



Evolving Micro-Segmentation for Preventive Security: Adaptive Protection in a DevOps World

Written by **Dave Shackleford**

January 2019

Sponsored by:

VMware

Intro to Micro-Segmentation

Network security has changed significantly during the past several decades. Basic network security started with packet filtering devices unable to perceive the state of each session; each packet was an isolated event. This approach allowed attackers to spoof traffic easily and bypass these simplistic controls, so security researchers began tying the state of the network traffic to the policy controls applied. With stateful filtering devices, we gained insight into legitimate sessions versus deliberate, malicious network traffic patterns. Before long, that wasn't enough, either. In hindsight, these attacks were often simple examples of traffic and protocol manipulation and became relatively easy to detect and block.

The attack landscape evolved to more application-centric attacks, and our network protection controls needed to advance and improve to keep up, ultimately leading to the creation of intrusion detection and intrusion prevention systems, as well as web application firewalls. As attacks grew more sophisticated, security professionals realized that we were continually failing to prevent many attacks because we couldn't properly evaluate application and user behaviors within our environments. This realization led to the creation of network behavior monitoring, as well as the "next gen" firewall industry.



Today, we're still struggling, even with all of these technologies. Attackers are still getting in, and frequently, all manner of malicious communication goes totally undetected. Some of today's attackers are very smart and have sophisticated ways to "blend in"—for example, they know huge volumes of traffic are coming in and going out on TCP port 443, and to many firewalls this just looks like traditional HTTPS traffic to websites, so attackers use this to their advantage. Many other traditional network controls are totally blind in these scenarios, even when attackers move laterally through the network to look for new systems and data to compromise.

It's time to rethink the way we're approaching network security today. Some of the things we need to address include:

- Looking at our entire environment as potentially untrusted or compromised, versus thinking in terms of "outside-in" attack vectors—increasingly, the most damaging attack scenarios are internal due to advanced malware and phishing exercises compromising end users
- Better understanding intended application behavior—from the processes running on workloads to the network traffic they generate—and doing our best to enforce these approved application behaviors
- Focusing on trust relationships and system-to-system relationships in general within all parts of our environment—most of the communications we see in enterprise networks today are either wholly unnecessary or not relevant to the systems or applications really needed for business

To start addressing network security in today's highly converged and cloud environments, we need to embrace the idea of software-defined micro-segmentation. Micro-segmentation is a model of defining network isolation policies allowing organizations to segment and control workloads based on application profiles and workload attributes. Its focus is on making network security more granular and controllable. With the growth in software-defined network stacks both internally and in the cloud, the capability to leverage a more configurable policy engine for controlling traffic is a reality today, and will change the way we define traffic policies and access controls going forward.

Internal Network Focus

A key component of this change is a shift in focus from traditional perimeter-based approaches to internal network segments. Most network security architecture designs and controls were heavily focused on "keeping bad things out" instead of controlling and monitoring traffic between internal segments and systems. Trying to adapt firewalls, routers, web application firewalls, intrusion prevention systems and other common control mechanisms to internal use cases were expensive and limited in success for many organizations, especially in physical data center environments. This area of security hygiene is also immensely improved within a cloud or software-defined data center. Using the network virtualization stack as a unified control and introspection plane, all workload traffic can be inspected and controlled in a single place. In addition, every compute asset has a software-defined perimeter policy that is attached to it and evaluated during interactions with other systems.

What is micro-segmentation?

Micro-segmentation is a model of defining network isolation policies allowing organizations to segment and control workloads based on application profiles and workload attributes.

Attack Surface Reduction

As organizations begin to explore how they'll implement micro-segmentation, they should think about how to start building a more sustainable access control and network monitoring strategy. The first step to managing this is to start with good "cyber hygiene" practices in the environment. These will include some of the following, shown in Figure 1.

- **System and application inventory discovery and maintenance.** To gather inventory data about systems and applications, we've traditionally relied on a mix of agent-based reporting and scanning tools. On the surface, this doesn't sound difficult to achieve, but server sprawl and the natural drift in versions of applications over time makes this challenging, especially in traditional physical data centers.





Figure 1. Top Three Cyber Hygiene Practice


- **Configuration management.** Defining configuration standards in our environments takes time, and these have to be revisited regularly to keep them up to date. Scanning tools are often used to assess configuration state, but implementing sound configurations and keeping them in place over time can be difficult, especially with special cases where incompatibility or applications breaking due to a configuration setting require exceptions.
- **Patching.** Even though we all know patching is important, the process of testing patches, rolling them out and tracking the exceptions (again) can easily lead to patches not getting installed in a timely manner. This problem is especially true for applications installed on workloads and spread out across large computing environments.

System and Application Inventory Discovery and Maintenance

While in-depth patching and configuration management discussions are beyond the scope of this paper, they both are critical to the first area of hygiene: inventory. Simply knowing what you have at any given time, what the intended state of the asset should be and its actual state, is paramount to building a sound base of cyber hygiene. Three important elements that apply to the discovery and inventory management phases of cyber hygiene include the following, pictured in Figure 2.

 **Mean time to detect/track.** How long it takes to discover/ detect compute workloads across the organization can have a major impact on the success of any inventory monitoring and management strategy, especially in highly dynamic cloud and DevOps scenarios (covered in more detail shortly).

 **Environment coverage.** The breadth of the environment regularly or continuously assessed for inventory changes or updates can affect how current the inventory is, especially with the older generation of scanning and agent-based reporting tools.

 **Asset criticality and grouping.** Identification of *specific* assets within the environment, ideally through some sort of tagging or naming mechanism, is invaluable when evaluating risk.

All aspects of asset tracking and evaluation are vastly simplified in a software-defined environment. In short, any virtual machines or instances (or containers running within these systems) are always linked to the underlying hypervisor, and through the hypervisor itself, APIs and other management and monitoring tools, these systems can all be queried and monitored continuously. Building on the theme of hygiene related to patching and configuration, as we plan a micro-segmentation strategy we will need to prioritize validating what services and applications are installed and available on each workload and also the versions of these apps in use (and related components such as libraries and drivers). Additionally, as the environment changes, we should ask:

- Can we reduce the workload attack surface?
- Do we have an up-to-date view of the workload configuration and status within the environment?



Figure 2. Important Aspects of Inventory Management

The entire cycle of discovery, asset evaluation, configuration and security posture, and monitoring workload state can be greatly facilitated by using software-defined infrastructure. Once we have inventory in place (and continuously updated), hypervisors and network virtualization tools can help us to enforce the desired state of not only the workloads themselves, but also the interaction between workloads that should be communicating for application environments to function properly. This dynamic, flexible model of micro-segmentation and application control is at the heart of the next generation of software-defined security, which we'll cover in the next sections.

DevOps and New Deployment Models

Over time, development teams and operations teams have had to collaborate more often and with much more rigor, sparking open discussions and much more integration in the overall deployment scenarios many organizations maintain. This trend led to a movement of sorts known as DevOps, which strives to foster open dialogue and intense collaboration between Development and IT Operations teams, leading to the possibility of “continuous delivery” of code or much more frequent code promotion than traditionally seen. Condensed data centers and cloud environments feel the effects, as many new features can be rolled out much more quickly.

The Value of Automation: Moving to Adaptive Micro-Segmentation

In order for security to keep pace with the DevOps teams and deployment models, we need to automate core security tasks by embedding security controls and processes into deployments and running production workloads. To successfully implement an adaptive micro-segmentation strategy in a fast-paced DevOps environment, controls will need to be defined for applications and workloads following standards and requirements and then applied automatically in several places:

- Within the workload template or image, or container image running within a virtual infrastructure
- Within configuration templates for “infrastructure as code” such as Amazon Web Services CloudFormation or Terraform
- Within a software-defined network security policy encapsulating any running workload and allows or denies specified traffic patterns
- Within a central policy engine/enforcement point to arbitrate network and application traffic between virtual workloads

With policies automatically applied after assignment, any workload will be dynamically protected at all times, and any updates to workload configuration or network status can be detected and controlled as well.

DevOps Goals

DevOps strives for a number of goals and focal areas:

- Automated provisioning. The more automated the provisioning of resources and assets, the more rapid the SDLC and operations model can operate.
- No-downtime deployments. As cloud services are based on service-oriented costing models, downtime is less acceptable.
- Monitoring. Constant monitoring and vigilance of code and operations will help to streamline and improve quality immensely. This goal is one of the foundations of DevOps.
- “Fail fast and often.” The sooner code flaws can be detected, the less impact they'll have in a working production environment. Rapid and almost constant testing needs to occur for this to happen.
- Automated builds and testing. More automation in the testing and QA processes will help speed things up and improve delivery times.

Enterprises need to adopt one overarching theme when designing a dynamic security architecture model: one of “zero trust.”

While there are many tools and controls available to help monitor internal workloads and data moving between hybrid cloud environments, enterprises need to adopt one overarching theme when designing a dynamic security architecture model: one of “zero trust.” In a nutshell, zero trust means no traffic and communications should be trusted, from both outside and inside the data center. Each asset should individually validate the others trying to communicate with it before consenting to share data.

In order to implement a zero trust model, organizations need to integrate security into the workloads themselves and move with the instances and data as they migrate between internal and public cloud environments. To create a more robust level of enforcement policy, security teams also need to better understand the actual behavior of the applications and services running on each system, as well as continuously evaluate the relationships between systems and applications. The implementation of these concepts results in “adaptive micro-segmentation,” which complements the zero trust concept.

Dynamic assets such as virtual instances (running on virtualization infrastructure technology) and containers are difficult to position behind “fixed” network enforcement points, so organizations can adopt a zero trust micro-segmentation strategy that only allows traffic to flow between approved systems and connections, and monitors the actual application and service behaviors within the workloads to adjust policy dynamically. Using a software-defined virtual backplane all communications and workload configuration elements are linked to, security and operations teams can accomplish this in a more scalable way. These elements—continuous monitoring of the network traffic between workloads as well as the internal services and application behaviors—are at the heart of adaptive micro-segmentation.

Adaptive micro-segmentation, which simply means using micro-segmentation with dynamic policy evaluation of both network traffic between workloads and the OS and application behaviors and components within the compute elements themselves, prevents attackers from using unapproved connections to move laterally from a compromised application or system regardless of environment. Essentially, this facilitates the creation of “affinity policies,” where systems have relationships and permitted applications and traffic, and any attempted communications are evaluated and compared against these policies to determine whether the actions should be permitted. This happens continuously, and effective micro-segmentation technology will also include some sort of machine learning capabilities to perform analytics processing of attempted behaviors, adapting dynamically over time to changes in the workloads and application environments.

This new micro-segmentation model also reduces the post-compromise risk when an attacker illicitly gains access to an asset within a data center or cloud environment, as the attacker will invariably exhibit some behavior (network connection attempts, changes in applications and services, or both) that is dynamically identified and counteracted with adaptive policy. Cloud design and operations teams (and often

To automate the implementation of an adaptive micro-segmentation strategy, organizations need to ensure visibility into network traffic, the workload and application configuration.

Benefits of an Adaptive Micro-Segmentation Strategy

A software-defined adaptive microsegmentation strategy offers:

- Detection and prevention of lateral movement scenarios on internal networks
- Improved visibility into application inventory and behaviors
- Reduction of post-compromise risk of data exfiltration and other negative consequences
- Centralized and more efficiently managed network policy control
- Compute/process level visibility and control dynamically updated with workload changes

DevOps teams) refer to this as limiting the “blast radius” of an attack, as any damage is contained to the smallest possible surface area and attackers are prevented from leveraging one compromised asset to access another. This method works not only by controlling asset-to-asset communication, but also by evaluating the actual applications running and assessing what these applications are trying to do, helping to identify and maintain the concepts of “cyber hygiene” mentioned earlier.

Next-Generation Security in the Software-Defined Data Center

As we shift toward a fully software-defined data center (SDDC), both in-house and across various public cloud environments, a number of things are changing in the realm of information security. Fortunately, all of these changes are positive and will help us finally get a handle on some of the industry’s more pressing and challenging problems.

We’re embracing software-defined security, which includes everything from configuration profiles defined within templates to network policy embedded in the virtual network stack across a hypervisor environment. The use of APIs and virtual appliances from vendors will only grow, eventually replacing many of the hardware-driven platforms we’ve been used to. We’ll see new skills emerging, focused on software-based security definitions and application security mapping and control, as well as automation tools and process updates that heavily rely on automation and orchestration platforms. Development, operations, and information security teams are blending and aligning more closely than ever before.

Additionally, we’ve seen new use cases and implementation methods for security arise due to the unification of the control and policy planes within the SDDC. Some of these use cases include the following:

- **Adaptive microsegmentation for application security.** This is probably the “killer use case” amongst all of them, as we’ve discussed in this paper. A dynamic “zero trust” policy engine adapting to changing workload applications and traffic patterns in DevOps environments will enable security and networking teams to construct much more effective and sustainable isolation and segmentation strategies that grow with your cloud strategy. By helping with hygiene (inventory and configuration) as well as access control (isolation and affinity policy between workloads), the SDDC facilitates highly granular application whitelisting models at all layers.
- **Software-defined DMZs.** In alignment with the previous use case, software-defined DMZs can be flexibly created and managed to encapsulate certain types of workloads of varying sensitivity. By adding a workload into a defined DMZ, it inherits the DMZ’s policies and can immediately adapt to the environment. Software-based DMZs can also be much more rapidly provisioned and updated compared to traditional physical network segments.

- **Security for virtual desktops.** As more organizations embrace Virtual Desktop Infrastructure (VDI), we will start to see the same benefits of micro-segmentation and application control applied in end-user compute environments. Today, many attacks begin with the end user and move laterally from end user desktops in the early stages of attack lifecycles. Few network environments are well-equipped to detect or control this kind of behavior and VDI with a security layer embedded in the virtual network greatly enhances security capabilities to control what traffic is allowed and trigger an early warning system of possible malware or attacker compromise.
- **Agentless anti-malware.** Offloading workload antimalware processing to a dedicated virtual appliance or other virtualization-compatible engine has been common for some time but will only continue to grow and advance as SDDC and cloud deployments proliferate. Integrating antimalware technology with the virtual control plane via APIs also facilitates more effective detection and response capabilities, such as automatically quarantining a suspicious or infected virtual machine.

There are certain to be many more security-oriented use cases that emerge as the SDDC and DevSecOps technologies and techniques take hold in modern data centers and cloud environments.

Conclusion

The state of network security today is most definitely in flux. Many organizations still have more traditional controls in place and may still be getting value from them in some ways. However, as our environments get more complex and dynamic workloads continue to evolve, static network segmentation and access control models and tools will continue to fail us more often than not. By defining policy rules based on applications, workloads and their relationships across environments, organizations will more effectively scale with their workloads as they progressively get more complex and distributed across internal and cloud environments.

Leveraging micro-segmentation tools easily implemented and centrally managed, will guarantee more flexibility and reduction of risk from leaning too heavily on any one cloud provider's infrastructure, as well. Look for zero trust technology supporting flexible policy creation with a central policy engine/controller, and also support for numerous types of enforcement within different types of workloads. A good zero trust model will start with the premise of denying all traffic and then only permit communications that are whitelisted, either explicitly or through some form of machine learning that evaluates behaviors within the environment and dynamically adapts policies to changes.

About the Author

[Dave Shackelford](#), a SANS analyst, instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

Sponsor

SANS would like to thank this paper's sponsor:

