ZERO TRUST STARTS AT THE *SILICON LEVEL*

Why endpoint hardware is the foundation of trust

AMD
together we advance_

EXECUTIVE SUMMARY

Zero Trust has become the default security posture for modern enterprises. But many implementations still treat endpoints as downstream recipients of policy rather than upstream sources of trust. In hybrid, distributed environments, that gap matters. The hidden risk isn't policy design but enforcement, dependent on endpoints unable to prove their integrity under attack.

This paper argues that Zero Trust is enforceable only when it starts at the silicon level. Hardware-rooted security provides a tamper-resistant foundation for device identity, boot integrity, isolation and runtime attestation – capabilities that software alone cannot reliably guarantee. By anchoring trust in hardware, endpoints become verifiable enforcement points rather than assumed participants.

For Zero Trust to move from architectural intent to operational reality – including AI workloads – trust must begin where enforcement begins: at the silicon layer.

THE CORE TENETS OF ZERO TRUST

Zero Trust is implemented through consistent architectural principles, such as:

- **Never trust, always verify.** No implicit trust for users, devices or workloads. Verification is continuous.
- **Least-privilege access.** Grant only the minimum access required; reduce blast radius by design.
- **Assume breach.** Architect for containment and resilience, not only prevention.
- **Context-aware policy.** Decisions are dynamic, informed by identity, device posture and risk.
- **Microsegmentation.** Limit lateral movement by isolating systems and workloads.

All these principles presuppose a trustworthy endpoint.

THE ZERO-TRUST IMPERATIVE AND THE ENDPOINT GAP

Endpoints remain one of the most targeted and operationally inconsistent layers in many security architectures. Research and field experience continue to show adversaries frequently choose endpoints as the path of least resistance, using them to pivot into more sensitive systems.[1]

Zero Trust intent is often expressed through centralized identity, access policies and monitoring. But **Zero Trust enforcement** must happen on the device during boot, at runtime and under attack.

THE "SOFTWARE-FIRST" ZERO TRUST FAILURE MODE

When firmware, boot processes or memory are compromised, software-based controls continue operating on an unverified platform. Policy may still be defined and monitored, but enforcement is no longer reliable. Hardware-rooted trust closes this gap by turning device integrity into a verifiable signal rather than an assumption.

WHY ZERO TRUST CANNOT BE SOLVED AT THE OS OR SOFTWARE LAYER ALONE

When integrity is compromised at the OS level, even well-designed software controls lose their reliability. This becomes critical in scenarios such as:

- **Firmware tampering and boot-chain manipulation.** Changes to early boot sequences can be stealthy, persist across reinstalls and evade OS-level telemetry.
- **Memory scraping and physical access attacks.** If data can be extracted from RAM, encryption at rest is not enough.
- **Device impersonation and policy drift.** If device identity and posture cannot be validated continuously, “trusted device” becomes a weak claim.

77 Bank overcame PC performance issues while improving Zero Trust security with AMD Ryzen™ PRO processors. Learn more here.

THE SILICON LAYER AS THE FOUNDATION OF ZERO TRUST

The silicon layer provides:

- **An immutable root of trust** for boot and measurement.
- **Protected execution environments** for security-sensitive operations.
- **Hardware-enforced isolation** independent of OS health or user behavior.
- **A chain of trust** that can be extended into runtime attestation.

This bridges the gap between "policy defined" and "policy enforceable." If the device can attest to its integrity – starting from the moment power is on – then higher-layer controls can make decisions based on a trustworthy signal rather than an assumption. Device-level encryption enables tailored access controls, which, in turn, can heighten data security.

AMD PRO solutions illustrate how these principles can be implemented in commercial endpoints today. Through hardware rooted in trust, memory encryption, firmware protections, standards-based manageability and platform stability, these systems demonstrate how Zero Trust intent can be translated into enforceable control at scale, even under attack.

SOFTWARE-ONLY ZERO TRUST

- Integrity inferred after OS load.
- Signals generated by software agents.
- Enforcement depends on assumed platform health.
- Can break silently under firmware or memory compromise.

HARDWARE-ANCHORED ZERO TRUST

- Integrity verified at power-on.
- Signals rooted in silicon.
- Enforcement persists under OS degradation.
- Enables reliable containment and recovery.

> "If you can't see or manage a device when the OS is down, you don't have Zero Trust."

TUSHAR OZA
Director of Product Management, Product Security Office, AMD

WHY ENFORCEABLE ZERO TRUST CHANGES RISK ECONOMICS

The value of enforceable Zero Trust shows up in three operational metrics:

- **Time to recover:** How quickly compromised endpoints can be returned to a trusted state.
- **Fleetwide blast radius:** How far an incident propagates before containment.
- **Audit defensibility:** The ability to demonstrate consistent enforcement and control during post-incident review.

Modeling suggests that platforms designed for enforceable Zero Trust can also deliver measurable operational benefits. Third-party analysis from Signal65 indicates that commercial laptops with AMD PRO solutions can produce up to **$53 million in combined modeled time value and up-front cost savings** in large enterprise scenarios, based on defined assumptions and configurations.[2]

The business case for hardware-enforced Zero Trust is not rooted in incremental security improvement but in risk reduction at scale. When endpoints can be verified, isolated, recovered and managed consistently, organizations reduce incident duration, limit blast radius and improve resilience under real-world conditions. This is where Zero Trust moves from architectural intent to measurable business value.

Proving Feasibility: Hardware-Enforced Zero Trust on Commercial Endpoints

AMD PRO provides a concrete example of how hardware-rooted identity, integrity, isolation and manageability can be implemented today on standard business PCs, making Zero Trust enforcement practical rather than aspirational.3

Hardware root of trust and platform integrity

At the heart of hardware-rooted trust are mechanisms that measure and authenticate boot components and support ongoing verification after boot. All AMD PRO solutions use the AMD Secure Processor as a dedicated on-chip security subsystem that anchors a hardware root of trust and supports secure boot flows.

The AMD Dynamic Root of Trust Measurement flow creates a chain of trust between firmware and runtime, enabling the OS/hypervisor to request remeasurement and attestation before executing sensitive operations.

RESULT: A verifiable integrity signal that higher-layer security controls can depend on.

Memory-level protection and isolation

Memory is a high-value target, especially in physical-access scenarios or advanced attacks designed to extract secrets from RAM. AMD Memory Guard provides full-system memory encryption (OEM enablement required), helping protect data in memory from physical DRAM attacks and certain classes of offline extraction.

AMD platforms also include additional protections aimed at reducing the exploitability of privileged levels. AMD Shadow Stack, for example, is designed to mitigate certain control-flow hijacking techniques, which can include return-oriented programming (ROP) attacks.

The "shadow stack," or a parallel stack, can mitigate ROP attacks by storing a copy of return addresses with specialized hardware. The copy is then compared with the normal program stack of returns, therefore "catching" disparities in the "jump and return" process.

RESULT: If attackers cannot extract usable data from memory, the endpoint is less likely to become a launchpad for broader compromise.

> **AI threats move faster than networks. Security intelligence has to exist where the attack happens — on the endpoint."**
>
> **TUSHAR OZA**
> Director of Product Management,
> Product Security Office, AMD

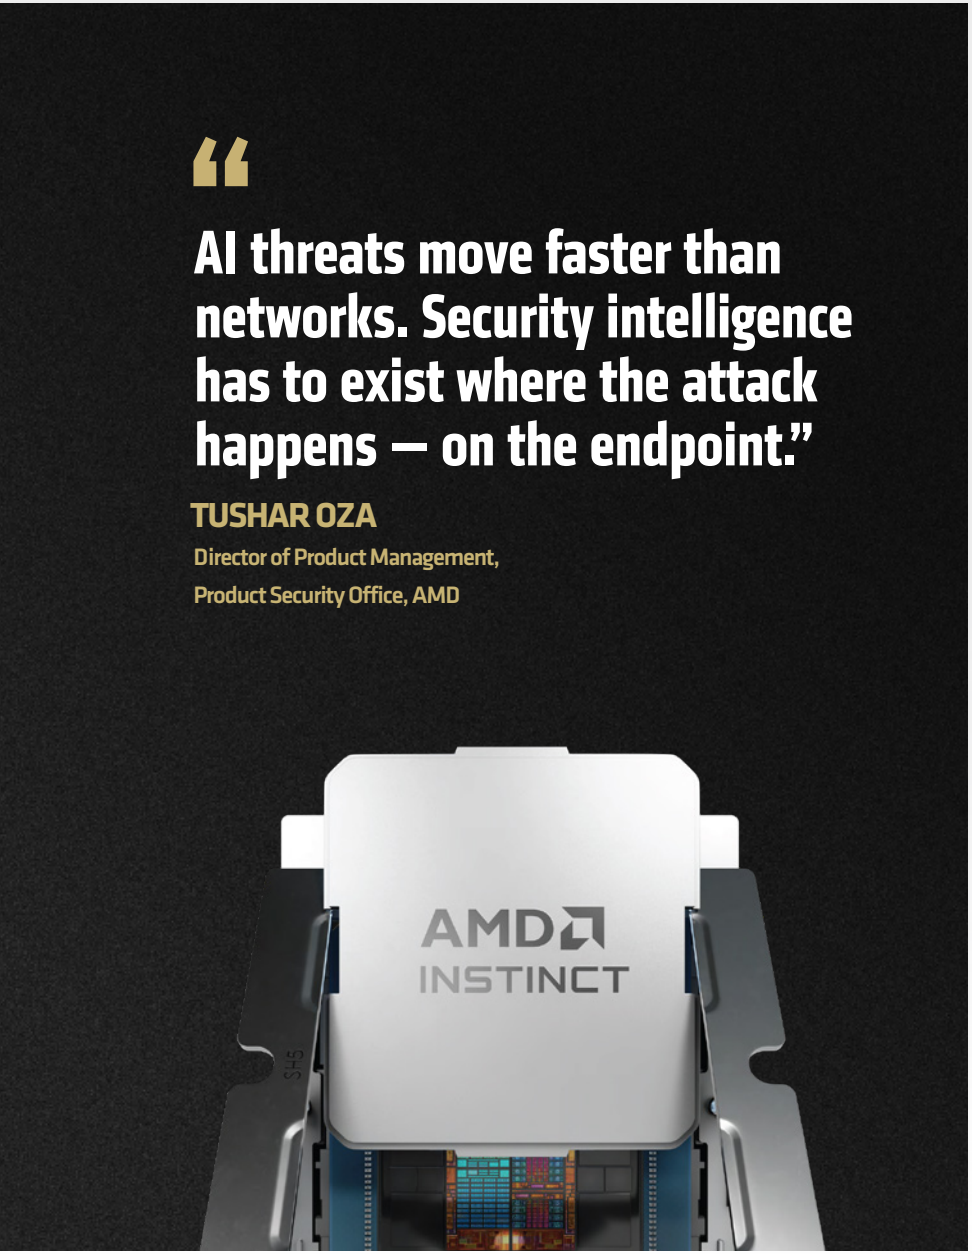

Enterprise security alignment and transparency

Zero Trust architectures in Windows enterprise environments often depend on a broader ecosystem. The AMD PRO platform aligns with modern OS security models, including support for Microsoft Secure-core PC–aligned protections and integration with technologies such as TPM 2.0 and Microsoft Pluton (OEM enablement required). Pluton, built into supported CPUs as a secure cryptography processor, is intended to strengthen the protection of credentials and keys and to receive updates through Windows Update.

RESULT: A security approach that integrates into an ecosystem rather than operating in isolation.

AI CHANGES THE THREAT — AND THE ENDPOINT

Security leaders are balancing AI-driven workflows with governance, privacy and compliance. This creates two parallel security requirements:

- **Protect AI workloads running on endpoints.** If AI models and data are executed locally, the platform must prevent tampering, protect prompts/outputs and enforce policy boundaries.
- **Use AI to protect endpoints.** Detection and response can become faster and more contextual when intelligent analysis occurs closer to the threat.

The shift is about completeness and latency: As endpoints become AI-capable nodes, some intelligence must move closer to where threats occur.

Why local AI matters for security

For local AI security, some pragmatic drivers include:

- **Speed of response.** On-device inference can support quicker local decisions.
- **Resilience in constrained environments.** Air-gapped or intermittently connected environments cannot depend on cloud inspection.
- **Privacy and compliance.** Some sensitive contexts benefit from local analysis where signals do not need to be continuously streamed off-device.

AMD Ryzen™ AI PRO processors add on-device AI acceleration – up to 55 TOPS in certain Ryzen AI PRO 300 Series configurations – creating headroom for local inference and emerging endpoint security models.[4]

None of this removes the need for cloud-scale intelligence. The strongest model is often hybrid: cloud for broad correlation and training; endpoint for local detection, triage and enforcement.

What becomes possible with AI-capable endpoints

AI-capable endpoints can support richer, more adaptive security behaviors, especially when paired with hardware-enforced isolation and integrity:

- **Anomaly detection** and behavioral analysis are executed locally.
- **Intelligent containment** that can operate even when connectivity is poor.
- **Higher-fidelity telemetry generation** without relying on constant cloud inspection.

AI-ready endpoints can become active security participants rather than passive policy recipients, provided the device remains verifiable and manageable under attack.

ZERO TRUST AT SCALE REQUIRES MANAGEABILITY, NOT COMPLEXITY

Endpoint management has become more complex in recent years, driven by factors such as device proliferation, remote work, and the frequency of OS updates.[5] Hardware-level manageability is a Zero Trust enabler because it helps preserve visibility and control even under OS failure conditions.

OPEN STANDARDS ENABLE CONSISTENCY

AMD PRO Manageability is built around open standards, including DMTF DASH, and supports modern encryption and authentication protocols such as TLS 1.2 and TLS 1.3.

This matters for two reasons:

- **Multivendor manageability.** DASH-based approaches support consistent management across diverse fleets, including mixed environments.
- **Faster incident response.** When organizations face large-scale endpoint disruptions, out-of-band capabilities can help IT teams diagnose and remediate faster at scale.

These capabilities align with modern enterprise expectations for secure, standards-based fleet management.

STABILITY AS A SECURITY ADVANTAGE

Drift undermines Zero Trust. If fleets run inconsistent firmware, drivers or configurations, enforcement becomes uneven and auditability weakens.

AMD PRO emphasizes predictable platform stability, including at least 18 months of planned software stability and hardware availability with an additional five years of software support after the final end-of-life ship date. This stability supports consistent enforcement and auditability across multiyear refresh cycles.

CONCLUSION: ZERO TRUST STARTS AT THE SILICON LAYER

With a hardware-enforced foundation, endpoints can act as verified enforcement points that:

- Prove identity and integrity from boot through runtime.
- Protect data in use and at rest.
- Enforce isolation and reduce lateral movement.
- Remain manageable and recoverable, even when the OS is impaired.

This reframes the endpoint from liability to the control plane.

With security threats outpacing device refresh cycles, the enterprise needs partners with experience layered into every piece of the computing infrastructure. Enabled by AMD PRO solutions, Zero Trust ambitions are possible with a security partner ecosystem and flexible deployments meant to complement existing tech stacks.

A hardware-enforced foundation for endpoint Zero Trust includes root of trust, memory protection, firmware integrity mechanisms, standards-based manageability and platform stability. These capabilities strengthen Zero Trust today while laying the foundation for reliable, AI-enabled security models.

Enterprises must start where trust begins: **at the silicon layer.**

Footnotes

1 Supporting Strong Cybersecurity Health for Next-Generation PCs (ESG / TechTarget).

2 Signal65 Lab Insights: AMD Ryzen™ AI PRO Commercial Value Analysis. “AMD RyzenTM AI PRO Commercial Value Leadership.”

3 AMD PRO Technologies Whitepaper, A LOOK AT AMD PRO SECURITY AND THE AMD FRAMEWORK FOR SECURE, MANAGEABLE, AND RELIABLE BUSINESS PCS (February 2025).

4 Top Reasons Enterprises Are Choosing AMD AI PC Solutions (TOPS + positioning).

5 AMD PRO Manageability Infographic / ESG Research Summary (June 2025).

AMD PRO TECHNOLOGIES

Help protect your business with a complete set of security features, robust manageability tools, and enterprise-grade stability and reliability.

LEARN MORE

AMD
together we advance_

Studio
by informa techtarget

Expert led.
Impact driven.

Studio is Informa TechTarget's global content studio offering brands an ROI rich tool kit: Deep industry expertise, first-party audience insights, an editorial approach to brand storytelling, and targeted distribution capabilities. Our trusted in-house content marketers help brands power insights-fueled content programs that nurture prospects and customers from discovery through to purchase, connecting brand to demand.

Learn more