Four Steps to Defeat a DDoS Attack

Millions of computers around the world are controlled by cybercriminals. These computers have been infected with software robots, or "bots", that automatically connect to command and control servers. The command and control servers then instruct the bots to carry out illicit activity, such as performing denial of service attacks, or harvesting application content. Building these networks of bots, or botnets, has become a lucrative business for botnet operators, who rent out their bots to the highest bidder.

One of the most dangerous botnet threats is the Distributed Denial of Service (DDoS) attack. Harnessing the aggregate power of thousands or tens of thousands of bots, DDoS attack can inflict tremendous damage on Websites, slowing down or even completely disabling them. And DDoS attacks are not isolated, but a regular issue for many organizations. According to recent survey of IT decision makers, 74% reported suffering one or more DDoS attacks in the past 12 months. Of these, 31% said that the attacks disrupted service.[1] Whether the motivation is political, financial or just random, DDoS attacks can be extraordinarily costly for the targeted organizations.

> *Like slick advertising executives, botnet operators and even bot malware creators promote their offerings with carefully fine-tuned messaging.*

[1] "The Trends and Changing Landscape of DDoS Threats and Protection," Forrester

Application DDoS

A Distributed Denial of Service (DDoS) attack is an attack initiated from multiple machines that is designed to disrupt normal operations. Traditional Denial of Service (DoS) attacks attempt to exploit server or application weaknesses to cause it to stop responding. DDoS attacks amplify the effects of DoS attacks by using thousands of machines to launch their assaults. These new attacks may not necessarily exploit vulnerabilities, they may just unleash a flood of requests, overwhelming the bandwidth and server processing power of the targeted site.

The End Game for DDoS

DDoS attacks have targeted a diverse range of organizations, from government institutions and banks, to social networking companies and even root name server operators. The motivations for DDoS attacks vary: financial, political, religious, entertainment, or even personal notoriety. Many organized cyber criminals use DDoS to extort money from online sites. Authorities convicted a Russian gang of blackmailing over 50 organizations, extracting over $4 million from British companies, typically online gambling sites.[2] In 2008, a wave of DDoS attacks brought down 10 online gambling sites, also purportedly targets of extortion schemes.

Hacktivism is another key motivation for DDoS attacks. Whether driven by national patriotism or the desire to squelch the opinions of an ideological foe, DDoS is the weapon on choice. Examples of hacktivism in action include DDoS attacks targeting Georgian Websites before the Ossetia War in 2008 and the Iranian government's Website during the 2009 Iranian election protests. Government Websites representing the US, Korea, Myanmar, Estonia, and many others have been targeted. In fact, a persistent DDoS attack on Burmese Websites during the Burma's 2010 national elections actually caused the entire country's Internet connectivity to go down. More recently, WikiLeaks has found itself in the center of a DDoS hacktivism war, named "Operation Payback". Hacktivists attacked the MasterCard, Visa and PayPal Websites in retaliation after these companies stopped processing donations to WikiLeaks. Imperva's ADC had tracked "Operation Payback" and had witnessed how this campaign had evolved. In the first stage of the campaign, individuals used a manually-tuned DDoS attack tool. The tool was later enhanced to become an automated DDoS attack tool, allowing any individual without any technical knowledge to participate in a full-fledged DDoS attack. In effect, participants were joining forces to form a "voluntary botnet". As "Operation Payback" continued, it had reached a stage where botnet farmers were donating the bots under their control as their contribution to the DDoS campaign[3].

DDoS Botnets-for-Hire

While the WikiLeaks-inspired "Operation Payback" attack used a combination of voluntary hackers and bots, almost all DDoS attacks are executed by criminal botnet services. DDoS rental fees typically start at $50 for small attacks, but some researchers have seen DDoS prices as low as $9. To attract customers, botnet owners advertise their services, continually seeking to outclass their botnet brethren. Owners promote their services in underground forums and mailing lists. In the case of the powerful IMDDOS botnet, the owners actually set up a public Website to showcase their offering.[4] On a message board, one botnet operator touted that his botnet offered "the best combination of quality and service" and special pricing for regular customers. Options included HTTP attacks, downloading flood, POST flood, and ping commands "tuned to perfection."[5] Like slick advertising executives, botnet operators and even bot malware creators promote their offerings with carefully fine-tuned messaging.

---

2 "Online Russian blackmail gang jailed for extorting $4m from gambling websites", http://www.sophos.com/pressoffice/news/articles/2006/10/extort-ddos-blackmail.html

3 "Operation Payback: How it Works", http://blog.imperva.com/2010/12/operation-payback-how-it-works.html

4 "Damballa Discovers New Wide-Spread Global Botnet Offering Commercial DDoS Services," Damballa, September 2010

5 "BlackEnergy competitor – The 'Darkness' DDoS Bot," Shadowserver calendar entry for December 5, 2010

DDoS 2.0

DDoS attacks traditionally are carried out by computer-based bots. The Imperva ADC uncovered a new breed of DDoS attacks in May 2010[6] that uses Web servers as payload-carrying bots. Imperva discovered a 300-server strong botnet that set a new standard for power, efficiency and stealth. Using a basic software program equipped with a dashboard and control panel, hackers could configure the IP, port, and duration of the attack. Hackers simply need to type the Website URL they wish to attack and then they can instantly disable targeted sites.

In fact, a single Web server is equal to 3,000 bot infected PCs. With such powerful attack weapons at their command, it is not surprising that DDoS rental services keep increasing the strength of their attacks.

Advanced Application DDoS Attacks

Many organizations witnessed an increase in application-based attacks in 2009 compared to previous years. While application-based attacks still only account for 26% of all DDoS attacks, they are more sophisticated and much more challenging to stop. There are several reasons why application-based attacks are the most dangerous type of DDoS. Network firewalls today can detect the majority of flood and network DoS attacks. Many ICMP and UDP flood attacks can also be identified using intelligent packet filtering and source and destination access control lists. However, application DDoS attacks usually bypass most traditional network security devices.

Application DDoS exploit vulnerabilities in application servers or application business logic. For example, application DDoS attacks may simply flood a Web application server with seemingly legitimate requests designed to overwhelm Web application servers. An attacker may also attempt to exploit an application vulnerability, such as sending Web requests with extremely long URLs. More sophisticated attacks exploit business logic flaws. For example, if an application's Website search mechanism is poorly written, it could require excessive processing by a back end database server. An application DDoS attack could exploit this vulnerability by performing thousands of search requests using wildcard search terms to overwhelm the back end application database.

"Slowloris" emerged as a perilous application DDoS attack in 2009. This attack disrupts application service by exhausting web server connections. In the Slowloris attack, the attacker sends an incomplete HTTP header and then periodically sends header lines to keep the connection alive, but never sends the full header. Without requiring that much bandwidth, an attacker can open numerous connections and overwhelm the targeted Web server. While multiple patches have been created for Apache to mitigate this vulnerability, it nonetheless demonstrates the power of more sophisticated DDoS attacks.

DDoS Mitigation Techniques

There are a number of measures that organizations can undertake to mitigate the risks of a DDoS attack. Organizations can:

1. **Over-provision bandwidth to absorb DDoS bandwidth peaks** – Although this is the most common measure to alleviate DDoS attacks, it is also probably the most expensive, especially since DDoS attacks can be ten times or even one hundred times greater than standard traffic levels. An alternative to over-provisioning Internet bandwidth is to use a security service to scale on-demand to absorb and filter DDoS traffic. DDoS protection services are designed to stop massive DDoS attacks without burdening businesses' Internet connections.

2. **Monitor application and network traffic** – The best way to detect when you are under an attack is by monitoring application and network traffic. Then, you can determine if poor application performance is due to service provider outages or a DDoS attack. Monitoring traffic also allows organizations to differentiate legitimate traffic from attacks. Ideally, security administrators should review traffic levels,

[6] "Security Advisory: DDoS Advisory – May 2010", http://www.imperva.com/resources/adc/adc_advisories_DDOS_Attack_Method_Payload-05182010.html

application performance, anomalous behavior, protocol violations, and Web server error codes. Since DDoS attacks are almost always executed by botnets, application tools should be able to differentiate between standard user and bot traffic. Monitoring application and network traffic provide IT security administrators instant visibility into DDoS attack status.

3. **Detect and Stop Malicious Users** – There are two primary methods to identify DDoS attack traffic: identify malicious users and identify malicious requests. For application DDoS traffic, often times identifying malicious users can be the most effective way to mitigate attacks.

   » Recognize known attack sources, such as malicious IP addresses that are actively attacking other sites, and identifying anonymous proxies and TOR networks. Known attack sources account for a large percentage of all DDoS attacks. Because malicious sources constantly change, organizations should have an up-to-date list of active attack sources.

   » Identify known bot agents; DDoS attacks are almost always performed by an automated client. Many of these client or bot agents have unique characteristics that differentiate them from regular Web browser agents. Tools that recognize bot agents can immediately stop many types of DDoS sources.

   » Perform validation tests to determine whether the Web visitor is a human or a bot. For example, if the visitor's browser can accept cookies, perform JavaScript calculations or understand HTTP redirects, then it is most likely a real browser and not a bot script.

4. **Detect and Stop Malicious Requests** – Because application DDoS attacks mimic regular Web application traffic, they can be difficult to detect through typical network DDoS techniques. However, using a combination of application-level controls and anomaly detection, organizations can identify and stop malicious traffic. Measures include:

   » Detect an excessive number of requests from a single source or user session – Automated attack sources almost always request Web pages more rapidly than standard users.

   » Prevent known network and application DDoS attacks – Many types of DDoS attacks rely on simple network techniques like fragmented packets, spoofing, or not completing TCP handshakes. More advanced attacks, typically application-level attacks, attempt to overwhelm server resources. These attacks can be detected through unusual user activity and known application attack signatures.

   » Distinguish the attributes, and the aftermath, of a malicious request. Some DDoS attacks can be detected through known attack patterns or signatures. In addition, many malicious Web requests do not conform to HTTP protocol standards. For instance, the Slowloris DDoS attack included redundant HTTP headers. In addition, DDoS clients may request Web pages that do not exist. Attacks may also generate Web server errors or slow Web server response time.

Summary

Over the past several years, DDoS attacks have become industrialized. Using off-the-shelf toolkits, automation techniques, and search engines, non-technical cyber criminals can build botnets of thousands or even millions of computers. Using botnets, malicious users can unleash destructive DDoS attacks on virtually any victim.

The aforementioned techniques are just a few of the measures that organizations can undertake to combat DDoS attacks. They should be combined with processes, such as developing an internal rapid response team that can quickly and adeptly analyze and address DDoS attacks. If organizations undertake effective security measures, they will be well equipped to fight DDoS attacks.

About Imperva

Imperva is the global leader in data security. Our customers include leading enterprises, government organizations, and managed service providers who rely on Imperva to prevent sensitive data theft by hackers and insiders. The award-winning Imperva SecureSphere is the only solution that delivers full activity monitoring for databases, Web applications and file systems.

To learn more about Imperva's solution visit http://www.imperva.com.

Imperva
Headquarters
3400 Bridge Parkway, Suite 200
Redwood Shores, CA 94065
Tel: +1-650-345-9000
Fax: +1-650-345-9004

Toll Free (U.S. only): +1-866-926-4678
www.imperva.com