Akamai

WHITE PAPER

Network Segmentation and Microsegmentation in Modern Enterprise Environments

Overview

The idea of segmentation for security is nothing new. Perimeter firewalls, along with VLANs and ACLs, are what most companies have traditionally used to segment and protect their IT infrastructure. However, times are changing. The increase in containerization, software-defined networking, the use of public and multicloud infrastructure, and the expansion of internet-connected devices have created a new set of security issues to address — one that needs a solution built for a heterogeneous IT environment with varying sets of security requirements. Plus, ransomware and nation-state-affiliated threat actors are now a risk to any business, and attackers are becoming more sophisticated at the same time that gaining visibility into your IT environment is becoming harder to grasp. Traditional perimeter security measures, as well as next-gen firewalls based on deep packet inspection or signature-based detection, struggle to keep up with the amount of traffic an enterprise data center experiences today. Let's look at how the right microsegmentation techniques are the best technology to address the shortcomings of other alternative network segmentation approaches.

> As hybrid cloud environments have become the norm, they demand a specific set of requirements above and beyond traditional perimeter security

Legacy firewalls are inadequate for east-west traffic

When looking to segment IT environments, an enterprise might first look to legacy perimeter security devices. Unfortunately, these devices were built to monitor traffic that moves from north to south, from client to server. This includes any traffic that comes to the data center from any external source. More recently, the amount of traffic within the data center moving from server to server, usually referred to as east-west traffic, has increased exponentially. This is in large part due to the growth of virtualization and converged infrastructure such as hypervisor, VPC, and container-based computing.

Perimeter security measures like traditional firewalls do nothing to protect your business from infected devices or to prevent attackers as they expand their foothold using east-west traffic. With the rise of TLS encryption and the easy hiding of malicious traffic piggybacking across open legitimate application ports, many attacks can get through even when traversing the firewalls. This leaves you unable to spot existing breaches and resolve or divert them. It also means that you cannot easily limit the dwell time that attackers have on your network. The longer the dwell time, the more catastrophic the breach. The Active Adversary Playbook 2022 from Sophos found that, while the average median dwell time was 15 days, small businesses and specific industries saw much longer average dwell times — up to 34 days.[1] The more time an attacker can go undetected in your network, the more damage they can do.

It is simply not possible to use enough virtualized firewalls to protect thousands of applications or workloads. Even if a virtualized solution could be created, it would be impossible to manage or control considering the ever-changing dynamic environments in which we now work. When it comes to hybrid cloud, for example, traditional firewalls are even more difficult to use, as they need to work across various environments, track workloads across different clouds, and be controlled from a single point. To try and solve these issues, several network segmentation approaches have appeared.

Three segmentation approaches to consider

With the understanding that firewalls, even when virtualized, are inadequate for protecting hybrid cloud data centers, enterprises look to apply segmentation within the east-west infrastructure in three basic ways. As we discussed, without strong segmentation policy and security measures, any port or server has access to communicate with any other. This means that if a server firewall gets breached, the attacker can move easily to any number of others in the network. The most effective way to limit connectivity between servers is by segmenting the network. There are three basic types of network segmentation, with microsegmentation being the technology that businesses can use to enforce increasingly granular policy and control. Users can combine the three kinds of segmentation policy listed below, building more granular policies for critical or risky applications.

Environment segmentation

This approach separates different environments from one another. In this way, businesses could segment the development arm of their company from the production environment, for example. This is the first, crucial stage in any segmentation strategy, which can then be followed by more granular policy creation.

Application segmentation

Taking segmentation further, "ringfencing" high-value applications takes each specific critical application and keeps it separate from the rest of the network. The best microsegmentation solutions will even allow this to be controlled at a process level.

Tier segmentation

The tightest form of segmentation is within the application itself. Here you could create policy for how communications are managed between tiers within the same application cluster, controlling traffic between web servers, application servers, and database servers, for example. This can also be controlled with process-level enforcement, if you choose.

Network segmentation method — network segmentation through VLANs

Most companies start by employing VLANs. These virtual local area networks allow businesses to allocate each segment with its own communication path, through either a firewall or access control lists (ACLs) on the router itself. While VLAN is a common choice for network segmentation, there are a lot of problems under the surface. Let's look deeper, taking stock of why VLANs are a subpar choice to meet today's security needs.

It's easy to see why many businesses choose VLANs as their segmentation method. It can be done with existing architecture, which makes it feel low cost and simple to deploy. However, it is a very rigid and complex segmentation approach, can be expensive to maintain, and requires downtime to implement.

In order to start using VLANs, you will need to familiarize yourself with the servers and dependencies in each segment, and then create the configuration you want for the network switch or switches you are segmenting. As this is completed by network engineers, and often involves multiple locations, this can take many days and cost a disproportionate amount of both time and money. Traffic may be interrupted or slow during the configuration time.

In an age when agility is a major competitive advantage and perhaps even a must-have, high costs and slow speed when it comes to change spells disaster for your bottom line. According to Forbes, adaptability is key to survival: "Disruption isn't new, but the speed, complexity and global nature of disruption is at a scale we've never seen before. ... It is not the biggest or more financially stable that will survive but the ones that manage to adapt to the exponentially accelerating pace of change."[2]

It's important to recognize that VLANs were not created with segmentation in mind. Initially built for reducing congestion, using them to control communications is not a smart way of leveraging this existing technology — it's in many ways a misuse. Considering this, it's not surprising that segmenting using VLANs comes with limitations.

- **Cloud technology** — VLANs and other traditional network segmentation policies cannot be extended to the cloud. If you use internal segmented firewalls (ISFW) or ACLs to control which users can access network segments, you will likely need to rely on SDN (software-defined networking) for the cloud. This is usually done through third-party software providers who use virtualized firewalls or subnets.
- **Containers** — Security remains a huge concern given the widespread adoption of containers in IT environments. As each container is run on the same kernel, an exploit could put all containers at risk. Isolation has been an ongoing struggle, and cannot be solved with the usual network segmentation methods.
- **Protocol restrictions** — The limit for VLANs is 4,096 segments, which constrains the ability to provide adequate segmentation in large data centers. More granular segmentation approaches do not have this limitation.

Network segmentation to application segmentation — introducing Layer 4 controls

Many of these issues have been improved by embracing application segmentation using security groups within cloud environments and hypervisor-based firewalls for on-premises virtualized environments. Traditional application segmentation implements Layer 4 controls, allowing you to isolate service tiers from one another, so that an application has a secure boundary. Each tier is limited to the level of access it needs to provide its full functionality, but no more. There is a clear separation between the tiers of an individual application, and the threat of a potential compromise is kept to a minimum.

Think about the tiers you might find in a standard business, from load balancers and databases to application servers inside of/outside of your own DMZ. Keeping these tiers separate allows for each to have its own security rules and capabilities. Application segmentation can support businesses in allowing the right controls for each tier, limiting their sensitive information and communications, while allowing broad user access where necessary. For example, a business can keep certain databases from communicating with the internet altogether, or ensure that if an attacker breaches a simple load balancer, they can't pivot to access more sensitive information on the database tier.

As the solution becomes more granular, application segmentation allows a business to segment a whole application cluster from other areas of the business. As discussed, this reduces the attack surface area and the ability of attackers to make lateral movements from one tier to another.

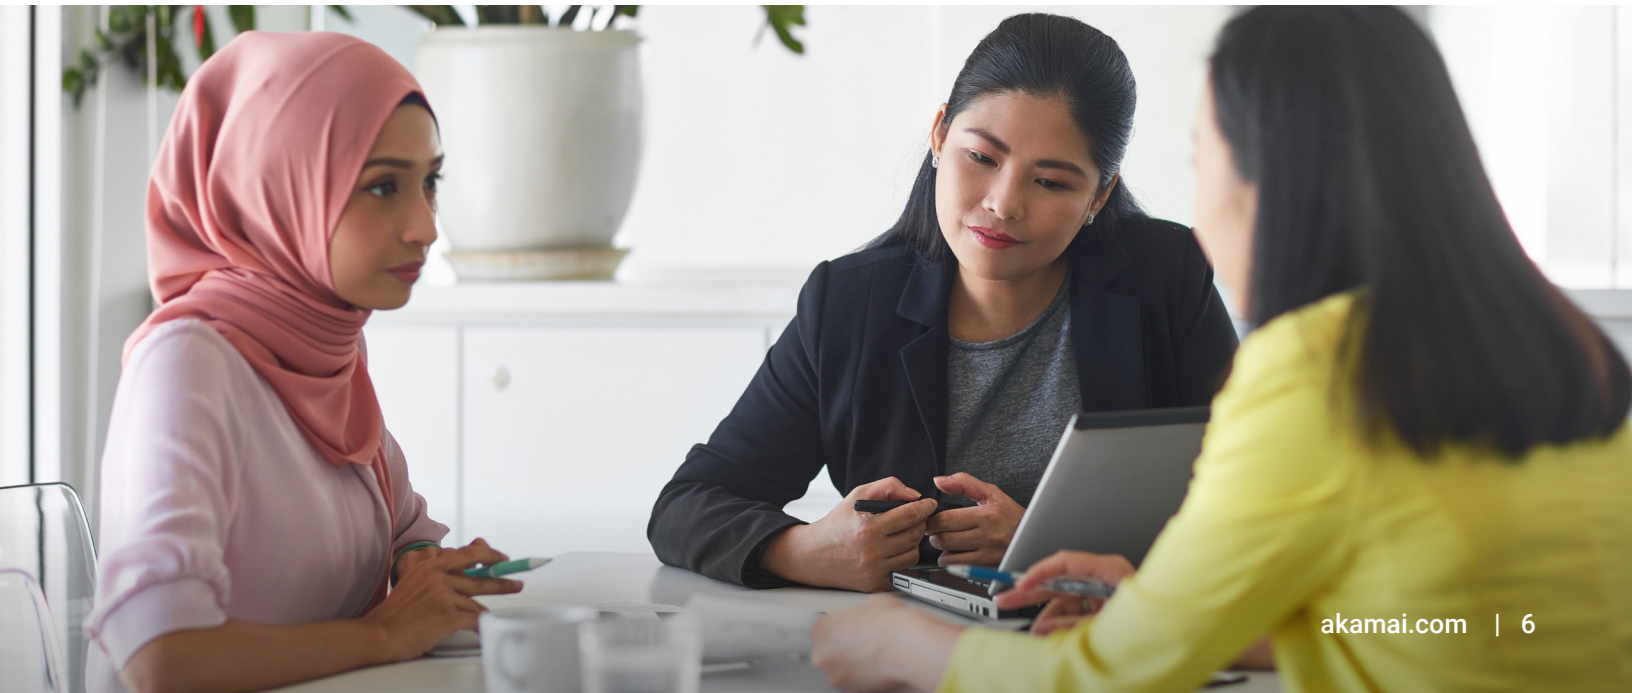

The limits of Layer 4 controls

Traditional application segmentation can lack depth, which has a direct impact on your visibility. The network layer, where routing happens, moves data among systems, assigning IP addresses and protocols that detail the path data segments take to their destination. Application segmentation often uses Layer 4 network controls, focusing on the way the data itself is delivered. Larger data segments are divided into smaller segments or blocks, ready to be put back together at their destination. Flow control allows this process to be sped up or slowed down dynamically, where the devices sending or receiving the information have need.

In today's threat landscape, controls to these layers are essential, but in certain cases you might want the ability to set policy at an even more granular level. Attackers have shown their ability to spoof IP addresses, and use piggybacking techniques on allowed ports to breach a network. Additionally, Layer 4 protection does not limit lateral movement within an application or a tier, which could still leave you with a larger attack surface than you would like.

One of the best examples of the need for more granular controls than just Layer 4 is in compliance initiatives. Traditional application segmentation techniques have to some extent allowed companies to satisfy some specific compliance regulations, such as keeping CDE separate for PCI-DSS, or protecting PHI for HIPAA. However, while Layer 4 techniques have in the past been accepted as effective means of showing compliance, the reality has shown that it may not do enough. According to the Verizon 2022 Payment Security Report, only 43% of companies are "fully compliant."[3] Worse still, even 100% compliance doesn't mean you're 100% secure. Although Layer 4 controls may cover you in terms of compliance, they don't reduce the attack surface enough to make a meaningful difference for security. Period. Attackers can ride an open Layer 4 port between two tiers with a separate process (Layer 7) and take everything they want.

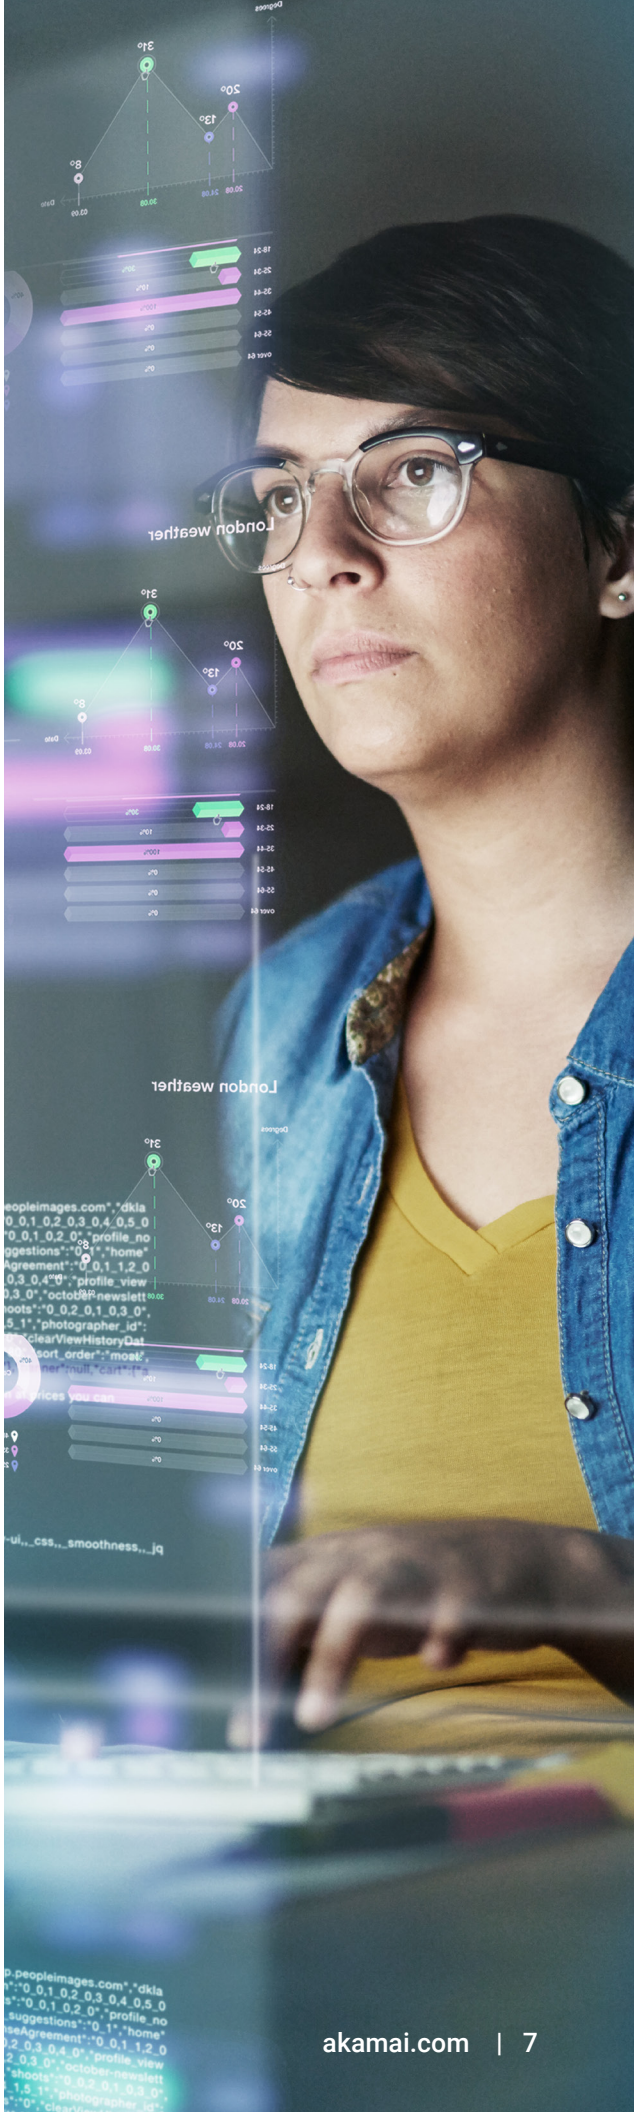

Segmenting in the dark — the lack of visibility in network and application segmentation

As enterprises are finding, while there is no doubt that application segmentation is a step in the right direction, it doesn't go far enough to solve all the issues inherent in a coarse segmentation approach. Another challenge that still needs addressing is visibility. Being able to see an accurate, real-time overview of your network is essential at each stage of your segmentation process, which is a limitation of many segmentation approaches.

Before you start, you'll want to visualize the application dependencies so that you can draw up accurate policy rules. After segmentation has been put into place, you'll need proof that your segmentation is working as intended, not only to confirm that your security posture is strong, but also to provide evidence of regulatory compliance where necessary.

Without real-time and historical visibility, there is no evidence for yourself or third-party stakeholders and regulatory bodies. Manually collecting this evidence is time-consuming and expensive to manage, and there's always the possibility of configuration errors and mistakes. A segmentation solution that can't provide this kind of visibility is simply not enough.

Microsegmentation up to Layer 7 — the application layer

In contrast, segmenting at the application layer (Layer 7) is highly effective at limiting lateral movement, even within an application cluster. Layer 7 is where network services integrate with the operating system. Protocols such as HTTP, FTP, TFTP, and SMTP are all Layer 7 protocols. The latest advances in microsegmentation technology are able to segment at this layer with far more depth than other solutions, enabling your business to visualize and control activity at Layer 7 as well as on the traditional Layer 4. This means that instead of relying on IP addresses and ports, specific processes and flows can be used when businesses configure their policies. This takes the benefits of segmentation far beyond a specific tier or even application cluster. It also allows you to spot potential threats with something as little as the wrong hash, even where the attacker is mirroring an authorized process or pathway.

When it comes to policy creation, segmenting to Layer 7 allows for very specific allowlisting rules or exceptions, where only exact processes or flows are allowed and all other communication is blocked by default. This can enforce the isolation of data between systems, but still allow for communication for necessary or business-critical data flows.

The best microsegmentation solutions provide the visibility businesses need to gain agility

With agents on every workload — hypervisor- or VPC-based, containers, bare-metal servers, or even IoT/OT systems — a holistic microsegmentation solution can provide your business with a full visual map of your entire IT infrastructure. With the truly intelligent solutions, this includes data center, cloud, multicloud, and hybrid cloud environments, and remote devices. Traditional application segmentation solutions struggle to get this all-in-one view, usually because they use a combination of network-centric technologies.

A comprehensive visual map of your environment should also show you which security policies are in place and being enforced in real time. At a glance, your engineers and security professionals should be able to see potential gaps to fix in your policy coverage, or what additional policies they need to implement or create from scratch.

Having this visibility also allows your business to prepare in advance for new software or updates to existing systems, by creating the rules for segmenting updated or new applications ahead of time, before they are ready for deployment. Once the updates are live, your security teams have the real-time information they need to detect and resolve application activity that is outside the norm, ensuring that no security risks go unnoticed or become active exploits. After the fact, your business has the contextual tools to compare an incident to historical data and understand the exact environment that allowed the anomaly to occur. Policies can be tightened, segmentation can be adapted, and you can detail the incident for compliance regulations or further study.

Employing the Zero Trust model

Another added benefit of microsegmentation is its ability to embrace the Zero Trust security model. Although the idea of Zero Trust was coined by Forrester all the way back in 2010, technologies such as microsegmentation are helping to make the concept a reality, and researchers and security experts continue to shout its benefits far and wide.[4]

The idea is simple: No traffic or user is trusted until proven so and approved, whether it's coming from an external source or an internal one, every time there's a connection attempt. Forrester's three main principles of Zero Trust[5] are all supported with strong, granular microsegmentation policies:

- All entities are untrusted by default
- Comprehensive security monitoring is implemented
- Least-privilege access is enforced

Zero Trust is on the opposite end of the spectrum from perimeter-only security, where you protect the entrances to your castle with a deep moat and assume that anything inside is cleared for entry. As most companies no longer have a contained network or data center, the idea of a "castle" is obsolete, and a least-privilege strategy like Zero Trust is the only way to ensure you can know and control who is inside at any given time.

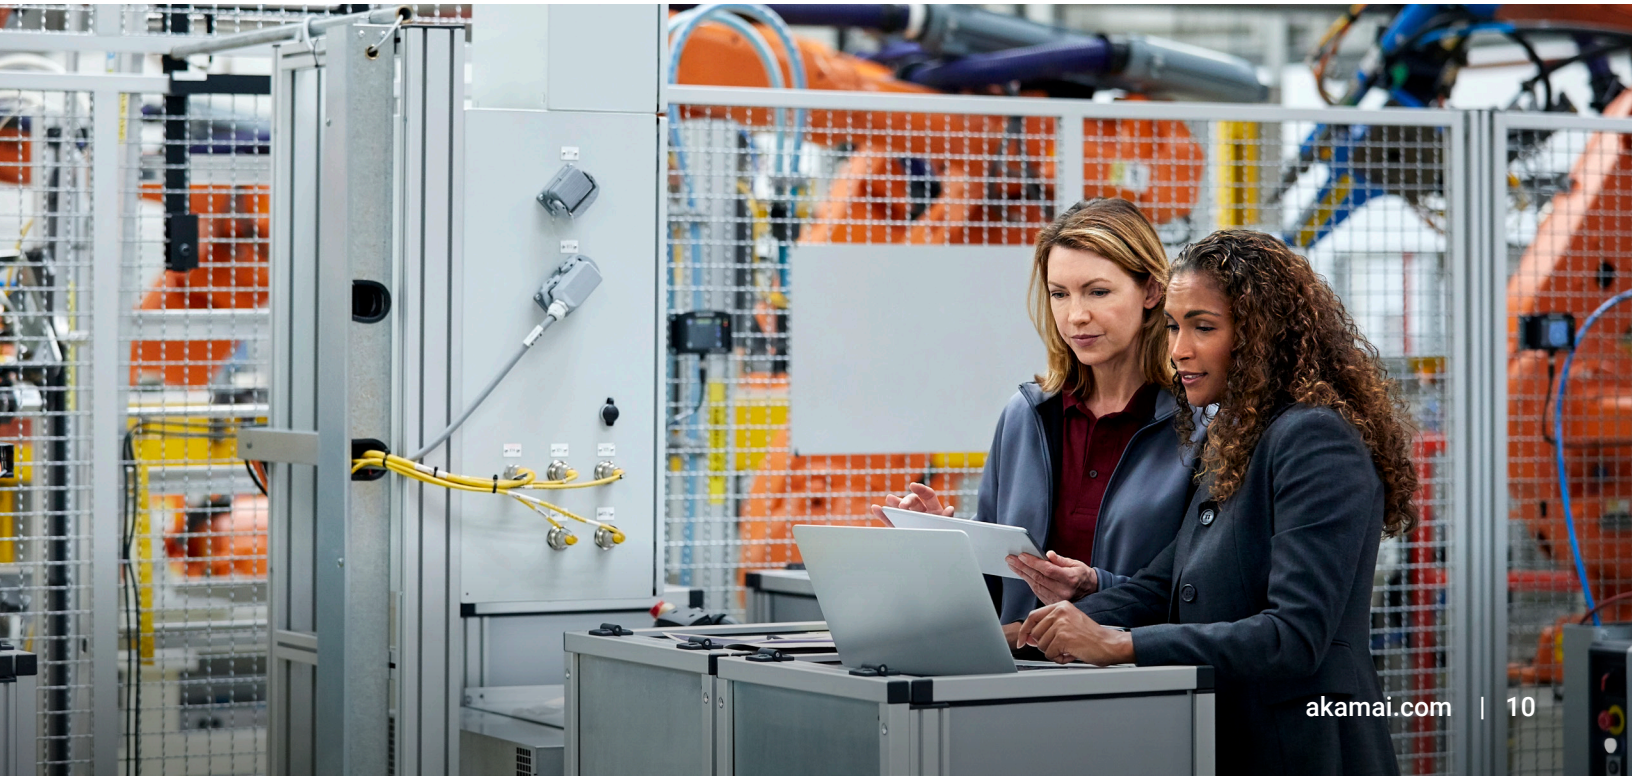

Future-proof your business with microsegmentation

Network segmentation can certainly go beyond perimeter security, and environment segmentation and application segmentation up to Layer 4 are important steps in building your segmentation strategy. But as IT environments become increasingly complex, you might find you need a segmentation solution that offers even more granularity with tier segmentation, and process-level enforcement to Layer 7 at the application and tier stages.

Modern businesses have moved beyond a self-contained infrastructure. They often rely on technology such as SDN in the cloud, containers, or bare-metal hypervisors. They work across various geographies and physical data centers.

The only way to protect yourself from external and internal threats is to employ a solution that inspects and controls all traffic, east-west as well as north-south, and — for crucial or risky applications — gives you more visibility than can be gleaned from Layer 4 alone. Microsegmentation up to Layer 7 at either the application or tier level gives you the ability to get an accurate view of your entire IT environment, and allows you to easily create and enforce granular security policies that follow the Zero Trust model. A good microsegmentation solution won't ask you to choose between security and agility, so make the choice that gives you the strongest overall security posture across your organization.

Please visit akamai.com/guardicore for more information.

1 Shier, John. 2022. "The Active Adversary Playbook 2022." Sophos. June 7.

2 Gonda, Rob. 2018. "Adaptability Is Key To Survival In The Age Of Digital Darwinism." Forbes. May 24.

3 https://www.verizon.com/business/reports/payment-security-report/

4 Holmes, David. June 2022. "Best Practices For Zero Trust Microsegmentation." Forrester. April.

5 Holmes, David and Jess Burn. Jan 2022. "The Definition Of Modern Zero Trust." Forrester. April.