Software-Based Segmentation

An inside-out approach to achieving security confidence

TABLE OF CONTENTS

Move on from yesterday's legacy firewalls

We get it. You're tired of your old, on-premises firewalls. IT environments and security requirements are now light-years ahead of what they were originally built for. And the cybersecurity landscape has evolved, too — attack methods have grown more sophisticated, and there are more cybercriminals than ever. A decades-old appliance architecture simply can't stand up to the latest malware, botnet attacks, phishing schemes, social engineering, and data extortion.

But even with their myriad problems — they are expensive, immobile, and lack visibility, to name just a few — the reality is that legacy firewalls are not going away anytime soon. They serve an important function at the perimeter handling north-south traffic, and provide a hard shell around the organization.

But firewalls cannot manage east-west traffic in on-premises data centers and in the cloud.

This is a job for software-based segmentation.

Did you know?

By 2031, ransomware is expected to attack a business, consumer, or device every 2 seconds.[1]

Solved!

3 problems with legacy firewalls

1. The problem: Lack of visibility

   The lack of visibility into the flow of data makes implementing and maintaining rules difficult. Because of this, firewalls often have extremely long rulesets, and they have a lot of rules that are overly permissive or not even necessary.

   ## The solution

   Look for solutions that integrate a visual map, asset classification, and application dependency mapping with policy creation and management.

Solved!

3 problems with legacy firewalls

2. The problem: Firewalls are hard to maintain

Application owners and firewall admins rarely know the appropriate IP ports and protocols that need to communicate. So managing firewalls becomes an iterative, troubleshooting process.

The solution

Instead of framing policies around the fixed network “plumbing” like IPs and ports, base them on meaningful attributes like the process an application uses, fully qualified domain names (FQDN), and user identity. This way, the attributes remain the same and your policies will keep working, even if you make a change to your data center or move your workload to the cloud.

Solved!

3 problems with legacy firewalls

3. The problem: Firewalls lack agility

Any changes you make to a firewall usually require scheduled downtime. When an application owner needs to make a change, they may wait a week or more for the change to be reviewed and implemented during a maintenance window.

The solution

Modern IT organizations have moved away from change windows to DevOps models in which applications are appearing and updating continuously. Find a technology solution that can be automated using the same DevOps tools that you're using for the applications themselves. This way, as applications continuously evolve, the security approach adapts along with it.

You can take it with you

Let’s talk about the traditional way of doing things. It is complicated. And it is not adaptable. The old-school approach to managing legacy firewalls bases segmentation on location — and that location cannot be easily changed. It is usually based on a hard-coded IP address or routed to a data center. This means you physically have to move whatever it is you want to secure behind the firewall, a process that is resource-intensive, risk averse, and slow. Cloud migration? Visibility? Adequate security? Forget about it.

Leave your legacy firewalls where they are. Take a deep breath and embrace something new. Software-based segmentation can be easily implemented alongside your existing firewalls, and it is adaptable. With software-based segmentation, you can actually make changes to your environment, data center, and network, and set policies based on what you see. And the workload and policies can show up anywhere — in the cloud, data center, wherever. Plus, you can apply and adapt your security policy without making changes to the network and with zero system downtime.

Reveal your internal segments

Would you trust something you cannot actually see? We didn’t think so. But this is precisely what you’re doing when it comes to establishing security policies behind a firewall. You cannot actually see what’s inside. It’s like looking at the building without being able to see the people inside.

Software-based segmentation is not based on chance. It breaks up the pieces so you’re totally aware of all activity that your workloads are involved in. Once you know what is inside your environment, you can form a plan and break up the segments into something meaningful and effective based on your specific use cases.

Security beyond the perimeter

Legacy firewalls simply were not built for change. While they serve an important purpose at the perimeter, like DDoS protection and traffic filtering and inspection, security inside the network is hard to pull off with firewalls. Why? They were deployed as natural choke points, which means every segmentation effort comes with operational roadblocks, like the need to change and remove networks and applications. This is tedious and resource-intensive.

Software-based segmentation can help you overcome these operational challenges and allow you to continue your security practices beyond endpoints and perimeters. First, it features a distributed firewall approach (versus a choke point). Second, it is workload-centric, which means it can collect data from the host system and then apply it to asset classification, and can take a more granular approach to rules, like process-level content and policies. Overall, software-based segmentation is a more adaptable, granular way to protect critical assets inside your network, and requires less effort and resources than firewalls.

4 segmentation basics

Segmentation is more important than ever before. Attack surfaces are bigger. Sophisticated attacks, like ransomware, move laterally after a breach, and you have to think about application dependencies beyond the perimeter. But segmentation is not a one-and-done approach.

Here is a look at four common types of segmentation, how they're different, and why you need them.

1. Environment segmentation

separates systems in different development environments, such as Development, QA, Staging, and Production. This is a broad version of segmentation where the end goal is to separate systems in different environments to ensure access is limited to only the necessary users and applications. A lot of compliance initiatives require the assurance that nonproduction systems cannot access production systems.

2. Network segmentation

is an architecture practice of splitting a network into multiple subnetworks, each being its own smaller network segment. Network segmentation gives IT operators a tool to better control network traffic, boost performance, and improve security.

3. Microsegmentation

is a more granular form of segmentation that's used to isolate workloads from one another and secure them individually. This includes the ability to set segmentation rules for elements such as processes, containers, users, domain names, and devices. This approach is superior at controlling east-west traffic and protecting against lateral movement.

4. Identity-based segmentation

expands beyond microsegmentation's ability to protect a single endpoint, device, workload, or container by enabling dynamic rules that assess identity — which can be the user, device, or context — as part of determining whether to allow communication. Identity-based segmentation policies can be based on granular settings — not just IP or port — such as tags, OS type, or application characteristics.

Myth vs. reality: 5 segmentation myths debunked

Myth 1

Segmentation projects are too difficult and take too long to complete.

Reality: Starting with visibility and a clear understanding of what is happening inside your environment reduces the time to complete a segmentation project from months to weeks or even days. Modern segmentation technologies can also use AI to accelerate the process even further.

Myth 2

Segmentation projects require network infrastructure changes and downtime.

Reality: Software-based segmentation decouples security from infrastructure, so segmentation can be performed independently from the underlying infrastructure without changes or downtime.

Myth 3

Segmentation blocks legitimate traffic in my network.

Reality: Visualizing your environment and using software-based segmentation policies makes it possible to see the effect that these policies will have on your business activities before real-time enforcement is activated.

Myth 4

Segmentation inhibits user access and introduces unnecessary latency.

Reality: Using distributed, software-based segmentation policies instead of forcing all traffic through specific firewall choke points eliminates network bottlenecks. And more precise policies that are application- and identity-aware reduce the risk of inadvertent user-access issues.

Myth 5

I can’t use the same segmentation tools in the cloud as I use on-premises.

Reality: If you decouple segmentation policies from infrastructure, the same policies used in the data center can also work in the cloud.

Reduce risk on the inside

Breaches will happen. And they can cripple your business, compromise your data, damage your brand, and cost you millions.

Still think firewalls can do it all? Think again. Once an attacker has breached a network, environment, or data center, it will use lateral movement to steal data and wreak havoc, like taking control of application servers or accessing database servers.

In fact, 70% of all attacks now involve attempts at lateral movement.[2]

While firewalls see lateral movement as legitimate traffic happening within a network, software-based segmentation stops it dead in its tracks. A critical component to your security program, software-based segmentation allows you to restrict lateral movement — and in the event of a breach, make it harder for an attacker to navigate the environment. You get a fighting chance at protecting data and critical applications, decreasing dwell time and even detecting the attacker. This approach is more scalable, easy to use and allows you to quickly implement segmentation without making changes to your network or systems.

Companies spent an average of $2.4 million in 2020 defending against an onslaught of malware and web-based attacks.[3]

Zero Trust doesn’t have to be complicated

Zero Trust is all about who does what to whom, and how they do it. In other words, having explicit control over who does what inside your network.

By giving a user access to anything inside the network, you’re automatically granting too much trust, and as a result, putting your entire organization at risk. First, employees often make mistakes, which could have serious security implications. Some even have malicious intent.

Plus, outside of VPN networks and devices, there are a lot of entry points to the data center you should consider. For example, attackers can get inside a network through the production server (like in the case of the SolarWinds breach), an internet-facing application that’s vulnerable, or a vulnerable VPN. In this case, you trust a server just because it’s within the network, but in practice, the attacker can access anything and move laterally without constraint.

To achieve Zero Trust in your production network, you have to block all activity that is not explicitly allowed.

This is something that legacy firewalls simply cannot do at a granular level, because it requires identifying attributes at a level deeper than IP addresses and ports.

Alternatively, software-based segmentation allows you to actually see what is happening in detail and create precise, human-understandable policies that include identity.

Your Zero Trust checklist: 6 ways to gain explicit control

Let's keep it simple. Trust should be based on the size of the segment — and the smaller the segment, the better, when it comes to protecting critical data, assets, and applications. Here are six steps to achieving Zero Trust without the operational complexity.

1 Identify your sensitive data using visualization labels.

2 Map the flows of your sensitive data using automated flow and dependency mapping.

3 Architect your Zero Trust microperimeters using the right tools for the rapid definition of any segmentation or microsegmentation policy.

4 Continuously monitor your Zero Trust ecosystem through real-time monitoring and analysis.

5 Embrace security automation and orchestration, and APIs and technology integrations.

6 Have capabilities in place to untrust someone or something — so that if you're under attack, you can easily untrust any machine with preset attributes, regardless of the user or segment.

The bottom line

By now, you are probably wondering how you can break up with your old-school solutions to strengthen your security posture inside your network.

No problem.

Leave your legacy firewalls where they are — they are good at protecting the network perimeter. But the benefits really stop there.

What matters most lives at the core of your organization — the digital assets, data, and applications that exist beyond the perimeter — the guts of your corporate infrastructure. Shifting your focus from the from the outside in and implementing software-based segmentation and a Zero Trust framework, will give you the visibility and control you need to detect and stop lateral movement, apply granular and adaptable policies, and stop cyberattacks like ransomware from propagating through your network.

Request a demo or **learn more** about how segmentation can help with ransomware, Zero Trust, cloud security, and more.

[1] Cybersecurity Ventures. 2022 Who's Who In Ransomware Report. Conceal, 2022.

[2] Kellerman, Tom, and Greg Foss. Global Incident Response Threat Report. VMware Carbon Black, Oct. 2020.

[3] "2023 Cyber Security Statistics Trends & Data." PurpleSec, 22 Feb. 2023.

Akamai protects your customer experience, workforce, systems, and data by helping to embed security into everything you create — anywhere you build it and everywhere you deliver it. Our platform's visibility into global threats helps us adapt and evolve your security posture — to enable Zero Trust, stop ransomware, secure apps and APIs, or fight off DDoS attacks — giving you the confidence to continually innovate, expand, and transform what's possible. Learn more about Akamai's cloud computing, security, and content delivery solutions at akamai.com and akamai.com/blog, or follow Akamai Technologies on Twitter and LinkedIn. Published 06/23