Sponsored by Arbor Networks

How DDoS Detection and Mitigation Can Fight Advanced Targeted Attacks

SEPTEMBER 2013

A SANS Whitepaper

Written by John Pescatore

©2013 SANS™ Institute

Introduction

Distributed denial of service (DDoS) attacks continue to harm enterprises around the world. The obvious damage caused by DDoS attacks is bad enough, including headline-grabbing, multigigabit/second volumetric attacks that crash critical business and government systems. More insidious is the use of DDoS as a component of advanced targeted attacks.[1] Many of these attacks include DDoS components designed to stay beneath the network security radar, mimicking legitimate user traffic to escape detection. Protective security services start failing or, worse, are blocked altogether. During the confusion caused by the DDoS, the real infiltration takes place: Malware and attacks infect web applications or dig deeper into the network during the confusion.

DDoS is no longer just a tool for political and social attacks to make a statement and shut down a site. It is now an insidious and easy-to-execute component of attacks that today's protective measures are often unable to detect.

This paper explores how DDoS is used as part of advanced targeted attacks (ATAs) and describes how DDoS detection and prevention tools and techniques can be used against ATAs as well.

[1] http://globalpublicsquare.blogs.cnn.com/2013/05/21/the-real-cyber-threat

Attacks Grow in Complexity

Just as enterprise use of technology changes constantly, cyberthreats are continually evolving. In recent years, one of the most dangerous new threats has been the growth of advanced targeted attacks (ATAs) in which the attackers use specially crafted techniques to penetrate a specific company or agency for a specific goal. ATAs use multiple vectors to accomplish their goal, combining custom malware, spear phishing, evasion and denial of service techniques to evade or disable existing security protections and steal customer or business critical information.

For the past several years, the Verizon Data Breach Investigations Report[2] has identified ATAs as causing the most financially damaging attacks against victim organizations.

Anatomy of an Advanced Targeted Attack

ATAs have four major characteristics that differentiate them from earlier generations of attacks:

1. They have a goal of stealing information from a specific person or organization for cybercrime or cyber espionage purposes, rather than for vandalism or simple "hacktivism."
2. They rely on stealth techniques to go unnoticed and can keep operating within the victim's enterprise for as long as possible.
3. They are highly customized, based on an examination of each target's infrastructure and defenses, and they can often defeat defense strategies that rely only on signature-based detection.
4. Multiple techniques are combined to create multistage attack "campaigns" that are designed to evade detection and to continue and expand attack operations once they have penetrated a target.

ATAs are generally executed in three distinct stages that are effective because they mirror the way legitimate businesses often conduct e-commerce, as shown in Figure 1.

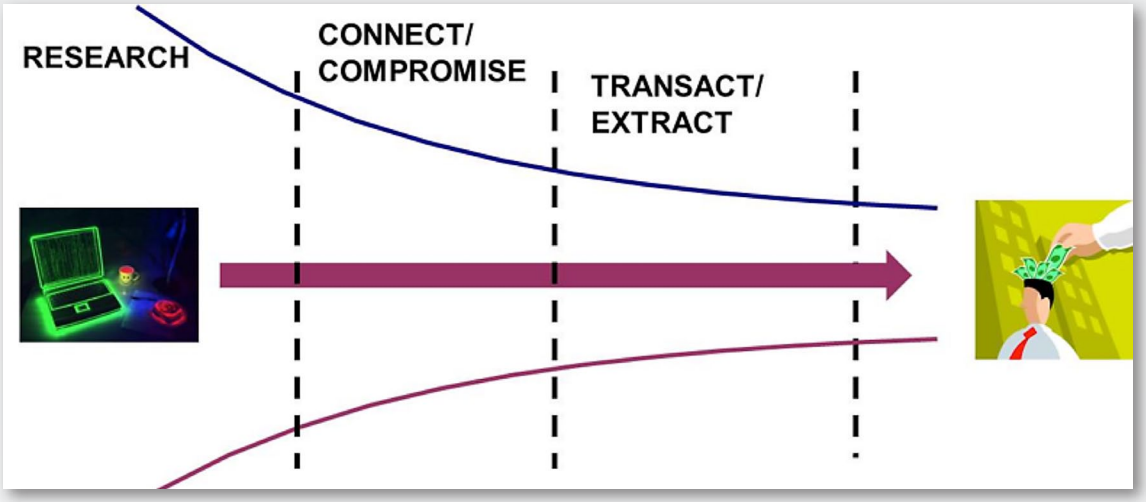

Figure 1. Stages of an Advanced Targeted Attack

[2] www.verizonenterprise.com/DBIR/2013

Attacks Grow in Complexity (CONTINUED)

Each stage of the attack depends upon the success of earlier stages, with research playing a major role. Here is how the three stages work together to escape detection, compromise their targets and extract valuable information.

1. **Research.** The attacker uses a number of techniques to determine the best way to penetrate and compromise the target. These include passive approaches, such as searching domain registration and network routing information to map out Internet connectivity, looking at corporate organization charts to obtain the names of managers and system administrators, and searching the web for other public information. This information is then used to search social media, job hunting sites and other locations to select the system and human targets. An attacker may follow such passive methods with more active approaches, such as IP address scanning and probing, web application scanning and account sign-up, or even initial penetration testing.

2. **Connect/Compromise.** Many ATAs begin with a spear phishing campaign that makes use of the research gathered on the target. The attacker uses the research information to send highly-customized and very believable emails to the human targets, who are often C-level executives or privileged users such as system administrators. The phishing emails attempt to steer the user toward external websites that are often legitimate sites that have been compromised by the attacker. When the user clicks on the link, the first stage of targeted malware (often called an advanced persistent threat or APT) is downloaded to their PC in what is usually called a "drive-by" attack.[3] A variant, which eliminates the initial phishing email and requires more patience on the part of the attacker, is called a "watering hole" attack.[4] In such an attack, a website the target has been observed to visit (such as the site of a local running club or an industry consortium) is compromised, and the attacker patiently waits for the target to go to the site and unknowingly download the targeted malware.

3. **Transact/Extract.** The first stage malware often waits until the host PC is connected to a private network, performs some initial surveillance and then communicates its findings to a command and control location to obtain instructions or to download the next stage of the malicious executable. ATAs often use network encryption to evade detection by intrusion detection systems and antivirus sensors. Sophisticated attacks can also use network flooding or denial of service techniques to overflow audit trails or cause floods of log events to hide the tracks of the actual attack.

   Because these targeted attacks often compromise a legitimate user's PC, they can take advantage of the user's privilege level to access databases, applications and other network nodes containing valuable and sensitive information without generating logs of unauthorized access attempts. The targeted information is then sent to the command and control center, often after being compressed and/or encrypted. This is why the stealthy nature of these attacks is so critical: The longer they are undetected and can send information of value out of the organization, the larger the hacker's profit.

---

[3] www.darkreading.com/attacks-breaches/cyberspies-target-victims-via-strategic/240000443

[4] http://krebsonsecurity.com/tag/watering-hole-attack

The Role of DDoS

Using DDoS helps ATAs hide among the "noise" of alerts and data organizations receive from their security monitoring systems. According to the SANS 2012 survey on Log and Event Management, respondents, for the first time, reported they were unable to detect active attacks in their networks. This was largely due to the advanced nature of ATAs and the resulting volume of security data they had to sort through.[5]

ATA research being conducted by attackers has also changed the nature[6] of DDOS attacks. Brute force, DDOS exploits that bombard targets with large amounts of data at high speed, continue to grow both in sheer bandwidth and frequency. Now, attackers have also begun to analyze the application logic running on the target's website. This allows the attacker to cause denial of service through "resource depletion" attacks, such as starting user authentication processes, site-wide searches or new account creation processes that consume high percentages of server CPU cycles or memory space without requiring large amounts of network traffic from the attacker. Such an attack achieves the attacker's goal (denial of service) but is much more difficult to detect and mitigate than a simple brute force attack.

Just as with ATAs, DDoS attacks have become multivector attacks. A good example is the DDoS techniques used in the "Operation Ababil" attacks against banks in early 2013. The Ababil DDoS attacks combined the following three techniques to increase their effectiveness:

- GET and POST application-layer attacks that use common commands to receive and send information via HTTP and HTTPS
- DNS query application-layer attacks, which try to disable, overwhelm or extract usable information from the domain name service that translates URLs into IP addresses
- Volumetric attacks, such as those that exploit the User Datagram Protocol (used on applications that require very high performance), the Internet Control Message Protocol (ICMP) used in network control and administration, or the SYN packets that request server connections to overwhelm network resources and deny access to legitimate users

[5] www.sans.org/reading_room/analysts_program/SortingThruNoise.pdf

[6] http://searchsecurity.techtarget.com/news/2240184214/DDoS-attack-trends-highlight-increasing-sophistication-larger-size

Gaps in Existing Network Defenses

As threats have advanced over the years, network defenses have also evolved. Most enterprises now deploy email, web and firewall/IPS platforms at their network boundaries to inspect inbound traffic and block clearly malicious connections. Intrusion detection systems are then used to inspect all traffic allowed by this first layer, looking for matches to signatures of known attacks, as shown in Figure 2.

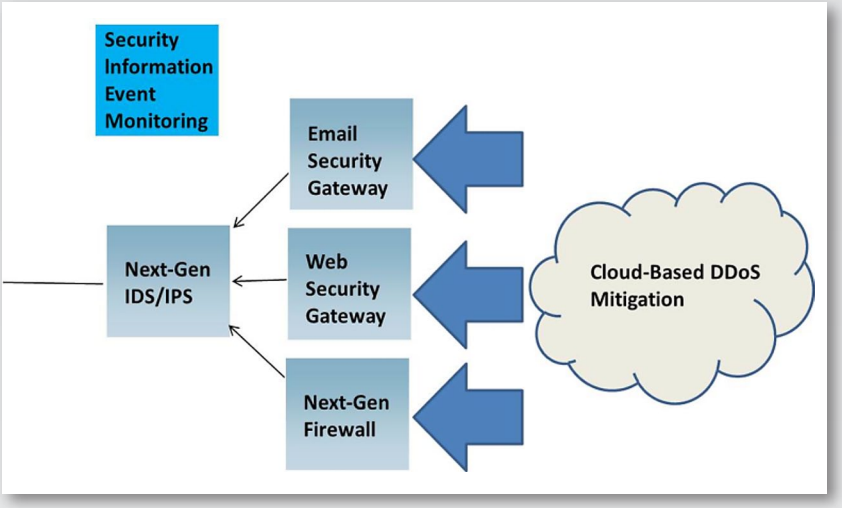

Figure 2. Modern Network Perimeter Security

Newer "next generation" versions of firewalls and IPS appliances have added techniques to detect more sophisticated attacks by their traffic patterns. However, a number of deficiencies allow ATAs to succeed against these as well:

- **Attacks from inside the network.** When a user's laptop is compromised while outside the perimeter and the user returns to his or her business network, there is no inbound traffic to inspect because that laptop is now inside the protected perimeter.
- **Evading detection.** When the evasion techniques used by attackers (such as encryption or fragmented packets) are specifically designed to blind pattern recognition and deep packet inspection techniques, the attacks can evade detection.
- **Inflexible defenses.** When DDoS techniques are added to the mix, high rates of traffic often cause inline network defenses to choose between failing open (allowing traffic to pass without inspection to maintain availability), or failing closed (blocking all traffic to maintain security but causing business disruption). Neither of these meets the optimal requirement of choking off only offensive traffic while keeping legitimate traffic flowing.

Operational gaps can also cause the typical perimeter defense to be ineffective against sophisticated ATAs. Many organizations have outsourced email, web and firewall security operations and monitoring to external managed security service providers or security-as-a-service providers to reduce staffing requirements. Often, ISPs are used for DDoS detection and scrubbing as well. These approaches can limit the amount of data available to the security staff to detect compromised PCs and servers operating on the internal network and make it harder to coordinate incident response. In some organizations, confusion over who within the organization is responsible for DDoS detection and mitigation can also result in unfocused or weak defenses.

Using DDoS Detection and Mitigation Techniques to Defend Against Advanced Threats

Because DDoS techniques can be used as part of multivector ATAs, DDoS mitigation can play a role in defending against them. The targeted nature of both ATAs and application-layer DDoS attacks means some common capabilities can be used to detect them both.

One example is the unusual traffic flow patterns caused by DDoS as well as ATAs, especially during their final extraction phase. Resource depletion DDoS attacks cause unusual access patterns to a web server and, often, unusual traffic flows between a web server and database server and/or application server. The ability to detect flow anomalies or suspicious patterns can thus help detect both types of attacks.

ATAs and application-layer DDoS attacks also both use evasion techniques such as transport encryption (Secure Sockets Layer) to evade detection. Where it is not possible to force network traffic through a friendly "man in the middle" decryption point, flow analysis is required to support analysis and detection.

Tools to detect both types of attacks must be able to detect both application and network traffic and make assumptions about their nature. For example, in the case of a custom application that is being exploited for a directed attack to steal customer data, the detection solution must be able to identify and highlight the fact the attack is producing new application traffic not normally seen on the network.

DDoS mitigation products combine a number of capabilities to detect both volumetric and application-level DDoS attacks. By monitoring and maintaining a history of network flows, these appliances can quickly spot such anomalies in network traffic. Used alone, this technique would generate high levels of false positives, so DDoS products can also identify applications and apply behavioral filtering to prioritize highly suspicious flows and events for each application.[7]

Looking at all internal network flows to indicate anomalies in the rate, direction and symmetry of traffic allows security analysts to look for the "footprints" the ATAs create. These same capabilities are necessary for detecting and disrupting the various stages of ATAs.

During the initial compromise phase, anomalies in the use of protocols or encrypted communications can quickly identify the start of an attack. Once an internal machine has been compromised, flow analysis can detect both internal reconnaissance attempts and communications out to command and control centers. Finally, if the attack has remained undetected until sensitive information is being extracted, flow and behavioral analysis can detect that extraction while it is underway.

---

[7] http://ijcns.uacee.org/vol1iss1/files/V1-I1-62.pdf

DDoS mitigation techniques essentially add the network security monitoring (NSM) capability shown in Figure 3 to an organization's other layers of information security. NSM is not only a key component of DDoS detection, but can be used by network operations groups to diagnose performance issues and isolate problem areas.

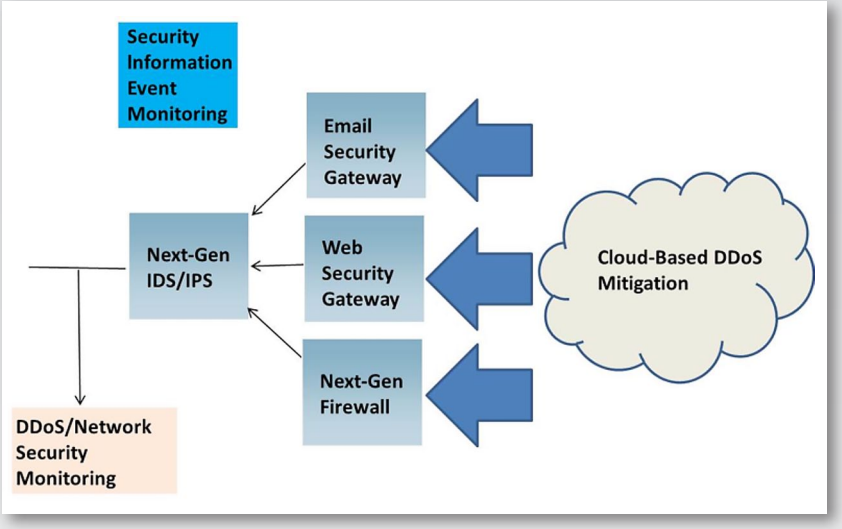

Figure 3. Example of a DDoS/Advanced Attack Mitigation Technique

By integrating DDoS detection and mitigation techniques into the network security architecture, enterprises can take advantage of the additional security intelligence to more quickly detect potential ATAs and take measured steps to prevent a compromise without disrupting legitimate traffic. Adding these capabilities also helps identify attacks that might be missed by other means. For example, a CFO's PC compromised with an ATP that attempts to infiltrate an internal server will likely go unnoticed by the external facing firewall, but it can trigger a DDoS NSM event. If the attack reaches the extraction phase, flow analysis performed by DDoS detection would cause an alert based on unusual outbound traffic. This is why it is important to integrate not only syslogs and alarms, but also the ability to see otherwise legitimate-looking traffic in a location or at a time that signals a possible attack.

Deployment Options

The DDoS filtering capabilities of operational devices such as load balancers, firewalls or intrusion prevention systems may provide acceptable prevention for small businesses. But high-volume attacks can overwhelm these devices, and sophisticated application-level DDoS attacks can often evade them. That is why larger enterprises, or even medium-sized businesses with revenue-critical systems that depend on high-availability Internet connectivity, need dedicated DDoS mitigation capabilities.

DDoS mitigation capabilities are available from most ISPs at an added cost or from third party "man in the middle" DDoS mitigation service providers. These services provide a high level of protection and can eliminate the "fill the pipe" problem, where high-volume attacks that are blocked on-premise still result in denial of service because the entire bandwidth of the Internet connection is consumed. However, many enterprises use multiple ISPs or find the cost of ISP or cloud-provided DDoS mitigation too high for certain geographies or websites. Relying solely on external DDoS mitigation also makes it more difficult to take advantage of the DDoS monitoring capabilities in combating advanced targeted attacks. For these reasons, many large enterprises use a hybrid configuration, where ISP or cloud-based protection is used as the first level of protection and local or customer premise-based DDoS prevention is used on site.

There are a number of options for deploying the local portion of your DDoS mitigation efforts. As you choose among them, note that being able to analyze flows across all network segments can provide the greatest level of visibility into, and protection from, ATAs—but at additional cost.

The most common placement for DDoS mitigation is at the front end of the traditional Internet DMZ, as shown in Figure 4.

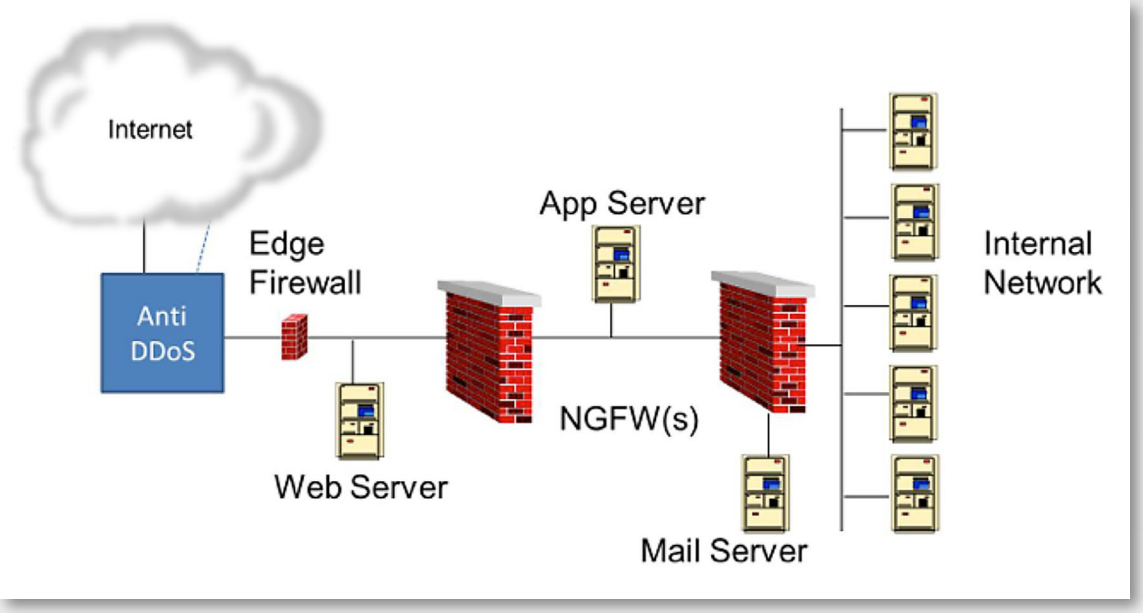

Figure 4. DDoS Mitigation at the Internet Edge

In this configuration, a DDoS mitigation appliance is used at the Internet edge of the DMZ, often in conjunction and coordination with cloud-based DDoS services. This provides visibility into all flows in and out of the enterprise. It also prevents DDoS attacks from choking operational web, application and mail servers, as well as critical security elements such as next-generation firewalls. Often DDoS attacks below a certain bandwidth are not detected by ISP-based services (as occurred in the Operation Ababil exploit) and need to be detected by network operations or network security staff using the DDoS mitigation appliance at the Internet edge of the DMZ. If the attackers step up the DDoS attack, at some point mitigation can be enabled at the service provider.

This configuration generally requires the lowest investment in local DDoS mitigation and can support detection and prevention of ATAs that are primarily "outside in." However, the lack of visibility into internal network flows limits its usefulness in dealing with other ATAs.

The ultimate targets of ATAs and DDoS attacks are often the mission-critical business servers in the data center. To defend them, DDoS detection and mitigation can be placed between the data center and the rest of the enterprise WAN (see Figure 5), where flow analysis and application awareness can be used to protect business servers from both forms of attack.

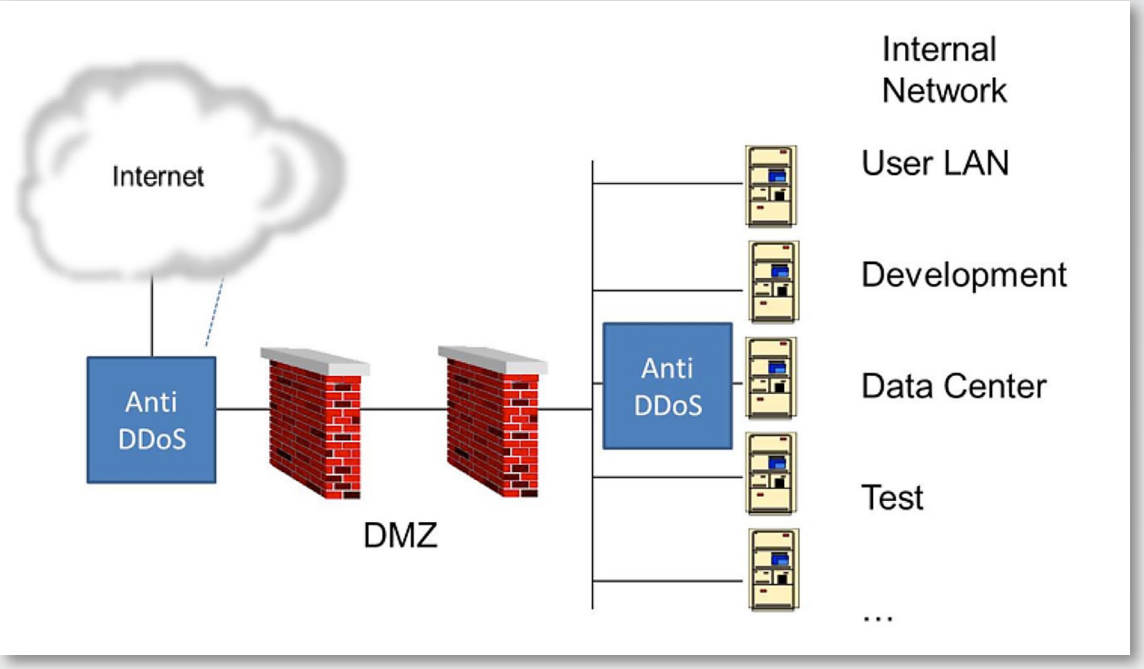

Figure 5. DDoS Mitigation in the Data Center

By having visibility into flows to and from the data center, this configuration also provides visibility into the later stages of ATAs, where compromised users' PCs are attempting to install APTs on data center servers, or where compromised servers are making unusual communication attempts to other internal servers or external locations.

The downside of this configuration is the cost of the additional anti-DDoS appliance. Companies that find cloud-based DDoS mitigation sufficient to protect DMZ-based servers can deploy a single DDoS appliance at the data center to gain the benefit without the cost.

Many enterprises are increasingly moving some business applications from their internal data centers to external cloud service providers, whether by purchasing applications in the form of Software as a Service (SaaS) or by running those applications on cloud-based Infrastructure as a Service (IaaS). In either scenario, enterprises have two choices.

One option is to route all traffic to and/or from the cloud through a "man in the middle" DDoS provider. The other is to use the DDoS mitigation capabilities provided by the SaaS or IaaS provider. Because they are completely dependent on the availability of their Internet connection for revenue, the enterprise-class providers have strong DDoS scrubbing capabilities, generally using a hybrid configuration of their own premise-based capabilities along with an ISP's or an external DDoS mitigation service's offerings. The better SaaS and IaaS providers also provide APIs or other mechanisms that allow customers to integrate their own data center or DMZ-based DDoS monitoring with that used in the SaaS/IaaS infrastructure.

Many organizations use security information and event management (SIEM) products to correlate log events across numerous security controls. The flow and behavior anomaly alerts from DDoS mitigation can provide valuable information to SIEM products to differentiate between low-level security events (for example, legitimate database extracts such as backups) and critical security events such as ATAs. Here again, it is vital the integration include insights into the traffic traversing the network, not just syslogs or alarms.

Summary

Today's DDoS attacks, as damaging as they are, sometimes mask even more threatening and dangerous advanced targeted threats. The good news is that some of the same tools that can detect the footprints of DDoS attacks can also find the telltale signs of ATAs. To be successful, though, enterprises need to coordinate their use of both types of logs and adjust their monitoring parameters correctly. They must also eliminate operational gaps, such as unclear lines of security authority, and limit the use of managed security services that can deprive enterprise security managers of the data they need to detect and stop DDoS attacks and ATAs.

About the Author

John Pescatore joined SANS in January 2013, with 35 years of experience in computer, network and information security. He was Gartner's lead security analyst for more than 13 years, working with global 5000 corporations and major technology and service providers. In 2008, he was named one of the top 15 most influential people in security and has testified before Congress on cybersecurity.

Prior to joining Gartner Inc. in 1999, John was senior consultant for Entrust Technologies and Trusted Information Systems, where he started, grew and managed security consulting groups focusing on firewalls, network security, encryption and public key infrastructures. Prior to that, he spent 11 years with GTE developing secure computing and telecommunications systems. In 1985 he won a GTE-wide Warner Technical Achievement award.

Mr. Pescatore began his career at the National Security Agency, where he designed secure voice systems, and the United States Secret Service, where he developed secure communications and surveillance systems—and the occasional ballistic armor installation. He holds a bachelor's degree in electrical engineering from the University of Connecticut and is an NSA-certified cryptologic engineer. He is an Extra class amateur radio operator, callsign K3TN.

SANS would like to thank its sponsor:

ARBOR® NETWORKS