

Trusted Execution Environments (TEEs: TDX/SEV-SNP) Remote Attestation

Current-source review, evidence-bounded adoption decision, explicit limitations, reproducibility controls, and preserved source research note.



PERMANENT ID

RES-023

EVIDENCE

A-

ADOPTION

CONDITIONAL ADOPTION

FRESHNESS

ACTIVE WATCH

Freshness review: 2026-09-17 / Next scheduled review: 2026-12-16

Required Research Fields

Each field remains independently reviewable; evidence strength does not imply production authority.

EVIDENCE GRADE	A-	Strength of the retained source set and technical basis.
SOURCE AUTHORITY	3 registered authorities	Intel Trust Domain Extensions Documentation; AMD SEV-SNP Documentation; Confidential Computing Consortium
FRESHNESS	ACTIVE WATCH	Reviewed 2026-09-17; dynamic sources require event-driven revalidation.
CONFLICTS	VISIBLE / NOT SILENT	Differences between specification, vendor behavior, research claims, and reproduced evidence must remain explicit.
LIMITATIONS	EXPLICIT	No certification, procurement approval, legal conclusion, or unbounded production claim is created by this report.
REPRODUCIBILITY	REQUIRED	Material claims require exact versions, representative data, failure cases, artifacts, and repeatable acceptance criteria.
ADOPTION POSTURE	CONDITIONAL ADOPTION	Adopt confidential-computing profiles for sensitive workloads where attestation, firmware trust, TCB scope, side-channel exposure, key release, and recovery are independently validated.
REVIEW TRIGGER	2026-12-16	Scheduled at 90 days; earlier on material source, vulnerability, implementation, or contradictory-evidence change.

Authority Register and Freshness Delta

Primary-source lineage is preserved; current-source changes are separated from unperformed implementation claims.

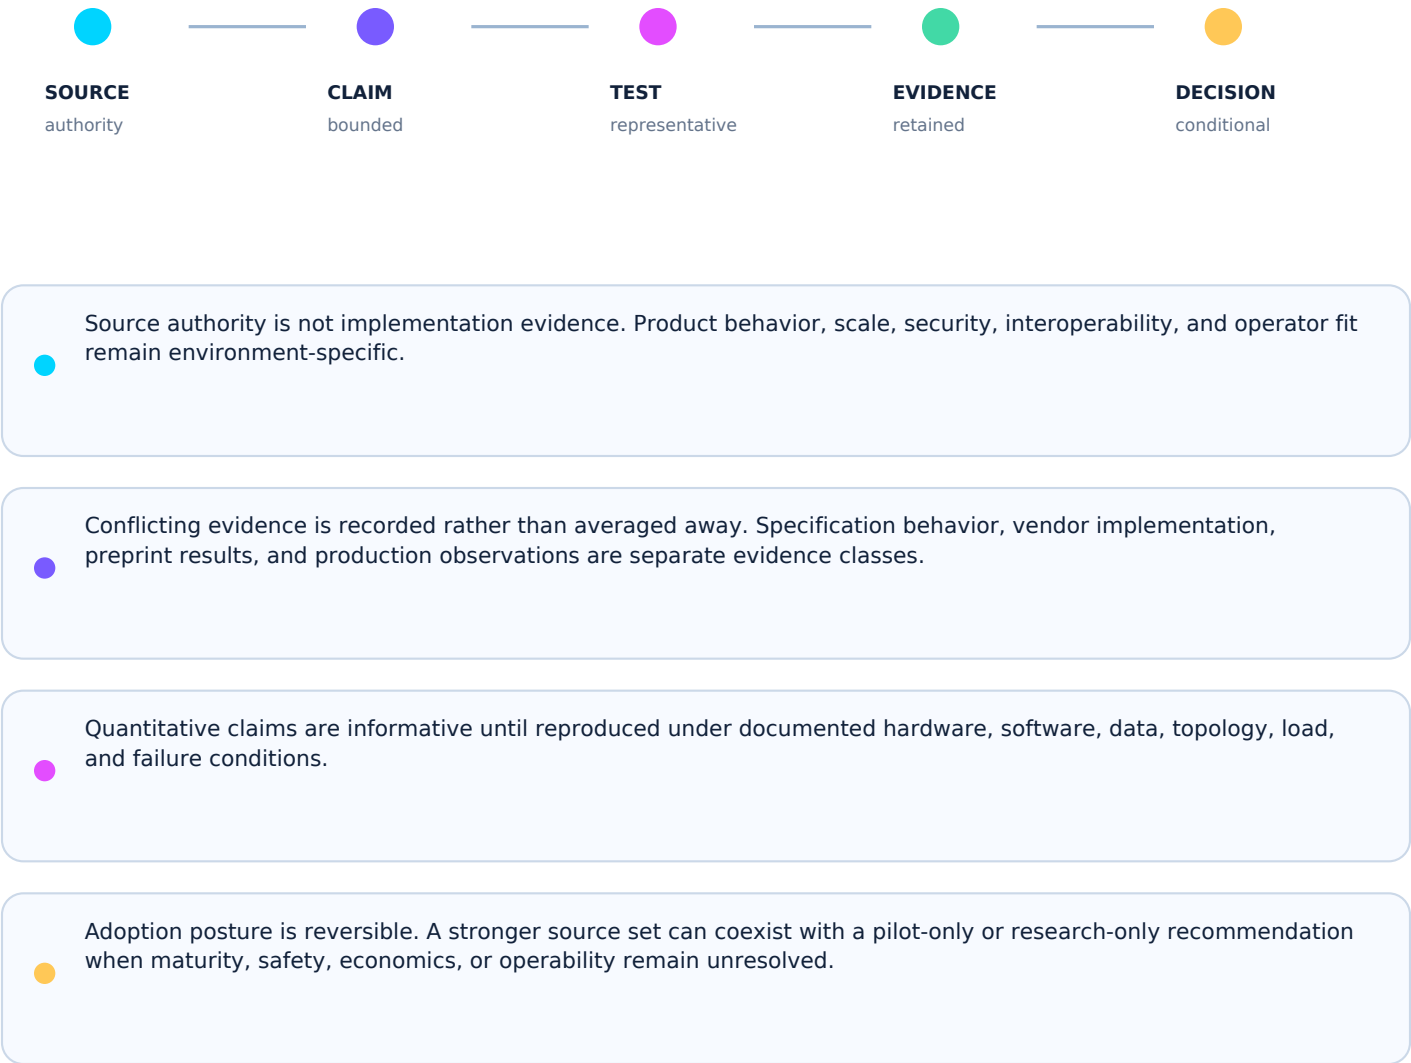
AUTHORITY 01	Intel Trust Domain Extensions Documentation https://www.intel.com/content/www/us/en/developer/tools/trust-domai
AUTHORITY 02	AMD SEV-SNP Documentation https://www.amd.com/en/developer/sev.html
AUTHORITY 03	Confidential Computing Consortium https://confidentialcomputing.io/

2026-09-17 FRESHNESS DELTA

Primary-source register retained and freshness controls advanced to the current review date. No new product or laboratory performance claim is introduced without reproduced evidence.

Freshness policy: scheduled review + immediate review on source supersession, material release, vulnerability, retraction, or contradictory evidence.

Evidence Must Survive Contact With Reality



Decision Gate and Validation Plan

ADOPTION POSTURE

CONDITIONAL ADOPTION

Adopt confidential-computing profiles for sensitive workloads where attestation, firmware trust, TCB scope, side-channel exposure, key release, and recovery are independently validated.

- 1 / SOURCE

Pin exact versions, publication state, retrieved artifacts, and source hashes.
- 2 / PROTOTYPE

Demonstrate the core mechanism with representative data and instrumented execution.
- 3 / FAILURE

Exercise malformed input, dependency loss, stale state, overload, recovery, rollback, and misuse.
- 4 / BASELINE

Compare against an incumbent, classical, or simpler architecture using the same workload.
- 5 / ACCEPT

Record acceptance criteria, residual limitations, owner, expiration, and revalidation triggers.

EXPIRATION / REVIEW TRIGGER

Next scheduled review: 2026-12-16 (90-day cadence). Review sooner after material source revision, breaking implementation release, critical vulnerability, retraction, benchmark contradiction, changed workload, or changed legal/safety boundary.

8. Complete Source Research Note

SOURCE-LINEAGE NOTICE

The following material preserves the complete original source note, excluding only the conversational wrapper and outer Markdown fence. Legacy status labels and absolute language are retained for traceability and do not override the candidate status, limitations, or adoption decision in this edition.

Document ID: RES-023 Version: 1.0.0 (Definitive Registry Edition) Status: Published Research Note Classification: Public Organization: Mythos Systems (MYTHOS AI) Owner: Aaron Stovall Effective Date: 2026-07-16 Applies To: Security architects, cloud platform engineers, and AI engineers designing sovereign cloud deployments, zero-trust AI inference, and confidential computing pipelines Companion Standards: MYTHOS-SEC-0009 (Confidential Computing & TEE Standard), MYTHOS-HW-0002 (Trusted Compute & DPU), MYTHOS-DAT-0003 (Data Sovereignty & Egress), MYTHOS-EMT-0002 (Sovereign AI)

The architecture of cloud data sovereignty has failed. Organizations attempt to secure proprietary AI model weights and sensitive user data by encrypting it at rest (disk) and in transit (TLS). However, to process the data, it MUST be decrypted into the server's system RAM. At that moment, the data is fully exposed to the cloud provider, the hypervisor, and any malicious insider or privileged malware on the host. "Encrypt at rest" is a legal fiction when the cloud provider holds the keys and the RAM.

Trusted Execution Environments (TEEs), specifically Intel Trust Domain Extensions (TDX) and AMD Secure Encrypted Virtualization - Secure Nested Paging (SEV-SNP), shatter this vulnerability. They introduce Confidential VMs (CVMs), where the memory and CPU registers of a virtual machine are encrypted by dedicated hardware on the physical CPU package. The hypervisor is demoted to an untrusted resource allocator; it cannot read the VM's memory.

This research note defines the architectural dynamics of TEEs. It dissects the mechanics of hardware memory encryption, the cryptographic measurement of boot states, and the Remote Attestation pipeline. Understanding these dynamics is a prerequisite for building the zero-trust, sovereign AI inference clusters mandated by the Mythos Standards, where an organization can mathematically prove that a cloud provider cannot read their AI weights.

To understand TEEs, we must move from software-based encryption to hardware-rooted memory isolation.

1.1 Hardware Memory Encryption (MKTME / AES-XTS)

Modern x86 processors feature a Memory Encryption Engine (MEE) integrated into the memory controller.

- The Mechanic: When a TEE (TDX Guest or SEV-SNP VM) is launched, the CPU generates a unique, ephemeral Data Encryption Key (DEK) for that specific VM. This key never leaves the CPU silicon.
- The Execution: When the VM writes data to RAM, the CPU encrypts it on the fly. When the VM reads data, the CPU decrypts it on the fly. If the hypervisor, the host OS, or another VM attempts to read that RAM, they see only ciphertext.

1.2 Intel TDX (Trust Domain Extensions)

TDX introduces a new mode: the Trust Domain (TD).

- The Architecture: A TD runs in a hardware-isolated enclave that spans multiple vCPUs. The hypervisor (KVM/ESXi) manages memory allocation (assigning pages), but it cannot access the content of those pages. A secure arbiter (the TDX Module) sits between the hypervisor and the TD, enforcing memory access rights.

1.3 AMD SEV-SNP (Secure Nested Paging)

SEV-SNP protects VMs by encrypting the page tables and the memory itself.

- The Architecture: It prevents the hypervisor from performing "malicious hypervisor attacks" (e.g., remapping a VM's memory page to read its contents or causing a denial-of-service by altering page tables). It uses a Reverse Map Table (RMP) to ensure that only the designated VM can access its own pages.

While TEEs protect data in-use, they introduce severe physical and operational constraints that standard VMs do not face.

2.1 The Attestation Cold-Start Penalty

Before a TEE can process sensitive data, the client must verify the TEE's hardware state (Remote Attestation).

- The Flaw: This cryptographic challenge-response protocol adds 5 to 15 seconds of latency to the VM boot process.
- The Impact: TEEs cannot be used for rapid, sub-second serverless cold starts without maintaining warm pools of pre-attested CVMs.

2.2 The Memory Overhead

Memory encryption introduces a physical performance tax.

- The Constraint: Every memory read/write requires an AES-XTS cryptographic operation. Additionally, TEEs require a "memory integrity tree" (to prevent replay attacks) that consumes additional RAM.
- The Impact: TEE workloads typically incur a 5% to 15% performance overhead compared to plaintext VMs. For ultra-low-latency LLM inference, this overhead must be factored into the SLO.

2.3 The Side-Channel Blind Spot

TEEs protect against logical memory reads, but they do not inherently protect against micro-architectural side-channel attacks.

- The Flaw: Attackers can measure cache timing, power fluctuations, or electromagnetic emissions to infer what the TEE is computing.
- The Mitigation: TEEs MUST be paired with OS-level mitigations (e.g., page table isolation, cache flushing) and hardware-constant-time cryptographic libraries.

The true power of a TEE is not memory encryption; it is the ability to mathematically prove to a remote client that the memory is encrypted and running the correct software.

3.1 The Cryptographic Measurement (The Quote)

When the TEE boots, the CPU hashes every component: the firmware, the bootloader, the kernel, and the initial ramdisk.

- The Mechanic: These hashes are recorded in Platform Configuration Registers (PCRs) or TDX Measurement Registers (TMRs). The CPU signs this measurement log with a hardware-attested key (derived from a burned-in AMD/Intel root of trust), producing a "Quote."

3.2 Remote Attestation (The Verification)

The client (e.g., the PANTHEON control plane) requests the Quote.

- The Verification: The client checks the signature against a public key infrastructure (e.g., AMD Key Distribution Service or Intel SGX PCCS). If the signature is valid, the client knows the Quote originated from genuine hardware.
- The Policy Match: The client then compares the measurement hashes in the Quote against a known-good database (e.g., a signed manifest of the approved OS image). If the hashes match, the client mathematically proves that the VM is running the exact, unmodified approved software.

3.3 Sealing AI Weights (The Payload Release)

Once attestation is successful, the client can securely provision secrets.

- **The Mechanic:** The client retrieves the TEE's public key (derived from the hardware measurement). The client encrypts the AI model weights (or database credentials) using this key and sends the ciphertext to the TEE.
- **The Absolute Guarantee:** Only that specific, verified TEE possesses the private key to decrypt the weights. If the cloud provider tries to copy the weights to another machine, or alters the OS to read them, the hardware measurement changes, the private key is lost, and the weights remain encrypted ciphertext.

The fundamental shift of TEEs is the abolition of implicit trust in the cloud provider.

- **Proprietary AI Protection:** An organization can deploy a 70B parameter proprietary LLM to AWS or Azure. The model weights are sealed to the TEE. The cloud provider cannot extract the weights, even with physical access to the server.
- **Privacy-Preserving Inference:** Users can submit encrypted prompts to the TEE. The TEE decrypts the prompt internally, runs the inference, and returns the encrypted result. The cloud provider cannot read the user's prompt or the model's output.
- **Verifiable Compute:** Regulators can audit the attestation logs to mathematically prove that a financial AI model was running on approved, unmodified hardware.

To deploy TEEs in production, the Mythos Architecture enforces strict attestation gates and operational boundaries.

5.1 The Zero-Trust Provisioning Gate

The Mythos control plane MUST NOT push AI weights or data to a CVM until a valid attestation Quote is verified.

- **The Mitigation:** The deployment pipeline (Clio/Morta) intercepts the deployment. The CVM boots with a generic OS. The CVM sends its Quote to the control plane. The control plane verifies the Quote against a SBOM/Reference Value. Only upon a 100% match are the weights streamed via a sealed TLS tunnel to the TEE.

5.2 Attested Observability (The Verifiable Telemetry)

If the TEE is a black box, how does the organization know what the agent inside it is doing?

- **The Mitigation:** The observability pipeline (Argus/OpenTelemetry) MUST sign the cognitive traces inside the TEE using the TEE's hardware key. The downstream SIEM verifies the signature. If the hypervisor attempts to alter the logs in transit, the signature breaks, alerting the SOC to a host compromise.

5.3 The TEE-FHE Hybrid (Maximum Security)

For Level 5 sovereign data (e.g., classified CUI), TEEs alone are insufficient because they rely on the integrity of the CPU firmware.

- **The Mitigation:** The architecture MUST combine TEEs with Fully Homomorphic Encryption (FHE). The user sends an FHE-encrypted prompt. The TEE receives the ciphertext. The TEE performs the LLM inference homomorphically (utilizing the TEE's hardware to accelerate FHE). The result is returned as FHE ciphertext. Even if the TEE's memory encryption is somehow bypassed, the data is still mathematically encrypted by FHE.

Trust in cloud providers is a liability, not a security strategy. Legal agreements and compliance audits cannot prevent a hypervisor zero-day or a malicious insider from reading plaintext RAM. Trusted Execution Environments, via hardware memory encryption and remote attestation, transform the cloud into a blind, untrusted utility. By sealing proprietary AI weights to verified hardware states, we achieve the ultimate mandate of data sovereignty: the compute happens in the cloud, but the ownership remains absolute.

Encryption at rest is a legal fiction; encryption in-use is a mathematical truth.

A hypervisor with read access is an insider threat; a demoted hypervisor is a utility.

A promise of security is a liability; a hardware attestation is a proof.
A sealed weight is a cipher; an attested TEE is the key.

The cloud provider sees only ciphertext; the math sees all.

- > Clouds may be shared.
 - > Memory sovereignty must be absolute.
-

End of Research Note RES-023

© 2026 Mythos Systems. This research is published for public reference. Attribution required.

9. Source Register

Source	Authority and Status	Freshness Control
Intel Trust Domain Extensions Documentation https://www.intel.com/content/www/us/en/developer/tools/trust-domain-extensions/overview.html	Primary vendor specification and guidance Living Published: Living	Validated: 2026-07-22 Review on material revision, withdrawal, supersession, or scheduled report review.
AMD SEV-SNP Documentation https://www.amd.com/en/developer/sev.html	Primary vendor documentation Living Published: Living	Validated: 2026-07-22 Review on material revision, withdrawal, supersession, or scheduled report review.
Confidential Computing Consortium https://confidentialcomputing.io/	Primary industry consortium Living Published: Living	Validated: 2026-07-22 Review on material revision, withdrawal, supersession, or scheduled report review.

10. Lineage, Review, and Publication Record

Record	Value
Original source file	RES-023_TEEs_TDX_SEV_SNP_Remote_Attestation.md
Original source words	1638
Upgrade type	Enterprise research report; source and freshness validation; limitations; adoption decision; reproducibility plan
Generated on	2026-07-22
Current status	Candidate Research Report
Publication authority	Formal Mythos approval not yet recorded
Next review	90 days or event-driven trigger, whichever occurs first