illumio

WHITE PAPER

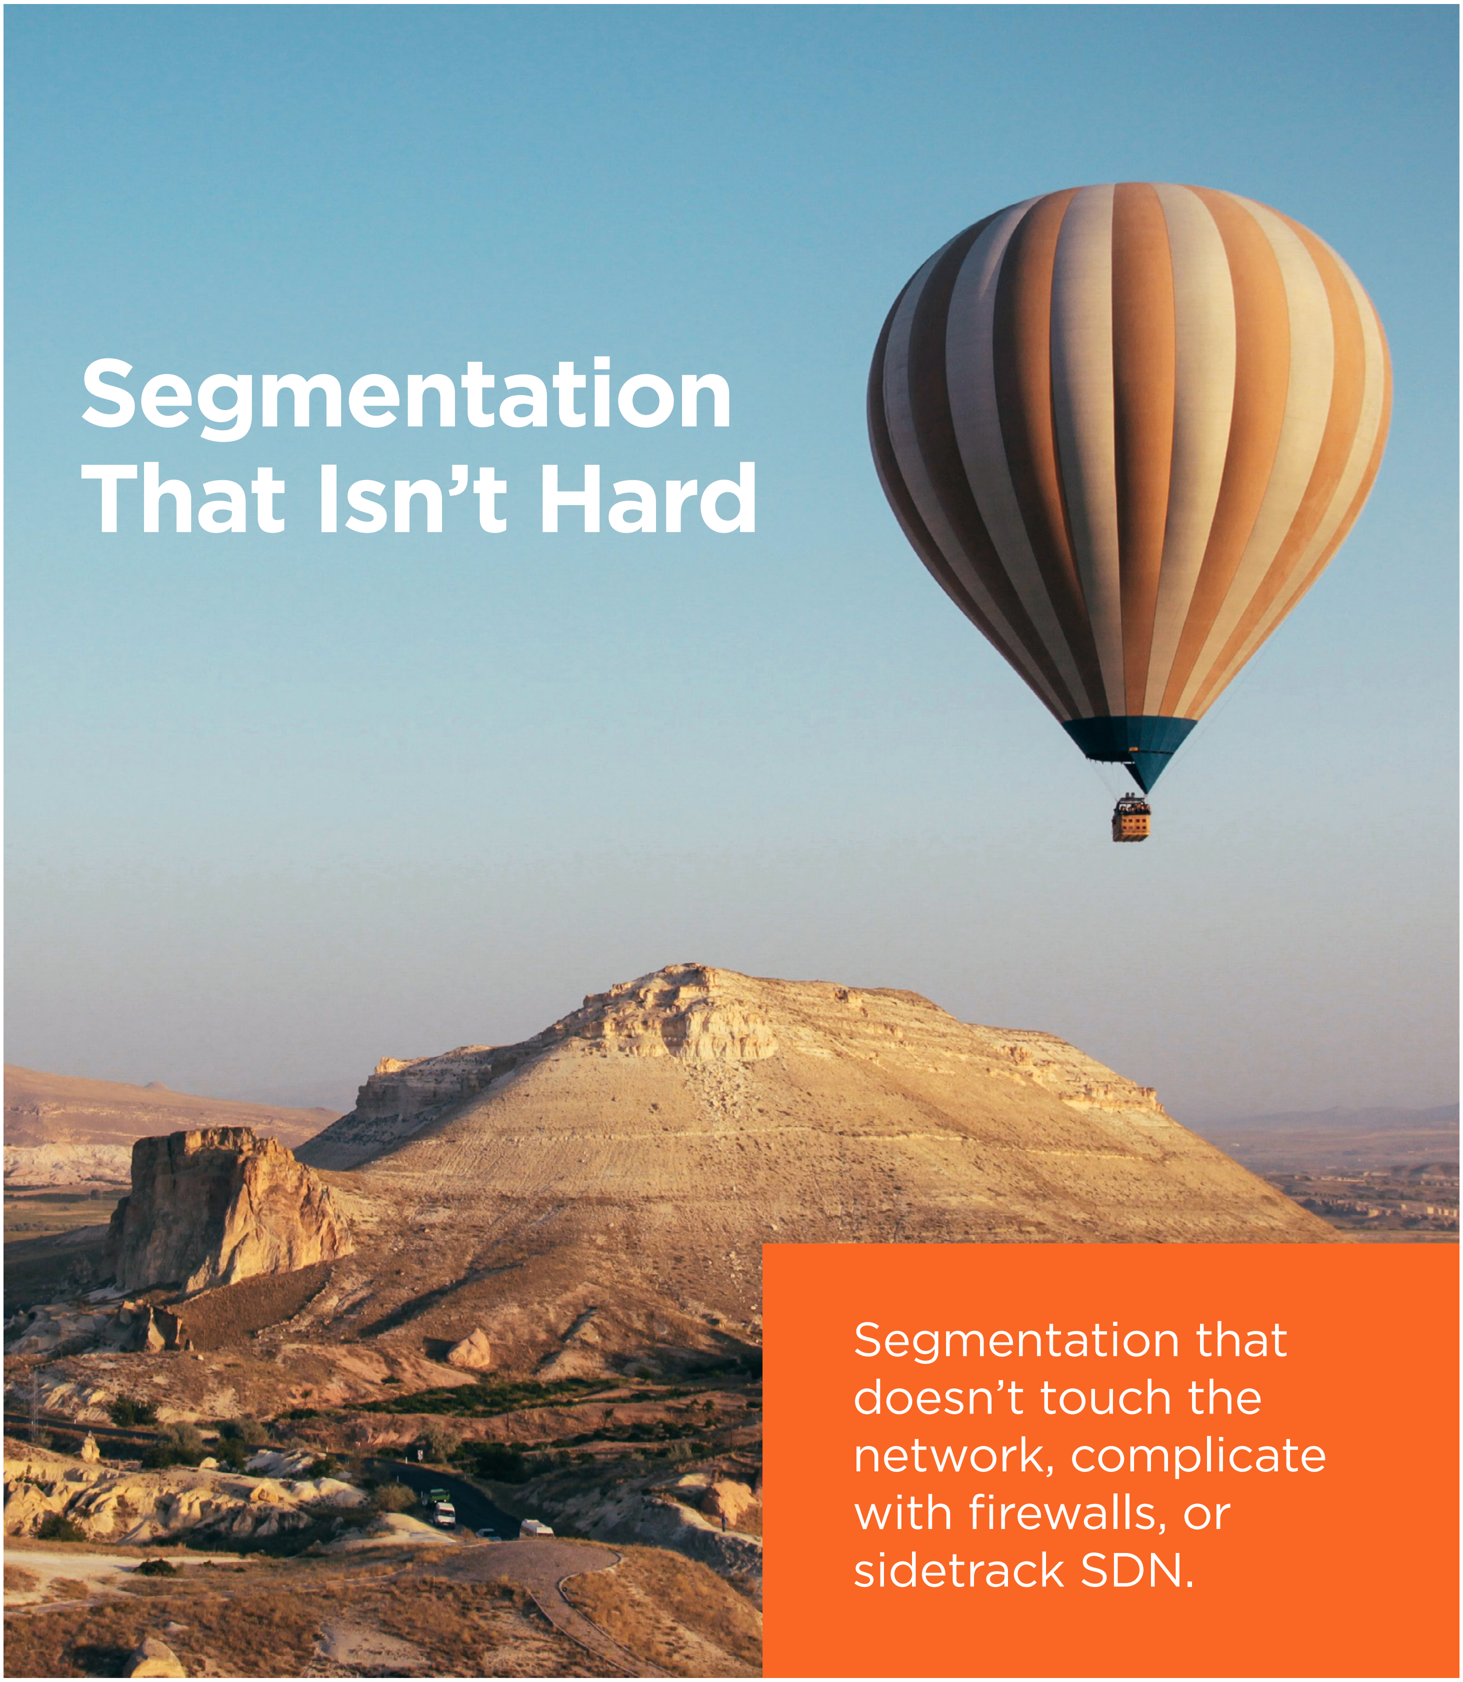

Segmentation That Isn’t Hard

Segmentation that doesn’t touch the network, complicate with firewalls, or sidetrack SDN.

Today's Challenges

It was an unfortunate triple whammy when a Florida municipality was hit with the Emotet trojan, that downloaded TrickBot, and later, downloaded Ryuk ransomware.

After being downloaded, TrickBot moved laterally throughout the municipality calling on exploitation of SMB vulnerabilities, brute forcing RDP (Remote Desktop Protocol), or via network shares.

This lateral movement is how a security incident on a single infected endpoint or workload turned into a devastating breach that drove headlines and resulted in ransom payments.

An occasional security incident is inevitable. However, breaches do not have to be. Motivated attackers will eventually find their way in. It might be a piece of novel malware attached to a well-written phishing email or containers left exposed to the internet.

This is a moment of truth. Stopping attackers moving laterally is in effect stopping attackers from pursuing their aim: a full-blown breach.

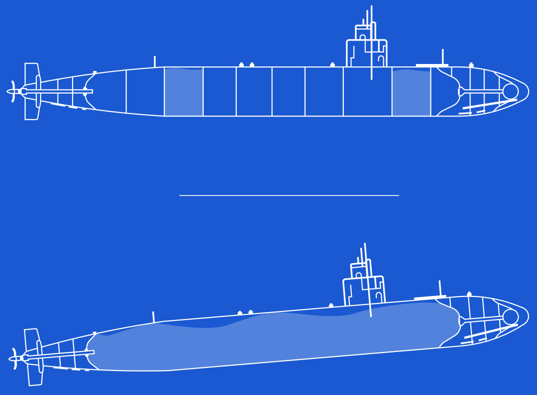

Why do we segment?

A submarine uses compartments to remain seaworthy in the face of a breach in one compartment. The breach is contained to that single compartment and the submarine does not sink – thanks to the compartmentalization (or segmentation) of the vessel.

A Tough Needle to Thread

How have we sought to stop this lateral movement?

Segmentation is a concept that has been around as long as we've been connecting networks. It offers better network performance through smaller broadcast domains and better security through smaller attack surfaces.

Let's not forget the primary purpose of a network: reliable, utility-like packet delivery to support a business and its applications. Segmentation, however, is about reliably separating and filtering to block unauthorized traffic from going where it should not.

Segmentation to date has programmed a firewall or the network infrastructure to understand what can connect to what and block everything else. Like a bouncer at the club, if you're not on the guest list, you won't make it past.

However, balancing packet delivery and security on the same network has been a tough needle to thread. On one hand, we seek reliable, utility-like packet delivery to run the business and, on the other, segmentation that blocks traffic to prevent lateral movement and breaches. As we attempt to thread this needle, the risk of mistakes or misconfigurations is very high.

This holds true even for software-defined networking (SDN). Similar to traditional networks, SDN is primarily designed for reliable packet delivery – not for enforcing the security of what should and shouldn't be allowed between two points on the network.

Hybrid and Ephemeral

Almost every organization has on-premises data center workloads. However, it turns out that most of those same companies also have cloud workloads working in conjunction. Naturally, many workloads in the cloud are containerized workloads given the speed and flexibility they offer for DevOps, coming and going as needed to support applications.

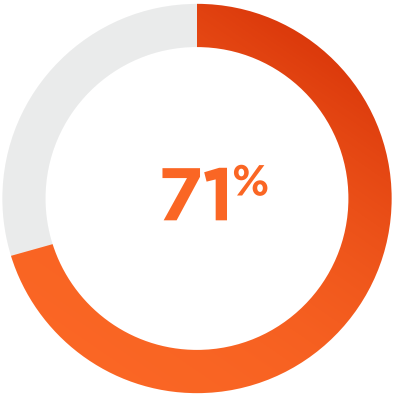

"71% of organizations rely on data centers and clouds working together"

- The State of Segmentation

Even if you can make segmentation work well enough on the network, it must also account for public clouds, third-party services, and APIs, making segmentation more challenging as we protect everything from attackers.

Approaches to Date: Hard

Why doesn't every organization have segmentation in place already? It is because there have been a number of common ways to do it, and they have often proven hard to implement and manage.

Segmentation via the network: it's very manual

Traditional segmentation began on the network, deployed through virtual LANs (VLANs) or subnets, relying on IP addresses to partition a network into smaller subnets. When we want to filter traffic between VLANs or subnets for segmentation, we introduce access control lists (ACLs) in the network infrastructure.

This approach will contain threats from spreading beyond a particular VLAN or subnet but creating ACLs has always been a manual effort requiring intimate knowledge of the traffic. If an ACL is added without proper scrutiny, a misconfiguration can inadvertently break an application when traffic cannot traverse a network control point, ruining reliable packet delivery and disrupting business.

The time it takes to write, approve, and provision ACLs is too slow for business today. If a container is spun up and down in seconds, why does a new ACL on the network take days or weeks? Did we mention that troubleshooting misconfiguration of ACLs is quite an undertaking?

> If a container is spun up and down in seconds, why does a new ACL on the network still take days or weeks?

Network segmentation does not adapt easily to change because networks are hard to re-architect to adjust. Reconfiguring a server or deploying a new subnet could take weeks due to the complexity of IP addresses.

Now businesses always want IT to operate and deploy faster, but segmentation is often too unwieldy, leaving us with security that slows down business.

Segmentation via firewalls: wrong tool for the job

Instead of using the network to enforce segmentation, deploying firewalls inside the data center is another option. Firewalls have been used to create useful coarse zones between areas like the campus user zone and the data center, but are sometimes stretched to segment traffic inside data centers to filter traffic between hosts.

On one hand, most IT teams are familiar with 5-tuple firewall rules from being deployed at the perimeter. However, segmentation becomes considerably more complex when the same firewalls are used for granular, internal micro-segmentation between hosts – often requiring a virtual firewall on every host, resulting in thousands of firewall rules overall. When we consider a landscape of static firewall rules trying to keep pace with dynamic cloud workloads and changing IP addresses, it really gets complex.

Firewall misconfigurations, just like on the network ACL misconfigurations, can break an application and harm business. Firewalls are pricey, whether a virtual firewall for every host or occasionally physical firewalls, bought in pairs, usually for multiple sites and often costing in the millions of dollars.

While firewalls are effective at separating trusted internal networks from the outside world and creating coarse internal zones, this does not mean they are the best tool for the job for granular, micro-segmentation.

We must wonder if it isn’t time to fire the firewall from its attempts at micro-segmentation and leave it at the perimeter for threat protection or to create coarse zones inside organizations. At the very minimum, call in some help for the firewall if it must attempt micro-segmentation to track workload to IP mappings.

Let SDN focus on its day job

Software-defined networking (SDN) is relied on for greater network automation and programmability through centralized controllers that are abstracted from the physical hardware of the network.

With SDN, we can deploy applications rapidly without having to think too much about the network.

SDN adds another layer of complexity because it relies on underlays, overlays, and tunnels to work. But as with traditional networks, SDN is ultimately tied to the infrastructure it resides on – the hypervisor or routers and switches.

Some network operators seek to coax segmentation from their SDN network overlay implementation by using it to create policies to funnel packets through a distributed set of firewalls.

What about working from a map that shows what must be protected? Automated segmentation policies? Cloud workloads? These are all shortcomings of SDN when it comes to segmentation. It is challenged to deliver visibility and consistent segmentation policy with workloads in multiple clouds without additional SDN technologies (or another segmentation solution) and more elbow grease. To a large extent, SDN technologies are useful up to the edge of their specific fabrics.

> SDN’s segmentation promise: All of the complexity, none of the visibility?

When we decide to push SDN beyond its core networking function, it is no longer in its element.

SDN should focus on its day job, network automation, rather than taxing it with something it was not designed for and is not adept at: segmentation across clouds and data centers.

Micro-segmentation easily picks up the security that SDN isn’t equipped to handle, making for a very productive co-existence of the two technologies.

Conscious Decoupling from the Network

To address these issues, we need a segmentation solution that is close to what’s being protected: the workload and its applications. This is why we must intentionally decouple security from the network with micro-segmentation (also referred to as host-based segmentation or security segmentation).

Micro-segmentation protects by using host workload controls, instead of the network, firewalls, or SDN.

Each workload operating system in the data center or cloud contains a native stateful firewall, such as IP tables in Linux or Windows Filtering Platform in Windows. Micro-segmentation manages and programs these at scale to enforce segmentation. It first uses their telemetry to build a map of the entire compute environment to then build automated segmentation policies. It also calls on easy-to-understand labels for policy instead of IP address complexity.

What happens when we segment without the network for enforcement? We get segmentation without complexity. Without expense. Without misconfigurations. Let's look at advantages of a decoupled approach.

Micro-Segmentation That Isn't Hard: Enter Illumio

Illumio has considered all of these challenges when approaching micro-segmentation, offering a modern solution that makes segmentation easier.

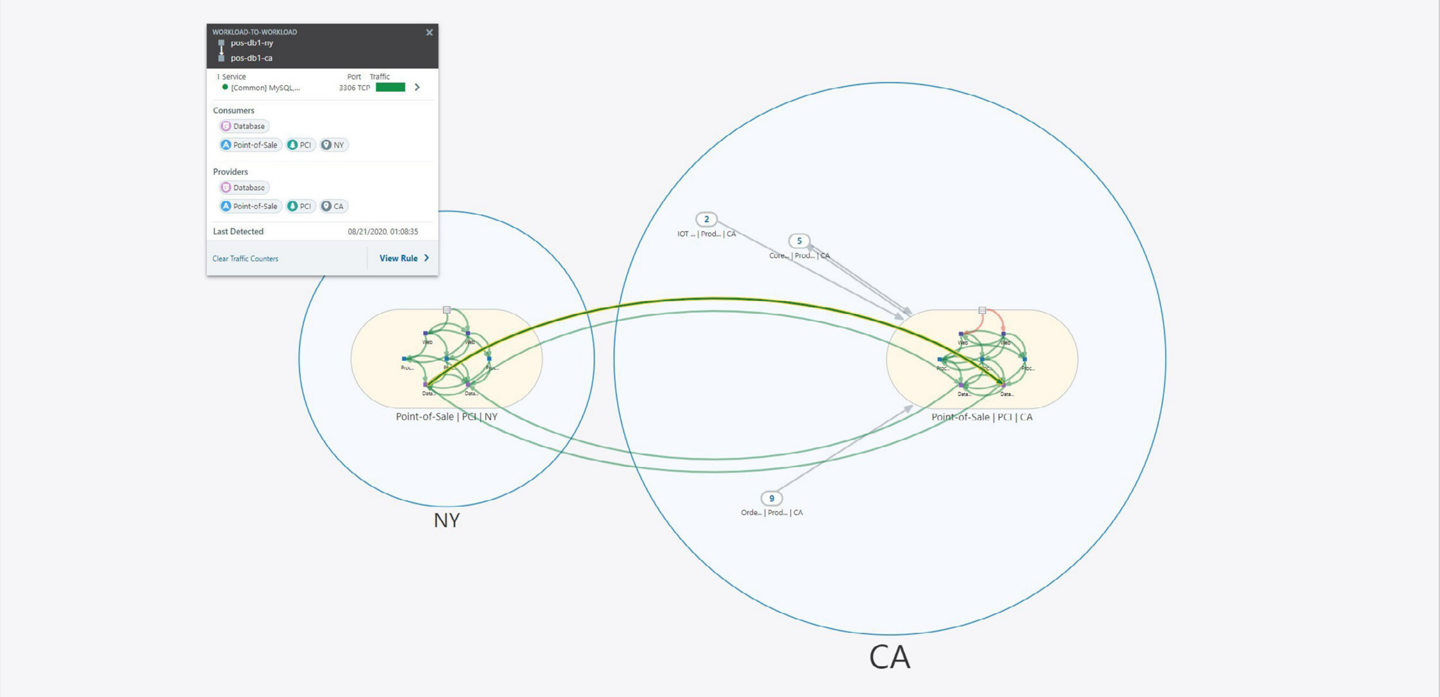

WORKLOAD-TO-WORKLOAD ×

pos-db1-ny

pos-db1-ca

1 Service	Port	Traffic
● [Common] MySQL,...	3306 TCP	🟩 >

Consumers

Database

Point-of-Sale | PCI | NY

Providers

Database

Point-of-Sale | PCI | CA

Last Detected 08/21/2020, 01:08:35

Clear Traffic Counters | View Rule >

Easier segmentation that starts with a map

Host-based segmentation uses workload telemetry to create a real-time map of cloud and on-premises compute environments and applications. This map is used to visualize application connectivity, allowing teams to clearly see what they must protect.

An advantage of using the host is the ability to see and enforce segmentation down to the process level, more granular than just specific ports. Permitting only specific services between particular workloads is true micro-segmentation.

Another key difference is that micro-segmentation uses human-readable labels – not IP addresses or firewall rules – to create policy. Illumio assigns four-dimensional labels to workloads (bare-metal servers, VMs, containers, or processes running on hosts) to identify and provide context for each workload: role, application, environment, and location.

With those labels, segmentation policy becomes as fast as a few clicks.

Faster segmentation

What would you rather do: write firewall rules for 70 hours or take the better part of a two-week vacation?

Fortunately, with micro-segmentation you don't have to choose since it is orders of magnitude faster. For example, to segment a single application, firewall-based segmentation needs 20 minutes of flawless, uninterrupted rule writing. Micro-segmentation? 20 seconds, maybe less. And it delivers even more granular segmentation.

When we scale firewall rule writing for segmentation out to 1,000 workloads, it amounts to writing segmentation policies for some 70 hours, non-stop.

How does Illumio make segmentation so much faster and easier? With our labels in place, we can automate policy creation, to make it a matter of a few clicks. Workflows make micro-segmentation as simple as selecting the granularity (or level of restrictiveness) of your segmentation policy. You can define traffic restrictions for workloads at the environment level (least granular), application level, role/tier level, or even by the process/service running on individual workloads (most granular).

By the way, initial micro-segmentation deployments are faster. They are measured in weeks, never quarters or years.

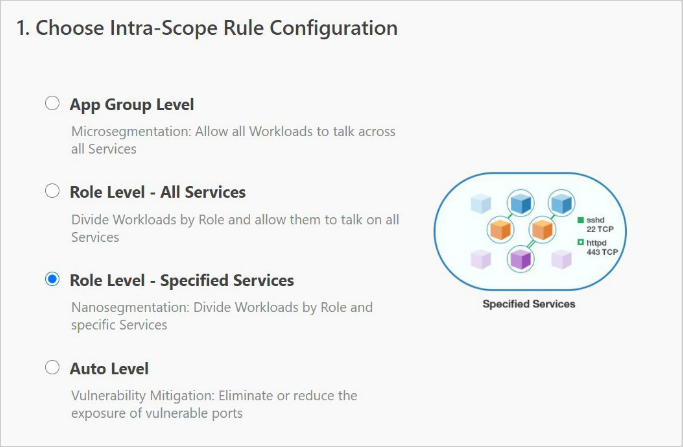

1. Choose Intra-Scope Rule Configuration

○ **App Group Level**

Microsegmentation: Allow all Workloads to talk across all Services

○ **Role Level - All Services**

Divide Workloads by Role and allow them to talk on all Services

● **Role Level - Specified Services**

Nanosegmentation: Divide Workloads by Role and specific Services

○ **Auto Level**

Vulnerability Mitigation: Eliminate or reduce the exposure of vulnerable ports

Safer segmentation with more uptime, less risk

Have you ever worried that a misconfigured segmentation rule on a firewall might break an application? You are not alone. Micro-segmentation makes misconfigurations a thing of the past with policy states that let you test policies before enforcement. The result? 100 percent confidence in segmentation and application uptime.

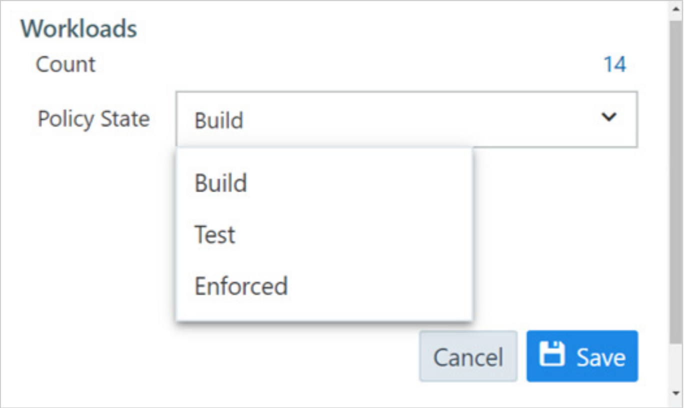

Workloads

Count 14

Policy State: Build

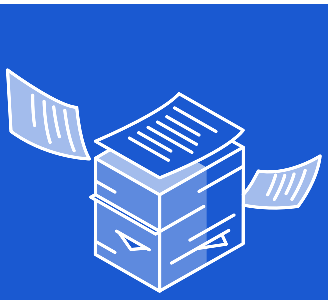

Would you rather?

Write firewall rules for 70 hours or take the better part of a two-week vacation?

Segmenting 1,000 workloads can take 70+ hours using firewalls.

More cost-effective segmentation

If you turn to firewalls for internal data center segmentation, you have to ensure the firewalls support the right amount of east-west throughput. You buy them in pairs, often for multiple locations. It all adds up. For larger organizations, this runs into the millions of dollars. Micro-segmentation software, pound for pound is much less expensive than firewalls. And it allows you to segment per workload. This means organizations only pay for the workloads they segment. For example, they can start small with a compliance initiative and only segment those workloads. They grow their deployments from there.

More scalable segmentation

Micro-segmentation can easily segment up to 100,000s of workloads, while keeping simplicity in place thanks to effective workload labels. You can't do that with a firewall.

The 22x Security Advantage

Why is segmentation worth it? It offers a better security posture that prevents attackers from reaching crown jewels. Red team specialists Bishop Fox recently set out to examine how effective micro-segmentation is, as it relates to "the generation of detectable events and time investment required for an attacker to traverse the network."

In a 1,000-workload environment, they concluded micro-segmentation makes it 22 times more difficult for an attacker to move laterally and reach crown jewels, dramatically deterring bad actors from reaching their target.

With this segmentation in place, it is easier to identify attempts by malware to either scan the network or to move laterally, with alerts that are generated as part of a security team's SIEM workflow.

BISHOPFOX

EFFICACY OF MICRO-SEGMENTATION
ASSESSMENT REPORT
JUNE 04, 2020

Discover more findings in

Efficacy of Micro-Segmentation: Assessment Report

SEGMENTATION THREE WAYS	Network/ Firewall	SDN	Micro-Segmentation
Application dependency map of data center and cloud	○	○	●
Easy-to-use workload labels	○	◔	●
Environment, tier, or application segmentation	◕	◕	●
Process-level segmentation	○	○	●
User segmentation	◑	◑	●
Holistic cloud and container support	◔	○	●
Start small, per workload deployment	○	○	●
Test policy before enforcing	○	○	●
Network / infrastructure independent	○	○	●
Cost	$$$$	$$$	$
Misconfiguration risk	High	High	Low
Number of policy rules	Many	Many	Few

◔ ◑ Partially Supported

Segmentation Is No Longer Hard

Many things in life are hard. Fortunately, segmentation is no longer one of them.

We rely on the network to deliver applications, but we have determined that the network is not the best option for delivering segmentation. The network cannot provide an interface to visualize and understand the connectivity of applications in order to design and maintain granular segmentation that protects them.

The network lacks the agility to adapt to change, and even with SDN, it is tethered to infrastructure that cannot adequately scale to keep up with the business's need for speed. The answer is to decouple segmentation from the network, not complicate with firewalls, or sidetrack SDN.

This allows us to protect applications wherever they run – because they do not live exclusively on our networks anymore, and enforcement must go wherever they do. Micro-segmentation ensures we do not miss the moment of truth: stopping attackers or malware moving laterally in an attempt at a headline-driving breach.

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

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.