The combinatorial explosion of input states makes FV computationally intensive, especially for large IP blocks or full-chip designs.
- 2. Specification Complexity**
Writing accurate and complete formal properties requires deep domain expertise. Poorly defined assertions can lead to false positives or missed bugs.
- 3. Tool and Resource Constraints**
FV tools require significant compute resources and specialized knowledge, which can limit adoption across teams.
- 4. Transmitter Verification Limitations in Simulation**
Simulation struggles to verify transmitter logic exhaustively. FV is often the only viable method to ensure transmitters do not generate malformed packets.
- 5. Collaboration Required**
Scaling FV across the industry demands collaboration among architects, designers, EDA vendors, and academia. Intel is actively working to lower entry barriers and promote community-driven solutions.

Negative Space Testing in Silicon Validation

Negative testing is a **quality assurance technique** used to evaluate how a system behaves when subjected to **invalid, unexpected, or malicious inputs**. Unlike positive testing, which confirms that a system works as intended under normal conditions, the goal of negative testing is to ensure that the system **fails gracefully** and **remains secure** when things go wrong.

In hardware security, particularly in silicon validation, negative testing is often referred to as "**negative space validation**"—testing the behavior of a design under undefined or unsupported conditions.



Negative Space Testing

Undefined, unsupported, or illegal behaviors

Examples:

- Illegal opcodes
- Malformed packets
- Out-of-bounds memory access
- Protocol violations

Not part of standard validation

Key Points of Negative Testing

1. Goal:

- To **identify vulnerabilities** and **edge-case failures**.
- To **validate robustness** against malformed or malicious inputs.
- To **ensure graceful degradation** or error handling.

2. Input Classification:

- Positive Inputs: Valid, expected inputs.
- Negative Inputs: Invalid, unexpected, or malicious inputs (e.g., illegal opcodes, unsupported packet lengths, malformed data).

3. Techniques:

- Software: Fuzzing, boundary value analysis, exception handling tests.
- Hardware: Formal verification (FV), simulation of undefined behaviors, architectural hardening of negative space.

4. Coverage Areas:

- Interfaces: Internal, external, and programmable interfaces.
- Protocols: Message formats, address spaces, encoding schemes.
- Security Domains: Isolation boundaries, debug access, firmware rollback, etc.


5. Tools & Methods:

- Software Fuzzing: Widely used in software but less effective for hardware.
- Formal Verification: Preferred in hardware to systematically explore negative input spaces.
- Negative Space Engineering: A structured approach to document and test undefined behaviors in hardware designs.

Security Hardening Results

Intel's hardening efforts have limited externally discovered processor vulnerabilities to just one in products released since the beginning of 2023. In contrast, external researchers have identified 10 vulnerabilities in AMD's latest processors.

Externally Found Processor Vulnerabilities in Products Released Since the Beginning of 2023



VULNERABILITY SPOTLIGHT: EntrySign

What: A medium severity (CVSS 6.4) vulnerability that allows an attacker to bypass AMD's microcode signature verification mechanism on Zen 1 through 5 processors. AMD used a 128-bit AES-CMAC with a secret key shared across CPUs. Intel platforms are not affected.

Impact: Potential for an attacker to forge valid microcode patches and installation of malware such as bootkits and rootkits that persist after reboots or OS reinstalls.

Engineering Decision: While AMD used digital signatures (RSA) to authenticate microcode updates, the process relied on a cryptographic function that was not robust enough to prevent collisions or falsification.

Secure Design Gap: The primary gap wasn't just a "bug" in the code; it was a fundamental architectural failure in how the CPU verified its most sensitive instructions.

CWE-327 (Use of a Broken or Risky Cryptographic Algorithm), CWE-321 (Use of hard-coded Cryptographic Key), and CWE-347 (Improper Verification of Cryptographic Signature) all apply to this issue.

Intel View: Intel products use an industry-standard SHA-256 secure hash for verification.

Who: EntrySign was discovered and disclosed by Google researchers Josh Eads, Kristoffer Janke, Eduardo 'Vela' Nava, Tavis Ormandy, and Matteo Rizzo in March 2025.



VULNERABILITY SPOTLIGHT: StackWarp

What: A low severity (CVSS 3.2) vulnerability, StackWarp is a microarchitectural vulnerability in AMD Zen-series CPUs that allows a malicious host to hijack a confidential virtual machine's control flow by exploiting a synchronization failure in the processor's stack engine. Intel platforms are not affected.

Impact: An attacker could hijack control flow to redirect execution to arbitrary locations within the guest VM, bypass security checks to escalate privilege, or gain full control over the confidential VM through remote code execution.

Engineering Decision: This is a microarchitectural flaw stemming from a hardware design choice meant for performance.

Secure Design Gap: The design is reliant on encrypted memory but did not account for the fact that an attacker can still influence the CPU's internal pointers (like the Stack Pointer). If an attacker can force the CPU to look at the "wrong" piece of encrypted memory, they don't need to decrypt it to cause damage. They can trick the VM into executing its own code in a way that gives the attacker control.

CWE-1264: Hardware Logic with Insecure De-Synchronization between Control and Data Channels: This is the most accurate "root cause" CWE for StackWarp. It describes a failure where hardware logic (the stack engine) becomes out of sync with the intended state of the system (the architectural stack pointer).

Intel View: Intel's microarchitecture does not share the same "deferral" logic or the specific thread configuration vulnerability found in AMD Zen microarchitecture. The Intel® TDX Module is architecturally required to save, scrub, and restore the entire CPU state (including general-purpose registers like the Stack Pointer) during every transition. The goal is that no "stale" or "desynchronized" data from the host's microarchitectural state can leak into or affect the guest's execution.

Who: StackWarp was discovered by Ruiyi Zhang, Tristan Hornetz, Daniel Weber, Fabian Thomas, Lukas Gerlach, and Michael Schwarz of CISPA, Youheng Lü of SCHUTZWERK GmbH.



Industry Collaboration & Leadership

Intel leads the silicon industry in security innovation by actively shaping global standards, driving open collaboration, and sharing breakthrough technologies, including its patented Memory Tagging Technology (MTT) to elevate platform robustness across the entire ecosystem.

Securing The Future Through Open Collaboration and Leadership

Intel is driving the evolution of trusted computing by actively contributing to and frequently leading global standards initiatives in hardware security, firmware integrity, and confidential computing. Rather than waiting for external guidance, Intel takes a proactive role in defining the architectural frameworks and security technologies that will form the foundation of trusted platforms for the foreseeable future.

CONFIDENTIAL COMPUTING

With Intel® SGX, Intel was the first silicon vendor to bring Confidential Computing technologies to market in 2015. Through these groups, Intel continues to drive innovation in this space through collaboration and leadership.

AI SECURITY

As a founding member of the Coalition for Secure AI, Intel recognizes that AI security requires deep industry engagement and continues to help drive advancements in this space.

SECURITY ASSURANCE

Solid product security assurance is the foundation security is built on whether it be software or hardware. Intel brings knowledge and leadership as it engages with industry groups that help to build and standardize that foundation. Examples include engaging with MITRE to define and introduce Common Weakness Enumerator (CWE) for hardware-specific vulnerabilities. Through the FIRST organization, multiple Intel employees participate in driving standards as well as creating a framework for those implementing or modernizing their Product Security Incident Response Team (PSIRT).

INTERNATIONAL STANDARDS

From vulnerability management to the TPM 2.0 specification, Intel has been engaged in critical international standards development for security for many years as a participant, co-author, and major contributor.

HARDWARE SECURITY

Intel is actively engaged with many efforts to drive greater hardware-level security not just for its customers, but also for those of its competitors. One example is sharing an Intel patent-pending technology called Memory Tagging Technology (MTT) that, once implemented, will have a dramatic impact on the world's most prevalent software exploit: memory safety violations. Intel is working with the x86 Ecosystem Advisory Group to ensure all customers benefit from this innovative silicon based technology.

FIRMWARE SECURITY

Intel's decision to release the EFI code as open source fostered industry-wide collaboration, accelerating UEFI adoption and enabling a more secure, standardized, and innovative firmware ecosystem.



References

1. As of December 2024, based on MITRE data report and [blog https://community.intel.com/t5/Blogs/Tech-Innovation/Artificial-Intelligence-AI/Intel-AI-PCs-Deliver-an-Industry-Validated-Defense-vs-Real-World/post/1650954](https://community.intel.com/t5/Blogs/Tech-Innovation/Artificial-Intelligence-AI/Intel-AI-PCs-Deliver-an-Industry-Validated-Defense-vs-Real-World/post/1650954).
2. As measured by [ABI Research](#)
3. Vulnerability analysis based on public data accessible from intel.com and amd.com.
4. Gartner Identifies the [Top Strategic Technology Trends for 2026](#)
5. ChkTag: x86 Memory Safety: <https://community.intel.com/t5/Blogs/Tech-Innovation/open-intel/ChkTag-x86-Memory-Safety/post/1721490>

intel security

Notices & Disclaimers

Intel technologies may require enabled hardware, software, or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

AI features may require software purchase, subscription or enablement by a software or platform provider, or may have specific configuration or compatibility requirements. Data latency, cost, and privacy advantages refer to non-cloud-based AI apps. Learn more at intel.com/AIPC.

All versions of Intel vPro® require an eligible Intel processor, a supported operating system, required connectivity technology, firmware enhancements, and other hardware and software. Remote management requires a network connection; must be a known network for Wi-Fi out-of-band management. Details at intel.com/vpro.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.