TCPWave DDI - DNS DDoS Attacks & Response Rate Limiting

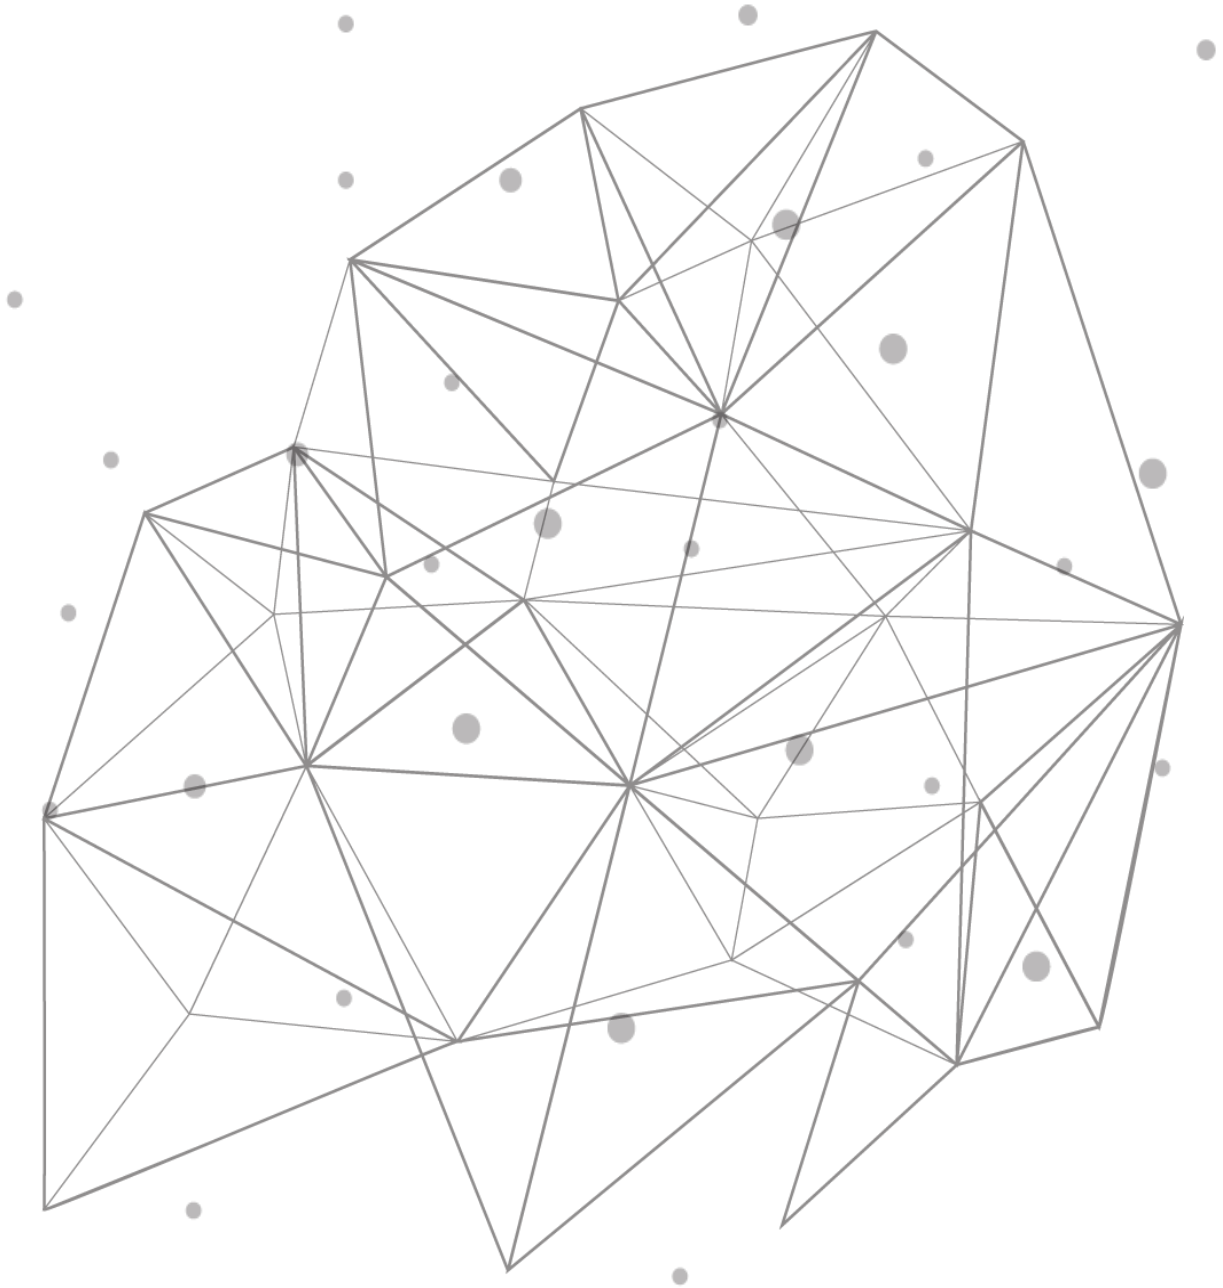

Introduction

Substantial development in technology is constantly extending the world to new milestones and more significant challenges, adding fuel to the vulnerabilities. The attacks have become one of the biggest threats to the internet, cloud platforms, and data centers. This whitepaper provides insights on categories of DNS DDoS attacks and their mitigation using TCPWave's advanced security feature – Response Rate Limiting.

About DDoS

Distributed Denial of Service attacks (DDoS) has been raging for more than 20 years since it was reported in 1998. This attack has been aggressive, destructive, and has had an enormous impact. It is a malicious action to interrupt the service and make it unavailable to legitimate users by suspending the services of its hosting server. The attack is launched from numerous compromised devices, referred to as a botnet. Once the malicious actors establish the botnet, they send the malicious instructions to each bot, which leads to total prevention of the web resource's normal functioning, a total 'denial of service.' The following figure illustrates one type of DDoS attack.

DDoS – Categories

There are three basic categories of attack:

- **Volumetric Attacks**: These attacks use high traffic to inundate the network bandwidth.
- **Layer 7 or Application Layer Attacks**: These attacks focus on web applications and are the most severe and sophisticated type of attack. This type of attack results in the application not being able to deliver data to legitimate users.
- **Protocol Attacks**: These attacks focus on exploiting the server resources.

Volumetric Attacks

ICMP Flood

ICMP (Internet Control Message Protocol) flood is also known as PING flood. In this DDoS attack, the malicious actors send ICMP echo request packets to the targeted server using multiple devices such as computers, IoT devices, etc. The targeted server then sends an ICMP echo reply packet to each requesting device's IP address as a response. This causes the target server to be inaccessible to the normal user.

IP Fragmentation

These attacks involve the transmission of fraudulent ICMP packets larger than the network's maximum transmission unit. As these packets are fake, the target server's resources are quickly consumed, resulting in server unavailability.

UDP Flood Attack

In a User Datagram Protocol (UDP) flood attack, the malicious actors send highly-spoofed user datagram protocol packets at a very high rate using a large source IP range. The target's network is overwhelmed by many incoming UDP packets. This type of attack consumes network resources and available bandwidth, exhausting the network until it becomes unavailable.

Application Layer Attacks

The following are the most common application layer attacks:

- HTTP(s) Flooding
- SYN Flood
- Slow Rate

HTTP Flooding

The malicious actor tries to send many queries to the system. At that instance, the system becomes overloaded and cannot resolve the queries of legitimate uses. This causes consumption of resources such as bandwidth, CPU, etc.

SYN Flood

This attack is based on the principle of a three-way handshake of TCP connection. The malicious actor sends SYN (synchronize) packets to every port on the server, which the server acknowledges by sending SYN-ACK (synchronize – acknowledge) back to the server. Still, the malicious actor does not return the ACK packet. Therefore, the server status is still waiting to be completed. This causes consumption of resources in terms of the number of connections, bandwidth, CPU, etc.

Slow Rate

A Slow Rate attack is also known as a Low and Slow DDoS attack. This attack relies on low traffic targeting the server resources and crippling the webserver preventing legitimate requests. It becomes difficult to distinguish from legitimate traffic.

Protocol Attacks

The Protocol attacks are also known as state-exhaustion attacks. These attacks consume the network resources and exhaust additional network equipment such as firewalls and server load balancers.

DNS DDoS Attacks

The Domain Name System (DNS) is one of the significant components of the internet. It resolves names into IP addresses and vice versa. It relies on the UDP protocol, making it vulnerable to DDoS attacks. The DNS DDoS attacks overwhelm the DNS servers, thus making the DNS service unavailable. Hence the organization's services are not up, leading to the halt of businesses. The most common DNS DDoS attacks are as follows:

- DNS Amplification Attacks
- DNS Reflection Attacks
- NXDOMAIN Attacks
- Phantom Domain Attacks

DNS Amplification Attacks

The attack aims to bombard the DNS server with fake DNS requests that consume the network bandwidth until the DNS service fails. A large volume of packets is generated, flooding the target site without the intermediate noticing. The malicious actors use tools to direct the requests to poorly protected services, which causes noticeably larger responses than the primary request.

DNS Reflection Attacks

This attack leverages the connectionless nature of the UDP protocol to abuse DNS servers that are configured as open resolvers. The malicious actors spoof the target's IP address and dispatches the request for the information, mainly via the UDP. The server responds by answering the target's IP address.

NXDOMAIN Attacks

In the NXDOMAIN attack, the malicious actors send a flood of queries to a DNS server to resolve invalid or non-existent records. The DNS server tries to resolve it but cannot find it. In the process, its cache gets filled up with bad requests, slowing the response for legitimate requests.

Phantom Domain Attacks

The malicious actors set up a bunch of “phantom” domains that do not respond to DNS queries. When phantom domain attacks happen, the DNS recursive server continues to query non-responsive servers, which causes resource consumption. The DNS server may drop legitimate queries when resources are fully consumed, causing performance issues.

Impact on Business Operations

- It prevents the users from accessing the data.
- Service/Application unavailability might fail the organization to meet its Service Level Agreements (SLA).
- Organizations depending on web traffic might face financial loss.

Mitigation Using Response Rate Limiting

With DNS, the organizations cannot block all the packets. To address the DDoS attacks against the DNS infrastructure, the applications should have the ability to separate the valid queries from malicious queries. The malicious actors excel at modifying their queries to look legitimate, so the way is to look at queries’ pattern, rate, and signature. This is the underlying idea behind response rate limiting (RRL).

Response Rate Limiting

RRL allows the network administrators to prevent the malicious actors from using the DNS appliances to initiate the DDoS attacks. RRL implementation is recommended for the authoritative servers. Users can also use RRL implementation for cache servers. This feature uses a token bucket scheme. Each combination of identical response and client identity has an account that earns a specified number of credits every second. A response debits its account by one and is truncated or dropped while the account is negative. Responses are tracked within a window of time, which defaults to 15 sec, but the users can configure to any value from 1 to 3600 seconds using the TCPWave IPAM DNS Option Template. You can also configure the below RRL settings in TCPWave IPAM on how requests are handled when the DNS appliances receive multiple requests.

Name	Description
Responses Per Second	This option defines the maximum allowed identical responses per second from any source IP address. Values: 0-1000. Default Value: 0, it means Responses Per Second has no limit, so Rate Limit section will not be generated in the named.conf. Rate Limit section will be generated in the named.conf only when the value is non zero.
Referrals Per Second	This option defines the number of referrals per second allowed before rate-

Name	Description
	limiting is triggered. It can be used with responses-per-second or nxdomains-per-second to cover error conditions in delegation-centric or mixed domains. Allowed values range: 0-1000. Default Value: 0 or the same value of Responses Per Second.
NODATA Per Second	This option is used to control rate-limiting based on the number of NODATA responses. Setting the value to 0 indicates that no limits are applied to this category.
NXDOMAINs Per Second	This option is used to control rate-limiting based on the number of NXDOMAIN responses.
Errors Per Second	This option allows the users to set the error limit (REFUSED, FORMERR, or SERVFAIL) responses.
All Per Second	This option counts all the responses sent to a client IP.
Window	The option specifies the period in seconds over which rates are measured.
QPS Scale	This option tightens defense during attacks. When the query per second rate exceeds the qps-scale value, then the responses per second, errors per second, nxdomains per second, and all per second values are reduced by the ratio of the current rate to the qps scale value.
IPv4 Prefix Length	Rate-limiting is performed per client IP address where the client IPv4 address is modified by the IP prefix to create an IP address block.
IPv6 Prefix Length	Rate-limiting is performed per client IP address where the client IPv6 address is modified by the IP prefix to create an IP address block.
Slip	The option discards the packets before sending a response with the DNS truncated bit set to 1.
Log Only	The allowed values are either Yes or No; setting the value to Yes, the rate-limiting function does not perform and will log when the rate-limit function is invoked.
Exempt Clients	This option defines all clients' address match list structure to whom the rate-limiting function does not apply.
Max Table Size	The rate limit information is defined in tables. The maximum size of 20K entries is defaulted for all rate-limiting tables covered by this rate-limit

Name	Description
	statement.
Min Table Size	The minimum size of 500 entries is defaulted at start-up for all rate-limiting tables covered by this rate-limit statement.

Solutions

TCPWave sets a high-security standard by offering robust, scalable, integrated, and thorough protection of DNS infrastructure. With the TCPWave's multi-layer defense mechanism, you can adequately protect the organization from the full spectrum of DNS DDoS attacks. For more information on how TCPWave's security features can meet your business needs, contact the TCPWave Sales Team.