Akamai

WHITE PAPER

Clearing the Path to Microsegmentation

A strategy guide for implementing microsegmentation in hybrid clouds

More clouds in the forecast

The migration of vast amounts of data and data processing to the cloud — or more precisely, to multiple clouds — is arguably the biggest change in enterprise computing in the past decade. More organizations are moving to public clouds, and typically, to public–private hybrid data center architectures. At the same time, they are leveraging infrastructure as a service (IaaS) in the quest for ever-greater agility. Technology analyst Gartner projects that by 2025, a little more than half of all IT spending in addressable market segments will have shifted from traditional solutions to the public cloud, compared to 41% in 2022, and total revenue spend on public cloud is expected to exceed $900 billion by 2025.[1]

The distinction between "the cloud" and "multiple clouds" is not trivial. Increasingly, enterprises are adopting multicloud platforms and service providers. One thing is clear: The idea of an enterprise data center as a single, secure physical space is headed the way of the dinosaur. Modern data centers are increasingly a heterogeneous mix of environments and technologies that combine physical servers, virtual machines, and containers in on-premises facilities, private clouds, and public cloud IaaS providers. And these disparate installations are not static — organizations are constantly shifting data and workloads among their various on-premises and cloud environments as traffic levels and processing demands dictate.

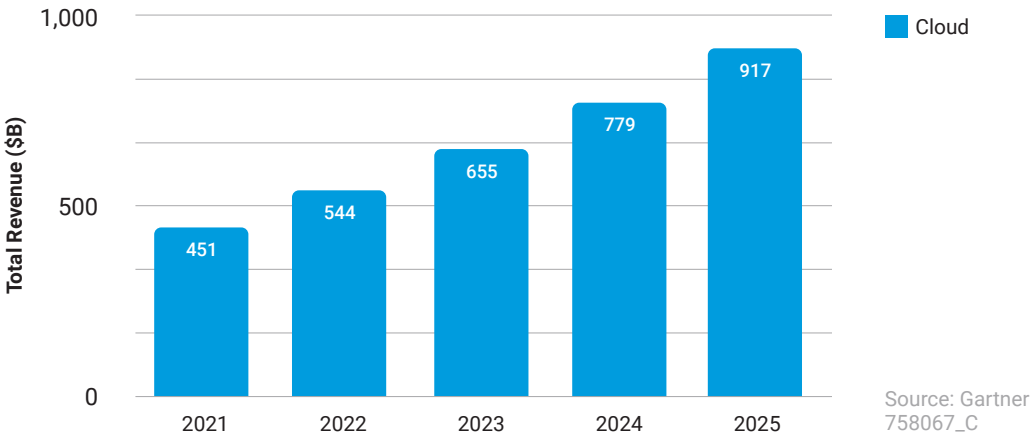

Worldwide public cloud services revenue forecast (in billions)

Increased complexity triggers new vulnerabilities and broadens attack surfaces

Cloud customers certainly benefit from the added agility, elasticity, and scalability that IaaS affords them — those perks are a big part of what makes the cloud so attractive. The trade-offs to this, however, are vastly increased management complexity, a loss of workload visibility across environments, and, in turn, an uncharted cybersecurity landscape. Working with multiple cloud providers means that security teams have to deal with widely varying security standards and capabilities. Traditional security tools designed for on-premises servers and endpoints simply can't handle the scale and complexity of the cloud. Newer tools provided by IaaS vendors may be effective in the provider's environment, but are of little value in a multi-provider infrastructure.

Moreover, even in this age of virtualization and "software-defined everything," the security mentality (and hence most of the investment) is still grounded in the perceived need to block attacks specifically at the point of entry. This isn't a knock on perimeter defenses — they are still very relevant to the IT security stack — but they don't perform as well when the perimeter is constantly shifting. Data and workloads are moving back and forth among public and private clouds and on-premises data centers, and the users accessing them are increasingly working from remote locations that may or may not have the appropriate security controls in place.

The sheer number of data breaches reported every year is enough to tell us that shrewd attackers are getting through perimeter defenses pretty much at will. And once inside, they find a relatively flat network where assets residing within the perimeter are virtually unguarded. For all the flexibility organizations have gained, the added complexity of managing and securing multicloud infrastructures has exponentially multiplied the attack surface; with little or no communication controls in place, each individual server becomes an attack surface in and of itself. As a result, attackers can spend more time moving laterally — and undetected — between east-west traffic workloads to find your most critical assets.

Network segmentation is a well-understood and established security practice, but nowadays it can be difficult to execute in dynamic IT infrastructures and at cloud scale, where workloads are communicating and often migrating across segments. Enterprise cloud customers have come to the realization they need to further segment their applications and workloads to tightly control communication flows in real time, and detect and thwart threats within the data center before they can do any damage. What's needed is a solution that reduces security complexity by working consistently across infrastructure boundaries to shrink the overall attack surface, enabling security teams to detect more threats more quickly and limit their spread.

That's where microsegmentation comes in.

Microsegmentation defined

Gartner defines microsegmentation as "the process of implementing isolation and segmentation for security purposes within the virtual data center." Further, microsegmentation "reduces the risk of a lateral spread of advanced attacks in enterprise data centers and enables enterprises to enforce consistent segmentation policies across on-premises and cloud-based workloads."[2]

Microsegmentation typically works by establishing security policies around individual or groups of applications, regardless of where they reside in the hybrid data center. These policies dictate which applications and components can and cannot communicate with each other. Thus, any attempt at unauthorized communication is an instant indicator of a threat. In the best case, microsegmentation technologies are infrastructure agnostic, so security policies can continue to protect their respective applications as they move among cloud environments.

Solution Areas for Segmentation

Infrastructure Segmentation

Secure application traffic within a particular infrastructure.

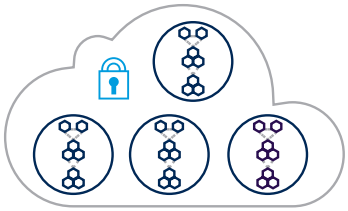

Application Segmentation

Secure traffic between applications and external networks.

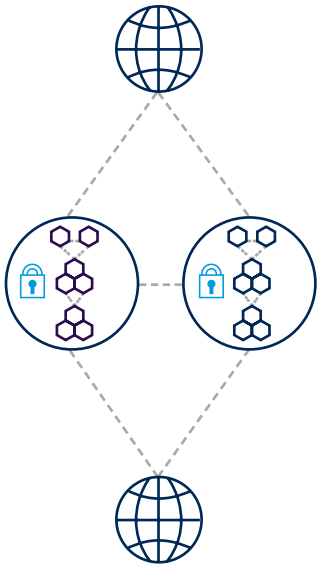

Microsegmentation

Rules that secure traffic within applications with additional context such as process-level attribution.

2 Gartner, "Technology Insight for Microsegmentation," March 2017; "Hype Cycle for Cloud Security 2017," July 2017

The case for microsegmentation

Today's dynamic data centers require enterprises to shift their attention from intrusion prevention and access management to the workloads and applications themselves. And that appears to be happening at an accelerating rate. Even back in 2017, Gartner started noticing a trend toward "increased focus on server workload protection from advanced targeted threats that bypass traditional perimeter and signature-based protection. Typically, these attacks are financially motivated and target server and application workloads as a way to get to sensitive data or transactions."[3]

A key driver of microsegmentation is the need to protect mission-critical applications and workloads. This may seem simply a matter of self-interest or good business, but in many cases, it is also mandated by security policies and regulatory requirements.

Security teams need to find ways to reduce the expanding attack surface within data centers, meaning reducing the vulnerability of servers running applications. Traditional authentication techniques such as signature blocking or application allowlisting are too easily subverted by sophisticated attackers. Microsegmentation enables teams to set and enforce strict, granular access and communication policies. It should also enhance visibility into application flows and enable teams to better assess their security posture.

Do you need microsegmentation?

Answering a few simple questions will help you ascertain your need for microsegmentation.

- Are you in a regulated industry, or do you need to comply with regulations governing the security of data and transactions?
- Do you have a hybrid infrastructure with workloads that span multiple clouds?
- Are you running applications in virtual machines or containers?
- Do you feel a loss of visibility and control of workloads?
- Can you tell, at any given time, that a threat is present or an attack is underway in your data center?
- Can you control security across your infrastructure through a "single pane of glass"?

The four main obstacles on the path

If security experts generally agree on the need for microsegmentation in today's dynamic data centers, why is it considered to be so daunting to implement efficiently and successfully? Organizations attempting to implement microsegmentation using conventional tools generally encounter four major obstacles:

1. **Lack of process-level visibility**
   This is likely the first impediment you will encounter — you can't secure what you can't see. Microsegmentation is about securing individual and groups of applications and workflow processes. Security teams need visibility into actual east-west traffic flows to understand them in context. Most tools do not give you that depth.

2. **Lack of hybrid multicloud support**
   Microsegmentation security policies have to be able to scale easily across on-premises and public cloud environments, and follow workloads as they move back and forth. Tools designed to work in one specific environment are ineffective in hybrid environments.

3. **Inflexible policy engines**
   As noted earlier, today's data centers are not static. Security measures cannot be either — the "set it and forget it" mindset won't cut it anymore. Unfortunately, existing tools from cloud providers don't allow the necessary flexibility to constantly scope, test, and refine rules. This challenge is compounded in hybrid infrastructures that require multiple policy tools.

4. **No integration with complementary controls**
   Done correctly, microsegmentation is not just about protecting processes, but also about catching attacks. However, single-function microsegmentation tools typically don't include breach detection capabilities, leaving it to the user to integrate tools and make them work together effectively. This patchwork approach carries a high risk of failure.

Unsuccessful projects are the norm, not the exception

Given these obstacles, it's not surprising that most microsegmentation projects tend to suffer from glacial implementation cycles, and they can run up costs, tax resources, and ultimately fail to achieve their goals. Organizations frequently stumble over figuring out what needs to be segmented (due to lack of visibility) and deciding how much segmentation is required. They may spend months building out spreadsheets of intricate rules for process-level communications, unable to recognize opportunities for grouping applications and streamlining policies. Too often, they err on the side of "over-segmentation" — setting too many discrete policies, resulting in too much security complexity, which is precisely what you are trying to overcome. As Gartner has noted, "… more than 70% of segmentation projects will have their initial design rearchitected because of over-segmentation."[4]

Over-segmentation runs the risk of slowing down applications, and ultimately, the business. But the pendulum can swing back too far the other way, toward not enough segmentation, and end up compromising your security posture.

Strategy for a successful microsegmentation journey

The path to implementing microsegmentation isn't a straight line — there are many twists and turns as you discover, understand, and control the communication flows in your environment. Security teams need flexibility when developing security policies to constantly incorporate new changes or additions without breaking applications. Many solutions offer inflexible policy creation engines — forcing security teams to implement incomplete or ineffective rules before they are ready.

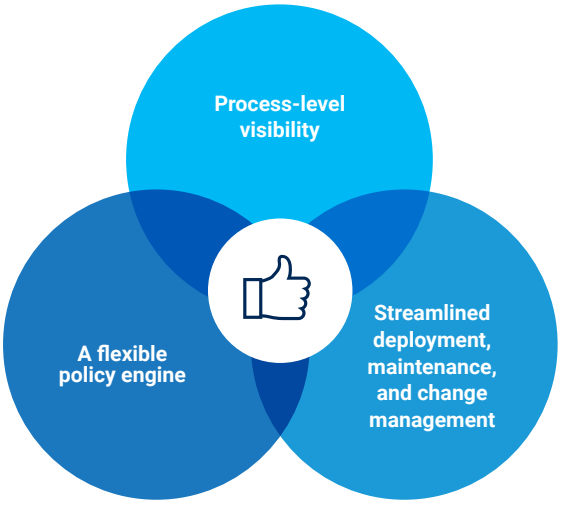

Quite simply, a successful implementation is one that overcomes or sidesteps the four main obstacles, avoiding undue complexity and reducing the risk of under- or over-segmentation by allowing a phased approach. This means having a solution that meets these requirements:

- **Process-level visibility:** Teams need the ability to reveal, collect, and normalize all east-west and north-south flows; tools that enable automatic discovery of applications and an understanding of their communication requirements; and the ability to filter on multiple application attributes to facilitate labeling and grouping of assets that can share policies.
- **A flexible policy engine:** You should be able to simultaneously design high-level best practice and compliance rules for large segments and more granular rules for microsegments. The solution should allow you to move gradually from alerting to enforcement. And it should enable you to establish policies that can work across all platforms, devices, and clouds.
- **Streamlined deployment, maintenance, and change management:** The system should make it easy to deploy, maintain, and modify rules as needed. It should incorporate built-in breach detection and incident response capabilities. Ultimately, your policies should be sufficiently well defined so that you can integrate them into automated deployment (CI/CD) tools for each new application launched.

Ideal solution capabilities

Of course, there are many microsegmentation tools on the market, and not all of them make it easy to follow this path. To ensure a smoother and more successful implementation, make sure you choose a solution with these capabilities:

- **Automatic application discovery**, with complete process-level visibility for bare-metal servers, virtual machines, and containers
- The ability to define **robust and extensive queries** to create contextual labels and groups of objects
- A **flexible policy engine** with intelligent rule design that helps you refine, strengthen, and maintain policies
- An integrated multi-method **breach detection capability** to find more threats more quickly and limit their spread
- **Hybrid infrastructure support** — one platform that works with any infrastructure — data centers, public and private clouds, and more

A solution with these core capabilities will put you on the most successful path to implementing microsegmentation, enable you to overcome the known obstacles and complexities, and prepare you to reap all the business advantages of a flexible hybrid cloud infrastructure without sacrificing security.

Hybrid data centers, multicloud platforms, and IaaS give organizations more flexibility, scalability, and agility than would be possible in a "closed" on-premises data center. But they also leave applications and workloads — the actual assets cyberattackers are targeting — more exposed and vulnerable. While microsegmentation is widely regarded as a best practice in protecting workloads in the cloud, enterprises are having a hard time getting it right. The good news is that you don't have to do it all at once. Today's advanced solutions, coupled with a phased, step-by-step approach, make the path to implementing microsegmentation much easier. And that means better security for your organization's most important assets.

Learn more about implementing microsegmentation successfully at akamai.com/guardicore

1 "Gartner Says More Than Half of Enterprise IT Spending in Key Market Segments Will Shift to the Cloud by 2025." Gartner, February 9, 2022.

2 Heiser, Jay. "Hype Cycle for Cloud Security, 2017." Gartner, July 17, 2017.

3 MacDonald, Neil. "Market Guide for Cloud Workload Protection Platforms." Gartner, March 22, 2017.

4 Young, Greg. "Best Practices in Network Segmentation for Security." Gartner, 28 July 2016.

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 security, compute, and delivery solutions at akamai.com and akamai.com/blog, or follow Akamai Technologies on Twitter and LinkedIn. Published 05/23.