OPAQUE

The Leading Confidential AI Platform

OPAQUE.co

2026

Where Your AI Stack Bleeds Data (46 Exposure Vectors)

Nothing is broken. Logs look clean. And your AI system is bleeding confidential data.

APPLICATION LAYER

8 **Access Control:** Source systems have RBAC/ABAC, vector DB doesn't.

7 **AI Derivative:** PII/MNPI deleted from database. Still in your embeddings, weights, snapshots, and caches.

CONTROL PLANE

6 **Agent Identity:** Neither agents nor the tools they use are verified before sensitive data is exchanged.

5 **Agent Output:** Agents route data autonomously. You can't audit every destination. One will be wrong.

4 **Policy Gap:** Config is NOT runtime enforcement, and agents will ignore.

COMPUTE PLANE

3 **Data-in-Use:** Inference runs in cleartext. Your ops team can see it.

2 **Operational:** Your APM logs every AI prompt (and context). By default.

1 **Proof Gap:** DORA, EU AI Act, and HIPAA now require runtime proof. Your SOC 2 describes intent, not what ran.

46 exposure vectors · 8 categories

Source: 2026 AI Data Leak Report | Developed with design partners: Anthropic, ServiceNow & Accenture | 200+ enterprise AI leaders | Validated with NVIDIA, Intel, AMD & Azure

OPAQUE

Control Every Agent. Prove Every Action.

One governance model — software to silicon, build to production.

THE STRUGGLE

- Setting & enforcing rules
- Uncontrolled data exposure
- Agent / model drift
- What ran and what it touched

Agents = Bionic Puppies

All power. No leash. Chews, occasional puddle.

CONTROL

OPAQUE
Agent Control

Control what agents *do* — not what they say.

- Deterministic, action-layer
- Signed agent identity
- Works with any framework

Built on AGT — the open standard you already use.

PROVE

OPAQUE
Confidential Core

BYO Container
Canvas to build your workload

Hardware Enforcement
Policies enforced in silicon

Verifiable Evidence
Cryptographic proof

TRUSTED BY
ANTHROP\C · encore CAPITAL GROUP · accenture · servicenow · EHS · TII Technology Innovation Institute

Zero trust agents with provable policy enforcement, verifiable by 3rd parties.

OPAQUE
Confidential AI Platform

BUILD FASTER WITH AGENT CONTROL.
SCALE SAFELY WITH CONFIDENTIAL CORE.

VERIFIABLE AI
REQUIRES BOTH.

OPAQUE

You set the rules. We prove they're enforced.

Hardware-enforced verifiable governance before, during, and after.

CONFIDENTIAL COMPUTING
(The Foundation)

Encrypted Runtime
Trusted Execution Environment (TEE) encrypts data in use

Hardware Attestation
Hardware-signed proof of what ran, where it ran, and what rules were enforced.

intel. AMD NVIDIA.

VERIFY THE AGENT BEFORE

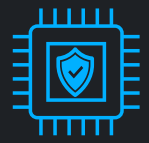

Workload Graph
Infrastructure
Silicon

Hardware attested and verified before the agent runs

Hardware does the security review. Not the security team.

ENFORCE POLICIES DURING

Policies bound to the agent at runtime, enforced by hardware

Your agents run on your most sensitive data. In production.

PROVE EVERYTHING AFTER

Runtime Evidence

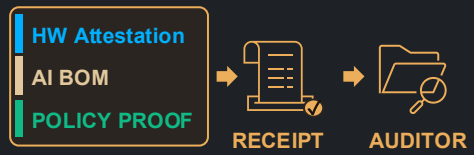

Verifiable by any third party

Deliver math to your auditor. Not promises.

OPAQUE Confidential AI

OPAQUE

Platform Architecture

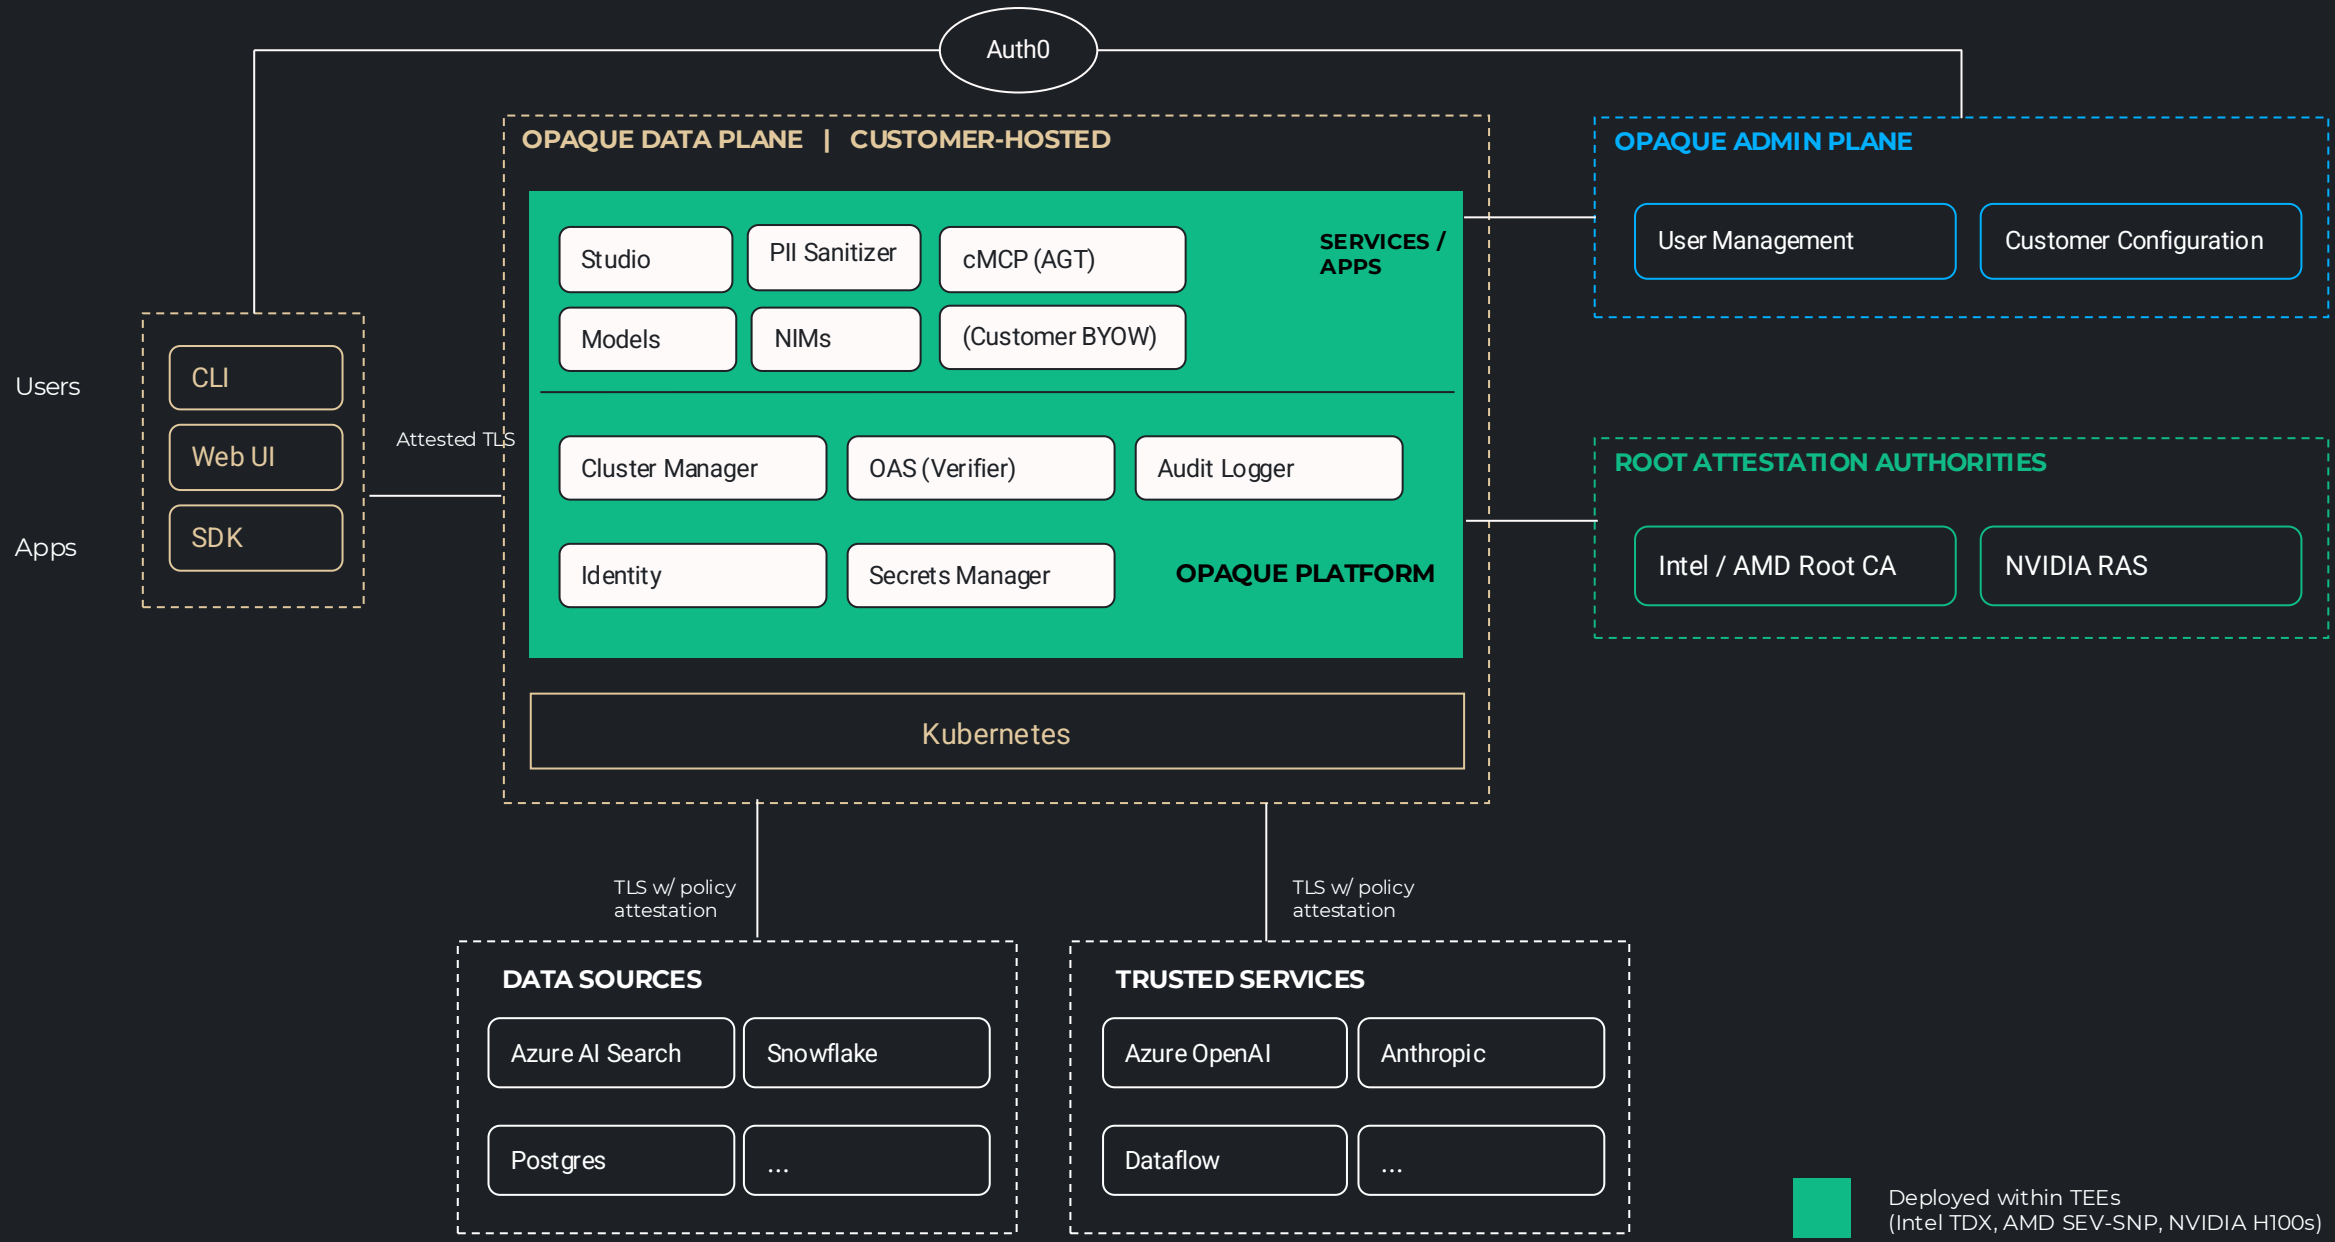

OPAQUE

DEMO

Don’t trust. Verify.

✓ Where it ran

Azure Attestation Document
Hardware-rooted evidence from the attestation service

Issuer	Microsoft Azure Attestation
Attestation type	Azure CVM (SEV-SNP)
Protocol version	2.0
Secure boot	Enabled

Platform configuration & OS
Confirms the exact operating system and environment running the workload

OS	Ubuntu Linux 22.04
VMID	1E22466A-0784-4F6E-BB3D-B9FB555F2AF3

Hardware
The hardware-rooted measurement proving this specific chip ran this workload

Hardware	Type	Verified via
AMD SEV-SNP	CPU / CVM	Microsoft Azure Attestation
NVIDIA H100*	GPU	NVIDIA Remote Attestation Service (NRAS)

Measured Boot & PCR Values
Cryptographically proves what actually booted, not just what was configured

PCRs Attested	0–7
PCR values	Hashes of firmware, bootloader, kernel,

Confidential VM (SEV-SNP) Guarantees
Hardware-enforced memory encryption and isolation for the VM

Memory encryption	Enabled (AMD SEV-SNP)
VM isolation	Hardware-enforced
Attestation root	AMD Secure Processor

✓ What ran

Workload Summary
The verified identity and attestation status of the workload

Workload	Exec-assist
Namespace	Acme-corp-prod
Overall tier	Affirming
Last attestation	2026-05-21T18:56:11Z
Workload identity	spiffe://acme/ns/corp-prod/workload/exec

Runtime Configuration
Confirms security-critical runtime features are enforced, not optional

Secure Boot	Enabled
TPM	Enabled
Console Access	Enabled

Cryptographic Keys
Post-quantum cryptographic keys, hardware-bound and attestation-tied

Field	Algorithm	Standard
Signing key	ML-DSA / SLH-DSA	FIPS 204 / 205
Encryption key	ML-KEM	FIPS 203
TPM Ephemeral key	Present	--

Workload Claims
Connects infrastructure trust to application-level identity and intent

Node	aks-pool1-32061274-vmss000000
Workflow ID	05d62a4e-...
Pod ID	Exec-assist-66b7d87bff-67jd5
Class ID	1E22466A-0784-4F6E-BB3D-B9FB555F2AF3

✓ Which rules held

Policies
Policies bound to this workload, as recorded in the manifest

Runtime Integrity Policy
Workload configuration locked at registration. Any deviation fails attestation.

Container image digest	Verified against RIM
Environment variables	Locked
Command-line arguments	Locked
Volume mounts	Locked
RIM hash	[cryptographic fingerprint of declared policy]

Network Policy
Enforces what the workload is allowed to communicate with

Firewall (netd)	NFTables rules — IP/port/protocol
Rule type	nftables NAT
Traffic redirect	All TCP → ATLS proxy (:15001)
aTLS mesh (fwd)	Peer attestation required on configured ports
Proxy mode	Attester (ATLS)
Effect	All traffic must transit the attested proxy — no plaintext fallback, no bypass

Pod Access Policy
Enforced at the KATA layer inside the TEE — not bypassable via Kubernetes controls

Shell access	Denied
Log access	Denied
Enforcement layer	KATA Agent (OPA, inside TEE)

* GPU attestation is illustrative

OPAQUE