Carnegie Mellon University
Software Engineering Institute

COMMON NETWORK SECURITY TOOLS AND CAPABILITIES

Timur Snoke
January 2019

Overview

Networks, systems, and applications need to be defended but the defense of a business or an organization is more than just capturing all the network traffic or hiring reverse engineers. A business must defend its profitability and limit its exposure to risks. The technological threats to an organization are growing at an increasing rate and the tools being developed to prevent or mitigate compromise, theft, and damage are becoming more sophisticated in an attempt to keep up. It is becoming increasingly difficult to understand both the elements within the enterprise that are at risk and just how porous the perimeter is. This document is intended to provide a basic understanding of the network security elements available and how their capabilities can be composed or coordinated to defend not just the network but the mission of the organization. It is not intended to cover the tools that perform incident/event management, merely classification and automated response in inspection points.

Network Security Stack

Monitoring organizational threats requires the instrumentation of inspection points at the ingress and egress points of the network as well as within the network. The combination of security appliances and services defining the security posture and its enforcement at an inspection point on the network can be referred to as a network security stack. The network security stack has to be tailored for where it is located, the volume of traffic being analyzed, and what the enforcing security policy is. The location of the security stack needs to be considered when designing the inspection point. An external network security stack has to be able to perform even with all the noise on the internet and an internal network security stack needs to be able to perform without impacting the volume of traffic in a high-performance network like a data center. Security policies should be in place to set expectations on how systems are to be administered and what end users can and cannot do with them. Much of the guidance for end users in an organization is covered in the Acceptable Use Policy, the guidance for how systems should be operated should be provided by the CISO. All three of these considerations; sensor placement, traffic volume, and guiding security policy, inform the design choices that will be used to select the appropriate combination of security controls.

There are two different policy enforcement approaches that are used by the network security stack, active or passive defense. The security policy should clearly indicate if the goal for each of the security controls is to passively monitor and alert or to monitor and actively respond. It is also important to be clear as to what the area of responsibility is for security controls and who is responsible to monitor the alerts.

Network security stacks can be external or internal. An external stack is placed outside the corporate network and is directly exposed to the internet ingress/egress point(s). An internal stack would be placed between security domains but not exposed to the internet directly, such as between the DMZ and the inside network or on the private connection between HQ and a remote office. The external network security stack has to be able to perform under the stress of all the latent noise on the internet without impacting the revenue-generating services of the corporation. This stress can be the result of network scans, backscatter from responses to address spoofing, attempted denial of service, or even unexpected surges in traffic for less malicious reasons. The internal network security stack may have to be a high-performance solution, as in the data center, or a simple blocking capability that explicitly denies traffic by default and only allows traffic by declared exception, as in the case for a guest wireless service. The external stack protects the network from the internet and the internet from the network while the internal network security stack protects the network from itself.

Both types of network security stacks have benefits and limitations. For example, an external network security stack will by its very nature not be aware of the traffic inside the encrypted tunnels that it is passing but it can observe when traffic volumes become congested and start to force traffic to be dropped and retransmitted as a result of interface saturation. An internal network security stack should understand better what the applications can and should be doing and be able to identify when a system is acting outside of the norm or when an internal client is making inappropriate requests from a local server. The external network security stack would not have high confidence blocking capabilities based upon the inside systems without knowledge: for example, an organization's primary ISP doesn't know how many web servers the organization has. The inside network security stack would have more insight into the nature of the communications and could identify features of concern that are closer to unintended application use and abuse: for example, a production system that produces web content should not also function as an NTP server. Both types of stacks provide value to an organization because they can report on the issues that they are best positioned to be authoritative of.

The network security stack can be replicated at similar inspection points across an enterprise. These network security stacks can work independently or be federated in providing capabilities. The federated security stack allows for security policy synchronization by using a common source for administration controls and also serve as a common destination for logs and alerts.

There are times when the network security stack is not able to be placed in the path of the traffic that is to be observed or there are many paths that need to be observed simultaneously by a shared capability. In these cases there are hardware-based and software-based solutions that can provide a means to inspect the traffic in a different location. Physical network taps can be placed in-line with data that needs to be inspected and will passively copy all traffic and direct the duplicate traffic to the network security stack. Some physical taps are designed to filter traffic and direct a subset of the observed traffic to specific collection capabilities. Span ports are a software solution that is configured in the network switching infrastructure to replicate the traffic traversing one set of ports or VLANs to go out another interface into the network security stack. There are network tap solutions that are able to direct the duplicated traffic to a network security stack that is not collocated with the inspection point. Span ports and network taps that redirect network security traffic to a remote network security stack is not

optimal for security functions like blocking because it is working on a copy of the traffic instead of the source.

A network security stack has four different capabilities that are tailored to support the defense of the inspection point for an organization, as shown in Figure 1. These capabilities are filtering, inspecting, proxying, and logging. Different tools are used to accomplish one or more of these functions in the network security stack. The most common tools include, but are not limited to, firewalls, proxy/load balancers, intrusion detection/prevention systems, and antivirus. These capabilities can exist in a single platform or require traffic to be bussed from one application/platform to another to facilitate the inclusion of non-native capabilities into the main security platform. These non-native capabilities can provide additional features which could be too computationally expensive to perform on the primary inspection platform. These non-native capabilities could exist on a purpose-built card or blade in a system or as a separate discrete system that has access to the data passing the inspection point. This not only allows missing functionality to be added, but can offload computationally expensive functionality to limit the impact of the inspection on network throughput or performance.

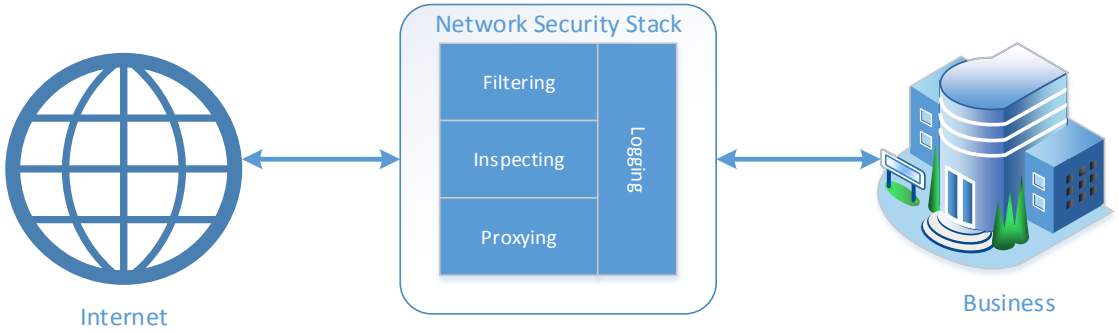

Figure 1 - Network Security Stack

The performance of the security stack is limited by the slowest worker process. With this in mind it is desirable to limit the noise that is evaluated by the security tools. A passive DNS collection capability does not need to see NTP traffic, for example, so some level of pre-filtering would be appropriate.

A solution that is often employed for security tools that are not able to perform at interface speeds is load balancing. Load balancers are used to distribute traffic across additional worker processes; this is also referred to as scaling horizontally. Instead of backing up the traffic to be inspected on a single worker process, adding more capacity enables traffic analysis at line rate.

For example, a network security stack could include a firewall and intrusion prevention system (IPS). The firewall may proxy public IP addresses for inside hosts to get internet access. It can also be configured to use an external service to lookup the reputation of outgoing web requests to enforce the organization's Acceptable Use Policy. Adding the capability of the external web filter to the firewall, where the web filter evaluates and returns instructions to block by the firewall, is sometimes called a sidecar. After passing the firewall, the traffic can then be fed through the IPS to inspect the filtered traffic and block/log/alert on triggered rules.

Adding the IPS behind the firewall is in effect stacking the capabilities and having the post-filtered data from the firewall be evaluated by the IPS. By adding a capability as a sidecar that utilizes an external capability or looks up information that is outside of the primary assessment platform it is possible to add/remove/modify additional features without having to completely re-engineer or disrupt the network security stack. Leveraging a load balancer allows additional worker processes to be added as needed to support the volume of traffic to be analyzed. The different configurations of network hardware in the network security stack are presented in Figure 2.

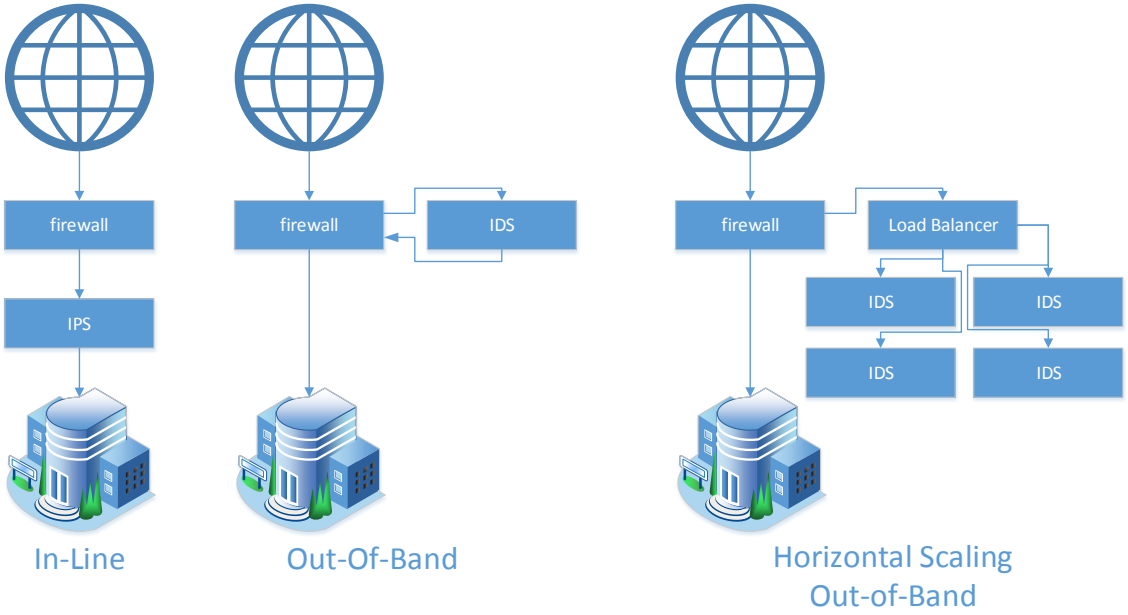

Figure 2 - Network Security Stack Hardware Configurations

The network security stack is the combination of capabilities that are orchestrated at an inspection point, internal or external, that provides the ability to observe and enforce compliance with corporate security policies. The following sections will go into more detail about the discrete components that can comprise the network security stack.

Network Security Stack Components

A network security stack is composed of different tools that are used to support the defense of the organization. These capabilities are coordinated to analyze and mitigate the traffic before it gets to assets that are to be protected. Network security stack can have components that perform one or more of the four different capabilities, (the following list is not comprehensive):

- Proxies - filter, proxy, and log
- Intrusion detection systems/intrusion prevention systems - inspect, filter, and log
- Network security monitoring systems - inspect and log
- Firewalls - proxy, inspect, filter, and log

- Network Access Controls - proxy, inspect, filter, and log

Proxy - A proxy is when one device acts as an intermediary for communications between endpoints, most often other host-based clients and servers. Network proxies are used for a variety of reasons: initially most involved efforts to accommodate scarce resources, to conserve IP addresses or internet connections. A network proxy makes it possible for a single internet connected device to share internet access to a group of internal hosts and appear to the outside world as being a single entity. Another application of a proxy would be as a load balancer that distributes the volume of client requests among a pool of worker processes that provide the same service. By its very nature the only way to determine what inside resource was being proxied is through proxy logs.

Proxies can conserve IP addresses by using network address translation (NAT) or port address translation (PAT) to proxy traffic between an internal host and the internet. With NAT an external routed IP address is mapped to an internal IP address; this can be a dedicated mapping one-to-one or to a dynamically allocated pool of publicly routed IP addresses that are assigned and re-allocated as needed. PAT on the other hand maps the internal IP addresses to a single publicly routable IP by mapping the ports used to the individual host. NAT is typically utilized when bi-directional communication is desired for services and applications, while PAT is utilized when clients need internet access but not to be accessed by the internet.

An advantage to using a proxy as a load balancer for a web application is that you can have an encrypted tunnel that terminates on the appliance and redirects traffic to the pool of presentation servers while creating the appearance of a single source of content, so each server doesn't need its own certificate. The down side to this is that the traffic becomes unencrypted and the content becomes exposed for the rest of its journey unless there is a second encrypted tunnel connecting the proxy to the presentation server, which removes the stated benefit of encrypting the traffic. The benefit is diminished because the dual encryption creates new overhead for key management and coordination as well as requiring security controls to protect the security controls because they are now putting the protected traffic at risk during un-encryption at rest.

IPS/IDS - Intrusion Prevention Systems/Intrusion Detection Systems are meant to find malicious content in traffic that traverses an inspection point. The primary difference between the two types of systems is that while both detect activity of interest, IDS only alerts but IPS can both alert and block the detected traffic. Another difference is that an IDS does not need to be in-line with the traffic; it can be placed out-of-band by forking a copy of network traffic to the appliance, sometimes called a one-armed deployment. IPS placed out-of-band are incapable of blocking traffic and therefore would just act as an IDS. Out-of-band deployments are often used when the volume of traffic traversing the inspection point is great enough that traffic inspection will introduce latency, in which case the IDS as a worker process can be scaled horizontally with a load balancer acting as a scheduler to distribute the traffic to multiple appliances or systems to provide the inspection capability. A load balancer would not be used to scale an IPS because it is not working on the original network traffic, but a copy. The potential for latency introduced by the load balancer lessens effectiveness of the IPS blocking capability because the response could traverse the network security stack before the connection is reset.

NSM - Network Security Monitoring provides analysis of the network traffic traversing the inspection point and performs network information collection. This information collection can include features that provide enriching details to inform network analysis, such as, artifacts of network traffic in the form of structured data or logs that can be ingested into a database. These artifacts can provide inventories of server and client applications and their associated version information, SSL certificates, and logs of a wide variety of network traffic types. The goal of an NSM is to both alert on network traffic of interest and provide contextual information around those alerts. This allows analysts to more quickly answer questions about what is happening and to determine without more manual research if action is required.

Firewalls - A firewall is an appliance or an application running on a system that is used to filter/restrict access to an endpoint or a service that is traversing a chokepoint between one security domain and another. This could be an appliance that is placed between the internet router and a private network or host-based firewall defending the interface between a network and a host. This traffic can be filtered by any one of a number of features that include but are not limited to: source IP, destination IP, port and protocol. Firewalls can also use blacklists as filters. In one application a web-aware firewall will pass a requested URL off to an external service to evaluate the destination. If the destination is determined to be "bad" the connected session gets terminated by the firewall and logged, effectively blocking the communications. The blacklist can be resident on the firewall and maintained by the administrator, or it can be automated as an external resource that gets polled and updated during some interval. Early firewalls were limited in how much inspection they could do so additional features were offloaded to purpose-built systems that are added as modules, sometimes called sidecars, to the firewall. More advanced firewalls exist: application and next generation.

Application firewalls - There are application-aware firewalls that understand how programs work and enforce compliance by evaluating the actual communications associated with a service instead of just evaluating IP addresses and ports. Some application-aware firewalls also will act as proxies to terminate an application session and instantiate a new session from the firewall to the intended destination. A use of this is for "break and inspect"–terminating encryption at the firewall to allow traffic inspection before re-encrypting with a new connection to the intended destination. These firewalls can also enforce application compliance and form validation so undesired inputs are not submitted to the destination application. Application compliance would seek to validate that the commands being sent past an inspection point are valid for the intended use of the application being protected. An example of form validation would be evaluating that the content for a field being submitted is appropriate before allowing it through to the application.

Next Generation Firewalls - Next Generation Firewalls, also called third generation firewalls, include additional capabilities beyond blocking traffic at the network layer. These features can include in-line deep packet inspection that is application aware, intrusion detection/prevention capabilities, SSL/TLS traffic inspection, website filtering, and other signature matching efforts such as antivirus. These features can be implemented in-line or out-of-band on the traffic traversing the inspection point.

Network Access Controls - NAC provides the ability to enable policy enforcement on end-points before allowing access to the network. As part of that authorization a NAC solution will evaluate and remediate end-points to enforce policy compliance. NAC is implemented at the network level with a

client-side assessment tool. If NAC is implemented and a computer tries to connect to the network, it is first placed into a quarantine while a NAC assessment takes place. If the computer is compliant, it is removed from quarantine, AKA granted network access. If not, the system is given the opportunity to remediate and allow access to the patches or files that would make the system compliant. The NAC client can evaluate multiple features about a system, such as the local users/administrators, patch level of the operating system, and the security controls installed and their update status.

Related host-based Security Components

Host-based security controls can be emulated in network security stacks to evaluate communications and files before they make it to the end-point by directly analyzing files or using a sandbox to emulate what would happen on the end-point. Host-based security components are frequently applications that are installed on end points and seek to defend hosts from being compromised. The goal is to provide the ability to identify malicious files or processes that are not authorized to run on a local system before they get to an end point. Here we are only discussing host-based controls often emulated in a network security stack, so the coverage is not exhaustive.

Antivirus - AV is often a host-based inspection capability that looks at software and related processes to detect, prevent, and remove malicious software from a system, anti-virus can be applied in network stacks as well to data in motion. Antivirus applications are also well positioned to identify software that has undesirable effects, such as: Remote Access Trojans, Worms, Cryptoware, Ransomware, and keyboard loggers. The challenge for Antivirus is that the volume of malicious binaries that could be detected is growing at an exponential rate and it is impossible to have static signatures available for all of them. A combination of approaches are used by different AV vendors to accommodate these circumstances based upon pattern matching, trending, and behavioral analysis.

Host-Based IDS - Host-based IDS are similar to their network based counterparts in their purpose of detecting and alerting on compromise or undesirable behavior. There are a variety of different approaches that are used to accomplish this task. Some emphasize capturing a known state for files and alerting on any change in the files stored on disk. A different strategy is providing a shim between the applications and the operating systems and will alert if an application attempts to engage in activities that are deemed inappropriate. The challenge with many of these approaches is that finding a deviant activity require knowledge of what normal is. Looking for changes in files is not trivial when you have temp files and cache files that are almost always in an unknown state. Every time an application gets patched the HIDS needs to be updated to support the changes that occurred. Understanding what is the appropriate behavior for a single application running on a system, is a challenge when most systems run multiple applications concurrently.

Summary

The network security stack is made up of a diverse group of tools that are used to monitor the traffic that is traversing an inspection point. The tools are selected and configured to support the security policies of the organization and mitigate/reduce risk to network assets, while generating network defense

alerts (or data). Future papers are intended to highlight placement and configuration of network security stacks to emphasize efficiencies.

Contact Us

Software Engineering Institute
4500 Fifth Avenue, Pittsburgh, PA 15213-2612

Phone: 412/268.5800 | 888.201.4479
Web: www.sei.cmu.edu | www.cert.org
Email: info@sei.cmu.eduCopyright

This material is based upon work funded and supported by the Department of Homeland Security under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center sponsored by the United States Department of Defense.

This report was prepared for the SEI Administrative Agent AFLCMC/AZS 5 Eglin Street Hanscom AFB, MA 01731-2100

DM19-0157