XENA NETWORKS

WHITE PAPER

Distributed Denial of Service

DDoS Attack Testing and Verification Solutions

OVERVIEW

First seen around 2000, Distributed Denial-of-Service (DDoS) attacks are a serious threat to businesses around the world. Attackers use multiple hosts to swamp targets with bogus traffic, paralyzing the network and potentially costing the victims millions of dollars. There are many security systems for preventing DDoS attacks – and they all need to be thoroughly tested and verified prior to being activated.

Xena offers a complete test solution for DDoS mitigation and network security with high-performance products and ample features. Going beyond generating DDoS traffic, Xena's solutions can help companies test their security products and operators test networks and detect flaws, thereby ensuring business continuity and preserve business integrity.

"Xena recreates complex traffic so client and server communicate in exactly the same order as the captured traffic to ensure realistic network scenarios for the DUT".

XENA NETWORKS

WHITE PAPER

Distributed Denial of Service

DDoS Attack Testing and Verification Solutions

Contents

INTRODUCTION 3

DDOS Attacks and Business Disruption 4

Understanding Different DDoS Attacks 5

UDP Flood Attack 6

Ping of Death 7

Ping Flood 8

Smurf Attack 9

ARP Spoofing 9

Teardrop Attack 10

TCP Sequence Prediction Attack 11

Xena DDoS Test and Verification Solutions 12

Importing Captured Traffic as Template 13

CONCLUSION 14

INTRODUCTION

DDOS – MAJOR NETWORK SECURITY THREAT

Distributed denial-of-service (DDoS) attack is a growing threat to business around the world with no sign of abating any time soon. It is estimated that DDoS attacks can make up as much as 10 percent of a country's total internet traffic according to 2016 Visual Networking Index report from Cisco Systems. With the prediction of 21% IP traffic compound annual growth rate (CAGR), by 2020, the global DDoS attacks can increase up to 17 million, 2.6-fold increase from 2015, as shown below.

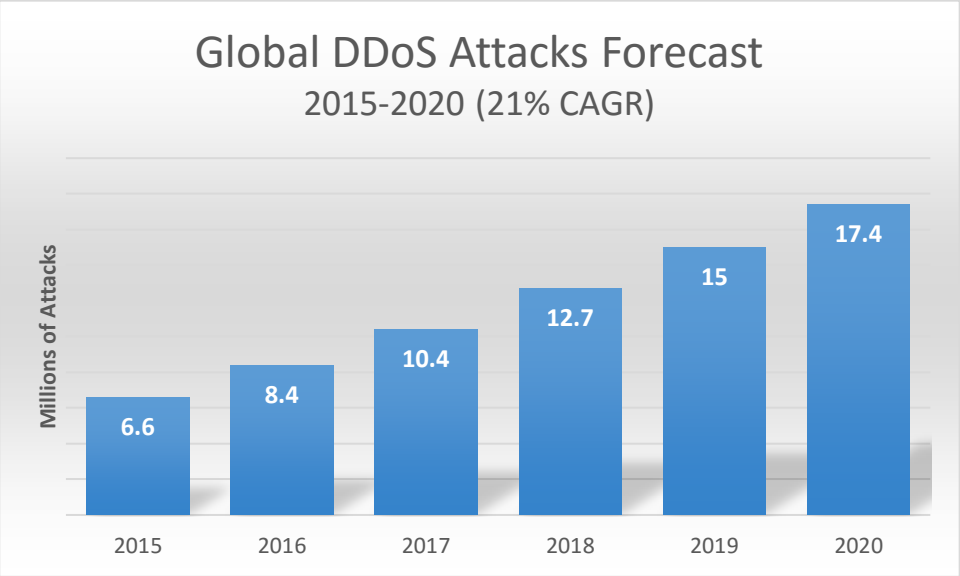

As the name suggests, DDoS attackers use multiple hosts to to overwhelm a target with bogus traffic. The network is paralyzed due to the overwhelmed servers, network links or devices (firewalls, routers, switches, etc.), and this can cost the victim millions of dollars. The average size of DDoS attacks is increasing progressively and approaching 1 Gbps with the peak size reaching 500 Gbps in 2015 - enough to bring most organizations offline completely.

DDoS attacks not only target individual websites at the edge of the network but also the network infrastructure, such as the aggregating or core routers and switches, or the Domain Name System (DNS) servers in provider networks. These attacks can cause more serious damages due to the size and scale of the victim network.

Thus, it is vital to properly deploy network security solutions into enterprises, telecom operators, ISPs, etc. More importantly, these solutions, devices, firewalls, or security software need be thoroughly tested and verified before roll out.

This White Paper provides examples of how DDoS disrupt business, the various types of DDoS attacks, and how Xena's innovative technology and high-performance product can be used to emulate DDoS traffic for mitigation testing and verification.

WHITE PAPER

DDOS Attacks and Business Disruption

Many companies are heavily dependent on the internet. Information is stored in the cloud for ubiquitous access and sharing. Mobile apps and websites are developed for user experience and expanding market shares. Financial transactions and mobile payments are processed online. Logistics is monitored and tracked through the internet. Businesses are moving from offline to online. For today's business, a top-notch customer experience is critical for success. That requires providing rich and responsive access to online services through the internet infrastructure.

DDoS attacks drastically impact access, no matter how much you invest in high-quality mobile experience for your users or in data security and availability. By flooding servers with garbage requests to consume network and computing resources, DDoS attacks can slow down or completely bring down server performance, preventing users from accessing the services they need. If the users cannot access their accounts or data on the servers, there is zero experience and zero usage. From the customer perspective, the services appear unstable and potentially insecure. This can result in a huge loss of customers and business.

In 2009, the popular gaming platform, Xbox Live, became a frequent target for DDoS attacks. The attacks had the ability to affect any of over 17 million users, kicking players off the network.

In April 2011, WordPress, a blog-hosting site with over 19 million blogs, was temporarily brought offline by a DDoS attack. Sensitive data was probably taken, according to the blog host.

In 2012, major US banks were disrupted by DDoS attacks. Website operations at Bank of America and at least five other major US banks used hundreds of compromised web servers to flood their targets with above-average amounts of Internet traffic. The DDoS attacks also caused disruptions at JP Morgan Chase, Wells Fargo, US Bancorp, Citigroup, and PNC Bank.

In 2015, the largest ever DDoS attack hit an Asian network operator, with Arbor Networks estimating traffic hit upwards of 334 Gbps.

In early 2016, BBC's website suffered a DDoS attack resulting in an extensive outage. In total, the entire domain of BBS - including its on-demand television and radio player - was unavailable for more than three hours.

Since DDoS attacks originate from multiple hosts across the internet and sometimes looks legitimate to the network operator, it is difficult to block once launched. Thus, the number one

priority is to test and verify any DDoS defense and mitigation systems before rolling out services to your customers.

Understanding Different DDoS Attacks

Any network attack that attempts to make a machine or network resource unavailable to users can be categorized as a DoS attack. When multiple attack sources are involved, it is called a DDoS attack. Here are the major types of DDoS attacks:

TCP SYN Flood Attack

A TCP SYN flood is a form of denial-of-service (DoS) attack where the attacker sends a sequence of SYN requests to a target in an attempt to consume enough resources to paralyze the server and make the system unresponsive to legitimate traffic.

A normal TCP connection (below) is initiated by a 3-way handshake between the client and server:

1. Client sends a SYN request message to the server.
2. Server acknowledges with a SYN-ACK message to the client.
3. Client replies with an ACK message

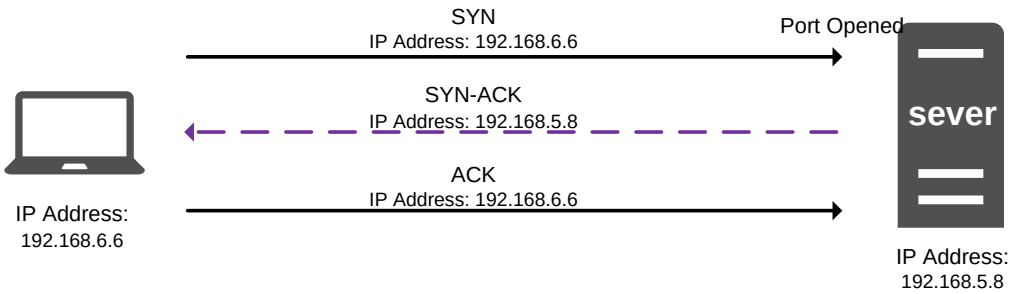

Figure 2: Normal TCP Connection

A SYN flood attack, (see Figure 3), works by deliberately not responding to an ACK to the server. The malicious client can either not send the expected ACK, or by spoofing the source IP address in the SYN message, causing the server to send the SYN-ACK message to a false IP address, which will not send an ACK because it never sent a SYN.

Either type of the attack can result in excessive use of resources on the server for the half-open connections created by the malicious client. At this point, the server cannot connect to any clients, whether legitimate or otherwise, because of the lack of enough resources.

Countermeasures to TCP SYN flood attacks are listed in RFC 49887 and solutions based on these countermeasures are commercially available. However, it is important to test those solutions, before they are brought online in the real world.

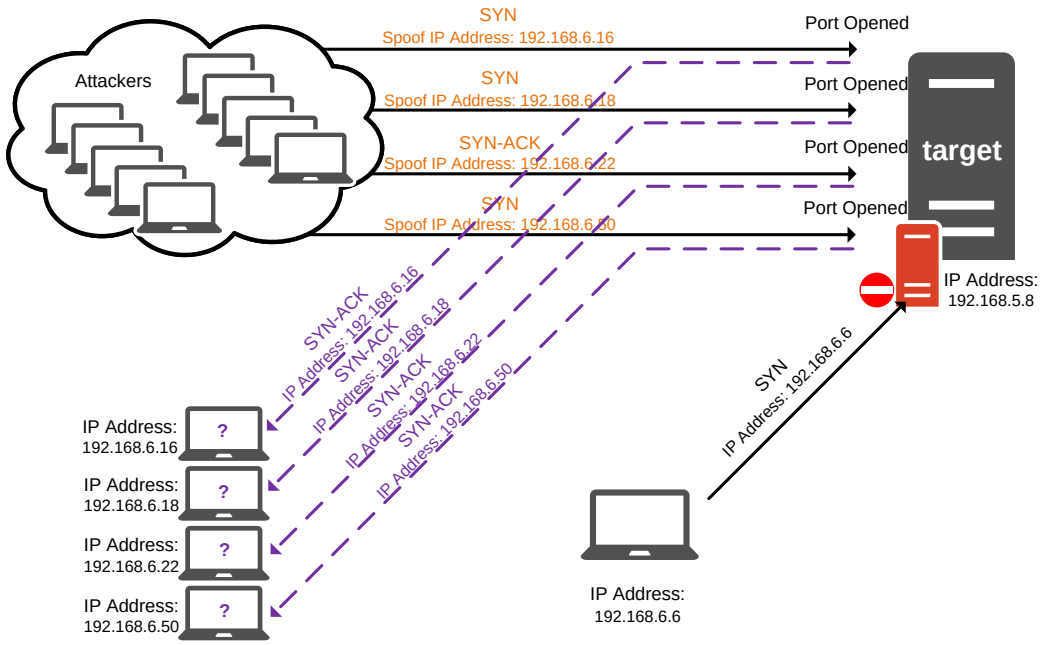

Figure 3: TCP SYN flood attack

UDP Flood Attack

This DoS attack uses the User Datagram Protocol (UDP), a connectionless computer networking protocol, to send a large number of UDP packets to random ports on a remote host. The target host checks for the application listening at that port, and replies with an ICMP Destination Unreachable packet if the host cannot find the application. This forces the target host to transmit excessive ICMP packets, eventually making it unreachable by other clients. The attackers can also spoof the IP address of the UDP packets to ensure the returned ICMP packets do not arrive (see below).

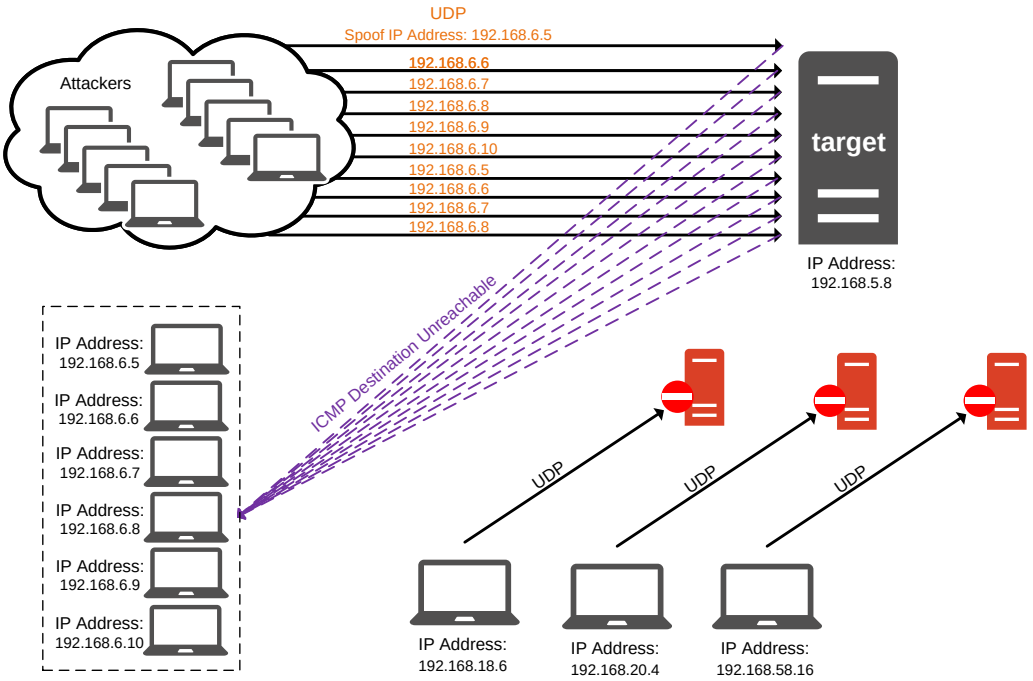

Figure 4: UDP flood attack

At the most basic level, most operating systems try to mitigate UDP flood attacks by limiting the rate of ICMP responses. UDP mitigation also relies on firewalls to filter out unwanted network traffic. The potential victim will not receive or respond to the malicious UDP packets if the firewall can stop them. However, as firewalls are stateful devices, i.e. can only keep a limited number of sessions due to memory constrains, they can also be vulnerable to flood attacks.

Ping of Death

Ping of Death is a type of DoS attack where the attacker tries to crash the targeted computer or system by sending malformed or malicious ping packets.

The correct size of a ping packet is typically 56 bytes. However, any IPv4 packet can be as large as 65,535 bytes. Some historical computer systems simply could not handle larger packets, and would crash if they received one. Since sending a ping packet larger than 65,535 bytes violates the Internet Protocol, attackers would usually send malformed packets in fragments. When fragmentation is performed, each IP fragment needs to carry information about which part of the original IP packet it contains. This information is kept in the Fragment Offset field, in the IP header. When the target system attempts to reassemble the fragments and ends up with an

oversized packet, buffer overflow could occur, causing a system crash and potentially allowing the injection of malicious code.

Ping of Death attacks were particularly effective because the attacker's identity could be easily spoofed. A Ping of Death attacker needs no detailed knowledge of the targeted machine, except for its IP address. The problem described above in fact has nothing to do with ICMP, which is used only as the payload. The problem lies in the reassembly process of IP fragments, which may contain any type of protocol.

One mitigation method is to use a firewall to block ICMP ping messages, or check each incoming IP fragment to make sure that the sum of "Fragment Offset" and "Total Length" fields in the IP header of each IP fragment is smaller than 65,535.

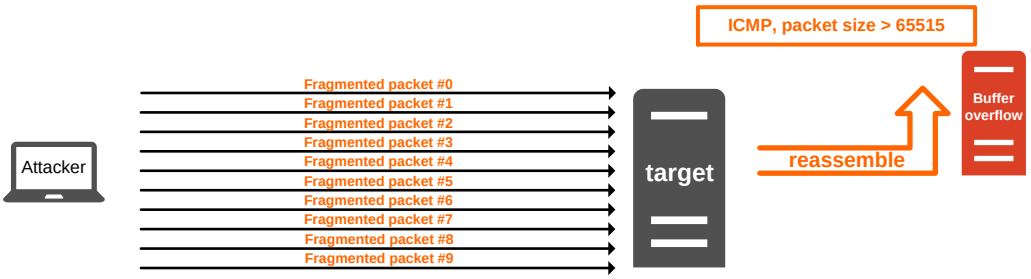

Figure 5: Ping of Death attack

Ping Flood

A ping food attack is a simple DoS attack where the attacker overwhelms the targeted host with ICMP Echo Request (ping) packets. This attack works most effectively by sending ICMP packets as fast as possible without waiting for replies. It is most successful if the attacker has more bandwidth than the victim. This attack seeks to overwhelm the targeted host's ability to respond – consume enough of its CPU cycles for a user to notice a significant slowdown – thereby blocking valid requests.

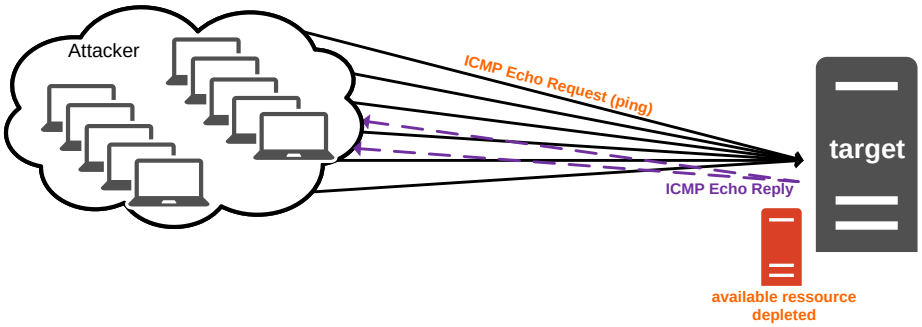

Figure 6: Ping flood attack

Smurf Attack

Smurf attack is a type of DDoS attack in which larger numbers of ICMP packets with the intended victim's spoofed source IP address are broadcast to a computer network using an IP broadcast address. By default, most devices on the network will respond by sending a reply to the source IP address. When the number of devices is larger, the victim's computer will be flooded with replying traffic. This will render the targeted host unresponsive.

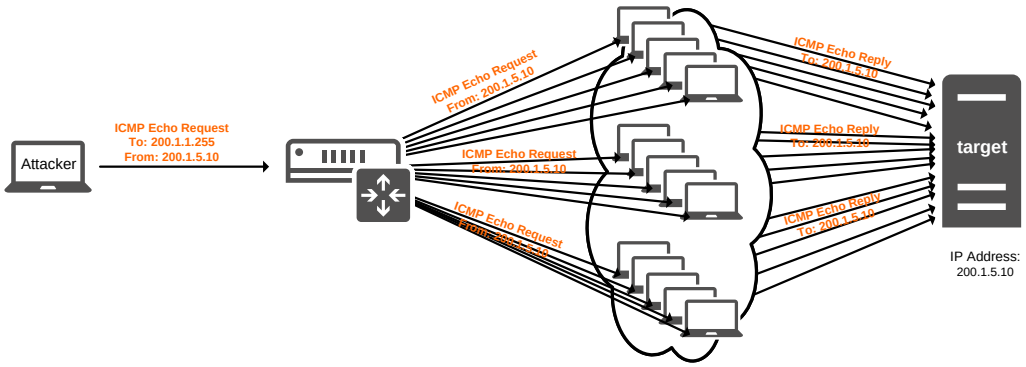

Figure 7: Smurf attack

ARP Spoofing

ARP spoofing, ARP poison routing, or ARP cache poisoning, are all types of attacks where an attacker sends spoofed Address Resolution Protocol (ARP) messages to a local area network (LAN). The aim is to associate the attacker's MAC address with the IP address of another host, such as the default gateway, resulting in any traffic bound for that IP address to be sent to the attacker instead.

ARP spoofing attacks are extremely easy to carry out as long as the attacker has control of a machine within the target LAN or is directly connected to it as the ARP protocol was designed for

efficiency not for security. ARP spoofing may allow an attacker to intercept data frames on a network, modify the traffic, or stop all traffic, as an opening for other attacks such as man in the middle (MITM), denial of service, or session hijacking.

Defenses against ARP spoofing include static ARP entries and ARP spoofing detection. Static ARP solution requests that the IP-to-MAC address mappings in the local ARP cache be statically configured so the host ignores all ARP replay packets. ARP spoofing detection generally relies on some form of certification or cross-checking of ARP responses. Uncertified ARP responses are blocked. These techniques may be integrated with the DHCP server so that both dynamic and static IP addresses are certified. This capability may be implemented in individual hosts or may be integrated into Ethernet switches or other network equipment.

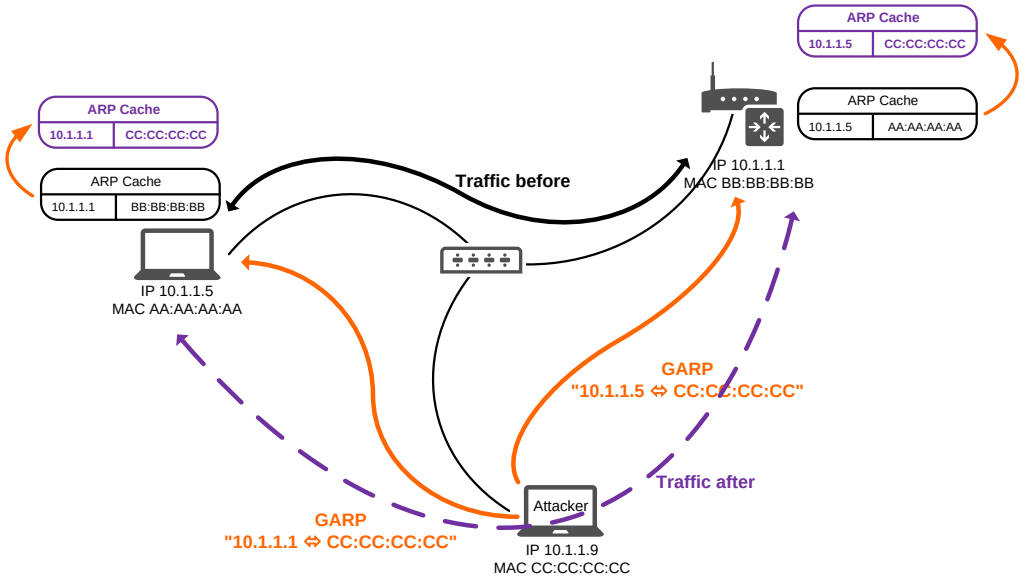

Figure 8: ARP Spoofing

Teardrop Attack

A teardrop attack is a DoS attack that involves sending fragmented packets to a target machine. Since the target receiving such packets cannot reassemble them due to a bug in TCP/IP fragmentation reassembly, the packets overlap one another, crashing the target network device.

IP fragmentation is the process of breaking up a single Internet Protocol (IP) datagram into multiple packets of smaller size because every network link has a characteristic size of messages that can be transmitted, called the maximum transmission unit (MTU). The IP layer in the network protocol stack is responsible for the transmission of packets between network

endpoints, which includes fragmentation of larger packets into small ones for the supporting datalink frames. The 13-bit Fragment Offset field in the IP header specifies the fragment's position within the original datagram, measured in 8-byte units.

The teardrop attack occurs when two fragments contained within the same IP datagram have offsets that indicate that they overlap each other in positioning within the datagram. This could mean that either fragment A is being completely overwritten by fragment B, or that fragment A is partially being overwritten by fragment B. Some operating systems do not properly handle fragments that overlap in this manner. As a result, the data packets overlap and quickly overwhelm the victim's servers, causing them to fail.

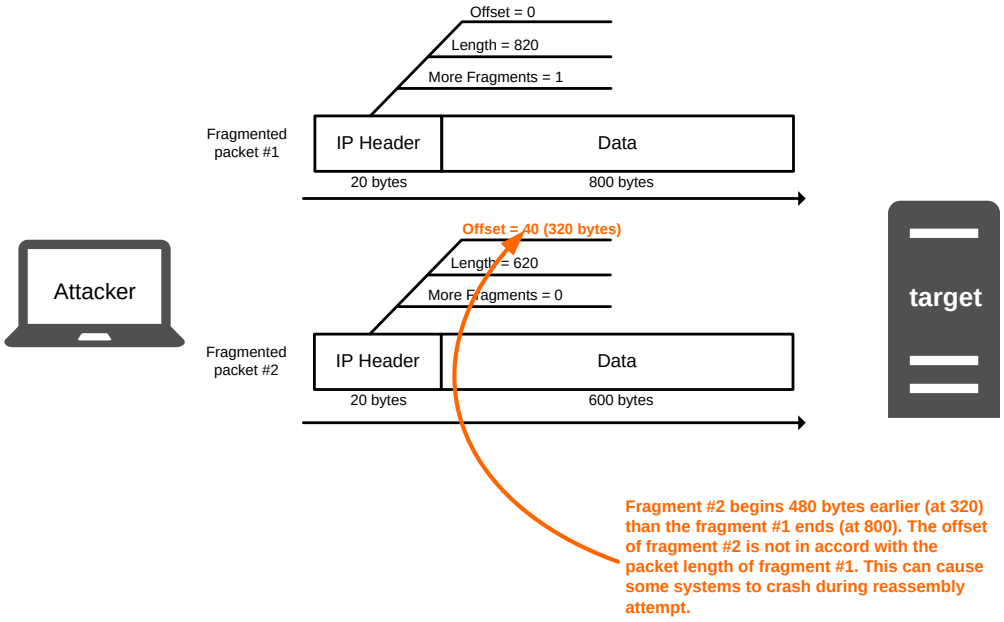

Figure 8: Teardrop attack

TCP Sequence Prediction Attack

A TCP sequence prediction attack is an attack that predicts the sequence number used to identify the packets in a TCP connection, which can be used to counterfeit packets. The attacker hopes to correctly guess the sequence number to be used by the sending host. If they can do this, they will be able to send counterfeit packets to the receiving host which will seem to originate from the sending host, even though the counterfeit packets may in fact originate from some third host controlled by the attacker.

If an attacker can deliver counterfeit packets of this sort, they can cause various sorts of mischief, including injecting data of the attacker's choosing into an existing TCP connection, and prematurely closing existing TCP connections by injecting counterfeit packets with the RST bit set.

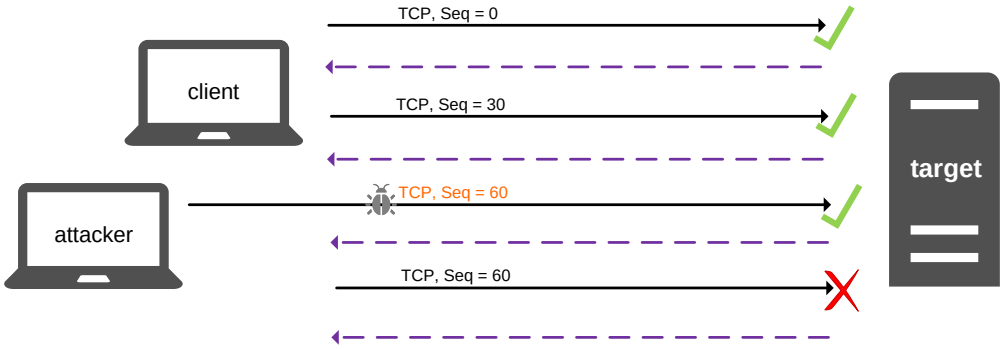

Figure 9: TCP sequence prediction attack

Xena DDoS Test and Verification Solutions

Xena offers a complete test solution for DDoS mitigation on e.g. firewalls, routers, and servers, based on its high-performance network test and measurement products. In terms of DDoS testing, Xena is capable of:

- Importing captured traffic as a template.
- Blasting user-defined packet streams from layer 2 up to layer 7.
- Transmitting and receiving traffic at rates ranging from a few Kbps to 100 Gbps.
- Real-time analysis and reporting.
- Offering ready-to-use port configuration files, step-by-step guidelines, and pcap examples

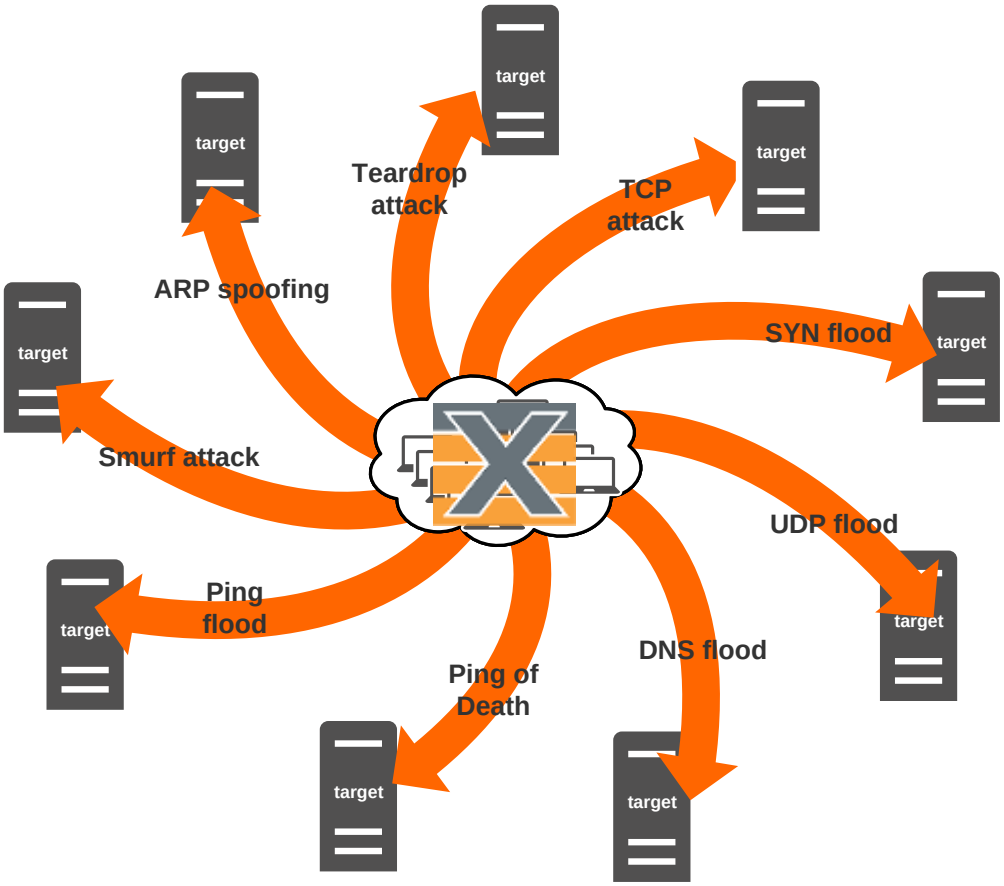

Figure 10: Using Xena to test for DDoS attacks

Importing Captured Traffic as Template

Users can emulate a simple DDoS attack at low transmission rate, e.g. ping the target from one IP address, and capture the traffic into a pcap file. This file can then be imported into the Xena platform which will parse and analyze the file and generate a packet template based on the user's selection. The simple attack packet thus becomes the template for a more complex DDoS attack.

Packets can be completely configured by users, or Ethernet, Ethernet II, VLAN, ARP, IPv4, IPv6, UDP, TCP, LLC, SNAP, GTP, ICMP, RTP, RTCP, STP, SCTP, MPLS, PBB, FCoE, IGMPv2/3, or fully user-specified. Any field in a packet template can be set to an invalid value for negative testing

Blasting User-Defined DDoS Attack Traffic

Xena's test solutions allow up to six field modifiers to be applied to any field in a packet, per stream. A field modifier can be set to increment or decrement or be random within a specified

range. The modifiers can be chained together so that a simple attack can grow into a complex distributed attack with various combinations. For instance, the user can add a modifier to the source IP address of the ICMP Echo Request (ping) packet and randomize it. This lets the user generate a large amount of distributed ping traffic towards a single destination – the ping flood attack is therefore ready.

Transmitting and Receiving Traffic at Different Rates

With well-defined traffic, users can select different bandwidth distribution, constant (uniform) or burst distributions. Traffic loads can be specified as percentages of line rate, frames per second, or Mbps. For example, if the link is a 100G link, the user can set the traffic load to 50% with constant bandwidth and generate a ping flood attack at 50 Gbps.

Real-Time Analysis

Packets can be captured and exported for further analysis using WireShark. Triggers and filters can be set up to trigger on specific events, and to capture packets meeting particular criteria. Multiple capture criteria can be specified using AND/OR expressions.

Ready-to-Use Port Configuration Files, Guidelines, and Pcap Examples

Xena offers port configuration files for the aforementioned DDoS attacks. User can download these configuration files from Xena's website and load to ports accordingly. The configurations can be adjusted as needed after loading, depending on different scenarios or requirements.

In addition to the port configuration files, there are step-by-step guides for users to download. These guides offer detailed information about each DDoS attack and how to configure DDoS streams using the pcap examples provided.

CONCLUSION

DDoS attacks will continue to grow in both scale and severity. Our ever-increasing dependence on the internet make businesses vulnerable to malicious attacks. As the cost of attacks rise, various DDoS mitigation and network security solutions, products, and systems are developed. The battle against DDoS attack shows no sign of diminishing. On the contrary, firewalls, routers and servers require constant upgrades to patch bugs and improve security. Thus, benchmarking test and verification of these products become vital.

Xena delivers a complete test solution for DDoS mitigation and network security with high-performance products and ample features. Xena also provides ready-to-use port configuration files, step-by-step guides and pcap examples to help users quickly learn and configure DDoS attack streams for testing.

Going beyond generating DDoS traffic, Xena helps companies test their security products, and operators to test networks and detect flaws, so that companies can ensure business continuity and preserve business integrity.

WHITE PAPER