



How to Build a Micro-Segmentation Strategy

Micro-segmentation reduces your attack surface and prevents the spread of breaches inside your data center and cloud environments

Overview

Deployed at the network layer, segmentation was first developed to improve network performance. But as cybersecurity experts have realized that a “perimeter only” approach to security is not working, it’s become increasingly clear that micro-segmentation is foundational to 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.

Contents

In this guide, you’ll learn:

Micro-segmentation and your security strategy	03
How to understand and shut down your attack surface	04
Building a micro-segmentation strategy in 5 steps	06
Sample micro-segmentation strategies	08
How to get started	09

Micro-Segmentation and Your Security Strategy

Why micro-segmentation? The perimeter doesn't stop all bad actors from making it inside data centers and cloud environments. Regardless of how many detection technologies organizations use, something is bound to get through – a new virus, phishing emails, or a bad actor working at a company.

Micro-segmentation reduces your attack surface, frustrates intruders, can be used as a compensating control against vulnerabilities, and hardens your data center.



Let's talk about breaches.

When intruders breach your perimeter, they most often enter at a low-value asset or environment – your development environment, a contractor's network, a low-value application, an unpatched server, or connected non-compute devices like HVAC systems or sensors. To cause damage to your organization, they first have to reach critical data or systems, and to do this, they move laterally through your environment. Unfortunately, intruders can often reach high-value targets because connectivity is wide open within most data centers and clouds. This means that if intruders find a way in, they often spend days and even months moving laterally inside data centers and cloud environments, undetected, until reaching their goal to cause harm and steal valuable data.

Micro-segmentation changes the game by helping you stop intruders from moving within your environment. Once an intruder is inside your data center or cloud, they can no longer move freely since micro-segmentation restricts all unauthorized communications.

Intruders typically:

- Manipulate servers, VMs, and containers.
- Take advantage of unpatched or under-protected servers.
- Leverage user accounts to increase privileges or run processes.
- Transmit data across network connections between servers.


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

But despite this risk, intruders spend months or years concealed inside compromised networks, regularly reaching high-value targets, and often don't get caught until well after the damage has been done. Why is that? Because most organizations do little to control the connections in the interior of their data centers and cloud. This lack of control has two major consequences.

First, it makes illicit lateral movement extremely easy because exposed application workloads, especially those with vulnerabilities, have an extremely large and penetrable attack surface. Intruders have a wealth of attack vectors to choose from to gain an entry point in the network. Examples include (but are not limited to) phishing attacks, weak or default passwords, stolen credentials, etc.

Second, it makes detecting 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, causing alert fatigue, and 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. It gives defenders a reliable platform to detect lateral movement without drowning in false positives.



How to Understand and Shut Down Your Attack Surface

The attack surface inside your data center consists of all the network connections that an intruder can use to move through your environment and reach 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. For example, by default, database servers open many different ports for services. If certain services are not being used, the open ports associated with these services provide an entry point for unauthorized users. This is not unique to databases and happens with other core services like domain controllers. Most data centers have hundreds of thousands or millions of potential connections.
- **An active connection:** a connection is active if traffic is currently flowing across it. Only a tiny fraction of potential connections is used for legitimate purposes.

Both potential and active connections are burdened with another factor: vulnerabilities. Hosts with vulnerabilities make it even easier for a bad actor to move from workload to workload since they can exploit potential and active ports more easily.

Every open port and active process in an environment is a potential connection that any other computer within that network can connect to. Legitimate traffic flows across these 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 organizations could close almost all of their interior attack surface without disrupting their operations.

Closing these unnecessary connections:

1. Makes an intruder's job harder by limiting their freedom of movement through the environment and increasing the likelihood that they set off an alarm.
2. Can be a way of implementing compensating controls to reduce the exposure of vulnerabilities.
3. Makes your job easier by limiting the number of attack vectors that you need to focus on so you can concentrate security resources where they will be most effective.
4. 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 is a security technique that enables assigning coarse- to fine-grained security policies to data center and cloud applications, down to the workload level. This approach lets you deploy security models deep inside a data center using a software-only approach.

If you're considering using micro-segmentation to improve your security, here are six important capabilities to look for in a solution.

Support for all environments and platforms

A micro-segmentation solution should work across all your data center and public/private cloud deployments: bare-metal, operating systems, hypervisors, containers, any network – physical or SDN. This lets you make infrastructure choices without being restricted by your micro-segmentation solution, and it lets you centralize policy management across all environments.

Application dependency mapping

It should provide a live application dependency map that shows how your applications connect and communicate. This is the first step to using micro-segmentation — it enables you to understand application dependencies and model security policies so you can effectively control communications.

Vulnerability mapping

It should include integrated vulnerability and threat data to show potential attack paths in real time. Vulnerability maps help application security, vulnerability management, and segmentation teams drive strategies to prioritize patching or, when patching is not possible, use micro-segmentation as a compensating control to eliminate unnecessary attack surface and reduce the exposure of vulnerabilities.

Security policy creation and management

Instead of using traditional firewall rules, a micro-segmentation solution should use high-level, natural language policies to describe desired application behavior – not infrastructure architecture, essentially decoupling network constructs from security policies. This lets you consolidate thousands of machine-readable firewall rules into dozens of human-readable policies based on labels, making compliance easier and empowering your security team to describe and enforce policy across today's increasingly complex, hybrid, distributed, dynamic environments. The solution should be scalable to accommodate large enterprises that have hundreds of thousands of workloads across the globe, with the ability to federate policies across multiple regions into a single administrative domain.

Adaptive and automated

Your applications shift constantly, and if your micro-segmentation doesn't adapt to those changes automatically, your security will be out of date within days – or hours. To keep up, your micro-segmentation solution should automatically respond to applications auto-scaling and moving across your infrastructure to ensure security stays intact and moves at the speed of the business.

Customizable micro-segmentation

Micro-segmentation policies should be customizable throughout your environment based on the asset you're protecting. For low-value assets, you might choose coarse-grained environmental segmentation, whereas for high-value assets, you might segment individual applications or specific application tiers (web front-end, application processing, database) with more granular policy or even with policy tied to specific ports and processes. Vulnerability-based micro-segmentation can be used to create compensating controls around vulnerabilities by constraining communications to only those flows that are required for operation, blocking paths that are not in use. Using a live application dependency map, you can ensure that micro-segmentation doesn't break your applications.

Building a Micro-Segmentation Strategy in 5 Steps

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

1. Identify high-value assets.

You first need to identify your highest value systems, applications, and data. These could be key applications that run your organization, 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 and enrich with vulnerability data.

Map the connections between your workloads, applications, and environments and add vulnerability scan data to see connectivity to vulnerable ports or to see if you have an exposed vulnerability with no traffic. 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, vulnerable, and most exposed 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 several types of segmentation:

- a. **Vulnerability-based micro-segmentation** ties real-time vulnerability data and application traffic to micro-segmentation policy. Vulnerabilities can be mitigated by using micro-segmentation as a compensating control that reduces East-West exposure and prevents the spread of breaches.
- b. **Environmental micro-segmentation**, the coarsest form of segmentation, separates 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, others will remain secure. It decreases attack surfaces significantly and is the easiest form of segmentation to implement. In most cases, you should deploy environmental segmentation across your entire data center.
- c. **Location micro-segmentation** separates your workloads based on the data centers and 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, you want to limit the access from other data centers.
- d. **Application micro-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.

- e. **Tier micro-segmentation** is more fine-grained than application segmentation. It divides the tiers within an application (e.g., the web, app, and database tiers). Because many intruders 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.
- f. **Process and service micro-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.
- g. **User micro-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 micro-segmentation strategy based on operational security requirements.

You won't use the same micro-segmentation throughout your environment.

In general, you'll want to apply more fine-grained segmentation to your high-value locations and more coarse-grained micro-segmentation to low-value locations.

To do this, identify the areas of your data center and cloud that you want to protect first, then assign appropriate micro-segmentation strategies to each one.

Vulnerability maps help drive patching and the use of micro-segmentation as a compensating control until you can patch. Even if a bad actor penetrates your perimeter, reducing the exposure of vulnerabilities dramatically reduces their ability to move within the environment.

Set a timeline for the various states of your micro-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 and areas that you identified as most vulnerable in steps one and two. Micro-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 and cloud 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




Environmental micro-segmentation

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




Process micro-segmentation

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



Tier micro-segmentation

Further protect high-value applications.



Vulnerability-based micro-segmentation

Use as a compensating control to reduce exposure of vulnerabilities.



Application micro-segmentation

Isolate applications in high-value environments.

For most data centers, we recommend:

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

- Use environmental micro-segmentation to separate out low-value environments. This lets you 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.
- Micro-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.
- Consider micro-segmenting the communications between servers in different geographic locations. For example, an organization with multiple data centers around the world (such as a global law firm) could segment these data centers to prevent a local intrusion from quickly spreading to other regions. This approach can also be used to address data localization requirements or the regional data controls imposed by the EU's General Data Protection Regime (GDPR).
- 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 targets – 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.
- Use process and service micro-segmentation to protect Active Directory and other core services. Rather than leaving potential connections open for the remaining services that are exposed, this technique closes connections for all but the services you actually use. It also limits connectivity, even for those services in use.
- Use vulnerability-based micro-segmentation as a compensating control to mitigate vulnerabilities, reduce East-West exposure, prioritize patching efforts, and prevent the spread of breaches.

How to Get Started

Micro-segmenting your environment starts with visualization. You must build that map, identify your most valuable assets and those that are the most vulnerable, and then develop, test, and implement a micro-segmentation strategy to defend them and shut down your attack surface.

Building and implementing a micro-segmentation strategy can be challenging, but Illumio can help. We can visualize your data center and cloud – without re-architecting the network. We can build your relationship graph and vulnerability map, work with you to develop a segmentation strategy that makes sense for your environment, and then help you implement it.

To get started, go to www.illumio.com for more information. Or, better yet, contact us for a [free trial](#).



Illumio enables organizations to realize a future without high-profile breaches by preventing the lateral movement of attackers across any organization. Founded on the principle of least privilege in 2013, Illumio provides visibility and segmentation for endpoints, data centers or clouds. The world's leading organizations, including Morgan Stanley, BNP Paribas, Salesforce, and Oracle NetSuite, trust Illumio to reduce cyber risk. For more information, visit www.illumio.com/what-we-do.



See what customers have to say about Illumio.

The GARTNER PEER INSIGHTS Logo is a trademark and service mark of Gartner, Inc. and/or its affiliates and is used herein with permission. All rights reserved. Gartner Peer Insights reviews constitute the subjective opinions of individual end users based on their own experiences and do not represent the views of Gartner or its affiliates.

Illumio, Inc. 920 De Guigne Drive, Sunnyvale, CA 94085, Tel (669) 800-5000, www.illumio.com. Copyright © 2020 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.

Follow us on: