INTEGRALIS
an NTT Communications Group Company

DDoS Protection Bypass Techniques

Version: 0.8

Presented by:

Allison Nixon and Christopher Camejo

Integralis, Inc.
60 Hickory Drive
Waltham, MA 02451

Tel: +1 (860) 761-2900
Fax: +1 (860) 761-0007
Web: http://www.integralis.us

INTEGRALIS
an NTT Communications Group Company

Table of Contents

Introduction

Integralis is a global information security consulting and managed security services firm. We offer a broad range of services to our clients, including analysis of threats and incident response. This document relates to recent threat research, and is intended to disclose bypass methods for cloud-based Distributed Denial of Service (DDoS) protection services.

While investigating illicit DDoS-for-hire services that were protected by DDoS protection services, Integralis has come across some techniques that make it possible to bypass certain commercial DDoS protection services. Initially we researched the bypass techniques shared in the criminal underground, and after understanding the underlying mechanisms we began to discover more methods. In our research, we have discovered that the majority of surveyed DDoS protected websites are not protected against these bypass techniques. We have continued investigating these techniques so that we can better advise our clients and we believe that the public needs to be aware of our discoveries because the issues described here may result in attackers bypassing their improperly (and sometimes properly, depending on vendor recommendations) configured DDoS protection services.

Since several companies offer different types of DDoS protection services these techniques may not apply to each and every DDoS protection service. If a company provides cloud based WAF or other filtering services the same bypass methods may apply to those protections too. Fundamentally, these flaws affect any service that uses DNS to route traffic. Services that use BGP to reroute traffic are not affected, nor are inline services.

This document provides a basic overview of DDoS techniques and profiles of some of the groups known to launch DDoS attacks, a discussion of the techniques available to bypass cloud-based DDoS protection services, a detailed description of a proof-of-concept tool Integralis has developed to demonstrate the scanning technique discussed in this document, and descriptions of possible techniques to mitigate these attacks.

DDoS in a Nutshell

What is DDoS?

Distributed Denial of Service (DDoS) attacks are a method of temporarily slowing down or making a target site or network inaccessible to its normal users through a flood of traffic from a large number of other hosts.

Traditional non-distributed Denial of Service (DoS) attacks rely upon a single or small number of hosts exploiting specific vulnerabilities in their target. These vulnerabilities could include programming flaws that allow the service to be crashed remotely via the network or design flaws in the programs or operating systems that make it easy for a single attacker to exhaust the resources of a server. These traditional DoS attacks can be mitigated by fixing the underlying flaws in the software or deploying Intrusion Prevention Systems (IPS) that can detect and block attempts to exploit these flaws.

The reliance of traditional DoS on specific flaws that can be mitigated is a sharp contrast to DDoS attacks that rely primarily on the sheer volume of traffic to overwhelm their targets regardless of any underlying vulnerabilities. The very nature of the Internet's design makes these attacks both possible and difficult to protect against.

While the underlying concept of flooding a target with traffic is simple, there are a number of advanced techniques that have evolved to amplify the effects of the flood and there are a number of different perpetrators behind DDoS attacks, each with their own motives and methods for acquiring control over enough hosts to launch an effective flood.

Attacker Techniques

Exploitation, Obfuscation, and Amplification

The most rudimentary DDoS attacks simply rely on the sheer volume of traffic to exceed the bandwidth capacity of the target's Internet connection. This may be trivial to accomplish when targeting small services hosted internally by companies but becomes much more difficult for services hosted at major data centers with high speed connections. In these cases it becomes necessary to resort to more advanced techniques.

One technique that is commonly in use is to try to exhaust the "connection table" of the target server. The TCP protocol that underlies most services on the Internet uses a "3-way handshake" exchange of packets to establish a connection. The connection table is how the operating system keeps track of all of the open connections to the other systems it is communicating with over the network and the table has a finite size (either limited arbitrarily by the programmer who designed it or by the amount of memory installed in the target machine). Once the table is full, the target can no longer handle any more connections and begins refusing any further attempts to connect to it. In this case, the goal of the attacker is to fill the connection table with his own connections to prevent legitimate users from establishing their own connections.

Filling the connection table can be easy depending on how the target is configured. The attacker can send a packet requesting that a connection be opened, thereby causing the target to add an entry to the connection table, and then intentionally ignore any replies. TCP is designed to be able to cope with packet loss so the target will retry sending the second packet in the connection handshake back to the

attacker before eventually timing out and closing the connection. The connection table becomes easier to fill if the target is configured for a large number of retries with a long timeout. Reducing the retries and timeouts is possible but this runs the risks of breaking connections to legitimate users who happen to be on slow links or on links with high packet loss.

One potential mitigation against this type of attack would be to look for IP addresses that are causing a large number of hung connections and preemptively cull them from the connection table before the timeout and retry limit is reached. Unfortunately, attackers still have the advantage in this situation as they do not intend to receive a reply to establish an actual connection and can easily spoof their source address. Attackers can also monitor the target to determine how long the timeout is set for and begin establishing actual connections which send and request real data albeit very slowly, ideally just fast enough to prevent the timeout from dropping the connection.

Instead of flooding the target directly attackers have also used “reflected” attacks whereby they leverage other legitimate public services on the Internet to flood the target with traffic. This is easy to accomplish with DNS, a service similar to a phone book that turns domain names like “www.integralis.com” into the IP addresses like 131.103.21.131 that operating systems need in order to establish connections across the Internet. DNS uses the connectionless UDP protocol rather than TCP and does not require connection tables and 3-way handshakes. In normal usage a system would send a single request packet to a DNS server which would then send back one or more reply packets depending on the size of the resulting answer. An attacker can easily spoof the source of DNS requests, intentionally sending queries forged to appear as if they are from his intended target to a large number of DNS servers across the Internet with requests designed to elicit the largest responses possible from these servers. This amplifies the attacker’s traffic and conceals his actual source address.

This is certainly not a comprehensive list of DDoS techniques, many combinations of these basic concepts of exploiting known limitations in software design, leveraging other services to amplify attacks, and concealing the attackers’ source addresses are possible. This is simply intended to demonstrate the hurdles facing companies that are facing DDoS threats and their reasons for increasingly turning to DDoS protection services.

Voluntary DDoS

The most basic DDoS attacks are launched by volunteers who willingly run DDoS software on their own personal computers. DDoS tools such as “Low Orbit Ion Cannon” (LOIC) are distributed across the Internet; users can simply download this software, enter a target IP address, and launch an attack. A single home user is unlikely to cause a serious outage but a large number of like-minded individuals can have an impact, especially on smaller sites that are not at major hosting centers.

These types of attacks are typically launched by Anonymous, a hacker group described in more detail below, as they are the only group that has the ability to muster a large enough number of volunteers to make this technique effective on any significant target. The disadvantage to voluntarily DDoSing a target from a personal connection is that it is not very difficult for authorities to track the traffic back to the perpetrators unless steps have been taken to conceal the source of the traffic. A number of perpetrators have been arrested as a result.

Cloud Services

The rise of cheap cloud hosting services has opened a new opportunity for small operators who want the capability to launch DDoS attacks but don’t have easy access to volunteer manpower or large

botnets. Major cloud service providers are typically hosted in top tier datacenters with enormous amounts of bandwidth and virtualized server capacity. Pricing is also normally usage-based and very reasonable. Attackers can set up accounts at these cloud service providers and pay to utilize these legitimate services to send a large volume of traffic to their target from a relatively small number of source machines.

Most cloud service providers prohibit these types of activities in their terms-of-service and some actively monitor for and block this activity, so these types of DDoS attacks are typically launched from providers who do not monitor their networks so strictly.

Botnets

The largest scale DDoS attacks are usually the result of vast botnets. These are worldwide networks of home computers and servers that have been compromised by malware and turned into "zombies" that await instructions. The largest of these networks use multi-tiered architectures where the attacker issues commands to "bot herders" that then forward the commands to the vast network of zombies. These networks are often used to relay spam messages but can also be ordered to unleash large-scale DDoS attacks. Compromised home computers usually come to mind when botnets are mentioned, but compromised servers can contribute much more bandwidth to an attack.

The scale and capacity of these botnets should not be underestimated. The Conficker botnet topped out at 10.5 million compromised zombie hosts. According to Prolexic, a leading DDoS protection service, DDoS attacks in the 1st quarter of 2013 were averaging 48.25 gigabits and 32.4 million packets per second, this exceeds the capacity of an OC-768 (a very large fiber-optic connection used for Internet backbones and trans-oceanic links) and will certainly swamp anything but the largest hosting provider.

Who is Doing it and Why

Anonymous

Many of the high profile DDoS attacks over the past few years have been perpetrated by Anonymous. This group does not have formally defined leadership, membership, or ideology. The group takes action when enough of the individuals who associate themselves with the Anonymous moniker agree that a particular issue warrants their attention and a critical mass is reached.

The group tends to take action on a variety of issues typically related to perceived instances of censorship on the Internet and bullying of individuals by powerful corporations or governments but has also acted on topics as varied as child and animal abuse and hate speech. Since anyone can attribute their activity to Anonymous, the group's activity is very unpredictable. Some of their major activities have included attempts to disrupt activities of the Westboro Baptist Church and Church of Scientology, providing communications support to the Arab Spring protests, attacking credit card companies and Internet hosting providers in retaliation for terminating their business relationships with WikiLeaks, supporting the Occupy movement both with a presence on the ground and through online attacks on banks, and attacks on various companies that publicly announced support for unpopular proposed laws having to do with online anti-piracy or Internet monitoring measures including SOPA, PIPA, and CISPA.

Making exact statements about the composition and capabilities of Anonymous is difficult due to its ad-hoc structure but analysis of previous activities seems to reveal that there is a core organization of a few dozen to a few hundred individuals with the capability to launch reasonably effective hacking attacks to

extract sensitive information from their targets while a large proportion of the group are only downloading and running tools provided by others.

The typical Anonymous "raid" is primarily a DDoS attack intended to knock the target offline causing them public embarrassment and lost revenue. A number of easy to use DDoS software packages are available for download on the Internet; unsophisticated individuals associated with Anonymous will simply download these packages and run them from their own personal computers when they see an announcement of a designated date, time, and target on one of the numerous message boards and chat rooms frequented by Anonymous. These individuals run a high risk of getting caught and a number of them have been successfully prosecuted for their activities in support of Anonymous. More sophisticated individuals will find other systems to run the DDoS software. If those individuals run a botnet, they can use their bots to obscure the true source of the attacks. Often there is talk of using proxies to obscure the attacker's identity, however it is difficult to push a large amount of traffic through public proxies.

As is the case with all DDoS attacks, these activities simply cause the target to be unavailable and do not directly expose any sensitive data to the attacker but DDoS attacks may also act as a sort of "suppressing fire" for further hacking activities that can result in compromised data. When a DDoS attack is launched the target's IT administrators may immediately become focused on keeping their systems online, distracting them from any alerts indicating the other more potent attacks. It is not known if this is an intentional tactic or a possible side-effect of individual uncoordinated attackers each using their own techniques. The major Anonymous attack against Sony in the spring of 2011 are a possible fit for this scenario with a large-scale DDoS attack launched against PlayStation Network. After 3 days Sony shut down the PlayStation Network and announced the compromise of 77 million user accounts, some including credit card data. Interestingly, in the aftermath, Sony announced the estimated cost of this attack at $171 million.

Organized Crime - Blackmail

Online organized crime operations, frequently operating out of Eastern Europe, set up large-scale botnets in order to support their activities including spamming, identity theft, and blackmailing sites with massive DDoS attacks. These activities are conducted as a for-profit business, in some cases raking in billions of dollars from illicit activities.

The prototypical example of this phenomenon is the "Russian Business Network" (RBN). The RBN builds botnets by distributing malware through affiliate programs similar to those used to promote commercial brands. Independent affiliates build websites and purchase ads to trick unsuspecting Internet users into downloading and installing fake antivirus software that make up the "zombies" within the botnet. One of the early operations of the RBN was to DDoS online sports betting operations set up in Costa Rica by a branch of the US mafia. These events led to the formation of Prolexic Technoligies, one of the earliest and currently one of the premier DDoS protection services[1]. The RBN eventually moved on to leasing access to their botnet to spammers as well as offering "bulletproof hosting" for child pornographers and other illicit operations. It is speculated that the RBN has a number of Eastern European political connections that help prevent the breakup of the network by government authorities.

The DDoS attacks launched by organized criminal operations are some of the largest yet seen owing to the enormous botnets they are capable of assembling. The sheer number of zombie hosts, their wide

[1] Menn, Joseph. *Fatal System Error*. Public Affairs Books, 2010

distribution across the Internet, the multi-level command and control architectures used, and the resiliency measures built into the networks' software all make it very difficult to block or quickly dismantle one of these botnets or their attacks.

Booters - DDoS for Hire

A recent phenomenon is the rise of small-scale DDoS for hire services openly advertising themselves as "stress testing" services. While legitimate stress testing services do exist, similar to commercially available penetration testing services, these "Booter" operations are unlike legitimate commercial operations in that they allow attacks to be launched without verifying the identity of their customer or whether their customer owns the targeted systems. These services are regularly advertised on hacking forums as a way to knock home Internet connections and small websites offline. Often these booters are used by gamers to knock their opponents offline, since most games hand a free win to a player if their opponent disconnects. Many small gaming websites and private servers are also targeted by booters. E-sports competitions are also affected and need to take countermeasures to prevent the effects of DDoS attacks because attackers target them for attention and to rig bets on games.

The services are typically run by small groups and payments are received from customers mostly via PayPal. Sometimes booters are powered by botnets but many booter owners simply rent out high bandwidth servers located at so called bulletproof hosting providers. The operations use many of the DDoS obfuscation and amplification techniques described in this document to amplify the impact of their available bandwidth. A victim being attacked by a booter would also not be able to determine the source of the attack.

Getting DDoSed is an occupational hazard of Booter operations. These services often use DDoS protection services themselves to stop attacks from competing booter services. Many of the techniques for bypassing DDoS protection services discussed in this document were developed or perfected while working to unmask illegitimate Booter services.

DDoS Protection Techniques

A number of anti-DDoS techniques exist that may or may not be effective depending on the scale of the attack and the techniques used by the attacker. Simple defensive techniques such as identifying and blocking IP addresses that are launching attacks on a firewall may block small and unsophisticated attacks but are easy to bypass, for example, by spoofing the source address of packets. The firewall may also protect a target server from receiving malicious traffic but the firewall itself has finite resources that can be exhausted. For large scale attacks the target site's Internet connection itself may be overwhelmed with a flood of traffic rendering any on-site protection techniques moot.

Various networking tricks can be used to redirect legitimate users and malicious traffic but most have drawbacks. Redirecting a hostname to an unaffected IP address via DNS during an attack is effective but takes time to propagate across the Internet and it is easy for the attacker to add the new host to his list of targets. Routing protocols like BGP can be used to reroute traffic away from affected networks but can only be used on entire IP blocks rather than individual servers and many companies do not have the capability to reroute their networks if they do not own their own IP range. DoS Defense Systems (DDS) are available that work similarly to Intrusion Prevention Systems (IPS), sitting inline on a network connection monitoring traffic for signs of attacks and blocking malicious traffic while allowing legitimate traffic through; these devices cost money and require administration overhead to manage and may also be rendered moot if the attack is sufficient to saturate the Internet connection that they sit behind.

A more effective technique than attempting to address DDoS attacks within an organization's network is to rely on an upstream provider to offer protection. Hosting companies typically have many times the bandwidth available to a single organization's Internet connection and should be able to stand up to much larger floods of traffic without major impact. These services can be inline, such as an Internet Service Provider offering a DDS on their end of a customer's Internet connection to prevent link saturation, or can be cloud hosted services that filter and redirect traffic using DNS or routing protocols. These redirection techniques are the primary subject of the bypass techniques discussed in this whitepaper.

Bypassing Cloud DDoS Protection

Principle

Many DDoS Protection services are “cloud-based”, that is the traffic intended for a customer’s servers is directed through the provider’s datacenter where it is filtered (referred to from now on as scrubbing servers) and valid traffic is forwarded to the customer’s servers (referred to from now on as the origin). These providers can use their man-in-the-middle position on the network to transparently provide other services such as WAF protection and content caching.

The primary principle behind bypassing cloud based DDoS protection services is actually quite simple: If the attacker can identify the IP address of the origin hidden behind the cloud-based DDoS protection service, he can directly attack the server. Many of the DDoS protection services rely on the customer simply pointing their domain names to the protection servers, accomplishing the rerouting through DNS. Finding the server’s real IP address through information leakage can be quite easy. There are a number of techniques for accomplishing this.

It is important to note that this bypass technique works when DNS is used to reroute traffic, but it will not work if BGP is used and it will not work if the inspection device is used in-line. Due to the nature of BGP, attack traffic sent to the origin IP address can still be routed through scrubbing servers.

Verifying the Origin IP

If you have the origin IP, then the simplest way to verify the identity of an origin is to manually resolve the domain for the secret origin IP in the hosts file of the attacking machine. After adding an entry for the domain and the secret IP, navigating to that domain in the browser can verify the identity of the origin.

Running Wireshark or some other packet capture program is an effective method to verify the traffic is indeed bypassing the DDoS protection service. When visiting the home page and running a packet capture, the outbound GET request should be going to the origin IP and should not be going to any IP that the domain would normally resolve to. After verifying that the traffic is going to the correct IP, the response headers should be verified. Many DDoS protection services add on headers in the response traffic. For example if you view a Cloudflare protected website, you will see a “CF-RAY” response header, and a Cloudflare specific cookie. If you view the origin directly, it will not normally have these headers.

Sometimes there are open reverse proxies out there on the Internet that may be mistaken for the origin, especially while scanning. Checking for the presence of these headers can help to eliminate false positives. Another way to detect the presence of an open reverse proxy is to send it a request with another host name that is protected by the same DDoS protection provider. The reason for the existence of these hosts is unknown but it is worth being aware of.

It should be understood that “name-based virtual hosting” might be in use by the target site. If this setting is enabled, the target site will not return the contents of its stored webpages unless the correct Host header is provided. Any attempt to scan for or verify the origin IP should therefore always contain the correct Host header.

Another technique to verify the origin IP would be to launch a large DDoS attack against a suspected IP. If a server has mitigations in place to prevent serving HTTP responses to requests coming from outside a cloud provider's IP ranges, there may be no obvious signs that it is the origin. Taking down the origin IP can be an effective method for verifying the origin IP if requesting the website through normal means would no longer be possible. The cloud provider would only serve up a cached copy of the webpage if the origin is unreachable.

Related Domains

Perhaps the simplest technique for unmasking the origin IP is to check similar domain names to see if they also host web sites that match the intended target site. For example, if "http://www.integralis.com" is hosted behind a DDoS protection service the attacker could try browsing to "http://ftp.integralis.com". This takes advantage of the fact that many web servers also host other services, FTP being one of the most frequent, that may have their own domain names that are not routed through DDoS protection services. Checking the IP of any known subdomains could reveal that the ftp.integralis.com domain name points to the actual IP address of the server hidden behind the DDoS protection service.

Outbound Connections

Some dynamic web services have the ability to establish outbound connections for various reasons. These can include retrieving files to be uploaded into the web site, such as avatar images on a message board or other application specific purposes. While inbound connections to the web server are routed through the DDoS protection service the reverse proxy appears to only operate one way. Outbound connections are not proxied through the DDoS protection service even though inbound connections are. An attacker can connect to the web site and cause it to take an action that will initiate an outbound connection to a server under his control. He can then check to see what IP address the connection was established from and verify whether or not this is the actual IP address of the server behind the DDoS protection service.

A variation on this technique is to cause the server to send an email. While the server may route the email through intermediate mail gateways rather than establishing a direct connection to the attacker's mail server, the mail headers normally include IP address information indicating the origin of the email and the chain of mail servers that it passed through. These headers could potentially reveal the actual IP address of the protected web server as well.

Servers Leaking IP Addresses

Web servers are complicated and there are a number of opportunities for them to leak information including their actual IP addresses. An attacker may leverage these to discover the server's true IP address if DDoS protection services do not take steps to search for the appearance of the protected server's actual IP address in the content or headers of web pages.

Web server error messages are often quite verbose and can reveal the server's actual IP address. Triggering one of these error messages can be as simple as requesting a non-existent page to trigger a 404 "Not Found" error, for example CVE-2000-0649 has always been seen as a security issue for leaking private IP addresses but it may leak public IP addresses as well. Other server and application specific vulnerabilities could result in disclosure of the public facing IP address. Since this issue is not typically

considered a security vulnerability, there is not much available information on public facing IP address disclosure vulnerabilities. More research needs to be done in this area.

Historical Data

Historical DNS data can be obtained through the ISC's passive DNS database ($7,500 for a copy of the data), and it may also be recorded and found in Google searches for the domain. Most sites that are located behind DDoS protection services were at one point available directly on the Internet and historical DNS records would reflect the server's actual IP address. An attacker utilizing one of these services could quickly identify the actual server.

Most DDoS protection services recommend to their customers that they change their IP address after setting up protection, however this instruction may not be followed due to negligence or operational requirements, as changing the IP address may break infrastructure.

DDoS Provider Specific Information Leakage Issues

In the past, DMCA complaints have been filed against Cloudflare customers for the express purpose of revealing the IP address of the origin[2]. Depending on the policies of a DDoS protection provider, abuse complaints may be used for information leakage purposes.

DDoS protection services also have a bypass mode where the customer can send the traffic to the origin without passing through the DDoS provider's network. This functionality may be triggered if the customer need to perform troubleshooting, and it may be triggered on some services if a DDoS attack exceeds the size of attack the customer paid for. This technique has been observed in use by the criminal underground[3] but Integralis was unable to test it first-hand since testing this behavior would involve attacking Cloudflare's network, and commanding that much bandwidth would either cost a lot of money or involve unethical botnet creation activities.

Scanning

A last resort method of identifying the actual IP address of server behind a DDoS protection service would be to scan a range of IP addresses in an attempt to find a directly accessible web site that matches the protected site. Given enough time this technique should eventually discover the web server hidden behind the DDoS protection service provided that server does not have mitigations in place that would prevent it from responding to an HTTP request coming from outside the cloud provider.

Scanning works on the same principle as verifying the origin IP of a site. Instead of manually resolving the DNS for one suspected IP, the tool simply does this manual DNS resolution and testing for every single IP in a given range.

---

[2] http://blog.cloudflare.com/thoughts-on-abuse

[3] http://webcache.googleusercontent.com/search?q=cache:ZqI3WvNcqOQJ:www.hackforums.net/showthread.php?pid=31747338+&cd=5&hl=en&ct=clnk&gl=us&client=firefox-a (this is a google cached copy of a hackforums thread, a user claiming to be able to bypass Cloudflare in this way)

Integralis has developed a proof-of-concept tool to demonstrate the feasibility of this technique discussed in more detail below.

Scanning Proof-of-Concept Details

Overview

The scanning process involves searching a range of IP addresses to identify live web servers and comparing the pages returned by these web servers against the DDoS protected version of the target site to find a match. Barring some issues with open reverse proxies, discussed below, a matching web site is indicative that the scanner has found the actual web site that the DDoS protection service is forwarding filtered traffic to.

Identifying Target Ranges

Scanning the entire Internet for potential matching web pages is possible as a last resort although it would take a significant amount of time and may generate abuse complaints. Initial searches could be conducted on much smaller ranges such as those belonging to the target company itself or the IP range of hosting providers that the target company is known to use.

WHOIS records tracking the ownership of domain names and IP addresses are easily queried and will likely reveal any IP ranges belonging to the target company. Scans of these ranges will likely reveal the target server if it is located on the target company's own network.

Checking historical DNS data and related domains can also give clues as to what IP ranges the origin may reside on. Victim.com may be behind DDoS protection, but ftp.victim.com may not be and if ftp.victim.com is not the same server as victim.com, the chance of them residing in the same hosting provider's range is very high. If a DDoS protection customer changed their IP address with their hosting provider after enabling protection (as is recommended by multiple providers), their origin server is still within the same IP range of the hosting provider. If their outbound mail is forwarded through other servers, those servers may be on the same hosting provider.

Tracking forum and social media postings by the site's administrators can also give clues. If the administrator recently "liked" or posted about a hosting provider, the IP ranges of that provider may be the most likely place to search for the origin IP.

Web Server Scanning

Retrieving candidate pages from the target IP range is the most time consuming part of the scanning process and limits the rate at which the scan can be conducted. Reducing the HTTP timeout, the amount of time the scanner will wait for a response from an IP address, can drastically reduce the scan time but runs the risk of missing web servers that do not respond quickly.

A multi-threaded process can also be used to request multiple pages simultaneously speeding the process further but this approach will inevitably hit its own resource limits (bandwidth, memory, CPU, number of open sockets, etc.). The proof-of-concept scanner running with 128 threads and a 3 second timeout is capable of scanning a /24 network (256 IP address) in a few seconds using this method while a /16 network (65536 IP addresses) takes about 15 minutes. This scanning concept can be split out across any number of hosts to increase the amount of resources available thereby reducing the scanning time and would be particularly potent if launched from a widely distributed botnet.

Page Matching

Once pages have been retrieved, a simple and fast method of identifying the target page is to simply search the retrieved contents for a snippet of unique text. The attacker can review the DDoS protected version of the page to find a unique string that is unlikely to occur on other web pages before launching his scan and in most cases a unique string should be easy to find.

The scanner must be pointed to a page which will contain unique content. It is not necessary to load the home page. Any arbitrary page will do. For the sake of efficiency, a user of this scanner should search for a page that does not redirect and responds with simple, unique text content.

Another interesting approach that may be useful when a unique string cannot be identified is for the scanner itself to retrieve an entire page from the protected site and compare the contents of this page against the candidate sites. An exact match may not occur in all cases due to dynamic content but a high percentage of correlation between the 2 pages indicates a candidate that the attacker should manually verify. The proof-of-concept tool utilizes both of these modes, accepting a search string on the command line or performing percent matching if the required libraries are installed.

Potential Issues

An issue that occurs during scans is false-positive matches resulting from open proxies within the scanned IP ranges. These open proxies will request and return a web page based on the site requested in the HTTP headers and will appear to the scanner as a name-based virtual host returning the actual web page. A simple manual check can be performed by the attacker to determine if the matching IP address will return content for other web pages known to be located elsewhere. Most of these servers we discovered in our research only served up Cloudflare protected domains and did not respond with the contents of webpages not behind Cloudflare protection. One way to detect these types of servers is to request an unrelated Cloudflare domain. If it returns a result, it is acting as a proxy. Manual verification checks quickly become infeasible for large search ranges that may include hundreds of open proxies. Any open proxy should request and return content from the DDoS protected version of a page. If the DDoS protection provider inserts headers into HTTP responses, the attacker can configure the scanner to use these headers to automatically rule out the false-positives. For example, Cloudflare protected HTTP responses contain the CF-RAY HTTP header.

Mitigation Techniques

Firewalling Origin Server

Firewalling a server off from direct Internet access, only allowing inbound connections from IP address associated with the DDoS protection service, is an easy but incomplete technique to mitigate some of the techniques that can be used to identify the actual IP address of a server behind a DDoS protection services.

Filtering the server in this manner would prevent attackers from checking related domains and scanning for matching web pages but would do nothing to defend against outbound connections from the web server, server information leaks, or historical data revealing the actual IP address. These issues would have to be addressed separately to keep the IP address hidden.

This mitigation cannot protect against an attacker using a DDoS attack as a method of verifying the origin IP. Using access rules only helps to maintain the secrecy (or arguably, security-by-obscurity) of the origin IP.

This mitigation may also harm availability. This is a non-standard configuration and in the event of an availability related problem, this may make troubleshooting more complicated- especially if the current administration is not aware of the configuration.

In-Line DDoS Protection

Real in-line DDoS protection is not vulnerable to these bypass methods. If the filtering appliance is inline, then knowing the IP address of the origin does not help the attacker in any way. The difference between inline and BGP based mitigation methods and DNS based mitigation methods, is that the former methods do not attempt to rely on security by obscurity while the latter does.

In-line DDoS protection may be achieved physically, by putting a filtering appliance directly in the path of the servers, or it may be achieved virtually, by using BGP to redirect traffic intended for the origin IP through the filtering provider instead.

Using BGP to reroute DDoS traffic poses its own limitations- the DDoS protection customer must own their own /24 or larger network range and the switchover adds latency and takes several minutes to take effect. The customer must also deal with privacy concerns because all traffic from that /24 will be rerouted. This may be a problem if other services run in that range. If this can be achieved, it appears this is the only effective “cloud based” method available currently.

If the hosting provider provides DDoS protection as a part of the service, and it is an “in-line” arrangement, this may be the best option for users who cannot afford to buy IP ranges or filtering appliances.

Using a physical appliance has bandwidth limitations, since it cannot help if a DDoS attack’s size exceeds the bandwidth of the customer’s connection. Privacy concerns are largely avoided with this method, since the data would never leave the customer’s premises.

Conclusion

DDoS attacks are not going away any time soon. Hacktivists will continue to use DDoS as a form of protest, organized crime will continue to use this as an effective blackmail method, and small "entrepreneurs" with their "stress testing" services will continue to operate. Attackers have been continually increasing their attack capacity and improving their techniques.

Attempts by organizations to perform their own DDoS protection using existing network security products such as firewalls and intrusion prevention systems are likely to be futile in the face of any sort of organized attack. The technology is not designed to deal with flood of traffic and the Internet link itself is the ultimate bottleneck. Organizations that are likely to be targeted by DDoS attacks and cannot afford downtime should instead turn to commercial protection services if they want to reduce damage from a coordinated attack.

Among the commercial services, those that use network tricks like DNS based redirection are relying on security through obscurity, protecting their customers systems so long as the attacker is unable to determine their target's real IP address. Mitigation methods exist to maintain the obscurity of the origin, but there are many sources of information leakage. It may not be possible to plug all the sources of information leakage and an attempt to try may be more expensive than simply switching to a more effective DDoS mitigation service.

In order to achieve resilient DDoS protection targeted organizations should seek out DDoS protection services that are truly in-line and cannot be bypassed.