WHITE PAPER–MARCH 2020

The Four Barriers to Micro-Segmentation

vmware®

Table of contents

Introduction	3
Policy discovery challenges	4
Limited-access controls	6
Reliance on agents	7
Lack of threat detection and prevention	8
Summary	9

Introduction

The growing number of sophisticated attacks on corporate information assets is a cause for concern. With the average cost of a data breach approaching $4 million[1], chief information security officers (CISOs) must look for ways to enhance their *enterprise security* posture. An increasing portion of these attacks falls under the category of advanced persistent threats (APTs). Attackers discover creative ways to penetrate the data center, dwell in it for months and use internal communication pathways to reach their desired targets—and compromise them.

The traditional security approach relied primarily on perimeter defense—securing the north-south traffic, but assuming that east-west traffic in the data center was inherently safe[2]. Thus, the traditional approach did little to thwart APTs.

Zero trust is one approach to improving data center defenses. This approach calls for the inspection of every traffic flow within the data center. Zero trust divides the data center infrastructure into smaller security zones. Traffic between the zones is inspected to ensure it adheres to the security policies defined by the organization.

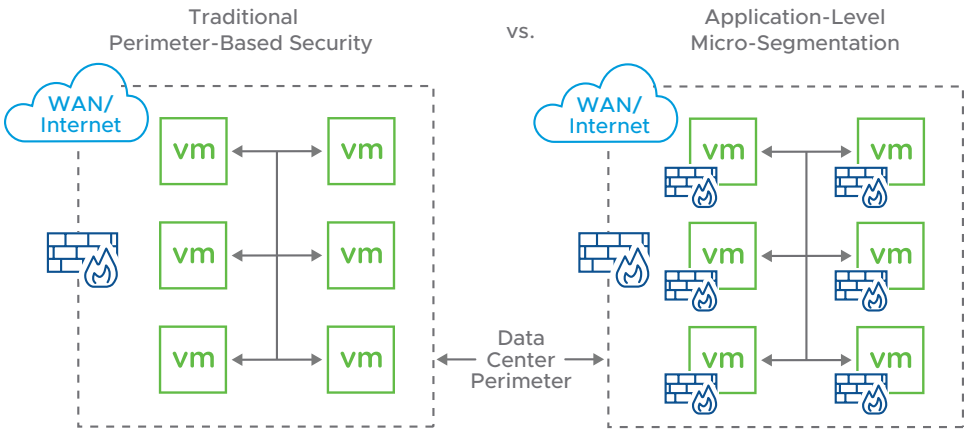

FIGURE 1: Perimeter defense vs. micro-segmentation.

Micro-segmentation is an approach to dividing the data center infrastructure into small zones (see Figure 1). Effectively, micro-segmentation allows fine-grain control of traffic flows between every workload, allowing the administrator to protect all east-west communication. In other words, micro-segmentation helps achieve zero trust.

So, deploy micro-segmentation and the problem is solved, right? Not quite. While the concept of micro-segmentation has been around for a while, organizations still face barriers when trying to apply it in practice. Let's examine some of the top barriers to micro-segmentation:

- **Policy discovery challenges** – Identifying the right micro-segments and configuring the proper security policies is an extremely daunting task, especially in a dynamic data center environment.
- **Limited-access controls** – Basing micro-segmentation solely on L4 attributes (e.g., IP addresses and ports) is not enough. The ephemeral nature of applications and flows requires more than that.

1. Verizon. "2018 Data Breach Investigations Report." April 2018.
2. SANS Institute. "Knock, Knock: Is This Security Thing Working?". March 2020.

- **Reliance on agents** – Some micro-segmentation implementations require the installation of extra software agents on each virtual machine (VM), causing complexity and introducing vulnerability.
- **Lack of threat detection and prevention** – Threats often masquerade as normal-looking traffic. Settling for basic traffic blocking rules isn't enough.

In the next sections, we examine these barriers and suggest ways to overcome them.

Policy discovery challenges

Micro-segmentation involves segmenting the data center network into small security zones and controlling the traffic flows between them. The implementation is straightforward if you know where to set the segment boundaries.

In many cases, administrators seek to apply micro-segmentation to existing data center architectures. Such brownfield topologies are challenging to analyze simply because they are the product of piecemeal evolution over the years. Architecture documentation is scarce, outdated or nonexistent. Creating a map of the current data center topology is a slow and error-prone process. Security policies applied based on incomplete information and guesswork are likely to leave gaping holes.

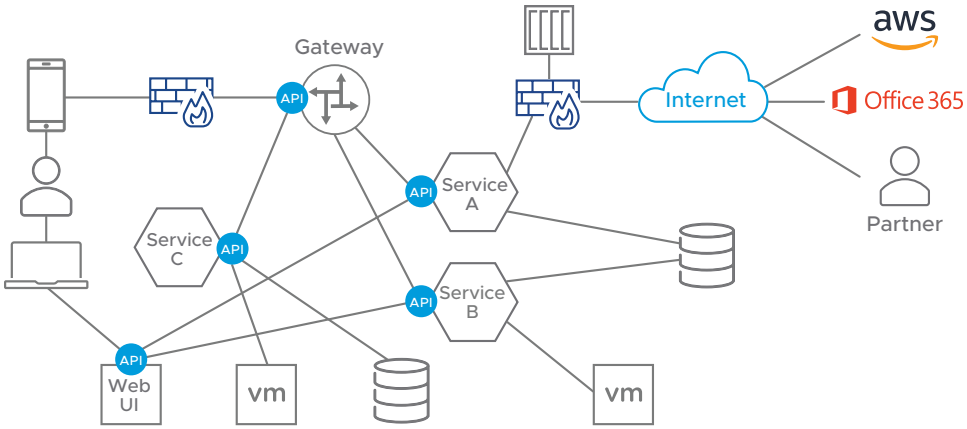

FIGURE 2: Modern apps architecture.

More recent data center architectures have challenges, too. Modern applications running in greenfield data centers are rarely monolithic and do not run on designated servers. Application sub-components span multiple dynamic resources. Some applications divide into tiers, others are built on microservices and some take advantage of cloud services (see Figure 2).

Understanding the current applications' topology and the communication flows between their sub-components is not easy. And because many applications are dynamic in nature, their topology and communication flow change over time.

Administrators need to construct a map of applications and flows that reflects the actual applications' topology and the communication flows between their sub-components to determine the correct micro-segmentation policies.

Building such a map involves the collection and analysis of information from multiple sources, such as configuration management databases (CMDBs), data center management platforms (e.g., VMware vCenter®) and traffic logs. The information collected is often incomplete, sometimes inconsistent, plus it changes over time.

Manually constructing accurate maps of applications and flows is time consuming and error prone. So, it's no surprise that administrators resort to a hit-or-miss approach, basing their micro-segmentation security policies on guesswork and stale information. This approach creates security gaps that leave the data center prone to attacks.

VMware recognized the policy discovery challenge and developed *VMware NSX® Intelligence™* to address it. NSX Intelligence is a *distributed analytics engine* that automates the policy discovery process as follows:

1. NSX Intelligence, aided by other information feeds (e.g., VMware vRealize® Network Insight™), collects information about the running applications and their communication flows. The collected information is analyzed centrally.
2. The result is a comprehensive topology map of applications and flows (see Figure 3). Administrators can easily see the actual application components and the communication flows between them. The map is hierarchical, making it easy to traverse large networks, eliminating the guesswork involved in understanding the topology.
3. Based on the map in Figure 3, NSX Intelligence automatically generates recommendations for micro-segmentation security policies. The recommendations are based on the observed flows, deducing which ones should be allowed and blocking the rest—effectively following the zero trust approach.
4. The administrator can apply these recommendations with a simple click. The selected policies are propagated to the NSX distributed firewall and applied to every node in the network.
5. NSX Intelligence validates the topology against the security policies and presents a color-coded compliance map. Administrators can immediately spot which flows are compliant and which aren't.
6. The administrator may add/modify policies and iterate the process.

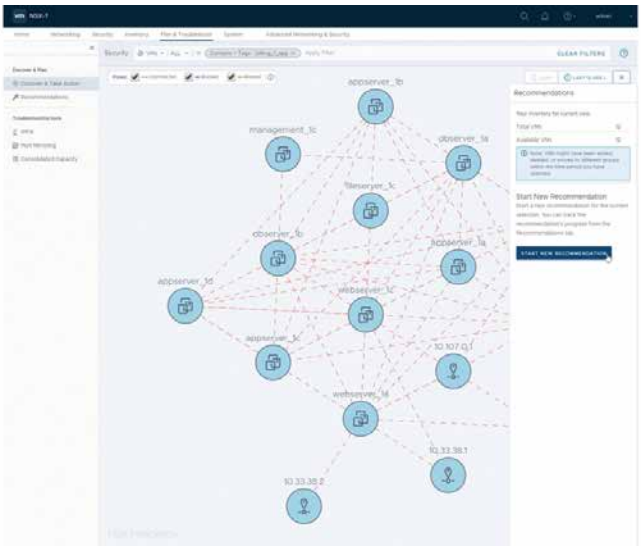

FIGURE 3: NSX Intelligence apps and flows map.

Limited-access controls

Understanding the actual topology of the data center helps determine where to apply micro-segmentation policies. Yet, the effectiveness of these policies depends on their level of sophistication. Just as with home protection, the ability to keep burglars out depends on the sophistication of the door locks.

Most micro-segmentation security policies rely on L4 network parameters—namely, source and destination IP addresses and port numbers. For example, to block (or allow) HTTP traffic between two nodes, it might be sufficient to block (or allow) TCP traffic between them on port 80.

While L4 parameters identify many standard traffic patterns, the last thing we should expect attackers to do is only to use standard traffic patterns. For example, many attacks have been traced back to innocent-looking port 80 traffic.

Furthermore, because modern applications and flows are dynamic, they often use ephemeral IP addresses and port numbers. Micro-segmentation security policies must be able to recognize such flows based on more than simple L4 parameters. We need better door locks.

But the challenge of recognizing good or bad traffic flows doesn't just stop with port numbers. For example, the convention for HTTPS is to use port 443. Can we assume that any traffic on port 443 is safe? Not quite. HTTPS uses different Transport Layer Security (TLS) protocol versions (e.g., TLS 1.0, 1.2, 1.3). Suppose we determine that we should only allow HTTPS traffic using TLS 1.2 or higher—how do we enforce that? Using port 443 as an indicator wouldn't be enough to distinguish between compliant and noncompliant HTTPS traffic.

If we want to detect HTTPS traffic, we must enhance our traffic analysis capabilities to look deeper into the traffic to detect HTTPS traffic regardless of the port the traffic uses. That is, the traffic analysis must be able to use application level (L7) characteristics.

The *VMware Service-defined Firewall* offers a variety of L7 capabilities for micro-segmentation (see Figure 4), including the following capabilities:

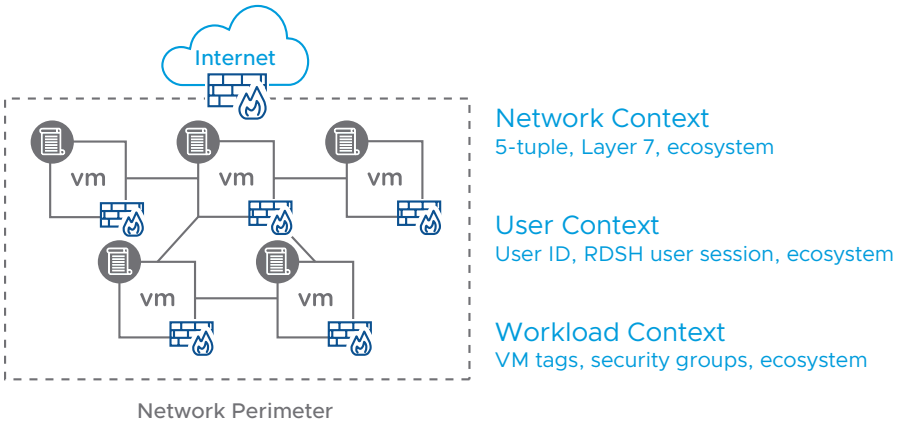

FIGURE 4: L7 micro-segmentation attributes.

- **Application identifier** – Security policies can make use of an application identifier (AppID) that uses L7 attributes to identify a particular flow. For example, the administrator can define a service named APP_HTTPS_TLS_V12, which designates HTTPS traffic that utilizes TLS 1.2. The Service-defined Firewall uses a built-in, deep packet inspection (DPI) engine to enforce the policy in real time. The AppID capability extends traffic analysis beyond simple port identification and makes the administrator's intention easier to define.

- **User identifier** – The growing adoption of virtual desktop infrastructure (VDI) allows multiple users to connect to a single server and initiate requests from it. Because not all users have the same access rights, it is often vital to allow (or block) traffic based on the actual user ID. For example, requests to an HR application should be allowed only for users who have access rights to this sensitive information. The Service-defined Firewall supports user ID attributes for defining micro-segmentation policies. The user ID is synched with information stored within Active Directory, allowing administrators to define security policies that allow/block traffic based on user IDs rather than IP addresses.
- **URL filtering** – The rapid adoption of software as a service (SaaS) and other cloud-based services requires fine-grain filtering of internet destinations. To enable such filtering, the Service-defined Firewall supports policy rules that can specify fully qualified domain names (FQDN), enabling administrators to allow traffic targeted at legitimate destinations (e.g., office365.com) while blocking suspicious targets.

VMware's L7 capabilities go far beyond L4 rules, enabling administrators to deploy more powerful locks throughout their data center network.

Reliance on agents

Understanding the topology and building the applications and flows map is important. So is the ability to define L7 policies capable of thwarting sophisticated attacks. But without the ability to do real-time traffic inspection and policy enforcement at every node, micro-segmentation wouldn't work.

Traditionally, firewalls have been tasked with inspecting network traffic and enforcing security policies. However, using traditional firewalls for micro-segmentation is complicated and prohibitively expensive. For a firewall to inspect traffic, the traffic has to pass through it; this can be done by either bringing firewalls closer to the traffic or the traffic closer to the firewall. The first alternative calls for deploying numerous firewall appliances throughout the network—a complex and expensive proposition. The second alternative calls for redirecting each node's traffic to the firewalls and back—a practice called hair-pining—that results in extra traffic and unnecessary latency.

Because traditional firewalls won't do, some *micro-segmentation* implementations install a software agent on each server and activate the firewall in the server OS using the agent. Let's examine the limitations of this approach:

- **Version complexity** – Most data centers have a plethora of OS flavors and versions deployed at any given time. A micro-segmentation solution that attempts to utilize the OS firewall would have to deal with the OS variety, changing OS versions and the different firewall capabilities on each OS. Managing different flavors of OS firewalls can quickly become a management nightmare.
- **Privileged access** – For an administrator to install and configure a software agent on an OS, they would need privileged access to the server, something that can be difficult to obtain. Furthermore, granting privileged access creates another security loophole, one that organizations want to avoid.
- **Limited capabilities** – Server OS firewalls come with their own set of limitations, and those vary from one OS to another. Most server OS vendors only provide L4 controls with no advertised roadmap to implement L7 capabilities.
- **Agent vulnerability** – Software agents typically run in user space rather than kernel space. Therefore, an attacker who gets root access to the host server can easily subvert the agent. Subsequently, the attacker can neutralize the agent and nullify security policies.
- **Agent fatigue** – Most security administrators already have tens of agents or more to manage[3]. Typically, they are not enthusiastic about managing yet another agent to achieve micro-segmentation. Micro-segmentation solutions that require agents add to the administrator's burden, taking time away from other security tasks.

3. Forrester Consulting. "To Enable Zero Trust, Rethink Your Firewall Strategy." February 2020.

In contrast, the Service-defined Firewall is built into the hypervisor (see Figure 5), eliminating the need to install and configure separate software on every VM. Because the NSX data plane functions are in kernel space, the firewall is immune to attackers who attempt to neutralize it.

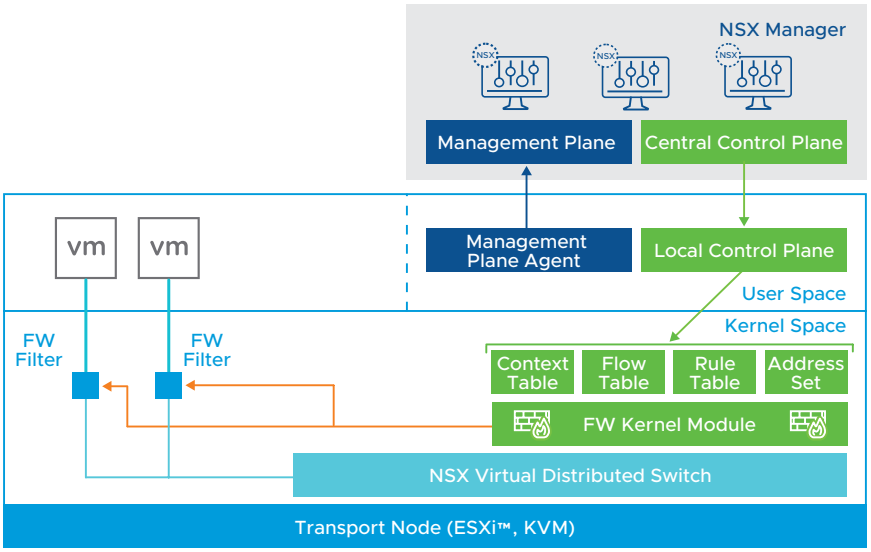

FIGURE 5: Built-in firewall functions.

The Service-defined Firewall is integrated into core networking functions of NSX, gaining native access to every packet and every workload—without extra configuration or traffic hair-pinning.

Security policies are defined centrally using the same management console as NSX. Once defined, policies automatically propagate to every node for enforcement.

By eliminating the need to install, configure and manage additional software agents, VMware has removed another significant barrier to micro-segmentation. Admins can focus on defining their micro-segmentation security policies and let the system do the rest—automatically.

Lack of threat detection and prevention

So far, we've reduced the complexity associated with defining security policies, added much-needed L7 capabilities and made the firewall functions native to the platform. These enhancements made micro-segmentation powerful and put it within easy reach. Yet, there's more to be done.

Hackers use masquerading techniques to embed threats within normal-looking traffic flows. Micro-segmentation identifies the types of traffic flows that should be allowed (or blocked) between segments. But without advanced inspection capabilities, micro-segmentation alone won't intercept hidden threats.

There are solutions designed to do just that: *intrusion detection systems and intrusion prevention systems (IDS/IPS)* excel at scanning traffic, using predefined signatures to identify threats. Once detected, the traffic carrying the threat is immediately blocked, preventing a potential attack.

Fortunately, *the NSX Distributed IDS/IPS™ solution* (see Figure 6) enables VMware to augment its micro-segmentation capabilities with threat detection tools.

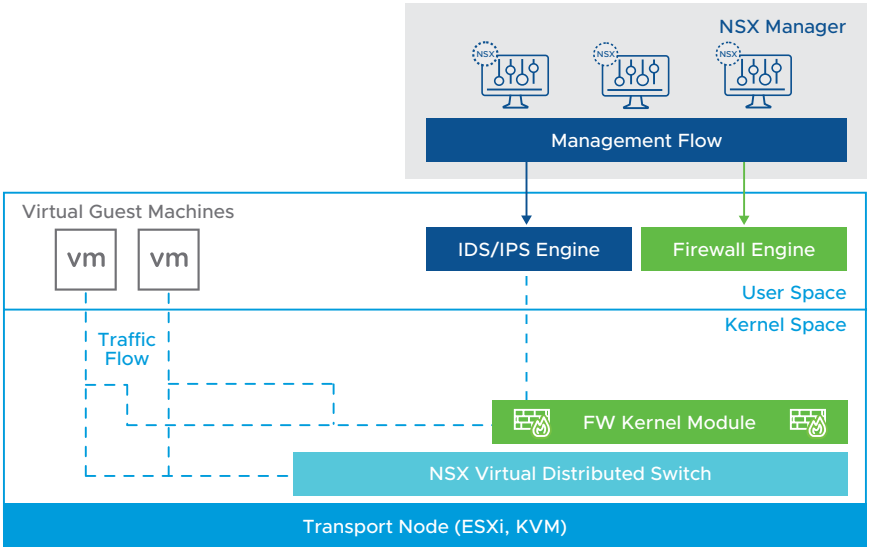

FIGURE 6: Built-in threat detection and prevention.

The distributed architecture of the VMware IDS/IPS solution offers several key advantages:

- Traffic inspection is done at every node, avoiding the need to deploy additional standalone appliances.
- Signature inspection is done in a distributed manner, offering ample capacity to inspect every flow.
- Because NSX is workload-aware, signature-based inspection is selective, based on the knowledge of running applications.
- IDS/IPS policies are centrally managed together with all other security policies, benefiting from the simplicity and scale of the NSX Manager™.

The addition of built-in IDS/IPS capabilities allows administrators to control traffic flows between segments, as well as to detect and stop hidden attacks.

Summary

Protecting corporate information assets is crucial, especially given the risks and costs associated with cyberattacks. *Micro-segmentation* is a viable strategy for improving data center security. However, traditional approaches to micro-segmentation pose significant limitations that impact its effectiveness and adoption.

With the introduction of *VMware NSX Intelligence*, VMware dramatically simplified and streamlined the creation of the required security policies. Additional L7 controls increase the effectiveness and management of security policies. The native (agent-less) firewall capabilities simplify deployment, and eliminate the overhead and risks associated with agent-based solutions. And finally, the built-in IDS/IPS capabilities ensure security policies can also address hidden attacks.

VMware pioneered micro-segmentation by embedding it into its NSX *virtual networking platform*. Micro-segmentation is a manifestation of VMware's *intrinsic security* vision, whereby security is built into the infrastructure rather than bolted on. VMware continues to lead the evolution of micro-segmentation. The capabilities described in this paper not only extend micro-segmentation, they also remove critical barriers to its adoption.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 vmware.com Copyright © 2020 VMware, Inc.
All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: 485849aq-wp-four-barrs-micro-segmntatn-uslet-Final 3/20