the Availability Digest

www.availabilitydigest.com

Anatomy of a DDoS Attack

April 2013

Question: What are botnets used for? Answer: Distributed Denial of Service (DDoS) Attacks.

Botnets are bad. The DDoS attacks that they can launch are even worse. The damage DDoS attacks can do to a company's public-facing Internet services, such as web sites, or to the Internet in general is massive. The last few months have seen several examples of the use of botnets to bring major corporations to their knees:

- In September, 2012, in retaliation for the anti-Islamic YouTube video "Innocence of Muslims, Islamic hackers launched massive DDoS attacks against several U.S. banks that took down their online banking portals for over a day each. In December, they repeated their attacks; and they have vowed to continue the attacks until the video is removed from the Internet.[1] Their attacks so far have generated up to 70 gbps of malicious traffic – enough to overwhelm most web sites.

- In March of 2013, Spamhaus was hit with the most massive DDoS attack yet reported – a malicious data rate of 300 gbps! Spamhaus is a firm that maintains a blacklist of spam-generating sites and sells the list to corporations, government agencies, and ISPs so that they can block traffic from these sites. One of the web sites on the blacklist is CyberBunker. It is CyberBunker that is suspected of launching the attack.[2]

Until these large attacks, most DDoS attacks generated about 10 gigabit-per-second (gbps) of malicious traffic. Clearly, the severity of DDoS attacks is increasing. So is the frequency of attacks. Prolexic, a DDoS mitigation firm, found in its surveys that DDoS attacks increased 53% from 2011 to 2012.[3] Prolexic has mitigated seven attacks that exceeded 50 gbps.

The concept of DDoS attacks is simple. Generate enough malicious traffic to a web site, and it will be unable to respond to legitimate requests. In effect, the web site has been taken down. DDoS attacks have been used for retaliation, for political statements, for competitive reasons, and even for ransom.

Botnets

A single system is not powerful enough to generate enough traffic to overwhelm most systems. It takes a concerted effort of many systems to do so. This is a botnet. A botnet is a controlled collection of infected

[1] Islamic Hacktivists Attack U.S. Banks, *Availability Digest*; October 2012.
http://www.availabilitydigest.com/public_articles/0710/bank_attacks.pdf
DDoS Attacks on U.S. Banks Continue, *Availability Digest*; January 2013.
http://www.availabilitydigest.com/public_articles/0801/more_bank_attacks.pdf
[2] History's Largest DDoS Attack?, *Availability Digest*; April 2013.
http://www.availabilitydigest.com/public_articles/0804/spamhaus.pdf
[3] Prolexic Quarterly Global DDoS Attack Report, *Prolexic White Paper*; 2013.

systems that can be commanded to take a joint action upon request. For DDoS attacks, this joint action is the generation of massive amounts of malicious data directed toward a victim system.

There are several classes of botnets:

- The earliest botnets were made up of infected PCs. Typically, a PC would be infected by a Trojan that entered the PC via a malicious email, a malicious web site, or an infected web site. The Trojan would then allow the botnet master to download the DDoS software into the PC. PCs cannot generate a great deal of traffic, if nothing else due to the bandwidth of its Internet connection. A megabit per second (mbps) is typical. Therefore, to generate ten gigabytes per second of traffic, the botnet must comprise ten thousand PCs.

- Some attacks are politically popular and generate a great deal of support among a class of people around the world. In this case, attackers have enlisted many people to voluntarily contribute the services of their PCs to the botnet. The Islamist hackers that attacked the U.S. banks in retaliation for the anti-Islamic YouTube video reportedly had access to hundreds of thousands of voluntarily provided PCs. Another example was an attack launched by supporters of Julian Assange, founder of WikiLeaks, when he was arrested for leaking classified material.

- The limited capability of a PC to generate DDoS traffic is solved to a great extent by using powerful servers instead. In this case, servers are infected with DDoS software, often through known security vulnerabilities in popular programs such as Joomla and WordPress. A powerful server with wide-band access to the Internet can generate a thousand times as much traffic as a PC.

DDoS Strategies

Layers in the Internet Protocol Suite

There are several different strategies that DDoS attackers use to overwhelm a web site by attacking different layers in the Internet protocol suite. The Internet protocol suite is divided into four layers:

- The *application layer* contains the higher-level protocols used by most applications for network communication. Examples of application layer protocols include the File Transfer Protocol (FTP) and the Simple Mail Transfer Protocol (SMTP).

- The *transport layer* establishes host-to-host connectivity Its responsibility includes end-to-end message transfer independent of the underlying network. End-to-end message transmission can be categorized as either connection oriented (TCP) or connectionless (UDP).

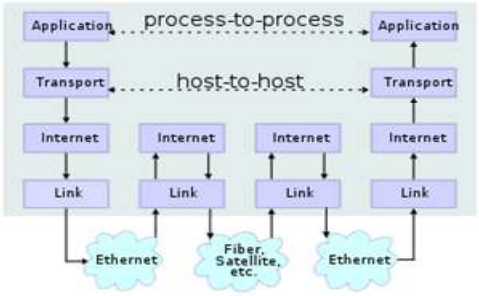

Wikipedia – Internet protocol suite

- The *internet layer* provides an unreliable, best-efforts datagram transmission facility between hosts located on different networks by forwarding the transport layer datagrams to an appropriate next-hop router for further relaying to its destination. The Internet Protocol (IP) is the protocol used at this layer.

- The *link layer* describes the protocols used to implement the local network topology needed to effect transmission of Internet-layer datagrams between Internet layers. It includes the physical routers, switches, hubs, firewalls, and other equipment required to implement local networks with protocols such as Ethernet and Fibre Channel.

Types of DDoS Attacks

DDoS attacks can be directed at any of these layers. In many cases, multiple layers are targeted in a single attack.

In general, there are three classes of attacks:

- Network level: The network is bombarded with traffic, consuming all the bandwidth needed by legitimate incoming requests. Such an attack can be as simple as sending massive numbers of pings to the web site or as sophisticated as an amplified DNS attack, in which UDP queries requiring large responses are sent to the web site's DNS server. In these cases, the source IP address can be spoofed; so there is no indication of the source of the requests.

- Infrastructure level: Network devices such as firewalls, routers, and switches are usually stateful. They maintain state in internal tables. If their state tables fill, they can handle no further traffic. DDoS attacks at the infrastructure level attempt to fill the state tables of network devices so that they are bogged down and cannot handle any more legitimate traffic.

- Application level: The web server applications are targeted. A simple example of this is a mass of login requests. The requests do not have to be successful – they just have to consume server resources as the login is processed and rejected. If an attacker has managed to obtain legitimate login names and passwords, it can do even more damage by invoking application functions such as searches that require a great deal of processing capacity.

Address Spoofing

In almost all successful DDoS attacks, the IP address if the sender is forged (spoofed) so that the locations and addresses of the attacking machines cannot be easily identified. The spoofed address can be randomly assigned for each malicious message to prevent filtering of the packets based on the source address.

UDP Flood (ICMP Flood)

ICMP is the Internet Control Message Protocol. It is used to send error messages indicating, for example, that a requested service is not available or that a host or router could not be reached.

A *UDP flood* (also known as an *ICMP Flood*) is a mass of UDP requests sent to a victim system that results in an equal mass of ICMP responses. For instance, requests may be sent to random ports on the victim machine. Most of these ports are likely not to be used, and the victim system will respond with a "destination port unreachable" ICMP message. Given a high enough rate of malicious messages, the target system will be so busy responding with ICMP messages that its ability to respond to legitimate traffic will be compromised.

Ping Attack

A *ping attack* is an ICMP attack in which the target system is flooded with ping requests (ICMP echo-request message). The target system must respond to each ping request with a ping-response message (ICMP echo-reply message), thus using up processing capacity.

Ping Smurf Attack

A ping attack can be launched from a source that has a limited bandwidth capability. This is an *amplified ping attack*. It depends upon the attacker finding systems in which the local networks are misconfigured so that a broadcast address is exposed. That network then serves as a *smurf amplifier.* A ping request

with a spoofed sending address is sent to the broadcast address. The ping request is broadcast to all systems on the local network, which respond with an echo-reply to the spoofed source address. The spoofed source address is that of the victim machine. Its bandwidth becomes overwhelmed with ping responses, preventing it from handling legitimate traffic.

A ping smurf attack is one form of a *reflected attack* in which a large number of computers respond to a forged request that has the victim's address as a spoofed address. The DNS Reflection attack described later is another example of this kind of attack.

SYN Flood

A SYN flood is a DDoS attack in which an attacker sends a succession of SYN requests to a target system in an attempt to consume enough server resources to render the system unresponsive to legitimate traffic.

A SYN message is used by a client to initiate a connection with a server host. Normally, the sever host will acknowledge the request by returning a SYN-ACK. The client responds with an ACK, and the connection is established.

In a SYN flood attack, the attacker sends multiple SYN requests to the victim server with spoofed source IP addresses. In this way, the attacker's ID is hidden and the attacker is not burdened with the servers' responses. A spoofed client will not return an ACK since it knows that it did not send a SYN.

The server will assign resources to the connection and will respond with a SYN-ACK. However, it will never get the ACK to complete the connection. It will eventually timeout after about three minutes and release the resources. The server eventually runs out of resources and is unable to handle any more connections, thereby denying service to legitimate users.

An advantage of a SYN flood is that it does not require a massive amount of malicious traffic to be generated. It can easily be carried out by a single machine. However, there are well-known and effective countermeasures to thwart a SYN attack.

GET/POST Floods

GET and POST are commands sent by an HTTP client (typically, a web browser) to an HTTP server requesting services. GET commands are used to retrieve data from a server (for instance, a picture), and POST commands are used to update data on the server (for instance, a form). Each of these commands use processing and disk resources of the server. Given a command rate sufficiently large, the server's processing resources will be used up, and the server will become unresponsive to legitimate traffic.

GET commands are easier to generate. It is only necessary to send a URL for a publicly available picture. POST commands require some specific knowledge of a form available on the server. However, a GET consumes far fewer processing resources, and the request rate for GETs has to be much higher than that for POSTs.

An attacker will often look for long responses to GET or POST commands, as that indicates that the server is doing a great deal of processing in order to satisfy the request.

DNS Reflection

DNS reflection is a particularly vicious type of attack since it can generate massive amounts of malicious data sent to the victim with only a relatively small amount of effort on the part of the attacker. DNS reflection depends upon DNS open resolvers.

An open resolver is a DNS server that will respond to a request from any source. The attacker sends DNS requests to resolve an address to multiple open resolvers with the spoofed source IP address of the victim. The DNS system will return its response to the victim, thus potentially overwhelming it. The reason that this is a particularly vicious attack is that the DNS response message is about 100 times as big (about 3,000 bytes) as the request message (about 30 bytes). Therefore, the attacker can generate a massive amount of malicious traffic by generating only 1% of that traffic.

DNS reflection was used in the recent attack on Spamhaus, referenced earlier, to generate a DDoS attack of 300 gbps. Years ago, DNS open resolvers were recognized as a security issue; and efforts were made to close all DNS resolvers. Unfortunately, network administrators are not particularly incentivized to update their DNS servers. There are still about 27 million open resolvers on the Internet, and their IP addresses are known.

Unintentional DDoS Attacks

DDoS attacks can be launched inadvertently. When Michael Jackson died, web sites such as Google and Twitter slowed down or crashed due to unanticipated traffic. This problem occurs sometimes when a URL is mentioned on TV and becomes popular.

Another cause of unintentional DDoS attacks is a focused attack gone wrong. In August, 2009, the operator of a game server in China attempted to shut down a group of competitors by launching a DDoS attack on a DNS server used by his competitors. Unfortunately, this was a major DNS server in China's network. As it became overloaded, it sent DNS requests to other DNS servers, which in turn overloaded. Much of China's Internet service was down for hours.

Legal action has been taken in at least one case. In 2006, Universal Tube and Rollform Equipment Corporation sued YouTube. A massive number of potential YouTube users accidentally typed the tube company's URL address, utube.com. As a result, the company had to spend a large amount of money to upgrade its bandwidth.

Legality

DDoS attacks are specifically outlawed by many countries. Violators in the U.K. can get up to ten years in prison for carrying out these attacks. The U.S. has similar penalties, as do most major countries. However, there are many countries from which DDoS attacks can be launched without penalty.

In January, 2013, the well-known hacker group Anonymous filed a petition in the U.S. courts to acknowledge DDoS as a legal form of protest similar to the Occupy protests.

Mitigation Strategies

Some protection from DDoS attacks can be provided by firewalls and intrusion-prevention (IPV) systems. However, these devises can be overwhelmed by very large or complex attacks if the volume of malicious traffic exceeds their capacity or if the sophistication of the attack is beyond the capabilities of their rule sets.

The next step is to use the services of a DDoS mitigation company with large data centers that can spread the attack volume over multiple data centers and scrub the traffic to separate bad traffic from

legitimate traffic. Prolexic, Tata Communications, AT&T, and Verisign are examples of providers of this service.

Summary

DDoS attacks are here to stay. They are motivated by too many factors – retaliation, political statements, attacking competitors, ransom –and are fairly easy to launch. There are even sophisticated tools that are publicly available to launch significant massive attacks (*itsoknoproblembro* is a popular toolkit for DNS reflection attacks). The defenses against DDoS attacks are at best limited.

Companies must prepare for the likelihood of losing their public-facing web services and must make plans for how they will continue in operation if these services are taken down. This should be a major topic in their Business Continuity Plans.