DDoS HANDBOOK

radware

THE ULTIMATE GUIDE TO EVERYTHING YOU NEED TO KNOW ABOUT **DDoS ATTACKS**

SHARE THE DDoS HANDBOOK in f t g+

Table of Contents

1 Introduction

Since the first denial of service (DoS) was launched in 1974, distributed denial of service (DDoS) and other DoS attacks have remained among the most persistent and damaging cyber-attacks. These attacks reflect hackers' frustratingly high levels of tenacity and creativity—and create complex and dynamic challenges for anyone responsible for cyber security.

While cyber-threats are by nature a moving target, this primer offers an overview to help detect and mitigate attacks. Radware's DDoS Handbook delivers:

- Brief history of DDoS attacks plus a roundup of recent cyber-attacks
- Overview of major attack types and tools
- Brief discussion of the ongoing evolution of enterprise security
- Actionable tools and tips for attack detection and mitigation
- Detailed vendor evaluation checklist for DDoS and cyber-attack detection and mitigation
- DDoS dictionary to help communicate about and address threats

Throughout the handbook, you'll also encounter some key findings and analysis from Radware's 2014-2015 Global Application & Network Security Report—one of the industry's leading pieces of research into DDoS and other cyber-attacks.

2 A Quick Look Back

In 2014, the DoS attack celebrated its 40th birthday. Born as the handiwork of a teenaged “computer geek,” these attacks have since exploded in quantity—and sophistication.

The Early Days

The first-ever DoS attack occurred in 1974 courtesy of David Dennis—a 13-year-old student at University High School, located across the street from the Computer-Based Education Research Laboratory (CERL) at the University of Illinois at Urbana-Champaign. David learned about a command that could be run on CERL’s PLATO terminals. PLATO was one of the first computerized shared learning systems, and a forerunner of many future multi-user computing systems. Called “external” or “ext,” the command was meant to allow for interaction with external devices connected to the terminals. However, when run on a terminal with no external devices attached the command would cause the terminal to lock up—requiring a shutdown and power-on to regain functionality.

Curious to see what it would be like for a room full of users to be locked out at once, David wrote a program that would send the “ext” command to many PLATO terminals at the same time. He went over to CERL and tested his program—which succeeded in forcing all 31 users to power off at once. Eventually the acceptance of a remote “ext” command was switched off by default, fixing the problem.

During the mid- to late 1990s, when Internet Relay Chat (IRC) first became popular, some users fought for control of non-registered chat channels, where an administrative user would lose his or her powers if he or she logged off. This behavior led hackers to attempt to force all users in a channel to log out, so hackers could enter the channel alone and gain administrator privileges as the only user present. These “king of the hill” battles—in which users would attempt to take control of an IRC channel and hold it in the face of attacks from other hackers—were fought using very simple bandwidth-based DoS attacks and IRC chat floods.

DDoS Attacks Spread

One of the first large-scale DDoS attacks occurred in August 1999, when a hacker used a tool called “Trinoo” to disable the University of Minnesota’s computer network for more than two days. Trinoo consisted of a network of compromised machines called “Masters” and “Daemons,” allowing an attacker to send a DoS instruction to a few Masters, which then forwarded instructions to the hundreds of Daemons to commence a UDP flood against the target IP address. The tool made no effort to hide the Daemons’ IP addresses, so the owners of the attacking systems were contacted and had no idea that their systems had been compromised and were being used in an attack.

Other early tools include “Stacheldraht” (German for barbed wire), which could be remotely updated and support IP spoofing, along with “Shaft” and “Omega”, tools that could collect attack statistics from victims. Because hackers were able to get information about their attacks, they could better understand the effects of certain types of attacks, as well as receive notification when an attack was detected and stopped.

Once hackers began to focus on DDoS attacks, DDoS attacks attracted public attention. The distributed nature of a DDoS attack makes it significantly more powerful, as well as harder to identify and block its source. With such a formidable weapon in their arsenals, hackers took on larger, more prominent targets using improved tools and methods.

By the new millennium, DDoS attacks captured the public’s attention. In the year 2000, various businesses, financial institutions and government agencies were brought down by DDoS attacks. Shortly after, DNS attacks began with all 13 of the Internet’s root domain name service (DNS) servers being attacked in 2002. DNS is an essential Internet service, as it translates host names in the form of uniform resource locators (URLs) into IP addresses. In effect, DNS is a phonebook maintaining a master list of all Internet addresses and their corresponding URLs. Without DNS, users would not be able to efficiently navigate the Internet, as visiting a website or contacting a specific device would require knowledge of its IP address.

From Script Kiddies to Geo-Political Events

As attack technology evolved, so have motivations and participants. Today, we no longer face only teenage “computer geeks” or “script kiddies” testing the limits of what they can do. While they still exist, they are not alone. Recent years have brought a continuous increase in the number of DDoS attacks—fueled by changing and increasingly complex motivations.

Timeline

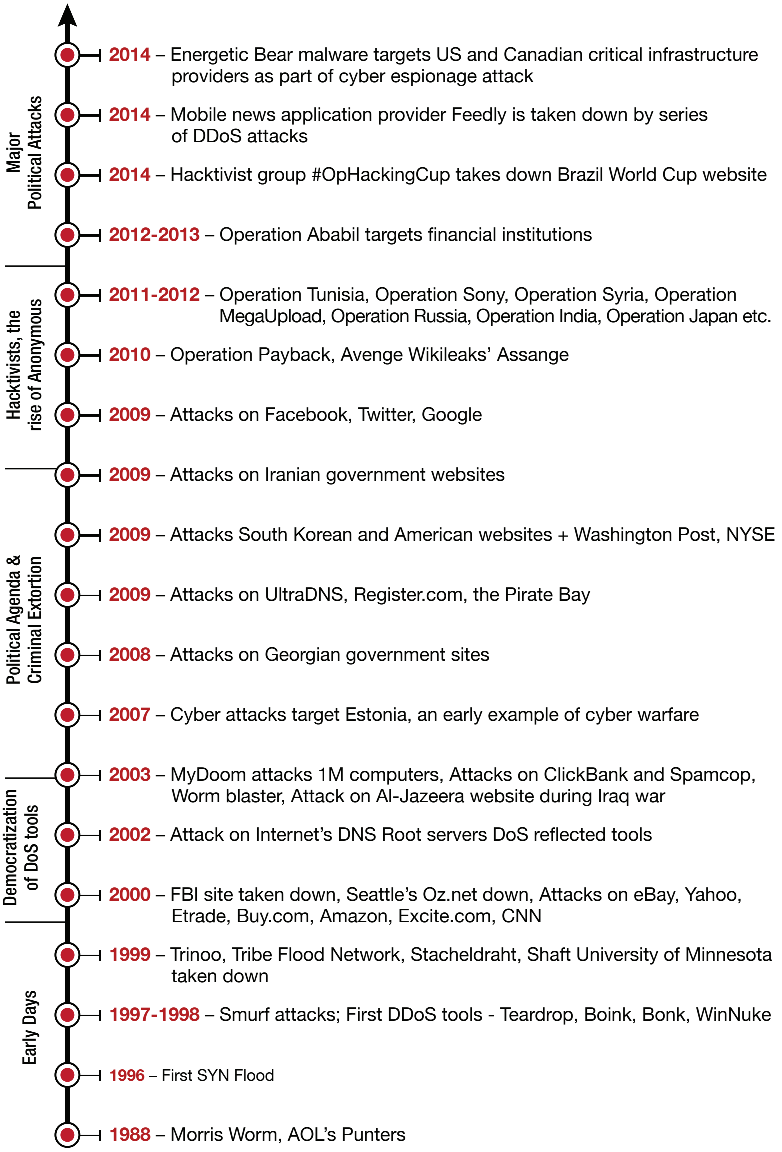

Major Political Attacks

2014 – Energetic Bear malware targets US and Canadian critical infrastructure providers as part of cyber espionage attack

2014 – Mobile news application provider Feedly is taken down by series of DDoS attacks

2014 – Hacktivist group #OpHackingCup takes down Brazil World Cup website

2012-2013 – Operation Ababil targets financial institutions

Hacktivists, the rise of Anonymous

2011-2012 – Operation Tunisia, Operation Sony, Operation Syria, Operation MegaUpload, Operation Russia, Operation India, Operation Japan etc.

2010 – Operation Payback, Avenge Wikileaks’ Assange

2009 – Attacks on Facebook, Twitter, Google

Political Agenda & Criminal Extortion

2009 – Attacks on Iranian government websites

2009 – Attacks South Korean and American websites + Washington Post, NYSE

2009 – Attacks on UltraDNS, Register.com, the Pirate Bay

2008 – Attacks on Georgian government sites

2007 – Cyber attacks target Estonia, an early example of cyber warfare

Democratization of DoS tools

2003 – MyDoom attacks 1M computers, Attacks on ClickBank and Spamcop, Worm blaster, Attack on Al-Jazeera website during Iraq war

2002 – Attack on Internet’s DNS Root servers DoS reflected tools

2000 – FBI site taken down, Seattle’s Oz.net down, Attacks on eBay, Yahoo, Etrade, Buy.com, Amazon, Excite.com, CNN

Early Days

1999 – Trinoo, Tribe Flood Network, Stacheldraht, Shaft University of Minnesota taken down

1997-1998 – Smurf attacks; First DDoS tools - Teardrop, Boink, Bonk, WinNuke

1996 – First SYN Flood

1988 – Morris Worm, AOL’s Punters

Figure 1

3 Recent History: Notable Cyber-Attacks of 2014

This section provides an overview of recent and notable cyber-attacks of 2014 with categorization for types of attacks: breach, outage, technical.

⚠ **Breach** 🔌 **Outage** ⚙ **Technical**

January to March

⚠ Yahoo! email service for 273 million users reportedly hacked, although the specific number of affected accounts is not released.
🔌 Bitcoin hit with code integrity issues and DDoS attacks.
⚙ Newly released NTP DDoS vulnerabilities uncovered.
🔌 UK Ministry of Justice, UK Government Communication Headquarters disrupted by DDoS attacks.
⚠ Credit card information of 350,000 individuals was stolen via Neiman Marcus, with more than 9,000 of the cards used fraudulently since the attack. Sophisticated code written by the hackers allowed them to spend months moving through company computers, undetected by employees.

April to June

⚙ Newly released Heartbleed vulnerability published.
⚠ Five Chinese nationals indicted for computer hacking and economic espionage of U.S. companies between 2006 and 2014.
🔌 Ukrainian/Russian cyber-war flared, targeting countries participating in the conflict.
⚠ According to the Department of Homeland Security, hackers accessed an unnamed public utility’s control system through a brute-force attack on employees’ log-in passwords.
🔌 Feedly’s 15 million users disrupted by numerous DDoS attacks.
🔌 In the same week as the Feedly cyber-attack, Evernote and its 100 million users faced a similar DoS attack.
🔌 Anonymous launched successful DDoS campaign against Boston Children’s Hospital, disrupting hospital and healthcare operations.
⚠ Credit and debit card information from 33 P.F. Chang’s restaurants was compromised and reportedly sold online.
🔌 DDoS hit sponsors and organizers of the 2014 World Cup, disrupting numerous broadcasts, news and marketing events.

July to September

- ⚙ Bash/Shellshock vulnerability released, affecting millions of network devices worldwide.
- ⚠ U.S. Investigations Services, a subcontractor for federal employee background checks, suffered a data breach in August, leading to theft of employee information.
- ⚙ New Tsunami DDoS vulnerability technique provided for powerful new volumetric DDoS capabilities for attackers.
- ⚠ Unnoticed until August, a June attack on J.P. Morgan Chase compromised contact information for 76 million households and 7 million small businesses. Hackers may have originated in Russia, with possible ties to the Russian government.
- ⚙ The FBI issued Brobot Alert, including a list of 1,492 URLs of confirmed infected Web sites, with the request that organizations help victims to remove the malware.
- ⚙ Google uncovered SSLv3 “Poodle” vulnerability, later updated to include Transport Layer Security.

October to December

- ⚠ Sony Pictures hit in much-publicized attack around the release of the movie The Interview. The attack disrupted movie production, movie revenue and employee/talent relations.
- ⚙ Open SSL vulnerability released, affecting millions of pieces of software and hardware devices worldwide.
- ⚠ Credit and debit card information from 395 Dairy Queen and Orange Julius stores compromised by Backoff malware.
- ⚠ Photos of 200,000 users reportedly hacked from Snapsave, a third-party app for saving photos from instant photo-sharing app Snapchat.
- 🔌 Over the Christmas holiday, Sony PSN and Microsoft Xbox live attacked for days, rendering them unable to serve millions of customers worldwide.

4 Attack Types

This section provides an overview of major attack categories, as well as a breakdown of specific attack types within each.

Attacks Targeting Network Resources

Attacks that target network resources attempt to consume all of a victim's network bandwidth by using a large volume of illegitimate traffic to saturate the company's Internet pipe. These attacks, called network floods, are simple yet effective.

In a typical flooding attack, the offense is distributed among an army of thousands of volunteered or compromised computers—a botnet—that simply sends a huge amount of traffic to the targeted site overwhelms its network.

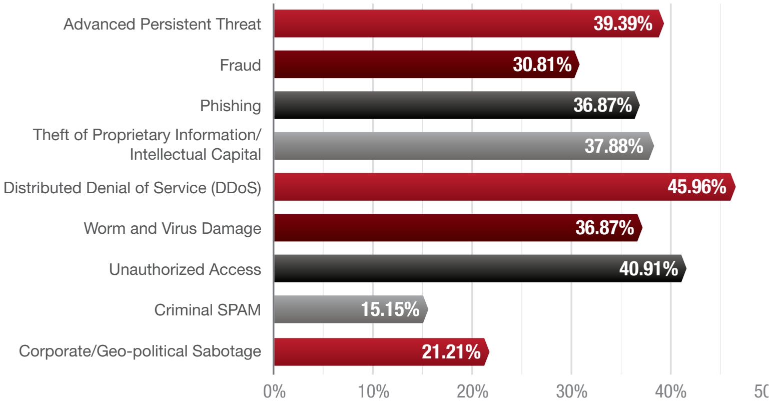

Figure 2: Attacks that will cause the most harm to business - Radware's 2014-2015 Global Application & Network Security Report.

In small numbers, requests of this manner may seem legitimate; in large numbers, they can be significantly harmful. A legitimate user trying to access a victim's site under a flooding attack will find the attacked site incredibly slow or unresponsive.

Types of Network Floods

UDP Flood: User Datagram Protocol (UDP) is a connectionless protocol that uses datagrams embedded in IP packets for communication without needing to create a session between two devices (in other words, it requires no handshake process).

A UDP Flood attack does not exploit a specific vulnerability. Instead, it simply abuses normal behavior at a high enough level to cause congestion for a targeted network. It consists of sending a large number of UDP datagrams from potentially spoofed IP addresses to random ports on a target server; the server receiving this traffic is unable to process every request, and consumes all of its bandwidth attempting to send ICMP “destination unreachable” packet replies to confirm that no application was listening on the targeted ports. As a volumetric attack, a UDP flood is measured in Mbps (bandwidth) and PPS (packets per second).

ICMP Flood: Internet Control Message Protocol (ICMP) is another connectionless protocol used for IP operations, diagnostics, and errors. Just as with a UDP flood, an ICMP flood (or Ping Flood) is a non-vulnerability based attack; that is, it does not rely on any specific vulnerability to achieve denial-of-service. An ICMP Flood can involve any type of ICMP message, such as a ping request (echo request and echo reply). Once enough ICMP traffic is sent to a target server, the server becomes overwhelmed from attempting to process every request, resulting in a denial-of-service condition. Like a UDP Flood, an ICMP Flood is also a volumetric attack, measured in Mbps (bandwidth) and PPS (packets per second).

IGMP Flood: Internet Group Management Protocol (IGMP) is another connectionless protocol. It is used by IP hosts (computers and routers) to report or leave multicast group memberships for adjacent routers. An IGMP Flood is non-vulnerability based, as IGMP is designed to allow multicast. Such floods involve a large number of IGMP message reports being sent to a network or router, significantly slowing and eventually preventing legitimate traffic from being transmitted across the target network.

Amplification Attacks: An Amplification attack takes advantage of a disparity between a request and a reply in technical communication. For instance, the attacker could use a router as an amplifier, taking advantage of the router’s broadcast IP address feature to send messages to multiple IP addresses in which the source IP (return address) is spoofed to the target IP. Famous examples of amplification attacks include Smurf Attacks (ICMP amplification) and Fraggle Attacks (UDP amplification). Another example of a type of amplification attack is DNS amplification, in which an attacker, having previously compromised a recursive DNS name server to cache a large file, sends a query directly or via a botnet to this recursive DNS server, which in turn opens a request asking for the large cached file. The return message (significantly amplified in size from the original request) is then sent to the victim's (spoofed) IP address, causing a denial-of-service condition.

Connection-Oriented Attacks:
A connection-oriented attack is one in which the attacker must first establish a connection prior to launching a DDoS attack. The outcome of this attack usually affects the server or application resources. TCP- or HTTP-based attacks are examples of connection-oriented DDoS attacks.

Connectionless Attacks:
A connectionless attack, on the other hand, does not require the attacker to open a complete connection to the victim, and therefore is much easier to launch. The outcome of a connectionless attack affects network resources, causing denial of service before the malicious packets can even reach the server. UDP floods and ICMP floods are examples of connectionless DDoS attacks.

Reflective Attacks:
An attack is reflective when the attacker makes use of a potentially legitimate third party to send his or her attack traffic, ultimately concealing his or her own identity.

Attack Motivations
Richard Clarke, former Special Advisor to the U.S. President on cyber-security, devised the "C.H.E.W." acronym to categorize and explain the origins of cyber-attack risks:

- **Cybercrime**
  The notion that someone is going to attack you with the primary motive being financial gain from the endeavor.

- **Hacktivisim**
  Attacks motivated by ideological differences. The primary focus of these attacks is not financial gain but rather persuading or dissuading certain actions or "voices."

- **Espionage**
  Straightforward motive of gaining information on another organization in pursuit of political, financial, capitalistic, market share or some other form of leverage.

- **War (Cyber)**
  The notion of a nation-state or transnational threat to an adversary's centers of power via a cyber-attack. Attacks could focus on non-military critical infrastructure or financial services.

Attacks Targeting Server Resources

Attacks that target server resources attempt to exhaust a server's processing capabilities or memory, potentially causing a denial-of-service condition. The idea is that an attacker can take advantage of an existing vulnerability on the target server (or a weakness in a communication protocol) to cause the target server to become so busy handling illegitimate requests that it no longer has the resources to handle legitimate ones. "Server" most commonly refers to a Website or Web application server, but these types of DDoS attacks can also target stateful devices, such as firewalls and intrusion prevention systems.

TCP/IP Weaknesses: These types of attacks abuse the TCP/IP protocol by exploiting some of its design weaknesses. They typically misuse the six control bits (or flags) of the TCP/IP protocol—SYN, ACK, RST, PSH, FIN and URG—in order to disrupt the normal mechanisms of TCP traffic. Unlike UDP and other connectionless protocols, TCP/IP is connection-based—requiring the packet sender to establish a full connection with his or her intended recipient prior to sending any packets. TCP/IP relies on a three-way handshake mechanism (SYN, SYN-ACK, ACK) where every request creates a half-open connection (SYN), a request for a reply (SYN-ACK), and then an acknowledgement of the reply (ACK). Attacks attempting to abuse the TCP/IP protocol will often involve sending TCP packets in the wrong order, causing the target server to run out of computing resources as it attempts to understand such abnormal traffic.

TCP SYN Flood: In the TCP handshake mechanism, there must be an agreement between each party for a connection to be established. If the TCP client does not exist or is a non-requesting client with a spoofed IP, such an agreement is not possible. In a TCP SYN, or simple SYN flood attack, the attacking clients lead the server to believe that they are asking for legitimate connections through a series of TCP requests with TCP flags set to SYN coming from spoofed IP addresses. To handle each of these SYN requests, the target server opens threads and allocates corresponding buffers to prepare for a connection. It then tries to send a SYN-ACK reply back to the requesting clients to acknowledge their connection requests, but because the clients' IP addresses are spoofed or the clients are unable to respond, an acknowledgement (ACK packet) is never sent back to the server. The server is still forced to maintain its open threads and buffers for each one of the original connection requests, attempting to resend its SYN-ACK request acknowledgement packets multiple times before resorting to a request timeout. Because server resources are limited and a SYN flood often involves a massive number of connection requests, a server is unable to time out its open requests before new requests arrive—causing a denial-of-service condition.

TCP RST Attack: The TCP RST flag is intended to notify a server that it should immediately reset its corresponding TCP connection. In a TCP RST attack, the attacker interferes with an active TCP connection between two entities by guessing the current sequence number and spoofing a TCP RST packet to use the client's source IP (which is then sent to the server). Typically a botnet is used to send thousands of such packets to the server with different sequence numbers, making it fairly easy to guess the correct one. Once this occurs, the server acknowledges the RST packet sent by the attacker, terminating its connection to the client located at the spoofed IP address.

TCP PSH+ACK Flood: When a TCP sender sends a packet with its PUSH flag set to 1, the result is that the TCP data is immediately sent or "pushed" to the TCP receiver. This action actually forces the receiving server to empty its TCP stack buffer and to send an acknowledgement when this action is complete. An attacker, usually using a botnet, can therefore flood a target server with many such requests. This overwhelms the TCP stack buffer on the target server, causing it to be unable to process the requests or even acknowledge them—resulting in a denial-of-service condition.

"Low and Slow" Attacks

Unlike floods, "low and slow" attacks do not require a large amount of traffic. They target specific design flaws or vulnerabilities on a target server with a relatively small amount of malicious traffic, eventually causing it to crash. "Low and slow" attacks mostly target application resources (and sometimes server resources). By nature, they are very difficult to detect because they involve connections and data transfer appearing to occur at a normal rate.

Sockstress: Sockstress is an attack tool that exploits vulnerabilities in the TCP stack—allowing an attacker to create a denial-of-service condition for a target server. In the normal TCP three-way handshake, a client sends a SYN packet to the server, the server responds with a SYN-ACK packet, and the client responds to the SYN-ACK with an ACK, establishing a connection. Attackers using Sockstress establish a normal TCP connection with the target server but send a “window size 0” packet to the server inside the last ACK, instructing it to set the size of the TCP window to 0 bytes. The TCP Window is a buffer that stores the received data before it uploads it up to the application layer. The Window size field indicates how much more room is in the buffer in each point of time. Window size set to zero means that there is no more space whatsoever and that the other side should stop sending more data until further notice.

In this case, the server will continually send window size probe packets to the client to see when it can accept new information. But because the attacker does not change the window size, the connection is kept open indefinitely. By opening many connections of this nature to a server, the attacker consumes all of the space in the server’s TCP connection table (as well as other tables), preventing legitimate users from establishing a connection. Alternately, the attacker may open many connections with a very small (around 4-byte) window size, forcing the server to break up information into a massive number of tiny 4-byte chunks. Many connections of this type will consume a server’s available memory, also causing a denial of service.

SSL-Based Attacks

Secure Socket Layer (SSL): a method of encryption used by various other network communication protocols—as it grows in prevalence, attackers began targeting it. Conceptually, SSL runs above TCP/IP, providing security to users communicating over other protocols by encrypting communications and authenticating communicating parties. SSL-based DoS attacks take many forms: targeting the SSL handshake mechanism, sending garbage data to the SSL server or abusing certain functions related to the SSL encryption key negotiation process. SSL-based attacks could also simply mean that the DoS attack is launched over SSL-encrypted traffic, which makes it extremely difficult to identify. Such attacks are often considered “asymmetric” because it takes significantly more server resources to deal with an SSL-based attack than it does to launch one.

Encrypted-based HTTP Attacks (HTTPS floods): Many online businesses increasingly use SSL/TLS (Transport Layer Security) in applications to encrypt traffic and secure end-to-end data transit. DoS attacks on encrypted traffic are on the rise, and mitigating them is not as obvious as might be expected. Most DoS mitigation technologies do not actually inspect SSL traffic, as it requires decrypting the encrypted traffic. HTTPS Floods—floods of encrypted HTTP traffic (see explanation below)—are now frequently participating in multi-vulnerability attack campaigns. Compounding the impact of “normal” HTTP Floods, encrypted HTTP attacks add several other challenges, such as the burden of encryption and decryption mechanisms.

THC-SSL-DoS: Hacking group The Hacker’s Choice (THC) developed this tool as a proof of concept to encourage vendors to patch SSL vulnerabilities. As with other “low and slow” attacks, THC-SSL-DoS requires only a small number of packets to cause denial of service for even a fairly large server. It works by initiating a regular SSL handshake, and then immediately requesting for the renegotiation of the encryption key. The tool constantly repeats this renegotiation request until all server resources have been exhausted. Attackers love to launch attacks that use SSL because each SSL session handshake consumes 15 times more resources from the server side than from the client side. In fact, a single standard home PC can take down an entire SSL-based web server, while several computers can take down a complete farm of large, secured online services.

Attacks Targeting Application Resources

Recent years have brought a rise in DoS attacks targeting applications. They target not only the well-known Hypertext Transfer Protocol (HTTP), but also HTTPS, DNS, SMTP, FTP, VOIP and other application protocols that possess exploitable weaknesses allowing for DoS attacks. Much like attacks targeting network resources, attacks targeting application resources come in a variety of flavors, including floods and “low and slow” attacks. Low and slow approaches are particularly prominent, mostly targeting weaknesses in the HTTP protocol—which, as the most widely used application protocol on the Internet, is an attractive target for attackers.

HTTP Flood: the most common DDoS attack targeting application resources. It consists of what seem to be legitimate, session-based sets of HTTP GET or POST requests sent to a victim's Web server, making it hard to detect. HTTP flood attacks are typically launched simultaneously from multiple computers (volunteered machines or bots). These bots continually and repeatedly request to download the target site's pages (HTTP GET flood), exhausting application resources and resulting in a denial-of-service condition. Modern DDoS attack tools, such as High Orbit Ion Cannon (HOIC), offer an easy-to-use means of performing multi-threaded HTTP flood attacks.

DNS Flood: is easy to launch yet difficult to detect. Based on the same idea as other flooding attacks, a DNS flood targets the DNS application protocol by sending a high volume of DNS requests. Domain Name System (DNS) is the protocol used to resolve domain names into IP addresses; its underlying protocol is UDP, taking advantage of fast request and response times without the overhead of having to establish connections (as TCP requires). In a DNS flood, the attacker sends multiple DNS requests to the victim's DNS server directly or via a botnet. The DNS server, overwhelmed and unable to process all of its incoming requests, eventually crashes.

"Low and Slow" Attacks: The characteristics of the "low and slow" attacks in this section relate particularly to application resources (whereas the previous "low and slow" attacks targeted server resources). These "low and slow" attacks target specific application vulnerabilities, allowing an attacker to stealthily cause denial of service. Not volumetric in nature, such attacks can often be launched with only a single machine. Additionally, because these attacks occur on the application layer, a TCP handshake is already established, successfully making the malicious traffic look like normal traffic traveling over a legitimate connection.

Slow HTTP GET Request: The idea behind a slow HTTP GET request is to dominate all or most of an application's resources through the use of many open connections, preventing it from providing service to users wishing to open legitimate connections. In this attack, the attacker generates and sends incomplete HTTP GET requests to the server, which opens a separate thread for each of these connection requests and waits for the rest of the data to be sent. The attacker continues to send HTTP header data at set, but slow, intervals to make sure the connection stays open and does not time out. Because the rest of the required data arrives so slowly, the server perpetually waits, exhausting the limited space in its connection table and thereby causing a denial-of-service condition.

Slow HTTP POST Request: To carry out a slow HTTP POST request attack, the attacker detects forms on the target website and sends HTTP POST requests to the Web server through these forms. The POST requests, rather than being sent normally, are sent byte by byte. As with a slow HTTP GET request, the attacker ensures that his or her malicious connection remains open by regularly sending each new byte of POST information slowly at regular intervals. The server, aware of the content length of the HTTP POST request, has no choice but to wait for the full POST request to be received (this behavior mimics legitimate users with slow Internet connection). The attacker repeats this behavior many times in parallel, never closes an open connection, and after several hundred open connections, the target server is unable to handle new requests—achieving a denial-of-service condition.

Regular Expression DoS Attacks: A special case of "low and slow" attacks is RegEx DoS (or ReDoS) attacks. In this scenario, the attacker sends a specially crafted message (sometimes called evil RegExes) that leverages a weakness in a library deployed in the server, in this case, a regular expression software library. This causes the server to consume large amounts of resources while trying to compute a regular expression over the user-provided input, or to execute a complex and resource-hungry regular expression processing dictated by the attacker.

Hash Collisions DoS Attacks: This kind of attack targets common security vulnerabilities in Web application frameworks. In short, most application servers create hash tables to index POST session parameters. Sometimes application servers must manage hash collisions when similar hash values are returned. Collision resolutions are resource intensive, as they require an additional amount of CPU to process the requests. In a Hash Collision DoS attack scenario, the attacker sends a specially crafted POST message with a multitude of parameters. The parameters are built in a way that causes hash collisions on the server side, slowing down the response processing dramatically. Hash Collisions DoS attacks are very effective and could be launched from a single attacker computer, slowly exhausting the application server's resources.

5 Attack Tools

Underscoring attackers' tenacity and creativity, a number of specialized attack tools has been created. Here are some of the most common—and threatening.

While it is possible to execute many types of DDoS attacks manually, specialized attack tools have been developed for the purpose of executing attacks more easily and efficiently. The turn of the century brought widespread use of the first DDoS tools—including Trinoo and Stacheldraht. However, these tools were somewhat complex and only ran on the Linux and Solaris operating systems. More recently, DDoS tools have evolved to target multiple platforms. They also have become more straightforward, rendering DDoS attacks much easier to carry out for attackers and more dangerous for targets.

What is the average security threat your organization experienced?

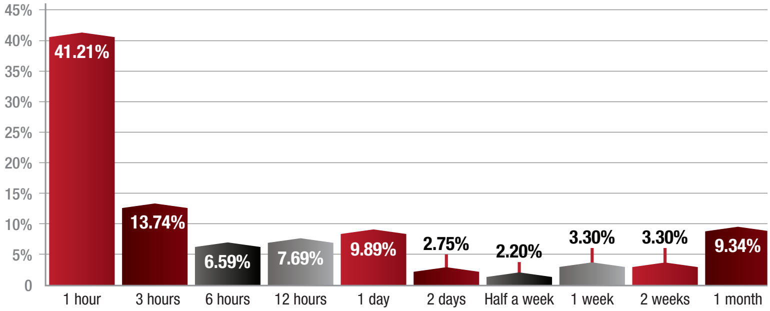

Figure 3: Average security threats -
Radware's 2014-2015 Global Application & Network Security Report.

Some of these newer DDoS tools, such as Low Orbit Ion Cannon (LOIC), were originally developed as network stress testing tools but were later modified and used for malicious purposes. Others, such as Slowloris, were developed by "gray hat" hackers whose aim is to direct the public's attention to a particular software weakness. By releasing such tools publicly, gray hat hackers force makers of vulnerable software to patch it in order to avoid large-scale attacks.

Of course, just as the network security and hacking world is continually evolving, so are the tools used to carry out DDoS attacks. Attack tools are becoming smaller, stealthier and more effective at causing a denial-of-service condition.

Low Orbit Ion Cannon (LOIC)

"Hacktivist" group Anonymous' first tool of choice—Low Orbit Ion Cannon (LOIC)—is a simple flooding tool that can generate massive volume of TCP, UDP or HTTP traffic in order to subject a server to a heavy network load. LOIC's original developers, Praetox Technologies, intended the tool to be used by developers who wanted to subject their own servers to a heavy network traffic load for testing purposes. However, Anonymous picked up the open-source tool and used it to launch coordinated DDoS attacks. Soon afterwards, LOIC was modified and given its "Hivemind" feature, allowing any LOIC user to point a copy of LOIC at an IRC server, transferring control of it to a master user who can then send commands over IRC to every connected LOIC client simultaneously. In this configuration, users are able to launch much more effective DDoS attacks than those of a group of less-coordinated LOIC users not operating simultaneously. In late 2011, however, Anonymous stepped away from LOIC as its DDoS tool of choice, as LOIC makes no effort to obscure its users' IP addresses. This lack of anonymity resulted in the arrest of various users around the world for participating in LOIC attacks, with Anonymous broadcasting a clear message across all of its IRC channels: "Do NOT use LOIC."

What is the maximum security threat your organization experienced?

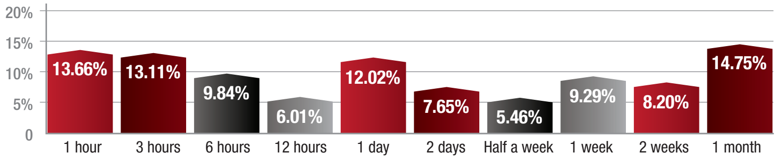

Figure 4: Maximum security threats -
Radware's 2014-2015 Global Application & Network Security Report.

How long can you efficiently fight a round-the-clock attack campaign?

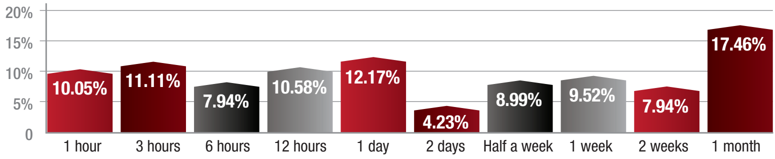

Figure 5: Maximum security threats -
Radware's 2014-2015 Global Application & Network Security Report.

High Orbit Ion Cannon (HOIC)

After Anonymous dropped LOIC as its tool of choice, High Orbit Ion Cannon (HOIC) quickly took the spotlight when it was used to target the United States Department of Justice in response to its decision to take down Megaupload.com. At its core, HOIC is also a simple application: a cross-platform basic script for sending HTTP POST and GET requests wrapped in an easy-to-use GUI. However, its effectiveness stems from add-on “booster” scripts—text files that contain additional basic code interpreted by the main application upon a user’s launch of an attack. Even though HOIC does not directly employ any anonymity techniques, the use of booster scripts allows a user to specify lists of target URLs and identifying information for HOIC to cycle through as it generates its attack traffic. That, in turn, makes HOIC attacks slightly harder to block. HOIC continues to be used by Anonymous all over the world to launch DDoS attacks, although Anonymous attacks are not limited to those involving HOIC.

hping

In addition to LOIC and HOIC, Anonymous and other hacking groups and individuals have employed a variety of tools to launch DDoS attacks, especially due to the Ion Cannons’ lack of anonymity. One such tool, hping, is a fairly basic command line utility similar to the ping utility. However, it offers more functionality than simply sending an ICMP echo request that is the traditional use of ping. Hping can be used to send large volumes of TCP traffic at a target while spoofing the source IP addresses, making it appear to be random or even to originate from a specific, user-defined source. As a powerful, well-rounded tool (possessing some spoofing capabilities), hping remains among the tools of choice for Anonymous.

Slowloris

Besides straightforward, brute-force flood attacks, many of the more intricate “low and slow” attack types have been wrapped up into easy-to-use tools, yielding denial-of-service attacks that are much harder to detect. Slowloris, a tool developed by a gray hat hacker who goes by the handle “RSnake,” is able to create a denial-of-service condition for a server by using a very slow HTTP request. By sending HTTP headers to the target site in tiny chunks as slow as possible (waiting to send the next tiny chunk until just before the server would time out the request), the server is forced to continue to wait for the headers to arrive. If enough connections are opened to the server in this fashion, it is quickly unable to handle legitimate requests.

R U Dead Yet? (R.U.D.Y.)

Another slow-rate denial-of-service tool similar to Slowloris is R U Dead Yet? (R.U.D.Y.). Named after a Children of Bodom album, R.U.D.Y. achieves denial of service by using long-form field HTTP POST submissions rather than HTTP headers, as Slowloris does.

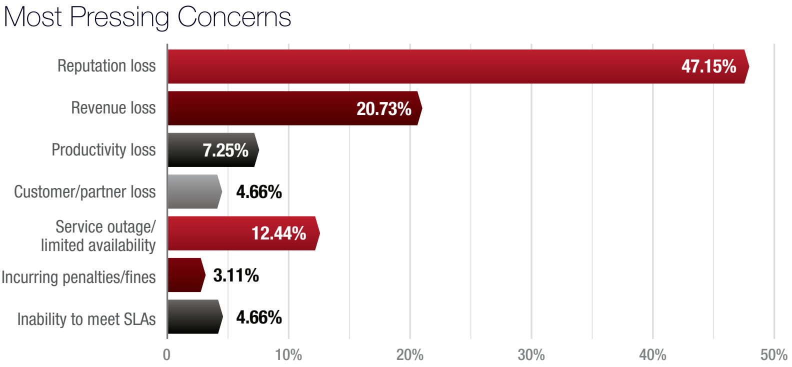

Figure 6: Business concerns due to cyber-attacks - Radware's 2014-2015 Global Application & Network Security Report.

By injecting one byte of information into an application POST field at a time and then waiting, R.U.D.Y. causes application threads to await the end of never-ending posts in order to perform processing (this behavior is necessary in order to allow Webservers to support users with slower connections). Since R.U.D.Y. causes the target Webserver to hang while waiting for the rest of an HTTP POST request, a user is able to create many simultaneous connections to the server—ultimately exhausting the server's connection table and causing a denial-of-service condition.

#RefRef

While all the aforementioned tools are non-vulnerability-based, #RefRef, another tool in Anonymous' arsenal, is based on vulnerability in the widely used SQL database software allowing for an injection attack. Using a SQL injection, #RefRef allows an attacker to cause a denial-of-service condition for a target server by forcing it to use a special SQL function (which allows for the repeated execution of any other SQL expression). This constant execution of a few lines of code consumes the target servers' resources, resulting in denial of service.

Unlike LOIC or HOIC, #RefRef does not require a vast number of machines to take down a server due to the nature of its attack vector. If the server's backend uses SQL and is vulnerable, only a few machines are needed to cause significant outage. While developing the tool, Anonymous tested it on various sites, easily causing outages for minutes at a time, and requiring only 10 to 20 seconds of a single machine running #RefRef. In one such attack (on Pastebin), a 17-second attack from a single machine was able to take the site offline for 42 minutes.

Botnets as a DDoS Tool

Regardless of the attack tool used, the ability to launch an attack from multiple computers—whether it is hundreds, thousands or millions—significantly amplifies the potential of an attack to cause denial of service. Attackers often have "botnets" at their disposal. Botnets are large collections of compromised computers, often referred to as "zombies," that are infected with malware allowing an attacker to control them. Botnet owners, or "herders," are able to control the machines in the botnet by means of a covert channel, such as IRC, issuing commands to perform malicious activities. Such activities may include distributed denial-of-service (DDoS) attacks, distribution of spam mail and information theft.

As of 2006, the average size of a botnet was around 20,000 machines, as botnet owners attempted to scale down networks to avoid detection. However, some larger, more advanced botnets—BredoLab, Conficker, TDL-4 and Zeus, for example—have been estimated to contain millions of machines. Large botnets can often be rented out by anyone willing to pay as little as $100 per day to use them. (One particular online forum ad offered the use of a botnet containing 80,000 to 120,000 infected hosts for $200 per day.) That accessibility enables anyone with only moderate technical knowledge and the right tools to launch a devastating attack. With this in mind, it is important to be aware of all recent attack tools, maintain up-to-date software on all servers and other network devices, and use some kind of in-house DDoS mitigation solution to protect against attacks as they continue to evolve.

6 Enterprise Security: Then and Now

Until recently, everything enterprises needed to protect—data centers, applications, databases—was nestled inside the perimeter. The basic rule? Secure an organization's perimeter and its assets are safe.

Today, the perimeter walls no longer exist, as enterprise applications move to the cloud. In short, assets are everywhere. How can an organization protect all enterprise assets—no matter where they reside?

In many organizations, the IT infrastructure resembles Figure 7. Data centers operate in multiple locations, while a growing portion of the infrastructure lives in the cloud. Dispersing the IT infrastructure introduces as many challenges as benefits. With the safe borders of the perimeter no longer protecting all enterprise assets, existing security measures need to be re-evaluated.

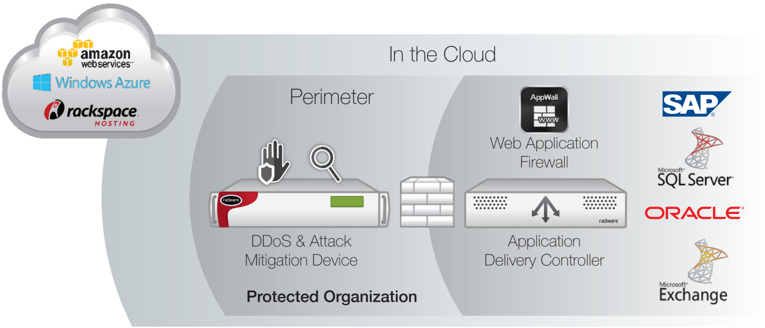

Figure 7: IT Infrastructure

While enterprise security is evolving, so are cybercriminals and attacks. Attackers and tactics are becoming increasingly sophisticated. It has become common knowledge that there is no way to prevent attacks—but there is a very strong need to mitigate them.

If an organization's security strategy does not take all of that into consideration, the organization and its users are at risk.

This chapter explores the challenges of protecting a dispersed IT infrastructure when it's unknown what part of it is going to be attacked or how such attacks will affect various assets. Some key considerations are outlined to address when revising a security strategy to reflect today's realities.

New Realities, New Challenges

Recent developments in information technologies and user mobility have transformed the IT infrastructure into a strong enabler for business agility and efficiency—while also introducing new security challenges to IT/security managers and enterprises who rely on Internet access for revenue generation and enterprise productivity:

- **The network perimeter is disappearing.** As enterprises have extended IT infrastructure to the public cloud, deploying new applications in the cloud or using it for disaster recovery, they now face the need to protect applications in the cloud as well as private data centers. This renders traditional security technologies inadequate and enterprises must build multiple skill sets and maintain a new set of management tools.

- **The Content Delivery Network market is expanding.** Content delivery network (CDN) solutions present new vulnerabilities, with hackers asking for dynamic content to overcome the powerful cache offloading mechanism that is the core of CDN solutions. Using this method, sophisticated attackers can build attack tools that go below the CDN radar and manage to saturate the application servers in the data center.

- **Data center virtualization is driving vulnerability to availability-based attacks.** Yes, private cloud technologies protect the confidentiality and integrity of application data. But they lack the ability to protect the physical infrastructure against availability-based attacks. Attacks targeting external applications impact the availability of internal critical applications—and an attack on a single application may endanger other applications on a shared infrastructure.

Layered IT Infrastructure Requires Layered Security Strategy

Traditional network and application security solutions typically combine detection and mitigation in the same system. The system operator sets rules (policies or profiles) and the system blocks (or allows) traffic that matches the pre-defined rules. In some cases, such as intrusion detection systems, the system will alert only upon suspicious traffic, and the operator is required to process the information manually. Traditional security solutions are also referred to as “point security solutions” as they prevent attacks inspected at their physical location. Subsequently, they have limited ability to mitigate sophisticated attacks, which require an overall network context awareness.

According to Gartner,[1] enterprises are overly dependent on blocking and prevention mechanisms that are increasingly ineffective against advanced attacks. Recognizing that security tools typically act as “islands of knowledge,” attackers are launching complex attack campaigns that exploit the lack of integration. Even organizations that invest in security information and event management (SIEM) solutions often become overwhelmed by the data generated from each tool. That, in turn, simply distracts operators, which can further prolong attack mitigation.

Organizations also find out that the ability to apply preventive measures is limited by the increasing complexity of security solutions. Often, they lack the product expertise required to select the right tool and location to apply the new rules.

The Age of the Integrated Hybrid Solution

Although it sounds like an oxymoron, the only holistic solution is a distributed solution. To fight complex attack campaigns and emerging threats, the distributed nature of a current IT infrastructure should be the core influence on the design of the security architecture.

In other words, if assets are dispersed in and accessed from multiple locations and devices, detection and mitigation tools should also be distributed. Detection coverage should be expanded to exist across all enterprise resources. More endpoints mean more types of detection tools, detection tools in different locations and, very possibly, tools from various vendors. Additionally, having multiple detection tools still requires a staff to manage and maintain them— and decide when, where and how to mitigate detected attacks.

Security Scenarios

Consider these scenarios, which correlate with some of the new security challenges organizations need to take into account when developing a security strategy:

1. Volumetric flood attacks are detected by an on-premise DoS/ DDoS protection device located at the perimeter. Once

[1] Designing an Adaptive Security Architecture for Protection From Advanced Attacks, Gartner, February 2014

detected, mitigation starts immediately. However, attack volume threatens to saturate the Internet pipe. Soon, the pipe has been saturated and the organization is losing business.

Is there another way? We see a growing number of DDoS mitigation solutions that provide a hybrid solution—mitigating the attack on-premise as long as there is no pipe saturation threat. Once such a threat appears, traffic is diverted to a cloud-based scrubbing center, with only clean traffic going back into the organization. The process is completely transparent, with no effect on user experience. (One caution: If detection and mitigation tools are from multiple vendors, the process may not necessarily be automated. Therefore, it may be time consuming and prone to human errors.)

2. Attacks based on true IPs masked by a CDN are resolved by the enterprise web application firewall (WAF). As the attack is mitigated close to the application (not at the perimeter), there is no guarantee that the attack has not reached other assets that are not protected by the WAF. In addition, if the attack volume increases, the WAF fails and the organization is vulnerable. The solution has to be scalable and therefore might be problematic if too many WAFs are implemented inline.

From Radware's perspective, the best possible solutions either have WAFs implemented out of path (making the solution scalable) or enable WAF to resolve the information (to be used in a network-wide context by configuring a blocking rule on the on-premise DoS protection device or, if attack volume increases – in the cloud). The result: an agile network that moves attack load from application devices, such as the WAF, to the perimeter or the cloud.

Having multiple detection and mitigation tools in different locations is inevitable. Operating, managing, maintaining and correlating them all—in "peace time" and especially when under attack—may seem like a mission impossible for an IT staff.

In an ongoing cat-and-mouse battle, attackers and security vendors both strive to be on the winning end. Already noted is the market trend toward hybrid solutions that combine on-premise detection and mitigation with cloud based scrubbing center mitigation. However, there is still a struggle with multiple detection tools, spread across various locations, very often, from different vendors.

Moving forward, we believe detection and mitigation solutions will evolve to become faster, more accurate and more automated—though they will still need to be managed and maintained. Thus, organizations will need a better understanding of the processes and improved visibility into attacks, before, during and after they occur.

From Location to Communication

What is the key factor for a successful security strategy that truly benefits from all the advantages of the advanced detection and mitigation solutions deployed in the organization? The answer seems to be a shift from location to communication. As the locations of detection and mitigation tools continually increase, the need for a coordinating mechanism—both human and machine based—grows, as well.

Distributing the detection and mitigation layers across all enterprise application infrastructures can deliver a global view of network behavior and the attacks state. Information collected from all detection tools needs to be correlated and analyzed to determine which mitigation process to use.

There is a real need for an automated central command and control system that manages all the tools by receiving ongoing information from all detection tools at all times—automatically controlling the mitigation process. Such a system would provide complete visibility and include sophisticated reporting features. Indeed, a single command and control center, receiving information from all detection tools (in peace time and under attack) will automatically choose the best mitigation process. This mastermind command and control center is constantly maintained and synchronized with legitimate traffic baselines and attack information in real time.

Why would such a system create the ideal solution for fighting current and emerging threats?

- It expands the detection coverage across all enterprise resources, whether on premises or in remote data centers (DR sites, private clouds and, to some extent, public clouds).
- It automates the mitigation by selecting the most effective tools and locations—whether in the data center, at the perimeter, at a scrubbing center or in the cloud.
- It offers unprecedented protection against current and future availability-based threats on all fronts.

Today's attack campaigns are complex. An organization can only protect against what it can detect. "Detect where you can, mitigate where you should" is the new mantra—and the best approach to stay ahead of attackers in perpetual cat-and-mouse chase.

7 What Lies Ahead: Predictions for 2015 and Beyond

As security professionals, many of us speak passionately about attack vectors, cyber-incidents or trends in information security. Just as often, we are asked to share our opinions and predictions. In reflecting back on 2014 — and looking to 2015 — we have five key predictions.

Prediction #1

Cyber Attacks Leading to Loss of Life.

For years, we've seen demonstrations of how attacks on all sorts of things—pacemakers, trains, automobiles and even aircraft systems—could one day lead to loss of life. Today, there's no doubt that cyber-attacks can and will turn deadly. It's no longer a question of "if" but "when."

Prediction #2

Rise in Cyber Ransoming & Hostage-Taking.

While there is a long history of cyber ransom activity, 2014 brought a new level of threat in criminal attacks. Nefarious groups have begun taking digital assets or services hostage—commandeering these resources until certain demands, which may or may not be financial, are met. In at least one case, this hostage-taking has led to business failure.

Prediction #3

More Critical Infrastructure Outages.

It's not hard to imagine how widespread cyber-attack disruptions could cripple a nation's critical infrastructure services—including power generation, water supply, cellular, telephone or television delivery services, or even police and first-responder networks. Even the world's most developed nations are not immune to this.

Prediction #4

Mass Adoption of Cyber-Attack Laws, Including Nationalistic Rules.

We believe that as governments face an increasingly dissatisfied, frustrated constituency—as well as growing threats around state-sponsored espionage—legislators will begin the process of writing laws on cyber-attacks. Such laws will likely aim to dictate network traffic flows, security levels at critical infrastructure companies and acceptable data processing domiciles. They will also provide guidelines on what constitutes acceptable Internet behavior.

Prediction #5

Reduced Sense of Urgency by Enterprise Managers.

Even as media reports and public awareness are at all-time highs, a certain sense of apathy or fatigue seems to have settled in among security decision makers. Perhaps many have grown disheartened and numb, believing that in the face of persistent attackers, a sense of urgency and doing the right thing will ultimately prove futile. We fear that business executives are increasingly abandoning rigorous exploration of how to secure endpoints and other points more effectively. We suspect that such execs are succumbing to the idea that becoming a victim—if they haven't already—is simply a foregone conclusion.

8 DDoS Mitigation Considerations

This section takes into consideration business and attack trends and provides a set of best practices for organizations to consider when planning for cyber-attacks.

Choosing a vendor. It is crucial to verify the vendor's experience and reputation. Is their technology market proven? Who are their clients and do they have MSSPs clients? Have their clients made it to the headlines due to being attacked? In addition, it is highly recommended to evaluate a single vendor that is able to provide a comprehensive detection and mitigation solution.

Attack coverage. Emerging threats bring with them new attack vectors. It is important to make sure that known attack vectors are mitigated by the offered solution and protection against SSL encryption attacks and various web-stealth attacks is included. Be certain to verify that the solution is a hybrid one in order to effectively handle pipe saturation risks with no disturbance to user-experience. Ensure the solution provides layered protecting covering attacks on network, servers and applications.

Real-time and post attack analysis. Visibility is critical in layered security architecture. Having a Security Information and Event Management system

Compliance Considerations

Targeting everything from financial services to power generation, cyber-attacks now threaten the fidelity and integrity of numerous industrial segments. As cyber-attacks have morphed into an existential threat to many countries, regulators have taken note. Among the most noteworthy initiatives:

- National Institute of Standards and Technology's (NIST) Cybersecurity Framework (US)
- Office of the Superintendent of Financial Institutions (OFSI) DDoS Memorandum (Canada)
- FFIEC Joint Statement Distributed Denial-of-Service (DDoS) Cyber-Attacks, Risk Mitigation, and Additional Resources (US)
- Securities and Exchange Commission Cyber Exams (US)
- Office of the Comptroller of the Currency Guidance (US)
- National Credit Union Administration Risk Alert (US)

(SIEM) integrated as part of a DDoS protection solution is extremely important. The fact that the IT staff can have full visibility and receive information in real-time, from all detection tools protecting the enterprise assets, is crucial. Advanced anti-DDoS solutions must be well integrated with SIEM systems that are able to aggregate, normalize, and correlate data from multiple sources. Real-time information, reports, automated analysis and processes provide visibility and insight during attacks and for post attack analysis and forensics.

Support under attack. It is important to verify in advance the vendor assistance offered when under attack. There are vendors that offer a team of experts to support clients under attack. Be sure this assistance lasts throughout the whole attack campaign and the team provides post attack analysis. Some vendors keep a team of researchers who provide periodic updates on the market and the new threats.

DDoS Do’s and Don'ts

Before an Attack - What to Consider Before Choosing a DDoS Protection Solution

Do's	Don'ts
1. Understand no organization is safe. It’s not about if you will be attacked, but about when. 2. Make sure detection tools are optimally located. Remember, you can only protect against what you can detect. 3. Make sure your security strategy is implemented into policies and procedures and that your staff is prepared with specifically defined roles and responsibilities. 4. Perform on-going tests and evaluations of your systems and of new technologies that are available in the market. For example: a. Verify whether your organization could benefit more from an out-of-path implementation of some of your detection tools. b. Evaluate the implementation of a hybrid solution to protect your organization during attacks that saturate the internet pipe. 5. Make sure your staff knows the DDoS do’s and don’ts and have an available easy-to-locate list of people to contact when under attack. If you are at risk of having a public website down, prepare an explanation and apology for an inconvenience message.	1. Don’t implement a solution just for compliance purposes. Understand your security risks and needs. 2. Don’t implement multiple detection tools from different vendors, unless these different tools are able to “communicate” with one another and pass relevant information for optimal detection.

During an Attack - How to Minimize Damage and Interference to Business

Do's	Don'ts
1. Contact the in-house and/or vendor’s Emergency Response Team to make sure best decisions are carried out. If you depend on an ISP vendor, contact them now.	1. Don’t panic. Manage it.
2. Define the detection point, attack type and tool, and decide on best mitigation process.	2. Don’t decide what to do before consulting your in-house/provider’s emergency response team.
3. Make sure every step of the attack is documented.	3. Don’t transfer traffic to the cloud scrubbing center unless you are close to pipe saturation.
4. Have a spokesperson ready to provide information to your customers during the attack (blog post, twitter, reporters).	4. Don’t ignore customers and make sure someone reassures them even during the attack.

After an Attack - What You Can Learn from the Attack and How to Prevent it From Reccurring

Do's	Don'ts
1. Perform a damage control analysis and review reports and forensics, learn what went wrong so you can better prepare for future attacks. Investigate everything.	1. Don’t think for one second that when the attack is over you can sit back and relax.
2. Optimize your security architecture. Make sure you analyze and evaluate every aspect of the attack. Adapt technologies, policies and solution strategies.	2. Don’t ignore your customers and press inquiries, address them and manage the crisis.
3. Notify customers/press with relevant details. Online businesses should consider a marketing campaign to win back the hearts of disappointed customers.	3. Don’t delay implementing the outcomes of the attack investigation, be it security strategy, technology solutions, policies, roles and responsibilities, and more.
4. Make sure your reports and forensics information is available in case it is needed for law enforcement investigation.	

Summary of Best Practices

When planning cyber-attack defense, be mindful of the C.H.E.W. threats, be demanding of vendors, and always consider the following tenets:

Timing is everything.
Organizations need to look at time to mitigate as a key success factor. With that in mind, ensure that the solution deployed provides the shortest time to mitigate.

Fill in the holes.
DDoS mitigation solutions need to offer wide attack coverage that can detect not just one attack vector, but also multi-vector attacks that hit different layers of the infrastructure.

Use multiple layers.
Resolve the issues of single-point solutions with cloud-based protection that blocks volumetric attacks plus an on-premise solution that blocks all other, non-volumetric attacks.

Mitigate SSL attacks.
SSL attacks remain a major threat. Look for SSL-based DoS/ DDoS mitigation solutions with a deployment that does not affect legitimate traffic performance.

Look for a single point of contact.
In the event of an attack, it's crucial to have a single point of contact that can help divert Internet traffic and deploy mitigation solutions.

9 Checklist: How to Evaluate a Vendor for DDoS & Cyber-Attack Mitigation

When evaluating a vendor for DDoS and cyber-attack mitigation, examine capabilities and strengths in two core competencies: detection and mitigation. Assess each vendor against these criteria—aiming to maximize capabilities in each of these areas.

How good is the vendor at detection?

Quality – This section evaluates the ability for the vendor to provide high-quality detection:

Type(s) of Detection Available
- Netflow
- Packet L7 Headerless
- Openflow
- Coverage of OWASP Vulnerabilities
- Packet L3/4
- Inputs/Signals from Other Mitigation Tools
- Packet L7 Header Required

Deployment Model Options
- In-Line
- Cloud Scrubbing Center – Asynschronous
- OOP – Synchronous
- Software Defined Networking (SDN)
- Hybrid Cloud Options
- Virtual Deployment Options
- Internal Scrubbing Center – Asynschronous
- Feeds from Partners/Works with Other Vendors' Signals

Time – This section evaluates the categories required for modern attack detection:
- Real-Time Options
- Signaling/Automatic Options (for Advanced Application Attacks)
- Signaling/Automatic Options (for Cloud Diversion)

Reporting & Response – This section evaluates the categories required for controlling and reporting modern attack detection:
- Real Time
- Detection Support Response – Real Time
- Historical

- Detection Support Response – On-Site Options
- Forensics
- Integrated Reporting with Cloud Portal
- Intelligence Reporting
- Ability to Discern Legitimate vs. (that is, can detect before attack) Illegitimate Traffic in Real Time

How good is the vendor at mitigation?

Quality – Does the vendor over-mitigate or under-mitigate the threats? How many technologies are leveraged to assist?

- Rate-Only
- HTTP Server-Based Protections
- Routing Techniques
- HTTP OWASP-Based Protections
- Rate Behavior Only
- Hybrid Signaling/Cloud Scrubbing Center Coordination
- Other Than Rate Behavior
- SSL Protections
- Heuristic Behavior
- HTTP Redirects
- Statistical Behavior
- JavaScript Challenge & Response
- Signatures – Static with Update Service
- Cloud Challenge Response
- Signatures – Custom Real Time

Time – How quickly can the vendor begin mitigation?

- Real-Time Options
- Automatic Options

Reporting & Response – How granular is the reporting? Can a user see if legitimate traffic is being impeded by the mitigation technique?

- Real-Time Displays
- Displays All Attacking Vectors Granularly
- Historical Mitigation Effectiveness Measures
- Mitigation Response Attack-Back Options
- Forensics & Detail Reports
- Mitigation Support Response – Real Time
- Emergency Response Options
- Mitigation Support Response – On-Site Options
- Displays Legitimate & Illegitimate Traffic
- Integrated Reporting with Cloud Portal

10 DDoS Dictionary

This dictionary focuses on network and application security terms with many DDoS-related definitions.

Advanced Persistent Threats (APT)
Category of cyber-security threats that seek to penetrate a network and gradually exfiltrate confidential or sensitive data from the network. These attacks are generally part of an attack with espionage as its core motive, and are often associated with state-sponsored attacks.

Always On
Security service delivery model that provides continuous application of security controls to all traffic flows. In the case of DDoS protection, “always on” generally refers to all traffic being inspected for detection of DDoS attacks, either via on-premise devices in-line or local out-of-path, or constant routing of traffic through cloud-based scrubbing services.

Availability Attacks
Availability attacks target a service in order to make it unavailable. Volumetric attacks are the most common availability attacks. However, any attack that renders a service unavailable is considered an availability attack. Such attacks include brute-force attacks on login pages, SSL encryption attacks and other stealthy methods that eventually cause severe service degradation or downtime. One of the main security challenges enterprises and service providers face is how to remain available even when under attack.

Bot/Botnet
A group of many (often thousands) of volunteered or compromised computers that send a huge amount of traffic to an attack target, seeking to overwhelm its network.

Distributed Denial of Service (DDoS)
A distributed denial-of-service (DDoS) attack is one in which two or more persons, bots, or other compromised systems, attack a single target—causing the system to slow down or shut down, thereby denying its users the ability to use it. During DDoS attacks, an

online service can be brought down by overwhelming it with traffic from multiple sources. Radware research suggests that the most common industries to experience such attacks are government and federal agencies, ISPs and hosting service providers, financial institutions and the gaming industry.

Denial of Service (DoS)

A denial of service attack is an attempt to make a machine or network resource unavailable to its intended, legitimate users. Denial-of-service attacks can disable a computer or a network for minutes or for days. Depending on the nature of the attacker and attacked party, such an attack can effectively disable an organization.

DNS Flood

Attack that targets the DNS application protocol by sending a high volume of DNS requests. Domain Name System (DNS) is the protocol used to resolve domain names into IP addresses; its underlying protocol is UDP, taking advantage of fast request and response times without the overhead of having to establish connections (as TCP requires).

Forensics

DDoS data forensics and post-attack analysis are crucial for a number of reasons. In the midst of an attack, forensics analysis is used to identify the attacking party and safely distinguish attack traffic from legitimate traffic. It also enables more accurate selection of the best mitigation tools to stop the attack.

Once an attack has been successfully mitigated, forensics are critical to understanding the attack origin, motivation and attack types and tools—whether for legal reasons or to enhance future preparation. Forensics also serve as a research tool, yielding a better understanding of DDoS trends.

HOIC

Tool commonly used to launch DDoS attacks that can send HTTP POST and GET requests wrapped in an easy-to-use GUI. Its effectiveness stems from add-on “booster” scripts—text files that contain additional basic code interpreted by the main application upon a user’s launch of an attack.

HTTP Flood

Common form of attack that consists of what seems to be legitimate, session-based sets of HTTP GET or POST requests sent to a victim's Web server, making it hard to detect. HTTP flood attacks are typically launched simultaneously from multiple computers (volunteered machines or bots).

Hybrid Mitigation

Combination of on-premise and cloud-based mitigation technology that delivers immediate mitigation of non-volumetric attacks with the availability of additional mitigation resources in the event an attack threatens to saturate the Internet pipe of the attack victim.

IP Spoofing

Tactic of creating Internet Protocol (IP) packets with a false source IP address, thereby concealing the identity of the sender which complicates IP-based attack blocking and attacker attribution.

Layer 3 and Layer 4 Attacks

Broad category of attacks that target the Network (Layer 3) and Transport (Layer 4) layers of the OSI stack model. Common attack vectors for Layer 3 and 4 attacks include TCP-SYN floods, UDP floods, and ICMP attacks.

Layer 7 Attacks

Broad category of attacks that target the Application layer (Layer 7) of the OSI stack model. Common attack vectors for Layer 7 attacks include SMTP attacks, DNS floods, and HTTP/HTTPS attacks.

LOIC

Tool commonly used to launch DDoS attacks that can generate massive volume of TCP, UDP or HTTP traffic in order to subject a server to a heavy network load. LOIC's original developers intended the tool to be used by developers who wanted to subject their own servers to a heavy network traffic load for testing purposes.

Low and Slow Attacks

Attacks that target specific design flaws or vulnerabilities on a target server with a relatively small amount of malicious traffic, eventually causing it to crash. "Low and slow" attacks mostly target application resources (and sometimes server resources) and are difficult to detect because they involve connections and data transfer appearing to occur at a normal rate.

On Demand

Refers generally to the availability of DDoS scrubbing services being available as needed, generally when volumetric attacks threaten to saturate inbound link capacity.

Out of Path

Security service architecture where security devices or services do not sit in-line with constant flow of traffic. Typically, in out-of-path architectures, the security device or service is connected to another device or service in the data path that redirects traffic to the out-of-path device based on certain traffic profiles or patterns. Out-of-path deployments reduce potential points of failure in the normal network traffic flow, but also reduce the ability of the security device or service to provide optimized service delivery.

Pipe Saturation

Internet pipe saturation can occur during attacks creating volumetric floods, which are often intended to flood the target by overwhelming bandwidth. Common attacks use UDP because it is easily spoofed and difficult to mitigate downstream. Out of state, SYN floods and malformed packets are also often seen. While many attacks aim at saturating inbound bandwidth, it’s not uncommon for attackers to identify and pull large files from websites or FTP shares as any means of saturating outbound bandwidth.

Scrubbing Center

A scrubbing center is a centralized data cleansing station where traffic is analyzed and malicious traffic is removed. Scrubbing centers are often used by large enterprises, such as ISP and cloud providers, as they often prefer to off-ramp traffic to an out-of-path, centralized data cleansing station.

When under attack, traffic is redirected (typically using DNS or BGP) to the scrubbing center. There, an attack mitigation system mitigates the attack traffic and passes clean traffic back to the network for delivery.

A scrubbing center should be equipped to sustain high volumetric floods at the network and application layers, low and slow attacks, RFC Compliance checks, known vulnerabilities and zero day anomalies.

Security Operations Center (SOC)/Emergency Response

A Security Operations Center (SOC) can be described as an enterprise IT “war room.” It is where a team of professionals continuously monitors, assesses and secures the enterprise data centers, servers, applications, networks, websites, endpoints and more.

DDoS attacks can last a number of hours or persist for days or weeks. Over such long, intense times, organizations look for a single point of contact to support the attack mitigation process: detecting the attack, applying the correct mitigation tools at the right time and when needed, and then diverting the traffic under attack to the cloud-based scrubbing center.

Be it an in-house SOC team or an external security vendor Emergency Response, an enterprise must have security professional services available 24/7 for hands-on attack mitigation assistance to successfully defend networks against cyber-attacks. Such professionals have the expertise required to fight prolonged, multi-vector attacks.

Service Degradation

Service degradation is a type of DoS/DDoS attack that disrupts a service by slowing the speed and response time of a network or website. At times, the attack is stopped at this stage; in other cases, the degradation is just the step before a service shutdown. Some hackers use service degradation attacks to evaluate the strength of the target they aim to disrupt before launching an actual attack.

Service Downtime/Shutdown

The term downtime is used to refer to periods when a system is unavailable—that is, when it fails to provide its primary function. A DDoS attack can cause a service shutdown, rendering the service unavailable. A service downtime can have severe financial consequences and in some cases even bring a business down. (Consider, for example, that in 2013 it was revealed that a five-minute outage costs Google $545,000 in revenue.)

SSL Based Attacks

Attacks that encrypt the malicious traffic to obfuscate its contents, bypassing certain detection methods. SSL attacks also consume greater computing capacity due to the need to decrypt and encrypt their contents.

Time to Mitigation

The longer an entity is under attack, the longer users suffer from unavailability and slow responses. This leads to frustration and dissatisfaction as well as a decrease in productivity. The time to detect, and more importantly, to mitigate is critical. Time to mitigate is a key decision factor for a DoS/DDoS mitigation solution. The sooner the mitigation starts, the sooner the organization's services resume.

Volumetric Attacks

Broad category of attacks that attempt to overwhelm the Internet pipe or other capacity limitations of the target. Volumetric attacks are challenging to protect against due to the need for significant bandwidth capacity to receive the traffic before scrubbing, and often require cloud-based scrubbing resources for mitigation.

Web Application Firewall (WAF)

Security product or service that applies a defined or dynamic set of security policies to transactions on a website. WAF's generally target common web attacks such as cross-site scripting (XSS) and SQL injection.

Web Stealth Attacks/Smokescreens

Web Stealth attacks are a set of vectors that include brute-force attacks (for example, attacks on the login page), file upload violations and SSL-encrypted application attacks, among others. These attack vectors are built on HTTP packets that conform to relevant Web traffic specifications, and thus cannot be detected by standard network security tools such as IPS, firewall and rate-limit-based DoS/DDoS protection tools.

Attackers use the evasive nature of HTTPS and other SSL-encrypted mechanisms as well as the asymmetric nature of these attacks to bypass network security mechanisms and attack servers deep inside the network topology. This is where they are most susceptible for resource saturation.

For More Information

Please visit www.radware.com for additional expert resources and information and our security center DDoSWarriors.com that provides a comprehensive analysis on DDoS attack tools, trends and threats.

Radware encourages you to join our community and follow us on: Facebook, Google+, LinkedIn, Radware Blog, SlideShare, Twitter, YouTube, Radware Connect app for iPhone®.

About the Authors

Radware (NASDAQ: RDWR), is a global leader of application delivery and application security solutions for virtual, cloud and software defined data centers. Its award-winning solutions portfolio delivers service level assurance for business-critical applications, while maximizing IT efficiency. Radware's solutions empower more than 10,000 enterprise and carrier customers worldwide to adapt to market challenges quickly, maintain business continuity and achieve maximum productivity while keeping costs down. For more information, please visit www.radware.com.

radware