

Decoupling Security from the Network: The Evolution of Segmentation

Overview

Segmentation has been around as long as we've been connecting networks, beginning from the earliest TCP/IP protocols designed to reliably deliver packets. But networks are about *connecting* things with utility-like reliability – whereas segmentation is about reliably *isolating* things.

Goals of this white paper:

- Explain the evolution of networks, SDN and host-based “security segmentation”
- Describe core competencies and trade-offs of segmenting on each
- Build a case for decoupling security from the network based on design logic and efficacy

Segmentation understands what can connect to what and enacts enforcement rules to limit everything else - like a bouncer at the club, if you're not on the guest list, you won't get in. These two objectives are diametrically opposed. Yet, we try to do both with the same equipment.

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

And even if you can make segmentation work with your network, the IT environment has grown beyond the data center to include public clouds, third party services and API's. Our environments are not only on the corporate network. The agile infrastructure necessary for DevOps means that workloads are dynamic, and certain application components are not inside the datacenter.

Endpoints are dynamic, too. What's needed is to secure closest to what's being protected. This requires us to decouple security segmentation from the network.

Enterprises are steadily moving to host-based segmentation to address these issues with traditional approaches. Before we can understand why they are turning to host-based segmentation, let's discuss how they got there, and why they're decoupling security segmentation from the network.

Segmentation on the Network

It's Very Manual

Traditional segmentation began on the network, and will always be deployed there through the use of virtual LANs (VLANs). When we want to filter traffic between VLANs, we introduce access control lists (ACLs).

Developing ACLs is a manual effort requiring intimate knowledge of the traffic – and when an ACL is added, if something has been overlooked, it can inadvertently break something, thereby ruining reliable packet delivery and disrupting applications. Therefore the time it takes to write, approve, and provision ACLs is painstakingly slow. Furthermore, when something *does* break, troubleshooting misconfiguration of ACLs is quite an undertaking.

Network segmentation does not adapt to change because classic networks cannot be easily re-architected to adjust. In a traditional network, reconfiguring a server or deploying a new subnet could take weeks to re-architect the network due to the complexity of IP addresses. Now consider that businesses always want IT to operate and deploy *faster and with more agility*.

Segmentation with Software-Defined Networking

The Need for Speed

With the rise of virtualization, speed became even more critical to IT and the business's competitive edge. This brought on the rise of the empowered developer who wants to be able to deploy applications rapidly without having to think about the network. But you still need to connect things, and you still need IP addresses.

The speed of IP address assignment became an obstacle. For example, Class C IP addresses (typically assigned one or two per physical box) quickly depleted with the proliferation of virtual machines. When IP addresses are depleted and more address space is needed, reconfiguring the network via traditional methods to take advantage of unused addresses is impossible.

Enter **software-defined networking (SDN)**.

Just like Uber allowed people to use idle cars, SDN is like an “Uber for IP addresses.” SDN increases the efficiency of unused IP addresses, allowing workloads and applications to be deployed faster. Tools like VMware NSX and Cisco ACI enable the network to be more agile by programatically making adds and changes - while maintaining the network’s core competency of reliable packet delivery. Yet fundamentally SDN is not a significant departure from the limitations of the network when it comes to segmentation.

SDN Segmentation: All of the Complexity, None of the Visibility

Software-defined networking vendors also try to use their products for security segmentation. SDN vendors create an overlay of networks that funnel packets through a distributed set of firewalls – or use the network itself for stateless filtering. The problem with SDN is that it 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.

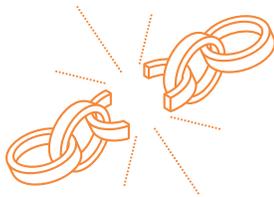
With the growth of public cloud adoption, workloads outside the traditional network are often of equal importance. These present a challenge to visibility and consistent enforcement of segmentation policy since your workloads may sit in a network that you don’t own.

Compounding this, enterprises typically leverage multi-cloud deployments for resiliency, to gain efficiencies and prevent vendor lock-in. Without a means of central management across the organization's different environments, it's not uncommon to see multiple separate segmentation solutions – none of which overlap or provide comprehensive visibility.

Security Limitations of SDN

Any network-based solution struggles to provide visibility and consistency within its own network, never mind the public cloud infrastructure where most enterprises are now running at least a portion of their IT. And once again, when core function goes beyond packet delivery to security, networks are not in their element.

Networks should be built for reliability. Isolating applications through network segmentation can potentially break the network especially if there is a lack of comprehensive visibility – thus creating a [Kobayashi Maru](#); a no-win situation for networking teams.



Segmentation: Decoupling Security From the Network

Traditionally when we planned travel, we thought of crossing the land (or sea). Both of these create friction, which ultimately slows us down. Sometimes the best answer to a problem requires a different approach. By instead moving through the air, decoupling travel from land, it allowed us all to move more quickly, without the resistance created by earth. That's why, when we need to cross an ocean, we switch from a car to a plane – decoupling travel from land. It gets us from point A to point B faster. It lets us go off-road to places cars can't operate. It's less manual since there is a pilot to "drive" us. It's safer – there are far fewer accidents in flight compared to driving. As a bonus, it gives us a bird's eye view of the terrain from the air, a perspective you only get at 30,000 feet.

Similarly, one solution to the restrictions of coupling security segmentation with network segmentation is simply to decouple them. It's safer – because enforcing security segmentation has no impact on the network, and therefore it cannot break the network. It's faster – because any small change in a 'coupled' network requires intensive planning (if not re-architecting). It turns out to be less “accident-prone.” It also gives us the agility to go off the grid across networks or clouds.

Host-based Segmentation: Enforce Close to the Application

How can we decouple security segmentation from the network? First, remember that the network is there to serve applications – applications are not there to serve the network.

Not only does removing security segmentation from the network free security from the limitations of the underlying infrastructure, it allows us to enforce security policy closest to what is being secured: the application. Applications consist of workloads processed at a host, so if you can secure the workloads, you remove any dependency on the network (other than reliable packet delivery).

Host-based segmentation has evolved to fill this gap, supported by a number of vendors including Illumio, Cisco, VMware, and others. In host-based segmentation, an agent enforces security policy by orchestrating native firewalls built into the major operating systems, primarily iptables on Linux systems and the Windows Filtering Platform (WFP) on Windows servers. These utilities monitor network traffic and enforce security rules for specific applications and workloads at the host, providing granular segmentation.

Network Segmentation

Segment through switches and firewalls with VLANs, ACLs, and Zones

Host-based Segmentation

Segment at the host – closest to the application it's protecting

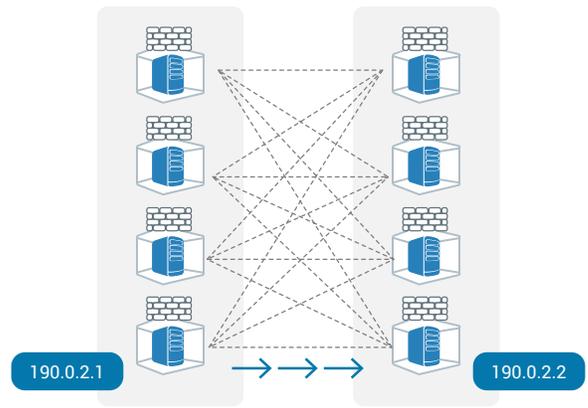
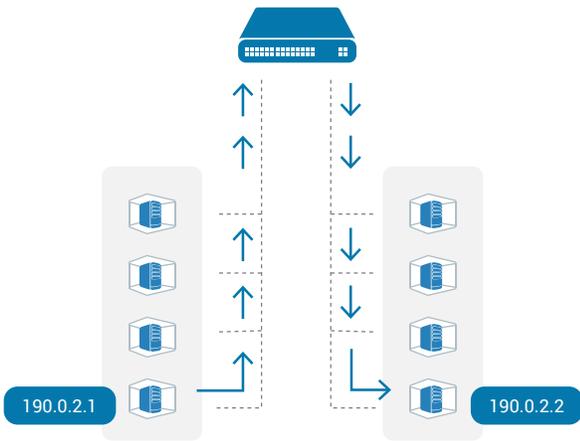


Figure 1: Network segmentation vs. host-based segmentation

Because iptables and WFP are part of standard operating systems, they are available on any host where those operating systems are deployed: physical servers, virtual machines, public cloud, and even container platforms. The ability to run in any environment is a key advantage of segmenting at the host. The idea is to support many disparate environments by orchestrating policy through centralized controls. This allows segmentation across highly distributed, heterogeneous environments with both consistency and granularity – a critical advantage as enterprises move off their network to cloud, and typically to multi-cloud to avoid vendor lock-in.

Review: Key Principles of Segmentation

What are some of the other key attributes that solutions should be evaluated on?

- **Scale:** Will the solution scale to meet the number of workloads you have in the data center and public cloud? In our experience, even 800 workloads will have a range of 180,000–360,000 unique connections. At a 10,000 workload threshold, that number mushrooms to 17M–43M rules (see Figure 3: Kirner’s equation). For any application you have, each one of those connections will require a rule.

$$R = \rho W^\sigma$$

Where Connectivity Size Factor $\sigma = 1.8 - 1.9$

And the Rules Per Edge Factor $\rho = 1.1$

$$R_{max} = 1.1W^{1.9}$$

$$R_{min} = 1.1W^{1.8}$$

Workloads	Total Rules
800	180K–360K
2,500	1.4M–3.1M
10,000	17M–43M

Figure 3: Kirner’s equation

- **Workflows:** Writing all the rules for those connections is another matter. Does the solution provide a set of workflows that help you successfully get to segmentation with an economy of resources and without operational disruption?
- **Agnosticism:** Does the solution have religion about where it will work? Or can it support all of your environments (including outside the network in the public cloud) with granularity and consistency, regardless of the infrastructure it’s running on?
- **Visibility promotes safety:** Real-time visibility provides a foundation for understanding application behavior and common service dependencies in order to create the ideal segmentation strategy without breaking applications (e.g., business applications, Active Directory, Exchange, DNS) – because you can’t protect what you can’t see.

- **Test before you enforce:** Does the solution provide feedback and workflows to ensure you can test your policies before going to enforcement? Ideally, your solution will allow traffic restriction rules to be run initially in a test mode that would expose volume of alerts and allow adjustments to be made as necessary. Model security policy and receive real-time visual feedback before going into production to eliminate the risk of breaking applications whenever new policies are enforced.

	Network (Cisco, PAN, Fortinet, Checkpoint)	SDN (Cisco, VMware)	Host-based (Illumio)
Environmental segmentation	●	●	●
Application segmentation	●	●	●
Tier segmentation	●	●	●
User segmentation	◐	◐	●
Process segmentation	○	○	●
Deployment: cloud, containers	◑	○	●
Simple, fast deployment	◑	○	●
Incremental, start small	○	○	●
Traffic visibility	○	○	●
Test policy before enforcing	○	○	●
Network / infrastructure independent	○	○	●
Cost	\$\$\$\$	\$\$\$	\$
Segmentation risk	High risk	High risk	Low risk
# policy rules	High	High	Low

Figure 2: Segmentation three ways

● Supported ○ Unsupported ◐ ◑ Partially Supported

Conclusion: Segmentation That Works – Anywhere

We rely on the network to deliver applications. But as IT scales in size, connectivity, and environments *outside* the network to public cloud, we cannot rely solely on the network to secure applications. The network is not the best option for designing, building, and delivering security segmentation.

The network cannot provide humans with 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 security segmentation from the network. 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.

Follow Us



About Illumio

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 © 2019 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.