SearchNetworking.com E-Guide

Signature vs. anomaly-based behavior analysis

News of successful network attacks has become so commonplace that they are almost no longer news. Hackers have broken into commercial sites to steal credit card information and into defense industry sites in search of top-secret military plans. Recent denial-of-service (DoS) attacks have made sites unavailable to legitimate users. Firewall and intrusion prevention systems across various enterprise networks routinely log hundreds of hacker attempts a day. To prevent successful attacks, two key detection approaches have evolved: signature-based and anomaly-based network behavior analysis (NBA). This E-Guide details each approach along with the pros and cons.

TechTarget
The Technology Media ROI Experts

SearchNetworking.com E-Guide

Signature vs. anomaly-based behavior analysis

Table of Contents:

Preventing hacker attacks with network behavior analysis IPS

News of successful network attacks has become so commonplace that they are almost no longer news. Hackers have broken into commercial sites to steal credit card information and into defense industry sites in search of top-secret military plans. Recent denial-of-service (DoS) attacks have made sites unavailable to legitimate users. Firewall and intrusion prevention systems across various enterprise networks routinely log hundreds of hacker attempts a day.

To prevent successful attacks, two key detection approaches have evolved: signature-based and anomaly-based network behavior analysis (NBA).

Signature-based intrusion protection and detection

Signature-based systems are extremely effective against attack types that have been detected in the past. They can be installed quickly and become effective immediately. These systems examine each incoming packet and compare its contents against a list of known attack mechanisms. False positives, legitimate activity that appears to be an attack, are rare. Generated reports are easy to understand because each incident indicates the type of attack that was detected.

While signature-based systems are effective against known attack types, they cannot detect zero-day attacks. Hackers understand that any new attack type will be quickly detected and countermeasures will be adopted by intrusion prevention vendors. They therefore launch attacks on a large number of sites as soon as a new attack method is developed.

Because of this, signature-based systems must be continually updated. Vendors collect and monitor attack reports from across the world. They also collect data from products installed at customer sites. When one customer experiences an attack, vendor staffs analyze it, develop a defense and distribute the update to all other customers' sites. While vendors can often detect new attack methods and devise a defense quickly, the first sites to be attacked have already been compromised.

Anomaly-based intrusion detection systems

Anomaly-based detection systems detect network activity that does not fit the pattern of expected behavior. The system must be configured, according to the product, with information on normal patterns of activity. For example, applications may legitimately access a single database record at a time. If the intrusion protection system detects access to a large number of records, the cause is likely to be an attack. Similarly, if a user with permission to access a restricted set of records begins to attempt access to other types of information, the user's workstation is likely to have been infected.

Unlike signature-based systems, zero-day attacks can be detected because the attacks do not have a pattern that is recognizable as legitimate to the anomaly-based intrusion system. All that is necessary is that something outside the ordinary is occurring. The downside is that anomaly-based systems must be carefully configured to recognize expected patterns of activity. Configurations must be updated when new applications are added or existing applications modified. False positives can occur when legitimate activity departs from its normal pattern.

Configuring IPS to defend against complex attacks

Attacks in which elements of the attack are spread across multiple commands such as HTTP messages for Web-based attacks present a difficulty for both signature-based and anomaly-based systems. For signature-based systems, the signature may be spread across a series of commands with no one packet matching an attack profile. Anomaly-based systems may fail to detect an attack that simultaneously targets several hosts. The sequence sent to each host may appear legitimate but may cause applications on the hosts to interact in such a way as to cause a breach.

Compounding the difficulty, not all of the packets may enter the network at the same point or gateway. Although enterprise networks often maintain more than one gateway to the Internet with intrusion prevention systems at each gateway, guarding all the gateways is not sufficient.

Viruses can penetrate a network through places other than gateways. Employees take home laptops for use on their minimally protected home networks. When they reconnect the infected laptop on the internal network, viruses enter the network without passing through an Internet gateway. Wireless networks are another vulnerable point and cannot be overlooked when implementing an intrusion prevention system. An outsider breaking in via the wireless LAN (WLAN) has also bypassed the network gateways.

Intrusion protection systems must also be installed at key points throughout the network (like a switch connecting network gateways to servers where applications run or connect to database servers) to detect these attacks. Systems must exchange information with each other and evaluate reports from sources such as router and host logs to correlate the sequence of packets to detect the attack.

While signature-based systems can be quickly installed and immediately become operational, designing, configuring and installing an anomaly-based system is more complex. The next section in this series explores the steps involved in configuring and installing an anomaly-based system.

Welcome to a New Era of Data Center Protection

Introducing the TippingPoint N-Platform — a groundbreaking network security platform for data center and core network deployments. Now you can integrate and automate multiple network protection services to help reduce the total cost of security and ensure business continuity.

The time is now to discover how the TippingPoint N-Platform can help secure your data center from today's evolving threats.

LEARN MORE AT WWW.TIPPINGPOINT.COM/N-TT

TippingPoint

Anomaly-based intrusion protection configuration and installation

Anomaly-based intrusion protection devices operate by detecting network activity that is out of the ordinary and unexpected, such as zero-day hacker attacks. Installing and configuring a system that will recognize unexpected activity requires an understanding of the activity that is expected.

Monitoring the network for a few hours is not sufficient. Patterns of activity change over the course of a day and at different times of the month. Sample expected behavior from normal day-to-day operations and any end-of-month or end-of-year activities. An accurate understanding of behavior requires analysis of each application during these periods.

To install anomaly-based intrusion protection, analyze network applications

The first step is to **determine which applications run on the network.** While it may seem that this step is unnecessary since the inventory of applications should already be up to date, that is not always true. Applications may have been running for years without any need for upgrade or support and may have been forgotten. A detailed inventory may never have been created or may not have been kept up to date. In any case, now is the time to create the inventory or update it.

Creating a profile of expected activity for each application is the next step. An accurate, detailed profile is based on an understanding of what the program does. For example, an application that processes customer credit and checks each time a purchase is made will deliver a single customer record for each transaction, while a program that analyzes monthly patterns of purchases is expected to return much larger blocks of data. An end-of-month accounting application will typically not be accessed mid-month. The profile should include a listing of the other systems and applications with which the application communicates. If user workstations connect with the application, document exactly which users and which workstations legitimately access the application.

After you create or update the network application profile, review the expected transaction rates. The application accessing customer records when a purchase is made will normally be executed at the rate of customer transactions. An attack may generate a rapid sequence of transactions. Each transaction may access just a single customer record, but the rate of transactions may indicate that an attack is under way.

Keeping the profiles up to date is a time-consuming task. Any change requires an update. Any time an application is added, an existing application is modified, new equipment is added, the network is modified, or transaction rates change to a significant degree, the change must be reflected in the profiles.

Simplify application profile updating by segmenting the network

Configuring all anomaly-based intrusion protection devices with profiles of all the applications is difficult. Doing so requires updating each device every time any application changes. The task can be made easier by grouping applications on the network so a single intrusion prevention device monitors network activity for a single application or small set of applications.

If multiple instances of an application are run, they should all be grouped on a single physical network link, subnet or virtual LAN (VLAN). In many cases, applications that interact intensively with each other should be grouped together. The intrusion-prevention device on that subnet or VLAN will be configured to recognize only the patterns of behavior expected for the single application or group of applications. Updating the configuration for that device need be done only when a change is made on that small set of applications. Responsibility for maintaining the configuration can be assigned to staff members responsible for the set of applications instead of requiring a central group to be responsible for monitoring all application changes and maintaining all configurations.

Virtualization would appear to make segmenting the network more difficult. Virtual machines (VMs) move from physical server to physical server as load increases or decreases or systems are taken down for maintenance. Grouping applications on a VLAN eliminates the difficulty. VMs maintain the same VLAN membership as they move. All that is required is to configure all of the switches for all of the VLANs in use.

Once installed and configured, anomaly-based intrusion protection is quite effective. But no technology is perfect. A cleverly constructed attack could remain within expected network behavior. False positives are possible. A sudden increase in sales may trigger a level of activity that appears to be an attack. The final section this series examines integrating anomaly-based protection with other technologies.

Application-specific network intrusion detection systems emerge

No single intrusion detection technique is 100% effective. Anomaly-based products must be carefully configured to recognize normal behavior but may still generate false positives. Signature-based products do not require extensive configuration, but they cannot detect zero-day attacks.

Vendors have responded by developing products that integrate intrusion detection system techniques in a single product. Some vendors have gone further by using an ongoing analysis of normal and abnormal network behavior to create new signatures as a broader intrusion detection technique.

In an effort to further integrate intrusion detection systems, users must consider application-specific protection devices.

Application-specific network behavior analysis tools: No more Web attacks

Web applications are frequently the entry path for serious attacks. E-commerce applications, for example, access internal databases with valuable information (e.g., customer lists and credit card numbers), so they are highly targeted. As a result, Web application firewalls integrate both anomaly-based and signature-based technologies to detect frequently used attack techniques.

Typical Web application attack techniques include:

- **SQL injection:** The exploitation of security vulnerability in the database layer of an application.
- **Cross-site scripting:** Malicious attackers inject client-side script into Web applications.
- **OS command injection:** Attackers execute OS commands through vulnerable Web applications and can obtain data or upload malicious programs.

Application firewalls use both anomaly- and signature-based intrusion detection techniques

Web application firewalls combine anomaly-based techniques with application-specific methods. For example, requests from a specific client to an e-commerce site are normally spaced at least several seconds apart. A rapid stream of requests, several per second, is likely to indicate an attack. Similarly, most Web applications deliver a limited amount of data in response to each request. A very large response probably indicates an attack that somehow was not caught when the incoming request was scanned.

Web application firewalls also use the signature-based approach, scanning incoming requests against a periodically updated signature list.

Anti-spam as an intrusion detection system: SPF and DKIM

Anti-spam products are generally not considered intrusion detection systems, but trojans and phishing attempts do constitute serious threats to network security. Most anti-spam products are signature-based, but application-specific techniques are also used.

Newly developed spam identification techniques have been integrated over the past few years, including the deployment of two standards: Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM). Neither technique actually detects spam; instead, they focus on ensuring email delivery from the correct sender -- without email spoofing.

SPF is an email validation system that enables administrators to register with a directory listing which hosts can send email from specific domains. The goal is to eliminate email spoofing. With DKIM, the sender can ensure the delivery of an email by placing an electronic signature on each message that can be verified by recipients.

Both standards are limited by the fact that use is optional. Not all senders include SPF or DKIM information. Many legitimate service providers have adopted one or the other protocol, so a reasonable technique is to consider any mail not utilizing one of them to be spam. The downside is that some mail could be erroneously tagged and blocked.

Anti-spam: Sender reputation as behavior analysis

Filtering based on sender reputation is a more recent addition. This technique depends on the fact that most spam comes from a limited set of sources. Anti-spam devices at customer sites compile lists of senders based on incoming email and report back to the vendor. The vendor collects and combines the customer input and then computes a reputation score for each sender.

The reputation score is based on a series of factors, including the identity of the sender's service provider, country of origin, daily email volume and inclusion of URLs known to link to infected Web pages. The vendor updates its database of senders, combining the latest reports with previous reports, periodically updating customers. Email from sources with extremely poor reputations is blocked, while email from questionable sources can be rate-limited.

Detecting abnormal network behavior in a VoIP environment

VoIP is another application that benefits from integrated anomaly and signature-based solutions. Hackers attempt to gain access to make free calls or to deny service to legitimate callers. They may also break into the system to make a large number of nuisance calls -- the voice equivalent of spam. VoIP security devices integrate methods specific to VoIP protocols plus both anomaly- and signature-based methods.

Combining prevention devices with host and workstation software

The scope, complexity and potential cost of attacks requires use of all of the techniques: anomaly-based, signature-based, and application-specific methods. The fact that attacks can enter the network without passing through an Internet connection point means that intrusion prevention devices must be installed at crucial internal locations as well as at gateways. Finally, host and workstation-based software provides an additional level of protection. Solutions are available that combine all of these components, share information among components, and together create a comprehensive, integrated defense.

About the author: *David B. Jacobs of The Jacobs Group has more than 20 years of networking industry experience. He has managed leading-edge software development projects and consulted to Fortune 500 companies as well as software startups.*

Resources from TippingPoint

TippingPoint

Securing the Next-Generation Data Center - Introducing the TippingPoint N-Platform

Three Waves, One Defense: How IPS Protects from an Ever-Broadening Threat

Eight Questions To Ask About Your Intrusion-Security Solution

About TippingPoint:
TippingPoint, an enterprise network security leader, offers a modern network security platform and intrusion prevention system (IPS) that secures next-generation data centers for enterprises, government agencies, service providers and academic institutions. Purpose-built to protect today's next-generation data center, TippingPoint arms security executives with new capabilities to stop threats faster and protect the highest, multi-gigabit bandwidth networks from ever-evolving, global security threats. TippingPoint is a powerful front-line defense that can be rapidly deployed, providing immediate protection at critical entry and isolation points in the network. With TippingPoint, your business is protected 24x7x365, with network security that is continually updated by Digital Vaccine® Labs (DVLabs), TippingPoint's cutting-edge team of top security researchers.