

How to Use Micro-Segmentation to Secure Government Organizations



How micro-segmentation reduces your attack surface, hardens your data center, and enables your cloud security.

Overview

Deployed at the network layer, segmentation was first developed to improve network performance. But as cybersecurity experts have become convinced that a “perimeter only” approach to security is not working, it’s become increasingly clear that interior segmentation is also the foundation of data center and cloud security. Organizations looking to secure the interior of their environments often need to be more dynamic and more granular than network segmentation will allow. Micro-segmentation makes this possible.

In this guide, you’ll learn:

- 1 How micro-segmentation fits into your security strategy
- 2 The principles of micro-segmentation
- 3 Five things to look for in a micro-segmentation solution
- 4 Real-world examples to help you plan
- 5 How to get started

Fitting Segmentation Into Your Security Strategy

Segmentation reduces your attack surface, frustrates intruders, and hardens your data center.

Let's talk about breaches.

2,260
breaches in 2015
alone

When intruders breach your perimeter, they most often enter at a low-value environment – your development environment, a contractor's network, or by piggy-backing on your employees' access. To cause damage to your organization, they first have to reach critical data or systems, and to do this they must move laterally through your environment. Unfortunately, intruders can often reach high-value targets today because most data centers are wide open. This means that if intruders find any chink in your armor, they can exploit that to cause harm and steal valuable data.

Micro-segmentation changes the game by helping you stop intruders that break into your low-value environments from making that critical move to your high-value assets. Once an intruder is inside your data center, they are operating on and through hardware, software, and network infrastructure that you control.

Dwell time:

95
Days


Intruders typically:

- manipulate servers, VMs, and containers;
- leverage user accounts to increase privileges or run processes;
- transmit data across network connections between servers.

Every one of those steps uses systems that you control, and that have alarms alerting you to the intrusions so you can stop it immediately. One alarm – a single mistake by the intruder – is all it should take.

But despite this risk, today intruders spend months or years concealed inside compromised networks, regularly reach high-value targets, and often don't get caught until well after the breach. Why is that? Because many

For most organizations,
as few as



of the potential
connections in their
data center are used for
legitimate traffic.

government agencies and private organizations today do little to control the connections in the interior of our data centers and cloud. This lack of control does two things.

First, it makes illicit lateral movement extremely easy, because exposed environments like this have an extremely large attack surface, so intruders have a wealth of attack vectors to choose from.

Second, it makes detection of lateral movement incredibly difficult, because defenders must spread their resources across the entire environment, and have few ways to identify movement, even though it's on their own systems. In fact, even tools that are designed to detect attackers inside environments – from malware detection to behavioral analytics – stumble because they generate thousands of alerts, many of which are false positives, letting intruders hide in the noise.

Micro-segmentation changes this equation by enabling organizations to reduce the number of ways that intruders can reach high-value targets, and by giving defenders a reliable platform to detect lateral movement without drowning in false positives.

Understanding and Shutting Down Your Attack Surface

The attack surface inside your data center is all of the network connections that an intruder can use to move through your environment and reach your high-value assets.

- Potential connection: A potential connection exists to any server with an open port/process from any other server in the same network unless it is explicitly blocked by a firewall. Most data centers have hundreds of thousands or millions of potential connections.
- Active connection: A connection is only active if traffic is currently flowing across it. Only a tiny fraction of potential connections are used for legitimate purposes.

Every open port and active process in an environment offers a potential connection that any other computer within that network can connect to. Legitimate traffic flows across these connections (“active connections”) during the ordinary course of business, but there are far more potential connections than any organization uses. In fact, for many organizations, less than 3 percent of their potential connections are active at any given time.

This means that most agencies could close almost all of their interior attack surface without constraining their operations.

Closing down these unnecessary potential connections:

1. Makes intruders’ job harder by limiting their freedom of movement through the environment and increasing the risk that they set off an alarm;
2. Makes your job easier by limiting the number of attack vectors that you need to focus on, so you can concentrate your other security resources where they will be most effective; and
3. Helps you move quickly to contain intruders when they do get inside, limiting the blast radius and reducing the cost and complexity of incident response and remediation.

What is Micro-Segmentation?

Micro-segmentation eliminates unnecessary network connections within your data center and cloud. It is distinct from network segmentation, which has been around for years and was originally developed to improve efficiency and reduce broadcast domains.

Everyone today is talking about micro-segmentation, but there is relatively little discussion about what micro-segmentation is and how to use it to effectively and practically improve the security of your organization. In fact, micro-segmentation is only one type of segmentation that can be used for security (we’ll discuss this more later). Most organizations don’t deploy segmentation identically across their entire environment – they match different types of segmentation to the different security requirements of the parts of their data center and cloud.

Security Segmentation:

The process of deploying different types of segmentation throughout your environment to increase your security without impacting your business process.

If you're considering using segmentation to improve your security, here are five reasons why you should consider micro-segmentation:

SUPPORTS ALL ENVIRONMENTS AND PLATFORMS

Micro-segmentation work across all your data center and public/private cloud deployments: bare-metal, operating systems, hypervisors, containers, any network – physical or SDN – and any public or private cloud.

APPLICATION-CENTRIC VISIBILITY

Micro-segmentation can build a live application dependency map of how your applications connect and how they communicate. It's the first step to using segmentation to control how they should communicate.

SECURITY POLICY CREATION AND MANAGEMENT

Instead of using traditional firewall rules, micro-segmentation uses high-level, natural language policies to describe desired application behavior – not infrastructure architecture. This lets you consolidate thousands of machine-readable firewall rules into dozens of human-readable policies, which makes compliance easier and empowers your security team to describe and enforce policy across today's increasingly complex, hybrid, distributed, dynamic environments.

ADAPTIVE AND AUTOMATED

Your applications shift constantly, and if your segmentation doesn't adapt to those changes automatically, your security will be out of date within days – or hours. To keep up, micro-segmentation automatically responds to applications autoscaling and moving across your infrastructure to ensure security stays intact.

CUSTOMIZABLE SEGMENTATION

Organizations customize their segmentation depending on the asset that they're protecting. For low-value assets, they might choose environmental segmentation, whereas for higher value assets, they might segment individual applications, or even ports and processes. Micro-segmentation enables you to use different types of segmentation throughout your environment as appropriate.

Building a Micro-Segmentation Strategy in 5 Steps

There are five essential steps to building a security segmentation strategy:

1 Identify high-value assets.

You first need to identify your highest value systems, applications, and data. These could be databases with data of national security value, sensitive personnel data, key applications that run your agency, communications platforms used by your employees for sensitive conversations, or critical industrial systems. Identifying the high-value assets enables you to focus your security efforts on what matters most. You can use fine-grained segmentation to protect these assets. For less valuable assets, more coarse-grained segmentation will be sufficient and less complex to implement.

2 Map your application dependencies.

Map the connections between your workloads, applications, and environments. Legitimate communications between your servers travel across these connections, but attackers can use them as well. Understanding which parts of your network are most connected will help you understand where segmentation can bring you the greatest benefit.

3 Understand the types of segmentation for security.


Your segmentation strategy should apply the right type of segmentation to provide the required security so you'll need to understand your options. There are seven types of segmentation:



a. Environmental segmentation, the coarsest form of segmentation, separates the environments within your data center. It is often used to isolate low-value environments from the rest of your organization, so any intruder that breaches that environment will be prevented from moving laterally to higher value environments. This could also be used to segment systems assigned to different customers, so if one is compromised, the others will remain secure. It provides large decreases in attack surface, and is the easiest form of segmentation to implement. In most cases, one should deploy environmental segmentation across your entire data center.



b. Location segmentation. Depending on your compute architecture, it may make sense to segment your workloads based on the data centers/clouds in which they operate. This could be useful if you operate in countries where you are required by law to store data locally, or if you have a particular data center that holds your most sensitive data, and want to limit the ability of devices from other data centers to access it.



c. Application segmentation, also called application ringfencing, separates individual applications, preventing cross-application communications—even within the same environment. Organizations often use application segmentation to give an added layer of security to their most valuable applications. In environments with many segmented applications, this greatly increases your security and throws up additional roadblocks for an intruder.



d. Tier segmentation is even more fine-grained than application segmentation and divides the tiers within an application (e.g., the web, app, and DB tiers). Because many intruders will first enter data centers via the web tier, this level of segmentation further isolates them, forcing them to cross even more data center segments in their search for high-value data.



e. Process and service segmentation, also called nano-segmentation, is the finest-grained form of segmentation and ensures that only active connections to other workloads are permitted. This fine-grained segmentation is most useful to protect high-value assets where restricting attacker movements is particularly important. No unnecessary potential connections are left open.



- f. **User segmentation**, prevents credential hopping—a common tactic wherein an intruder or insider tries to use acquired credentials that permit them access to a high-value application. This ensures that when a particular user is logged in to a workload, that workload is only permitted to contact servers that the user is permitted to access.

4 Map your segmentation strategy based on your operational security requirements.

You won't use the same segmentation throughout your environment. In general, you'll want to apply more fine-grained segmentation to your high-value locations, and more coarse-grained segmentation to lower value locations. To do this, identify the areas of your data center and cloud that you want to protect first, then assign appropriate segmentation strategies to each one. Micro-segmentation is particularly useful for addressing this stage, because it allows you to deploy segmentation of differing granularities to different parts of your environment.

Set a timeline for the various states of your segmentation strategy. You may decide to begin with the lowest-risk environment first, so you can test out your approach without risking business interruption. Be sure to prioritize those high-value assets you identified in stage 2. Segmenting those assets will give you the greatest security increase for your effort.

5 Test and deploy your strategy.

Since micro-segmentation changes the data center itself, it's essential to make sure the strategy is aligned with the way the data center functions, and isn't breaking anything. The ability to test and model your segmentation strategy before you deploy it is an essential final step to deploying any security strategy.

TIP:
Quantitative and qualitative approaches

Most organizations use a qualitative approach to identify their high-value assets, calculate their attack surface, and then reduce that attack surface through segmentation. Quantifying the attack surface of your different applications and environments will help you develop a micro-segmentation strategy that is optimized for your data center and cloud.


Sample Micro-Segmentation Strategies

For most data centers, we recommend:



ENVIRONMENTAL SEGMENTATION

to wall off the most exposed, least valuable environments (e.g., the development environment).



APPLICATION SEGMENTATION

to isolate applications in high-value environments.



TIER SEGMENTATION

to further protect high-value applications.



PROCESS SEGMENTATION

for core services or other particularly valuable workloads or clusters of workloads.

Here are a few more ways you can optimize your segmentation strategy to secure specific characteristics of your data center and cloud:

- (1) Use environmental segmentation to separate out low-value environments. This enables you to maintain the flexibility of less sensitive environments, but contains their exposure so intruders that enter the environment can't jump from them to high-value targets.

- (2) Segment large applications based on the role or tier of workloads (e.g., segmenting the web, database, and application servers from each other). This approach avoids the complexity of attempting to segment the entire application by workload or process, but still significantly reduces the ability of attackers to move freely through the application.
- (3) Isolate subject-matter data stores from each other within policy-focused organizations. For instance, a law-enforcement agency might segment servers holding data on investigations based on the district of the investigation, and the State Department might segment servers holding data based on their region of the world. By segmenting data-stores based on subject, organizations can prevent the rapid spread of threats throughout policy environments.
- (4) Consider segmenting the communications between servers in different geographic locations. For example, an agency with multiple data centers (for example, U.S. Embassies) in regions around the world could segment these data centers to prevent a local intrusion from quickly spreading to other regions. This could also apply to a domestic agency with multiple locations around the United States (for example, the Transportation Security Administration).
- (5) Cluster heavy processing platforms like Hadoop on a dedicated, non-routable network. “Tier” the application by making the internal processing machines—the true high-value target—accessible only from the external-facing machines, and controlling access to the external-facing machines as you normally would. This forces attackers to take multiple steps to reach the valuable data inside your Hadoop cluster, giving you more opportunities to identify and stop them.
- (6) Use process and service segmentation to protect Active Directory and other core services. Rather than leaving potential connections open for the remaining 50 services exposed, use process and service segmentation to close the connections for all but the services you actually use, and to limit connectivity even for those services in use.



How to Get Started

Segmenting your compute starts with visualization. You must build that map; identify your most valuable assets; and then develop, test, and implement a segmentation strategy to defend them and shut down your attack surface.

Building and implementing a segmentation strategy can be challenging, but Illumio can help. We can visualize your data center and cloud, without needing to install anything. We can build your relationship graph, work with you to develop a segmentation strategy that makes sense for your environment, and then help you implement it.

If you'd like to get started, go to www.illumio.com for more information. Or, better yet, contact us for a [live demo](#).

About Illumio

Follow Us



Illumio, the leader in micro-segmentation, prevents the spread of breaches inside data center and cloud environments. Enterprises such as Morgan Stanley, BNP Paribas, Salesforce, and Oracle NetSuite use Illumio to reduce cyber risk and achieve regulatory compliance. The Illumio Adaptive Security Platform® uniquely protects critical information with real-time application dependency and vulnerability mapping coupled with micro-segmentation that works across any data center, public cloud, or hybrid cloud deployment on bare-metal, virtual machines, and containers. For more information about Illumio, visit www.illumio.com/what-we-do or follow [@Illumio](https://twitter.com/Illumio).

Illumio, Inc. 920 De Guigne Drive, Sunnyvale, CA 94085 Tel (669) 800-5000 www.illumio.com

Copyright © 2018 Illumio, Inc. All rights reserved. This document is protected by U.S. and international copyright and intellectual property laws. Illumio's products and services are protected by one or more U.S. and international patents listed at <https://www.illumio.com/patents>. Illumio® is a trademark or registered trademark of Illumio, Inc. or its affiliates in the U.S. and other countries. To review a list of Illumio's trademarks, go to <https://www.illumio.com/trademarks>. Third-party trademarks mentioned in this document are the property of their respective owners.