AI-Based Microsegmentation Design

Authors:
Mark Viglione
Michael Hoehl

January 16, 2023

1. Executive Overview

Several data protection regulations including Cybersecurity Maturity Model Compliance (CMMC) and Payment Card Industry Data Security Standard (PCI DSS) require network layer controls to define the boundary for internal systems included in scope for compliance. Companies aspiring for Zero-Trust depend on this same foundation of network design. Further, Cyber Insurance rates are impacted by the effectiveness of network separation of systems of different business risk levels. This internal network enclave design, required in both physical data centers and virtual cloud environments, is called Microsegmentation.

The old approach of placing a firewall only at the perimeter between the Internet and internal networks is no longer sufficient for compliance or modern secure computing environments. Without Microsegmentation, obtaining and sustaining compliance can be impossible. The number of systems that must comply becomes elastic when unrelated systems share the same network. Therefore, controlling scope is essential to reduce audit expense and avoid audit fatigue as a result of recurring compliance validation of an unnecessarily large number of systems.

Microsegmentation is a daunting effort for organizations large and small. It is often a manual, time consuming and expensive process to implement and maintain. A significant amount of discovery and research are required to understand the current network flows. Traditional approaches also require multiple teams (e.g., firewall admins, application owners, cybersecurity, etc.) working together to collect the required data to successfully segment their network. On top of that, the process can be especially challenging for organizations that have legacy or specialty systems (e.g., Industrial Control Systems, Hospital Clinical Systems, etc.) that are unable have a software agent installed on them to collect data for baselining and firewall ruleset creation/updates. In many cases, the volume of research data generated is extremely large. This raw data must then be transformed into new network design and firewall policies. Organizations need a new approach to implementing and sustaining Microsegmentation that is less challenging. This paper outlines a modern Artificial Intelligence based solution that saves time, saves money, increases credibility, and ultimately reduces risk.

2. Why Microsegmentation?

2.1. Reduce exposure and risk

A fundamental requirement to protect data is to build a network architecture that logically separates systems of different levels of trust and data classification. To accomplish this, information systems of similar value and function are placed into unique network enclaves or zones. Pinch points are then created between each enclave with the use of firewalls. This zone-based network design is commonly referred to as Microsegmentation. This network design makes it easier for IT Operations and Cybersecurity Teams to restrict network access as well as trend network traffic. The baselining provides valuable insight into what is normal and what is an irregularity. A popular mantra for many threat hunters when performing digital forensics is, "to best find evil it is necessary to know normal for the environment." An enclave or zone-based network design helps incident responders and threat hunters rapidly recognize anomalous activity on the network that could indicate harmful actions and intruder presence. This design approach also provides IT Operations teams valuable insight into unplanned change and service disruption.

Isolating similar systems into network zones also allows integration of additional access, detective, and preventive controls (e.g., intrusion detection, data loss prevention, malware scanning, etc.). In the event of a cybersecurity incident, these same pinch points serve to contain threats and limit the lateral propagation of attacks across multiple internal networks (e.g., from enterprise networks to manufacturing networks).

Without Microsegmentation, organizations are not able to meet compliance or contract requirement obligations and may not conduct business with certain customers (e.g., DOD, Bank, etc.). This can represent a significant lost business opportunity. Further, Cyber Insurance rates are substantially higher when Microsegmentation is not in place. Ultimately, Microsegmentation opens the doors to business opportunity and reduces direct costs.

2.2. Implementing Microsegmentation is a compliance requirement for many industries

Modern data protection compliance obligations typically apply to anything that transports, stores, or processes specific regulated data. Several data protection compliance standards require firewalls and similar network layer controls to define scope boundary to systems that must meet compliance standards. These zone-based compliance standards include Payment Card Industry Data Security Standard (PCI DSS) for retail industry, Cybersecurity Maturity Model Compliance (CMMC) for Department of Defense supply chain, and Society for Worldwide Interbank Financial Telecommunications (SWIFT) for Banking. Without Microsegmentation and firewalls, the systems included in scope for compliance become elastic and unintentionally may include unrelated systems sharing the same network VLAN/subnet that do not transport, store, or process the regulated data. Sustaining compliance standards can be very costly. Further, recurring compliance validation for an unnecessarily large number of systems can result in audit fatigue.

Payment Card Industry (PCI) Security Standards Council mission is to provide guidance to protect payment cards and advises, “Scoping involves the identification of people, processes, and technologies that interact with or could otherwise impact the security of Cardholder Data Environment (CDE). Segmentation involves the implementation of additional controls to separate systems with different security needs.” (PCI SSC, 2016). Therefore, scoping considerations include all system components included in or connected to the CDE.

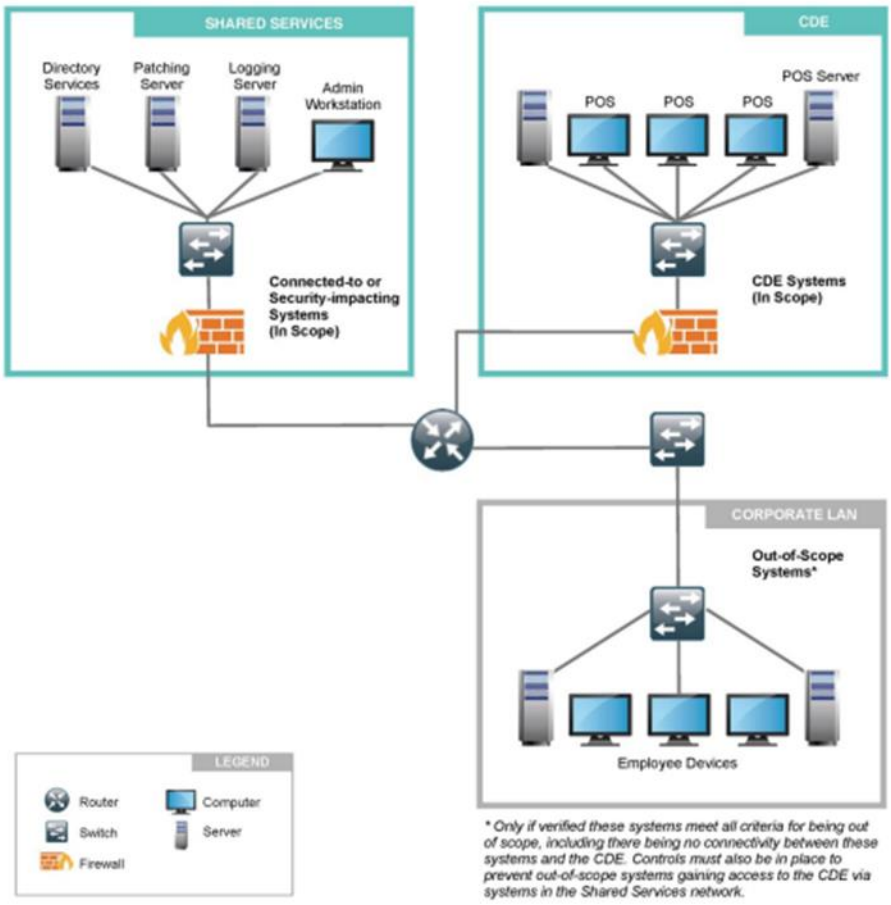

The Cybersecurity Maturity Model Certification (CMMC) model is designed to protect Federal Contract Information (FCI) and Controlled Unclassified Information (CUI) that is shared with contractors and subcontractors of the U.S. Government. CMMC is based on National Institute of Standards and Technology (NIST) Special Publishing 800-171, which provides the following guidance: “organizations may limit the scope of the security requirements by isolating the designated system components in a separate CUI security domain. Isolation can be achieved by applying architectural and design concepts (e.g., implementing subnetworks with firewalls or other boundary protection devices and using information flow control mechanisms).” (Ross, Ron, 2020). Both compliance standards emphasize that out-of-scope assets must be physically or logically separated. This means separate networks.

In addition to the requirement for separate networks, all the systems that co-exist within the same network must meet the compliance requirements—even if they do not transport, store, or process the regulated data. PCI SSC provides clear direction with the

following scoping concepts (PCI SSC, 2016) that are universally accepted for other compliance obligations (e.g., CMMC):

- Systems located within the CDE are in scope, irrespective of their functionality or the reason why they are in the CDE.
- In a flat network, all systems are in scope if any single system stores, processes, or transmits account data.
- Similarly, systems that connect to a system in the CDE are in scope, irrespective of their functionality or the reason they have connectivity to the CDE.

3. Why is sustaining Microsegmentation a challenge?

3.1. Firewalls require recurring weeding and feeding

No network environment remains static. Changes are a natural and common occurrence in today's computing environment. All of these changes are not always anticipated and presented to the firewall admin to consider well in advance. Routine system upgrades, functionality enhancements, high-availability/cluster fail-over changes, and scaling of services up/down all cause changes in network traffic patterns. This dynamic environment condition subsequently results in required changes to the firewalls. Therefore, firewall rulesets must be reviewed routinely to ensure they align with the current environment and intended safeguards. This is addition to the periodic review of firewall rules performed for demonstrating compliance.

Proper administration and operation of a firewall after implementation is necessary to sustain the intended preventive control. Firewalls are not intended to be set and forget. The best designed networks and firewalls can be defeated by a cyber-attack when routine administration and monitoring is not performed. Unlike some technology (e.g., routers and switches), a firewall that is "broken' does not necessarily result in a poor user experience. When a router or switch is down, the Helpdesk typically gets a surge in calls from employees or customers raising awareness of a problem. For firewalls with poorly designed rulesets, a cyber-criminal will not alert an organization to the security weakness. Like all systems, firewalls also can have defects. Routine updating and patching are also essential for firewalls.

A firewall without routine monitoring is the same as a building surveillance system without someone responsible for looking at what the camera has captured. A web application firewall that is not monitored for occurrence of Cross Site Scripting or SQL Injection attacks will only serve as a postmortem forensic tool to explain how the compromise occurred after the harm is done. A network firewall without routinely reviewed firewall rules will not meet compliance requirements. As mentioned earlier, firewall performance also needs to be monitored to identify potential resource issues so they are addressed prior to being overwhelmed.

3.2. Continuous firewall performance trending and network traffic baselining

In addition to routine rule review, network traffic presented to firewall and the resulting performance of the firewall should be baselined. Rerouted network traffic can have the unintended consequence of tilting a firewall that was previously capable of handling traffic volume. High frequency of hits on a previously idle rule, change in typical bytes/hour, emergence of denied traffic, surge in connection attempts, etc. all have an impact on the firewall performance and effectiveness.

As mentioned earlier, no network remains static. Unfortunately for the firewall admin, firewalls are not a once-and-done configured technology. Knowing what is normal with the use of baselining traffic and trending performance is essential for firewall admins. For Cybersecurity, this knowledge improves speed to identify indications of compromise or attempts at data exfiltration. For Operations, this knowledge helps identify well intended changes to systems and applications (e.g., system updates, failover, patching, etc.) with unintended results. Ultimately, this historical perspective is vital for first responders to rapidly determine if they are facing a friend or foe.

Lastly, baselining and trending is a fundamental IT Operations service to proactively forecast future necessary financial investment. As new systems are being planned or brought on-line, does the compute capacity of the firewall need to be increased too? This advanced insight is critical when there are only a few change windows a year (e.g., manufacturing environments, customer-facing systems, etc.) or opportunity to ask for investment once a year.

4. Why isn’t Microsegmentation already in place?

Though easily stated and understood, executing Microsegmentation is very challenging for most organizations. It can be challenging to introduce firewalls into a production environment that did not have the requirements for network segmentation and restriction when originally implemented. Existing production systems that now require firewall restrictions may not have the necessary network related technical details documented or understood. The project team and subject matter experts that implemented the applications and systems may no longer be available. Updates and patches may have introduced new functionality and data flows. Documentation created as part of the original implementation project may no longer reflect the current as-built design of all the systems or applications added to scope as time has passed. A firewall ruleset design approach that solely depends on Application and System teams providing change requests for essential network communication requirements of firewalls may not lead to a successful outcome.

To reduce this risk, many organizations attempt to document network flows using tools like WireShark or TCPDump. This results in a huge amount of data that must be combined and then reviewed. Interpreting the data requires a highly skilled individual that understands how to read packet captures and create traffic maps. In some cases, an external consultant is engaged to perform this service. Unfortunately, this individual can only provide limited context as they do not understand the operating environment.

5. Solution

5.1. Overview

The solution is an agentless based Artificial Intelligence (AI) solution capable of collecting and analyzing internal network flow data. The solution baselines traffic patterns and classifies internal systems. After learning what is normal, the solution advises on network design and necessary firewall ruleset creation for implementing Microsegmentation. It is important to note that the solution only uses header data from the captures. Analytics do not ingest or use full packet data. The captured data is transformed into flows so that sensitive data is not exposed.

The solution is broken down into three parts:

a) Microsegmentation implementation assistance
   i) Intelligence backed by Machine Learning (ML) to help organizations redesign their internal networks to zone-based compliant networks.
b) Firewall ruleset maintenance assistance
   i) Once Microsegmentation has been successfully implemented in the environment, the firewall still needs to be updated as new network patterns emerge in the segmented enclave. This solution detects these new patterns and advises on necessary changes.
c) Ongoing network traffic baselining and cybersecurity monitoring
   i) To help detect new operational changes as well as indications of compromise

5.2. Deployment Options

There are multiple options for deploying the solution depending on the environment. The required data for analysis can be acquired via any of the following methods.

4.2.1. A SPAN or MIRROR port (create a copy of selected packets passing through a device)

4.2.2. Network TAP (ingest data from a specific Test Access Point)

4.2.3. PCAP file

5.3. Client Dependences

In order for the solution to properly work, there are a few dependencies on the client side that need to be set up prior to deployment for a successful implementation.

- A Windows 10/11 host (capable of running Docker and Pktmon)
- Docker needs to be installed and running on the Windows 10/11 system
- An approved remote monitoring tool/service needs to be installed on the Windows 10/11 system for maintenance and troubleshooting
- Secure internet access (to ensure the data can be sent to the processing engine)
- Network flow data (collected via the aforementioned Deployment Options)
    - SPAN port
    - Access to a TAP
    - PCAP file to be stripped of sensitive payload information

References

Baykarapci, Surkay (2020) *What are the Firewall Requirements for PCI DSS? - PCI DSS GUIDE*, https://www.pcidssguide.com/pci-dss-firewall-requirements/ (Accessed January 1, 2022)

PCI Security Standards Council (2016) *Information Supplement: Guidance for PCI DSS Scoping and Network Segmentation*, https://www.pcisecuritystandards.org/documents/Guidance-PCI-DSS-Scoping-and-Segmentation_v1.pdf (Accessed January 2, 2022)

Ross, Ron etal. (2020) *Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations (NIST SP800-171r2)*, https://doi.org/10.6028/NIST.SP.800-171r2 (Accessed January 3, 2022)

Scarfone, K. and Hoffman, P. (2009) *Guidelines on Firewalls and Firewall Policy, Special Publication (NIST SP800-41)*, National Institute of Standards and Technology, Gaithersburg, MD, [online], https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=901083 (Accessed January 1, 2022)