Akamai

WHITE PAPER

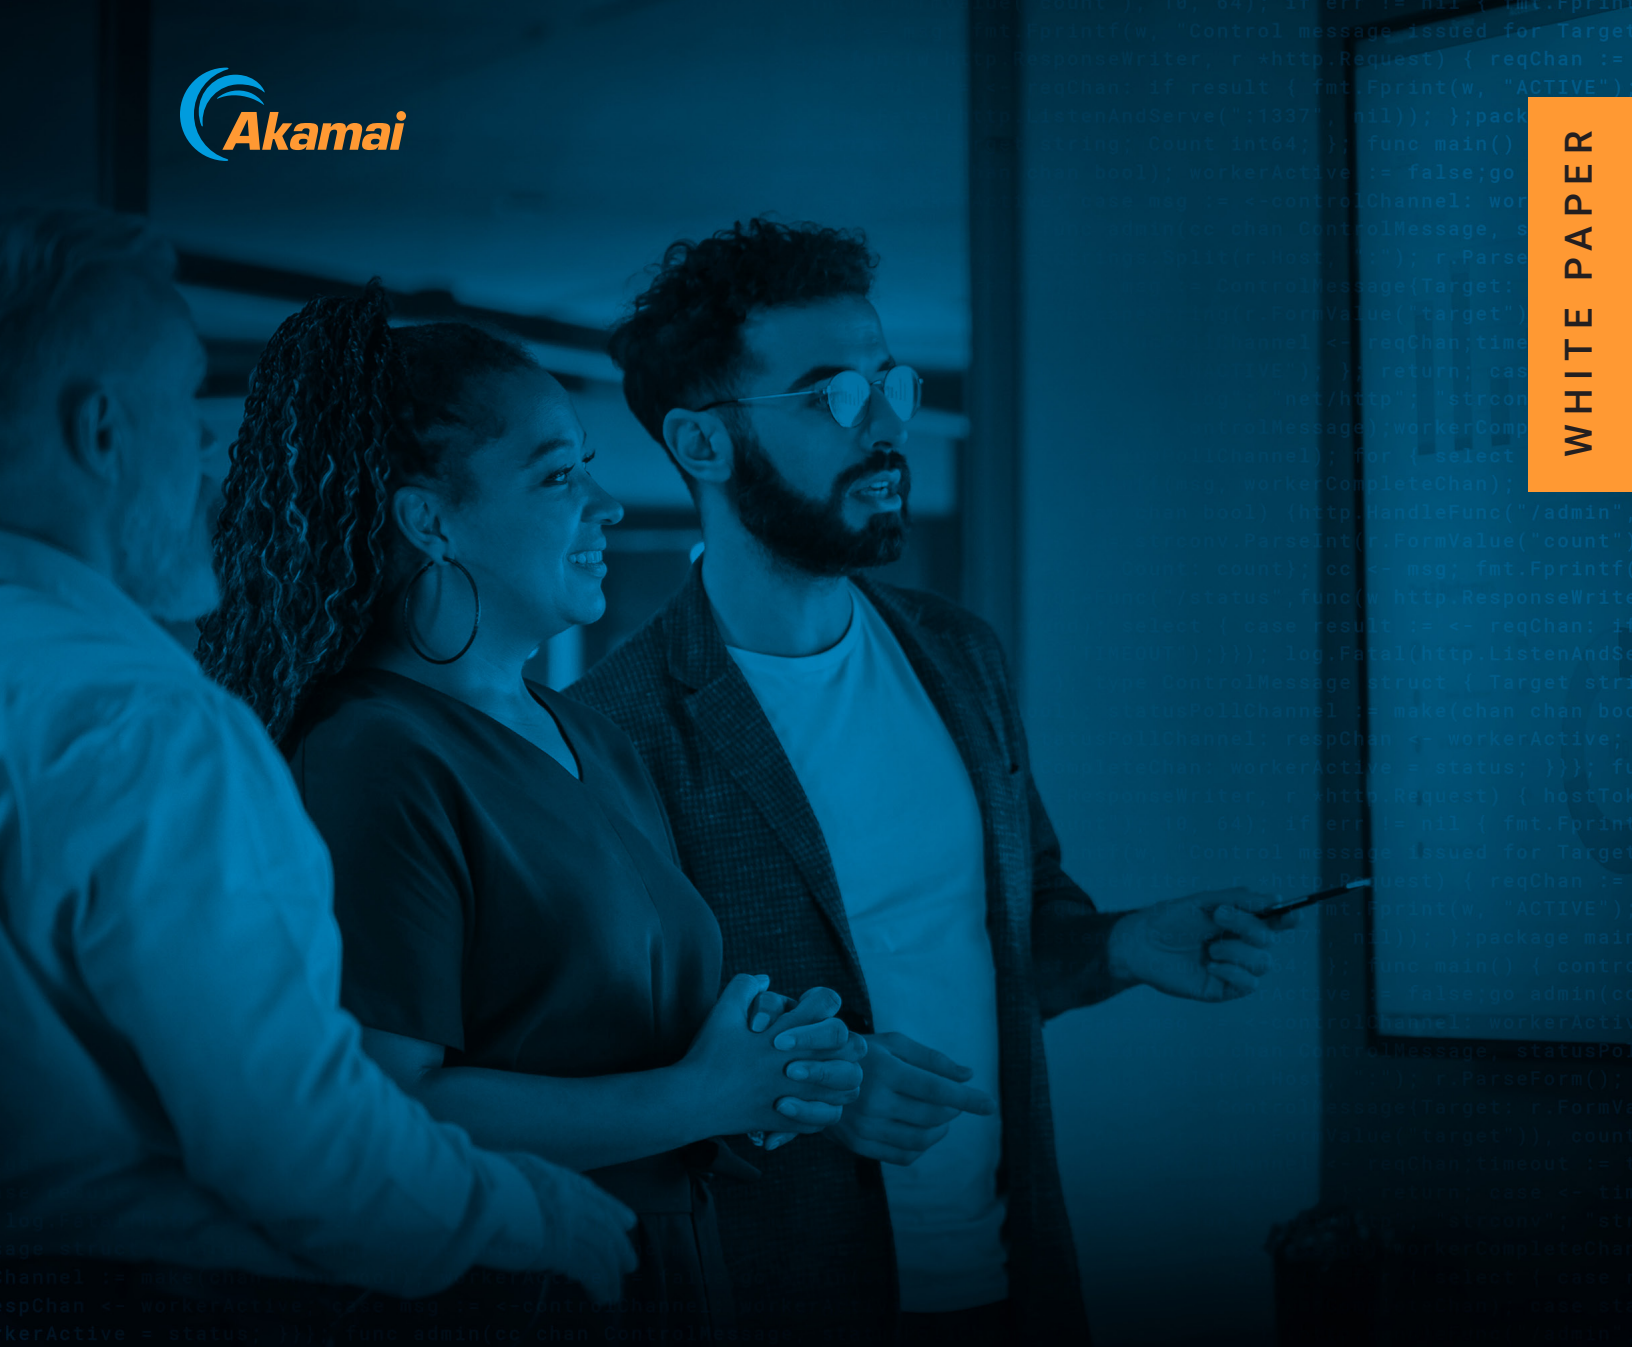

A Blueprint for Zero Trust Network Access

Who should read this guide?

Network architects, security engineers, CTOs, CISOs, and other IT and security decision-makers will all benefit from reading this guide.

For those responsible for scoping, configuring, deploying, implementing, and managing a Zero Trust Network Access (ZTNA) project, this guide provides a comprehensive review of potential benefits and the differences among the different systems. The guide includes:

- The limitations and security flaws in legacy approaches to application access and why ZTNA is required
- The components of ZTNA and how it works
- How Akamai Enterprise Application Access and Akamai MFA can deliver ZTNA quickly and easily

As the business world changes and cyberthreats mount, companies are taking a fresh look at their cyber defenses. Many have come to realize that the traditional networking architecture — which relied on a centralized location where all parties could access applications — leaves them vulnerable. This castle-and-moat approach to security — protecting the perimeter while assuming everyone within it is safe — leaves companies at risk of cyberattacks in today's landscape of mobile connections and the cloud. Instead, forward-thinking companies are turning to the concept of a Zero Trust architecture to protect vital assets. A core principle of any Zero Trust project is protecting the network. This white paper details how traditional hub-and-spoke approaches to network security are no longer sufficient and how shifting to ZTNA can better defend critical assets and serve as the key linchpin to a comprehensive Zero Trust architecture.

The pace of change for businesses has never been faster

How businesses operate and use technology is evolving and at an ever-faster rate. The evolution of computing has driven a rapid transition from hosting business applications in on-premises data centers to using multiple public clouds, private clouds, or a hybrid approach (both on-premises and public/private cloud).

Business model evolution has also spurred increased collaboration among entities and the need to provide partners and suppliers with access to applications and resources.

Finally, as businesses continue to embrace remote or hybrid work, users are now accessing business applications and resources from anywhere, on both managed and unmanaged devices.

With these changes, legacy approaches to managing application access are no longer sufficient, and companies must now adopt a new approach that allows secure access regardless of where the applications are hosted or where the users are located.

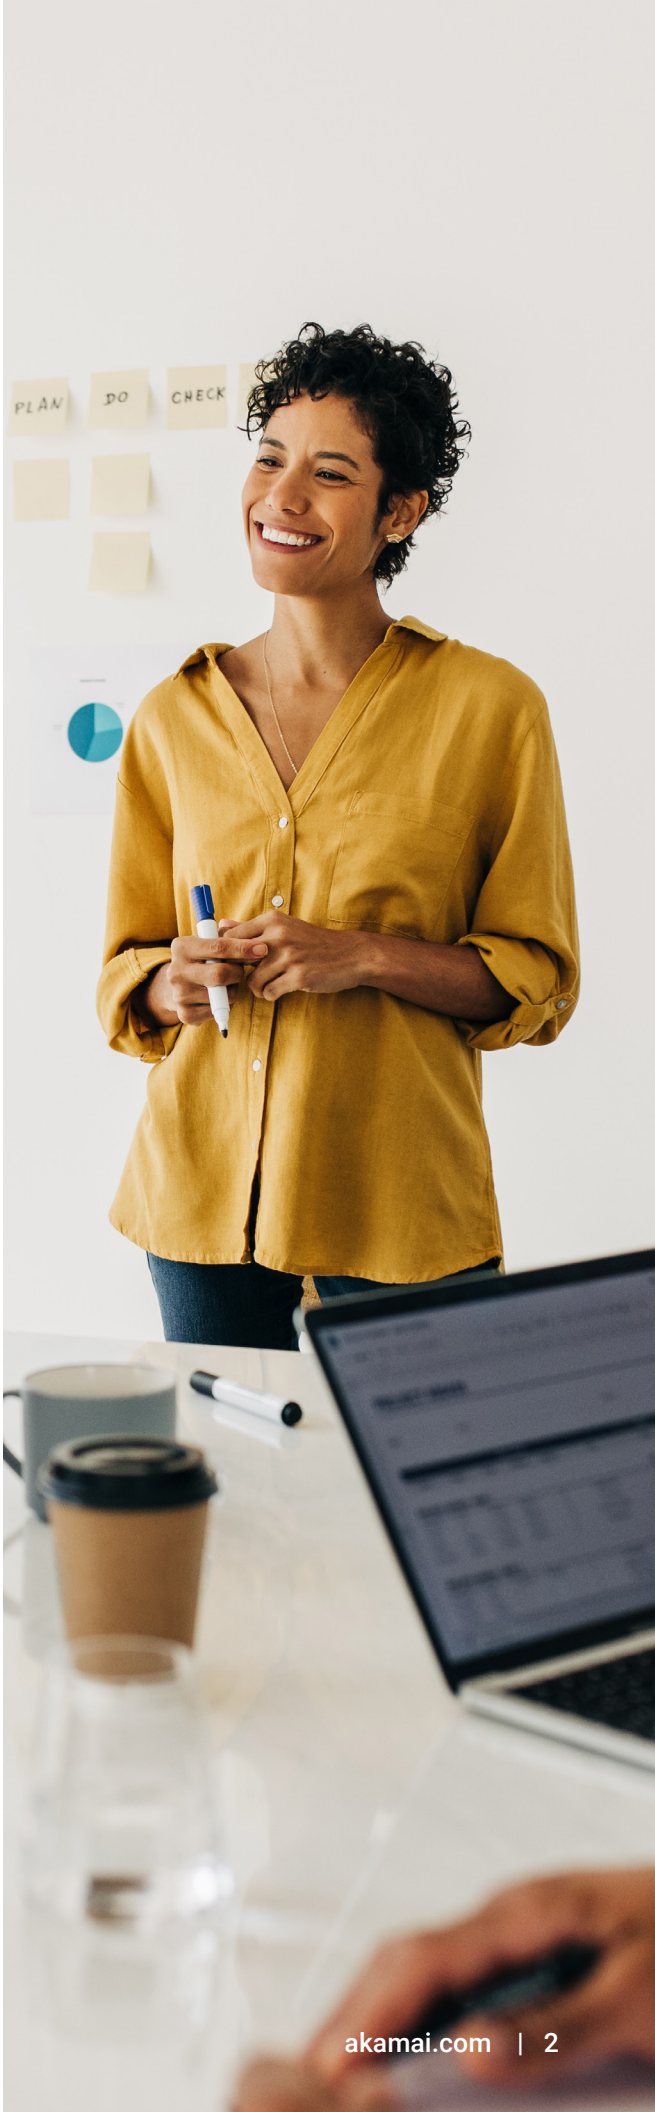

Legacy application access

For more than 20 years, companies have relied on firewalls to build a strong security perimeter and they've trusted the users who are inside that perimeter. This is akin to treating networks like castles with moats: thick walls and heavily secured gates form the perimeter to protect the castle (or, in this case, the network) and only users with the right credentials are allowed access. Once inside, users can then access specific applications based on their identity, which is delivered through identity provider (IdP) solutions such as Microsoft Active Directory.

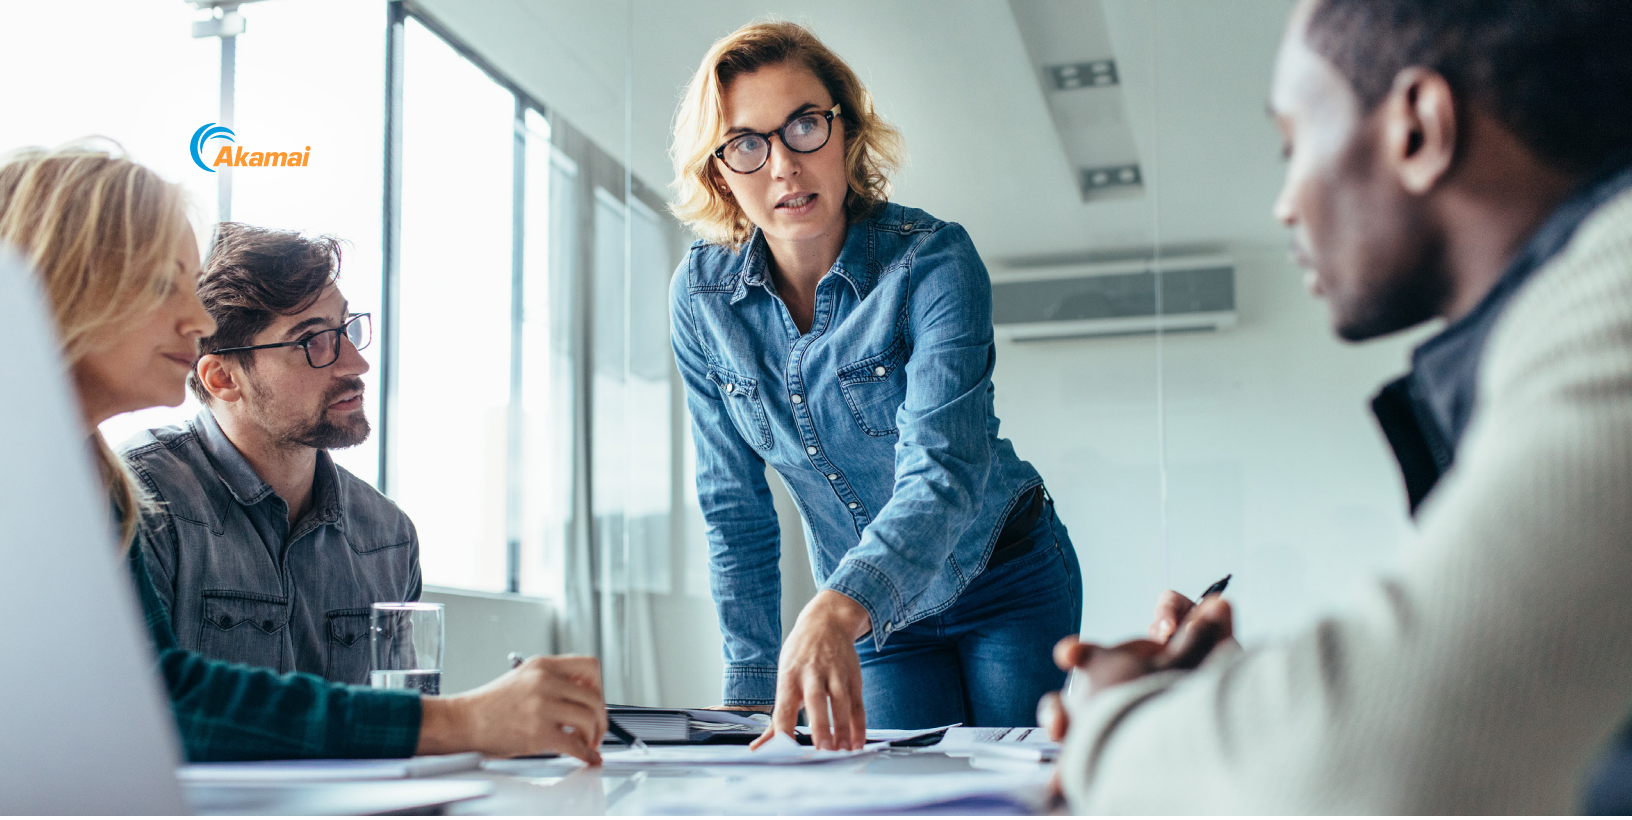

However, with flat networks, users actually have IP access to the entire network, which means they can discover other servers and applications. For example, if the IdP is configured correctly, a user might be able to find the server on which the payroll application is hosted, but when they try to log in to the application they will be denied access.

To fix this unfettered lateral movement problem, companies partitioned applications via virtual local access networks (VLANs) into separate segments behind a firewall and enforced now-archaic IP range-based rules for individual users or groups. This process is brittle and very prone to errors. Consider a scenario in which someone is doing maintenance and moves machines to a new rack or needs to re-IP them to a new range. Suddenly, users are locked out and the support calls come rolling in. Or perhaps a software upgrade requires changes to an application's architecture and users are redirected to another machine as part of the workflow. That machine may then be inaccessible to certain users or groups because the firewall rules were not updated.

This architecture is exceedingly complex and requires a very high degree of communication among application owners, network administrators, and security groups during any changes to ensure zero downtime.

We know what often happens when that coordination fails. Administrators want to follow best practices, but in times of desperation, they add the dreaded IP ANY/ANY ALLOW rule as a quick fix to allow affected users to access everything until the underlying problem can be diagnosed and repaired. Often, there isn't time to go back and revert these changes, however, and these quick fixes diminish a company's security posture over time.

VPNs add complexity, performance, and security challenges

For remote users, a virtual private network (VPN) typically provides access to on-premises applications that are hosted inside the perimeter, which then provide direct tunneled access to the company's network.

To manage user access to applications, companies often add dedicated application delivery controllers or use the access controls that are built into their VPN solutions. The goal is to align application access permissions regardless of where the user is located. If a user is denied access to the CRM application when they are inside the perimeter, then they should be denied access when they are connected through the VPN. Although that is the goal, the complexities of synchronizing application permissions between the two use cases and quick fixes can lead to users acquiring unintended application access.

Application access for contractors, partners, and suppliers

Companies also often use VPNs to allow remote application access for contractors, partner companies, or suppliers. For example, a company may allow outside access to its finance systems to allow suppliers to submit invoices. Allowing third-party application access through a VPN introduces additional security risks because the company no longer holds end-to-end security. If a third-party device with VPN access becomes compromised, attackers can gain access to the company's network.

VPNs and performance

The same trade-off happens with performance. In a VPN's simplest form, all traffic is directed back toward data center infrastructure. This can result in extremely slow access to internet properties and software-as-a-service (SaaS) applications due to hairpinning, which effectively doubles traffic.

To overcome this performance burden, administrators often deploy split tunnels, again marking which IP ranges should travel down the VPN and which should egress directly to the internet. This can be simple and effective when you only have one internal perimeter. However, it begins to get much more complex as you add multiple data centers and virtual private cloud providers. Administrators must then determine whether they are going to install VPN aggregators in every data center and how they are going to manage multipoint split tunnels effectively.

It's not that VPNs don't provide value. Far from it, in fact. Site-to-site access for multiple data center infrastructures is one case in which they shine. However, network-level access is not the correct paradigm for users who access applications, because network-level access enforces an unnatural compromise between simplicity and security/performance.

Network-based application access is good news for attackers

So far, we have focused on the risks and challenges associated with granting network-level access to all employees. However, this approach also exposes businesses to another risk: Cybercriminals who exploit stolen user credentials or a security vulnerability also have the potential to gain unfettered network-wide access. For example, if an attacker obtains VPN access using compromised employee credentials, they can then move laterally across the network, to find, access, and attack high-value targets.

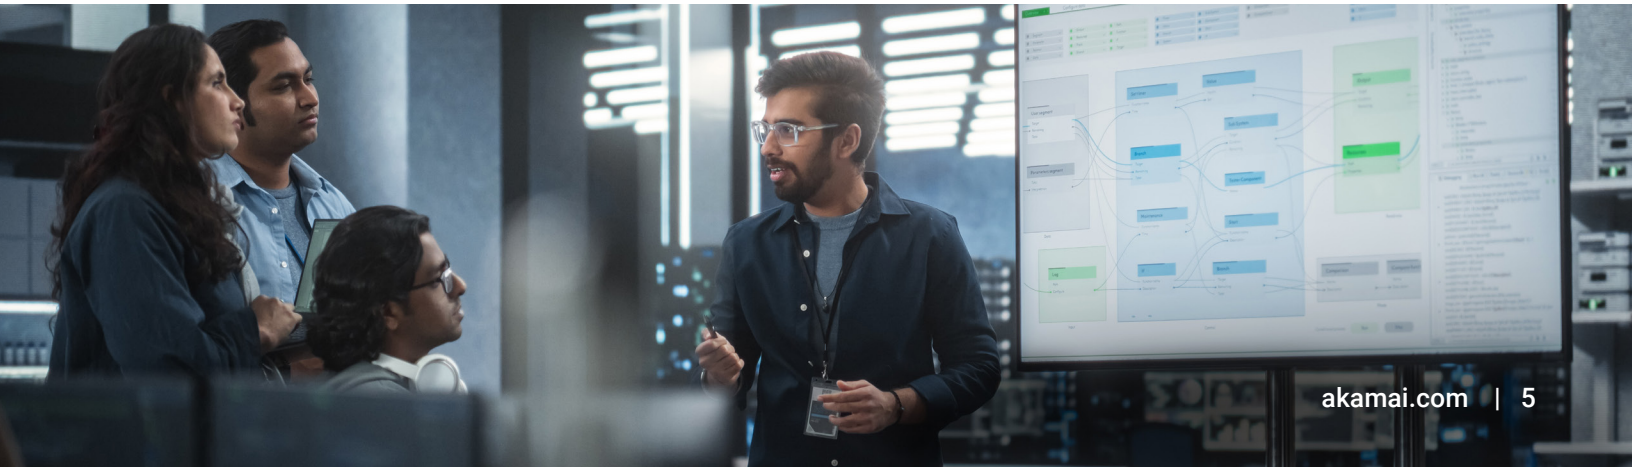

These approaches open the possibility of a catastrophic breach

It's theoretically possible to manage application access securely and with minimal friction using these approaches. You may already be using some combination of them. The problem is that to implement them well, maintain them, and provide proper security and performance over their lifetime is often far too operationally complex to always get correct. In many cases, companies convince themselves that because employees can access their applications, everything must be working optimally. They then find themselves caught off guard when one of these quick fixes results in a catastrophic breach or degrades performance so severely that there's either an outage or employee productivity is significantly limited.

A Zero Trust approach to application access

Given the inherent flaws in perimeter security approaches and the specific challenges they present in managing access to applications, the emerging Zero Trust cybersecurity model offers a better alternative. First introduced by Forrester Research in 2010, it is a framework that companies are using to transform their IT infrastructure, security policies, and business processes.

The principle behind it is quite simple, but very powerful: Trust is not an attribute of location. You shouldn't trust something simply because it is behind your firewall. Instead, any action, no matter where it occurs, should only be trusted if it has been explicitly allowed. Ultimately, only that which *should* happen *can* happen. Remove all implicit trust for actions that are not required because they create risk but not value.

This requires strong authentication and authorization, and systems shouldn't transfer data until trust has been established. In addition, analytics, filtering, and logging should be employed to verify behavior and to continuously watch for signals of compromise.

This fundamental shift defeats a vast amount of the security compromises we have seen in the past decade. No longer can attackers exploit weaknesses in your perimeter and then harvest your sensitive data and applications because they made it inside the castle. Now there is no moat to cross for access. There are just applications and users that must mutually authenticate and verify authorization before access can occur.

Zero Trust Network Access

ZTNA is an architecture built on these principles that grants secure access to applications and resources on the basis of strong authentication, authorization, and context. A ZTNA architecture provides access only to the applications that users need to do their job, not to the entire network. With a ZTNA approach, it no longer matters where users are located; there's no longer the concept of inside or outside the perimeter. Where an application is hosted is irrelevant — on-premises, public cloud, or private cloud — because authenticated users only get access to applications that they have been authorized to use.

For example, an employee in sales will only have access to applications that relate to their sales role, not to human resources or finance applications.

How Akamai's ZTNA works

Akamai Enterprise Application Access and Akamai MFA allow you to move to a ZTNA architecture, which can be an important and critical step in your journey toward Zero Trust.

Enterprise Application Access is an Identity-Aware Proxy (IAP) in the cloud. It's a flexible and adaptable service with granular decision-making based on real-time signals, such as threat intelligence, device posture, and user identity information. Akamai MFA is a multi-factor authentication service that provides the strongest levels of authentication to ensure that a user requesting access is who they claim to be.

To get started, you run a small virtual machine called the Enterprise Application Access connector behind the firewall, but with connectivity to your applications. It does not need to be, nor *should* it be, inside your DMZ. Its address should be on private IP space and not directly reachable from the internet. In fact, it should look exactly like any other application you would place behind the firewall.

To support multicloud environments, a connector can be deployed inside your on-premises data center or in a private or public cloud.

The Enterprise Application Access connector immediately establishes an outbound encrypted connection to the IAP on Akamai Connected Cloud. Once connected to the IAP, the connector downloads its configuration and is ready to service connections. The connection between the connector to the IAP is outbound, which allows you to close all inbound firewall connections, making the applications nearly invisible on the public internet.

The IAP performs all the pre-processing that happens before a user is connected to the application, including authentication, authorization, and device security and posture checks. When a user attempts to access an application, they are directed to Akamai via a DNS CNAME and connected to the IAP. Assuming your end user and their device pass all checks, they are then routed for authentication, multi-factor authentication, and single sign-on, after which device identity functions are performed.

After the user and machine are authorized, the connection from the end user is stitched together with the outbound connection from the Enterprise Application Access connector. Traffic from the user session flows through this stitched IAP, which then connects to the requested application or service. At that point, a complete data path is established, and all access decisions are then continuously and dynamically enforced on the basis of identity, device, and user context.

There are distinct and significant advantages to this method of access. The activities that are most performance- and security-sensitive take place at the edge, closest to the end user, where Akamai has more than 4,200 locations across 134 countries.

Additionally, the sensitive ingress path into the application happens over a reverse application tunnel, effectively removing the IP visibility of the perimeter and reducing the risk of volumetric attacks.

Because Enterprise Application Access can integrate directly with a company's identity infrastructure even if it uses multiple directories and identity service providers, the ZTNA service can be deployed quickly with no need to change the existing identity infrastructure or architecture.

For legacy applications that do not support modern authentication protocols, Enterprise Application Access has an IdP bridge capability that provides authentication to SAML-based IdPs and translates the authentication token into the authentication protocol supported by the legacy applications.

What makes IAP-based approaches like Enterprise Application Access so appealing is that they provide application-level access. With application-level access, performance and security are *decoupled* from complexity.

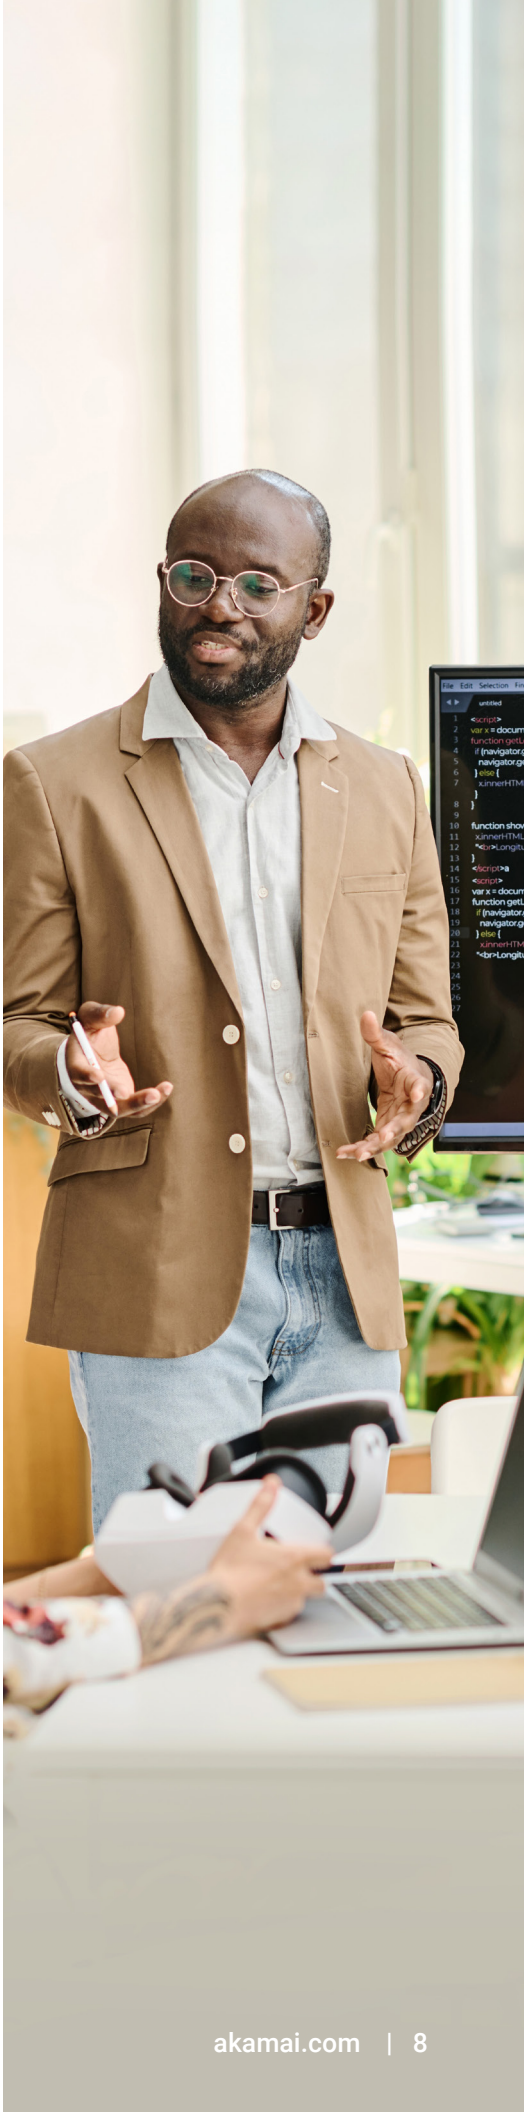

You simply take all applications that have locality with one another (all hosted in the same data center or same virtual private cloud, for instance), place them into a private network IP space or a restricted VLAN, and place an access proxy in that microperimeter. That's it — you're finished.

Application owners set their own security policies on the access proxy — policies about who can access what and why — and, even more compelling, users can be anywhere. There is no distinction between on-premises and off-premises because there is no network perimeter that includes the end users. An employee working in a coffee shop is equal to an employee working in your office. All that matters is whether the user is authorized and whether the machine is safe.

With application-level access, you get best-in-class performance, despite the ease of deployment and use. Users simply go to the internet to access applications directly, no matter where they are hosted or where they appear, allowing the internet to route packets to their destination without having to go through aggregators or intermediaries that aren't in their path.

In fact, with application-level access, internal networks often dissolve into simple guest Wi-Fi. Remember, for Zero Trust to truly be effective, you cannot treat internal users any differently than external users. No one is trusted by default.

ZTNA's desired end state

All users, whether they are on- or off-premises, should be required to access all applications through identity-aware access proxies, regardless of where the applications are hosted. These proxies should perform not only standard authentication, but also use phish-proof multi-factor authentication, such as Akamai MFA. Additionally, there should be robust device posture capabilities that obtain device criteria to allow access to specific applications.

We strongly believe that ZTNA does not end with authentication and authorization. To support Zero Trust principles, all the parameters that are checked at the initial authentication and authorization stage should be continuously monitored during the activation session. Any detected changes should trigger an action — for example, re-authenticate the user, remove access to the application, or limit access to the application.

One crucial security system that should be layered on top of your access proxies is web application and API protection (WAAP), which will ensure that end users are not launching application-level attacks (intentionally or inadvertently) toward your internal applications. You can leverage other advanced systems such as human/bot detection for non-API sites to help ensure that malware is not masquerading behind valid endpoints. It is at the IAP where Akamai can layer in WAAP, bot detection, behavioral analytics, and caching. This is designed to provide best-in-class performance, as well as the ability to keep potential threat actors as far away from your physical locations, applications, and data as possible.

As you bring your applications online and make them accessible through access proxies, distributed denial-of-service (DDoS) prevention becomes even more important. You should align yourself with providers who can absorb attacks against your microperimeters and access proxies, allowing continued operation under intense loads.

And, finally, to ensure performance is best in class for your applications and that users not only accept this shift in access but also champion it, your access proxies should be fronted by networks that can provide performance benefits. Specifically, content delivery networks and internet routing overlays should be part of your arsenal to not only make access available, but also make it more performant than prior methodologies ever allowed.

Threat protection

Solutions like Akamai Enterprise Application Access can protect your applications from malicious actors. But what about protecting users from inadvertently becoming those very actors through compromise, such as through a device infected by malware or credentials stolen via a phishing link and landing page? This is where prevention and detection become crucial for web traffic.

One approach is to deploy a cloud-based DNS firewall solution such as Akamai Secure Internet Access. This product inspects every DNS request that users make and applies real-time threat intelligence so that benign requests are resolved as normal but any requests to malicious domains are proactively blocked. This reduces the risks of employees' devices being compromised with malware or ransomware or falling victim to a phishing attack.

Summary

Traditional hub-and-spoke networking architectures, along with the castle-and-moat security perimeter they utilize, simply cannot effectively provide performance or security in today's cloud-and-mobile world. This is a problem all companies must begin to address or they will be left vulnerable. The failure to transition to safer enterprise security architectures is the number one cause of corporate breaches today, and the number of breaches is only going to increase. Simply put, you are not safe behind the perimeter, because the perimeter itself no longer exists.

Next steps

How do you start transitioning to a Zero Trust Network Access architecture?

Akamai's cloud security services can be combined to build a comprehensive ZTNA architecture, not only enabling safe application access in a multicloud world, but also leveraging the cloud to remove the need for internal corporate networks almost completely.

By utilizing our advanced distributed IAP and our phish-proof multi-factor authentication, along with the power of Akamai Connected Cloud, you can finally move to a perimeterless world in an extremely easy way — phasing in applications, reducing your migration risk profile to near zero, and leveraging Akamai's extensive history of proven performance and security solutions.

As you continue on your Zero Trust journey, you can rest assured that Akamai will be there with you at each step, helping you to transform your network into an architecture that not only provides access to your applications and data, but also does so in an easy-to-manage way while maintaining the highest levels of security and performance.

Learn more about meeting your business needs with the Akamai Zero Trust portfolio.

Akamai powers and protects life online. Leading companies worldwide choose Akamai to build, deliver, and secure their digital experiences — helping billions of people live, work, and play every day. Akamai Connected Cloud, a massively distributed edge and cloud platform, puts apps and experiences closer to users and keeps threats farther away. Learn more about Akamai's cloud computing, security, and content delivery solutions at akamai.com and akamai.com/blog, or follow Akamai Technologies on X, formerly known as Twitter, and LinkedIn. Published 02/24.