WHITE PAPER

A Blueprint for Building a Zero Trust Architecture

Table of contents

Introduction

The concept of Zero Trust has been around since 2009, when Forrester Research first promoted it, warning organizations that it was time to overhaul the traditional method of granting unfettered access to any user or application that passed the network perimeter. Instead, every device, user, and network flow should be verified before full access is granted. In the ensuing years, thanks to many factors, the urgency to adopt the concept of Zero Trust has only grown.

Today's hybrid workforce operates from various locations, with BYOD programs enabling employees to access corporate applications and resources using both managed and unmanaged devices. Applications are hosted everywhere — cloud, on-premises, and hybrid. The net result of these changes is that the network perimeter as it was no longer exists. Ransomware attacks have grown in frequency and sophistication, increasing the chances of an attacker breaching your defenses and increasing the costs once they do. The average cost of a data breach in the United States is the highest in the world, US$9.36 million, according to the IBM Cost of a Data Breach 2024 Report. Additionally, the growth of network-connected devices, like Internet of Things (IoT) devices, and the additional requirements for network access by partners and customers have combined to significantly expand a company's attack surface.

Amid this evolving cybersecurity landscape, network and security software vendors have rushed in to brand their existing products as Zero Trust or to introduce new products, while consultants and analysts are introducing new acronyms and market definitions. It has left security teams struggling to explain these sometimes complex concepts and make purchasing decisions that set the foundation for a shift to a Zero Trust strategy.

This white paper is designed to provide security teams with a blueprint for making investments in Zero Trust technology by identifying where to start and by outlining the key differentiating factors.

Hybrid work, cloud apps shatter the network security paradigm

When, how, and where people do their work have moved beyond the four walls of an office.

As a result, the network perimeter no longer exists — at least not in any recognizable form. Your users are just as likely to be outside the proverbial moat as they are to be inside. And the applications they use as software as a service (SaaS) and multicloud implementations are proliferating. With advanced and persistent threats, you are highly likely to inadvertently let the malicious actors have full access to your most valuable assets once they are inside the network. Once they're inside, if you don't have a comprehensive Zero Trust program in place, the malicious actors have free rein.

This isn't just theory. It's evident in the widespread — and costly — data breaches in recent years, the vast majority of which happened as a result of abused trust inside the network perimeter.

Meanwhile, the applications that were designed to live inside a network perimeter often have the worst security profiles. After all, if you were a developer who assumed that only authorized employees with good intentions could reach your system, would you have been as defensive as the coder today who knows vast armies of hackers will try to exploit their internet-based application?

The solution to these challenges, across the marketplace, is Zero Trust.

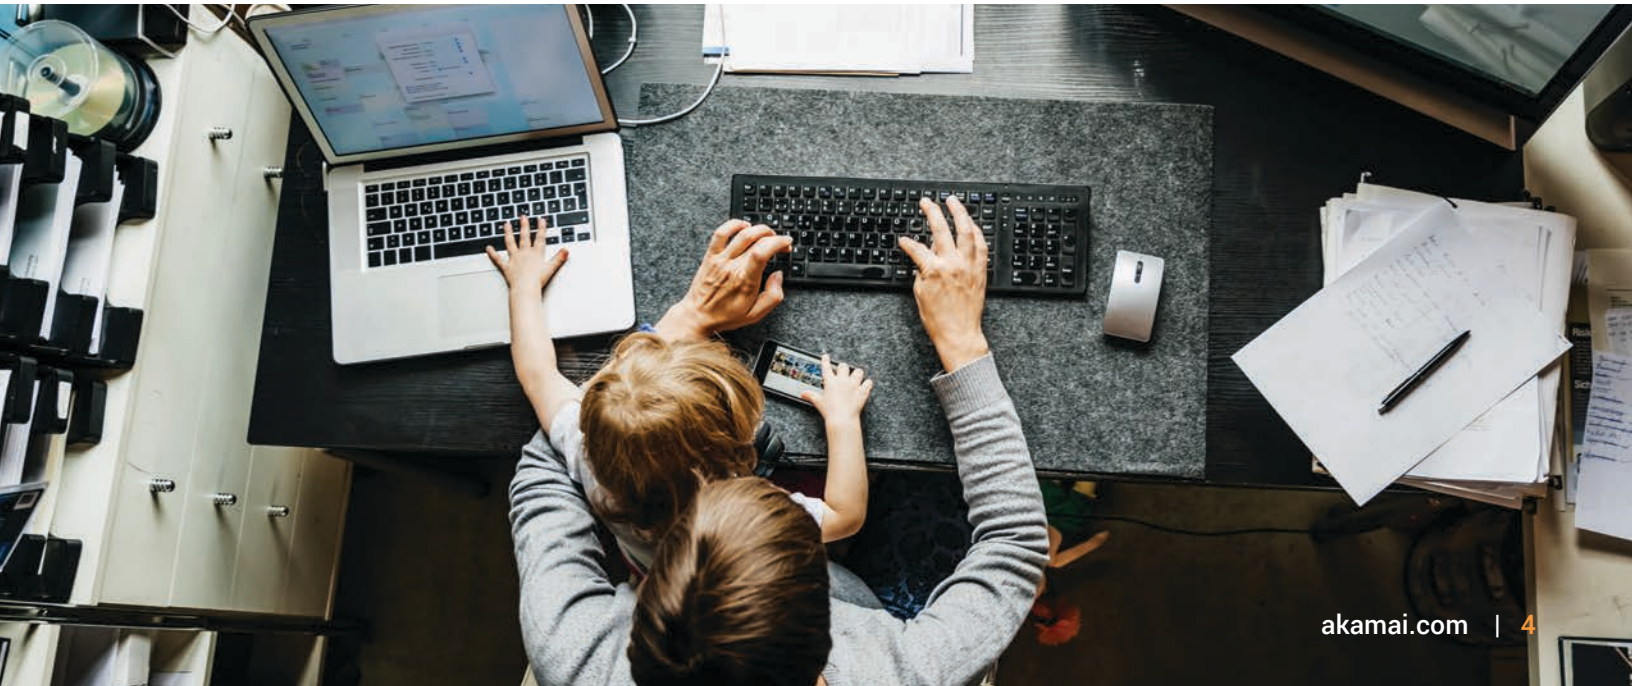

A Zero Trust security architecture

The principle behind Zero Trust 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. Organizations need to remove all implicit trust for actions that are not required. For example, giving all users from your accounting group access to the financial system, when only a handful need it, creates risk but not value.

The method of proof for this is strong authentication and authorization, and systems should not transfer data until trust has been established. In addition, analytics and logging should be employed to verify behavior and to continually watch for signals of compromise.

This fundamental shift defeats a vast amount of the compromises we have seen in the past decade. No longer can attackers spend time exploiting weaknesses in your perimeter and then exploit your sensitive data and applications once they've made it inside the moat. Now, there is no moat. There are just applications and users, each of which must mutually authenticate and verify authorization before access can occur.

Traditional security architecture

Modern reality

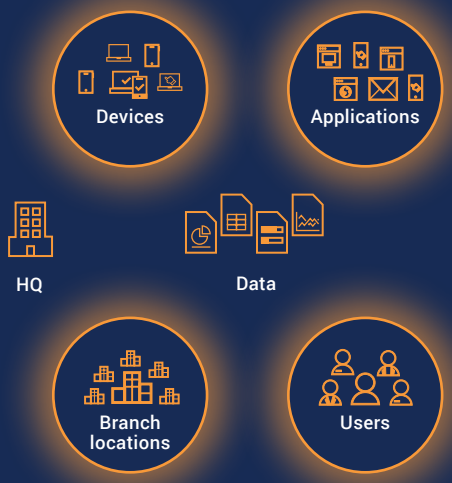

How does an organization create a Zero Trust architecture?

First, all companies need to map out a strategy for their existing landscape and determine if and when they need to hire new talent for their workforce. An entire paper could be dedicated to this key step in the process, but the actual products that can help fulfill a Zero Trust strategy should be driven by three goals.

1

Don't trust any entity; constantly verify.

"Don't trust and constantly verify" sounds far easier in the abstract. If you simply cut off all access to all systems and all data, you have locked down your network. The real challenge is to constantly verify without creating massive business disruptions, particularly when most systems were designed with implicit trust in mind. You need broad visibility and control for all types of access, and simple and practical means of enforcing and maintaining policy.

2

Once verified, ensure you're providing minimal access.

In a Zero Trust environment, once a user has been verified, they must be given access to only what is required by their role.

3

Continuously monitor for threats.

As most industry experts will tell you, Zero Trust is an ongoing exercise. Threat actors are becoming increasingly sophisticated as they try to breach a company's defenses and the organization must continuously monitor, verify, and limit access. One of the advantages of a Zero Trust model is that it is not focused on what threat actors are doing, but rather, what the business itself is doing. With a true Zero Trust policy in place, attackers are hard-pressed to subvert all the things that your business needs to run at once. Ideally, you'll be able to stop every attack at some point in the chain. That includes the ability to stop attacks that have not yet been conceived. You won't care whether or not it's a zero-day attack, Zero Trust can help mitigate it.

The dark side of Zero Trust

However, as an organization embarks on implementing Zero Trust, it must also consider the flip side of all this distrust and the limits on access. A fundamental aspect of Zero Trust is restricting access, primarily through allowlisting. This is the practice of dictating what can happen; everything else is denied by default. However, by decreasing an attacker's ability to carry out their malicious campaign, an organization can increase the likelihood of accidentally preventing someone from being able to do their job. Alternatively, repeated checks of workloads and devices can cumulatively add up to delays and frustrations. A Zero Trust strategy that keeps people from effectively doing their jobs is no kind of strategy at all.

A strong Zero Trust strategy will therefore strike a balance between security and access. It also needs to strike a balance between what can effectively be accomplished and the resources — both budgetary and personnel — of your security team.

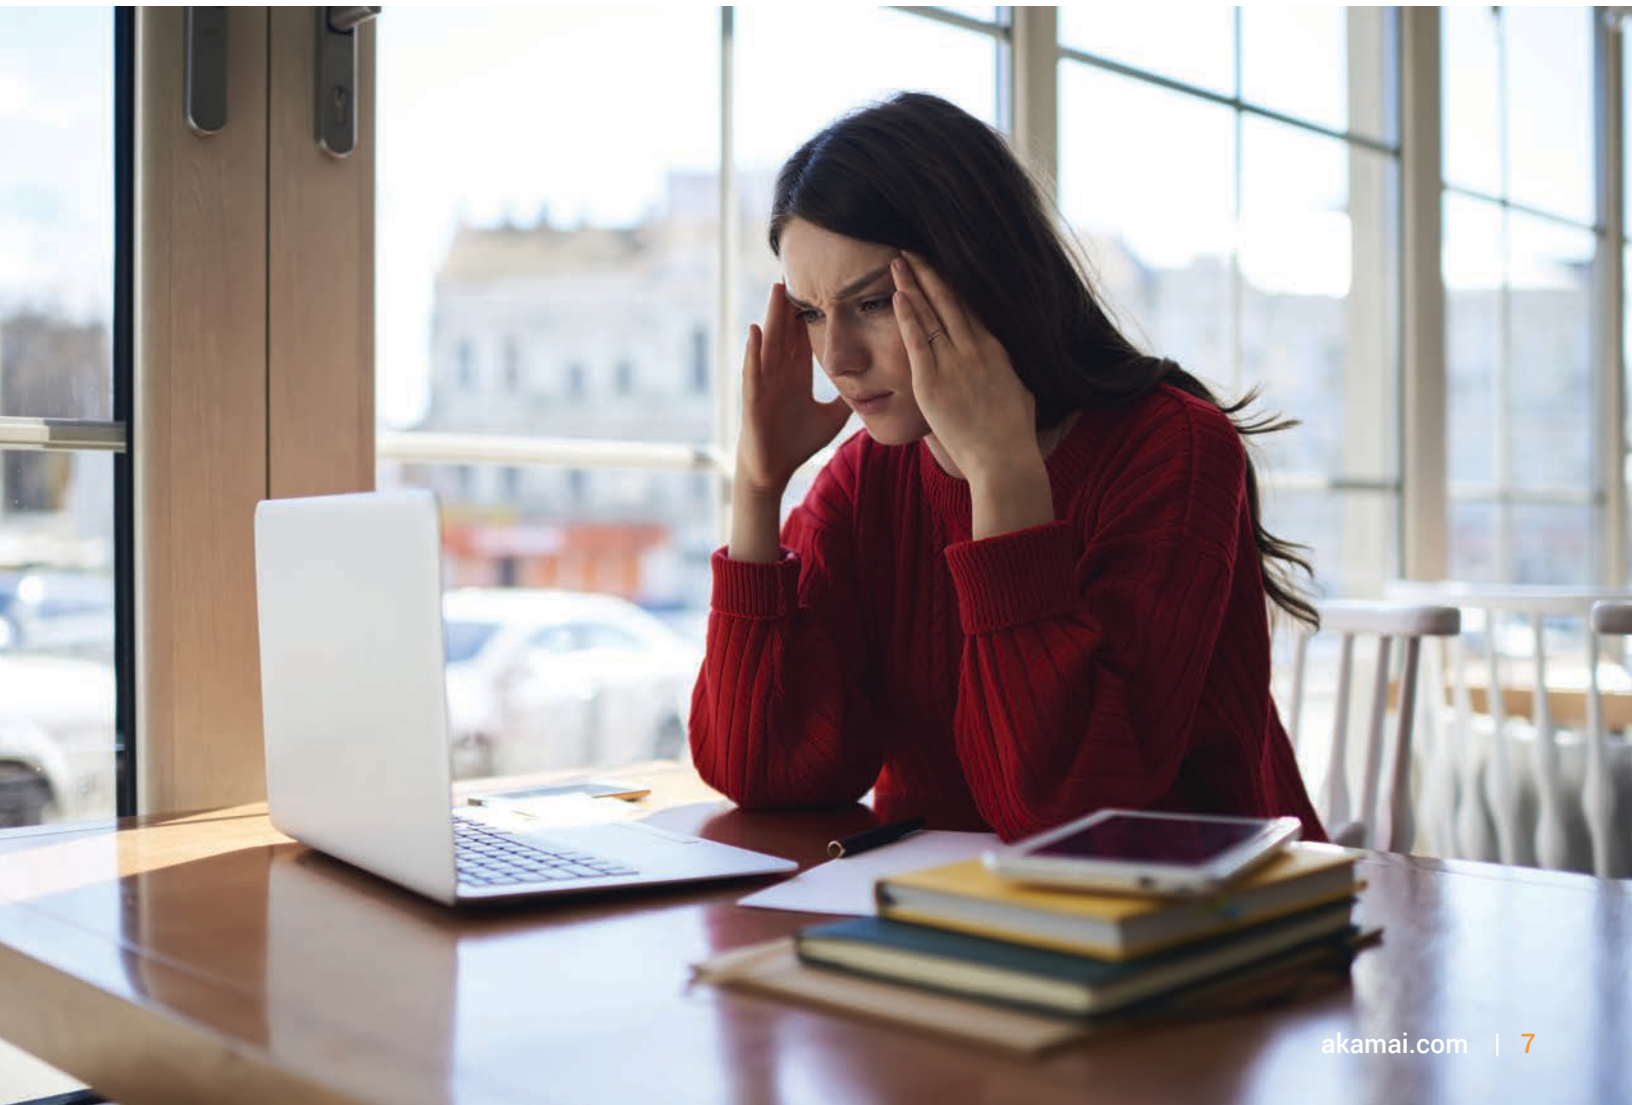

Elements of Zero Trust

It's been 15 years since Forrester first outlined the concept of Zero Trust. Many organizations are just now beginning their Zero Trust journey, and they are confronting a convoluted marketplace of software products. Some products have been around for years and address parts of a Zero Trust architecture; other new products have emerged; and plenty of software providers have been quick to rebrand their offerings under the Zero Trust moniker. Moreover, as many analysts and industry observers will inform you, "Zero Trust is not a product, it's a comprehensive strategy" and "Zero Trust is not a destination, it's a journey." Yet, these oft-repeated claims do little to help those who are confronting purchasing decisions for Zero Trust technology solutions, and, in fact, can create more confusion.

Because there is no single product that provides a company with Zero Trust and because individual organizations will have different priorities and vulnerabilities, the starting point is going to be different for every company. Yet, thanks to technological advancements and industry consolidation, companies are now able to obtain the tools required to implement a Zero Trust policy from a single source. The analyst firms are beginning to recognize this, as well.

The principles of Zero Trust

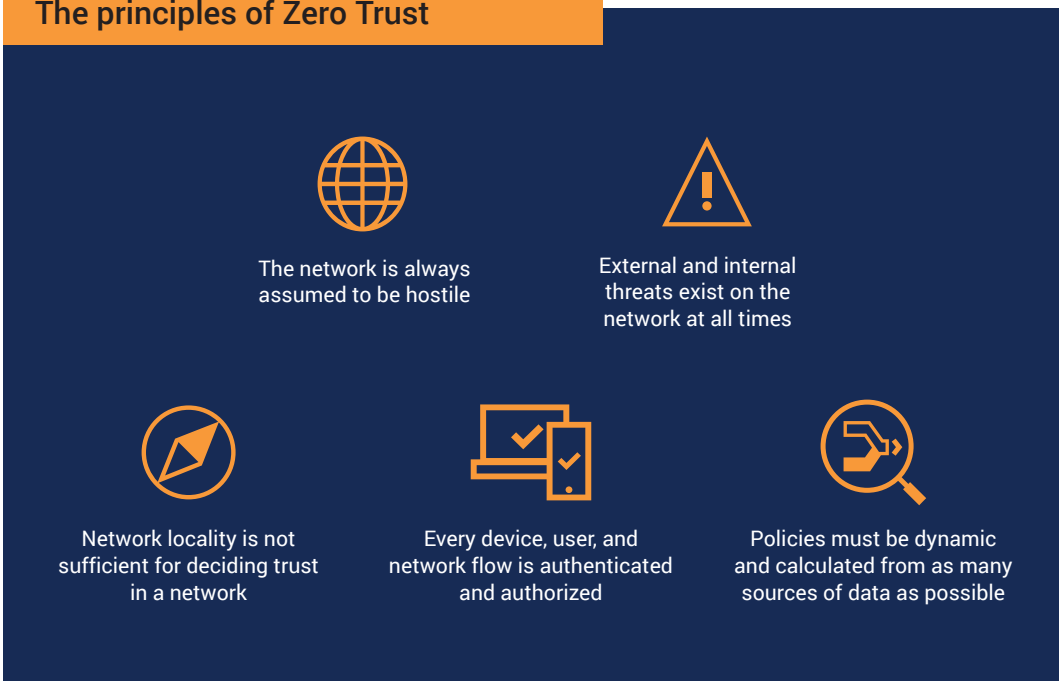

Gartner tracks what it calls the Secure Service Edge (SSE), a combination of secure web gateways, cloud access security brokers, and Zero Trust Network Access (ZTNA). In its report, What Are Practical Projects for Implementing Zero Trust?, Gartner includes microsegmentation, as well (which it calls workload-to-workload segmentation), recommending that "Organizations looking to move to practical implementation should focus on two primary projects: user-to-application segmentation (ZTNA) and workload-to-workload segmentation (identity-based segmentation)."

Similarly, IDC breaks Zero Trust down to secure access and segmentation, which it defines as a comprehensive view of emerging and legacy technologies used to protect computing systems, resources, and data through logical segmentation, access control, and threat detection.

However, assembling these separate systems into one cohesive strategy becomes a core challenge. What are the key elements that CIOs, CISOs, and other security professionals should be looking for as they construct a Zero Trust architecture that works for their organization?

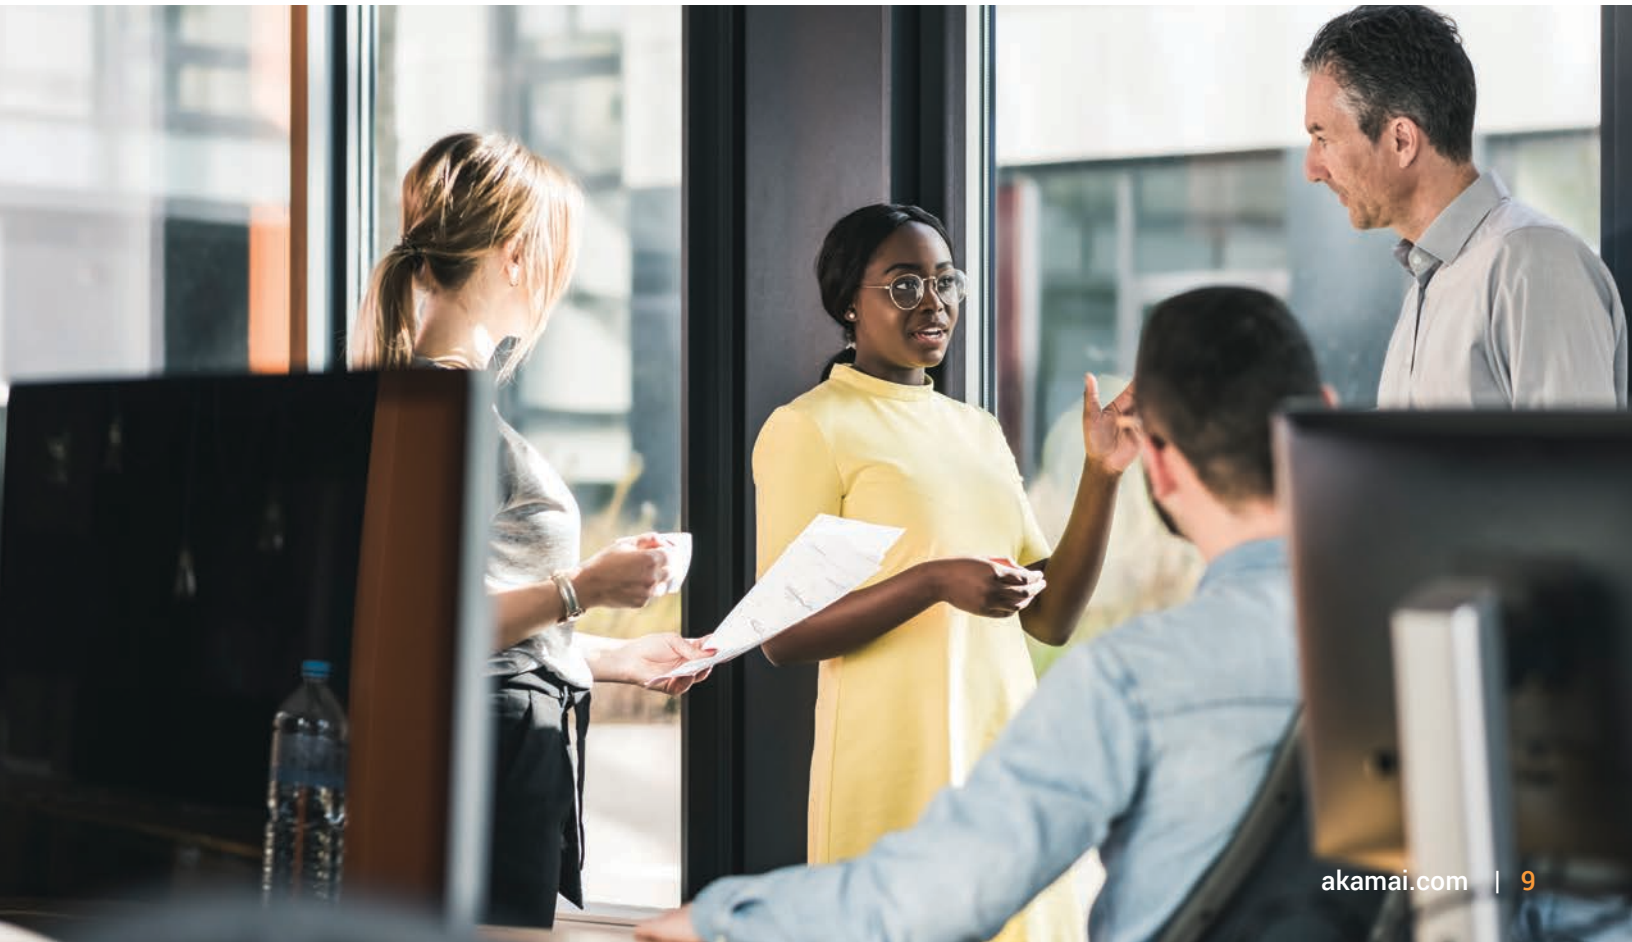

Zero Trust Network Access

Sometimes confused with the overall approach to Zero Trust, ZTNA is a fundamental part of the technology stack. Secure access is the key initial step in any Zero Trust framework. Unfortunately, like so many elements of the process, it quickly becomes more complex than it sounds. Secure access is not a binary decision. Providing the right level of access to the right application for the right users at the right time has become far more complex as users and applications have become more widely distributed. In fact, the very definition of a user can now include customers, suppliers, and partners, as well as employees. Meanwhile, applications can include legacy apps, SaaS, or mobile apps and require access to and from the data center, internet, or cloud environments.

An effective ZTNA solution will verify the identity of the user and the health of their device, and verify that they can access the applications they need, no matter where they are, which reduces the possible attack area and improves flexibility and monitoring. For decades, organizations relied on virtual private networks (VPNs) supported by identity providers to provide access. Those VPNs, designed for a different era, are no longer sufficient for the size and scope of today's distributed workforce. ZTNA has evolved to become more than just a replacement for VPNs and now grants access based not only on verifying the identity of the user and their device but also attributes like time and date, geolocation, and device posture to grant the appropriate level of trust.

Key considerations for purchasing Zero Trust Network Access solutions

As companies begin to replace their older VPNs with more sophisticated identity management solutions, there are a number of areas to consider. Today's more advanced solutions should combine identity and access management, application security, multi-factor authentication (MFA), and single sign-on, all with management visibility and control under a single interface. Organizations pursuing Zero Trust initiatives should look for solutions that can address their current needs but also scale with the business, allowing them to quickly onboard employees from a merger or acquired company, enable manufacturing or production in different markets or geographies, easily add and remove contractors to adapt to changing business needs, and move applications to the cloud cost-effectively without sacrificing security.

Organizations should seek solutions that can integrate directly with existing identity infrastructures, even if they include multiple directories and identity service providers. This allows the ZTNA service to be deployed quickly with no need to change the existing identity infrastructure or architecture.

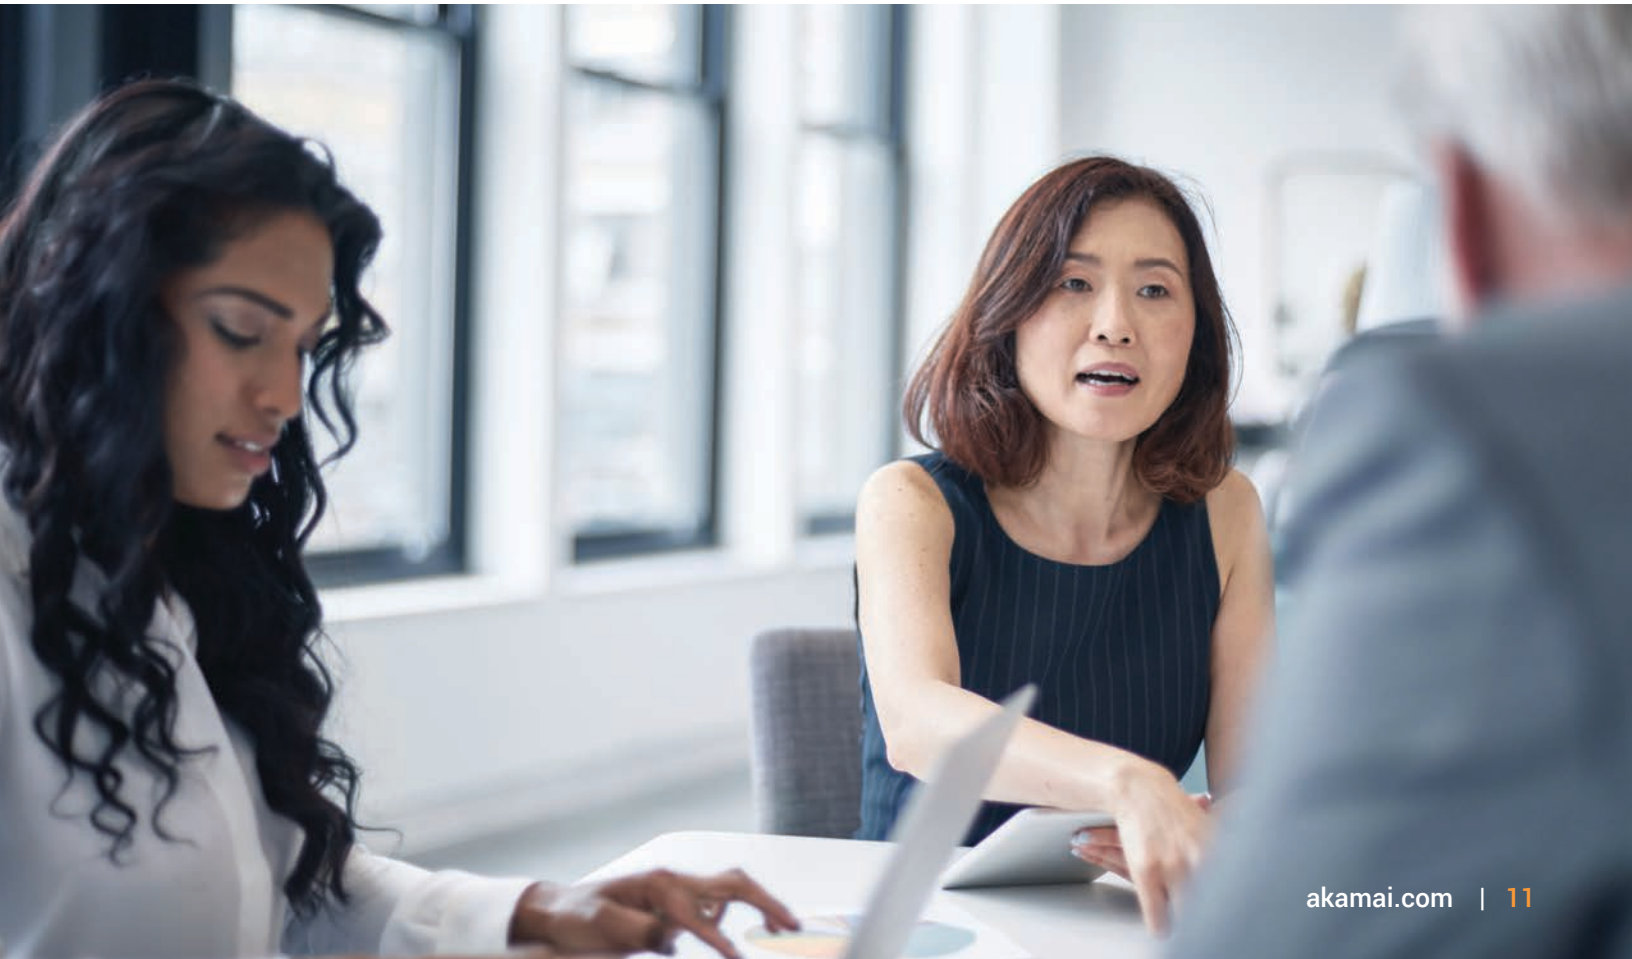

Look to the edge

There is also a significant differentiator among the products in the marketplace that the Zero Trust purchasing teams may not take into consideration, but definitely should. Solutions that are combined with edge cloud platforms can offer additional benefits by acting as an identity-aware proxy that abstracts connectivity to the edge platform, ensuring that all authentication is done at the edge and away from the data center. Although some companies turn to access proxy architectures run within the DMZ, this fails to take advantage of the cloud's ability to better absorb attacks, provide bandwidth for caching, and autoscale as needed.

An identity-aware proxy built in the cloud can scale on demand, run CPU-heavy resources, and absorb attacks. Moreover, it sits on a private IP address that is not directly reachable from the internet. The activities that are most performance- and security-sensitive take place at the edge, closest to the end user. 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.

> **Solutions that are combined with edge cloud platforms can offer additional benefits by acting as an identity-aware proxy.**

Multi-factor authentication considerations in building a Zero Trust blueprint

The rise of hybrid work and the need for greater access mean most organizations have already embraced MFA and have some sort of solution in place. It's important to recognize, however, that the combination of enterprise-wide access and MFA is greater than the sum of its parts. MFA is central to the concept of trust because it requires you to have more than just a password. You need a second verification to ensure that you're not falling prey to one of the most commonly abused areas of trust. It's also important to remember that not all MFA solutions are created equal.

When evaluating MFA solutions as part of a Zero Trust strategy, organizations should look for solutions that are:

Integrated with identity management and enterprise access

Compliant with FIDO2 to ensure that user credentials are decentralized, isolated, and encrypted on users' personal devices, which is particularly important in fending off phishing attacks

Able to verify users via their smartphone without relying on a physical key

Microsegmentation

There is no perfect state of Zero Trust. Inevitably, there will be gaps that the most persistent attackers may find and exploit. Any comprehensive approach to Zero Trust will therefore require microsegmentation. Today, most networks either have no segments or very few segments. Indeed, organizations have traditionally protected their critical applications with firewalls, but that can prove difficult for a number of reasons. Firewalls basically require you to enforce network policy, creating a choke point. You need network connections to go through a firewall, which gets expensive quickly, is unaware of many of the risks in modern network traffic, and is extremely difficult to change. Instead, organizations are turning to software-based microsegmentation, which simplifies many of these labor-intensive processes.

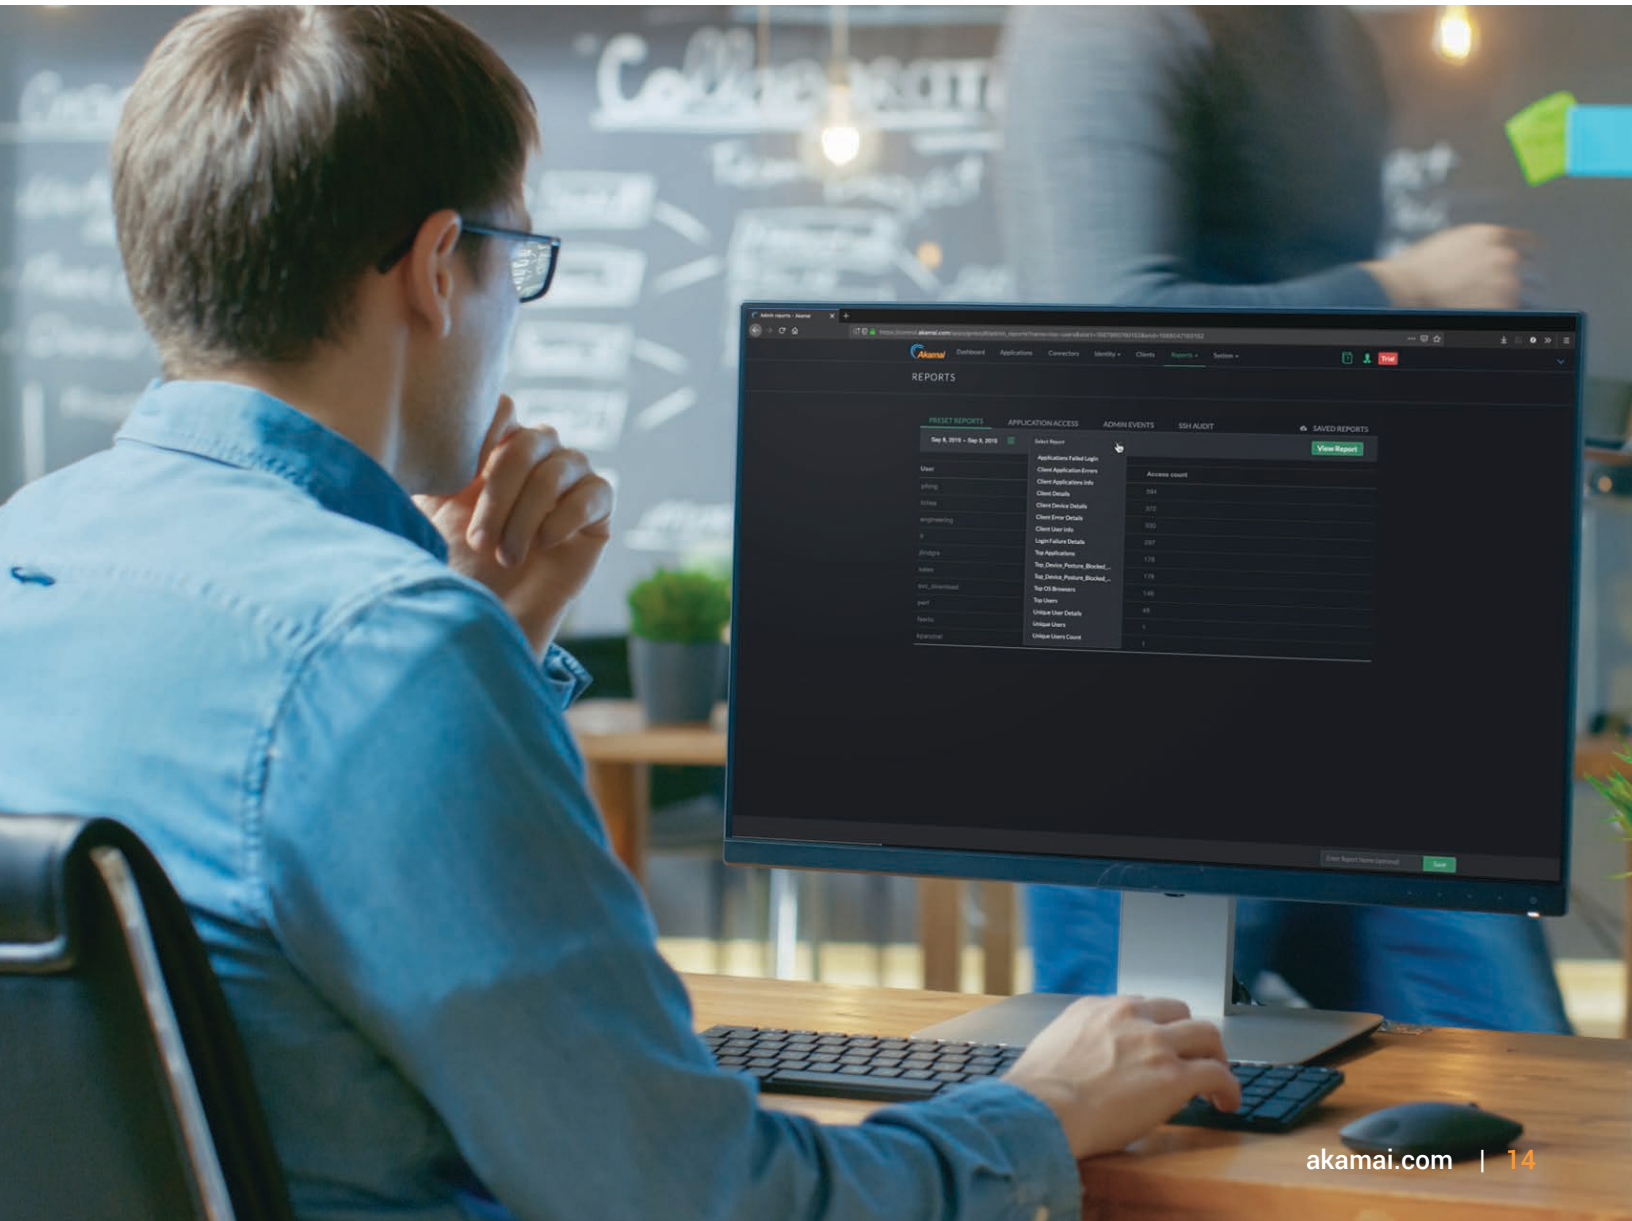

Differentiators in microsegmentation

Although it's a core requirement of any Zero Trust initiative, microsegmentation has often been considered separately from core ZTNA solutions. And, though microsegmentation is sold by both security platform providers and as a stand-alone solution, there are some core differences buyers need to understand.

Where can I deploy it? Microsegmentation solutions that were built as network tools instead of with a security-first approach and those built for on-premises systems should raise a red flag for prospective buyers. Today's tools should be deployable in the cloud, in on-premises environments, on devices (including those on which you can't install agents), and among containers across hybrid environments. That's typically going to require cloud-based software. If a microsegmentation solution can only cover 80% of your environment, that's not sufficient.

How much visibility does it provide? Although microsegmentation solutions restrict access, too much restriction can break business processes and lead to calls from the COO. Microsegmentation requires a sophisticated understanding of your environment. Which servers can access which servers? Can you define policies between a Kubernetes cluster and a Windows 2008 server? A lot of tools don't have agents that go back as far as 2008 or are as forward-thinking as enforcing policy on Kubernetes. Your microsegmentation software must be able to address these sorts of complexities if you're going to deploy Zero Trust effectively.

Additionally, microsegmentation software buyers need to consider the granularity of the policies the product will support. Most systems will enforce policies at the application layer across ports and processes. More sophisticated products can enforce policies at the microservices layer. For example, attackers can use some of the services of svchost, like Task Scheduler, to move laterally throughout the network. However, companies can't block svchost outright because it does too many important things. A microsegmentation solution that enforces policy at the microservices layer can make a difference there.

How difficult is it to implement? How easy it is to express the policy you need now and, just as important, what you will need in the future should be core considerations for any microsegmentation solution. Whether you're in a planning phase or there is a threat to your environment that you need to lock down, make sure the engine you invest in will easily support both.

Starting with allowlisting in a microsegmentation project can be intimidating for security teams thanks to the risks of incorrectly denying a needed application or service. A sophisticated microsegmentation solution should come with denylisting templates that teams can roll out quickly and easily to establish some quick wins for the project. Once that's accomplished, organizations can continue the journey toward comprehensive allowlisting protection that includes accurate dependency and contextual inventory mapping capabilities.

> **Microsegmentation solutions that were built as network tools instead of with a security-first approach and those built for on-premises systems should raise a red flag for prospective buyers.**

DNS firewall

In a Zero Trust environment, it's not just people who can't be trusted, but the internet itself. Employees need access to the internet and, as SaaS and mobile applications, cloud services, hybrid work, and IoT devices spread, so too does an organization's attack surface. Protecting the organization and users against threats such as malware, ransomware, phishing, and data exfiltration becomes exponentially more difficult. Organizations have limited resources to manage security control point complications and complexities, and security gaps in legacy on-premises solutions.

Enforcing Zero Trust between a person and the internet requires a DNS firewall, which becomes a central feature of any Zero Trust initiative.

Core Zero Trust requirements of DNS firewall investment

While seemingly straightforward, there are requirements that technology buyers must consider when investing in a DNS firewall. Many organizations have deployed on-premises DNS firewalls but now need to extend that protection to users no matter their location. Similar to identity management, providers that have robust edge platforms typically have stronger DNS security thanks to the threat intelligence garnered from the extended platform. Decision-makers should carefully consider these core requirements.

DNS inspection. Providers should be able to provide real-time inspections of all domains with sophisticated threat intelligence and automatically block malicious domains. Solutions also need to be effective across all ports and protocols to protect against malware that does not use standard web ports and protocols. The quality of DNS inspection can vary greatly across providers, and buyers should look for those with experience in the marketplace and established customer success.

Protection for all devices. Providers should have agents for devices that will be used on and off network, such as laptops, smartphones, and tablets.

Flexible DNS onboarding. Providers should have multiple methods to forward DNS requests to the DNS firewall to deliver maximum flexibility and cover all use cases.

DNS exfiltration identification and blocking. DNS exfiltration, especially low-throughput variations, can enable attackers to exfiltrate data over the DNS channel. Look for providers that have both inline and offline DNS exfiltration detections based on proprietary detection algorithms.

Threat monitoring

The final piece of core Zero Trust technology is threat monitoring. Although the assumption of Zero Trust is that nothing is implicitly trusted, organizations need to remain vigilant to uncover ongoing and emerging attacks, as well as potential risks (like misconfigurations or overly permissive access rights). As security teams evaluate the software on the marketplace, they should review the following three considerations for effective threat monitoring.

Key considerations

Effective algorithms
Sophisticated algorithms with a track record of success based on user and network activity anomalies, executable analysis, log analysis, and more should be part of any threat monitoring service.

Strong signal detection
Although software and artificial intelligence are vital tools in threat monitoring, Zero Trust decision-makers should still evaluate the internal expertise of the vendors with whom they're working. Threat monitoring services need to be able to separate the good signals from the bad to help avoid alert fatigue and provide immediate notifications of any incident. Organizations should also expect regular reports with analyses of any high-profile campaigns.

Experienced staff
Teams should include people with a broad range of backgrounds, including offensive, incident response, and data science, and should be available 24/7. This is an area where content delivery providers can add a substantial benefit. The insights from monitoring hundreds of terabytes per second contribute a unique perspective to any signal detection.

Where to begin?

A Zero Trust initiative is never complete, so for those considering the software, hardware, and hiring requirements, the primary question is often, “Which technology do we begin with?”

As with so many things, the answer is going to depend on a company’s individual needs, risk assessments, and relative strengths and weaknesses. For many industry observers, the answer is to begin by implementing ZTNA. Indeed, protecting the organization against malicious north-south traffic can be a prudent starting point. Yet, there are also those who believe that an east-west approach with microsegmentation, specifically software-defined microsegmentation, is the better route.

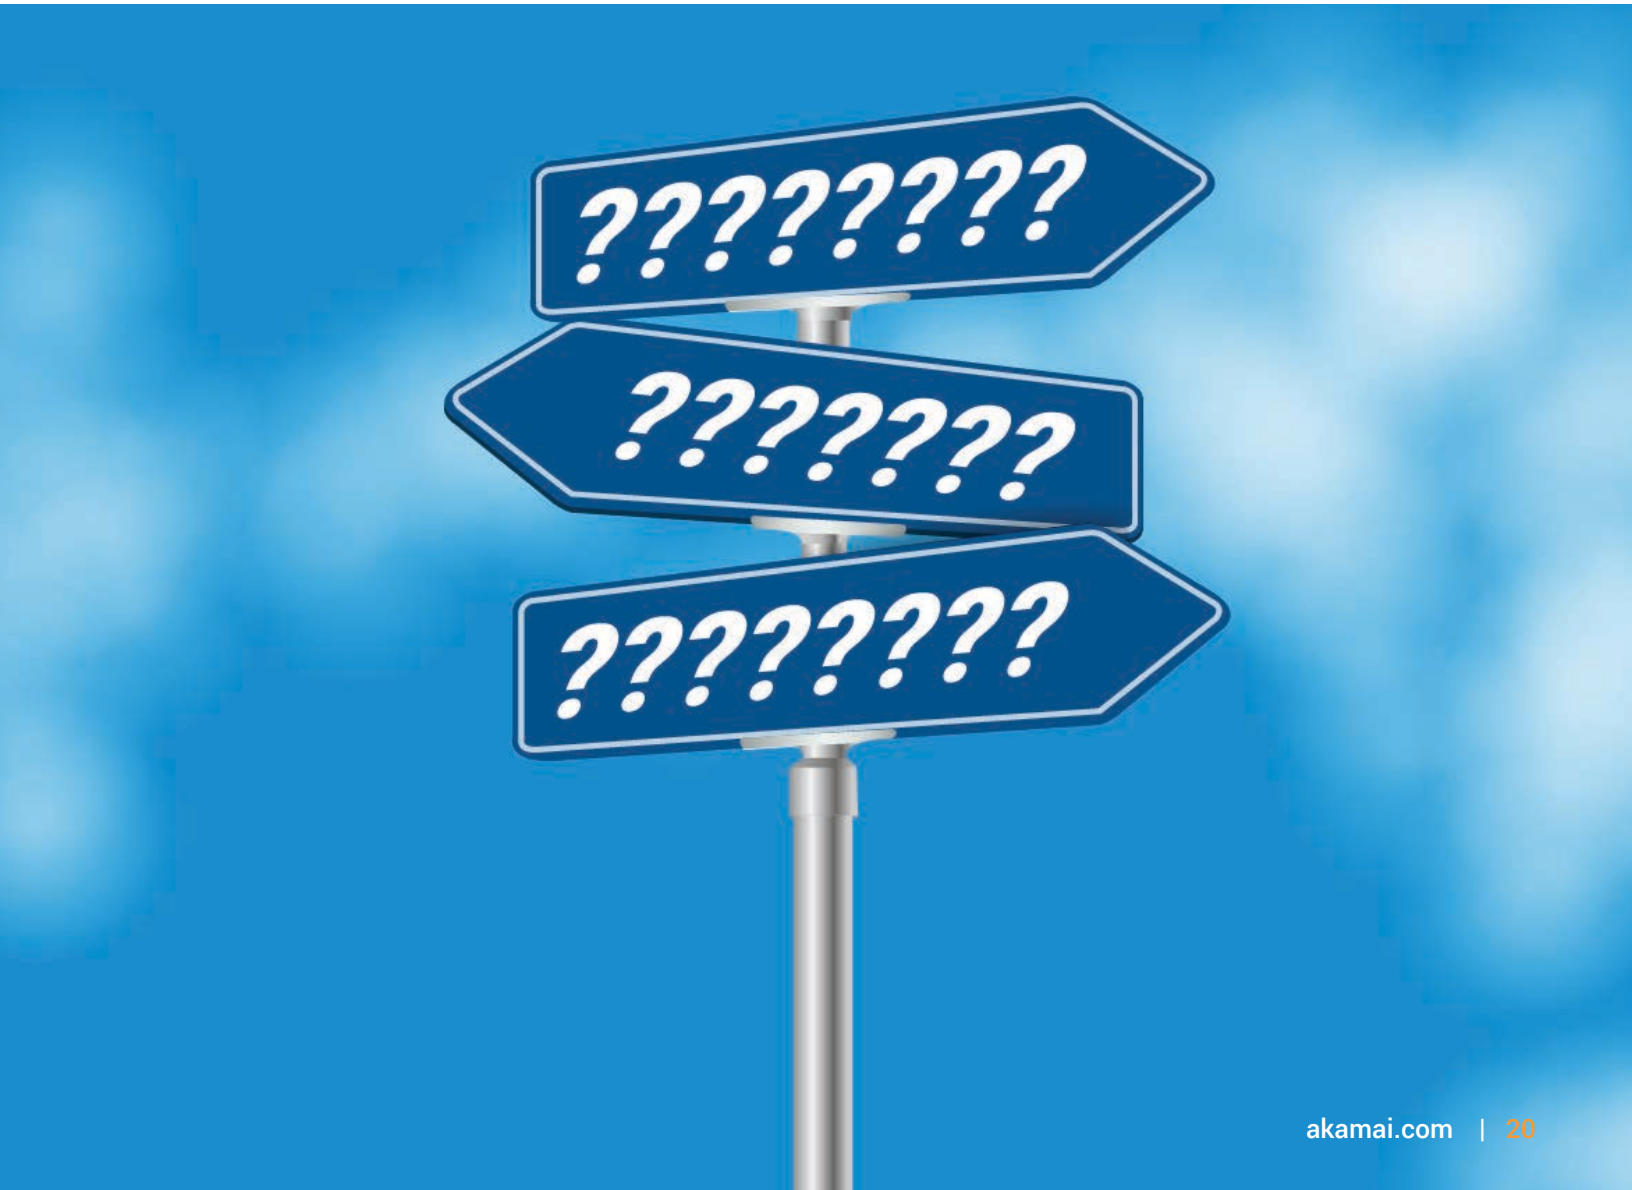

The case for starting with microsegmentation

If you believe, as most experts do, that there is no perfect defense and a malicious attack will eventually make its way through, then you want to be able to protect your most valuable assets. This is what microsegmentation offers. One reason organizations may be reluctant to start with microsegmentation is the perception of complexity.

First, microsegmentation is not an all-or-nothing approach. Like Zero Trust itself, it can be undertaken in stages. Organizations can begin by identifying their most valuable assets. Focus on what's critical. Ensure that if someone gets into your system, they can't bring your business down. The importance of an asset can be based on the data inside that asset or on the existing level of protection.

In many cases, you will want a microsegmentation solution that will cover your legacy systems, as these systems are often running business-critical applications and are especially vulnerable. There are some microsegmentation solutions that don't support securing those legacy systems.

Second, software-defined microsegmentation removes much of the perceived complexity. You will not need to deal with hardware or to call on your network architects and security architects repeatedly. You'll just roll out the software, which lowers the barrier to entry significantly.

Once a microsegmentation initiative has begun, the early benefits are clear and can help push the rest of the project forward. For example, you will now have a source of truth for what's happening in your environment. You can get that right away without even enforcing policy, and once you do, you'll have a great understanding of how flows are happening. Additionally, once an organization begins application ringfencing, you can quickly and easily lock down critical applications so they're only communicating over specific ports and processes.

Alternatively, a quick win might be targeting threat-specific policies. Sophisticated microsegmentation platforms will have denylisting built in. That means you can quickly create a policy to stop unnecessary connections between remote desktop services and the internet. Organizations can quickly close off the kind of vulnerability that led to the Colonial Pipeline attack, for example.

Whatever the starting point, the key to any ongoing Zero Trust journey is balance — world-class identity management with poor segmentation or poor web access protections does not produce good security.

Platform vs. specialized tools

As with many technology decisions, buying Zero Trust software often comes down to the choice between individual specialists and a platform that combines multiple components. The impact of Zero Trust across security teams, integrators, architects, and analysts and their need to maintain policy across multiple consoles, different agents, and multiple integrations offers a compelling case to go the platform route. This is particularly true in a tight labor market with a shortage of skilled cybersecurity professionals. Managing solutions from multiple vendors can increase personnel costs significantly as solutions that do not effectively communicate with one another create false positives, which burden end users and can require additional support and training.

Additionally, the proverbial "one hand to shake" when it comes to support and contract negotiations provides a compelling case for implementing Zero Trust with a platform provider.

Ideally you should look for a single provider with a flexible approach — one that offers a comprehensive platform for Zero Trust alongside individual point products. This flexibility makes it easier to achieve Zero Trust while enjoying the benefits of a single provider.

> **One reason organizations may be reluctant to start with microsegmentation is the perception of complexity.**

Revisiting the elements of Zero Trust

Know who your users are.
Ensure they're verified.

Protect your assets.
Authenticate/authorize all transactions.

Protect your users.
Prevent malware from infecting users.

Conclusion

Ultimately, most organizations concerned with protecting against cyberattacks recognize the need to begin moving to a Zero Trust architecture sooner rather than later. Many have already begun their journey, either gradually or more suddenly, as a response to the rise in remote work. Yet, as attackers become more sophisticated, threat surfaces spread, and more and more constituents demand remote access, the need for a comprehensive portfolio of solutions that work together only grows.

For details on specific elements of Akamai's approach to Zero Trust, visit akamai.com/zerotrust or talk to one of our experts.

About Akamai Security

Akamai Security protects the applications that drive your business at every point of interaction, without compromising performance or customer experience. By leveraging the scale of our global platform and its visibility to threats, we partner with you to prevent, detect, and mitigate threats, so you can build brand trust and deliver on your vision. 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 10/24.