A Quick Guide to DDoS Attacks

Practical advice from Wanstor

wanstor

Introduction

A Denial of Service (DoS) attack is an attempt to make one or more systems unavailable to the intended user(s), such as access to online resources e.g. a company website.

A successful DoS attack consumes all available network or system resources, usually resulting in a slowdown or server crash.

Whenever multiple sources are coordinating in the DoS attack, it becomes known as a **DDoS attack**.

wanstor

Wanstor regularly observes two methods of DDoS attacks:

Standard and Reflection

A Standard DDoS attack occurs when attackers send a substantial amount of abnormal network traffic directly to a target server or network.

One of the ways an attacker can accomplish this is by using a botnet to send the traffic. A botnet is a large number of victim computers, or zombies, connected over the Internet, they communicate with each other and can be controlled from a single location.

When an attacker uses a botnet to perform the DDoS attack, they send instructions to some or all of the zombie machines connected to that botnet, thus magnifying the size of their attack, allowing it to originate from multiple networks and in some cases from multiple countries.

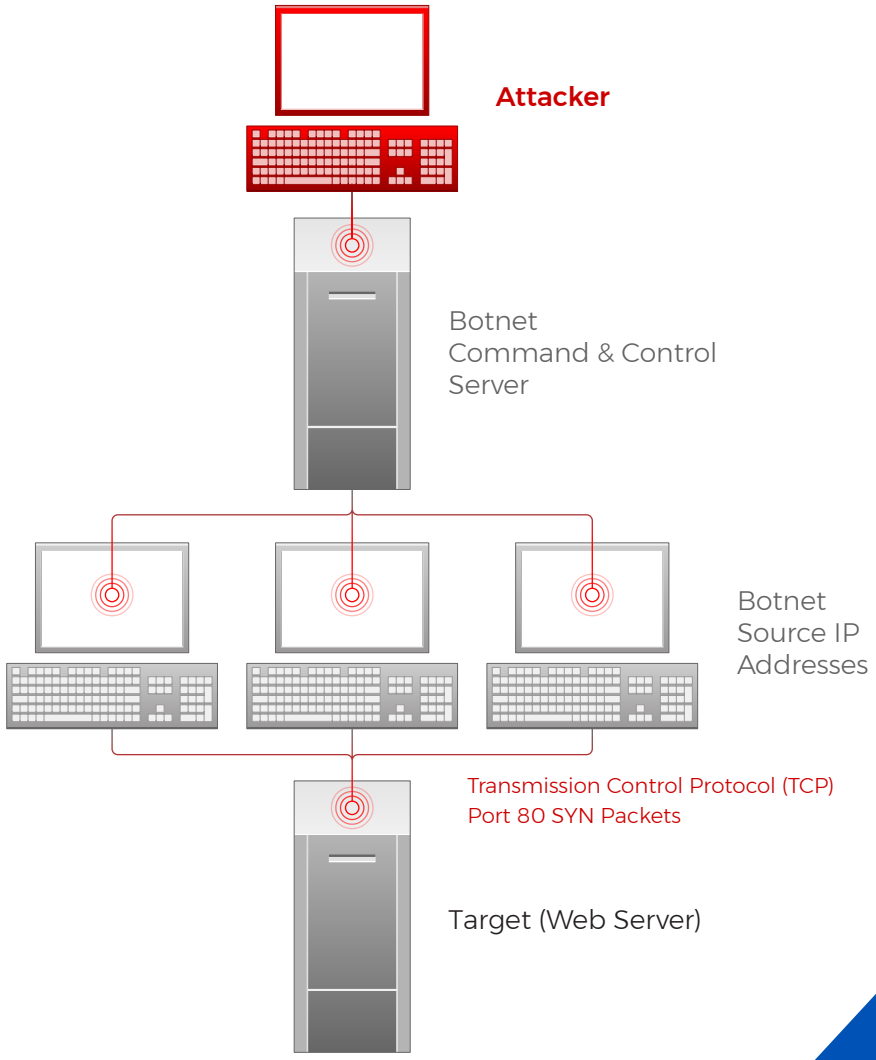

Figure 1 (opposite): Example Standard DDoS SYN Flood

wanstor

A Reflection DDoS attack occurs when attackers spoof their IP address to pose as the intended victim and then send legitimate requests to legitimate public-facing servers.

The responses to these requests are sent to the intended victim and originate from legitimate servers. In addition to these methods, a technique used by attackers to increase the effectiveness of their attack is called Amplification.

Usually used in conjunction with Reflection attacks, Amplification occurs when the response that is sent to the victim is larger than the request that is sent from the attacker. The attacker is able to orchestrate this by requesting a large amount of data from a third-party system.

As the diagram demonstrates, this will probably occur when the attacker spoofs its IP address, pretending to be the victim, and requests all known data from a public server.

This results in the attacker sending a request that is small in size, but results in the public server responding to the victim with a large amount of data.

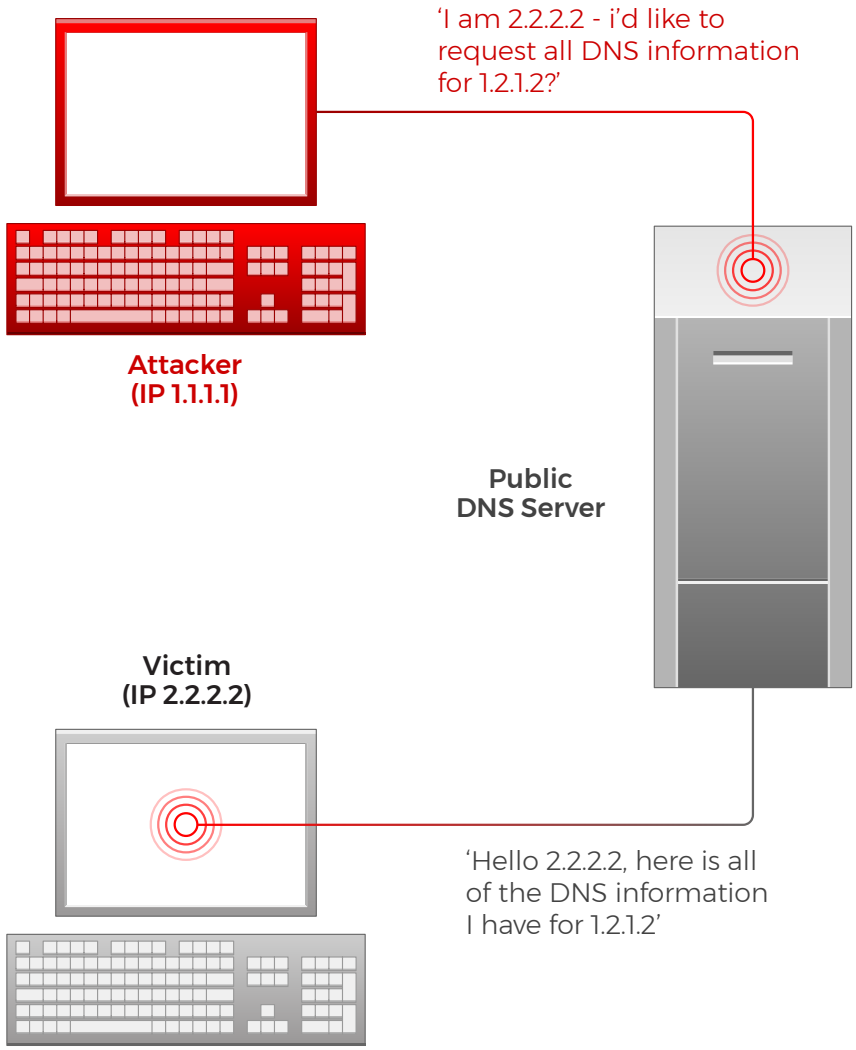

Figure 2 (opposite): Example DNS Reflection DDoS with Amplification

In addition to the use of botnets, some tools are freely available online that cyber criminals can use to perform DDoS attacks. Most of these tools were originally designed to be stress testers and have since become open source tools used to conduct DDoS attacks by amateur cyber criminals.

Tools to be aware of include:

Low Orbit Ion Cannon (LOIC) and the High Orbit Ion Cannon (HOIC). These tools can be downloaded, installed, and utilized by anyone who wishes to be a part of an ongoing DDoS attack.

With the goal of consuming all available bandwidth allocated to the target, the LOIC sends significant amounts of Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) traffic, while the HOIC specifically sends HTTP traffic.

Other examples of tools that can be used to perform DDoS activities and IT teams should watch out for include Metasploit, Pyloris, and Slowloris.

Generally, it is accepted that the main purpose behind a DDoS attack is the malicious consumption of resources. Different attackers may use different techniques to generate the traffic necessary for an effective DDoS. A lone actor with a botnet at their disposal may use that botnet to orchestrate the attacks.

However, botnets are also available for hire, with operators charging minimal fees for short duration attacks.

A group of actors working together may choose to use the same type of free tool, rather than trying to gain access to a botnet.

Attacks like these are usually less successful, as it is difficult to coordinate enough attackers for the effect to be noticeable.

Standard DDoS Attack Types

SYN Flood

An SYN Flood attack is probably the most common form of DDoS attacks observed by Wanstor. It occurs when an attacker sends a succession of TCP Synchronize (SYN) requests to the target in an attempt to consume enough resources to make the server unavailable for legitimate users.

This works because an SYN request opens network communication between a prospective client and the target server. When the server receives an SYN request, it responds acknowledging the request and holds the communication open while it waits for the client to acknowledge the open connection.

However, in a successful SYN Flood attack, the client's acknowledgement never arrives, consuming the server's resources until the connection times out.

A large number of incoming SYN requests to the target server exhausts all available server resources and results in a successful DDoS attack.

Wanstor Recommends

+ To identify an SYN Flood, investigate network logs and locate the TCP SYN flag. TCPDump or Wireshark may work for this.
+ TCP SYN packets are normal and are not necessarily indicative of malicious activity. Instead look for a large number of SYN packets, from multiple sources, over a short duration.
+ If you identify an attack, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.
+ To help minimize the impact of successful SYN Flood attacks, define strict *TCP keepalive* and *maximum connection* rules on all perimeter devices, such as firewalls and proxy servers.
+ On some firewall appliances, you can enable *SYN cookies* to help mitigate the effects of a SYN Flood. This forces a firewall to validate the TCP connection between client and server before traffic is passed to the server. When attackers never send a final acknowledgment of the open connection, the firewall drops the connection.

Slowloris Attacks

While Slowloris is a DoS tool easily accessed by hackers, the term Slowloris is also used to describe a type of DoS attack that attempt to establish multiple TCP connections on a target web server and maintain them for as long as possible by sending partial requests, very similar to a SYN Flood.

ESSYN / XSYN Flood (SYN Flood Variation)

An attack designed to target entities using firewalls which works when a large number of unique source IP addresses all attempt to open connections with the target destination IP. Each new connection from an exclusive source IP creates a new entry in the firewall state table, eventually creating more unique connections than there is space for in the firewall's state table.

Once the table is full, the firewall will not accept any additional inbound connections, denying service to legitimate users attempting to access the destination IP.

PSH Flood (SYN Flood Variation)

A Push (PSH) Flood involves sending large numbers of TCP packets with the PSH bit enabled. PSH packets function to bypass packet buffering, allowing efficient transfer of data by ensuring they are filled to the maximum segment size when multiple packets are sent over a TCP connection.

If the PSH bit is enabled, it indicates the packet should immediately be sent to the application. In normal circumstances, this does not present an issue, however when a significant number of PSH packets are sent to a target server, there is a potential to overload its resources, creating a DoS situation.

UDP Flood

Very similar to an SYN Flood, with attackers using a botnet to send significant amounts of traffic to target servers. This attack is much faster and, rather than attempting to exhaust server resources, it seeks to consume all available bandwidth on the server's network link, denying access to legitimate users.

This works because servers receiving a UDP packet on a network port, such as 50555 / UDP, check for an application listening on that port. If nothing is listening, it replies to the UDP packet sender with an Internet Control Message Protocol (ICMP) Destination Unreachable packet.

During an attack, large numbers of UDP packets arrive with various destination ports, forcing the server to process and respond to each, leading to the consumption of all available bandwidth.

Wanstor Recommends

+ To identify a UDP Flood, investigate network logs and look for a large number of inbound UDP packets over irregular network ports coming from a large number of source IP addresses.

+ Many legitimate services use UDP for their network traffic. Common UDP ports are 53 (DNS), 88 (Kerberos), 137/138/445 (Windows), and 161 (SNMP). When investigating a DDoS attack, look for UDP traffic with high numbered network ports (1024+).

+ If you identify an attack, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.

+ To minimize the effect of UDP Flood attacks, define strict rules on your perimeter network devices, like firewalls, to allow only inbound traffic on ports that are required.

Slowloris Attacks

While Slowloris is a DoS tool easily accessed by hackers, the term Slowloris is also used to describe a type of DoS attack that attempt to establish multiple TCP connections on a target web server and maintain them for as long as possible by sending partial requests, very similar to a SYN Flood.

Wanstor Recommends

+ To prevent remote SMBLoris attacks, configure the border firewall to block all ingress traffic over ports 445 and 139.

+ To prevent internal SMBLoris attacks, set an artificial rate limit for the number of connections local devices can have open.

ICMP Flood

An ICMP Flood occurs when an attacker uses a botnet to send a large number of ICMP packets to a target server in an attempt to consume all available bandwidth and deny legitimate users access.

This attack works when a large number of sources can send enough ICMP traffic to consume all available bandwidth of the target's network

An example of this could be the *ping* command. This command is primarily used to test network connectivity between two points on a network.

However, it is possible to supply this command with different variables to make the ping larger in size and occur more often. By using these variables correctly, and with enough source machines initiating the traffic, it is possible to consume all of the available bandwidth.

Wanstor Recommends

+ To identify an ICMP Flood, investigate network logs and look for significant inbound ICMP traffic from many sources.
+ Depending on the tools used to investigate logs (such as WireShark) you can identify ICMP packets by protocols displayed in the graphical user interface. When analyzing ICMP traffic, you will notice that no port information is available, as ICMP does not use network ports like TCP or UDP.
+ If you are using a tool that displays the network protocols as numbered values, ICMP is protocol 1.
+ There are also ICMP type and code fields that identify what ICMP traffic is being sent or received.
+ If you identify an attack, try to leverage your upstream network service provider in order for them to stifle activity before it reaches your network.
+ To mitigate some damage from ICMP Flood attacks, block ICMP traffic at perimeter network devices such as routers and set a packet-per-second threshold for ICMP requests on perimeter routers. If inbound ICMP traffic exceeds this threshold, excess traffic is ignored until the next second. Packet-per-second thresholds effectively keep your network from being overrun with ICMP traffic.

Smurf Attack (ICMP Flood Variant Using Reflection)

A Smurf attack is an alternative method of undertaking an ICMP Flood attack.

Attackers use the target's IP address as their own, called spoofing, and then send ICMP ping requests to the broadcast IP address of a public network on the Internet.

The broadcast IP address of a network will then send any traffic that it receives to all other IP addresses within its network.

When the ICMP ping request is received by the broadcast IP address, it is then forwarded on to all live computers on its network.

Each of those computers think that these ping requests are coming from the target IP address and send their responses to the target rather than back to the attacker.

The result of this is a large number of unsolicited ICMP ping replies being sent to the target of the DDoS, resulting in the consumption of available bandwidth.

HTTP GET Flood

A HTTP GET Flood attack occurs when an attacker generates a significant number of continuous HTTP GET requests for a target website in an attempt to consume enough resources to make the server unavailable for legitimate users.

In this case, the attacking IP addresses never wait for a response from the target server, despite the server attempting to respond to all incoming requests.

This results in connections being left open on the web server. A large enough number of incoming HTTP GET requests to the target web server eventually exhausts all available server resources and results in a successful DDoS attack.

HTTP POST Flood (HTTP GET Flood Variation)

Another HTTP Flood incorporates the use of the HTTP POST request instead of GET. This attack works because it forces the web server to allocate more resources in response to each inbound request.

A large number of these requests could tie up enough server resources as to deny legitimate users access to the web server. Flood attack.

Wanstor Recommends

+ To identify an HTTP GET Flood, investigate network logs and look for inbound traffic from a significant number of source IP addresses with a destination port of 80 and a protocol of TCP. Packet data should also begin with *GET*. We recommend using either TCPDump or WireShark.

+ HTTP GET requests are normal and are not indicative of malicious activity. Look for numerous identical GET requests from multiple sources over a short period. The same source IP addresses should re-send the same GET requests rapidly.

+ If you identify an attack, leverage a DDoS mitigation service provider for the best results in mitigating this activity.

+ It is difficult to set up proactive security measures to block against this attack, as legitimate traffic is used to carry it out. Often, rate based protections are not sufficient to block this attack, and the source IP addresses of the attack are part of a large botnet, so blocking every source IP is not efficient and may include legitimate users.

+ One solution that may help mitigate this type of attack is to use a Web Application Firewall (WAF). HTTP Floods often exhibit trends that a correctly configured WAF filters and blocks without blocking legitimate access to the web server.

Reflection DDoS Attack Types

NTP Reflection Attack with Amplification

When the attacker uses traffic from a legitimate NTP server to overwhelm the resources of the target, NTP is used to synchronize clocks on networked machines and runs over port 123 / UDP.

An obscure command, monlist, allows a requesting computer to receive information regarding the last 600 connections to the NTP server.

An attacker can spoof the target's IP address and send a monlist command to request that the NTP server send a large amount of information to the target. These responses typically have a fixed packet size that can be identified across a large number of replies.

Since the response from the NTP server is larger than the request sent from the attacker, the effect of the attack is amplified.

When an attacker spoofs the target's IP address and then sends the monlist command to a large number of Internet-facing NTP servers, the amplified responses are sent back to the target.

This eventually results in consumption of all available bandwidth.

Wanstor Recommends

+ To identify an NTP Reflection Attack with Amplification, investigate your network logs and look for inbound traffic with a source port of 123 / UDP and a specific packet size.
+ Once identified, leverage your upstream network service provider, supplying attacking IP addresses and packet sizes used in the attack. Upstream providers have the ability to place a filter at their location forcing inbound NTP traffic using the specific packet size identified to drop.
+ Along with remediating inbound attacks, implement the following preventative measures to ensure that your NTP servers are not used to attack others.
+ If you are unsure whether or not your NTP server is vulnerable to being utilized in an attack, follow the instructions available at OpenNTP: **http://openntpproject.org**
+ Add firewall rules restricting unauthorized traffic to NTP servers.

wanstor

+ Upgrade NTP servers to version 2.4.7 or later, which removes the monlist command entirely, or implement a version of NTP that does not utilize the monlist command, such as OpenNTPD.

+ If you are unable to upgrade your server, disable the monlist query feature by adding *disable monitor* to your ntp.conf file and restarting the NTP process.

DNS Reflection Attack with Amplification

A Domain Name System (DNS) Reflection attack occurs when an attacker manipulates the DNS system to send an overwhelming amount of traffic to the target.

DNS servers resolve IP addresses to domain names allowing users to type an easily remembered domain name into browsers, rather than remembering the IP addresses of websites.

A DNS Reflection attack occurs when an attacker spoofs the victim's IP address and sends DNS name lookup requests to public DNS servers. The DNS server then sends the response to the target server, and the size of the response depends on the options specified by the attacker in their name lookup request.

To get the maximum amplification, the attacker can use the word *ANY* in their request, which returns all known information about a DNS zone to a single request.

When an attacker spoofs a target's IP address and sends DNS lookup requests to a large number of public DNS servers, the amplified responses are sent back to the target and will eventually result in the consumption of all available bandwidth.

Wanstor Recommends

+ To identify if a DNS Reflection Attack with Amplification is occurring, investigate network logs and look for inbound DNS query responses with no matching DNS query requests.

+ DNS queries are normal and are themselves not indicative of an attack.

+ If you identify an attack, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.

+ Along with remediating inbound attacks, disable DNS recursion, if possible, by following the guidelines provided by your DNS server vendor (BIND, Microsoft, etc.). In doing so, this makes sure that your DNS servers are not used to attack others.

+ To discover if any of your public DNS servers may be used to attack others, use the free test at *openresolverproject.org*.

CLDAP Reflection Attack with Amplification

A Connection-less Lightweight Directory Access Protocol (CLDAP) Reflection Attack with Amplification occurs when an attacker sends a CLDAP request to an LDAP server, using a spoofed sender IP address.

CLDAP is used to connect, search, and modify shared Internet directories. It runs over port 389 / UDP.

A CLDAP Reflection attack occurs when a cyber-threat actor spoofs the victim's IP address and sends a CLDAP query to multiple LDAP servers. The LDAP servers then send the requested data to the spoofed IP address.

This unsolicited response is what results in a DDoS attack, as the victim's machine can't process an overabundance of LDAP / CLDAP data at the same time. The amplification is due to the number of times a packet is enlarged while processed by the LDAP server.

LDAP UDP protocol responses are much larger than the initial request with an amplification factor of 52, and can peak at up to a factor of 70.

Wanstor Recommends

+ To identify a CLDAP Reflection Attack with Amplification, investigate network logs and look for inbound traffic with a source port of 389 / UDP.
+ Once identified, try to leverage upstream network service provider and provide them with the attacking IP addresses and the packet sizes used in the attack. Upstream providers have the ability to place a filter at their level.
+ Create a DDoS protection plan.
+ Implement ingress firewall rules that restrict unauthorized use of the LDAP server.
+ Auditing policies can be used to provide reporting of network services that are potentially exploitable as reflection attacks.

Wordpress Pingback Reflection Attack with Amplification

A function of WordPress sites is called the Pingback feature, which is used to notify other WordPress websites that you have put a link to their website on your site.

Sites using WordPress automate this process, and maintain automated lists linking back to sites that link to them. These *pingbacks* are sent as Hypertext Transfer Protocol (HTTP) POST requests to the */xmlrpc.php* page, which is used by WordPress to carry out the pingback process.

By default, this feature downloads the entire web page that contains the link that triggered the pingback process.

An attacker can then locate any number of WordPress websites and then send pingback requests to each of them with the URL of the target website, resulting in each of those WordPress websites sending requests to the target server requesting the download of the web page.

A large number of requests to download the web page can eventually overload the target web server.

Wanstor Recommends

+ To identify a WordPress Pingback Reflection attack with Amplification, investigate your network logs and look for a large number of inbound TCP traffic over port 80 from a large number of sources. The traffic appears as HTTP GET requests for random values such as *?5454545=6767676*, bypassing the cache and forcing a full-page reload for every packet.

+ If you identify an attack, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.

At present, there is no way to prevent this inbound traffic as on its own it is normal web traffic. However, there is a way to make sure your WordPress websites are not used to attack others.

To do this, WordPress offers a tool that is available for download that disables the pingback feature of XMLRPC. Download the tool at this link:

http://wordpress.org/plugins/disable-xml-rpc-pingback/

SSDP Reflection Attack with Amplification

The Simple Service Discovery Protocol (SSDP) is commonly used for the discovery of Universal Plug and Play (UPnP) devices.

UPnP is a series of networking protocols that allows networking devices to discover and connect with one another, without user intervention.

Using SSDP, Simple Object Access Protocol (SOAP) is used to deliver control messages to UPnP devices. A SSDP reflection attack occurs when an attacker spoofs the victim's IP address and sends crafted SOAP requests to open UPnP devices on the Internet.

These devices then send their responses to that victim IP address.

Depending on how the attacker crafted the request, the response could be amplified by a factor of 30 from a single request.

When an attacker spoofs a victim's IP address and sends crafted SOAP requests over SSDP to a large number of public UPnP devices, the amplified responses are sent back to the victim, eventually resulting in the consumption of all available bandwidth.

Wanstor Recommends

+ To identify if an SSDP Reflection Attack with Amplification is occurring, investigate network logs and look for inbound source port 1900 / UDP (SSDP) traffic from a large number of source IP addresses.

+ Once an attack is identified, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.

Microsoft SQL Reflection Attack with Amplification

Microsoft SQL is a popular application used to manage relational databases. Database servers using MS SQL are sometimes left on external IP addresses so that they can be accessed remotely over the Internet.

A MS SQL reflection attack occurs when an attacker spoofs the target's IP address and then sends crafted requests to public-facing MS SQL servers using the MS SQL Server Resolution Protocol (MC-SQLR), which listens on port 1434 / UDP.

The response from the database server contains information about the database instances running on the server as well as how to connect to each one.

Depending on the configuration of the database server, and the number of database instances on the server, the response to the client request could be amplified by a factor of 25 for a single request.

Attackers can send scripted MC-SQLR requests, spoofing the target's IP address, to a large number of public-facing MS SQL servers. Amplified responses are sent back to the target, prospectively resulting in consumption of all the target's available bandwidth.

Wanstor Recommends

+ To identify if a MS SQL Reflection Attack with Amplification is occurring, investigate network logs and look for inbound source port 1434 / UDP (MC-SQLR) traffic from a large number of source IP addresses. In some instances, it may be possible to identify a particular payload signature.
+ If you identify an attack, try to leverage your upstream provider in order for them to mitigate the activity before it reaches your network.
+ If possible, block inbound connections to port 1434 / UDP or filter connections to allow only trusted IP addresses.

General Recommendations and Mitigation Strategies

At Wanstor we believe IT teams should use the following recommendations for DDoS mitigation to reduce the impact of attempted DDoS attacks and enable a faster response when successful DDoS attacks occur:

+ Establish and maintain effective partnerships with your upstream network service provider. Consider how they will assist during a DDoS attack. The faster a provider can implement traffic blocks and mitigation strategies at their level, the sooner your services will become available for legitimate users.

+ Provide attacking IP addresses to your upstream network service provider to implement restrictions at their level. Keep in mind that Reflection DDoS attacks typically originate from legitimate public servers. It is important to ascertain to whom an IP belongs when examining network logs during an attack.

+ Enable firewall logging of accepted and denied traffic to determine where the DDoS may be originating.

+ Define strict *TCP keepalive* and *maximum connection* settings active on perimeter devices like firewalls and proxy servers. This assists in prevention of SYN Flood attacks.

+ Consider port and packet size filtering by the upstream network service provider.

+ Establish and regularly validate baseline traffic patterns (volume and type) for public-facing websites.

+ Apply all vendor patches after appropriate testing.

+ Configure firewalls to block, as a minimum, inbound traffic sourced from IP addresses that are reserved (0/8), loopback (127/8), private (RFC 1918 blocks 10/8, 172.16/12, and 192.168/16), unassigned DHCP clients (169.254.0.0/16), multicast (224.0.0.0/4) and otherwise listed in RFC 5735. This configuration should also be requested at the ISP level.

+ Tune public-facing server processes to allow the minimum amount of processes or connections necessary to effectively conduct business.

+ Configure firewalls and intrusion detection/prevention devices to alarm on traffic anomalies.

+ Configure firewalls only to accept traffic detailed in your organisation's security policy as required for business purposes.

+ Consider setting up Out-of-Band access, Internet and telephony, to an incident management room to ensure connection in the event of a DDoS attack that disrupts normal connectivity.

To find out more information about Wanstor's IT Security solutions and how they can help protect your organisation from attack, please contact us on 0333 123 0360, email us at info@wanstor.com or visit us at www.wanstor.com