

Zero Trust Architecture in Computer Networks and Hardware

Sivabalaji Sivasankarapandian*

ABSTRACT

This paper discusses Zero Trust Architecture in software networks and hardware architecture. It highlights the importance of implementing zero trust principles at both the hardware and Network layer architecture. It analyzes the various components of ZTA and how Zero-Trust is handled in Hardware level and Software Level.

Keywords – ZTA, PE, TPM, PA, TLS, HMAC

Date of Submission: 15-06-2025

Date of acceptance: 30-06-2025

I. INTRODUCTION

The cybersecurity industry uses broad concepts to address security challenges, including multifactor authentication, least privilege access, segregation of duties, and automated correlation of indicators of compromise. Zero Trust Hardware is essential in strengthening defenses.

Zero Trust, introduced in an analyst report, promotes network architectures that go beyond traditional firewalls. This model is crucial for users accessing public cloud workloads, where strong security controls are necessary as neither the user nor the cloud can be fully trusted.

Though Zero Trust Architectures often focus on software-based security for endpoints, networks, cloud infrastructure, and containerized applications, hardware implementation is less discussed. Zero Trust Hardware ensures the integrity of endpoints, servers, and connected devices, making it a vital part of modern cybersecurity strategies.

II. ZERO TRUST ARCHITECTURE

Zero Trust Architecture (ZTA) is a security model based on the principle of "never trust, always verify." This approach assumes that no user or device should be trusted by default, regardless of their location—whether inside or outside the network. Instead, each access request must be authenticated and authorized based on various factors, ensuring that the principle of least privilege is upheld.

III. ZERO TRUST ARCHITECTURE COMPONENTS




Fig 1. ZERO TRUST ARCHITECTURE COMPONENTS

The component descriptions:

a) Policy Engine (PE):

This component is responsible for making the final decision to grant access to a resource for a given subject. The PE utilizes enterprise policies as well as inputs from external sources (such as CDM systems and threat intelligence services) as inputs to a trust algorithm (refer to Section 3.3 for more details) to grant, deny, or revoke access to the resource. It operates in conjunction with the policy administrator component. The policy engine makes and logs the decision (whether approved or denied), while the policy administrator executes the decision.

b) Policy Administrator (PA):

This component is tasked with establishing and/or terminating the communication path between a subject and a resource (via commands to relevant PEPs). It generates any session-specific authentication tokens or credentials used by a client

to access an enterprise resource. The PA relies on the decision of the PE to ultimately permit or deny a session. If the session is authorized and the request authenticated, the PA configures the PEP to allow the session to commence. If the session is denied or previously granted access is revoked, the PA signals to the PEP to terminate the connection. Although some implementations may treat the PE and PA as a single service, here they are distinguished into two logical components. The PA communicates with the PEP during the creation of the communication path via the control plane.

c) Policy Enforcement Point (PEP):

This system is responsible for enabling, monitoring, and eventually terminating connections between a subject and an enterprise resource. The PEP interacts with the PA to forward requests and/or receive policy updates. While the PEP is a single logical component within ZTA, it can be divided into two separate components: the client side (e.g., an agent on a laptop) and the resource side (e.g., a gateway component controlling access in front of the resource), or it may function as a single portal component that serves as a gatekeeper for communication paths. Beyond the PEP lies the trust zone (see Section 2), which hosts the enterprise resource.

IV. ZERO TRUST ARCHITECTURE IN HARDWARE

Zero Trust

Zero Trust is a term that is commonly discussed, but it can be slightly misleading. It refers to zero implicit trust rather than zero trust entirely. Nothing should be trusted solely based on its network location or claims from the developer, which is crucial in today's diverse and hybrid cloud computing environments. All interactions must be verified and all access to data must be authenticated and authorized, resulting in explicit trust.

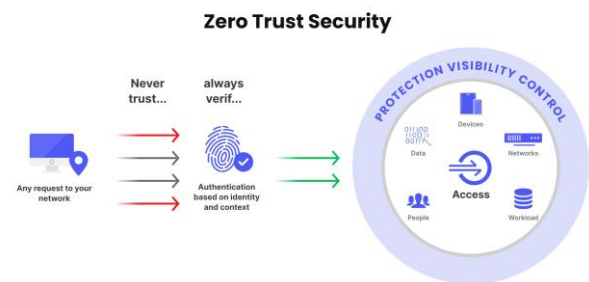


Fig 2. Zero Trust Security

Interactions typically involve services and users but can also include how a system is initially designed. For example, the mathematical equations behind encryption algorithms have been verified and proven over time by multiple third parties—not merely because a developer endorsed them. Each software stack does not need to perform its own mathematical proof each time the algorithms are used, and these components are trusted due to the established explicit trust. This trust is extended upward and outward into the rest of the stack and architecture, allowing those algorithms to:

- Create Transport Layer security (TLS) connections between services to encrypt data as it flows across the internal network.
- Encrypt data as it resides on disk.
- Use cryptographic digests and hash-based message authentication code (HMAC) to create challenge/response systems.

Trust is not always permanent. As technology advances and new vulnerabilities are identified, an algorithm or implementation might be deemed trustworthy one day and lose that trust the next. Several cryptographic algorithms that were once considered secure are now known to be flawed (such as DES, MD4, MD5, SHA). Understanding our explicit trust roots can help mitigate potential harm by knowing what needs to be replaced when trust is lost.

Roots of Trust

When trust is anchored in a solid foundation, it enables the establishment of more intricate relationships between services that do not require reliance on trust. This fundamental component used to build trust among other components is known as a root of trust.

Every system possesses roots of trust, although they are often overlooked. A securely designed system must explicitly define its roots of trust to prevent vulnerabilities from arising due to oversight. In the contemporary landscape of cloud computing, hybrid cloud environments, and edge computing, physical security of systems cannot always be guaranteed. Roots of trust must be fortified against physical and environmental tampering, as well as against systems that target your code. While some individuals may fully trust their cloud provider and its personnel implicitly, those with heightened security awareness should base their trust on explicit evidence.

A large modern software system should incorporate multiple roots of trust, such as encryption algorithms, secret management systems, and TLS certificate authorities. Security-sensitive applications benefit from hardware-based roots of trust, ideally equipped with remote attestation, as they offer superior tamper resistance and tamper evidence compared to software solutions. Although software's flexibility is one of its key advantages, this trait is detrimental to security, especially when serving as a root of trust.

Trusted Platform Modules

One approach used to extend trust up through a software stack and help protect it against physical and virtual threats is with a Trusted Platform Module (TPM)—a cryptographic sub-processor, which is usually hardware but can be virtualized, that is designed to provide certain cryptographic guarantees while being resistant to physical tampering. TPMs are fairly ubiquitous, as they're present in many phones, routers, servers, laptops and even cloud computing offerings. A given TPM can be tied back to its manufacturer via a certificate chain, proving it's an authentic device (as long as the manufacturer protects their private certificates), while also containing an encryption key that's unique to this particular TPM. This certificate chain and the cryptographic functions allow it to be used to enhance the security of a given system with things such as disk encryption, measured boot and file integrity measurements.

A system using a TPM as a root of trust can make cryptographic guarantees about its state that other systems can build on. For instance, because we can make hardware-backed assertions about the state of a given system—that it hasn't been tampered with

at boot or run time—we can tie its authentication and authorization to those guarantees before it tries to access sensitive information. Now it doesn't matter that we don't physically control the resource, because as long as the TPM is as secure as possible, we can have greater confidence that the layers we build on top of it also have a high degree of security.

A Trusted Platform Module (TPM) is a dedicated chip that functions as a hardware root of trust, while a Trusted Execution Environment (TEE) uses a different approach. CPUs with TEE capabilities provide higher integrity guarantees for data and code, and greater confidentiality guarantees for data, particularly within a specific area of system memory used for general purpose computation. This method, known as confidential computing, protects data in use from unauthorized access, including access from privileged levels such as the hypervisor or operating system, and safeguards code and data from tampering. This enhances the security posture for applications or workloads running in the confidential environment by removing implicit trust from lower stack levels.

A TEE must be hardware-based and capable of attestation for it to qualify as confidential computing. Therefore, a hardware root of trust is essential, and attestation is a critical piece of the security guarantees provided by any TEE. Attestation provides verifiable information that enables trust decisions regarding the TEE. The format and content of a TEE's attestation can vary depending on implementation, but ideally, it should establish a chain of trust from the hardware root—specifically the CPU and its hardware keys—to both the TEE operating on the CPU and the CPU manufacturer. Similar to the TPM, the TEE's CPU should be linked to its manufacturer through a certificate chain to verify authenticity. Additionally, signatures from the CPU's hardware keys should be traceable, typically through intermediary keys, to the TEE instance, indicating that the TEE is functioning correctly. Each link or signature in the chain of trust from the manufacturer, through the hardware root, to the TEE should be auditable and verifiable in a well-formed TEE attestation, minimizing the need for implicit trust except in the hardware root.

Applying Zero Trust in Hardware

Adding hardware security to an enterprise zero trust plan protects endpoints and other tangible computing systems from direct network access. The

main functional requirements for hardware security in zero trust can be grouped into three areas:




Fig 3. Zero Trust in Hardware

• Hardware-Level Visibility

The main goal is to enhance the accuracy and coverage of hardware metadata in a zero trust session. This ensures only approved devices can engage in a zero trust network by collecting Layer 1 data from all peripherals within the target environment.

• Hardware Identity Management

The second goal is to utilize hardware metadata for identity decisions. Physical and electrical-level data can reveal device identities, using unique identifiers like fingerprints.

• Hardware Access Control

The third goal is to decide if a hardware device can access a workload based on cyber risk. This includes identifying rogue or impersonating devices, often done using machine learning analysis on central servers.

V. ZERO TRUST SECURITY PLATFORMS

General Zero Trust Security Platforms & Comprehensive Solutions:

A) Check Point Harmony SASE:

This platform provides an integrated approach to Zero Trust, combining secure access, cloud services, and multi-layered security. It supports a secure access service edge (SASE) model and offers scalability for different environments.

B) Akamai Guardicore Platform:

This platform merges microsegmentation, Zero Trust Network Access (ZTNA), multi-factor authentication (MFA), and threat hunting into a unified system. It delivers detailed visibility and access control for users and applications.

C) Microsoft Azure Entra ID (formerly Azure AD):

A cloud-based identity and access management (IAM) service, it enables secure access to external and internal resources, supporting features such as MFA, SSO, and conditional access policies.

D) Okta Identity Cloud: A versatile cloud-based IAM solution, it provides robust authentication and access control, including adaptive MFA and centralized management.

E) StrongDM: Specializes in securing privileged access management (PAM) and offers continuous verification and detailed access control for users, devices, and applications.

F) Palo Alto Networks Prisma Access: This solution offers secure remote access and protection for distributed workforces, integrating ZTNA and secure web gateways (SWG) with scalability.

G) SEPIO PLATFORM

Sepio leverages unique physical layer visibility and hardware fingerprinting technology to enhance security in line with Zero Trust security Principles. Let us explore SEPIO PLATFORM in detail.

SEPIO PLATFORM

The Sepio HAC-1 platform offers comprehensive cybersecurity functions for hardware devices and Internet of Things (IoT) systems, focusing on enhanced asset visibility for enterprise teams and other organizations operating networks. Typical customers of Sepio include banks, insurance companies, critical infrastructure operators, government agencies, Internet service providers, and various other organizations of differing sizes and scopes. The Sepio platform facilitates Hardware Access Control (HAC) within a zero-trust environment, providing in-depth visibility into deployed hardware assets essential for mitigation, policy enforcement, third-party integrations, and other zero trust controls. The HAC-1 platform primarily emphasizes physical layer visibility, hardware access control support, and protection against rogue devices.

Sepio's commercial HAC-1 platform achieves accurate, real-time visibility of deployed hardware devices through a unique fingerprinting algorithm based on observable physical layer characteristics. These physical and electrical signals

are collected and analyzed using machine learning models to create a device fingerprint, thus generating a unique identifier for each specific hardware device. The collective view formed by these device fingerprints provides significant value by offering detailed visibility that surpasses traditional IT inventories, which often do not distinguish between devices, or software-based information, such as MAC addresses and device names, which may be unreliable.

VI. CONCLUSION

The complete implementation of zero trust will include various other considerations, predominantly software-related. However, incorporating hardware visibility, identity management, and access control into any network setup can be beneficial. In the next section, we review how Sepio Systems provides a commercial platform that addresses these objectives.

REFERENCES

- [1] M T. Sarkorn and K. Chimmanee, "Review on Zero Trust Architecture Apply In Enterprise Next Generation Firewall," 2024 8th International Conference on Information Technology (InCIT), Chonburi, Thailand, 2024, pp. 255-260
- [2] S. R. Pokhrel, G. Li, R. Doss and S. Nepal, "Toward Decentralized Operationalization of Zero Trust Architecture for Next Generation Networks," in *IEEE Journal on Selected Areas in Communications*, vol. 43, no. 6, pp. 1998-2010, June 2025
- [3] Sina Ahmadi. Zero Trust Architecture in Cloud Networks: Application, Challenges and Future Opportunities. *Journal of Engineering Research and Reports*, 2024, 26 (2), pp.215-228.
- [4] B. G. Jung, Y. -S. Yoo, K. Kim, B. -S. Kim, H. Lee and H. Park, "ZTA-based Federated Policy Control Paradigm for Enterprise Wireless Network Infrastructure," 2022 27th Asia Pacific Conference on Communications (APCC), Jeju Island, Korea, Republic of, 2022, pp. 1-5
- [5] M. Nasiruzzaman, M. Ali, I. Salam and M. H. Miraz, "The Evolution of Zero Trust Architecture (ZTA) from Concept to Implementation," 2025 29th International Conference on Information Technology (IT), Zabljak, Montenegro, 2025, pp. 1-8