



Sponsored by Symantec

Layered Security: Why It Works

December 2013

A SANS Analyst Whitepaper

Written by Jerry Shenk

No Silver Bullet *PAGE 2*

Defense In Depth *PAGE 5*

Key Layers *PAGE 6*

The Human Layer *PAGE 11*

Introduction

Since the dawn of civilization people have thought they were more secure than they really were. Babylon once felt so secure behind its massive walls and sturdy gates that its king feasted while the Persian army laid siege to the city. The Persians entered early the next morning by damming the river and wading along the riverbed, right into the city. No one had considered that the river that passed beneath the walls—the city’s water supply—could also be its fatal weakness; the king paid for his mistake with his life.¹

Our modern digital kingdoms are no different. We may think one layer of security will protect us—for example, antivirus. Unfortunately for that approach, history has proven that, although single-focus solutions are useful in stopping specific attacks, the capabilities of advanced malware are so broad that such protections inevitably fail.

Organizations operating in the digital world today need layers of security so that an email message that gets through the firewall will get stopped by the mail server’s antivirus; and if it makes it through that, then it should be stopped by the workstation’s antivirus. If the hostile program actually secures a toehold on the workstation, it should be detected when it runs on the workstation because it’s doing things that are suspicious or unexpected. Look for connections to sites on the Internet with known relationships to hostile activity, and block such sites by egress filtering on the firewall.

Attackers are leaving no stone unturned, prying into web applications, operating systems and even deeper in the hardware. They’re taking advantage of conventional endpoints and mobile devices, slipping past and through network security, and even taking advantage of the human element operating the devices.

In this paper, we discuss how the layered model is more relevant than ever and how a layered approach to security provides better protection of your organization’s IT assets.

¹ www.livius.org/ct-cz/cyrus_1/babylon07.html

No Silver Bullet

Security companies often have marketing pieces that suggest that the installation of their product will solve all your security problems. There is no such thing as a silver bullet, and it takes many technologies and processes to provide comprehensive risk and security management. Trusting the latest thing will not end up well if your organization finds itself under a targeted attack. Even if you aren't specifically targeted, assuming you are safe is a dangerous way to operate. Instead, organizations should continually be checking their systems for vulnerabilities, learning about new threats, thinking like attackers and adjusting their defenses as needed.

Protecting the Network

Firewalls are designed to let some traffic in based on type, port and destination. Take email for an example: Blindly allowing email traffic through a firewall could permit an email message designed to compromise your systems (say, through a sophisticated—or maybe not so sophisticated—phishing attack) to reach your mail server. Firewalls have evolved by incorporating deeper analytics, such as IPS and traffic inspection services, enabling them to look deeper into packets and better detect malicious traffic.

Such improvements make the attacker's job more difficult, requiring an adjustment of tactics to make the outbound connection look like HTTP traffic, which most firewalls blindly allow to pass without inspection. Some malware uses DNS traffic or another accepted protocol. It is nearly impossible to make a firewall that is, by itself, tight enough to stop all possible ways of connecting a back door and still have a network that fulfills the user's needs. Nevertheless, such devices are one of the easiest starting points when implementing a layered defense.

Protecting the Computer

Antivirus software is still also a valuable part of a security plan, but it should not be the only technology used to protect networks and devices. Antivirus started out with a single mission: to protect computers from malware by relying on comparisons to file signatures. Testing and attack tools such as BackOrifice emerged and, over time, were updated to employ the concept of polymorphism; these tools use different—but functionally equivalent—programming logic and encoding mechanisms to change their file signatures and avoid detection.

In response, antimalware started looking for anomalies and indicators of bad behavior, thereby improving endpoint protections; but attackers can still circumvent detection. For example, cloud-based virus testing has become a big thing among attackers: The Google-owned VirusTotal file-checking service—designed to validate a file by running it past dozens of scanners—is one way that attackers can test their wares against several antivirus tools at one time. Attackers are getting better at selecting their targets and reconnoitering their victims to try to find out what type of antivirus software they use in order to try and disable it.

Protecting Other Layers

As attack methods have evolved, more choices of information security and risk management tools have been developed to fill the void. For example, endpoints—particularly mobile ones—use protections beyond antivirus to centrally manage things such as what endpoints can access or what applications they can use.

Today we have a whole ecosystem of layered security—endpoint protection, application protection, network protection and end-user controls—which must carry us through a variety of use cases. One set of guidelines coming into increasingly common use is the Critical Security Controls, listed in Figure 1.



Figure 1. The Critical Security Controls

No Silver Bullet (CONTINUED)

Tackling all of these from scratch simultaneously is impossible, but most organizations already have some of the controls in place, although they may not all be fully automated or integrated.

In this paper, we focus most of our discussion on layers involved in detecting and blocking attacks. The key to an effective security posture is to have multiple tools using different techniques to block attacks, alert if there is a breach, provide reporting and analysis for follow-up, and finally, remediate the vulnerabilities that allowed the hostile activity to take place.

Defense In Depth

To assess your defense-in-depth strategy, it's important to start with identifying the type of data you have that might be of interest to attackers, then determining where that data resides and evaluating what your level of vulnerability is. This, in turn, will help you apply the right resources to the right systems.

When you think of layered security, it's important to consider the exploit paths taken by attackers and malware. For example, most attacks start with some sort of targeted (or even not targeted) phishing attack against a user. When the user falls for the phishing or clicks a malicious "drive-by" site, the endpoint gets infected. Attackers use these initial infections as launch points to get deeper into the organization, where they can access data with real value: Admin passwords (the "keys to the kingdom"), financial account information and customer databases with personal data are a few of the favorite targets. Attackers may also want to engage compromised computers in other illicit activity, including DDoS and phishing relays.

In other words: The target is everybody; the attackers throw a wide net hoping for any takers. Organizations need to widen their security nets to protect and defend against opportunistic attacks and infections.

The term *layered security* describes a defensive strategy featuring multiple defensive layers that are designed to slow down an attacker. The military calls this *deep defense* or *defense in depth*. The goal is slowing an attack and causing enemy casualties. In the digital world, this means causing delays for the attackers and detecting them before they can do serious damage. In some cases, a properly placed layer may act as a strong enough deterrent to cause the attacker to look for an easier target. However, some attacks have built-in contingencies to look for multiple vulnerabilities in a variety of systems.

If an organization is not aware of an attack, the layers of security may only slow the attacker, but a determined attacker may eventually circumvent the layers of security. Consistent monitoring and reporting is a key requirement for good, layered security. Adding a management layer on top of all of the security and vulnerability layers is another requirement and should include integrated workflow for remediation of detected vulnerabilities.

Key Layers

For this paper, we've treated digital security as having five layers related to protection, detection and remediation of events. These layers don't always correspond to a physical system; some of these are more likely to be services. Some security systems will span multiple layers; and there could be good reason to further divide these layers, but these five are a sound starting point. Organizations that get a good handle on these five layers will be in a much better situation than most to repel and discover attacks. Figure 2 shows a representation of this model that treats the various layers as slices through an onion.




Figure 2. Layered Approach to Security

Network Controls

When we discuss network security, firewalls are one of the first things that come to mind; but IDS, IPS and other systems are also valued components of a defense in depth. As noted earlier, firewalls are often configured to let some traffic pass with ease (such as web traffic on ports 80 and 443), so attackers will attempt to hide hostile traffic by making it look like legitimate traffic. They can even encrypt this traffic to make deeper analysis with an IDS or IPS more difficult.

Data loss prevention (DLP) systems and in-line decryption appliances can examine network traffic in a fashion similar to IDS systems. They use deep packet inspection to determine if any tagged documents are transferred out of the network or if any other data policies are being violated.

Antivirus

The antivirus layer is one that organizations often apply as their single control against vulnerabilities and attacks; they're assuming that it makes them safe. Antivirus is a good part of any security plan, but it is not a guaranteed way to avoid security problems. Attackers are quite aware that antivirus tools are widely used and rely heavily on signature-based detection, so they try to make sure their code doesn't match the signatures.

Antivirus scanning is typically implemented at multiple points throughout a network, generally on email and file servers, workstations and mobile devices. Many organizations scan mail at both the email server and the client: If an email server detects a virus-infected email, it can delete or quarantine the message so that it will never arrive at the final destination. If the infection is detected after delivery to the client, it can often be deleted before it ever runs.

It is also possible to scan files as they move through the network. This might take the form of an IDS/IPS, or other hardware placed at network chokepoints. In some cases, file transmissions can be interrupted so that the workstation never gets the file; in others, network administrators can be alerted to the delivery of a hostile program.

For about a decade, antivirus has also deployed the use of *heuristics*—looking at what applications can do and watching for suspicious activity in addition to signature-based detection. Endpoint security takes this farther and can now monitor for suspicious activity, which can take many forms. One example that would certainly be suspicious would be a case of somebody who receives an email message with a PDF attachment that, when opened, attempts to download an executable from a website.

Reputation

Another way some organizations try to keep their computers free of malware is to check the “reputation” of a file, which can be done in a number of ways. Every computer file has a unique checksum—a relatively short mathematical value for the file. Another reputational characteristic of a file is how widely it has been used. Such assessments create a context for the file, indicating whether it is known to be good or bad or whether it is an unknown risk that should be monitored closely.

All files have checksums, including hostile files. So if a file is unknown to you, but the checksum matches a piece of previously discovered malware, then you can be certain that it's the same malicious file—even if the file name is different. This checksum can also be used to find out how widely used a program is. For example, just because your copy of Excel doesn't match the checksum listed in the sidebar doesn't mean that it's a Trojan

Creating Checksums

The math to create a checksum is complicated; but in simple terms, every version of a file has a value that allows it to be easily compared with the checksum value of another file. For example, the checksum for the 64-bit version of Excel (2013 version 14.0.7106.5003) has an MD5 checksum of f8d69729bfe6e26b672b78d368954b31. If an attacker were to get this copy of Excel and modify it so that it would carry a backdoor, that would create a detectable change to the checksum. It is nearly impossible to modify a file in any way and have it retain the original checksum.

horse; it could be a newer or older version. No matter what version of Excel you are running, you should not be the only one running it—if you are, that is indeed suspicious and would give the file an unknown reputation.

Another point of reputation checking is the IP address or domain name from which the file came. For example, the risks of downloading a file from a .ru site (Russia) or connecting to a .cn site (China) is of concern to a lot of organizations, and the related IP addresses and domains might have a bad reputation already attached to them.

So far, it's been tough for attackers to get past file reputation checking that relies on whitelisting as well as blacklisting. An attacker can avoid a list of naughty checksums by simply rewriting the code a little; but then the checksum wouldn't show up in the lists of known good files. To take it a little further, if you find that you have the only copy of a particular executable ... well, that's probably not a good thing.

Behavioral Analysis

In the words of Dr. Eric Cole, a SANS Faculty Fellow, "Prevention is best, but detection is a must." If the prevention doesn't prevent, the next best thing is to notice that an attack took place. Breach investigation reports show that a majority of breaches are not discovered for at least a month, many much longer than that: By then, of course, the trail could well be cold.

There are a number of ways to automate the detection of abnormal behavior, through DNS lookups or on the computer itself, but the first task is to learn and understand what "normal" is. Some behavior can easily be established as normal by establishing a baseline for network traffic. Such a baseline takes time to develop and must be updated as applications and user profiles change.

Baselining Normal Behavior

The source of the traffic baseline can be firewalls, routers, switches, flow collectors and network taps. Because firewalls and routers pass traffic through them, they are an ideal place to start. Most enterprise firewalls are capable of generating logs about every connection they allow or block. Routers can also generate traffic logs, but they often don't provide the same level of detail.

This is where an understanding of network flows—detailed descriptions of the connections being made through the firewall or router, as well as within the organization—are helpful. What's typically most interesting, from a security perspective, is the flow between an internal host and one on the Internet. Network flows can be gathered from most firewalls and routers as well as from quite a few switches and some other devices. There are different types of flow collectors, but they all collect the source and destination IP addresses and the source and destination ports, along with the time and duration of the connection and how many bytes were transmitted. This information can then be processed to show visual patterns of overall traffic as well as detailed subsets of this information; it can also be saved as text for further processing and review.

Analyzing Unusual Activity

With the baseline in hand, we can dive into deep analysis, which should include monitoring the size and frequency of connections and could include the frequency of blocked connections (for those using egress filtering—and all organizations should do so). Some sources for analysis and what they mean include:

- **High-Bandwidth Traffic.** Some illicit transfers can occur over a single, very large transfer at a high rate of speed; this can often be detected by relatively basic traffic pattern analysis, noticing a spike in traffic and then digging deeper to find the host or hosts involved with the transfer.
- **Stealthy Traffic.** It is more difficult to detect long, slow connections, but firewall logs often report the size and duration of a transfer. Flow data can also be searched for long connections; some long connections may be legitimate, such as an off-site backup or data replication for disaster recovery purposes, but any unknown long connections should be reviewed. Critical Control 14 recommends that connections longer than 10 hours be examined to see whether they serve a legitimate business purpose. A security-conscious organization should keep a list of regular large data transfers so they can quickly identify abnormal ones.
- **Web Traffic.** Firewalls typically don't block web traffic because it is approved traffic; however, attackers often use it as a communications or delivery channel. If we initially look at legitimate web traffic, we will see a certain level of irregularity in the sites being connected to and the associated file transfers. There is also some regularity in the timing of connections such as normal work hours or lunch breaks. Illicit traffic may have a very distinct, regular pattern as the malware tries to "phone home." If workstations are left on during nonworking hours, these regular outbound connections can be especially evident. Of course, some malware tries to avoid this detection by randomizing the outbound connection—that's why we need multiple layers of detection.
- **DNS Traffic.** Malicious DNS and countermeasures are briefly discussed in Critical Controls 11 and 19, because hostile programs often use DNS to identify their control servers. The software could connect directly to the IP address of the control server, but because IP addresses can be easily and quickly blocked, attackers have a better chance of maintaining a toehold on the compromised system if they can change the control server's IP address quickly and frequently. In such a scenario, the compromised computers need DNS to find the server. DNS addresses can also be blocked quickly, but savvy attackers will rapidly change the DNS records for their control servers, sometimes cycling through names in less than an hour. This forces the malware to repeatedly perform DNS lookups, many of which will fail. This high number of failures gives an opportunity for discovery. Although some DNS lookup failures are normal, if the number of failed DNS lookups becomes excessive, you should examine the host(s) generating those lookups.

We can also perform behavior analysis at the workstation or server. Some of today's antivirus software includes aspects of this technology—the basic idea is that certain things are normal and other activities are suspicious. A 2011 vulnerability in which executable PHP code was included in a JPEG file provides a great illustration of questionable behavior.² Of course, picture files shouldn't be connecting to the Internet, formatting drives or running scripts when they are viewed; so if the behavior analysis software detects an image file doing this or executing other commands, the software can be shut down and deleted, and an alert can be sent to network administrators. As malware increases in sophistication, the ability to enhance defenses by adapting real-time behavior monitoring is increasingly important.

Remediation

As we noted above, "Prevention is ideal, but detection is a must." However, remediation is where the real work begins. Once something malicious is detected on your network, it is imperative that you detect and fix it before it gets more deeply embedded or starts sending your data out or putting personally identifiable information (PII) at risk. Network-based IDS/IPS and firewalls may be able to limit the scope of infection or identify systems requiring remediation. Detection can be the result of behavior analysis, log monitoring and reports from end users or third parties. Critical Control 14 recommends automated maintenance, monitoring and analysis of audit logs. Most attacks can be detected with effective log monitoring; processes crashing, ports being blocked, files being transferred and user accounts being manipulated are some of the key indicators of a system compromise.

Any successful malware deployment should generate an alert, and the offending software must be removed from the network and analyzed. The exploited vulnerabilities must be fixed as soon as possible. Critical Control 5 recommends malware detection and reporting within one hour, with an ongoing goal of reducing that response window. At this level, speedy detection and remediation will reduce the likelihood that data is compromised and the likelihood that the security breach will expand to other computers.

² <http://php.webtutor.pl/en/2011/05/13/php-code-injection-a-simple-virus-written-in-php-and-carried-in-a-jpeg-image>

The Human Layer

This really isn't a layer like the others we've discussed, but it may be the most important and vulnerable. When it comes down to it, many malware infections and targeted attacks have been successful because someone clicked on a link, opened a suspicious email or visited a malicious website.

Employees can and should be regularly trained to use common sense and avoid falling for too-good-to-be-true offers or even unsolicited offers that look like legitimate business requests, because employees are easily targeted, thanks to social networking tools such as LinkedIn.

Educate all employees and remind the bosses as well as your end users of the dangers of malicious advertisements. Policies must cover all devices that employees are using for work purposes, including their own, and should include a combination of education and technology on the endpoint and, often, on the network in the form of network access controls.

One ideal way to encourage common sense is to catch somebody doing something right. If Joe in Accounting asks you about an email that you supposedly sent him asking for his password, congratulate him right away for being suspicious, then congratulate him publicly in your company newsletter. Another way to reinforce secure practices is to automate endpoint security, enabling those systems to actually tell the user when they're violating policy—for example, visiting a potentially malicious website, clicking on a phishing link, or loading an unapproved application—so the systems that users interact with every day become teachers instead of mere network police.

Indeed, this is where all of your security intelligence comes together. If your systems recognize new attacks starting on user systems, if you see something in logs or hear about some attack that's following the latest trending news, learn from that. Educate employees and start monitoring that type of activity more closely.

Conclusion

Today's attackers strike across multiple layers. That means that our security must also be layered. The Critical Controls lists 20 areas required to protect networks, endpoints, applications and personally identifiable data, while regulatory requirements also provide frameworks for layered security and compliance policies.

Most IT pros know that there is no silver bullet, and they don't put all their trust in a single security solution. Attackers launch their attacks at whatever vulnerability they find most promising, and some particularly motivated attackers will research the vulnerabilities of a network to build attacks that can work through network defenses. You need to build different layers of defense to cover all known vulnerabilities. To thwart determined attackers, the defensive layers must be capable of slowing the attack down and enabling network monitors to detect the attack before damage is done.

Layered security will continue to grow as new threats tax current defenses. Putting all these layers together—and combining them in a coordinated efforts—is going to be in high demand and presents an opportunity for vendors to standardize and consolidate layers as they mature and become part of the infrastructure. The ultimate responsibility lies with organizations to determine their key assets, identify where vulnerabilities lie and design security that will protect them from attack and detect any breach.

About the Author

Jerry Shenk currently serves as a senior analyst for the SANS Institute and is senior security analyst for Windstream Communications, working out of the company's Ephrata, Pennsylvania location. Since 1984, he has consulted with companies and financial and educational institutions on issues of network design, security, forensic analysis and penetration testing. His experience spans networks of all sizes, from small home-office systems to global networks. Along with some vendor-specific certifications, Jerry holds six Global Information Assurance Certifications (GIACs), all completed with honors: GIAC-Certified Intrusion Analyst (GCIA), GIAC-Certified Incident Handler (GCIH), GIAC-Certified Firewall Analyst (GCFW), GIAC Systems and Network Auditor (GSNA), GIAC Penetration Tester (GPEN) and GIAC-Certified Forensic Analyst (GCFA). Five of his certifications are Gold certifications. He also holds the CISSP certification.

SANS would like to thank its sponsor:

