Distributed Denial of Service Attacks

The Biggest Threat on the Internet

By: Craig Kitching

For: BA559,

Professor Shaw

Viruses, spyware, and all sorts of other internet malware exist in cyberspace to try to turn a computer into a bedding ground for illegal activity. With all of the malicious content out there it is easy to get overwhelmed, it is easy for some individuals and companies to lose sight of exactly how powerful some of these threats can be. In this report, I will discuss the largest threat on the internet today, the Distributed Denial of Service attack. Specifically, I will be discussing this form of attack with a focus on 'BotNets,' which are the main tool producing much of the malicious and viral content seen on the internet today. The prevention of such attacks, and the mediation once attacks occur, can be the difference between the life and death of a real life organization. The proliferation and magnitude of cyber attacks can have affects reaching much further out than simply the internet, and understanding how truly devastating a network lurks out there can help a corporation or even an individual be prepared.

Denial of Service. Vs. Distributed Denial of Service

It is important to understand the difference of a denial of service (DoS) attack and a distributed denial of service (DDoS) attack. Wikipedia classifies the attacks as: "A denial-of-service attack (DoS attack) or distributed denial-of-service attack (DDoS attack) is an attempt to make a computer resource unavailable to its intended users. Although the means to, motives for, and targets of a DoS attack may vary, it generally consists of the concerted, malevolent efforts of a person or persons to prevent an Internet site or service from functioning efficiently or at all, temporarily or indefinitely.” The main difference between a DoS and a DDoS then is really the magnitude.

Imagine, for instance, if one computer sent a request out to a web server for information. This request is relatively small, and could be sent thousands of times a minute. This would cause additional traffic on the side of the server, and might slow down the service for other users. With today’s advanced technology however, and amounts of bandwidth, it is unlikely that no matter how many thousands of requests that computer made it would truly impact the performance of an organization.

Distributed Denial of Service

Now imagine the same computer in the above example was not one computer, but five thousand. Five thousand different computers making the same requests of a site, thousands of times each per minute, now the site has a serious problem on their hands. The problem with the structure of DoS attacks is that they disguise themselves as legitimate internet traffic. The server has to respond to the traffic because, at the initiation of the attack, it is not sure whether this is a legitimate user or some type of malicious hacker. If it were to simply turn off that type of traffic, it would essentially be rejecting all of its actual customers that are trying to use its web services.

Because of this, DDoS attacks can be hard to see coming, and can be even harder to stop depending on how creative the attacker can be. It is, at its core, nothing more than a brute force attack on a web server. The web server has to do something with all of the information coming at it, and even if it discards the attacking packets of information, it will still be slowed down significantly.

DDoS and DoS come in a variety of methods, and have been broken down into the following categories:

ICMP Floods

ICMP floods, also colloquially called ‘smurf attacks’, are used by abusing misconfigured computer networks. The network must be configured in such a way that the host machine, or router, does not distinguish its local machines by individual network address. That means that any traffic sent to a particular port on its broadcast address (the IP address that it exposes to the internet) is mirrored back to every computer on its internal network. A typical router will not do this, and distinguishes requests based on its own routing tables, but misconfigured networks can produce the above scenarios. What this leads to is any request being sent to the broadcast, wide area network (WAN), IP will be mimicked by every machine on its Local Area Network (LAN).

Therefore, if an outside machine sends a request with a spoofed (fake) source IP address, all of the machines inside of the LAN network will respond to the request. As an example, assume a hacker wanted to go after Google (not a very realistic target). The hacker would send a request to this misconfigured network for a packet of information, but the hacker would tell the network that he was actually google.com. Once the network receives the request, it reflects it to every computer inside of its LAN, and all of those computers send the requests packet back to google.com. While Google has a large enough internet presence it would not be an easy target for an ICMP flood, it serves as a good example for how such an attack can get very large, very fast.

Teardrop Attacks

Teardrop attacks are very specific attacks that only target older operating systems (Windows 3.1, 95, NT and Linux prior to 2.1.63). It exploits a vulnerability in the way it handles packets, or requests, from servers. Basically, a mangled overlapping request can be sent to the host machine, which will cause it to try and respond in an incomprehensible manner and then subsequently crash. This is not a common form of attack today as these operating systems become phased out.

Peer-to-Peer Attacks

Peer-to-Peer attacks take advantage of an early vulnerability in some peer-to-peer (p2p) sharing software like Kazaa or E-Donkey. The way that the attacks worked was to trick the peer-to-peer clients into trying to connect to the victim's computer instead of the peer-to-peer server. The connection would of course not work, but with the hundreds of thousands of users on p2p networks attacks over 500,000 connections strong were made. Many servers will crash with only a few thousand attempted connections a second, so this type of attack would make short work of nearly any network. This form of attack is also dying out as p2p networks fall out of favor and also become more secure.

Application Level Floods

Application level floods come in many different shapes and sizes, and are the most common type of DDoS today. At its core, application level floods either exploit a buffer overload error in a particular application, allowing an attacker to fill the victim's hard drive with faulty data, or involve simply overpowering a connections bandwidth. The brute force method, the later of the two mentioned, is achieved when the attacker has more bandwidth then the victim, and requests more information than the victim can accommodate, slowing or terminating their connection. On a large scale, this type of attack is done today with botnets, which will be discussed later.

Nukes

Another older form of DoS attacks, Nukes were a simpler version of the Teardrop attacks described above. It was achieved by sending out of sequence and fragmented packets to a host machine causing it to severely slow down or crash. This form of attack was used primarily in games and instant messaging and could be controlled by just one machine. Graphical user interfaces (GUIs) were developed to assist in the 'nuking' of other clients. Nearly all instant messaging and games servers now have a filtering system that makes this type of attack no longer work.

Reflected Attacks

This type of DDoS attack involves the same mechanisms as an application level flood, but in reverse. A large number of requests are sent out to thousands of machines, and the source IP on the request is set to the victim's IP address. This means that thousands of machines will now try to reply to this victim’s computer, slowing or crashing its network interface.

Unintentional Attacks

Unintentional attacks are not really an attack at all, but simply an overwhelming unforeseen amount of traffic that was not intended to take down a web server, but does. These types of attacks occur frequently when smaller web sites become national superstars via the media, and then receive millions of additional web requests. The sites were not built to sustain such traffic and as a result will frequently be shut down. This has the same footprint of a DDoS, but was completely unintentional. While press is usually seen as a good thing, the additional bandwidth costs and hardware costs that can come with an unintentional DoS attack can be substantial and could really be cause for concern for smaller company websites.

The above list is only a small sample of the variety and magnitude of DoS attacks, but will serve as a great place to get started as we delve deeper into the underworld of the internet, and what it means to those of us at the surface.

What is a Botnet?

So how is it exactly that a simple packet request from one computer to another can cause all of this catastrophe? As discussed above, magnitude is the key to any DDoS, without it an attack is simply any ordinary request from a server. While network vulnerabilities years ago made it possible for one individual computer to crash another, networks have been becoming much more sophisticated over the years and it is very unlikely that there are many exploits that exist today that can simply crash a computer via the internet packets it receives. This means that as far as denial of service attacks go, attackers are greatly limited in the types of attacks they can use. However, the tried and true brute force method still reigns supreme. But in order to successfully execute such an attack, thousands upon thousands of machines must be employed. This is where a botnet comes into service.

Botnets are large groups of computers tied together by one common thread, a malicious virus that is installed in them. These viruses simply sit and wait, until they receive a command to attack or initiate another form of action. These viruses can affect computers in a variety of ways, but most commonly it is through email. Attachments are disguised with executable files attached, such as an imagine might display a .jpg extension but in reality it is a .jpg.exe executable. Once the victim launches the picture, which may even display a real graphic so the user is ignorant to the installation occurring, a program is installed on their machine. The machine is now labeled as a ‘zombie’. It has no direction or initiative on its own, but instead it is one of many that is linked to a central controller.

Methods of Control

The most common form of control is Internet Relay Chat (IRC). The infected computer will connect to an IRC server and join a specified IRC channel, all set up within the virus. Once in that channel, which is very similar to a normal chat channel in any instant messaging program, it will broadcast its IP address and what version of the vulnerability it is currently running. This entire procedure happens in the background of the victims operating system, the victim never knows that it is connected to any IRC channel or that it is exposed on the internet. Eventually, an administrator of the bots will join the channel and send out a command. This command will be echoed across all bots on the current server. This command can be malicious in nature, such as initiating a DDoS attack, but it can also be used to update the bot software.

The nature of the zombie machines is such that the virus must stay under the radar and not be noticed by the user. This means that it must frequently change its name and its code to not only update abilities, but also be able to avoid detection from anti virus and anti spyware scanners. Older bots were subject to a number of anti virus mechanisms and can be caught easily with today's software, however newer bots have been developed that are much more evolutionary and are very difficult to detect by today's software.

Coming of the Storm

Botnets have been around for over a decade, but have really only come to mainstream attention a little over a year ago when the media reported the largest botnet known at the time, Storm. Storm was estimated, at its largest size, to be over 1.5 million computers strong. It was installed on PC's via the 'Storm Worm', a Trojan virus that was installed through email spam, as explained above. The storm worm at one time comprised "8% of all malware installed on Microsoft Computers" (Wikipedia). The Storm botnet was said to be large enough to take entire countries off of the internet grid, and contained the computing power of some of the top supercomputers in the world. Bradley Anstis, of the United Kingdom security firm Marshal, said, "The more worrying thing is bandwidth. Just calculate four million times a standard ADSL connection. That's a lot of bandwidth. It's quite worrying. Having resources like that at their disposal—distributed around the world with a high presence and in a lot of countries—means they can deliver very effective distributed attacks against hosts." (Tung)

The botnet has a variety of functions, but one of its primary functions is to keep itself alive. It does this by proliferating, sending copies of itself out through malicious emails from the host’s computer. It was estimated that at least 6,000 of the computer on the botnet were dedicated simply to sending out copies of the virus. While not a large percentage of the entire botnet, it was a huge amount of spam email. A record 57 million virus containing email messages was sent on August 22nd, 2007 alone (Wikipedia). So what was the primary function of the rest of the botnet? The primary function was making money.

How to Sell a Botnet

Have you ever wondered who sends those thousands of spam emails users receive on a daily basis? Storm had a variety of uses for prospective buyers, but this was certainly one of the largest ones. Infected computers could be used to send out thousands of spam emails a day, millions in a month. Over 1.2 billion virus containing emails alone had been sent out within its first year of operations. This does not even include your typical spam email messages asking for money or other more sultry services. Hundreds of billions of spam email messages are sent out every year, and many of them come from Storm. On top of all of that, it is estimated that only 10-20% of the strength of the network is actually even being used.

According to Matt Sergeant, chief anti-spam technologist at MessageLabs, "In terms of power, [the botnet] utterly blows the supercomputers away. If you add up all 500 of the top supercomputers, it blows them all away with just 2 million of its machines. It's very frightening that criminals have access to that much computing power, but there's not much we can do about it." (Gaudin) This computer power can also be harnessed to attack competitors, slowing down their web servers to a crawl. In fact, it can even be used to destroy companies. A DDoS attack, for instance, costs roughly a nickel a bot on a 10,000-bot network – about $500 for a fairly effective assault. (Wired)

The Tale of Blue Security

Blue Security was an anti-spam company out of Israel. They ran a controversial registry for do-not-spam use, basically using a bot network. When a subscriber signed up for Blue Security they downloaded a piece of software (called Blue Frog), which was a bot, and loaded it on their computer. Every time they received a piece of spam mail the bot would automatically, and in the background, send an opt-out email to the sender. This created a lot of anti-spam spam mail, and was controversial because it was basically using a bot to try to defeat the botnet. The CEO of the company boasted that there were

500,000 registered users, and many of the top spammers had stopped targeting their customers.

On May 1st, 2006, an anonymous hacker (the CEO came to believe it was someone by the name of PharmaMaster, but no identity was ever confirmed) began to send threatening spam to Blue Security subscribers. They read:

> "Unfortunately, due to the tactics used by Blue Security, you will end up receiving this message or other nonsensical spams 20-40 times more than you would normally," one message read. Another predicted: "Soon, you will be found guilty of computer crimes such as DDoS attacking of Web sites, conspiracy, and sending mass unsolicited bulk email messages for everything from Viagra to porn, as long as you continue to run Blue Frog." (Wired)

Not long after that a massive botnet began an attack on Blue Security, taking the website completely offline within hours. The CEO fought on however, and had Tucows redirect his traffic to his TypePad blog to try to explain to users what was happening. Unfortunately, this redirected all of the DDoS traffic as well. Six Apart, the company that ran TypePad, was now in the line of fire. Tucows was also suffering its own damages transferring the DDoS traffic to the destination address; this was causing their own DNS servers to crash meaning they could not supply bandwidth to any other customer.

The attack persisted for three days, the first day being a struggle just to find out what was going on and the next two preparing to make a decision about it. This was to be one of the longest botnet attacks in history. What would the conclusion be? Tucows dropped Blue Security. The CEO of Tucows said they could no longer stay in the fight losing their own customers because of some fighting between Blue Security and a spammer, so they refused to provide bandwidth to the site and took down Blue Security’s DNS rendering its URL useless.

Blue Security was getting desperate now, and turned to Prolexic. Prolexic was a hosting company that had become very famous during 2004 for protecting a number of gaming sites from DDoS attacks. Their filtering and massive bandwidth proved effective, they weathered a massive storm of bot attacks (over 3gbps of traffic) and after several days the attacks began to slow; the spammer seemed to be getting bored. CEO Keith Laslop of Prolexic guessed that the attacks were over. He was wrong.

A little over a week later the hacker decided to try a different route. Instead of targeting Prolexic, he targeted their DNS provider, UltraDNS. UltraDNS was swarmed by the botnet and went offline, taking much of Prolexic’s subscribers (Blue Security included) with it. Prolexic dropped Blue Security. Reshef, the CEO of Blue Security, looked around for other options, but was warned that anywhere he set up shop would simply be swarmed by the botnet and forced to drop him. With that, Blue Security closed its doors for good.

As Scott Berinato of WIRED summarizes: “In the space of two weeks, an unidentified assailant had carried out a series of devastating attacks. He mowed down Six Apart, Tucows, several of their top-tier service providers, Prolexic, UltraDNS, and hundreds of thousands of Web sites, along with millions of Internet users who rely on their services. And he put Blue Security out of business. Why? Who knows. "Someone decided to get rid of Blue Security, and he did," Evron says. "It's as simple as that." (Wired)

Today's Botnet

Unfortunately, botnets are not just a thing of the past. Today, Storm has mostly faded from the internet, but it has been replaced by an even faster growing botnet known as Kraken. Today, Kraken is 495,000 bots large, but it is growing very quickly and expected to be in the millions by only this summer. It has infected PC's in 50 of the top Fortune 500 companies.

What will Kraken be used for, what are its targets? No one knows. All that is for certain is there is a lot of money to be made by some hackers out there, and one can be assured that the Blue Security incident will not be the last of its kind.

Defeating the Botnet

For a small company like Blue Security, suffering through the wrath of such a large network can mean the end of a business. It is an unfortunate truth that there is not much that can be done to prevent it. The botnet, much like a human virus, is incurable. Once it turns its eye upon a company there is nothing outside of brute force bandwidth that can stop it. For a big company like Ebay (who was attacked by a botnet back in 2001), Google, or Microsoft bandwidth might not be an issue. Those companies have the funds available to fight back. A small to mid sized firm does not.

Proaction instead of reaction is the most important thing when dealing with the botnet. DDoS attacks are only as strong as the number of computers that can be controlled. Anti-Virus measures are one way of keeping the bots off of a computer, but the newer scripts are being built specifically with these measures in mind. It really starts with the person sitting behind the PC, and that is an important thing for every company to keep in mind. User policies and user controls can help limit the amount of infections and subsequently the strength of these malicious hackers.

Detection, once such DDoS attacks begin, is also of importance to any firm that thinks it may be the target of such attacks. Having someone on staff that is tech savvy enough to see what these spammers might be instituting against the company is the first step to beginning to regain control of your bandwidth.

The first step, however, is education. All companies need to be aware of the threats that exist out there on the internet, and need to know that simply installing Norton Anti-Virus on a computer is not an effective way to keep your company out of trouble. User policies and rules are a great step in ensuring your companies safety, but even then there exist outside threats such as the botnet that are really outside of your companies control. Solid public relations and a very clean, upright presence on the internet are the only deterrents for some hackers. Do not make enemies that you cannot defeat, because you do not want to stir up a hornets nest.

Gavi Evron is on the forefront of bot control. He used to work with the Israeli government as their head of Internet Security, but now works out of his home leading the bot-busters. “"We need to start doing this as more than an afterthought," he says. Researchers, vendors, and officials must come together to build a broad defense, a combination of technical, legal, regulatory, and social fortifications capable of turning back the bot tide. Last summer, he posted an impassioned plea for such an effort on his botnet discussion board. "We have fallen too far behind for this to go on," he wrote. "All it takes is some good people to make a change." (Wired)

For the time being, however, there is no cure for the bot blues. Education is the first step, and through it hopefully we can all be more cautious internet users. "After learning about bots, you might think, 'I feel hopelessly outgunned and outmatched,'" says Peter Tippett, CTO of security consultancy Cybertrust. "You are."

Work Cited

Tung, Liam. "Storm worm: More powerful than Blue Gene?", ZDNet Australia, September 12, 2007. Retrieved on 2007-10-10.

Gaudin, Sharon. "Storm Worm Botnet More Powerful Than Top Supercomputers", September 6, 2007. Retrieved on 2007-10-10.

Berinato, Scott "Attack of the Bots", WIRED, http://www.wired.com/wired/archive/14.11/botnet.html?pg=4&topic=botnet&topic_set=

"Botnet" Wikipedia, http://en.wikipedia.org/wiki/Botnet

"Storm botnet" Wikipedia, http://en.wikipedia.org/wiki/Storm_botnet

"Kraken botnet" Wikipedia, http://en.wikipedia.org/wiki/Kraken_botnet

"Denial of Service" Wikipedia, http://en.wikipedia.org/wiki/Denial_of_service

"Kraken" Palomides, http://palomides.net/viewtopic.php?t=34778&highlight=kraken

"Move Over Storm," http://www.theregister.co.uk/2008/04/07/kraken_botnet_menace/