

THE DEFINITIVE GUIDE

to Micro-Segmentation

By John Friedman



About the Author

Jon Friedman is a Managing Consultant at CyberEdge Group, a premier research and marketing consulting firm serving the needs of high-tech vendors and service providers. Jon has more than 20 years experience in industry analysis and marketing, working with more than 40 software, computer, and IT services companies. He has a BA from Yale and an MBA from Harvard.

Table of Contents

- i Foreword
- iii Introduction
- iv Chapters at a Glance

1 Segmentation for Security and Agility

- 2 Two Giant Challenges
- 5 Network Segmentation Has Maxed Out
- 6 Two Principles of Micro-Segmentation
- 8 Segmentation for Smart People

2 Good Fences Make Good Data Centers: Controlling Access

- 10 The $(N \times (N-1)) / 2$ Problem
- 12 The Name of the Game: Isolating Applications and Workloads
- 16 Enforcement Points – Three Approaches
- 19 The Visibility Challenge
- 21 Granular Segmentation: Who Benefits?

3 Stuff Changes: Dynamic Micro-Segmentation

- 23 Why Firewalls Are Inflexible
- 25 Dynamic Micro-Segmentation
- 27 Whitelists vs. Blacklists
- 28 Making Micro-Segmentation Work
- 31 Advantages of Dynamic Micro-Segmentation

4 Use Cases

- 37 Segmenting Development Environments
- 38 Micro-Segmenting Applications
- 38 Enforcing and Demonstrating Compliance
- 39 Moving Applications to Public Cloud Platforms
- 40 Securing the Hybrid Cloud
- 41 Scaling Applications on Cloud Platforms
- 42 Improving Incident Response and Mitigation
- 43 Levels of Segmentation

5 Implementation

- 47 Big Bang Not Required
- 47 Select the Project Team
- 48 Train the Team
- 49 Create Design Documents and a Project Plan
- 51 Install the Micro-Segmentation Solution
- 52 Integrate Logs, Events, and Alerts
- 53 Prioritize Application Groups
- 57 Discovery and Visibility
- 60 Model Policies
- 62 Test Policies
- 63 Be Ready to Fix Problems
- 64 Extend and Refine

6 Selecting the Right Solution

66 Enforcement Points

67 Visualization

68 Policy Modeling

69 Testing

69 Intelligence and Automation

70 Final Thoughts

Foreword

The Paradox of Can and Should

The explosive growth of IP networking has let a genie out of the bottle. The more an organization or a nation is connected to the Internet, the more vulnerable it becomes. But the less it is connected, the less able it is to compete in our globalized economy. The challenge is finding the right balance between “can” and “should”: the fact that we *can* connect every person, every server, every app, and every cloud vs. the reality that we *should* control how applications and users communicate with each other.

Unfortunately, from a security perspective, the network does not know whether it should forward packets. How can organizations address this challenge in their networks? By building a new cybersecurity foundation for data centers and clouds that supports innovation, but also offers protection against the lateral spread of attacks enabled by overly permissive connectivity.

Adaptive segmentation offers the optimal balance between connectivity and speed. By creating “watertight compartments” around application tiers, apps, or environments, segmentation reduces risks to an organization by limiting the damage a bad actor can inflict. This approach can protect applications regardless of the infrastructure on which they are running, whether it’s owned by the organization or leased from cloud providers like Amazon Web Services or Microsoft Azure.

Today security control and network connectivity are one and the same. Adaptive segmentation separates them. This allows IT and security professionals to decouple security controls from connectivity.

Cybersecurity professionals are tasked to make organizations both safe and agile. This dual mandate means selecting technologies that reduce the complexity of operations, but without compromising security. Therefore, security must not only leverage existing investments; it must also support future development and emerging technologies. That is why micro-segmentation needs to be smart and adaptive.

This book is a simple and practical guide for security and IT professionals interested in creating a new security foundation as they move down the micro-segmentation path. It outlines the capabilities you should look for, the benefits you can expect, and the next steps to take so you can find the right balance between can and should.

Alan S. Cohen, Chief Commercial Officer, Illumio

Introduction

For most IT professionals, “network segmentation” is a familiar but somewhat passé concept. That’s why “micro-segmentation” sounds at first like a minor incremental improvement.

Nothing could be further from the truth. Conventional network segmentation focuses on network performance and management. Micro-segmentation, in contrast, addresses critical issues related to security and business agility. It is a new and powerful approach to reducing risk and adapting security to dynamic IT environments.

As the name implies, micro-segmentation protects applications at a more granular level than conventional network segmentation. Equally important, it is adaptive; security policies “follow” applications automatically as they are moved and scaled.

With micro-segmentation, security becomes more dynamic, scalable, and consistent inside and across cloud and data center environments. Applications can be deployed faster, data centers can be reconfigured with fewer errors, and security and network administrators can spend less time on routine tasks like writing firewall rules.

The purpose of this guide is to explain the core concepts of micro-segmentation, describe the most common use cases, and provide practical advice on how to roll out micro-segmentation in a complex organization.

We hope that by reading this guide you will not only understand micro-segmentation better, but also see how it can make your organization more secure and more agile.

Chapters at a Glance

1 Segmentation for Security and Agility

Reviews two major challenges for IT security organizations, describes why micro-segmentation is different from conventional network segmentation, and explains how micro-segmentation addresses the two challenges.

2 Good Fences Make Good Data Centers: Controlling Access

Discusses why controlling the application environment is so difficult, and reviews several approaches to enforcing micro-segmentation.

3 Stuff Changes: Dynamic Segmentation

Outlines the differences between policies based on network constructs and policies based on applications, and explains how segmentation can be dynamic.

4 Use Cases

Explores important use cases of micro-segmentation, and explains how it can be applied at progressively more granular levels.

5 Implementation

Describes best practices for rolling out a micro-segmentation solution in stages.

6 Selecting the Right Solution

Enumerates criteria for choosing the micro-segmentation solution that best fits your organization.

1

Segmentation for Security and Agility

In this chapter:

- Understand two major challenges facing IT security organizations
- Assess the weaknesses of conventional network segmentation
- Learn why micro-segmentation is different and how it addresses the two challenges

“Love your neighbor as yourself;
but don’t take down the fence.”

— Carl Sandburg

IT security organizations today face a double whammy. On one side, persistent attackers are finding new ways to hide in plain sight inside the data center — inside application environments — and exfiltrate sensitive data. On the other, business managers want security people to get out of the way as they fire up new applications and technologies. These pressures create trade-offs between tightening up security (at the expense of flexibility) and pushing ahead quickly (but assuming more risk).

But there are a few technologies that have the potential to improve both security and business agility at the same time. Micro-segmentation is one of those rare technologies.

In this chapter, we introduce micro-segmentation by describing how it addresses two critical issues that confront IT organizations today.

Two Giant Challenges

Stopping Lateral Movement

We can no longer rely on perimeter defenses to keep the bad guys out, and are not doing so well catching them inside the data center either.

Most IT security professionals are familiar with frameworks such as Lockheed Martin's Cyber Kill Chain®. These frameworks explain how hackers can establish a beachhead inside a corporate network by exploiting a vulnerability or acquiring legitimate user credentials. From there, they explore the network to find additional weaknesses, and finally exfiltrate (or destroy) confidential information.

Statistics show that it is extremely difficult to reduce the “dwell time” of attackers once they have a foothold inside the data center. Virtualization and cloud technology exacerbate this challenge. It is hard to protect applications that can be executing anywhere, with pieces being moved around continually.

In this environment, limiting lateral movement within the data center becomes a top priority for IT groups. If a cybercriminal compromises the credentials of an employee who uses application A, can we make sure he can't reach applications B, C, and D? If a hacker uncovers the password of a system administrator in location X, can we make sure she has no way to connect to systems in locations Y and Z?

Your Credentials Are Out There

63% of confirmed data breaches involve weak or stolen passwords.

Credentials from employees of **97%** of the top 1,000 global corporations are available on social media or dark websites.

60% of organizations cannot detect attacks that use compromised credentials.

Sources: Verizon 2016 Data Breach Investigation Report, Digital Shadows_ Compromised Credentials study, Rapid7 2015 Rapid Detection and Response.

Making Security Dynamic

Once upon a time, computer systems, networks, and applications were relatively static. Data flowed along the same network connections, day after day. Applications ran on the same systems, month after month. Static firewall rules and access control lists (ACLs) could be used to limit traffic between systems across network segments.

But today's IT environments are far more complex and dynamic:

- Multi-tier applications often span multiple data centers and cloud platforms.
- Applications are released and enhanced daily.
- Software modules are spun up in new virtual machines, then moved across different physical servers.
- New instances of modules are created on demand so applications can scale up and down.
- Organizations continuously reconfigure data centers as they move applications, leverage cloud resources, and consolidate workloads.

Most areas of IT have been able to keep pace with these rapid changes by automating. They have introduced new concepts and automated tools in areas like virtualization, service orchestration, continuous deployment, and software-defined networking (SDN).

In contrast, many IT security processes are still manual. As a result, they act as a brake against technical and business innovation. For example, before deploying a new application, a security team may require weeks to analyze the current and proposed application architectures, document existing security policies, create new firewall rules, and test the resulting data flows.

Now consider multiple applications that need to be migrated. A project to consolidate two data centers might take many months or even over a year.

IT security faces a giant challenge to become dynamic and automate change so it can support, rather than impede, business agility.

Network Segmentation Has Maxed Out

Conventional network segmentation can restrict lateral movement by attackers, but it is neither granular nor flexible.

Today most data centers are divided into large zones, as illustrated in Figure 1-1. Traffic between the zones runs through a few choke points, typically firewalls. Traffic filtering policies are configured on each firewall.




Figure 1-1: Conventional network segmentation

This form of segmentation does provide protection against *some* lateral movement. An attacker who penetrates one zone cannot easily reach the other zones. However, conventional network segmentation has serious limitations:

- The zones are large, so an attacker who establishes a beachhead in one can target many applications within that zone. There is little support for East-West segmentation
- The high cost of firewalls and rearchitecting the network makes it expensive to create more zones.
- Firewalls are usually configured with thousands of rules, and even small modifications to applications and infrastructure may require manually reviewing, modifying, and testing hundreds of rules.
- Data centers and cloud platforms use different security tools and access rules, making it extremely hard to manage access for applications that cross them.
- Conventional firewalls and NGFWs can't effectively monitor or control traffic between virtual machines on the same hypervisor. Application ID technologies have limited applicability as you have to classify the traffic or send it off to your vendor to classify them.

Two Principles of Micro-Segmentation

Micro-segmentation is a security technique that enables fine-grained security policies to be assigned to applications, down to the workload level. It is built around two key principles: granularity and dynamic adaptation. The application of these principles makes micro-segmentation fundamentally different from conventional network segmentation.

Granular Segmentation

The trick to making segmentation granular is finding monitoring and enforcement points that are widely available and inexpensive. Monitoring and enforcement can be performed:

- In the network
- In the hypervisor layer of virtualization platforms
- On hosts
- On cloud platforms

Note that a little extra granularity can go a long way. You can strengthen security significantly just by fencing off a few high-value applications, or separating development and test systems from production environments.

In Chapter 2 we will discuss how to make segmentation more granular, and in Chapter 4, what use cases you might want to address first to get immediate benefits.

Dynamic Segmentation

In today's dynamic environments, granular segmentation would be impractical if administrators have to manually create and manage security rules on every enforcement point. More enforcement points would lead to less flexibility, more opportunities for error, and impossibly long workdays for security and network administrators.

“Dynamic segmentation” makes a granular approach to segmentation feasible by combining abstraction, intelligence, and automation. In the context of dynamic segmentation:

- **Abstraction** is the ability to express security policies in terms of application concepts (such as web, application, and database tiers) rather than in terms of network constructs (such as IP addresses, subnets, and VLANs).
- **Intelligence** is the ability to detect when changes are made to applications or the infrastructure, and then reconfigure policies to adjust for the changes.
- **Automation** is the ability to rapidly deploy new and revised security policies to monitoring and enforcement points, without human intervention.

Depending on the implementation, micro-segmentation can increase your adaptiveness in other ways as well, such as helping provide consistent security across data centers and cloud platforms. In Chapter 3 we will explain how to make segmentation dynamic.

Segmentation for Smart People

Micro-segmentation is segmentation for smart people. It addresses the challenge of stopping lateral movement by dividing IT environments into controllable compartments. It makes security dynamic by allowing security rules to be expressed in terms of application concepts, and reconfigured automatically when applications and infrastructure components change.

Also, micro-segmentation can be very economical if it uses existing infrastructure components, rather than expensive firewalls and NGFWs, as enforcement points.

2

Good Fences Make Good Data Centers: Controlling Access

In this chapter:

- Understand why controlling application environments is so difficult
- Explore the idea of isolating applications and workloads
- Review pros and cons of different enforcement points

“The gaps I mean,

No one has seen them made or
heard them made,

But at spring mending time we
find them there.”

— Robert Frost (Mending Wall)

We have said that stopping lateral movement is one of the key benefits of micro-segmentation. Now we will look at why controlling applications is so difficult, and consider different ways to make micro-segmentation cost-effective.

The $(N \times (N-1))/2$ Problem

Access has several meanings in IT environments. Often it refers to permissions for specific individuals to interact with specific network and system resources. This is not exactly the type of access we are going to focus on here.

In this section we discuss access in terms of software tiers and services communicating with each other, as in: “in a three-tier web application, the web server needs access to the application server, and the application server needs access to the database server.” In this context, controlling access means allowing or blocking traffic between the software modules and services in applications.

The great challenge for this type of access control is that the number of potential connections among software modules expands almost exponentially. As shown in Figure 2-1, if you have four software modules, there are six potential connections between them. Add two more modules, and there are 15 connections. Go to 10 modules and you have 45 connections.




Figure 2-1: Connections between 4, 6, and 10 software modules

The equation for calculating potential connections is $(N \times (N-1))/2$. This equation tells us that 50 modules and services have 1,225 connections between them, 200 modules have 19,900 connections, and so forth.

In a data center or cloud platform the situation is even more complicated. You can have software modules running on different physical servers, communicating via multiple services over multiple ports, or modules in different virtual machines on the same server, communicating via multiple services and ports, or... you get the picture.

The Name of the Game: Isolating Applications and Workloads

From this point on we are going to discuss segmenting environments or micro-segmenting applications and workloads. *Environments* in this context could be development and production environments, or distinct customer environments if you are a service provider. *Applications* are collections of workloads that address a business or technical need. *Workloads* are discrete operating system instances running one or more software modules, services, or containers.

A New Way of Looking at Segmentation

We can look at a computing environment as a collection of applications and services, which are made up of workloads and the connections between them. The workloads or containers can be executing on physical (“bare-metal”) systems or in virtual machines, in data centers, on public cloud platforms like Amazon AWS, Google Cloud Platform, and Microsoft Azure, and in private clouds (Figure 2-2).

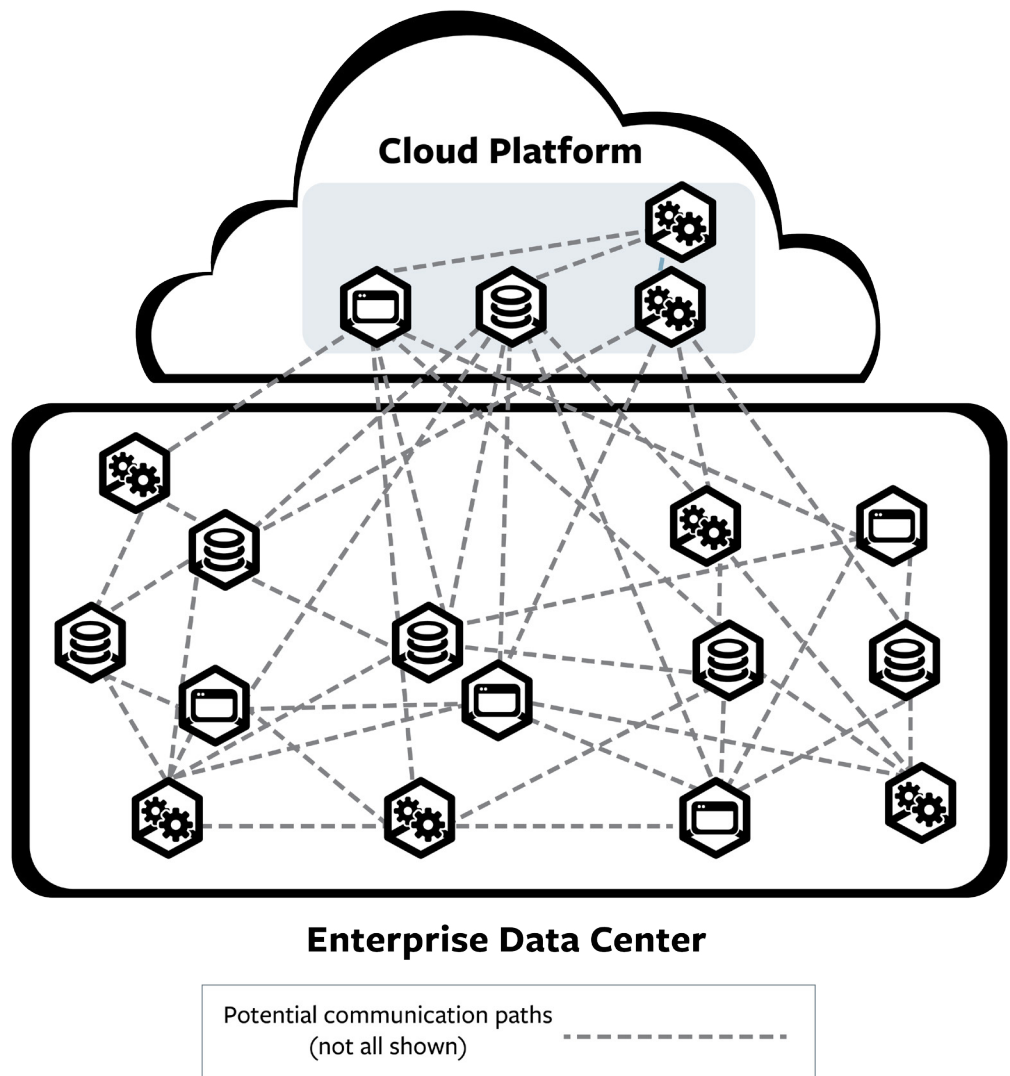


Figure 2-2: A computing environment viewed as a collection of applications and workloads, and the connections between them

From this perspective, micro-segmentation can be seen as the process of isolating environments, applications, or workloads by allowing traffic to flow only through approved connections.

In effect, micro-segmentation prevents attackers from using unapproved connections (that is, connections that are not part of the application design) to move laterally from a compromised application.

An Example: Fencing Off a Three-Tier Web Application

Let's consider the simple scenario illustrated in Figure 2-3.

What if your accounts payable application had access to the customer database used by the customer service system? In that case a cybercriminal could use social engineering to steal credentials from one of your vendors, gain access to the accounts payable system, and move laterally to reach the customer database, causing a massive data breach.

What if the accounts payable application could be accessed from the human resources (HR) system? Hackers impersonating an employee might find a way to circumvent the controls in the accounts payable system and start issuing checks to themselves.

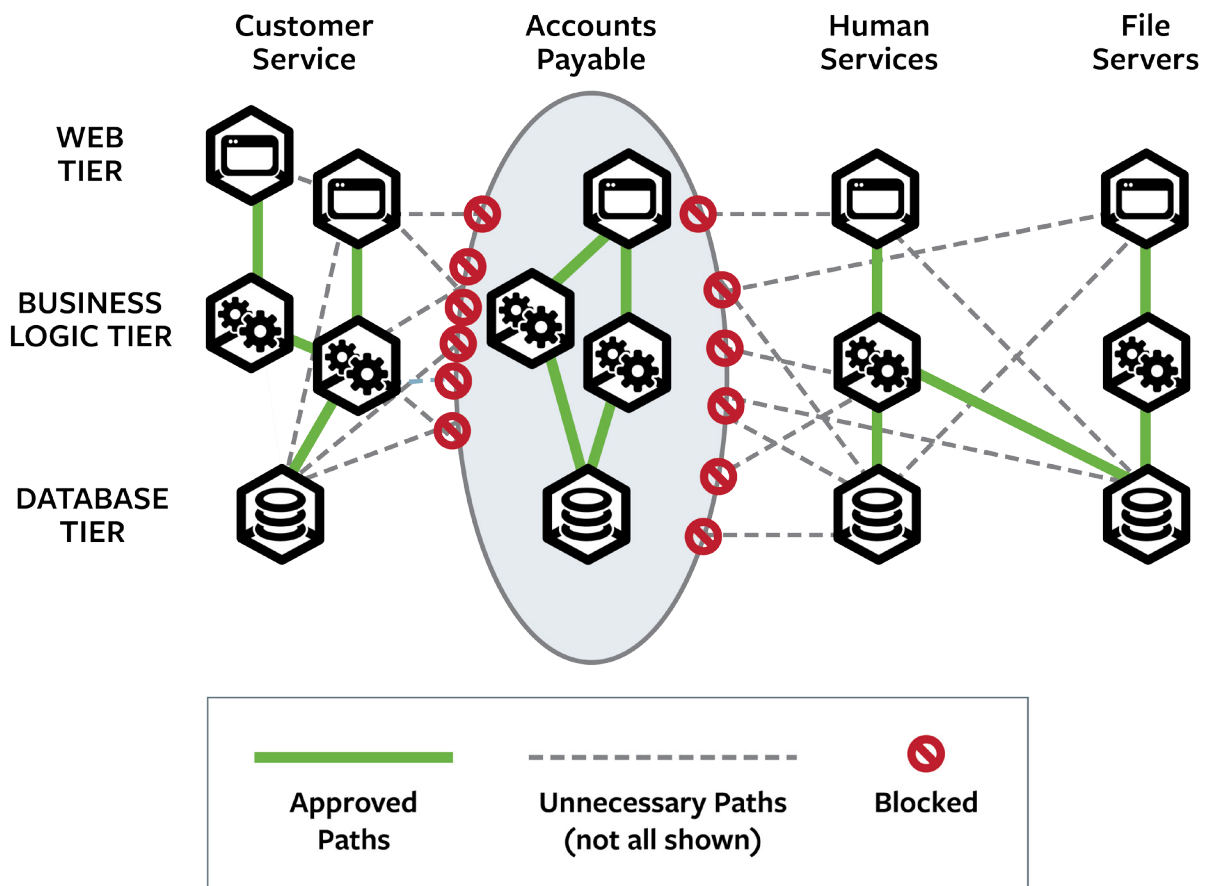


Figure 2-3: Fencing off a web application by limiting traffic to connections actually needed by the application

However, if you could fence off the accounts payable software by allowing traffic to flow only along the communication connections needed by that application, you would prevent lateral movement from the customer service and HR systems, and dramatically lower your risk of a career-ending outcome.

Enforcement Points – Three Approaches

As we discussed in Chapter 1, conventional firewalls operate as choke points that manage communication into and out of data centers, and also between zones within data centers and cloud platforms. But zones and security groups typically contain many applications, so there are always many open connections going in and out. An attacker who finds one open connection in can reach all of the applications and workloads inside that zone.

Network switch vendors have added features that allow network administrators to divide zones into subnets and virtual LANs (VLANs). However, these subdivisions still contain many applications, and it can be hard to keep track of all of the policy rules that apply to each one. However, network segmentation is not the same as security segmentation. The former is focused on network performance and the latter on security isolation.

In theory, you could isolate your applications by deploying more firewalls. But firewalls are too expensive and too inflexible to provide granular segmentation at a reasonable cost.

Fortunately, IT security vendors have found alternatives that can act as ubiquitous and inexpensive enforcement points for security policies.

Network Equipment

One approach is to use network equipment, such as switches, routers, and load balancers, as enforcement points for security policies.

Most medium and large organizations have deployed hundreds of network devices, so they can divide data centers and networks into far more zones than would be possible with firewalls and NGFWs. Typically, the only capital expense is for a handful of policy controllers (systems used to create and manage policies on the enforcement points).

An advantage of these solutions is that frequently they are integrated with other offerings from the network equipment vendor that provides them, for example, that vendor's network service orchestration product. One example is Cisco's Application Centric Infrastructure (ACI).

One disadvantage of approaches based on network equipment is that their policy models rely on networking constructs such as IP addresses and VLANs, so administrators often have to perform "unnatural acts" (from the perspective of the applications) to define and enforce desired policies.

In addition, adding or changing network devices can entail effort and delay while network administrators analyze complex traffic steering requirements and reconfigure networks.

Finally, micro-segmentation solutions provided by each network equipment vendor are limited by the scope of that vendor's devices; i.e., they can't be extended to cloud platforms or data centers that use another network equipment vendor's devices. They also might not work with management products from third parties.

Hypervisors

Hypervisor-based micro-segmentation involves adding firewall-like capabilities to the hypervisor layer of a virtualization

platform. Then the hypervisor can enforce security policies for traffic to and from workloads running in virtual machines.

Using hypervisors as enforcement points provides granular segmentation for all applications and workloads deployed on a common platform, as well as integration with other hypervisor-based management tools from the same vendor.

The downside of hypervisor-based micro-segmentation is that it provides no protection for software running on physical servers, on hypervisors from other virtualization vendors, or on cloud platforms that don't offer that specific brand of hypervisor. This can become challenging in container environments.

Host-based Software

A third approach to micro-segmentation takes advantage of software built into major operating systems. The most important of these are iptables on Linux systems and the Windows Filtering Platform (WFP) on Windows servers.

These utilities monitor network traffic and apply security rules for specific applications and workloads, providing granular segmentation. Because they are part of standard operating systems, they are available wherever those operating systems are deployed: on physical servers, on virtual machines, and even on platforms with container technologies like Docker.

Another advantage of host-based micro-segmentation solutions is their ability to acquire and use information about individual applications and services, completely independent of any network constructs. We will discuss this further in Chapter 3.

The main challenge for host-based solutions for micro-segmentation is deploying software agents to many distributed systems.

The Visibility Challenge

You may have noticed that in our discussion of enforcement, we left out one major issue: visibility. How do we know which connections are needed by our applications? Differentiating these from inactive and unapproved connections can be a formidable challenge.

The $(N \times (N-1)) / 2$ equation tells us that even a medium-sized organization can have a huge number of potential connections between applications and workloads. Research shows that less than 5 percent of those potential connections are typically approved and in use.¹

To be effective, a micro-segmentation solution needs to include a tool or utility that identifies the roughly 5 percent of potential communications that are in use.

Ideally, this tool should display information in a way that helps administrators:

- Identify the applications that are most important to protect
- Determine what connections are needed by those critical applications
- Use “whitelisting” to lock unapproved communications and create security policies that minimize the chance of attackers using the approved connection

¹ eWeek: [Illumio Launching Effort to Reduce Data Center Attack Surface.](#)

Figure 2-4 shows an example of the output of such a tool.

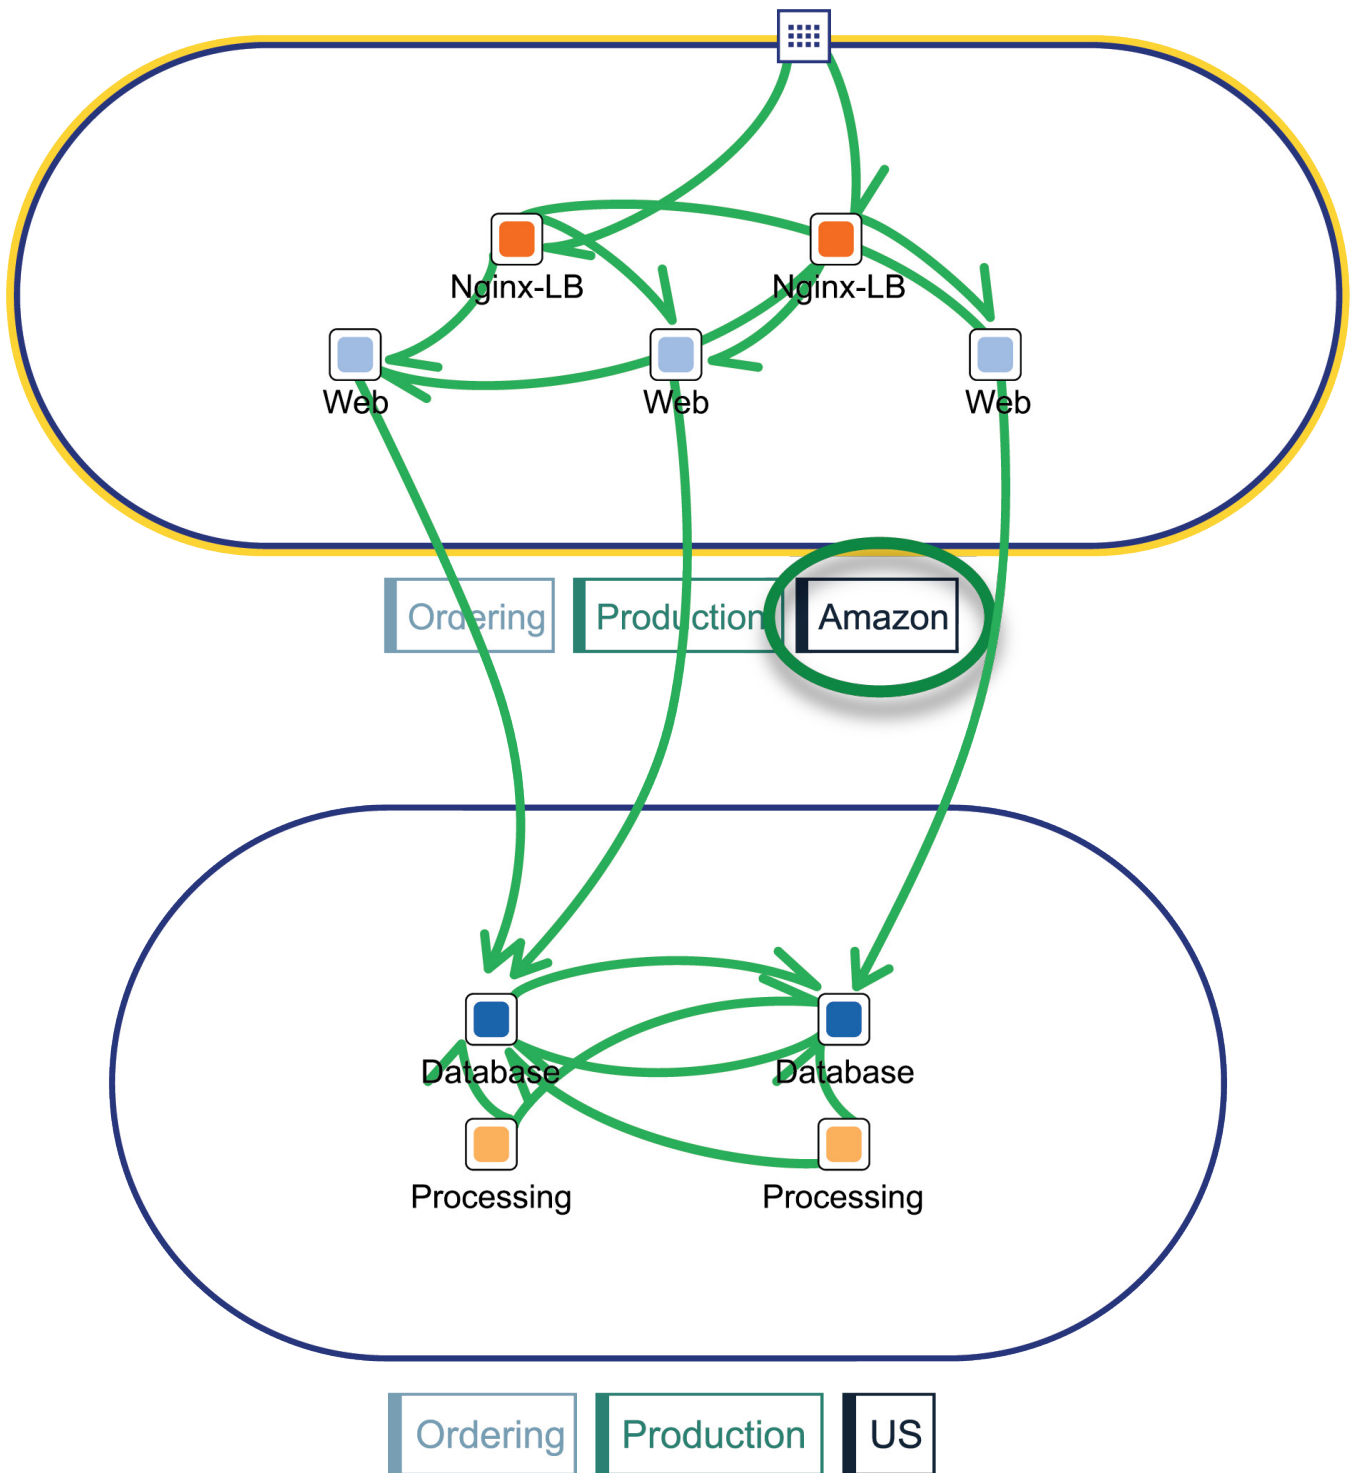


Figure 2-4: A diagram showing active connections between workloads in a hybrid cloud/data center environment. The connections displayed should not be blocked, and should be assigned security policies. Source: Illumio

Granular Segmentation: Who Benefits?

Making segmentation more granular offers advantages in several areas related to security and compliance. Let's look at how it can help people with various roles in the IT organization do their jobs better.

Application Owners

Micro-segmentation can isolate applications from each other. It allows security policies to be defined on an application-by-application basis, insulated from other applications and from issues related to the underlying infrastructure. That means that “application owners,” the software developers or administrators with the deepest knowledge of each application, can fully participate in designing and updating the security policies for their applications.

The team developing the application therefore has control over the speed at which the application can be secured.

Compliance and Risk Officers

Many industry standards and government regulations require certain types of sensitive information to be protected by access controls. Micro-segmentation can put a fence around environments and specific systems that hold protected data, and prove that the strongest measures are being used to comply with data protection requirements.

By reducing opportunities for lateral movement, micro-segmentation also diminishes the risk inherent in any single penetration of a data center or cloud environment. This is sometimes referred to as limiting the “blast radius” of a breach, because it contains the damage to an area that can be as small as a single application, or even a single workload.

SOC and Incident Response Teams

Most micro-segmentation solutions can generate alerts based on attempts to use unapproved connections. These alerts provide early warning of targeted attacks and clues about the attackers' methods. This information can help Security Operations Center (SOC) and incident response teams separate actual advanced attacks from events that are merely suspicious, block ongoing attacks, and reduce vulnerabilities.

3

Stuff Changes: Dynamic Micro- Segmentation

In this chapter:

- Understand the difference between policies based on network constructs and policies based on applications and workloads
- Learn how to make micro-segmentation adaptive
- Review the advantages of dynamic micro-segmentation

“There is nothing permanent
except change.”

— Heraclitus

Why Firewalls Are Inflexible

Firewall rules are notoriously difficult to keep up to date. They are hard to reconfigure when an application is modified or a data center is reconfigured. Two factors contribute to firewall inflexibility.

Too Many Rules in One Place

Several thousand rules are configured on the average organization's firewall. It is almost impossible for firewall and network administrators to keep track of them all. When applications and infrastructure are updated, extensive research and analysis are required to know which rules to change and what order they need to be in.

Also, busy administrators rarely take the time to identify and remove outdated and unnecessary firewall rules. In many organizations, administrators do not remove *any* rules for fear of inadvertently blocking the connections of an active application. As a result, rules accumulate and administrators are hard-pressed to figure out which ones are safe to modify.

Policies Based on Network Constructs

Another factor making firewalls inflexible is that their rules are based on network constructs such as:

- Source and destination IP addresses
- Source and destination ports
- Communication protocols
- Subnets, VLANs, and zones

The first problem with this approach is that IP addresses are rarely statically assigned. They change, especially in the cloud. Policies based on IP addresses can quickly become obsolete.

The second problem is that there is no intrinsic relationship between application security requirements and network constructs, which makes it difficult to convert the former to the latter.

You can't just say: "The web front end of this application needs to access the business logic, which needs to access the database." Someone has to look up the IP addresses of the workloads, the protocols the application wants to use to communicate between workloads, and the ports the application is designed to use. Only then can the firewall rules be written to allow the necessary data flows.

If you want to move the web front end to the cloud six months later, someone (maybe not the same person who created the rules) has to go back, find all of the relevant rules, and modify them to account for new IP addresses, WAN connections, different protocols, and different firewalls.

Do you want to prevent Development's test environment from accessing live production databases? Block software in Costa Rica from accessing files in Tokyo? Then someone needs to perform more research, analysis, and testing.

The third problem is that organizations may attempt to try this through coarse-grained rules such as subnets and zones, but as the applications evolve, new applications are introduced and other applications are decommissioned – the extra IP space is whittled away and their ACL catalog becomes more porous, or the applications evolve and break down the zoning strategy.

Dynamic Micro-Segmentation

Dynamic micro-segmentation solutions take a different approach. They define rules based on applications, workloads, and the relationships among them. That allows them to adjust rules automatically as applications and environments evolve.

With dynamic micro-segmentation, you can define a set of policies for communication between the web front end and business logic of an application, rather than between two IP addresses. If the web front end is moved, that set of policies can be made to follow the workload automatically to a different server, data center, or cloud platform.

Some micro-segmentation solutions have even more advanced options. For example, multiple workloads can share a label and thereby gain the same policy.

The first capability addresses some of the issues mentioned above. All workloads classified as “test” can be blocked from accessing workloads classified as “production.” Applications with a “Latin America” location label can be granted or denied connections to applications with an “Asia headquarters” location label.

The second capability is particularly useful if you want to scale applications on virtual platforms. When your orchestration software detects increased demand on a workload and fires up a new instance, that instance will be automatically assigned the security policies of the original workload, and can immediately use the same connections to related workloads.

Note:

The micro-segmentation product category is fairly new, and not all solutions have the dynamic capabilities mentioned here or elsewhere in this guide. Please evaluate specific products carefully to make sure they have the features that are important to your organization.

Whitelists vs. Blacklists

Another difference between firewalls and micro-segmentation solutions is that the former typically use blacklists while the latter primarily use whitelists.

The default setting on most firewalls is any/any. They rely on administrators to create hundreds of rules to block traffic to and from suspicious IP addresses, traffic that uses unapproved protocols for a specific server, etc.

Most firewalls are forced to take this approach because rule creation is so difficult and error prone. To avoid “breaking” active applications, it is safer to add restrictive rules slowly and carefully.

In contrast, many micro-segmentation solutions restrict network traffic to connections that have been explicitly enabled by administrators. This whitelist approach has several major advantages:

- Control over network traffic is much better, because far fewer connections are open to reach an application. There are an unlimited set of connections; however, there are a limited number of open ports.
- Rule maintenance is much easier, because there are far fewer rules to update, and more of the updating is done automatically.
- Rule evaluation and testing have little or no performance impact on the enforcement points, because there are far fewer rules to evaluate.
- A malformed rule only impacts the individual workloads rather than an entire data center.

Of course, administrators must be very careful that micro-segmentation solutions don't break active applications. But a few micro-segmentation solutions have connection visualization and policy test capabilities that make this challenge very manageable, as we will discuss in Chapter 5.

Making Micro-Segmentation Work

In this section we will look at how micro-segmentation solutions work, keeping in mind that different products have different feature sets.

Architecture

Most micro-segmentation solutions have two basic architectural components, as illustrated in Figure 3-1.

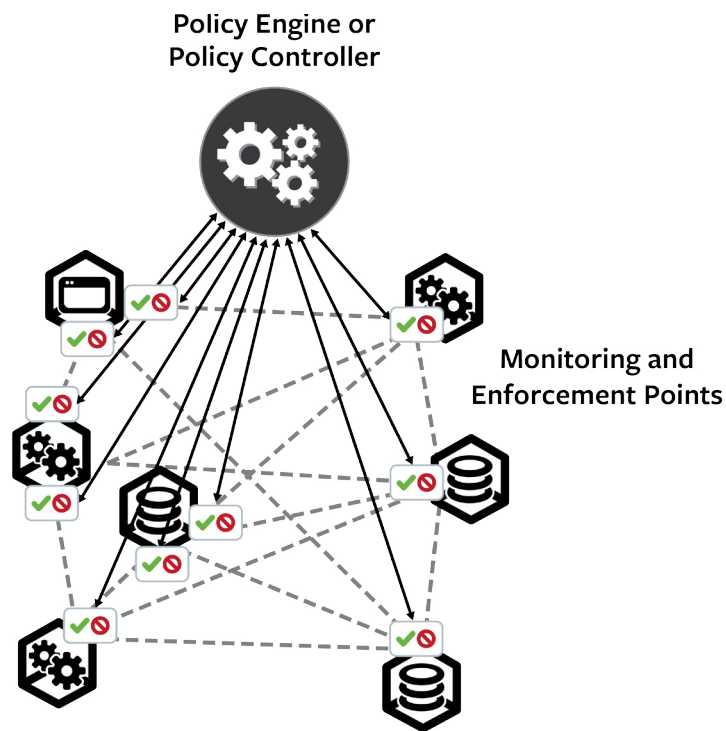


Figure 3-1: Basic components of a micro-segmentation solution

The policy engine, or policy controller, is a central point that provides visibility into connections and control over security policies. Among other tasks, it:

- Discovers applications and workloads by collecting data from the monitoring and enforcement points
- Identifies active connections
- Helps define, model, and validate security policies

The monitoring and enforcement points collect data on information flows to give administrators visibility into traffic along approved and connections. They also need to be able to inspect and evaluate network packets so they can apply security policies to allow or block network traffic.

As discussed earlier, the most promising monitoring and enforcement points have proven to be network devices, hypervisor layer software, and existing host-based utilities. When the monitoring and enforcement points are based on devices or operating system software, they typically require software modules or software agents to manage the security policies. Depending on the product, these modules or agents can be pre-loaded as part of the operating system image, or distributed with scripts, orchestration tools, or other automated software distribution tools.

Translating Natural Language Policies Into Enforceable Rules

As we have mentioned, one of the great advantages of micro-segmentation solutions is that security policies can be abstracted from the network and defined in terms of applications and workloads, and their relationships.

However, applications need to communicate by sending data packets to each other over networks. The networks understand IP addresses and network segments, not the relationships between application modules and services. That means that actual policy enforcement must be performed using rules that are very similar to firewall rules.

To enable this, the policy engine translates natural language policies into enforceable micro-segmentation rules.

Disseminating Rules to the Enforcement Points

After natural language policies have been created and translated into enforceable rules, the policy engine is responsible for disseminating them to the enforcement points based on the policies designated by the security team.

Revising Policies Automatically

Policy engines detect changes to applications. For example, they detect when workloads are moved to different servers in a data center, or migrated from the data center to a cloud platform. The policy engine re-computes policies to account for the changes, translates them again into enforceable rules, and disseminates the updated rules to the monitoring and enforcement points.

Advantages of Dynamic Micro-Segmentation

Dynamic micro-segmentation helps security become as dynamic as other areas of IT. Let's review some of the ways it improves business agility and allows organizations to take advantage of technologies like virtualization and cloud computing.

Deploy and Enhance Applications Faster

As we discussed earlier, firewall rules are notoriously difficult to track and manage. Firewall and network analysts often need weeks to research, analyze, modify, and validate firewall rules in order to secure connectivity for a new application or enhancements to an existing one.

Dynamic micro-segmentation reduces “time to delivery” for new and enhanced applications by:

- Increasing visibility into connections between environments, applications, and workloads
- Allowing connectivity requirements to be expressed in terms of application logic rather than network constructs
- Translating natural language policies into enforceable rules

Involve Application Teams in Designing Secure Connectivity

Today, application teams usually have only a peripheral role in designing secure connectivity. They outline application requirements, which IT security analysts translate into

connectivity policies, which IT administrators translate into firewall rules. Something is always lost in the translation, resulting in gaps between business requirements and technical implementations.

Some micro-segmentation solutions define policies in natural language. These solutions allow business people and technologists to work together to understand business risk – because policies are defined in the language of the applications – and to define secure connectivity policies. These policies are then translated automatically to enforceable segmentation rules (layer 3/layer 4 packet filters). Because this process avoids handoffs and misunderstandings, it is faster, more reliable, and less stressful for everyone involved, as illustrated in Figure 3-2.

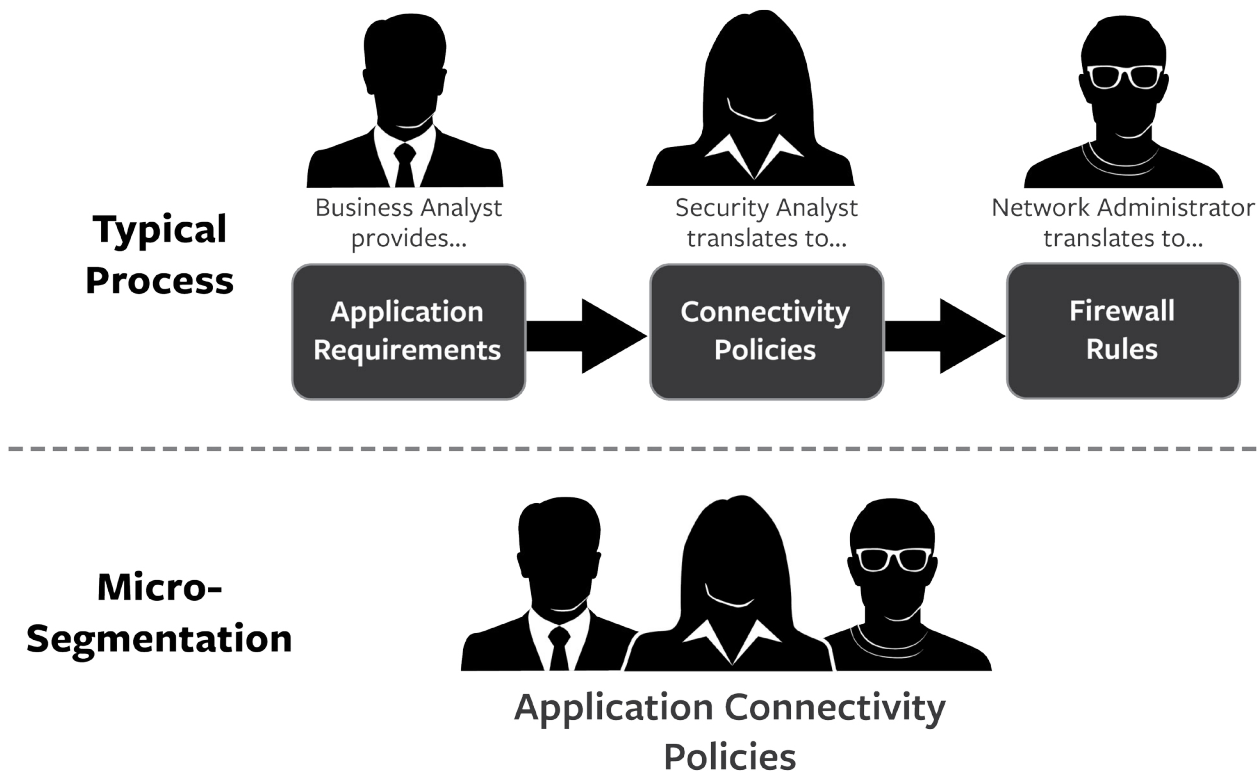


Figure 3-2: Policy-based segmentation allows business people and technicians to collaborate in defining application connectivity policies.

Support Fast-Changing Infrastructures

Many organizations have invested heavily in virtualization platforms, container technology, and orchestration software that make it easy to move applications and workloads around data centers on the fly. They can quickly add and reconfigure software, servers, storage networks, network devices, and other components of the computing infrastructure. They can spin up and replicate hundreds of virtual machines and containers. But they still must wait for firewall rules to be updated manually.

Some micro-segmentation solutions detect changes to the infrastructure, automatically revise policies, and distribute policies to the enforcement points. These capabilities allow organizations to enjoy the return on their investments in IT automation in minutes, rather than hours or days.

Manage Hybrid Cloud and Distributed Data Center Environments

Most organizations operate hybrid environments, with some applications deployed on public or private cloud platforms, others deployed in corporate data centers, and still others distributed across sites. These organizations are forced to use multiple security tools.

Some micro-segmentation solutions can be used across all cloud platforms and data centers, which simplifies management and reduces costs. A single solution also means that policies for complex applications can be changed more quickly and reliably.

Scale Cloud Applications

One of the greatest advantages of virtualized and cloud platforms is that they can scale up applications quickly by firing up new instances of application workloads on demand. However, this flexibility is diminished when administrators have to manually create or modify fine-grained rules to allow traffic to flow to and from the new instances.

Some micro-segmentation solutions detect when a new instance of a workload is fired up, match that instance to the original workload, and then automatically apply the same policies. This means that secure connectivity is provided for the new instance immediately, so there is no waiting to scale the application.

Plantronics: Increasing Velocity and Innovation with Micro-Segmentation

Plantronics is a global leader in audio communications for businesses and consumers, from Bluetooth® headsets, to personal speaker systems, to gaming solutions and unified communications. To support its business users, the company needs to enhance applications and deploy technology in hours, days or weeks, not months or years.

Plantronics invested heavily in cloud technology and virtualization to create a dynamic computing environment. However, the IT staff found that every time they added a virtual machine or another service, they had to spend a huge amount of time updating static security configurations.

Plantronics decided to deploy the Illumio Adaptive Security Platform (ASP)[™] as its micro-segmentation solution. The ASP allows the IT team to create a profile for a virtual machine, then automatically apply the profile to new instances as they are instantiated. This process has eliminated the wait time to reconfigure static security rules.

Today Plantronics can deploy new applications with a couple of mouse clicks, while providing better protection from threats. In addition, the IT organization can quickly supply computing resources to applications, even when the software is deployed on fractional parts of dozens of virtual servers, because the ASP applies the correct security policies to every running process.

You can read more about Plantronics and micro-segmentation at:
<https://www.illumio.com/success-story-plantronics>.

4

Use Cases

In this chapter:

- Learn about use cases such as segmenting development environments, micro-segmenting applications, enforcing compliance, moving applications to the public cloud, and securing the hybrid cloud
- Explore how micro-segmentation can be applied progressively to different levels of segmentation

“In theory there is no difference between theory and practice. In practice there is.”

— Yogi Berra

So far, we have discussed what micro-segmentation is and how it works. Now it is time to look at how micro-segmentation can be used to address specific technical and business problems.

Segmenting Development Environments for Better Security

Most IT groups maintain different environments to develop, test, and stage applications before they are released onto live production systems. Except for a few well-defined processes for promoting releases from one stage to the next, these environments should be isolated from each other.

In practice, however, many organizations are not careful about enforcing this separation. That carelessness can have serious consequences.

For example:

- Developers and software testers can violate compliance requirements by pulling credit card numbers, Social Security numbers, protected health information, and other confidential data from production databases for testing.
- Attackers who penetrate the production environment can work their way back into the development system and steal software programs and privileged user credentials.
- Developer applications can interfere with production applications and break them if there is no segmentation.

A micro-segmentation solution can separate these environments, ensuring that only approved connections are open within and across environments.

Environmental segmentation is extremely valuable for:

- Companies in regulated industries that must protect certain types of information from any unauthorized access, including access by developers and testers

- Software vendors, computer service providers, and other firms where software under development is critical intellectual property
- Organizations that do software development and testing on public cloud platforms, and are concerned about protecting their code from other “tenants”

Micro-Segmenting Applications

Micro-segmentation can put a fence around applications and limit the ability of attackers to move laterally between applications. This capability is especially valuable for high-value software such as:

- Applications that store confidential customer and employee information, intellectual property, and company financial data
- Mission-critical software, where downtime results in lost revenue, damaged reputation, or lower productivity
- Unique applications that give the organization a competitive edge

Enforcing and Demonstrating Compliance

Micro-segmentation can help organizations in regulated industries provide the highest levels of protection to workloads that store or process payment card information, account and Social Security numbers, protected health information, addresses and additional contact information, passwords and access credentials, and other data specified by regulations and industry standards.

This protection is especially important for organizations subject to regulations and standards such as PCI DSS, HIPAA, FISMA, Sarbanes-Oxley, and the EU GDPR.

Besides making it harder for attackers to access protected information, micro-segmentation can help organizations:

- Demonstrate to auditors that security controls are in place
- Apply a consistent set of controls across cloud and data center environments

Moving Applications to Public Cloud Platforms

Many IT managers find that migrating major applications to the cloud, or consolidating data centers, involves the same planning and execution skills required for the D-Day invasion of Normandy. Researching, analyzing, planning, and testing firewall rules often represent a significant portion of this work.

Micro-segmentation solutions can speed up migration and consolidation by:

- Discovering and visualizing applications, workloads, and connections in the current environment
- Creating security policies at the application level, abstracted from network details
- Modeling and testing those policies based on the parameters of the new environment
- Automatically updating and distributing rules to the enforcement points when applications are moved to the new environment

By automating parts of the transition process, micro-segmentation solutions reduce the risk of missing a detail and breaking an application during the cut-over.

Securing the Hybrid Cloud

Organizations with hybrid cloud environments (infrastructures that include both cloud platforms and corporate data centers) face the challenge of using two or more sets of security and management tools. This situation creates extra work and makes it very hard to provide consistent protection for applications or organizations to span the two types of environments.

A micro-segmentation solution can allow organizations to:

- Discover and visualize data flows between workloads that cross cloud and on-premise locations
- Provide seamless protection for applications that span clouds
- Implement uniform security policies across clouds (including multiple data centers and cloud platforms from different service providers)
- Move application workloads between locations without breaking the applications or requiring major projects to analyze and rewrite security policies

Scaling Applications on Cloud Platforms

Are you a fan of cloud computing? If you are, it is probably because you liked the fact that public cloud platforms such as AWS, Google Cloud Platform, and Microsoft Azure are designed to create workloads and scale applications on demand by launching workloads on virtual machines.

However, security can be a barrier on cloud platforms if you have to create security rules manually to allow connectivity to the new workloads.

Fortunately, some micro-segmentation solutions can detect new workloads when they are created on a cloud platform, recognize that they are copies of existing workloads, and automatically apply appropriate security policies.

Figure 4-1 illustrates how a cloud platform can meet a spike in demand by spinning up multiple instances of an application's web and database modules. The green lines represent connections that have been enabled to allow the new modules to exchange data with other components of the application, either in the cloud or a corporate data center.

The red lines in Figure 4-1 indicate connections that are not allowed by policy. The micro-segmentation solution can either block these connections, or keep them open and generate alerts to the SOC when they are used.

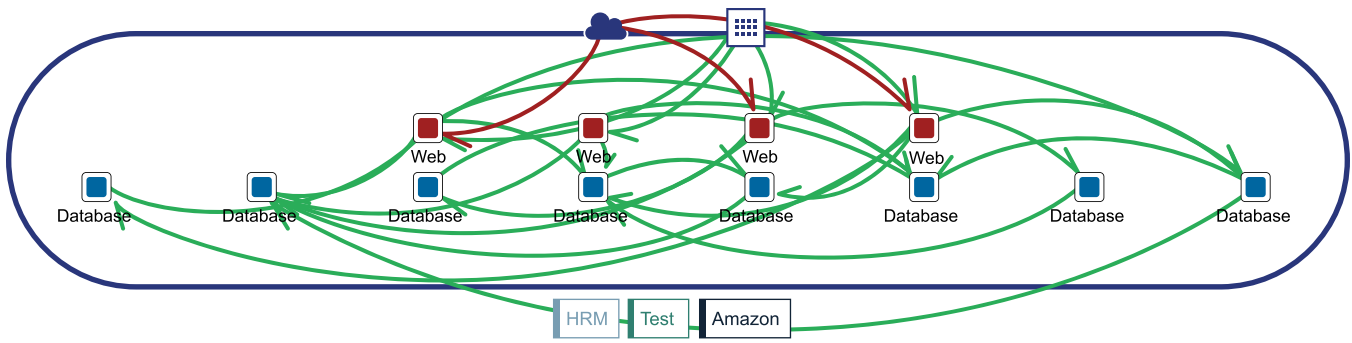


Figure 4-1: A visualization of multiple workloads spun up to meet a spike in demand, and the communications between them. Source: Illumio

Improving Incident Response and Mitigation

In addition to generating alerts about traffic on connections that violate policy, micro-segmentation solutions can help incident response teams by:

- Providing information from logs that reveal attackers' tactics, techniques, and procedures (TTPs)
- Provide telemetry as to what specific application has a policy violation (not just IP addresses)
- Help analysts visualize information flows and block or apply restrictive policies to those that are rarely used or provide little value


As mentioned earlier, micro-segmentation solutions also aid incident response teams (and the entire organization) by limiting the blast radius of attacks.

Levels of Segmentation

There are some situations where it makes sense to progress over time from simpler coarse-grained to more-granular levels of segmentation, as illustrated in Figure 4-2. Use cases include progressively improving protection for high-value applications, and tightening up access to data covered by compliance standards and regulations. Let's discuss how this progression might go.



Environmental Segmentation



Application Segmentation



Server and Tier Segmentation



Process Segmentation

Figure 4-2: Four levels of segmentation. Source: Illumio



Environmental Segmentation

Separating the environments used to develop, test, and stage software releases from production environments is a “no-brainer” for many organizations. It can be done quickly, and instantly improves both security and compliance.



Application Segmentation

Isolating applications from one another provides the obvious advantage of preventing attackers who have compromised one application from moving laterally to reach others. It also reduces the exposure of an application to a bad actor.

Many conventional network segmentation schemes do not take advantage of this insight. If you refer to Figure 1-1, you will see web servers from multiple applications in the same DMZ, and database servers from multiple applications in the same internal zone. Grouping application modules that way simplifies some management functions, but creates target-rich zones for hackers.



Server and Tier Segmentation

Servers running different tiers of an application can be separated from each other. Each tier obviously needs to communicate in places with adjacent tiers, but there are benefits to restricting communication to connections.

Examples of useful server and tier restrictions include:

- Preventing front-end web servers from bypassing the business logic tier and directly accessing the database
- Preventing business logic workloads from accessing each other (unless there is a business reason)

- Preventing servers from accessing each other using unusual ports and protocols
- Controlling directional flows; for instance, if a service is designed only to answer queries, it should not be allowed to initiate contact with other systems.

Even if attackers create a beachhead in one part of an application, these restrictions may limit their ability to reach targets (like customer data or intellectual property) in other parts of that application.



Process Segmentation

Some applications, notably Microsoft Active Directory, use large, dynamic port ranges during operation. Conventional firewalls, as well as micro-segmentation solutions that rely on network equipment, cannot adjust for dynamically changing ports. This inflexibility creates a dilemma for administrators:

- Either leave the entire port range open (and available to attackers); or
- Risk breaking applications

Micro-segmentation solutions that employ hosts as enforcement points can monitor real-time activities related to individual processes on each workload, and the ports they use. They can use this information to restrict traffic to currently active ports. This is referred to as *process segmentation* (or sometimes *nano-segmentation*), and it gives administrators the ability to enforce policy without compromising security or application functionality.

5

Implementation

In this chapter:

- Review an implementation plan
- Consider how micro-segmentation can be rolled out in stages
- Understand how discovery, policy modeling, and testing can reduce risk

“In real life, strategy is actually very straightforward. You pick a general direction and you implement like hell.”

— Jack Welch

A well-planned implementation process is especially important for micro-segmentation. A good plan leads to better outcomes, sooner. Also, the right techniques can reduce the risk of blocking connections needed by applications.

Big Bang Not Required

An organization-wide “big bang” implementation is not recommended for micro-segmentation. You should roll out micro-segmentation gradually to different applications – as the organization is ready.

You can also minimize risk by using a four-step process that moves from visibility, to policy modeling, to testing, and only then, to enforcing the policy rules and blocking unnecessary connections.

Although a staged implementation plan can take a while to complete, it has the advantage of providing benefits as soon as you start making application workloads and connectivity visible.

Select the Project Team

It is important to select a project team with the knowledge, skills, and available time to make the implementation a success. The key members typically include:

- An executive sponsor
- A security or project architect, to establish overall direction and strategy
- A project manager
- A tech lead
- Project management – either internal or external (possibly from the micro-segmentation vendor)

In addition, the team needs to be able to draw on resources from:

- Security architecture and security operations
- Network engineering and the network operations center
- The systems administration, directory, and virtualization groups
- The group that manages IT process automation and orchestration
- Application teams

Train the Team

Make sure the entire project team is educated about:

- The functions and value of micro-segmentation
- The features and operation of the micro-segmentation tools that will be used
- The structure and activities of the implementation plan

If possible, send the entire team to a one- or two-day overview training class from your micro-segmentation solution vendor or an independent consultant. Training will give all team members a solid grounding in the core concepts, architecture, and policy language of the product you are implementing, and will allow them to begin contributing immediately to the project.

Create Design Documents and a Project Plan

After training, the project team should create design documents and a project plan. Try to involve everyone from the beginning in developing the plan. Organizations that delay involving team members until the start of “their” part of the project often fail to uncover critical tasks or project blockers until well into the deployment process.

That is because each member of the cross-functional team will have visibility into different issues that may arise. You need to leverage their collective knowledge about applications and software services, network and application architectures, virtualization and cloud computing, automation and orchestration tools, and existing policies and procedures. It is much easier to create a complete, robust implementation plan if you take advantage of everyone’s expertise and experience early in the project.

The design documents and project plan should address:

- Labeling and metadata development, import, and management
- A complete initial security policy
- Automation required for bulk installation
- Integration with reference data such as configuration management databases (CMDBs) and security information and event management (SIEM) systems

- Workflows and internal processes that may need to be tweaked to accommodate the new security model
- Pre-production preparation, including testing and certification.

The team should make policy decisions. For example, who is authorized to create and modify segmentation policies, and what procedures should be followed before enforcement is turned on in each application area?

The team also needs to decide on the sequence in which applications will be protected. That sequence can then be turned into a roadmap showing the dates when major tasks will be completed and enforcement for different application areas will “go live.”

All this work should culminate in a deployment plan that lists the dates, dependencies, and responsible parties. Make sure that the project team and executive management agree on the plan.

Note:

As you create the plan, avoid the temptation to try to obtain all possible benefits in the initial coarse-grained segmentation deployment. Focus on providing an improved level of segmentation compared to the baseline, and a stable beginning to a label-based, meta-driven model. Your implementation will evolve as the organization develops institutional knowledge and habits around a new way of securing and automating the infrastructure.

Getting Started

Allow extra time for creating and deploying the first dynamic micro-segmentation security policies.

For the team, this is the first step in implementing an entirely new security model. Activating the first policies will be a learning experience. Also, that activity will be happening in parallel with operational integration, label work, application owner education, and other tasks.

For the organization, there will be heightened sensitivity to outages, interruptions, and any mistakes. In most places, the most important consideration will be: “Don’t break anything!”

Given this environment, the best initial policies are rarely the same as the tightest security policies. Start with visibility and monitoring, or with a basic level of enforcement. You can refine and tighten the policies as you gain experience.

Install the Micro-Segmentation Solution

You will need to install and configure the micro-segmentation tools and related infrastructure, and distribute the micro-segmentation software to the monitoring and enforcement points.

Start by setting up a QA or “pre-prod” version of the solution. You can use it to help everyone in the organization become familiar with the software. The QA environment can also be used as a test and development location for automation development, logging integration, and other tasks.

Some pre-production testing, integration, and certification should occur before you deploy the software in bulk. Work through these processes carefully; they are primarily instituted for safety in complex environments, but they also build confidence and expertise across the organization.

Depending on the type of micro-segmentation solution you have, the method you use to distribute the software to the monitoring and enforcement points can be:

- Pre-installing software into operating system images
- Using scripts or orchestration tools to distribute software agents to hosts

Integrate Logs, Events, and Alerts

Before conducting a large-scale production rollout, be sure to fully integrate existing log and event management tools with the operational and security logs of the micro-segmentation solution. You will want full visibility into the environment from the first installation onward.

Integrate your micro-segmentation solution with network devices, data center management tools, and orchestration tools that can help it monitor traffic, detect changes in the infrastructure, and apply policies accurately to new and relocated applications and workloads.

Also, set up reports so you can monitor micro-segmentation-related activities and events and look for anomalies and suspicious behaviors.

Create rules for generating alerts, so you and your SOC or incident response team can respond faster to anomalies and attempts to exploit unnecessary or little-used connections.

Other common tasks to be completed during the operational integration period include:

- Security hardening any micro-segmentation controllers and admin consoles
- Tuning and optimizing controller performance
- Fine-tuning change management procedures to work with the new solution
- Applying role-based access control (RBAC) to administrative and policy-writing accounts
- Building operational run-books and other guidance for ops teams

Prioritize Application Groups

To develop a sequence for staging your implementation of micro-segmentation, first divide your applications into groups. You may want to group applications that are managed by the same team, or related by some combination of factors.

Rank the application groups based on:

- Business value of the data and processes to be protected
- Business value of improving agility and flexibility
- Ease of micro-segmentation

Ease of micro-segmentation depends on factors such as how self-contained the application group is, the number of applications

within it, and the complexity of their architecture. A consultant or micro-segmentation solution vendor can help you make that determination.

You can then assign each application group to a stage in the rollout. The rankings should ensure that the first stages will produce high value and go relative quickly, so your implementation can show early “wins.”

A typical sequence is shown in Figure 5-1.




Figure 5-1: Typical stages in a micro-segmentation rollout plan

Core Services

Every organization has a collection of core networking and computing services. These include DNS and DHCP servers, Active Directory and other organization directories, VPN concentrators, SIEM, log management, network and systems monitoring products, and backup services.

Because these services are essential to running a data center, protecting them provides great value to the organization. They are also good candidates for the initial stage because it is usually easy to identify IT administrators who have a detailed knowledge of their communication flows.

High-Value and High-Risk Applications

The next stage (or stages) of the micro-segmentation rollout typically involves micro-segmenting individual applications that are mission critical or otherwise essential to the operations of the organization. Applications that use or store confidential information (particularly information covered by regulations and industry standards) should also be micro-segmented.

Micro-segmenting shrinks the security perimeter from a subnet, VLAN, or zone to a single application. When all the applications in a subnet have been micro-segmented, the number of open connections in the environment is typically reduced by 90 percent or more.

Micro-segmenting an application provides a great deal of impact for relatively little work. One line of policy per application can close off most of the potential for malicious East-West movement. Application micro-segmentation also promotes business agility. Because all connections can be allowed *within* the microsegment, any changes application owners and developers make to the application's internal communication will always work.




Figure 5-2: Micro-segmentation isolates the application from East-West traffic, while allowing free communication internally. Source: Illumio

The Production Environment

In Chapter 4 we discussed the value of tightly controlling flows across development, test, staging, and production environments. Because environmental separation is relatively easy to achieve and brings significant benefits, it is also a good candidate for an early stage of the rollout.

Most of the benefit comes from isolating the production environment from the others, so you could consider doing that in an early stage and coming back to the other environments later.

Key Projects

Most organizations' IT groups have a few key projects where micro-segmentation can play an important role in improving security, increasing agility, and even ensuring completion on time.

These projects include migrating applications to cloud platforms or across data centers, consolidating data centers, delivering new applications that span cloud platforms and data centers, and using cloud platforms to make applications scalable on demand.

Discovery and Visibility

As micro-segmentation is rolled out to each application group, it is critical to follow a four-step process within that group: discovery and visibility, policy modeling, test, and enforcement (Figure 5-3). Proceeding through the four steps in sequence maximizes your ability to isolate applications and workloads, and minimizes the risk of inadvertently blocking critical connections.




Figure 5-3: Minimize risk by proceeding from discovery and visibility, to policy creation, to policy testing, to enforcement

The first of these steps, discovery and visibility, involves gathering information on:

- The applications and workloads
- The active connections between them

Basic information can be gathered from sources like:

- Configuration management databases (CMDBs)
- Orchestration tools (e.g., Chef, Puppet, Ansible)
- System inventories and IP lists
- Active Directory and other organization directories
- Traffic and event logs
- Firewalls and SIEMs
- Load balancers and network devices
- Informal data stores maintained by administrators (which are usually spreadsheets)

Use the data you collect to catalog and classify the applications and workloads.

The best source of information about active connections may be visibility tools provided by your micro-segmentation solution vendor (see Figures 2-4 and 4-2). An accurate map of these flows is very important because:

- If you block active communication flows, someone (actually, a lot of people) will complain.
- If you leave open too many unnecessary connections (which make up 95+ percent of the potential connections), you defeat the purpose of implementing micro-segmentation.

Model Policies

Policy modeling involves defining applications and workloads, the connections between them, and the rules for allowing or blocking traffic along those connections. Because most micro-segmentation solutions take a whitelist approach, you don't need to create rules to block unnecessary ports.

At one level, the rules may be very similar to conventional firewall rules. But most micro-segmentation solutions (although not all) designate sources and destinations by names or other unique identifiers, not by IP addresses. That means that if an application or workload is moved to a different system or data center, or to a different virtual machine or container, its policies can follow it. It doesn't matter if the IP address changes, or if changes are made to network connections, or proxy servers, or LAN segments; the policies stay with the pair of applications or workloads.

In addition, most micro-segmentation solutions allow rules to be stated in a natural language format. That means that someone can create a rule that says, in effect, "Allow the web front end of the application to connect to the business logic module" without needing any knowledge of the network topology.

Some micro-segmentation solutions allow policy creators to associate labels or tags with workloads and apply policies automatically to all workloads with a specific label or tag, or a specific combination of labels or tags. With that capability, your micro-segmentation solution can apply the same policy automatically to every database server, or to every database server in a certain location (e.g., in North America, or on AWS), or to every database server for the customer service application that is running in the production environment.

That capability also means that the same policy can be applied automatically to new instances as soon as they are created. If your cloud platform starts creating more copies of the customer

service web module to handle more users, the micro-segmentation solution can automatically apply appropriate policies to the new workloads.

Start by creating general policies for all workloads. Even if you have a visualization tool that displays detailed process and port information for every flow in the managed environment, resist the temptation to use all of that information right away. Researching every unknown port and process could require months of effort.

After creating broad default policies across the board, you can circle back and write more granular rules for selected critical systems. Be selective about where to dive deep. Not all applications and workloads are equal, and those with less risk can wait a little bit longer. Those with critical importance should be given more granular security policies first.

Nobody (and No Policy) is Perfect

Unless you have perfect knowledge, you will never have a perfect policy. That's fine because, in practice, all policies are based on a mix of perfect and imperfect knowledge.

Your goal is not perfection, but rather a significant and measurable reduction in attack surface. This goal is attainable, and puts the organization on a path to move from a strong position to a stronger one. Where information is missing, or will take too long to obtain, keep moving — with the confidence that improvement is still occurring.

Test Policies

After modeling policies, it is wise to test them to make sure you don't interfere with applications. There are three common methods:

1. Eyeballing the policy definitions
2. Simulating communication flows for the applications
3. Observing actual application behavior over a period of time

Of these methods, observing actual application behavior is likely to be the most effective, provided the observation period lasts enough to catch unusual but legitimate actions. The eyeballing method can be useful when policies are expressed in natural language, but is less likely to be effective when policies are expressed in technical constructs.

During the testing phase, validate all reporting and event management workflows. Ensure that the operations and security teams receive the reports, dashboards, and alerts they need to perform their jobs. This is also a good time to conduct final operational training for help desk teams and other groups that respond when a problem occurs.

Testing serves other purposes as well. It can turn up unexpected interactions between applications and workloads, and identify probes and indicators of compromise (IoCs) from attackers. Some micro-segmentation solutions allow you to start generating alerts during the test phase so your SOC or incident response team can be notified immediately of unexpected connection attempts.

Be Ready to Fix Problems

At last, it is time to start protecting applications and workloads by enforcing the policies you have created. But before you push the button:

- Notify everyone who might be affected of the time and date when enforcement will start. Take extra time with the application and system administration teams to make sure all of their concerns are addressed.
- Let everyone know how to contact you if anything goes wrong.
- Be ready to fix problems, or even turn off enforcement, if problems occur.

In fact, there are several steps you should take to be ready for the big rollout.

Turn on enforcement in stages. Begin with less-critical applications and workloads, work with them until they are stable, and then move on to more sensitive systems. This approach will minimize the impact of any early issues, and allow the entire organization to adjust and learn how to manage micro-segmentation.

Have a “Tiger Team” ready for rapid response. If you have done a good job of testing, you may not see any issues at first. But when nobody is looking, a communication pattern will change, most likely in the middle of the night. Be sure the security, operations, and application teams are on alert for several days after the move to enforcement.

Train multiple teams to resolve issues. Teach help desk or other “first responders” how to make simple fixes like changing policies to resolve an inadvertently blocked flow and get the application back online. More complicated issues can be escalated to the security or application teams for final resolution. Having multiple teams able to respond helps ensure that new flows and missed rules can be addressed almost immediately.

Refine incident and policy handling processes. As the team remediates policy gaps:

- Application owners will learn that they have more visibility than ever before, but also that they need to communicate changes better.
- The security team will develop a fresh understanding of how critical applications work.
- The operations and security teams will refine procedures for commissioning and testing new systems.

Extend and Refine

Don't disband the project team at the end of the official implementation period. Stuff will keep happening. Put together a post-rollout program that includes plans to:

- Create policies for new applications as they are deployed (try to participate in the planning and development phases of the application lifecycle so there are no last-minute surprises).
- Periodically review each application area and tighten up security policies as needed.
- Circle back to high-value areas where you have only segmented at the application level and drill down to segment servers, tiers, and individual processes (see Chapter 4).

- Monitor communication flows to identify legitimate traffic that was inadvertently blocked.
- Monitor communications flows to identify anomalies and IoCs.

Many teams start with the goal of “going broad” by quickly implementing basic policies across the entire organization. That is a good approach. However, when you progress to the refinement stage you will need to change the organization’s mindset to focus on accuracy instead of speed. That means going a bit slower, being more methodical, and putting more effort into research and coordination.

6

Selecting the Right Solution

In this chapter:

- Examine factors that will help you select the most appropriate micro-segmentation solution

“When you make a choice, you change the future.”

— Deepak Chopra

Which micro-segmentation solution is the best fit for your organization? In this chapter, we look at some of the selection criteria.

Enforcement Points

In Chapter 2 we looked at three different types of enforcement points. There are pros and cons for each.

The first type of micro-segmentation solution uses network devices. This type can allow you to leverage your existing investment in these devices, and to integrate micro-segmentation with other security and management tools from your network equipment vendor. On the other hand, these solutions tend to use more network constructs for their policy modeling, making them harder to learn and less flexible. You are also limited to locations where that vendor's devices are found, which usually do not include cloud platforms.

Other micro-segmentation solutions use hypervisors. These can give you extremely granular segmentation for workloads deployed in virtualized environments, but don't work with software running on physical servers, on other vendors' virtualization platforms, or in the cloud.

A third type of micro-segmentation solution leverages host-based software such as iptables on Linux systems and WPF on Windows servers. These solutions provide extremely granular segmentation for both physical servers and virtual environments, but you must deploy software agents to the target systems.

Visualization

It is extremely important to have good tools for visualizing applications, workloads, and the connections between them. Look for visualization tools that are tightly integrated with the rest of the micro-segmentation solution and that make it easy to:

- Understand the components and tiers of applications
- Visualize connections between applications and workloads
- Drill down to find details about workloads, connections, and policies

- Create and test policies
- Identify suspect connections and find better ways to enforce policies

Policy Modeling

Natural Language

Natural language policies offer several huge advantages:

- Policies can be created faster.
- Policies can be created with fewer errors because their meanings are clearer.
- Application teams can collaborate with technologists to create policies, reducing misunderstandings and errors when converting requirements into rules that can be enforced on the network.

The more natural language, and the fewer network constructs, the better.

Labels and Tags for Workloads

If you can assign labels or tags to workloads, then you can apply the same policies to all workloads with the same label or tag, or the same combination of them. That is a tremendous time-saver, as well as a way of ensuring consistency. Labels and tags enable your micro-segmentation solution to apply exactly the right policy set to workloads created on the fly on virtual platforms.

Labels and tags also allow the same policy to be applied automatically to new instances as soon as they are created in virtual machines. If your cloud platform starts creating copies of

the web module to handle more traffic, the micro-segmentation solution can automatically apply appropriate policies to those new workloads.

Testing

Look for solutions that allow you to test policies by simulating communication flows, observing actual application behaviors, or both.

You will get even more value from the testing tool if it allows you to monitor communication flows on an ongoing basis, detect anomalies, and generate alerts for suspicious events.

Intelligence and Automation

Intelligence and automation increase flexibility and agility, reduce opportunities for human error – and get you out of the office earlier.

Look for micro-segmentation solutions that:

- Discover and catalog applications and workloads
- Monitor active communications flows, and turn that data into intuitive diagrams
- Detect changes in applications and networks, and adapt by intelligently reconfiguring policies
- Automatically distribute policies and policy updates to enforcement points
- Distribute policies to newly created workloads (typically in conjunction with services orchestration tools)

Final Thoughts

There aren't many ways to prevent cybercriminals and hackers from roaming freely inside data centers. There are few good responses to the challenge of making IT security more flexible and agile.

Micro-segmentation provides a solution to both of these challenges. It is one of the few technologies that can improve security and agility at the same time.

We hope this guide has helped you understand why. We've discussed granular segmentation and its benefits. We've outlined dynamic micro-segmentation and its advantages. We've reviewed use cases where micro-segmentation can pay for itself quickly by preventing data breaches and enabling faster innovation. We've recommended implementation steps that can speed up deployment and minimize risk.

We also urge you to try a micro-segmentation solution yourself. A lot of the details of how they work and what they can achieve are only understandable when you see them in action. If you watch a demonstration, or arrange for a trial in your own environment, you will gain a much better grasp of the potential of this powerful technology.

© 2017 Illumio. All Rights Reserved.