



Zero Trust Hardware Access (ZTHA)

Superintelligence That Makes
Zero Trust Complete

EXECUTIVE PERSPECTIVE

Zero Trust has a hardware blind spot

Hardware identity is often assumed precisely where Zero Trust says trust must be earned.

Zero Trust has matured around users, applications, workloads, and network access. Yet the hardware layer is still frequently trusted using identifiers and signals that a device can report, inherit, or spoof. A MAC address, IP address, certificate, traffic profile, or asset record may be useful - but none of them alone proves what a device physically is.

Sepio Zero Trust Hardware Access (ZTHA) closes that gap. It extends 'never trust, always verify' to every connected physical asset by combining trafficless discovery, physical-layer evidence, Hardware AssetDNA™, contextual risk analysis, and policy-driven enforcement.

Hardware Superintelligence is Sepio's domain-specific intelligence layer: advanced AI and machine learning applied to massive, multidimensional hardware data to establish device truth, understand risk, and drive Zero Trust decisions at scale.

The shift: from visibility to understanding

SEE

Continuously discover managed, unmanaged, dormant, hidden, and non-traditional assets.

KNOW

Verify hardware identity using physical-layer characteristics and Hardware AssetDNA™.

DECIDE

Combine identity, posture, vulnerabilities, context, and anomalies into risk-aware trust decisions.

ACT

Translate verified hardware intelligence into policy, containment, isolation, and remediation.

The core premise

Zero Trust cannot be complete if the device requesting access has never been truly verified. Sepio provides the hardware intelligence layer that enables organizations to answer six questions continuously: What is connected? What is it really? Does it belong here? Is it operating as expected? What risk does it represent? What should happen next?

WHY NOW

The attack surface is becoming more intelligent - and more physical

AI accelerates both defense and adversarial capability, while hardware diversity expands the number of places trust can fail.

Modern environments mix enterprise IT, OT, IoT, IIoT, IoMT, peripherals, embedded systems, remote endpoints, and AI-enabled edge devices. Many cannot run conventional agents. Some remain silent for long periods. Others sit behind unmanaged infrastructure or expose identities that can be copied or manipulated.

Three forces are converging



AI-ACCELERATED THREATS

AI lowers the barrier to vulnerability discovery, reconnaissance, social engineering, and attack automation. Hardware compromise and supply-chain manipulation become more attractive where software controls are mature.



DEVICE COMPLEXITY

More device types, more vendors, longer lifecycles, and more unmanaged assets make static inventories and signature-driven identification increasingly fragile.



IDENTITY ASSUMPTIONS

Many controls still trust self-reported identifiers, enrollment records, traffic behavior, or network location. Sophisticated attackers target those assumptions.

The hardware blind spots that matter

- ✓ A legitimate MAC address presented by an illegitimate or substituted device.
- ✓ A dormant device that produces little or no useful traffic for conventional monitoring.
- ✓ An unmanaged switch, hub, tap, bridge, or transparent device that changes the physical environment without presenting a normal endpoint identity.
- ✓ A USB or peripheral device whose declared identity does not match its physical characteristics.
- ✓ An OT or medical device that cannot safely host an agent or tolerate active probing.
- ✓ A device that is genuine but appears in the wrong location, role, segment, or operational context.

The result: As software trust controls improve, attackers look for the assumptions beneath them. Hardware identity is one of those assumptions.


THE INTELLIGENCE LAYER

From Hardware AssetDNA™ to Hardware Superintelligence

Physical-layer evidence becomes more powerful when AI can interpret it across identity, posture, risk, and context.

Sepio observes hardware through signals that are difficult to reduce to a single spoofable identifier. Depending on the deployment model, these can include Layer-1 physical characteristics, network infrastructure telemetry, host-side device attributes, connectivity context, and other non-payload metadata.

Hardware AssetDNA™ converts those observations into a persistent, high-confidence representation of device identity. Hardware Superintelligence builds on that foundation by applying AI and machine learning to recognize patterns, identify outliers, correlate context, and convert raw device evidence into security decisions.



What makes this different from generic “AI-powered security”?

Data advantage. The intelligence is grounded in hardware-specific, real-world observations rather than relying only on logs, packets, or software telemetry.

Domain-specific models. AI/ML is applied to hardware identity, classification, anomaly detection, posture, and risk - not as a general-purpose chatbot layer.


Explainable evidence. Security teams can trace decisions back to identity, context, risk indicators, and policy rather than accept an opaque score.

THE FRAMEWORK

The Zero Trust Hardware Access framework

Five connected pillars move from hardware reality to enforceable trust.

ZTHA is designed as a continuous decision loop. Each pillar contributes evidence to the next; Hardware Superintelligence runs across the entire loop, learning from device identity, environmental context, posture, risk indicators, and outcomes.



**HARDWARE SUPERINTELLIGENCE - CONTINUOUS
LEARNING & CONTEXT ACROSS THE ENTIRE ZERO
TRUST HARDWARE ACCESS LOOP**

A continuous hardware trust loop

- ✓ **DISCOVER** establishes what is really connected.
- ✓ **VERIFY** establishes what each asset actually is.
- ✓ **VALIDATE** determines whether the asset belongs and whether its current state aligns with policy.
- ✓ **ASSESS** prioritizes risk using context and evidence.
- ✓ **ENFORCE** converts the decision into action through Sepio and the existing security ecosystem.

The key principle: Enforcement should be based on verified hardware identity + posture + contextual risk - not on identity claims alone.

PILLARS 1-2

Discover everything. Verify what it really is.

Zero Trust starts with complete hardware reality and defensible identity.



PILLAR 1: **DISCOVER**

Complete asset truth

- Continuously discover assets across IT, OT, IoT, IIoT, IoMT, network infrastructure, and peripherals.
- Find unmanaged, silent, dormant, hidden, MAC-less, and non-traditional assets that network-centric methods can miss.
- Use trafficless techniques to reduce operational overhead and avoid dependence on packet payload inspection, decryption, SPAN ports, or active probing.
- Build an inventory and connectivity view based on observed reality rather than procurement records or self-reported identity.

Outcome: a continuously refreshed, normalized view of the connected hardware estate.



PILLAR 2: **VERIFY**

Hardware identity beyond claims

- Use Hardware AssetDNA™ and physical-layer characteristics to establish a high-confidence representation of device identity.
- Correlate physical evidence with registered information, role, location, network context, and expected function.
- Identify mismatches between what a device claims to be and what its physical evidence indicates.
- Detect impersonation, substitution, spoofing, anomalous hardware, and device types that do not fit known-good patterns.

Outcome: a defensible hardware identity record that can support downstream trust and policy decisions.

Why Hardware Superintelligence matters here

At enterprise scale, identity verification is not simply a fingerprint lookup. Different models, revisions, interfaces, operating contexts, and deployment patterns create enormous variability. Hardware Superintelligence helps separate meaningful deviation from normal variation - improving classification, anomaly recognition, and prioritization while preserving the physical evidence behind the decision.

PILLARS 3-5

Validate continuously. Assess context. Enforce intelligently.

Identity becomes operationally useful when it drives posture, risk, and policy.



PILLAR 3: **VALIDATE**

Policy & posture

A genuine device can still be in the wrong place, connected through an unexpected path, or operating outside policy. Sepio continuously evaluates identity, configuration, vulnerabilities, location, connectivity, and role against expected posture.

Outcome: A continuously refreshed view of whether each device remains aligned with policy and expected behavior.



PILLAR 4: **ASSESS**

Contextual hardware risk

Sepio combines verified identity, posture, vulnerabilities, context, anomaly indicators, and asset criticality into prioritization - not another flat vulnerability list, but an understanding of what matters now and why.

Outcome: An explainable risk view that directs investigation, remediation, isolation, acceptance, or immediate action.



PILLAR 5: **ENFORCE**

Policy-driven action

A Zero Trust decision must change the environment. Sepio can trigger notifications, workflows, containment, blocking, quarantine, or isolation through policy and existing network/security controls.

Outcome: Verified hardware intelligence becomes measurable risk reduction, not passive visibility.

HOW SEPIO DELIVERS ZTHA

Hardware Superintelligence, grounded in physical truth

Sepio adds a hardware-native intelligence layer without forcing organizations to replace their existing Zero Trust stack.

1. Trafficless discovery

Sepio can discover network-connected assets using switch and router counters, telemetry, and Layer-1 characteristics rather than relying solely on mirrored traffic or deep packet inspection. This reduces infrastructure overhead, avoids payload decryption, and makes the approach suitable for sensitive and availability-focused environments, including OT.

2. Hardware AssetDNA™

Physical and contextual device characteristics are transformed into Hardware AssetDNA™, enabling high-confidence identification that goes beyond MAC/IP identity and other easily replicated claims.

3. Hardware Superintelligence

AI/ML models analyze device evidence and context to improve classification, identify anomalies and rare assets, recognize identity inconsistencies, and prioritize risk. The intelligence layer is designed specifically around connected hardware - the data Sepio has been collecting and normalizing at scale.

4. Risk intelligence

Sepio combine, risk Indicators, posture, vulnerabilities, device rarity, connectivity context, and identity confidence to turn hardware observations into prioritization. Examples include unexpected port speed, AssetDNA anomalies, unsupervised infrastructure, rare devices, threats, and vulnerability exposure.

5. Integrated enforcement

Sepio enriches the broader security ecosystem with trusted hardware identity and risk context. Integrations with NAC, firewalls, SIEM/SOAR, ITSM/CMDB, and other security platforms allow organizations to translate ZTHA decisions into automated remediation and enforcement using the controls already in place.



DEPLOYMENT

SaaS, on-premises/private cloud, and air-gapped architectures.



COVERAGE

Network and host perspectives across IT, OT, IoT, IIoT, IoMT, and peripherals.



OPERATIONS

Designed to augment existing controls rather than create another isolated visibility console.

AI & SUPERINTELLIGENCE

Why more hardware data changes the Zero Trust decision

The advantage is not AI in isolation; it is AI applied to a uniquely hardware-centric evidence base.

Cybersecurity AI is only as useful as the data it can learn from. Most security platforms are trained and optimized around software events, packet flows, logs, identities, or endpoint telemetry. Sepio adds another dimension: the physical and connectivity characteristics of the hardware itself.

A multidimensional hardware evidence base

- Physical-layer and electrical characteristics used for device differentiation.
- Network infrastructure telemetry and connectivity attributes without requiring packet payload inspection.
- Host-side peripheral and USB characteristics where the Host Security Agent is deployed.
- Known-good asset patterns, deployment context, organizational role, and historical observations.
- Risk Indicators, vulnerabilities, anomalies, and remediation outcomes that add security meaning to identity.

What AI/ML can do with that evidence



CLASSIFY

Improve device-type and identity classification across model, revision, vendor, and environment variability.



DETECT

Surface rare, anomalous, substituted, or inconsistent hardware that deserves investigation.



CONTEXTUALIZE

Distinguish benign variation from deviations that become meaningful in a specific location, role, or topology.



PRIORITIZE

Combine multiple weak signals into a stronger, explainable risk decision and route action accordingly.

Superintelligence is the outcome, not the buzzword. The objective is a continuously improving understanding of hardware reality that helps security teams make better Zero Trust decisions faster and at greater scale.

COMPLETING THE ARCHITECTURE

Sepio strengthens - rather than replaces - the Zero Trust stack

Hardware Superintelligence becomes trusted input to policy engines, enforcement points, and operational workflows.

NIST SP 800-207 emphasizes that Zero Trust decisions should use multiple sources of enterprise knowledge and that authentication and authorization apply to both subjects and devices. Sepio contributes the missing hardware-specific evidence: what the device physically is, whether that identity is credible, whether it belongs in context, and what risk it represents.

ZERO TRUST FUNCTION	WHAT EXISTING CONTROLS KNOW	WHAT SEPIO ADDS
Asset awareness	CMDB, EDR/MDM inventory, NAC records	Trafficless discovery of unmanaged, hidden, silent, and non-traditional hardware.
Device identity	MAC/IP, certificates, agent identity, enrollment	Physical-layer evidence and Hardware AssetDNA™ to challenge self-reported identity.
Policy decision	User/device posture, business rules, network context	Identity confidence, risk indicators, device context, anomaly intelligence.
Policy enforcement	NAC, firewall, segmentation, SOAR workflows	Verified hardware intelligence to trigger risk-based containment, blocking, isolation, or remediation.

A particularly important extension for NAC

When devices cannot support 802.1X, many NAC deployments fall back to MAC Authentication Bypass (MAB). That makes the MAC address a critical trust input - and a predictable target for impersonation. Sepio can add independent hardware identity and risk evidence to extend policy beyond simple MAC-based trust, enabling stronger control for unmanaged and non-802.1X assets.

OPERATIONAL OUTCOMES

What changes when hardware trust becomes evidence-based

What changes when hardware trust becomes evidence-based



Rogue or substituted device

A device presents a familiar network identity but its physical characteristics are inconsistent with the expected asset. ZTHA surfaces the mismatch and can trigger containment.



Dormant or silent asset

A device generates little useful traffic, yet network infrastructure evidence indicates its presence. ZTHA maintains visibility without requiring payload analysis.



Unmanaged infrastructure

An unexpected switch, hub, bridge, or other network component changes the environment. Sepio identifies the hardware and provides context for policy and remediation.



Peripheral and USB risk

Host-side hardware is evaluated beyond VID/PID or self-reported descriptors, improving visibility into suspicious, substituted, or policy-violating peripherals.



OT / IoT environments

Organizations gain hardware visibility and risk context without intrusive scanning or dependence on endpoint agents across every device.



Supply-chain and integrity questions

Unexpected physical identity or component-level deviations can become a signal for deeper investigation rather than remaining invisible to software-centric tools.

CONCLUSION

Superintelligence That Makes Zero Trust Complete

Zero Trust is strongest when hardware is verified as rigorously as users, applications, and workloads.

The fundamental Zero Trust question is not whether a device is inside the network, but whether it should be trusted at all. That requires evidence independent of the identity the device presents.

Sepio Zero Trust Hardware Access brings that evidence to the hardware layer. Trafficless discovery establishes what is connected. Hardware AssetDNA™ establishes what it physically is. Hardware Superintelligence interprets identity, posture, context, anomalies, and risk at scale. Policy and integrations then convert that intelligence into action.

The result: Zero Trust no longer stops at the user, credential, application, or network session. It extends to the physical device itself.

Six questions every connected asset should answer

What is connected?

Is it operating as expected?

What is it really?

What risk does it represent?

Does it belong here?

What should happen next?

Sepio provides the hardware-native intelligence to answer those questions continuously - and make Zero Trust decisions based on reality, not assumptions.

Reference: NIST SP 800-207, Zero Trust Architecture (2020).

READY TO COMPLETE ZERO TRUST?

Discover how Sepio Hardware Superintelligence brings verified hardware identity, contextual risk, and policy-driven enforcement to your environment.



SEPIO