TechTarget
Security Media

SearchSecurity.com
SearchFinancialSecurity.com
INFORMATION SECURITY®
SearchSecurity.co.UK
SearchMidmarketSecurity.com
INFORMATION SECURITY DECISIONS

SearchMidmarketSecurity.com

Pocket E-Guide

Key technologies in a network perimeter intrusion defense strategy

A lot of terms are thrown around when talking about security and intrusion defense, and it can be confusing when the meanings aren't clear.

This Pocket E-Guide cuts through the noise and gives you independent, expert advice from security guru Joel Snyder as he defines these terms in the context of a perimeter intrusion defense strategy.

Sponsored By: ArcSight™

TechTarget®
The Technology Media
ROI Experts

Intrusion defense strategy: exploring key technologies in a network perimeter

This tip is part of Network perimeter security, a lesson in SearchMidmarketSecurity.com's Intrusion Defense Security School. Visit the lesson page or our Security School Course Catalog for additional learning resources.

A lot of terms are thrown around when talking about security and intrusion defense, and it can be confusing when the meanings aren't clear. Before we get knee-deep in our discussion of intrusion defense, let's define a few terms. I'm not going to attempt to provide the definitive answers to what each of these terms means in all contexts, but I will define them in the context of a perimeter intrusion defense strategy.

Antivirus

Antivirus is probably the most common term in perimeter security, but even such a simple term is open to multiple definitions. Viruses, Trojan horses, worms and malware are terms used collectively to describe malicious (or perhaps benign but unwanted) software in one incarnation or another. The nature of the malware is what differentiates a virus from a worm, for example. But when we say "antivirus," we're talking about detecting the presence of any of these types of unwanted software, not just viruses.

A virus is malicious software that infects other applications. When the application is launched by the end user, the virus is activated, and it both infects additional applications and does its evil deed, such as deleting random files on your disk or popping up Viagra advertisements on your Web browser. Unlike viruses, worms are both self-contained and self-propagating; no one needs to launch an application once a worm is loose. A Trojan horse is malicious software that is disguised as a legitimate application; it doesn't propagate itself.

Of course, hackers aren't worried about these different classifications -- they are just as happy to use a Trojan horse to carry a worm payload that also infects applications like a virus. The term "blended threat" is often used to describe these hybrids.

The reason it's important to understand these differences is that antivirus, particularly at the perimeter, may have multiple functions to catch different types of malware at different points in its lifecycle. In addition, you may see the term "antivirus" used in different contexts.

At its simplest, antivirus looks for malware in transit. The most popular malware self-propagation technique is via e-mail, and thus the scanning of e-mail for malicious attachments is a critical part of an antivirus strategy. However, malware can also be left on Web sites, and perimeter-based antivirus scanners look in Web data streams for malware as well.

The problem with both of these techniques is that they're not guaranteed to be able to reconstruct the payload of the virus from the packets that they see. Encrypted e-mail and Web sessions are obviously an issue, but Web traffic on non-standard ports or viruses in certain Web applications may also get past virus scanners. For this reason, any antivirus strategy at the perimeter can only complement antivirus on the desktop.

A second perimeter antivirus technique involves searching for virus misbehavior signatures. For example, the famous Code Red worm propagated to IIS Web servers by sending a particular URL. Perimeter-based antivirus can discover the presence of a Code Red-infected machine by its behavior. This technique is only useful for identifying malware after an infection has already occurred. However, knowing that someone is infected is only slightly less useful than keeping them from being infected in the first place. There are also intrusion defense tools that specifically look for malware propagation signs and use that information to help isolate infected systems. These are known generically as "network behavior anomaly detection" (NBAD) systems.

Antispyware

Knowing what you've just read about viruses, it's fairly easy to see that spyware (sometimes called adware) is another class of malware and can be detected by the same techniques. Spyware is most commonly propagated by "drive-by downloads," in which a user visits a Web site and, as a side effect, downloads additional software to their PC. Sometimes the software is downloaded without the user's knowledge, or the Web site might attempt to deliberately confuse the user into allowing the download to occur by bypassing a browser security protection. Users even download and install spyware deliberately, usually because they've been led to believe by a deceptive promotion that the software will somehow improve their Internet experience.

For reasons best left to conspiracy theorists, antispyware is often handled separately from antivirus. However, the detection techniques of looking for file signatures (especially in Web pages, more so than e-mail messages) and looking for behavior anomalies are equally relevant for antispyware and antivirus. And, like antivirus, a desktop detection and prevention strategy must complement a perimeter-based defense. Over time, we can expect that antispyware and antivirus software will merge into a single tool, although the turmoil and churn in the market right now has caused many enterprises to have to buy both tools to fight this growing plight.

Antispam

Compared to malware detection, spam detection is both much more difficult and much simpler. Because anyone with a functioning brain can generate a spam message, the rate of new spam creation is tremendous. At the same time, spam can only travel via e-mail, so the task of redirecting traffic to go through a spam filter is simpler than catching all potential virus propagation or activity traffic.

Another difference between antispam and antivirus techniques is that both the false positive and false negative rates for antispam software are much higher than that of antivirus software. In other words, a lot more spam gets through (than viruses) and a lot more messages are misclassified as spam (than viruses). As a result, antispam features such as end user-based quarantine and per-user sensitivity settings and whitelists are often critical to end-user satisfaction.

Antiphishing

As a variety of unwanted or malicious mail, phishing mail can be detected in the same manner as spam mail. At a technical level, any antispam engine can also be an antiphishing engine. At a marketing level, the idea of selling a product as combating both menaces is irresistible, and thus a number of products are specifically called out as handling both types of traffic.

As antispyware is related to antivirus, so antiphishing is related to antispam. All competent antispam engines are also good at handling phishing attacks.

A variant on antiphishing based on network behavior is gaining currency with some perimeter defense vendors. Because phishing e-mails generally ask the reader to click on a link and provide information to a Web page, the idea is that you can help "infected users"-- those who are deceived by the phishing e-mail into visiting the decoy Web site -- by catching those clicks with an NBAD system. It turns out that this technique would have worked for standard spam (since users are being directed to Web sites to order drugs, watch porn and buy stocks) as well, but since the damage to the user was not considered significant, the technique wasn't commonly used.

Now, a number of perimeter intrusion defense products are using the URLs in e-mail as well as URLs clicked upon by end users to detect incoming spam/phishing mail and outgoing responses. In essence, this is the same technique used to detect propagating or attacking malware, but turned to the specific task of protecting users from responding to phishing e-mails.

Intrusion detection and prevention systems

Although these product families sound like they should be related, they have few similarities. An intrusion detection system (IDS) watches data flow at one or more points in the network, providing alerts and forensics on suspect or malicious traffic. Critical aspects of an IDS are an alerting system, and a data store for forensics and logging purposes.

An intrusion prevention system (IPS) is an in-line device that blocks malicious traffic. Some earlier IPSes were not in line. They would detect malicious traffic and then mitigate the effects of the traffic; for example, by flooding the sender and receiver with TCP Resets, or by changing access list rules in a firewall or router. However, contemporary IPSes all look about the same: in-line traffic evaluators looking for some reason to drop a packet or reset a connection. Because an IPS is looking for malicious traffic, it can both be less discriminating and more careful than an IDS.

For example, an IDS might detect attempts by a worm to propagate within a network, and then alert on those attempts. However, the IDS should classify the attempts by the system being attacked, its level of importance to the organization and its vulnerability to the particular attack attempt. An IPS does not have the same level of complexity. When an IPS sees a definite attack attempt, it can simply block the attack. It doesn't matter whether the system being attacked is vulnerable or not, important or not, or even if the system exists. The IPS can safely block definite attack attempts. However, the IPS must not block legitimate traffic. Since an IDS alerts while an IPS blocks, most IPSes have only a few hundred enabled signatures (to avoid false positives) while IDSes often have thousands of attack signatures.

Not every IPS uses signatures to identify attacks, and even those IPSes that use signatures may use a combination of other techniques to help identify (and block) malicious traffic. The area of NBAD systems closely overlaps the function of IPSes, and these products are often seen as solving similar problems, although using very different techniques.

DoS/DDoS defense

IPSes may also be "rate-based," meaning that they are looking for unusual volumes of traffic. These systems are also marketed as DoS and DDoS (denial-of-service and distributed denial-of-service) defense tools. Rate-based IPSes may be combined with signature-based IPSes. However, because they defend against very different kinds of attacks, they are usually placed at different points in the network and protect different types of systems. For example, a rate-based IPS is most often used in front of large Web server farms or large e-mail servers, while a signature-based IPS is placed directly inside the corporate firewall (or as part of the corporate firewall) to protect end users or more generic types of servers.

IPSes and IDSes that implement behavior anomaly detection or look for specific virus-oriented or phishing-oriented behaviors may be sold as antivirus or antiphishing tools. Although this is a very useful aspect of these kinds of products, it is important that you not consider an IPS or IDS a "first line of defense" against malware or spam.

Content filtering

Content filtering tools can use a number of different techniques, all aimed at the same goal: to restrict unwanted content from corporate desktops. Content filtering is almost always used in the context of Web browsing, although the idea can be extended in other ways. Most content filtering uses a categorize-and-block approach. Web URLs are inspected by a content filter before leaving the corporate network and compared against large databases. A URL might come back as unknown or might fit into a category, such as "sports" or "gambling." Based on this categorization, the network manager might elect to block traffic to those types of Web sites, either entirely or based on some other more restrictive criteria, such as time of day or user authorization information.

Some content filters actually attempt to analyze the content returned with an eye towards catching traffic that the content filter doesn't properly categorize. This has gone to absurd lengths. For example, products have been brought to market that attempt to analyze the content of images, specifically for the purposes of blocking pornographic images.

Content filtering is not a particularly reliable technique, and it generally won't keep a determined user from downloading inappropriate information. However, it is a commonly used tool in environments where some sort of filtering is required (such as grade schools) or in environments where a stated security or usage policy should be backed up by some technology enforcement (such as in a customer-facing retail setting).

Application control and bandwidth management

What content filtering is to Web browsing, application control and bandwidth management are to all other types of applications. Both of these are techniques used to block or control certain types of network usage. Application control is generally a part of advanced firewalls, while bandwidth management may be integrated into firewalls and other infrastructure devices such as routers, or may be handled through independent devices.

Like content filtering, application control is often used in environments where a technical enforcement must accompany a stated usage policy. For example, if an enterprise wanted to ban Skype voice-over-IP usage, application control could be used to enforce this ban.

Regulatory controls

As a broad category, the concept of regulatory controls is large enough to deserve its own article. However, most regulatory controls at the perimeter fall into one of three subcategories: leak protection, auditing and logging, and compliance.

Leak protection is the most difficult to do properly. Borrowing from the world of IDSes, leak protection tools attempt to monitor and manage the flow of sensitive information out of an organization by watching at the perimeter. Depending on the regulatory regime, this may range from protected private information (such as personal health data) to corporate sensitive financial data.

Auditing and logging tools are more passive and designed to help organizations comply with requirements to audit access (for example to corporate financial information) or maintain long-term records (for example of all instant messaging traffic outside of the organization).

Compliance tools are more active in ensuring that connections outside of the organization follow a policy that complies with the applicable regulatory regime. The most common example would be encryption of sensitive information. For example, a compliance tool might watch e-mail communications between a hospital and an insurance company, and either block any unencrypted communications, or step in and add encryption as required to comply with policy.

About the author:
Joel Snyder is a senior partner with consulting firm Opus One in Tucson, Ariz. He has worked in IT for more than 25 years.

Resources from ArcSight, Inc.

World-Class Protection for the Mid-Size Organization

Case Study: Long Term Care Partners

ArcSight Express: Security Expert "In a Box"

About ArcSight, Inc.
ArcSight (NASDAQ: ARST) is a leading global provider of security and compliance management solutions that protect businesses and government agencies. ArcSight identifies, assesses, and mitigates both internal and external cyber threats and risks across the organization for activities associated with critical assets and processes. With the market-leading ArcSight SIEM platform, organizations can proactively safeguard their assets, comply with corporate and regulatory policy and control the risks associated with cyber-theft, cyber-fraud, cyber-warfare and cyber-espionage.