



NATIONAL SECURITY AGENCY CYBERSECURITY INFORMATION

SEGMENT NETWORKS AND DEPLOY APPLICATION-AWARE DEFENSES

A CHANGE IN TACTICS

During the last few years, the world has seen a surge of ransomware, which is malicious code that encrypts data, making it inaccessible. Ransomware can cripple networks unless a ransom is paid or the encryption key is obtained. The critical service disruptions from many of these incidents are amplified by lateral movement, the ability to maneuver across a network using vulnerable internal protocols or privileged features of devices not accessible outside of the network. Some of these ransomware cases involved SMB (Server Message Block) worms that took advantage of vulnerabilities to rapidly spread via the SMB protocol and disrupt large areas of a flat network. Some relied on encrypted malicious traffic and protocols which could have been isolated. With granular network segmentation and application-aware defenses, these types of attack can be contained to minimize damage. Further, analysis of encrypted traffic is becoming increasingly important as the majority of network traffic is now encrypted, allowing malicious activity to pass through it or blend in with it.

Using application-aware firewalls, Virtual Local Area Networks (VLANs), Virtual Routing and Forwarding instances (VRFs), Virtual Private Networks (VPNs), and Software Defined Networking (SDN) to segment network traffic within the network reduces the attack surface and makes it harder for the adversary to move laterally through the network. Combining this with Network Function Virtualization (NFV) in virtualized environments can achieve a more granular level of segmentation—or micro-segmentation—of the network, further restricting an adversary's movements. Failure to segment networks and deploy these application-aware defenses can result in extensive adversary proliferation and malicious insider access within a network.

NETWORK SEGMENTATION

Segmentation is the practice of dividing a network into sub-networks (segments) of devices that share similar security requirements. Generally, segmentation is done by separating access to the most sensitive and vulnerable services on the network, such as directory services, file-share services, and network management. The network can be further segmented by user groups and determinations made on the level of access each user group requires. Once a network is properly segmented, appropriate application-aware defenses can be utilized to isolate and secure each network segment. Such isolation is essential to blocking an adversary's lateral movement through the network and instilling the principle of least privilege to every network device. Limiting device communication between segments also enables better monitoring and visibility into an adversary's attempts to spread from one segment to another.

Physical segmentation uses the configuration and placement of physical devices like routers and switches to create network segments based on functional importance and/or levels of access. Each of these segments can be configured to be physically isolated or connected together using application-aware detection devices. With proper filtering rules and policy implemented at the border of each segment, physical segmentation is the most secure method of segmentation and provides the most protection against adversary lateral movement. However, this technique can be expensive as it requires the use of separate infrastructure for every subnet and requires extensive out of band network management.

Virtual segmentation bears similarities to physical segmentation, but utilizes software features on network devices to perform the segmentation. These software features include VLANs, VPNs, VRFs, and VLAN Access Control Lists (VACLs). VLANs logically separate hosts and endpoints by creating separate subnets and utilize network layer and application aware filtering devices to prevent segments from communicating with each other based on a configured policy. When deploying VLANs, VACLs can be implemented to restrict, allow, or redirect the flow of intra-VLAN traffic on switching devices. Utilizing a combination of VACLs, network layer and application-aware filtering devices provides a defense in depth strategy.

VRFs can also be deployed to virtually segment devices by creating separate routing domains on layer 3 devices. In order for subnets in separate VRFs to communicate, routes between VRFs must be deliberately shared. One benefit of VRFs over VLANs is that VRFs are not susceptible to VLAN hopping attacks. Last of all, VPNs provide a secure communication channel—or tunnel—through a network. A VPN establishes a one-for-one, or gateway-to-gateway connection between an endpoint and a network resource, or two network segments. Encrypted VPNs can provide confidentiality, integrity, and authentication security techniques that make it difficult for an adversary to access these resources. All methods of virtual segmentation require attention to the details of implementation and the filtering policies between segments to ensure proper segmentation is maintained.

As you segment your network, pay attention to how protocols like SMB, NetBIOS, and RPC services are configured. Constraints on these protocols will limit the damage of pass-the-hash and SMB relay attacks. Ensure these protocols are only implemented in accordance with approved policy. This is one way of better defending privileges and accounts.¹

APPLICATION-AWARE FIREWALLS

As attackers become better at obfuscating malicious traffic, firewall manufacturers have added more functionality to inspect deeper into the packet and traffic flow. Traditional firewalls would block or allow data traffic based on packet headers, making assumptions about how traffic will be treated based on its port numbers. Application-aware firewalls perform deep packet inspection and attempt to identify the applications that are traversing the firewall. Blocking traffic that is not properly formatted for the type of applications that are permitted on the network helps prevent adversaries from abusing allowed application protocols. This approach to whitelisting for network applications does not rely on generic ports as the filtering criteria. Additionally, some application-aware firewalls provide automatic signature updates and can filter dynamically changing applications like peer-to-peer connections, which can evade traditional firewall rules.

In addition to those features specific to application-aware firewalls, some offer the ability to control traffic based on things such as specific parts of a website or user actions. When all of these features are utilized together they give network administrators the ability to apply policy at a very granular level. As application-aware firewalls are deployed between network segments, these granular policies can be configured in such a way to dramatically limit an adversary's lateral movement. As a best practice, every network segment should be filtered on the inbound and outbound direction with only the permitted applications allowed in or out. To provide the highest level of security, consider deploying additional features such as network appliances that perform virus scanning, and Data Loss Prevention (DLP).

SOFTWARE DEFINED NETWORKING

Traditional routing strategies typically describe the actions of the “forwarding plane,” which carries data traffic, the “control plane,” which decides how to route or handle data traffic, and the “management plane”. Software Defined Networking (SDN) works to decouple the control plane from the forwarding plane so that control of data flow over the entire network can be directly programmable from a central location. In other words, SDN moves from a distributed control configuration to a logically centralized controller.

The term SDN has taken on many different meanings and become a buzz word within the industry. For the purposes of this paper SDN will be broken into two different technologies: automated provisioning and control plane abstraction. Automated provisioning is defined as any solution that will automatically connect to network devices and reconfigure the device based on a set policy in the controller (e.g. using standard protocols such as SSH and SNMP). On the other hand, with control plane abstraction every device forwards traffic based upon a set of rules that are pushed to the device from a central controller most commonly using a protocol known as OpenFlow. In summary, with automated provisioning every device maintains its own control plane, and with control plane abstraction every device depends on a central controller for their control plane. Deployed properly, both methods provide a flexible way to perform network segmentation and traffic can be dynamically directed to an application-aware firewall when traversing segments. Each method does come with security concerns and caution should be taken when deploying SDN.

¹ For more information on this, please refer to Defend Privileges and Accounts, one of NSA's Top 10 Cybersecurity Mitigations.

Some of the security concerns that need to be addressed with SDN revolve around the security of the controller and the security of the control protocols. With both SDN technologies discussed it is paramount that the central controller have the appropriate security controls in place to limit network accessibility and provide strong authentication of all administrators.

The control protocols should always utilize encryption to provide confidentiality, integrity, and authentication. If the control protocols do not provide strong authentication and integrity checking, then it is likely that an adversary will be able to make unauthorized changes to the network. When deploying the control plane abstraction methodology administrators should physically separate control plane traffic from user traffic by implementing an out of band management network.

After properly securing the controller and control protocols, SDN can quickly adjust network configuration. One of the major benefits to SDN is that it is highly programmable and the network can dynamically adjust based on policy, the application in use, or which user is logged on. With the growth of network device malware and new powerful SDN application programming interfaces (APIs), great care should be taken when programming an SDN network so as not to introduce unexpected behavior (such as security vulnerabilities) with poor development practices.

Application-aware segmentation—using SDN techniques—individually directs data traffic by linking the application's communication preferences and paths together with policies that determine how and where that data should flow. The SDN policies can even be configured to inspect the application traffic and make determinations on the validity of application traffic. In other words, once an application has been identified, pre-determined policies can dictate how software should segregate the data by placing it into the appropriate VLANs, redirecting the traffic to application-aware defenses, or apply filtering rules to limit access to specific segments. This allows for multiple segmentation strategies to be applied simultaneously.

MICRO-SEGMENTATION OF THE APPLICATION LAYER

The increased use of Virtual Machines (VM) and NFV makes it possible to micro-segment the network in virtualized environments such as cloud data centers. Micro-segmentation is the deployment of a virtual firewall at every single virtual network interface where network traffic enters and leaves each virtual machine. The firewall residing on each virtual machine serves to isolate every network resource and endpoint from each other even when residing on the same subnet.

Direct communication between virtual network resources can be controlled at the most granular level. The virtual firewalls can be dynamically created or torn-down on the fly as network resources are provisioned or modified. Because micro-segmentation is software-based, administrative control can be centralized through a single SDN-like controller. Just as in SDN, central policies can determine whether to deny or permit applications from flowing through or to a device. Host-based firewalls, with or without a host intrusion detection system, are another option that can be used on the virtual machine. However, these systems tend to be more vulnerable due to the fact that they are required to operate on the host operating system. If this operating system becomes compromised by either vulnerabilities in the software or misconfiguration, the host-based firewall may be easily disabled by the attacker.

CONCLUSION

To protect against attacks such as pass-the-hash, SMB worms, adversary lateral movement, and malicious insiders, network segmentation and application-aware defenses are essential to a network's security. The deployment of application-aware defenses such as firewalls and proxies, coupled with emerging techniques like SDN and micro-segmentation, create very granular control over data flows—allowing administrators to reduce the adversary's attack surface and enhance overall network security.



DISCLAIMER OF WARRANTIES AND ENDORSEMENT

The information and opinions contained in this document are provided "as is" and without any warranties or guarantees. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement, recommendation, or favoring by the United States Government, and this guidance shall not be used for advertising or product endorsement purposes.

CONTACT INFORMATION

Client Requirements and General Cybersecurity Inquiries
Cybersecurity Requirements Center (CRC), 410-854-4200, email: Cybersecurity_Requests@nsa.gov