Non-flood/non-volumetric Distributed Denial of Service (DDoS)

Prepared by:
Gabriel Garrido, Principal Security Consultant

Table of contents

1. Introduction

Over the last two decades, both Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks have been growing in frequency, complexity and volume.

Traditionally, these attacks are associated with botnets and large amounts of traffic aimed at disrupting Internet-facing services. However, while the goal of these attacks remains the same, the variety and complexity has evolved and coupled with the increase of available network speeds, these attacks are a growing threat.

DoS attacks were almost anecdotal 25 years ago, however, they have now become a real problem for online trading companies and service providers.

Today, if an organisation is subject to a DDoS attack for a period of time it can have a significant impact on revenue, brand reputation and stock market valuation as well as the cost of handling and recovering following an attack.

This whitepaper aims to provide an overview of non-volumetric DDoS attacks, addressing the techniques used to carry out such attacks and the defences or mitigations needed to improve system resilience when under attack.

A DDoS attack can have significant impact on revenue, brand reputation and stock market valuations.

2. Definitions

Denial of Service (DoS) attack

These attacks occur when an attacker takes action that results in the target computer systems being unable to operate as usual.

Distributed Denial of Service (DDoS) attack

DDoS attacks occur when an attacker carries out attacks from multiple sources.

Flood/volumetric DDoS

Flood and volumetric attacks seek to overwhelm the targeted system in order to exhaust the connection limit or the processing power to achieve DoS conditions.

Non-flood/non-volumetric DDoS

By focusing on specific resources and exhausting them, low volume attacks can take out a piece of the infrastructure while being under the radar of any volumetric monitoring. Resource examples include CPU usage, connection table sizes, HTTP session limit and complex back-end database queries.

3. Categories

There are four properties that define and categorise DoS attacks. These are:

IP protocol: Connection vs connectionless

Connection-based attacks require a connection between the client and the server. The connection is either required to launch the attacks or the connection mechanism is abused to cause the DoS attack, for example HTTP slow requests or HTTP session exhaustion.

On the other hand, connectionless attacks do not require a session that is established between the attacker and server. All ICMP/UDP-based DoS attacks are a good example of connectionless attacks since these are connectionless protocols.

However, there is a concept of state when it comes to connectionless protocols. For example, a DNS request for “www.example.com” using UDP is expecting a UDP reply that says it is at IP address “w.x.y.z”.

Volumetric vs non-volumetric

The aim of volumetric attacks is to exhaust bandwidth or processing resources. This is done by flooding the victim with network traffic regardless of the protocol used. Amplification attacks are volumetric.

The goal of non-volumetric attacks is the same as volumetric attacks but without flooding the victim with large amounts of packets. Most non-volumetric attacks are under 100MBps [1].

Direction vs reflection

Direction and reflection attacks are two options used when attacking a system.

A direct attack is where the assailant strikes the victim directly and a reflected attack on the other hand relies on the attacker spoofing the victim’s IP address in order to send packets to multiple third-party reflectors so the third-parties reply directly to the victim.

Network layer vs application layer

Attacks can be aimed at different OSI layers. Layers three-four (network) and seven (application) are the main targets.

4. Real-world examples

There is not enough public incident data to analyse the impact of non-volumetric attacks in the real world. However, the efficiency of these has been well documented.

According to Digital Map Attack [2], all recent DDoS attacks are a combination of all available techniques. However, only volumetric-based attacks tend to hit the news.

As occurrences mutate over time, it becomes more difficult to put the appropriate defences in place. The use of multiple DDoS techniques in a single incident increases the probability of causing real damage. This is because volumetric attacks tends to attract attention and is the first symptom to be addressed when an incident arises.

DDoS attacks are often used as a smokescreen. This involves getting everyone to focus on an attack while sensitive information such as logins, passwords and credit card details are being siphoned off via a discovered exploitable vulnerability in the underlying infrastructure.

5. Techniques

This section covers the most relevant non-volumetric attacks as well as the tools used and the best mitigations measures.

5.1 Spoofed source attack

IP packets contain a source IP address which can be easily crafted and faked. The receiver will process the packet and interpret the source IP as the sender.

A source spoofing attack is usually combined with other techniques to make it more complex. However, by sending a burst of source-spoofed packets from the same spoofed IP to a server, an attacker can cause the source IP to get blocked on a destination firewall. This attack emulates a flood attack behaviour however it is not aimed at causing disruption to the third party service but instead, to force the third party to block the source IP. As such, it is not considered a volumetric attack.

In this scenario, an attacker wants to prevent the victim from accessing a third party service or the third party accessing the victim. If the third party service is protected by firewalls or any IDS/IPS mechanism, the attacker can launch an attack by spoofing the source address with the victim's IP address. When the firewall detects the attack and blocks the IP address, it will also block the victim's IP address and as a result, the victim will not be able to access the third party service.

Examples of the victim could be a payment provider or DNS server.

Tools: Hping2 [3], Nmap [4], Metasploit [5].

5.2 Low & Slow

There are three main Low & Slow attacks. These are Slow POST, Slow HTTP Headers and Slow Read.

A HTTP Slow POST DoS attack follows the principle of low bandwidth and slow action. It aims to exhaust the allowed connection for web servers without affecting bandwidth or other services. However, the victim will still have bandwidth and CPU power to operate as this attack does not target those elements.

The main approach of this attack consists of the attacker sending a valid HTTP POST header with a 'content-type' of the message that is being sent. When the server receives the request, a connection is opened, waiting for the number of bytes to be declared in the content-length header. The attacker then sends the message at the slowest rate possible (slightly under timeout, in most cases one byte every 100 seconds). The server will then have to wait until all of the messages are transmitted, which can take a long time. If the attacker opens all possible connections, other users won't be able to connect to the server.

Slow HTTP Header attacks work on the same principle. However, instead of sending a POST request on the timeout limit, the HTTP header is used. The attacker sends HTTP headers in partial messages with enough intervals for the connection to keep alive and the request not to be completed.

Slow Read attacks send legitimate application-layer requests but read responses very slowly, thus trying to exhaust the server's connection pool. Slow reading is achieved by advertising a very small number for the TCP Receive Window size and by emptying the clients' TCP receive buffer slowly at the same time. By doing this, it naturally ensures a very low dataflow rate.

The traffic generated by the attacker looks legitimate, making it complicated to detect and therefore more possible to bypass some of the protection systems.

Tools: PyLoris [6], Slowloris [7], R.U.D.Y [8], Torshammer [9], High Orbit Ion Cannon (HOIC)

5.3 Application attack

Application attacks relate to application functionalities and are difficult to detect since the attacker is using legitimate traffic. They also have a very wide spectrum as any action that can prevent a valid user from accessing the application is an application DoS [10].

These are aimed at specific application functionalities and it is the abuse of these functionalities that leads to a DoS condition.

Examples of these attacks include:

- Account lockout: An attacker with a list of valid users, forces an account lockout by introducing wrong passwords. As a result, users have to request their account to be unlocked.
- File upload: An attacker exhausts the server’s disk capacity with uploaded files and as a result, the server becomes unusable.
- Database exhaustion: An attacker fills the database with fake records.
- Logic errors: A logic flaw in an application such as an infinite loop in redirections might be invoked to deliberately cause a DoS.

It is difficult to be specific and categorise all potential attacks because different applications and web services have different functionalities. The programming languages used to implement web applications services will also have their own potential security issues and quirks.

Tools: Nikto [11], Accunetix [12], WebInspect [13], Burp [14], OWASP ZAP [15].

5.4 DNS poisoning

When a domain name server does not correctly validate DNS responses from authoritative servers, it can store incorrect responses in the DNS cache table. As a result, any other user trying to resolve that name will be served a poisoned response.

In this scenario, the attacker sends a query to the target DNS server then sends the spoofed reply on behalf of the authoritative DNS server for that domain immediately after. Every DNS request contains a Query ID so the server can differentiate between the different requests and responses. If the target server does not contain the record for the solicited query, it will send a query to the authoritative server. At this point, the attacker can guess the ID and authoritative server, therefore, by sending the poisoned response using the spoofed IP address and ID, the target server will take it as the authoritative response and ignore the genuine response when it arrives. The target server now has a bad DNS record that will get served to anyone asking for it.

Tools: ARPwner [16], Metasploit (DNS BailiWicked Host Attack) [17]

5.5 Black Nurse

Black Nurse [18] exploits how different network device (mainly firewalls) vendors process a specific ICMP packet (Type 3 Code 3). Researchers at TDC-SOC-CERT [19] observed that ICMP Type 3 Code 3 packets were taking significantly more time to be processed than other types. As a result, the attack exploits the processing overhead to cause a service disruption.

This attack can be successful with a low volume of traffic that is between 15Mbps and 18Mbps. This would be suffice to exhaust the network device CPU causing a DoS situation. Most devices return to normal once the attack has ended.

Tools: HPing3 [20]

5.6 SSL-based DDoS attacks

A number of vulnerabilities regarding the SSL/TLS protocol have been discovered, however, only three attack vectors have been proven to have a significant impact. These attacks are brute force, with and without cryptography and on SSL renegotiation.

The conclusions on the effectiveness of these attacks are based on the server-based DoS vulnerabilities in an SSL/TLS Protocols master thesis paper that was published in 2012 [21].

The most popular of the attacks is the abuse of the SSL renegotiation functionality. This is because the required CPU time on the server is larger than the time required by the client to do the renegotiation. While the server is still processing the renegotiation request, the client can keep sending additional renegotiation requests aiming to cause DoS. This is done by exhausting the CPU time on the server.

The brute force attack with a cryptography tool focuses on executing the maximum number of concurrent SSL connections with the SSL server. This tool/technique has not been proven to be as efficient as the abuse of SSL renegotiation, but its results were better than brute force without cryptography.

The most lightweight tool (sslsqueeze) exploits the brute force without cryptography and has a minimal impact on live operation hosts.

Tools: sslsqueeze [22], thc-ssl-dos [23]

5.7 Fragmentation attack

Fragmentation attacks target the packet reassembling process of the underlying OS or network device.

Well known attacks such as Ping of Death and Teardrop fall into this category and most of the documented attacks are variations of these two. When the victim host tries to reassemble the received packets, the specially crafted packets exploit vulnerabilities in the reassembling process, resulting in the host crashing, hanging or rebooting.

Tools: HPing3 [24], Nemesy13

5.8 Shrew attack

The shrew attack is a DoS attack on the Transmission Control Protocol (TCP). By abusing two protocol features, Round Trip Time (RTT) and Retransmission Time Out (RTO), an attacker can impact the actual transmission rate of data. The victim will be throttled to near-zero throughput while the attacker will have a low average rate, making it difficult for counter-DoS mechanisms to detect.

This attack is based on theory and at the time of writing it has not been observed in the wild. No tools have been identified apart from a proof of concept. [25]

Tools: Proof of concept available on the reference site.

6. Classification

The table below shows a classification of the attacks described in the previous section.

The parameters to classify the attacks are based on the definitions provided at the beginning of this document.

Attack	Connection/ connectionless	Volumetric/ non-volumetric	Direction/ reflection	Layer
Spoofed source	Both	Both	Reflection	Network
Low & Slow	Connection	Non-volumetric	Directed	Network
Application	Connection	Non-volumetric	Directed	Application
DNS poisoning	Connectionless	Non-volumetric	Reflection	Network
Black Nurse	Connection	Non-volumetric	Directed	Network
SSL-based	Connection	Non-volumetric	Directed	Application
Fragmentation	Connection	Non-volumetric	Directed	Network
Shrew	Connection	Non-volumetric	Directed	Network

7. The cloud & non-volumetric DDoS

Cloud providers are moving towards providing DDoS mitigation for the cloud. However, most of these mitigation defences are volumetric based. They are the simplest metrics used to analyse attack behaviour and they are also unlikely to detect the slow volume attacks that look like legitimate types of traffic.

Today, it makes little difference as to whether applications are running on real hardware, or out in the cloud when under a non-volumetric type attack. Therefore, appropriate defences need to be put into place on the applications and servers.

8. Defences

It is possible to harden servers against these types of attacks by keeping all software up-to-date and ensuring secure configuration. It is also worth considering installation of DDoS Mitigation technology sat between servers and the Internet.

8.1 AntiSpoofing

In 2000, the IETF published a common practices document (BCP 38 [26]) about defending against DoS attacks which employ IP source address spoofing.

This should be employed at borders with the Internet so that internal users cannot initiate spoofed IP attacks. However, this does not help with spoofing a target's IPs and causing reflective attacks against a target. It also doesn't help with spoofing a third party's IP address when attacking a target so that the target can block the IP address, stopping any legitimate access to the third party.

8.2 Low & Slow

Web servers should be fully patched and configured to run an up-to-date Operating System (OS).

Low and Slow attacks are difficult to defend against for a number of reasons. The traffic generated is very similar to legitimate traffic, which makes it very difficult for network devices to detect them. The amount of traffic required is also very small, therefore, it will be unnoticed by most anti-DoS defence mechanisms.

Recommendations and mitigations include:

- Define minimum acceptable values and reject slower connections
- Set up an absolute connection timeout when possible
- Use of a reverse proxy

Note that adjusting timeout parameters can result in dropping legitimate connections.

8.3 Application attack

Defending against application attacks is difficult. This is because different applications can have different vulnerabilities that can lead to a DoS condition. However, following best practices can help reduce the risk of exposure, for instance:

- Use secure coding methodologies and train developers in secure and defensive coding
- Use a Web Application Firewall (WAF)
- Perform penetration testing of applications before exposing them to the Internet

8.4 DNS poisoning

DNSSEC would be the better solution as the security extensions for DNS add the security controls to prevent this attack, however, DNSSEC implementation rate is marginal.

Without DNSSEC, DNS Servers should be fully patched and running an up-to-date OS. Modern DNS servers such as bind from version 9.50-P1 onwards include security checks that will prevent DNS poisoning attacks. These mitigations include:

- Source port randomisation
- Cryptographic nonces

It is also important to control who is allowed to make recursive DNS requests against DNS servers (i.e. Internet IPs disabled, internal IPs allowed).

8.5 Black Nurse

The recommendation from TDC [27] is to deny ICMP type 3 messages sent to the WAN interface of Cisco ASA firewalls in order to prevent the Black Nurse attack.

8.6 SSL-based DDoS

The most effective method to defend against SSL renegotiation attacks is to disable SSL renegotiation support on the Web Server.

Most DDoS protection providers do not inspect SSL traffic. Inspecting SSL traffic would require decryption, analysis then re-encrypt which would result in an unacceptable overhead due to the computational cost of these operations (in addition to potential privacy impact issue). Mitigations to brute force, with or without cryptography, are in line with any other flood mitigation. The overhead in processing is not as big to take down a service with a single connection and as a result, the behaviour of an attack using brute force techniques will be similar to a flood attack.

8.7 Fragmentation

IP fragmentation attacks are mitigated in several different ways, depending on the type and severity of the attack. Most mitigation methods ensure that malicious data packets never reach their target destinations. The most common involves inspecting incoming packets for violations of fragmentation rules (e.g. using a router or a secured proxy).

Disabling allowed fragmented packets may help, but this may cause issues with VPN tunnels.

8.8 Shrew

According to the attack publishers, there are two ways to mitigate this attack. One is throttling the requests in order to control the short bursts and the other relies on DoS detection algorithms in network devices.

9. Conclusion

DDoS attacks are increasing at significant rates and having significant impact. A recent report [28] cites $2.5M per attack, this includes volumetric attacks.

While hardening and keeping software up-to-date helps, it may make sense to implement additional DDoS mitigations through appliances and services whose focus is on DDoS mitigation.

10. Further reading

Spate of email extortion campaigns threatening DDoS attacks reported in US - https://www.nccgroup.trust/uk/about-us/newsroom-and-events/news/2015/august/spate-of-email-extortion-campaigns-threatening-ddos-attacks-reported-in-us/

The Register: DDoS, the cloud and you - https://www.nccgroup.trust/uk/about-us/newsroom-and-events/in-the-media/2016/july/the-register-ddos-the-cloud-and-you/

What DDoS disasters can teach us about securing the Internet of Things - https://www.nccgroup.trust/uk/about-us/newsroom-and-events/blogs/2016/october/what-ddos-disasters-can-teach-us-about-securing-the-internet-of-things/

Digital Smokescreen Diversion Use in Sophisticated Attacks - https://www.nccgroup.trust/uk/about-us/newsroom-and-events/blogs/2015/august/digital-smokescreen-diversion-use-in-sophisticated-attacks/

Why economics and the free market likely contributed to Friday's Internet service disruption - https://www.nccgroup.trust/uk/about-us/newsroom-and-events/blogs/2016/october/why-economics-and-the-free-market-likely-contributed-to-fridays-internet-service-disruption/

10.1 Further reading resources

NCC Group's DDoS Assured https://www.nccgroup.trust/uk/our-services/security-consulting/managed-and-hosted-security-services/vulnerability-management-and-detection/

NCC Group's DDoS Secure https://www.nccgroup.trust/uk/our-services/security-consulting/technology-solutions/defence-in-depth/

11. References

[1] https://security.radware.com/ddos-knowledge-center/ddos-chronicles/ddos-attacks-history/

[2] http://www.digitalattackmap.com

[3] http://www.hping.org/

[4] https://nmap.org/

[5] https://www.metasploit.com/

[6] https://motoma.io/pyloris/

[7] https://github.com/llaera/slowloris.pl

[8] https://code.google.com/p/r-u-dead-yet/

[9] https://github.com/dotfighter/torshammer

[10] https://www.nccgroup.trust/globalassets/our-research/uk/whitepapers/the-new-ddos-battleground-white-paper-final.pdf

[11] https://github.com/sullo/nikto

[12] https://www.acunetix.com/

[13] https://saas.hpe.com/en-us/software/webinspect

[14] https://portswigger.net/burp/

[15] https://www.owasp.org/index.php/OWASP_Zed_Attack_Proxy_Project

[16] https://github.com/ntrippar/ARPwner

[17] https://www.metasploit.com/

[18] http://blacknurse.dk

[19] http://www.netresec.com/?page=Blog&month=2016-11&post=BlackNurse-Denial-of-Service-Attack

[20] http://tools.kali.org/information-gathering/hping3

[21] https://www.jbisa.nl/download/?id=17683062

[22] https://github.com/mmgaggle/sslsqueeze

[23] https://www.thc.org/thc-ssl-dos/

[24] http://tools.kali.org/information-gathering/hping3

[25] http://www.cs.northwestern.edu/~akuzma/rice/shrew/

[26] https://tools.ietf.org/html/bcp38

[27] http://blacknurse.dk/

[28] http://www.zdnet.com/article/the-average-ddos-attack-cost-for-businesses-rises-to-over-2-5m