SANS

Analyst Program

Whitepaper

SANS Protects: The Network

Written by **Matt Bromiley**

August 2022

corelight

Introduction

One element found in every enterprise around the world is instrumental to linking various asset types together, whether they are on-premises or in the cloud, local or remote. The network itself—the links between assets within an environment—is what unites clients and servers, employees and customers, and enables business-to-business communications. Enterprise networks are the backbone of business and commerce—we cannot expect devices or users to be able to communicate with one another unless there is a medium for that communication to occur.

All of that means, however, that the network attracts adversaries looking to execute their attacks and achieve their objectives. Malware must communicate, adversaries like to move laterally, and data needs to be exfiltrated to be used against the victim. In all these cases, the network is the enabling medium. Thus, it is also a key point for observing and detecting suspicious or malicious activity. Often, it is the earliest.

In this SANS Protects paper, we look at threats to enterprise networks and ways that your organization can overcome or mitigate them. Our SANS Protects papers focus on threats and mitigations, helping organizations consider elements of security that they should be implementing today. Our goal is to start or enhance the conversation about what assets are in your environment and the corresponding protections that are in place.

Given the recent success we've seen among adversaries, the network is an asset that simply must be protected. We highly recommend that organizations place network detection and response technologies at the top of their list of security priorities. Some of our key takeaways from this whitepaper include these:

- Adversaries often find success in misconfigured networks, but simple changes such as network segmentation can prohibit widespread attacks.
- Network connectivity is vital for nearly all malware, and this creates a choke point that defenders can use to their advantage.
- The enterprise perimeter is changing, and the ease of network access is a critical component to the growth of breaches. With the right visibility, security teams can gain a huge advantage over adversaries.

Adversaries continually invent new ways to abuse or take advantage of an enterprise's network. Security teams are often looking for new and novel ways to detect and respond to these threats—and we believe this need will continue to be a constant one. The technology is available, and we must use it to keep up with whatever threats adversaries may throw at us. This paper can help you to analyze your current network monitoring tools and capabilities or to provide guidelines for an upcoming integration.

Threats to Enterprise Networks

The key thing to remember about threats to an enterprise network is that they may not be threats to the network itself. For example, the first type of threat we'll look at is ransomware. Ransomware is not a network threat, but it relies entirely on the network for success. For that reason, we must understand how we can utilize the network to detect attacks such as ransomware. Network-based attacks specifically take advantage of network devices or network protocols to affect an organization. In this section we look at both. When deploying technologies to assist, however, ensure that your security team knows whether they are stopping a network threat or a network-based threat.

Know Your Threats

Some threats are threats *to* the network, whereas others are threats that *use* the network. Understanding the difference is key to deploying efficient detection, response, and countermeasure capabilities.

Ransomware

Ransomware is the first threat that we will examine from a network perspective. Note, however, that it is often an all-encompassing threat. Ransomware includes servers, workstations, networks, data, and user accounts. It can include an entire Active Directory or an organization's entire cloud footprint. Either way, the attack would not be possible without the network. The necessity of the network in a ransomware attack is both a strength and a weakness for defenders.

Figure 1 depicts a ransomware attack. We'll use this as a guide to view the criticality of the network in the execution of these attacks.

The key to adversaries' success is the ability to move between an initial access point or compromised system to the rest of the environment. As we can see, at Step 1, a malicious email can target one or many users, though it is just an initial entry vector for the adversaries. They have compromised one or a small handful of systems. However, as we see in Step 2, flat or misconfigured networks or easy-to-obtain administrative credentials allow the adversaries to view and access other systems in the environment, furthering the scope of impacted systems and the depth of their attack.

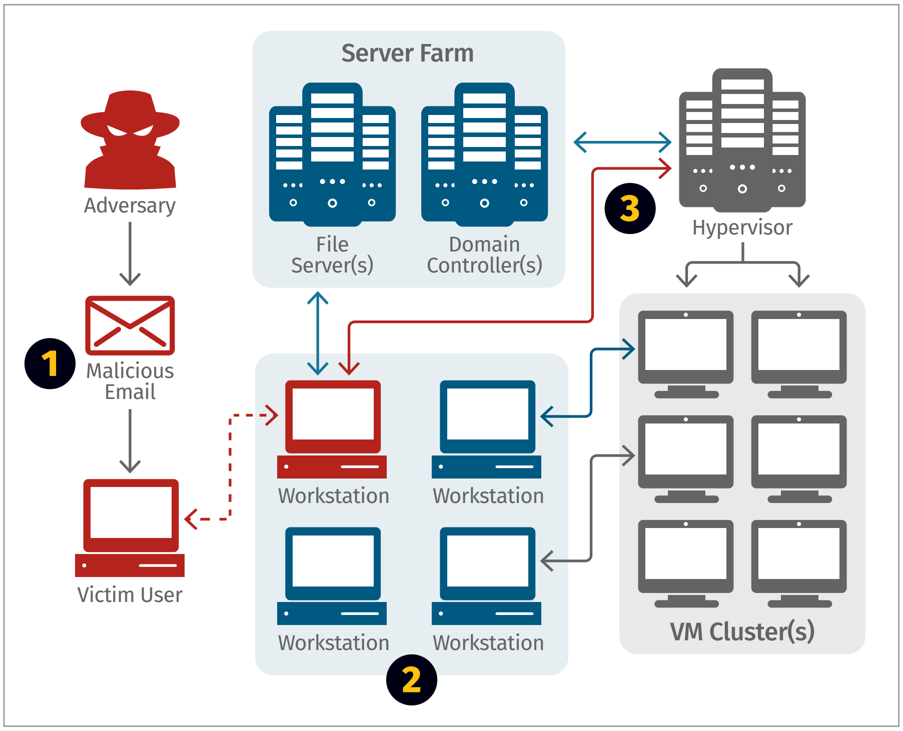

Figure 1. High-Level Ransomware Attack

With the correct access, moving from one to many compromised systems is an easy feat. Step 3 highlights a common goal of a ransomware threat actor: going after systems critical to business operations, such as file servers, or systems key to recovery operations, such as hypervisors and backup servers. Within an enterprise, these systems *must* be able to communicate with one another. That link provides adversaries an easy opportunity for compromise and/or impact.

Perimeter Network Devices

Another inherent threat to the network is the fact that networks are both internal and external for every organization. By this, we mean that the same "pipes" move data internally, from workstation to workstation, and serve as an entry or exit vector. Compare this to the roles between an internal- and external-facing server; their roles may be different, but their placement within an organization makes them susceptible to different types of attacks.

Say Goodbye to the Perimeter

The security industry seems to be moving away from the concept of a perimeter, or at least a "traditional" network perimeter. Organizations are starting to embrace models such as secure access service edge (SASE) and zero-trust network access (ZTNA) to limit how much damage an adversary can cause at the perimeter, even with valid credentials or an access mechanism.

Therefore, just as we must consider lateral movement as an internal network threat, we also must consider any adversary coming from the outside as an external network threat (or a perimeter network threat).

Disabling and Destructive Attacks

Adversaries will often look to the network to cripple, disable, or even cause destruction to a victim organization. When we think of disabling or destructive attacks, one of the first that come to mind are DDoS attacks. These are a straightforward protocol abuse. They take advantage of a particular network protocol and point a significant number of resources at a victim organization or a victim IP space—so many resources, in fact, that the victim organization cannot keep up. Figure 2 depicts a typical DDoS attack.

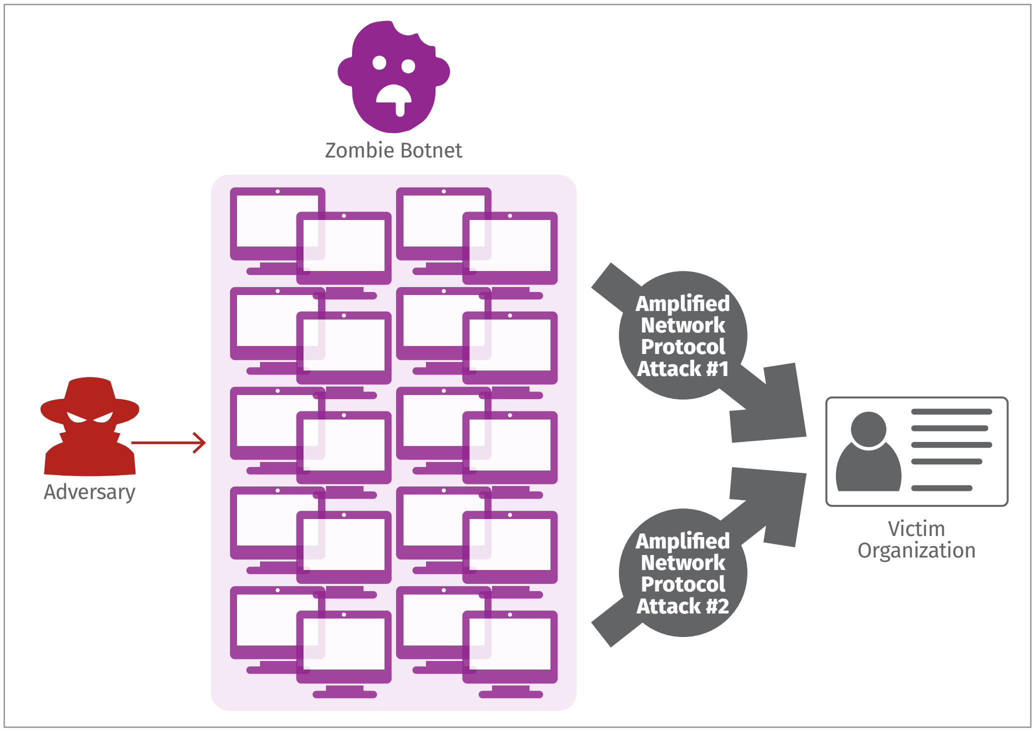
Zombie Botnet
Adversary
Amplified Network Protocol Attack #1
Amplified Network Protocol Attack #2
Victim Organization

Figure 2. DDoS Attack Involving a Zombie Botnet Controlled by an Adversary

This high-level diagram depicts how powerful a series of interconnected systems, such as a botnet, can be. With a *single* command, an adversary can instruct thousands or millions of zombie systems around the globe to launch network protocol attacks, such as DNS amplification attacks, that are designed to overwhelm the victim organization. Furthermore, they can sustain the attack with little retribution, because they exist behind multiple veils of anonymity.

DDoS attacks are no joke. They can bring a business to its knees or shut down network traffic for a period of time. In October 2020, Google suffered a DDoS attack from three Chinese ISPs that lasted for six months and peaked at 2.5Tbps. And we must always consider whether the DDoS attack is the focus of the adversary or if it is a distraction while the adversary takes advantage of another weakness or entry vector in the environment.

Flat or Minimal-Depth Architecture

The architecture of an enterprise's network can also pose a threat to network security. Adversaries have been particularly successful in enterprise networks that are either flat or maintain minimal depth. Again, this architecture itself may not be a threat to the network, but the way that networks are set up inside an organization can determine the success of an attack. Depth, ease of movement between segments, and other considerations should be factored into an organization's risk profile to determine how easily an adversary would be able to execute their attack.

Let's use ransomware once again as an example of how dangerous a flat network can be. In a ransomware attack, as shown in Figure 1, an adversary's primary objective is obtaining access. A very close No. 2 is to find credentials that would allow the adversary to move laterally throughout the network. In an unsegmented network, this is as easy as gaining access to the right administrator or domain to easily move from one system to hundreds or thousands. Even minimal network architecture is easily thwarted by adversaries.

Legacy Protocols

In our SANS Protects Enterprise Email paper, we explored how old or legacy protocols are a danger to email security.[1] Unfortunately, enterprise networks are also susceptible to attacks based on the use of legacy protocols. A quick point of clarification here: When we say "legacy protocol," sometimes we are talking about protocols that are plaintext and so can be easily intercepted by adversaries and other times we are talking about protocols that should have been phased out long ago but still are widely used by organizations despite the danger of abuse by adversaries.

[1] "2022 SANS Protects: Enterprise Email," SANS.org, www.sans.org/white-papers/2022-sans-protects-enterprise-email, July 2022 [registration required].

HTTP is one example of a widespread protocol that sits at the crux of legacy and usefulness. There are plenty of reasons why HTTP continues to serve as a valuable, though vulnerable, protocol. It is in plaintext, unencrypted, and easy to observe. We can observe HTTP in full-packet captures or by enabling verbose logging on a suspect server. However, its use in driving API traffic and serving up information cannot be understated—it is a necessary protocol. Therefore, we build systems around this risk and insert protections, such as encrypting traffic as quickly as possible to prevent snooping on sensitive information, where possible.

FTP is another protocol that sits at this perilous junction. It's a legacy plaintext protocol that probably should be phased out because we have a secure protocol (SFTP) that can be used to encrypt and hide traffic from suspicious users or unauthorized listeners. If a secure protocol exists, why do we continue to see use of unencrypted protocols transferring sensitive information? The answer is relatively simple: We continue to observe *systems* built around legacy protocols still in operation without any future hope of encryption. Many times, this continues in an organization until an adversary forces a change.

Network Infrastructure Vulnerabilities

Lastly, threats are also delivered in the form of vulnerabilities to network architecture devices or software. In recent years, we have seen a plethora of vulnerabilities invade the information security industry. Unfortunately, network devices, such as VPNs and remote access tools, have their own set of vulnerabilities. For example, in May 2022, CVE-2022-20742, a severe vulnerability found in Cisco's ASA appliance and Cisco FTD software, could allow an unauthenticated adversary to read or modify data within an IPsec IKEv2 VPN tunnel.[2]

We must also consider potential vulnerabilities in the hardware that helps direct and facilitate network traffic. These vulnerabilities can include routers, switches, hardware firewalls, or any other device that may serve as one of the first hops in an enterprise network's ingress path. Load balancers fall within this risk category. In May 2022, a 9.8-severity vulnerability (CVE-2022-1388) impacted F5 BIG-IP load balancers and firewalls. Vulnerable instances would allow an adversary with access to the management interface to execute commands with root privileges.[3]

Protecting Enterprise Networks

Protecting enterprise networks can have multiple phases. All that's needed is a simple block to prevent a port from being opened or taken advantage of. One easy example would be disabling RDP access from external IP addresses. Conversely, some network defenses may be intricate setups that span hybrid enterprise models in multiple countries. For this reason, we have examined protections that we'd expect to find across any type of network—though applicability and ease of implementation will vary from organization to organization.

[2] CVE-2022-20742, CVE, https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2022-20742

[3] CVE-2022-1388, CVE, https://cve.mitre.org/cgi-bin/cvename.cgi?name=2022-1388

Network Segmentation

The most important recommendation for all organizations is network segmentation. Earlier in this paper, we underscored how vulnerable an organization with a flat network is to an adversary that gains access. That adversary is one correct credential away from being able to compromise the entire network. Network segmentation is an exercise in healthy network architecture and one of the best steps to prevent this type of one-to-many relationship from occurring.

Keep Network Privileges in Check

Network segmentation effectiveness can be bolstered by ensuring that users have least-privilege access in a network. The ZTNA security model, for example, provides users only the access they need to complete their job, using only the required resources. ZTNA can also limit the ability of a user (and thus an adversary who compromises that user's credentials) to move around a network, thus assisting segmentation in keeping systems safe.

Network Encryption

Regarding the use of legacy protocols, it is imperative that organizations move toward encrypted traffic as much as possible. Earlier, we discussed protocols such as HTTP and FTP that are easily accessible and captured. An organization that still utilizes a protocol such as plaintext FTP should consider moving to a more secure and/or encrypted version without delay.

Sometimes this shift can be difficult to accomplish, but businesses may find new options that will serve their business needs as well as adopting a more secure protocol would. For example, an organization that maintains its own FTP server for receiving and sending data may find that cloud storage meets its needs and offers multiple levels of encryption, data resiliency, and backup of data, if need be.

Traffic Inspection and Detection

With good network segmentation in place that makes lateral movement throughout an environment difficult, another valuable technology is one that offers traffic inspection and threat detection. Such technology is perhaps some of the oldest and most reliable security hardware or software available. It monitors network traffic and detects threats for security purposes.

Two technologies that inspect network traffic and detect threats are Snort and Suricata, but a plethora of proprietary network signatures have been created by the various network monitoring vendors in the space. There is a reason why an organization might deploy its own network monitoring and detection capabilities or gets in the practice of building its own signatures: Every enterprise network, while using the same protocols as others, has its own niches and nuances that only its security team knows about. Custom detections enable a team to create signatures optimized for the organization.

Custom detections are also a way of guaranteeing that your network segmentation and other preventative measures are working correctly. For example, if two systems should not have the ability to talk to each other per segmentation rules, then a custom detection focused on source and destination pairs would offer a high-fidelity "this shouldn't happen" detection. As another example, if your organization denies remote desktop protocol usage, then an easy `port 3389` would trigger an event that the security team must investigate.

Heuristic Profiling and Analysis

In addition to the detection of malicious traffic via signatures, it's a good idea to utilize an enterprise's network traffic, or metadata of that traffic, to provide a historical profile of traffic throughout the network. This approach can sometimes differ from malicious traffic detection. It is easy to see how security analysts and network engineers might benefit from the same types of traffic analyses.

For example, if an enterprise is using an unencrypted protocol inside of the environment (such as FTP when it could use an encrypted one such as SFTP), security analysts might want to know just how much traffic is being moved via that unencrypted protocol. Network engineers would want to know the same.

Network heuristics can also be used to threat hunt and find anomalies in the overall environment. It can sometimes be difficult to threat hunt inside an environment if you limit your searches to packet capture (PCAP) or a very specific request (such as requests to an IP). Instead, security teams will find greater success threat hunting through metadata and large statistics about network traffic inside of the environment.

Table 1 provides a list of suspicious network threat activity inside an environment that may benefit from PCAP and/or network metadata.

Table 1. Intelligence Requirements Year over Year

Hunt	PCAP	Network Metadata
Port/Protocol Combinations	Traffic content (second step)	High-level flows (first step)
Encrypted Traffic Signatures (Certificate/JA3/JA3S)	N/A	Certificates/JA3/JA3S signatures
Non-Application Layer Protocols	Traffic content, if applicable (second step)	High-level flows (first step)
Directionality Anomalies	Traffic content, if applicable (second step)	High-level flows (first step)
Lateral Movement	Traffic content, if applicable (second step)	High-level flows (first step)
New Domains/IPs	Traffic content, if applicable (second step)	High-level flows (first step)

Patch Management

An oldie but a goodie, patch management is an extremely effective step that any security team should ensure it is staying current on. Patch management is a task vital to nearly every area of information security, and network devices are no exception—although in practice they often are. Weak patch management presents a severe threat to any organization.

Earlier, we discussed vulnerabilities that may exist in network infrastructure devices such as VPNs. When it comes to vulnerabilities in such devices or a server, that twofold threat should be addressed by patching the vulnerable device to remove the vulnerability, especially if the device is external facing.

Figure 3 shows part of a Shodan server search, detecting millions of external-facing VPN servers. This simple scan pinpoints what countries the servers are in, what ports are accessible, and, perhaps most concerning, what vulnerabilities have been identified in the servers as well.

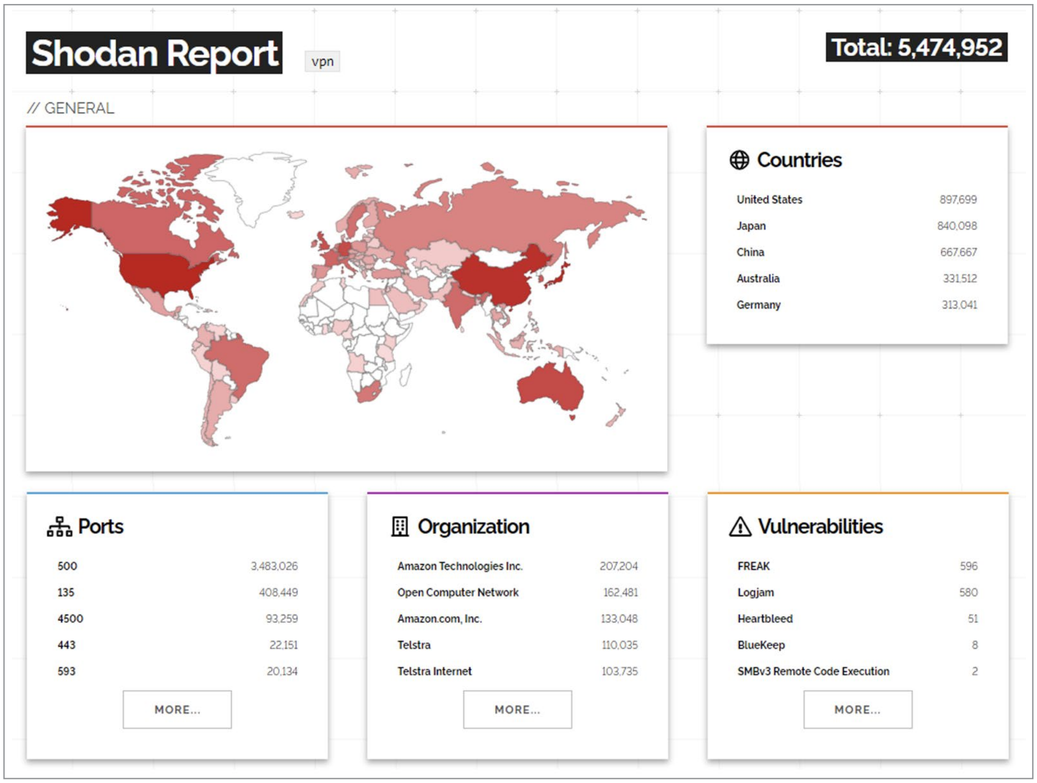

Shodan Report vpn

Total: 5,474,952

// GENERAL

Countries

United States	897,699
Japan	840,098
China	667,667
Australia	331,512
Germany	313,041

Ports

500	3,483,026
135	408,449
4500	93,259
443	22,151
593	20,134

Organization

Amazon Technologies Inc.	207,204
Open Computer Network	162,481
Amazon.com, Inc.	133,048
Telstra	110,035
Telstra Internet	103,735

Vulnerabilities

FREAK	596
Logjam	580
Heartbleed	51
BlueKeep	8
SMBv3 Remote Code Execution	2

Figure 3. Snippet of a Shodan Scan for the vpn Tag

Closing Thoughts

Networks are ubiquitous in enterprise architectures. They are the links between all the various asset classes that organizations seek to secure. For that reason, adversaries cannot launch an attack on an organization without crossing the lines of network architecture. Organizations must take steps to ensure that the network can detect adversaries inside the environment, and they must implement protections to stop adversaries from abusing its networks.

In this whitepaper, we looked at some of the common threats to enterprise networks and defenses that can be implemented to mitigate those threats. In many cases, threats depend on adversaries abusing the network as part of a larger attack. For example, we examined ransomware as an incident that is heavily dependent upon the network but also affects other assets inside the organization. We also looked at attacks on the network itself, such as network infrastructure vulnerability or disabling attacks such as DDoS that are designed to bring the network down. Given its criticality, the network remains a medium that must be protected.

Luckily, the protections that are available for enterprise networks today far exceed those of many years ago. Gone are the days when implementing a firewall and perhaps enabling network segmentation were the most that organizations would (or could) do. These days, technology exists that can offer smart detections, deep packet inspection, anomaly detection, and the capability to stop attacks with automated playbooks. It is imperative that organizations utilize as much technology as possible to prevent network-based attacks to stop adversaries in their tracks.

Sponsor

SANS would like to thank this paper's sponsor:

corelight