



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Does Network Micro-segmentation Provide Additional Security?

GIAC (GSEC) Gold Certification

Author: Steve Jaworski, jaworski.steve [at] gmail.com

Advisor: Mohammed Haron

Accepted: June 2017

Abstract

Network segmentation is a concept of taking a large group of hosts and creating smaller groups of hosts that can communicate with each other without traversing a security control. The smaller groups of hosts each have defined security controls, and groups are independent of each other. Network micro-segmentation takes the smaller group of hosts by configuring controls around individual hosts. The goal of network micro-segmentation is to provide more granular security and reduce an attacker's capability to easily compromise an entire network. If an attacker is successful in compromising a host, he or she is limited to only the network segment on which the host resides. If the host resides in a micro-segment, then the attacker is restricted to only that host. This paper will discuss what network and network micro-segmentation is, where it applies, any additional layer of security including levels of complexity.

Introduction

Network segmentation is a defense in depth strategy that helps prevent attackers from moving laterally within an organization's network. (AlgoSec, n.d.) A lateral movement within a network is commonly called east-west traffic. Segmentation creates choke points within the network and is separated by a security control and is also a way to reduce the attack surface of an organization (Northcutt, n.d.). Breaking up a network with VLANs and Layer 3 network interfaces also segments the network but will not provide effective security. By default, the network will permit all traffic between VLANs. Access Control Lists (ACLs) on the Layer 3 network interfaces can be used to restrict traffic, but they are limited in their capabilities. The ACL's on routers are not the same as firewall rules, depending on the network vendor they may slow down traffic, and they can become too complex to manage if too many exceptions are required. Layer 3 ACLs are most effective when they are small and used for explicit denies. For effective network segmentation, a dedicated security control should exist between hosts and networks.

Network segmentation does not always reduce the attack surface enough. Network micro-segmentation is a more granular approach to preventing lateral movement between hosts where it is more challenging to use traditional hardware security controls (SDX Central, n.d.). One network segment could still contain hundreds, if not thousands of hosts. If a malicious actor compromises a host in a segment, the goal of network micro-segmentation is to prevent or detect the compromise of additional systems. By limiting the attacker's scope, the security team should have a better chance of detecting and remediating the breach, limiting further damage.

Cisco had the Enterprise Strategy Group (ESG) write an analyst report titled "Cisco: ACL Survey" in 2015 asking questions regarding Application Centric

Infrastructure. The survey interviewed 154 IT professionals from companies in North America with more than 500 employees. Some questions in the survey asked if network segmentation can prevent server compromise. For the organizations that experienced a compromise, 47% of the 88 respondents acknowledged that the attacker was able to move laterally from one data center to another. Only 35% of the 154 respondents agreed, some degree of further network segmentation would “definitely” prevent a server compromise (Oltsik, 2015). Another 42% of those surveyed agree that “probably” network segmentation would prevent a server compromise (Oltsik, 2015). Network segmentation is important because 77% of the 154 respondents believe that network segmentation can help prevent server compromise.

Segmentation Details

What is Network Segmentation?

The simplest form of network segmentation is a host connected to the Internet with a boundary device in between, typically a firewall. This common traffic path is also known as north-south traffic. The role of the firewall is to protect the host from the Internet. Adding a security control at the Internet connection is the first step in reducing the attack surface of the organization's environment. In Figure 1 below, the protected host is behind the firewall. Inbound traffic that is not explicitly permitted is blocked by the firewall from the Internet.

Figure SEQ Figure * ARABIC 1. Most Common Network Segmentation

As organization deployed firewalls, attackers learned they only needed to find another weakness behind the organization's security control. They exploited the discovered vulnerability, then, by pivoting, they gain access to a more secure system.

The pivot is when the attacker can laterally move between hosts because a security control was not in place. Lateral network movement is also known as east-west traffic. In response, the organization then moved hosts that were publically accessible into the zone of the firewall. The segmented zone of the firewall is popularly known as a DMZ. Segmentation at this level helped better separate the private and public assets of the organization. Even with a configured DMZ, having multiple hosts in the same segment makes them susceptible to lateral network movement.

Without segmentation, both the client and host can communicate with each other without traversing a security control. The risk in Figure 2, shows that both the client and protected host reside on the same network. An attacker cannot compromise the protected host because the firewall is restricting access to only the permitted HTTP service and the web application does not have any discovered vulnerabilities. The next approach is for the attacker to compromise the client using social engineering. The attacker convinces the user to access a website that successfully exploits an unpatched vulnerability on the client machine. Since the attacker has access to the internal network, he or she can pivot, also known as moving laterally, to the protected host. Without a security control in front of the protected host, the attacker can compromise it with other methods.

Figure SEQ Figure * ARABIC 2. No Internal Protection

Then, the organization moves the protected host into a firewall security zone to reduce the attack surface, as seen in Figure 3. The client has to traverse the firewall just like the public users on the Internet. In the event the client is compromised, the attacker will still have the same restrictions as if he or she were coming from the Internet. A new risk has surfaced in the firewall security zone, which is that all hosts in that zone can communicate with each other. In Figure 3, there are two protected hosts. One protected host contains public data, while the other contains confidential data. The Internet side of

the firewall is configured to allow anyone to access the public-protected host and restricts access to specific source IP addresses for the confidential-protected host. Unfortunately, the public-protected host has a vulnerability that can be easily exploited granting the attacker shell access. The attacker now has the capability to attack the confidential-protected host from the public-protected host. The lateral network traffic risk identified between the client and protected host in Figure 2 is the same risk for hosts in the same segment behind the firewall. The organization neglected to identify the value of the data they were trying to protect.

Figure SEQ Figure * ARABIC 3. Internal Segmentation

To better segment the data, the organization took it one step further by creating multiple firewall security zones, also called DMZs. Typically, web applications are made up of multiple servers performing different functions. Their physical structure is not as simplistic as just having two protected hosts shared with the Internet. Consider the following example: the outermost DMZ typically holds the front-end interface the user's accesses. Then, if the application has a middleware design, there is a middle DMZ that houses that layer of the application process. The innermost layer of the DMZ typically contains the database that supports the application. See Figure 4 for a depiction of this topology. The idea behind this level of segmentation is if an attacker compromises the first layer of the application, the organization's defense team has time to restore and harden that layer before the attacker could compromise the middle layer.

Figure SEQ Figure * ARABIC 4. Segmented Perimeter

Also, not all applications have a middleware design. Many web applications have the interface and logic which a user interacts with, commonly called the front end. The back end is often the database that stores the information for the application. This design only requires two security zones, the front end layer and the back end layer.

With the perimeter hardened, attackers start going after users who have more access to the internal network. Phishing is a popular technique used by attackers to convince a user to click on a link, load a piece of software, or provide his or her credentials. The attacker will send the victim an email with instructions that appear to be from a trusted source. If the victim falls for the Phish, the attacker now has control of an internal host. From that point, the attacker can access any other host on the network the victim has access to. Often the attacker can escalate their privileges to an administrative level, granting the attacker even more access to an organization's valuable data. Many networks trust internal users without additional security controls. Lacking this extra level of security allows the attacker more time to stay on the organization's network undetected. Internal networks are mostly east-west traffic flows. In response, the organization's defense team starts segmenting the data center the same way as the perimeter, as illustrated in Figure 5.

Figure SEQ Figure * ARABIC 5. Internal Segmentation

While segmentation is an improvement for security, it has introduced other issues into the environment such as ease of scalability and cost. Segmentation of north-south traffic is fairly scalable, but east-west traffic is not. The concept, "firewall on a stick", helps segment east-west traffic, when Host A and Host B are in different segments but only communicate via the boundary device. While the segmentation may be effective, the boundary device may have scalability issues due to resource limitations. Cost increases by purchasing a boundary device with more capacity to handle an increase in traffic. While this solution may be effective for small environments, what about large environments? The organization still needs to consider redundancy for their workloads and applications. Depending on the organization's operational requirements, they may choose to duplicate their entire infrastructure at a remote facility. Virtualization is an effective method to reduce hardware and operating costs. With the introduction of

virtualized hosts, the complexity continues to increase for network segmentation. The cost and complexity of routing virtual machines to the physical infrastructure through a security device will become overwhelming. Figure 6 depicts an example of “firewall on a stick”. The switch in the graphic is configured with VLANs for each physical host and each virtual guest. For one of the physical hosts on the left to communicate with one of the virtual guests on the right, the network traffic has to traverse the internal firewall. This is where scale starts to become an issue as more physical or virtual hosts are added to the infrastructure. The internal firewall has to be able to handle the entire load of systems and network traffic is not optimized because all traffic has to route thru a single point.

Figure SEQ Figure * ARABIC 6. Firewall on a Stick

Segmentation can secure north-south network traffic fairly easily and can be cost effective in most physical and virtualized environments. Segmentation of east-west traffic introduces complexity and cost into both physical and virtualized environments. In response to reducing the complexity of the cost of segmenting east-west traffic, the concept of micro-segmentation was developed.

What is Micro-segmentation?

Micro-segmentation is the result of trying to protect hosts that reside in the same security zone. The security zone could be a single subnet, VLAN or broadcast domain. Hosts that can communicate with each other directly without traversing a security control are candidates for micro-segmentation (Bigelow, 2016). Network micro-segmentation places a security control in front of each host. Figure 7. shows an topology of three physical hosts running four virtual guests. Each virtual guest has a firewall running on the hypervisor kernel, not the virtual guest itself. For any communication to occur between virtual guests, the network traffic must pass thru the firewall. If two virtual

guests are running on the same physical host, the traffic must pass through the firewall, which provides the east-west protection. If one virtual guest is running on one physical host and another virtual guest is running on a different physical host, again network traffic must traverse the firewall, providing the north-south protection.

Figure SEQ Figure * ARABIC 7. Virtualized Micro-segmentation

Micro-segmentation is easier to deploy because of two technologies that enable network layer abstraction at the hardware level (Bigelow, 2016). The first technology is Software Defined Networking (SDN). Virtualization of the network means the control plane is separated from the data plane on the network switch or router. The data plane is also known as the forwarding plane in the network, which is responsible for moving packets. The control plane no longer resides on the network equipment itself and operates on a dedicated server called a controller. The central controller instructs the routers and switches how to move packets using what is called the southbound API (SDX Central, n.d). Also, SDN can integrate applications using what is called the northbound API. Programmers have the ability to allow their applications to instruct the network on how to function (SDX Central, n.d.). The most significant benefit of SDN, is the technology can understand the requirements of the application being served to users. The network can optimize itself to meet the performance requirements of the application.

The Software Defined Data Center (SDDC) is the second technology that makes it easier to deploy micro-segmentation. SDDC abstracts all the infrastructure layers in the data center, not only the network layer. The SDDC divides the physical hardware into four components: compute, storage, network, and hardware (Rouse, n.d). This technology allows multiple applications, operating systems, and different network configurations to function on a single piece of hardware. If the organization needs more computing power, additional hardware can be added to balance the virtual resources

across the hardware layer.

What is the additional security?

Forrester Research introduced the model of the “Zero Trust Model” in April of 2013. They submitted a paper “Developing a Framework to Improve Critical Infrastructure Cybersecurity” to NIST (The National Institute of Science and Technology). Forrester reiterates the common security issue of organizations network perimeters being a hard candy shell and the internal network being a soft chewy center for attackers to exploit with ease (Kindervag & Ferrara, 2013). The Zero Trust Model introduces three concepts: securing all resources no matter the location, access control follows the least privileged model, and all traffic is logged and inspected. The first concept, all resources are secured no matter the location is more plausible to implement due to capabilities offered in micro-segmentation products. Due to the wide acceptance of virtualization platforms, organization can easily move guest systems from server to server in a data center and between data centers.

As previously discussed, when trying to protect east-west traffic of resources within the same security zone, it can become very cumbersome and costly to move traffic outside to the zone to an inspection point, and then back in. Both Cisco and VMware have a blog battle touting whom as a truer network micro-segmentation offering. VMware argues their NSX their solution is better because they do not require additional hardware and provide network traffic efficiency (Germain, 2016). Cisco argues their SDN Application Centric Infrastructure is superior to VMware’s NSX offering because they protect any endpoint, whether it’s physical or virtual. (D'Agostino, 2016). No matter what vendors claim about their products, the organization needs to determine what it is trying to protect and what services are offered by the micro-segmentation vendor benefit them. For example: an organization is trying to protect a web application that runs in a virtual machine that is only available to select business partners. The select business

partners are restricted by an IP address and use a VPN tunnel that terminates on the organization's perimeter firewall. Once the business partner is in the network, there are no other restrictions to the web application. Also, the organization wants to be able to move the application between three data centers around the world for redundancy; also, there are internal users that access the web application.

The first option is to setup a virtual firewall that resides within their virtual infrastructure. This type of virtual firewall mimics a hardware firewall that protects physical hosts and provides full Layer 7 inspection. With this option, three virtual firewalls are required, one for each data center. Since, uptime is important to the organization, two virtual firewalls are required per data center. The total number of virtual firewalls now stands at six. This configuration will be complex to manage because the web application needs to move from data center to data center. Firewalls maintain a state table to track the connections it has permitted and are currently active. If the web application moves from one virtual firewall to another, where the state table is not synced, all active connections to the application are dropped. In a single data center, the two virtual firewalls are configured in a high-availability pair (HA), which synchronizes the state table. The application can move freely within the data center and users are not disconnected. Moving the application from data center to data center will be more challenging. Either the organization will have to synchronize firewall state across data centers or configure the application to handle network interruptions. Network interruptions will create a bad user experience, including productivity loss as users wait for the application to become available again.

The second option is to consider using a network micro-segmentation firewall product. The firewall follows the virtual machine as it migrates from one physical host to another and between data centers. This type of firewall is not restricted to running on a physical host as a guest machine. It ties itself into the kernel of the virtualization product. However, the capability of this micro-segmentation firewall only offers Layer 4

inspection services. The firewall only keeps state, restricts what IP addresses are permitted/denied, and on which ports. As the application moves across servers within a data center or from data center to data center, there will not be an issue with maintaining firewall state. Since the firewall state table is maintained, network interruptions are prevented, reducing downtime and provides the user a better experience.

The organization's security team will get the most value from having full application Layer 7 inspection capabilities with the virtual firewalls but have to deal with configuration and operational complexity. The security team along with management will have to choose between two options, the Layer 7 inspection or the network micro-segmentation Layer 4 functionality. Before making a decision, the security team should require the management team to determine the value of data. The organization will need to determine if the cost of the security controls exceeds the value and remediation costs of losing the data.

Besides virtual and micro-segmentation firewalls, other options could include a host-based firewall on the virtual machine with or without a host intrusion detection system (HIDS) component. These technologies have their challenges when it comes to licensing and administration costs. This technology can be less secure, because it is required to run on the guest operating system itself. If the guest operating system is compromised by either an application vulnerability or misconfiguration, the attacker can easily disable the host-based firewall.

Network micro-segmentation does not protect against virtual machine escape. A virtual machine escape is when a vulnerability is exploited on the guest, which grants access to the hypervisor running on the physical host. The attacker now has a backdoor to all the other guests running on the physical host (Rouse, 2016). An escape vulnerability called "virtualized environment neglected operations manipulation" (VENOM) was discovered in open source virtualization products XEN and KVM (Geffner, 2015). VMware, with its commercial hypervisor products, also has

had escape vulnerabilities discovered (Goodin, 2017). Trying to escape out of a virtual machine is not easy to do, but the discovered vulnerabilities have shown it is possible. Virtual Machine sprawl and poor configuration management are considered a higher risk than an attacker successfully breaking out of a virtual machine (Savage, 2015).

Micro-segmentation complexity

Determining how to protect an application or operating system can become quite challenging. Introducing micro-segmentation can add complexity to an already complex environment. VMware NSX offers a grouping/tagging type system to help deploy micro-segmentation within their virtual platform (Miller & Soto, 2015). The concept of creating groups identifies a certain workload and or system type. Then rules are added to the defined groups. When a virtual machine matches the identified criteria, the hypervisor applies the configured rules. If the virtual machine migrates to another physical host, the rules are still enforced. Rule enforcement is always enabled, even when creating new virtual machines. The security team does not have to modify rules every time a new hosts or service comes online. Reduction in administration time occurs with guest deployment because information security policies are in place. VMware is not the only micro-segmentation vendor. Before choosing to deploy any security product, organizations should define their requirements and evaluate multiple vendors.

Protecting a group of web applications hosts should be very easy for micro-segmentation, by only allowing network traffic inbound on TCP port 80 and 443. However, there is more to operating and maintaining web applications servers. Most web applications have some database backend, which means that the web servers need to be able to initiate a connection to the database. There are most likely other applications for tracking activity, performing transactions, securing the OS, or managing the OS, to name a few. Many of these operations require dedicated ports. Some of the connections may

be initiated by the web application hosts themselves or the hosts outside initiate a connection inbound. After some time in operation, the ruleset may start to look a little like Swiss cheese. To be the most secure, the rules to be as granular as possible; but for ease of administration, more relaxed rules may fulfill the need. Corporate politics may come into play when deciding on the level of security required. Security teams typically want to mitigate as much risk as possible. For example; there is that one developer that insists on having direct access to the web hosts and management has high visibility for these particular systems. The information security team has outlined to management the increased risk of allowing direct access to the application, but the access is still approved.

Side effects of micro-segmentation are policies that can create too much granularity. Another concern is the consistency of deploying the policies (Bigelow, 2016). Complexity will vary by organization. The larger the organization, the more likely there will be various competing priorities between departments. To try to appease everyone, the security team may start creating too many complex rules and policies. The organization then may have to maintain configurations for individual hosts instead of groups of hosts. The more granular the security requirements, the more time is required to analyze and make changes, to prevent creating new security issues. Network micro-segmentation does require ongoing administration. It is not, configure once and forget.

Conclusion

Network micro-segmentation does provide additional security but is not a replacement for traditional security controls. It is another layer of security and helps to continue reducing the attack surface for east-west network traffic. Proper patching, disabling unnecessary services, configuring policies, and changing system defaults on hosts and applications are just a start in securing hosts and applications. Organizations may or may not choose segmentation between every host. The Forester Research's "Zero

Trust Model” requires network micro-segmentation to meet the requirement of security all systems regardless of location. An organization needs to be cognizant of how to deploy micro-segmentation rules. Well-written security policies and guidelines can help keep rules from becoming too granular and complex to manage. Vendor evaluation is a must before deployment. Organizations need to make sure the vendor can operate with their existing infrastructure and avoid a proprietary solution. With a proper risk and gap assessment, organizations can make the most effective decision regarding the level of network segmentation.

References

- AlgoSec. (n.d.). *Network Segmentation*. Retrieved May 21, 2017, from www.algosec.com: <https://www.algosec.com/network-segmentation/>
- Bigelow, S. J. (2016, March). *Microsegmentation lets software define network security*. Retrieved April 9, 2017, from Tech Target Search DataCenter: <http://searchdatacenter.techtarget.com/feature/Microsegmentation-lets-software-define-network-security>
- D'Agostino, F. (2016, January 7). *ACI Surpasses VMware NSX Again with Micro Segmentation & End-Point Granularity*. Retrieved June 19, 2017, from Cisco Blogs: <http://blogs.cisco.com/datacenter/aci-surpasses-vmware-nsx-again-with-micro-segmentation-end-point-granularity>
- Geffner, J. (2015, May 21). *VENOM*. Retrieved May 30, 2017, from CrowdStrike: <http://venom.crowdstrike.com/>
- Germain, B. (2016, January 5). *VMware NSX and Split and Smear Micro-Segmentation*. Retrieved June 19, 2017, from VMware Blogs: <https://blogs.vmware.com/networkvirtualization/2016/01/vmware-nsx-and-split-and-smear-micro-segmentation.html/>
- Goodin, D. (2017, March 17). *Virtual machine escape fetches \$105,000 at Pwn2Own hacking contest*. Retrieved May 30, 2017, from Ars Technica: <https://arstechnica.com/security/2017/03/hack-that-escapes-vm-by-exploiting-edge-browser-fetches-105000-at-pwn2own/>
- Kindervag, J., & Ferrara, E. (2013, April 8). *Developing a Framework to Improve Critical Infrastructure Cybersecurity*. Retrieved April 30, 2017, from NIST: http://csrc.nist.gov/cyberframework/rfi_comments/040813_forrester_research.pdf
- Miller, L., & Soto, J. (2015). *Micro-segmentation for Dummies*. Hoboken: John Wiley & Sons, Inc.
- Northcutt, S. (n.d.). *The Attack Surface Problem*. Retrieved May 21, 2017, from Security Laboratory: Defense In Depth Series: <https://www.sans.edu/cyber-research/security-laboratory/article/did-attack-surface>
- Oltsik, J. (2015). *data-center-virtualization/application-centric-infrastructure/white-paper*. Retrieved May 21, 2017, from Cisco.com: <http://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-734499.pdf>
- Rouse, M. (2008, March). *Defintion hairpinning*. Retrieved April 9, 2017, from Tech Target Search Unified Communications: <http://searchunifiedcommunications.techtarget.com/definition/hairpinning>
- Rouse, M. (2016, April). *Definition Virtual Machine Escape*. Retrieved May 30, 2017, from TechTarget Whatis.com: <http://whatis.techtarget.com/definition/virtual-machine-escape>
- Rouse, M. (n.d.). *Definition SDDC (software-defined data center)*. Retrieved April 30, 2017, from Tech Target Search Converged Infrastructure: <http://searchconvergedinfrastructure.techtarget.com/definition/software-defined-data-center-SDDC>
- Savage, M. (2015, May 7). *Top 11 Virtualization Risks Identified*. Retrieved May 30,

2017, from Network Computing: <http://www.networkcomputing.com/data-centers/top-11-virtualization-risks-identified/2062567936>

SDX Central. (n.d.). *How Does Micro-Segmentation Help Security? Explanation*. Retrieved May 21, 2017, from SDX Central: <https://www.sdxcentral.com/sdn/network-virtualization/definitions/how-does-micro-segmentation-help-security-explanation/>

SDX Central. (n.d.). *What are SDN Northbound APIs?* Retrieved April 30, 2017, from SDX Central: <https://www.sdxcentral.com/sdn/definitions/north-bound-interfaces-api/>

SDX Central. (n.d.). *What are SDN Southbound APIs?* Retrieved April 30, 2017, from SDX Central: <https://www.sdxcentral.com/sdn/definitions/southbound-interface-api/>

Does network micro-segmentation provide additional security?	PAGE * MERGEFORMAT 18
--	------------------------

AUTHOR * MERGEFORMAT Steve Jaworski, jaworski.steve@gmail.com

--	--	--