



# Clearing the Path to Micro-Segmentation

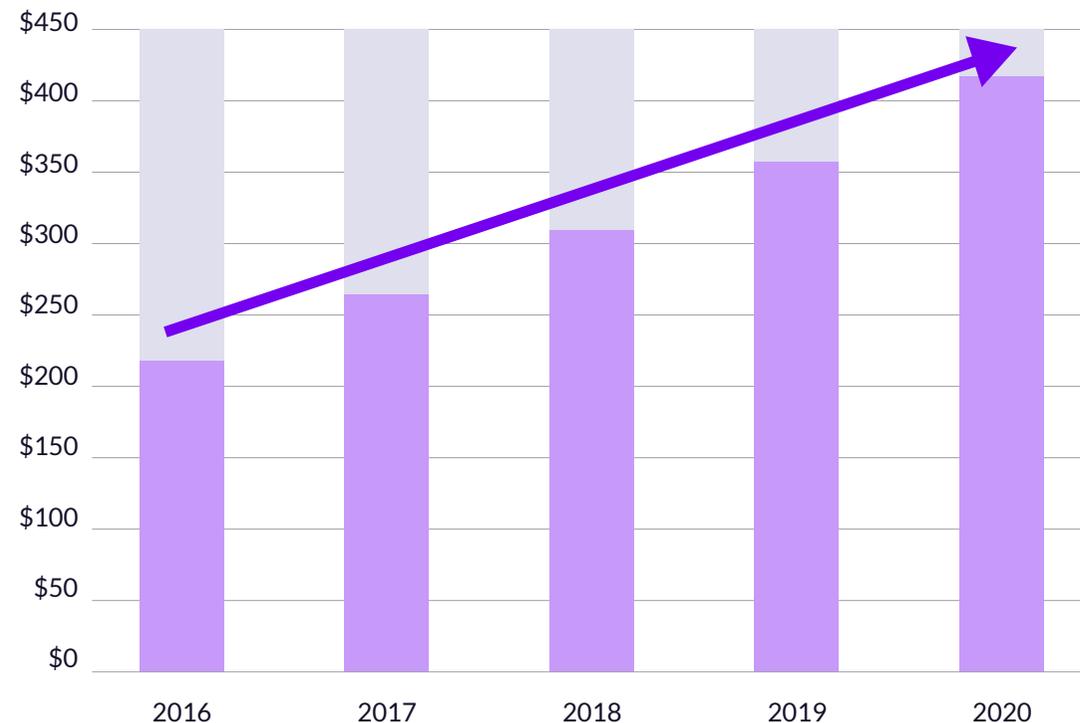
A Strategy Guide for Implementing  
Micro-Segmentation in Hybrid Clouds

# More Clouds in the Forecast

The migration of vast amounts of data and data processing to the cloud — or more precisely, to clouds — is arguably the biggest change in enterprise computing in the past decade. More and more organizations are moving to public clouds and, more 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 the public cloud services market will reach \$246.8 billion by the end of 2017, up 18% from the previous year, and that most of that growth will come from cloud IaaS.<sup>1</sup>

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

## Worldwide Public Cloud Services Revenue Forecast (in Billions)



<sup>1</sup> "Gartner Forecasts Worldwide Public Cloud Services Revenue to Reach \$260 Billion in 2017," October 12, 2017, <http://www.gartner.com/newsroom/id/3616417>



# Increased Complexity Triggers New Vulnerabilities and Broadens Attack Surfaces

Cloud customers certainly benefit from the added agility, elasticity and scalability that IaaS affords them — that is what makes it attractive. The tradeoffs, however, are vastly increased complexity, a loss of workload visibility 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-premise servers and endpoints simply can't handle cloud scale and complexity. 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 at the point of entry. This isn't a knock on perimeter defenses — they are not obsolete by any stretch — but they become problematic when the perimeter is constantly shifting, as data and workloads move back and forth among public and private clouds and on-premise data centers.

The sheer number of data breaches reported every year is enough to tell us that wily attackers are getting through

perimeter defenses pretty much at will. And once inside, they are finding the assets residing within the perimeter virtually unguarded. For all the flexibility organizations have gained, the added complexity of multiple-cloud infrastructures has 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.

Network segmentation is a well-understood and established security practice. The problem is that it is difficult to execute in dynamic 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 in order to detect and thwart threats within the data center in real time before they can do any damage. What's needed is a solution that reduces security complexity by working across infrastructure boundaries to shrink the attack surface, so that security teams detect more threats more quickly and limit their spread.

**That's where micro-segmentation comes in.**

# Micro-Segmentation Defined

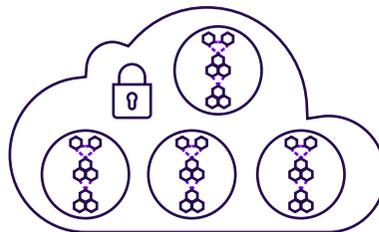
Gartner defines micro-segmentation as “the process of implementing isolation and segmentation for security purposes within the virtual data center.” Further, micro-segmentation “reduces the risk of a lateral spread of advanced attacks in enterprise data centers and enables enterprises to enforce consistent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

Micro-segmentation 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 can and cannot communicate with each other. Thus, any attempt at unauthorized communication is an instant indicator of a threat. In the best case, micro-segmentation technologies are infrastructure agnostic and security policies continue to protect their respective applications as they move among cloud environments.

## Solution Areas for Segmentation

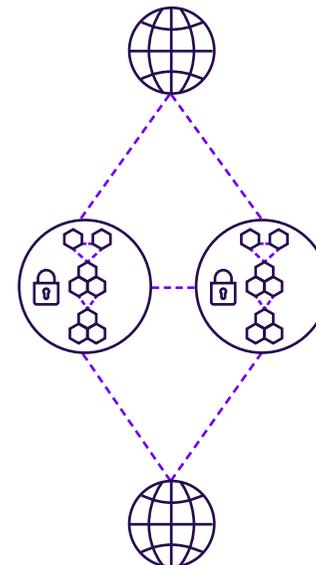
### Infrastructure Segmentation

Secure application traffic within a particular infrastructure.



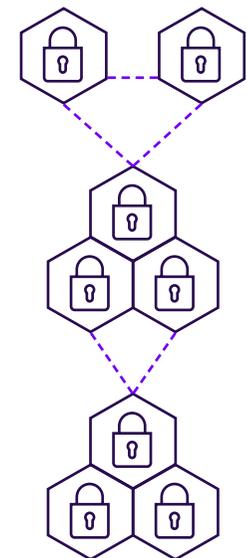
### Application Segmentation

Secure traffic between applications and external networks.



### Micro-Segmentation

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



<sup>2</sup> Gartner, “Technology Insight for Microsegmentation,” March 2017; “Hype Cycle for Cloud Security 2017,” July 2017

## DO YOU NEED MICRO-SEGMENTATION?

Answering a few simple questions will help you ascertain your need for micro-segmentation.

- 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 Case for Micro-Segmentation

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. Gartner recently noted 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.”<sup>3</sup>

A key driver of micro-segmentation 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 whitelisting are too easily subverted by sophisticated attackers. Micro-segmentation enables teams to set and enforce strict access and communication policies. It also restores visibility into application flows and enables teams to better assess their security posture. vulnerability of servers running applications. Traditional authentication techniques such as signature blocking or application whitelisting are too easily subverted by sophisticated attackers. Micro-segmentation enables teams to set and enforce strict access and communication policies. It also restores visibility into application flows and enables teams to better assess their security posture.

<sup>3</sup> Gartner, “Market Guide for Cloud Workload Protection Platforms,” March 2017

# The Four Main Obstacles on the Path

If security experts generally agree on the need for micro-segmentation in today's dynamic data centers, why is it so daunting to implement efficiently and successfully? Organizations attempting to implement micro-segmentation 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. Micro-segmentation is about securing individual and groups of applications and workflow processes. Security teams need visibility into actual east-west traffic flows in order to understand them in context. Most tools do not give you that depth.

3

## Inflexible Policy Engines

As noted earlier, today's data centers are not static. Security measures cannot be either – “set it and forget it” won't cut it. 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.

2

## Lack of Hybrid Multi-Cloud Support

Micro-segmentation security policies have to be able to scale easily across on-premise and public cloud environments and follow workloads as they move back and forth. Tools designed to work in a specific environment are ineffective in hybrid environments.

4

## No Integration with Complementary Controls

Done correctly, micro-segmentation is not just about protecting processes, but also about catching attacks. However, single function micro-segmentation 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

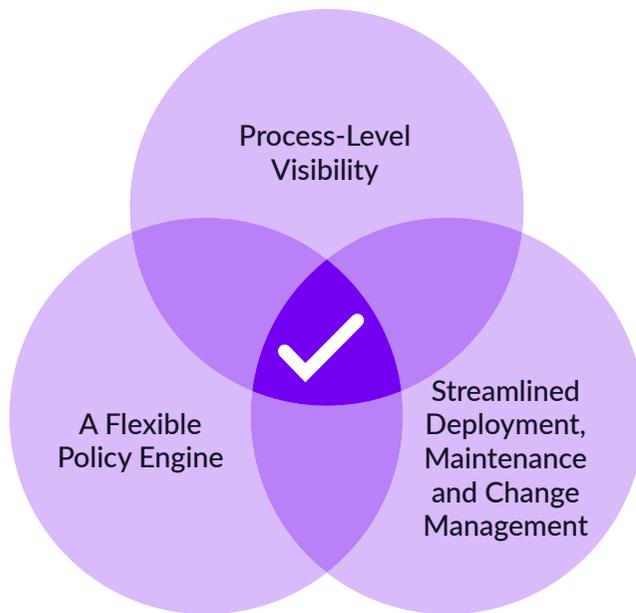
Given these obstacles, it's not surprising that most micro-segmentation projects tend to suffer from glacial implementation cycles, run up costs, tax resources, and ultimately fail to achieve their goals. Organizations frequently stumble on figuring out what needs to be segmented (due to lack of visibility) 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 to group applications and streamline 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."<sup>4</sup>

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 macro-segmentation, and end up compromising your security posture.

## Strategy for a Successful Micro-Segmentation Journey

The path to developing a micro-segmentation policy isn't a straight line, there are many twists and turns as you discover, understand and control application flows. Security teams need the flexibility when developing policies to constantly incorporate learnings as they tighten policies without breaking applications. Many solutions offer in-flexible policy engines – forcing security teams to implement rules before they are ready.

Quite simply, a successful implementation is one that overcomes the four main obstacles, avoids undue complexity, and reduces 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 micro-segments. 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 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 integration breach detection and incident response capabilities. Ultimately, your policies should be sufficiently well defined that you can integrate them into automated deployment (CI/CD) tools for each new application launched.

<sup>4</sup> Gartner, "Best Practices in Network Segmentation for Security," July, 2016



# Micro-Segmentation Solution Requirements

Of course, there are many micro-segmentation tools on the market, and not all of them make it easy to follow this path. To ensure a successful process, make sure you have these capabilities:

- **Automatic application discovery**, with complete process-level visibility for bare-metal servers, VMs 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.

A solution with these core capabilities will put you on the path to successful micro-segmentation, equipped to overcome the obstacles and complexities, and prepared to reap all the business advantages of a flexible hybrid cloud infrastructure.

Hybrid data centers, multi-cloud platforms and Infrastructure-as-a-Service give organizations more flexibility, scalability and agility than would be possible in a “closed” on-premise data center. But they also leave applications and workloads – the actual assets cyber-attackers are targeting – more exposed and vulnerable. While micro-segmentation 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 you don’t have to do it all at once. Using today’s advanced tools and following a phased, step-by-step approach, make the path to micro-segmentation much easier. And that means a more certain degree of security for your organization’s most important assets.



## About Guardicore

Guardicore is an innovator in data center and cloud security that protects your organization's core assets using flexible, quickly deployed, and easy to understand micro-segmentation controls. Our solutions provide a simpler, faster way to guarantee persistent and consistent security - for any application, in any IT environment.

[www.guardicore.com](http://www.guardicore.com)

Copyright 2019