Akamai

TD SYNNEX
Public Sector

DLT™

Least Permissive Trust: What Zero Trust Wishes It Could Be

A three-volume ebook series from Akamai

Volume 1: Rethinking Zero Trust Architectures

Federal agencies and departments — including the Department of Defense (DoD) — have long been at the forefront of cybersecurity innovation, responding to rapidly evolving threats and an increasingly complex technology landscape. Among the most significant advancements in recent years is the adoption of the **Cybersecurity and Infrastructure Security Agency's (CISA's) Zero Trust architecture (ZTA)** as the core framework for securing federal systems. The central idea of Zero Trust is straightforward but profound: **trust no one, verify everything**.

While federal organizations are making great strides in implementing the five CISA pillars of Zero Trust — identity, devices, networks, applications and workloads, and data — this approach is arguably insufficient to meet today's complex threat landscape. What is needed is a shift in the vernacular and mindset. Federal organizations should stop thinking of these technologies as pillars and instead think of them as ecosystems.

These ecosystems must of course allow access, but the access should be the least permissive possible, and it should be constantly validated. Every request should be challenged, and policy and attributes should be mapped to ensure the least privilege is granted to authorized entities, resulting in a strategy of **Least Permissive Trust**.

This ebook series is designed to help educate and inform federal organizations on the practical application of Least Permissive Trust. The comprehensive strategy discussion is presented in three volumes:

- **Volume 1:** Rethinking Zero Trust Architectures
- **Volume 2:** Identity, Application Access, and Microsegmentation
- **Volume 3:** Cross-Pillar Capabilities and Implementation Guidelines

Breaking down the CISA Zero Trust model

Traditional cybersecurity models relied on perimeter-based defenses, assuming that users and devices could be trusted once they were inside the network. This approach has proven insufficient, as attackers have become more sophisticated, exploiting internal weaknesses and gaining unauthorized access once past the perimeter. Recognizing this, the U.S. government shifted to **Zero Trust**, which assumes that every user, device, and application — even those inside the network — must be continuously verified.

In support of this model, CISA developed its Zero Trust Maturity Model, which is structured around five key pillars: **identity, devices, networks, applications and workloads,** and **data** (Figure 1).

Identity

Devices

Networks

Applications and workloads

Data

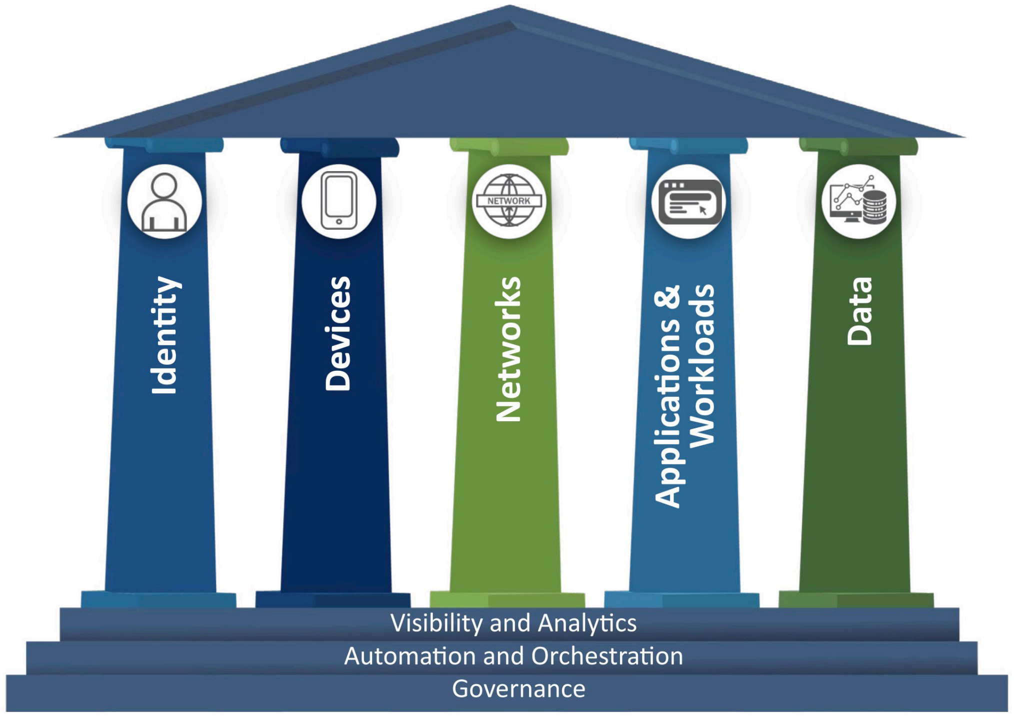

Fig. 1: CISA's Zero Trust Maturity Model Pillars (Source: CISA)

These five pillars serve as the foundation for implementing Zero Trust principles across federal agencies. However, while the pillars provide a strong foundation, they can also inadvertently create **technology silos**, where security measures are isolated rather than integrated across the organization. This ebook examines these pillars in detail, highlighting their strengths and challenges, and introduces the need for an evolved approach — **Least Permissive Trust** — to overcome the fragmentation that can arise from a pillar-based model.

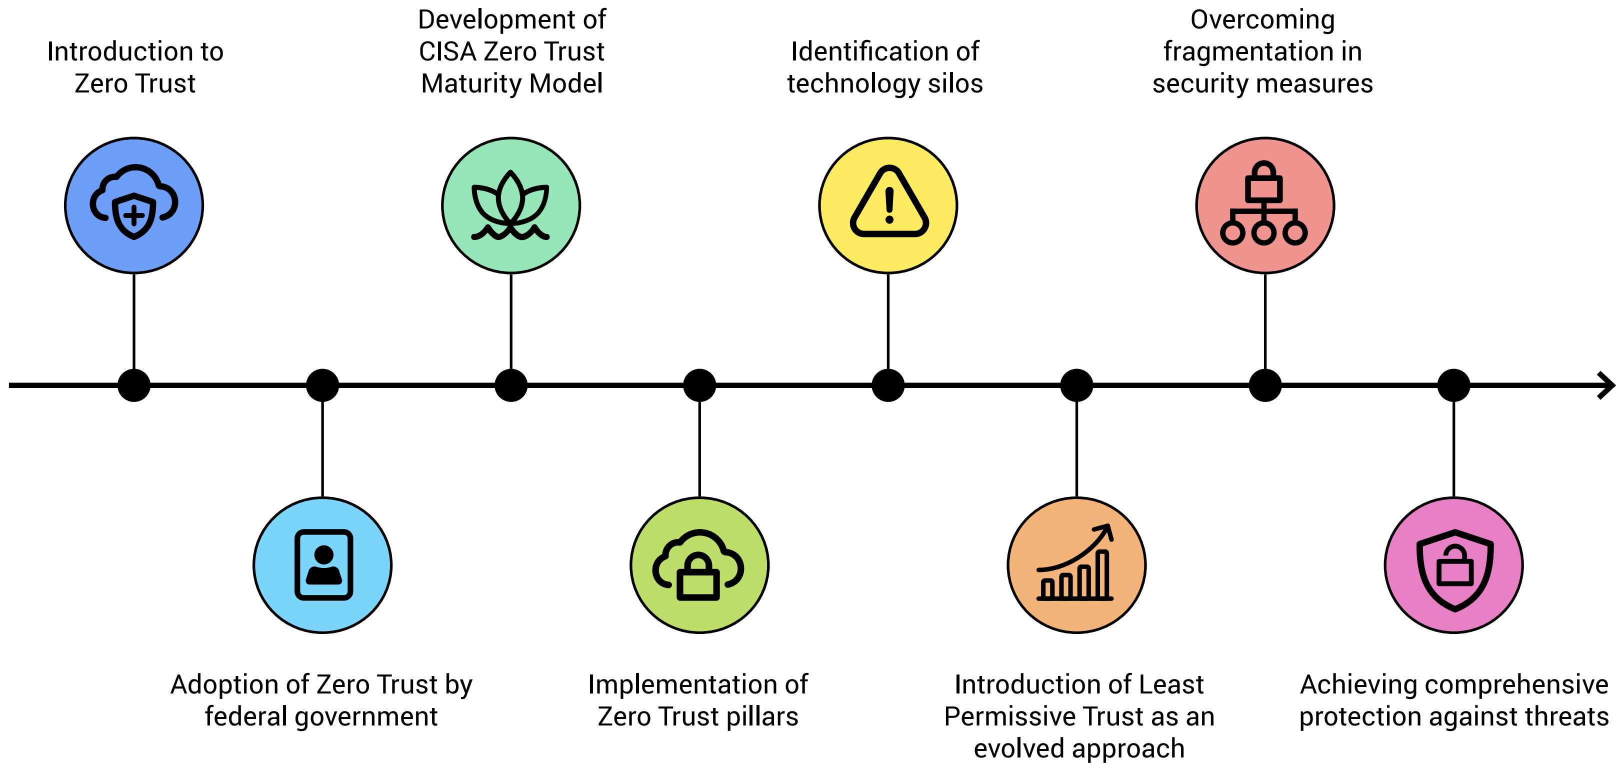
Introduction to Zero Trust
Adoption of Zero Trust by federal government
Development of CISA Zero Trust Maturity Model
Implementation of Zero Trust pillars
Identification of technology silos
Introduction of Least Permissive Trust as an evolved approach
Overcoming fragmentation in security measures
Achieving comprehensive protection against threats

Fig. 2: Evolving federal cybersecurity — from Zero Trust to Least Permissive Trust

CISA's five Zero Trust pillars

The CISA Zero Trust pillars are designed to cover every aspect of federal cybersecurity, ensuring comprehensive protection against both internal and external threats. Each pillar addresses a specific domain within the Zero Trust architecture, providing guidelines for how federal agencies should secure their systems.

Identity pillar

The **identity pillar** focuses on verifying that all users and entities accessing a network are who they claim to be. This involves **identity, credential, and access management (ICAM)** systems, which ensure that users are continuously authenticated, authorized, and monitored. Identity systems are essential in enforcing the principle of **least privilege**, where users are only granted the minimum access required to perform their tasks.

However, in practice, the identity pillar often becomes siloed, focusing solely on verifying credentials at login without considering the broader context, such as device health, network security, or behavioral anomalies during a session. This can lead to overprovisioned access rights and potential vulnerabilities if identities are trusted throughout a session without continuous re-evaluation. The **static nature** of traditional identity systems can conflict with the dynamic, real-time needs of a Zero Trust framework.

Devices pillar

The **devices pillar** ensures that every device accessing the network, whether managed or unmanaged, is authenticated, monitored, and secured. This includes agency-owned devices as well as **bring-your-own-device (BYOD)** policies. Device health, configuration, and compliance checks are critical to preventing compromised or unauthorized devices from accessing sensitive data.

While device management is crucial, it can also create silos when treated as a standalone pillar. Federal agencies may secure devices by using endpoint detection and response (EDR) tools, but without integrating device security with **identity management** or **application access controls**, a compromised device could still be used to access sensitive systems. For example, even if a device is deemed secure at login, its security posture can change during a session, requiring dynamissions and access. Without integration, the device pillar may fail to prevent unauthorized actions if device health initial authentication.

Networks pillar

The **networks pillar** focuses on securing the flow of information between systems, applications, and users. Zero Trust assumes that **all network traffic is untrusted**, regardless of whether it originates inside or outside the network. This means that encryption, authentication, and monitoring must be applied to every communication within the network. The principle of **least-privilege access** is enforced by ensuring that users, devices, and applications can only communicate with specific resources based on verified identities and real-time risk assessments.

Despite its importance, the networks pillar can also contribute to siloed security operations if not properly integrated with other pillars. For example, network segmentation might effectively isolate traffic between different departments or systems, but without integration with identity management and application controls, attackers could still move laterally within the network by exploiting overprovisioned identities or compromised devices. **Microsegmentation**, a key feature of solutions like Akamai Guardicore Segmentation, addresses this issue by enforcing granular, real-time controls over internal network traffic, but without broader integration, network security alone is insufficient to stop advanced threats.

Applications and workloads pillar

The **applications and workloads pillar** addresses the need to secure access to applications, whether they are hosted on-premises, in the cloud, or in hybrid environments. The goal is to ensure that only authorized users can access specific applications and that application interactions are continuously monitored. This includes protecting application workloads from unauthorized access, preventing lateral movement, and enforcing policies such as **Zero Trust Network Access (ZTNA)**.

Applications are often the primary target of cyberattacks, as they house sensitive data and provide access to mission-critical systems. However, when application security operates in isolation, vulnerabilities arise. Without real-time integration with **identity verification, device posture,** and **network controls**, attackers can gain access to applications by using stolen credentials or compromised devices. Solutions like Akamai **Enterprise Application Access** address this by integrating identity verification and access control with application-specific policies, but these must be part of a broader, cross-pillar security strategy to be truly effective.

Data pillar

The **data pillar** focuses on protecting sensitive information, ensuring that data is encrypted both in transit and at rest, and that access is restricted to authorized users. Data protection is a central concern for federal agencies, especially given the increasing threat of **data breaches** and **insider threats**. Zero Trust principles demand that access to data be continually reassessed and that data be monitored for unauthorized access or anomalies.

Data security often operates independently of other pillars, leading to gaps in protection. For example, if a user is granted access to data based solely on their identity but their device or network connection is compromised, sensitive information could still be exfiltrated or tampered with. **Cross-pillar integration** is essential to ensure that data access is contingent not just on identity but also on the security of the device, network, and applications involved in handling the data.

The fragmentation problem: Technology silos in Zero Trust

While each of these pillars plays a critical role in the Zero Trust model, treating them in isolation can lead to **technology silos**. These silos occur when security controls within one pillar do not communicate or interact with controls in other pillars, resulting in fragmented security policies and inconsistent enforcement. For example:

- A user might be authenticated at login (identity pillar), but their device could become compromised during the session (devices pillar), and without cross-pillar integration, the system may not revoke or adjust the user's access permissions.
- Network segmentation might limit access between systems (networks pillar), but if identity and application security policies are not aligned, attackers could exploit identity weaknesses to bypass segmentation controls.

This **lack of cohesion** creates opportunities for attackers to exploit gaps between pillars by moving laterally between systems, escalating privileges, or exfiltrating data under the guise of legitimate access. Additionally, technology silos increase operational complexity for security teams, who must manage multiple disparate systems and tools, each with its own policies and interfaces.

The case for Least Permissive Trust: Moving beyond pillars

The solution to these challenges lies in moving beyond a strict pillar-based approach and adopting a Least Permissive Trust model. This evolution of Zero Trust focuses on:

Cross-pillar integration

This ensures that identity, device, network, application, and data security controls work together in real time, dynamically adjusting permissions and access based on the most up-to-date information from all pillars.

Granular access controls

Implementing fine-grained policies helps ensure that users, devices, and applications only receive the minimal permissions necessary at any given moment. This reduces the risk of overprovisioning and minimizes the potential attack surface.

Continuous verification

Instead of verifying users or devices once (at login, for example), Least Permissive Trust ensures that access is continuously reassessed based on changing conditions, such as device health, network security, and user behavior.

Akamai's **ICAM solutions**, **Enterprise Application Access**, and Akamai **Guardicore Segmentation** provide the tools needed to achieve this model, enabling **dynamic, real-time security controls** that span identity, network, and application security. These technologies break down the barriers between pillars, ensuring that every access request is evaluated holistically and that security policies are enforced consistently across all layers of the architecture.

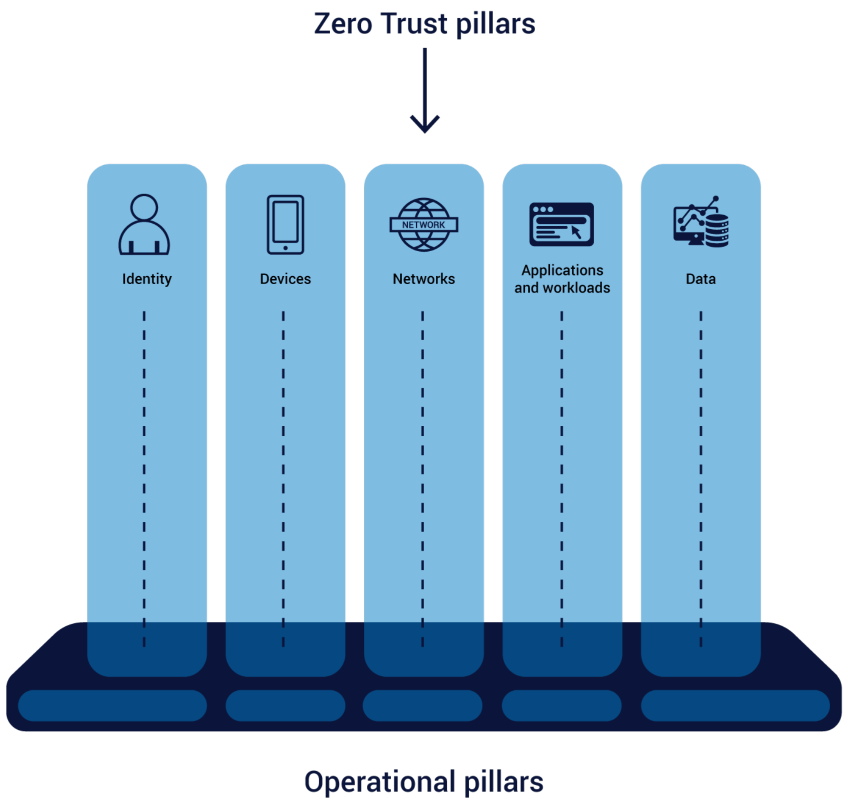

Fig. 3: Technology silos increase vulnerabilities and operational complexity

The path to unified security

The CISA Zero Trust pillars provide a solid foundation for securing federal agencies, but the potential for technology silos poses a significant risk. To overcome this, agencies must evolve toward **Least Permissive Trust**, a unified approach that integrates security controls across all pillars, ensuring dynamic, context-aware access control. By leveraging Akamai's advanced solutions, federal agencies can break down these silos, reduce their attack surface, and ensure that permissions are continuously updated based on real-time risk assessments.

Shifting from Zero Trust to Least Permissive Trust

The concept of Zero Trust has become a dominant framework in cybersecurity, emphasizing a “never trust, always verify” model. However, as this approach matures, a critical evolution is taking place: the shift toward **Least Permissive Trust**. Least Permissive Trust enhances Zero Trust by addressing its limitations, focusing on dynamic, integrated controls that minimize trust across all layers of an organization’s architecture. By restricting permissions to the absolute minimum required for any task at any given time, Least Permissive Trust ensures tighter security, reduced risk, and increased flexibility.

While Zero Trust pillars such as **identity, devices, networks, applications,** and **data** provide strong foundational security, they often operate in silos. These silos can prevent the fluid, real-time enforcement of security policies across an enterprise. Least Permissive Trust seeks to break down these silos and create a unified security posture, dynamically adjusting access and privileges based on real-time data. The result is a more flexible, adaptable, and secure environment — one that can effectively combat modern cyberthreats while allowing federal agencies and the DoD to maintain high operational efficiency.

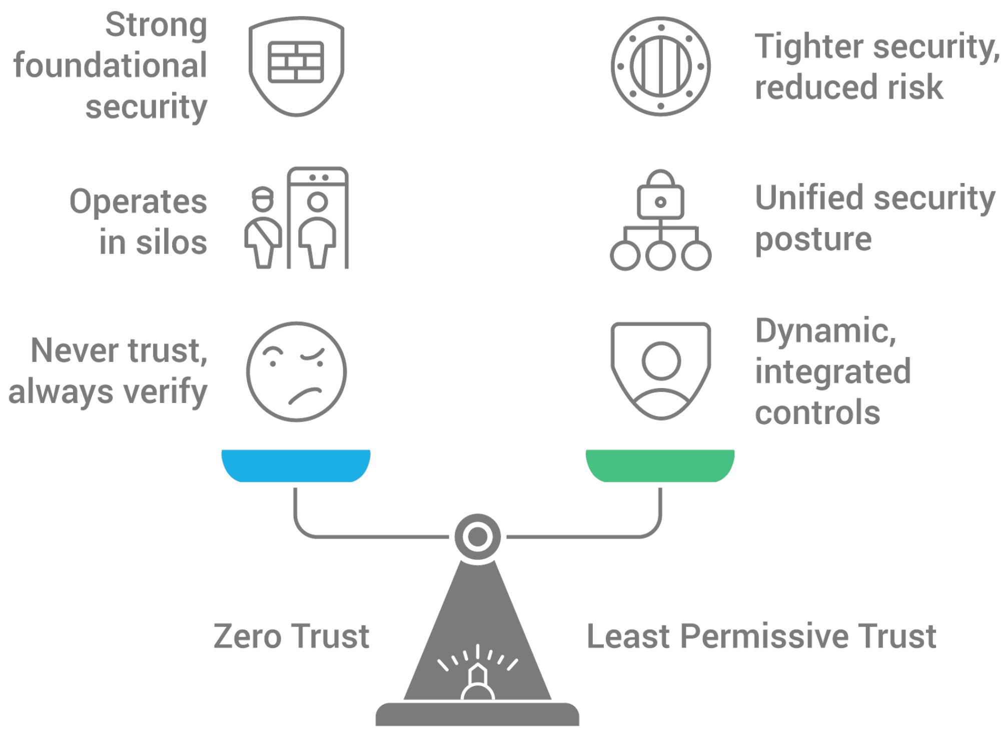
Strong foundational security
Operates in silos
Never trust, always verify
Tighter security, reduced risk
Unified security posture
Dynamic, integrated controls
Zero Trust
Least Permissive Trust

Fig. 4: Zero Trust vs. Least Permissive Trust

What is Least Permissive Trust?

At its core, Least Permissive Trust is an evolution of Zero Trust that emphasizes **minimal access, dynamic enforcement, and continuous adjustment of permissions** based on real-time contextual information. Traditional Zero Trust models focus on eliminating implicit trust and enforcing strict access controls, but they often do so in a static or isolated manner. Least Permissive Trust builds on these principles but introduces more granularity and flexibility in how access is granted, based on factors like user **behavior, device health, application sensitivity,** and **network risk**.

This approach ensures that users, devices, and applications receive the **least amount of trust possible** to complete their tasks, and only for the minimum necessary time. Unlike traditional role-based access models that provide broad permissions based on a user's role, Least Permissive Trust dynamically adjusts access levels, even during a session, to minimize exposure.

For federal agencies, where classified and sensitive data is routinely accessed, and for the DoD, where national security is at stake, this level of granularity is essential. Least Permissive Trust guarantees that even if one layer of defense is compromised, an attacker's ability to move laterally or escalate privileges is severely restricted.

Why shift to Least Permissive Trust?

While Zero Trust pillars provide strong security, they can lead to fragmentation and inefficiency when not integrated. For example:

- **Identity systems** might authenticate a user, but without real-time integration with device health and network analytics, that same user could access sensitive applications from a compromised device.
- **Network segmentation** can isolate traffic, but without dynamic policy adjustments, a malicious actor within a trusted network segment might still have access to critical systems.
- **Data protections** such as encryption might secure information in transit or at rest, but without integrating with identity and application management systems, unauthorized users could still access and manipulate that data.

Least Permissive Trust addresses these gaps by implementing **continuous validation** across all systems and layers. It focuses on the real-time adjustment of access privileges, ensuring that permissions are revoked, limited, or expanded dynamically based on **contextual risk assessments**.

In federal environments, this model is particularly powerful. As agencies face increasing cybersecurity threats — from nation-state actors to insider threats — Least Permissive Trust helps reduce the attack surface, increase visibility, and adapt to emerging risks in real time.

Minimizing the attack surface: Dynamic adjustments

One of the key advantages of Least Permissive Trust is its ability to **minimize the attack surface** by enforcing the most restrictive policies possible across all layers of the organization. This concept is supported by Akamai's suite of technologies, including Akamai's ICAM solutions, Enterprise Application Access, and Akamai Guardicore Segmentation. These technologies enable organizations to **continuously assess and adjust** the level of trust assigned to users, devices, and applications, based on ongoing behavioral analysis, device health checks, and network conditions. Key features of Least Permissive Trust include:

Dynamic policy enforcement

Instead of relying on static access controls, policies are enforced in real time, adjusting as new information is gathered. For example, if a user logs in from a trusted device but later connects to a suspicious network, their access can be downgraded or restricted until further validation.

Granular permissions

Access is granted at the most granular level possible, ensuring that users can only interact with the specific resources required for their role, and only for the duration of the task. This prevents overprovisioning and limits the impact of potential breaches.

Real-time threat detection

Integrated monitoring and analytics tools continuously assess the environment for signs of compromise or anomalous behavior, triggering immediate policy adjustments.

For example, if a federal employee accessing a classified system displays unusual behavior — such as trying to download large amounts of data at odd hours — the system can immediately adjust their access level, notify security teams, and initiate further investigation. This real-time responsiveness is key to preventing breaches before they cause damage.

Akamai's role in enabling Least Permissive Trust

Akamai's technologies play a critical role in enabling Least Permissive Trust across federal agencies and DoD environments. Each solution is designed to dynamically assess risk and adjust trust levels in real time, minimizing the potential for unauthorized access.

1 **ICAM solutions**

Akamai's ICAM solutions provide a robust foundation for dynamic identity verification. They enable continuous monitoring of user credentials, using phishing-resistant multi-factor authentication (MFA) and context-based access controls. The platform assesses risk not only based on who a user is but also based on what device they're using, where they're accessing the system from, and their behavior over time. For instance, even if a user has successfully authenticated, suspicious activity on their device could trigger a revalidation or limitation of access, thereby enforcing Least Permissive Trust principles.

2 **Enterprise Application Access**

Enterprise Application Access delivers ZTNA by integrating application-specific policies that grant the minimum necessary access. This ensures that users only interact with the specific applications they need, and access is continuously verified throughout the session. Enterprise Application Access also integrates with Akamai's ICAM solutions to dynamically revoke or adjust permissions based on real-time risk analysis, providing an extra layer of protection in sensitive federal environments.

3 **Akamai Guardicore Segmentation**

Akamai Guardicore Segmentation enables granular, software-defined segmentation across network environments. By controlling traffic between workloads, it ensures that even if an attacker breaches one segment of the network, they cannot move laterally to access other sensitive systems. Microsegmentation allows Least Permissive Trust policies to be enforced at the network level, ensuring that permissions are tightly controlled, even within trusted network zones.

Continuous monitoring and threat intelligence

A crucial element of Least Permissive Trust is its reliance on **continuous monitoring and real-time threat intelligence**. Unlike static models that rely on periodic reviews of access rights, Least Permissive Trust requires continuous assessment of threats, network conditions, and user behavior to dynamically adjust access permissions.

Akamai's integrated solutions provide visibility across all pillars — identity, devices, networks, applications, and data — allowing for the real-time assessment of risk. For federal agencies, this means that threats can be identified and mitigated before they escalate into full-scale breaches.

Behavioral analytics

Akamai's solutions continuously analyze user behavior to detect deviations from normal activity. When a potential threat is identified, access permissions are automatically adjusted, and security teams are alerted to the suspicious activity.

Threat intelligence

Akamai leverages its extensive threat intelligence network to stay ahead of emerging threats. This intelligence feeds into the real-time decision-making process, enabling federal agencies to adapt their security posture as new threats evolve.

Device posture assessment

Akamai continuously evaluates device health and compliance status before and during access sessions. This real-time assessment examines factors like patch levels, endpoint protection status, disk encryption, certificate validity, and device integrity to create dynamic trust scores. When device posture deteriorates or suspicious changes occur, the system automatically restricts access privileges while alerting security teams, preventing compromised devices from becoming attack vectors even with valid credentials.

How to implement Least Permissive Trust in federal agencies

Transitioning to a Least Permissive Trust model requires careful planning, but the benefits are clear: improved security, reduced attack surface, and enhanced flexibility in managing access.

To implement this model, federal agencies can follow these key steps:

Assess current security posture

Begin by evaluating existing Zero Trust implementations and identifying areas where permissions might be too broad or not dynamically enforced. Look for opportunities to integrate cross-pillar capabilities, such as linking identity management systems with network monitoring tools.

Adopt dynamic access controls

Leverage Akamai's ICAM solutions and Enterprise Application Access to implement dynamic, context-based access controls that can continuously adjust permissions in real time. This will prevent overprovisioning and ensure that access is always minimized.

Integrate microsegmentation

Use Akamai Guardicore Segmentation to apply granular controls within the network. This will prevent lateral movement and ensure that even if one segment is compromised, the rest of the network remains secure.

Automate policy enforcement

Implement automation and orchestration tools to enforce security policies in real time. Automation is key to ensuring that Least Permissive Trust principles are consistently applied without overburdening security teams.

Monitor continuously

Deploy continuous monitoring solutions to assess threats, user behavior, and device health in real time. Integrate this data across all security pillars to ensure that permissions are dynamically adjusted based on the latest threat intelligence.

What’s next?

In summary, federal agencies and departments are encouraged to evolve their Zero Trust efforts toward a unified Least Permissive Trust strategy that integrates security controls across all pillars, ensuring dynamic, context-aware access control. By leveraging Akamai solutions, federal organizations can break down silos, reduce their attack surface, and ensure that permissions are continuously updated based on real-time risk assessments.

To learn more about Least Permissive Trust, read Volume 2 of this ebook series: **Identity, Application Access, and Microsegmentation.**

Contact Akamai today to learn more about our comprehensive security solutions.

Akamai

TD SYNNEX
Public Sector

DLT