WHITE PAPER

Beyond Perimeter Defense: Implementing Zero Trust in Federal Agencies

Akamai

TD SYNNEX
Public Sector

DLT

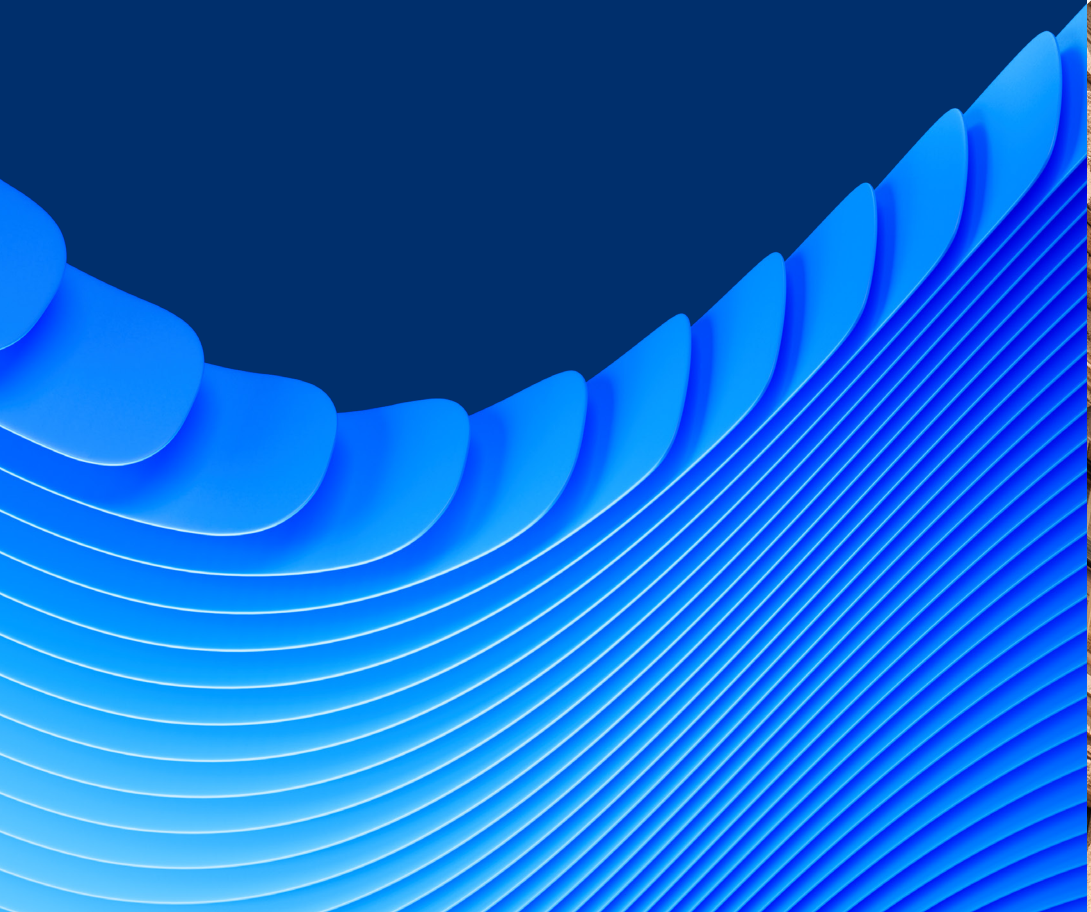

Contents

Introduction

The promises of the cloud include: “Be agile!” “Dynamically spin up servers to match demand!” “Rapid maintenance!” “Eliminate CapEx infrastructure costs!” But with no solid ground to plant your fence post: Where is your perimeter?

Information security has always had a physical component. Not too long ago, federal government information systems were in a physical location on-premises. As you connected your systems together, there was a physical boundary between the internal network and the outside world. No problem; you put in a firewall.

So, where is the perimeter when most of your assets are in the cloud, your applications are no longer actually in the building, and much of your infrastructure and software and most everything else is “as a service”?

Despite the evolving cloud-scape of information technology, the basic principles of security have not changed. The objectives of confidentiality, integrity, and availability are still valid, and the principle of least privilege access remains a core tenant to cybersecurity.

What changes is how you apply those principles, and where you apply our baseline controls and enforcement when physical boundaries are no longer relevant. You can no longer rely on physical, on-premises controls as part of your data security plans; instead, you need to secure the assets and systems, wherever they may lie.

Zero Trust to the rescue

To address cybersecurity in this new perimeter-less world is the Zero Trust security model, which dictates that strict access controls and authentication are always being enforced, regardless of where users or data are located.

The most concise statement on Zero Trust security is the National Institute of Standards and Technology Special Publication for Zero Trust architecture, NIST SP 800-207, that states the overall goal of Zero Trust is to “prevent unauthorized access to data and services coupled with making the access control enforcement as granular as possible.”

Zero Trust security became a government-wide initiative in 2021 after the White House issued Executive Order (EO) 14028, Improving the Nation’s Cybersecurity, required agencies to adopt Zero Trust cybersecurity architectures. In response, the Office of Management and Budget (OMB) released a Zero Trust strategy document with timelines and goals to have been completed at the end of fiscal year 2024.

The Zero Trust Maturity Model

The OMB strategy document refers U.S. government agencies to what has become the definitive guidebook for Zero Trust maturity within the DOT-Gov world, the Zero Trust Maturity Model (ZTMM) created by the Cybersecurity and Infrastructure Security Agency (CISA).

The ZTMM divides cybersecurity into five key areas, called pillars: **Identity, Devices, Networks, Applications and Workloads**, and **Data**. It then defines what Zero Trust maturity should look like within each of these areas.

Although the ZTMM is very helpful in defining Zero Trust principles and providing a roadmap to achieving Zero Trust in each of the security pillars, in practice it encourages technical silos between different teams working toward the same goal.

Technical silos increase risk, response time, and cost

Technical silos are created as individual project teams begin designing Zero Trust architecture for a given pillar. Buying managers might then begin solution procurement and even implementations. Suddenly it's discovered that the security controls implemented for one pillar do not communicate or interact with controls in other pillars.

For example, a user might be authenticated at login (Identity), but during their active session, the user's laptop or mobile device might become compromised (Devices). Because there has not been a unifying set of controls between Identity and Devices, the system may not revoke or adjust the user's access permissions.

Technical silos present a clear and present danger by increasing:

- Risk
- Incident response time and time to mitigation
- Cost

Risk

Silos increase risk by creating opportunities for attackers to exploit vulnerabilities among pillars, such as moving laterally between systems, escalating privileges, or exfiltrating data under the guise of legitimate access.

Incident response time and time to mitigation

Silos increase incident response time and time to mitigation by adding complexity for security teams, who must manage multiple disparate systems and tools, each with their own policies and interfaces, to mitigate active threats.

Cost

Silos increase costs by encouraging procurement of solutions and services for each pillar rather than looking at solutions and services for the entire organization. Because point solutions are often procured separately by multiple stakeholders, the added costs of integrations with other point solutions can be overlooked, which can result in increased costs as implementation plans are adjusted after procurement.

Federal CIOs need to take an agency-wide, holistic approach to their Zero Trust modernization initiatives.

Eliminate the silos with Least Permissive Trust

The Least Permissive Trust approach is the evolution of the Zero Trust model that addresses cross-pillar integration and granular access controls. The objective here is to attain Zero Trust maturity without technical silos.

This approach ensures that the **Identity, Devices, Networks, Applications and Workloads**, and **Data pillars** work together in real time with fine-grained policies that provide entities (i.e., users, devices, systems, etc.) with the minimal permissions necessary at any given moment.

Least Permissive Trust enhances Zero Trust by ensuring dynamic enforcement and continuous adjustment of permissions according to factors like user behavior, device health, application sensitivity, and network risk. More simply, Least Permissive Trust is looking at data from multiple pillars in real time.

Least Permissive Trust implements 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.

Focus on the objective of Zero Trust

When you begin your Zero Trust journey, it’s important to note that the CISA’s ZTMM is a guideline, but it is by no means the only way to proceed. CISA states, “This ZTMM is one of many paths that an organization can take in designing and implementing their transition plan to zero trust architectures ….”

In other words, focus on the objective of Zero Trust rather than the steps that any one organization has outlined.

Determine who, what, where, and how

It seems simple, but the first move in the Zero Trust journey is to determine:

- **Who** — Who (that is, which employees, contractors, agencies, public, etc.) needs access?
- **What** — What data, assets, and information does your agency have?
- **Where** — Where does the data reside?)
- **How** — How should the data be accessed and via what methods?

Who

The identity of users and entities is key to any security architecture, as the identity is what enables access to the systems, applications, and data. Evolving Zero Trust to the Least Permissive Trust model makes this even more important.

Rather than authenticating a user at login, Least Permissive Trust sees identity, credential, and access management (ICAM) as a dynamic, contextually aware process that takes into account changing risk factors, such as behavioral anomalies or device health.

What

What are the data sources, assets, applications that are being accessed? You need to understand what is being accessed in order to set policies for access decisions. If you do not know what data needs to be protected, you will have a hard time securing it from cyberthreats.

It's important to know where data resides, and from where it is accessed. While the front-end applications for interacting with data may be in one cloud, the actual data may be in another cloud.

Where

This is about the environment (e.g., What cloud service provider (CSP) is hosting sensitive data?) or internal environmental designations (such as production versus development) more than the physical location.

How

How should systems be accessed (including which protocol should be used and which direction the data access is flowing)

Pulling it all together

As you look for Zero Trust solutions, it is important that you look for solutions that will help you bind the pillars together without creating silos. Starting with an ICAM solution will provide dynamic identity verification with contextual authentication; identity verification is not just based on who the user is, but also where they are, what device they are using, and how they are behaving.

You need a ZTNA solution that seamlessly integrates with your ICAM, and a microsegmentation solution that can be applied to all your environments. Avoid solutions that are environmentally specific or require policy enforcement choke points.

You want solutions that can be integrated so that policy configurations, updates, and enforcement can be managed and applied to your organization as a whole.

Ideally, a unified platform solution is going to be the most cost-effective and the most secure.

This is because a platform solution will be less costly to deploy than multiple point solutions, be less costly to maintain than multiple point solutions, and will provide thecross-pillar visibility and enforcement that is required for true Least Permissive Trust.

Akamai Least Permissive Trust

Akamai plays a critical role in enabling Least Permissive Trust across federal agencies and Department of Defense (DoD) environments. Each solution is designed to dynamically assess risk and adjust trust levels in real time, minimizing the potential for unauthorized access.

Identity, credential, and access management (ICAM)

Akamai's ICAM solution provides a robust foundation for dynamic identity verification using phishing-resistant multi-factor authentication (MFA) and continuous monitoring of user credentials. The platform assesses risk not only according to who a user is, but also on what device they're using, from where they're accessing the system, and their behavior over time to provide context-based access controls.

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.

Akamai Enterprise Application Access

Enterprise Application Access delivers Zero Trust Network Access (ZTNA) but goes a step further 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 can be integrated with ICAM to provide real-time risk analysis — and dynamic revocation or adjustment of permissions, if necessary. This cross-pillar integration is key to getting to Least Permissive Trust necessary for the protection of sensitive federal environments.

Akamai Guardicore Microsegmentation

This solution enables granular, software-defined segmentation across network environments. By controlling network traffic between workloads and endpoints, Akamai Guardicore Segmentation 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.

As a software-based microsegmentation platform, Akamai Guardicore Segmentation is able to be deployed across all environments. Akamai Guardicore Segmentation provides visibility and observability for the entire agency, not just one environment or one network. It allows agencies to view mappings and system dependencies among all environments — systems, clouds, and on-premises — including legacy systems that are still in use.

Artificial intelligence (AI) helps agencies quickly identify dataflows and functional dependencies. AI makes creating new policies easy, enabling IT staff to use a common language to create new policies rather than set policies using coded languages.

Akamai's integrated solutions provide visibility across all pillars

Continuous monitoring and real-time threat intelligence is critical to breaking down the technical silos to go beyond Zero Trust to Least Permissive Trust. Rather than using periodic reviews of access privileges, Least Permissive Trust requires continuous assessment of threats, network conditions, and user behavior to dynamically adjust access permissions.

When these core solution sets are integrated, they provide visibility across all pillars — Identity, Devices, Networks, Applications and Workloads, and Data — allowing for the real-time assessment of risk. Akamai's use of machine learning, real-time behavioral analytics, and extensive threat intelligence empowers federal agencies to identify and mitigate threats before they can escalate into full-scale breaches.

Six steps to Least Permissive Trust for federal agencies

Transitioning to a Zero Trust model or a Least Permissive Trust model requires careful planning, but the benefits are clear: improved security, reduced attack surface, and enhanced flexibility in managing access. Moreover, having a unified platform solution is going to provide better security at a lower cost than trying to integrate multiple point solutions.

To implement the Least Permissive Trust model, federal agencies can follow these six key steps:

1. **Assess the current security posture.** Evaluate existing Zero Trust implementations and identify areas where permissions might be too broad or not dynamically enforced. Look for opportunities to integrate cross-pillar capabilities, such as linking ICAM systems with network monitoring tools.

2. **Adopt dynamic access controls.** Use Akamai's ICAM and Enterprise Application Access solutions to implement dynamic, context-based access controls that can continuously adjust permissions in real time.

3. **Integrate microsegmentation.** Use Akamai Guardicore Segmentation to apply granular controls within the network.

4. **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.

5. **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.

6. **Break the silos.** The Zero Trust journey should not be siloed. Zero Trust must be applied to the whole organization and then you'll be able to bind your controls together with Least Permissive Trust.

Book a demo