

Denial of Service Problems and It's Solutions



Geetanjali Jindal

Research Scholar,
Deptt. of Computer Science,
Faculty of Science,
Tantia University,
Sri Ganganagar, Rajasthan



Kalpna Midha

Associate Professor,
Deptt. of Computer Science,
Faculty of Science,
Tantia University,
Sri Ganganagar, Rajasthan

Abstract

Behind a Client is a person that orchestrate an attack. A Handler is a compromised host with a special program running on it. Each handler is capable of controlling multiple agents. An Agent is a compromised host that runs a special program. Each agent is responsible for generating a stream of packets that is directed toward the intended victim.

Attackers have been known to use these four programs to launch DDoS attacks:

1. Trinoo
2. TFN
3. TFN2K
4. Stacheldraht


In order to facilitate DDoS, the attackers need to have several hundred to several thousand compromised hosts. The hosts are usually Linux and SUN computers; but, the tools can be ported to other platforms as well. The process of compromising a host and installing the tool is automated. The process can be divided into these steps, in which the attackers:

1. Initiate a scan phase in which a large number of hosts (on the order of 100,000 or more) are probed for a known vulnerability.
2. Compromise the vulnerable hosts to gain access.
3. Install the tool on each host.
4. Use the compromised hosts for further scanning and compromises.

Because an automated process is used, attackers can compromise and install the tool on a single host in under five seconds. In other words, several thousand hosts can be compromised in under an hour

Keywords: Denial, Service Problem, Solution.

Introduction



A denial-of-service attack (DoS attack) or distributed denial-of-service attack (DDoS attack) is an attempt to make a computer or network resource unavailable to its intended users. Although the means to carry out, motives for, and targets of a DoS attack may vary, it generally consists of the concerted efforts of a person, or multiple people to prevent an Internet site or service from functioning efficiently or at all, temporarily or indefinitely. Perpetrators of DoS attacks typically target sites or services hosted on high-profile web servers such as banks, credit card payment gateways, and even root nameservers. The term is generally used relating to computer networks, but is not limited to this field; for example, it is also used in reference to CPU resource management.

One common method of attack involves saturating the target machine with external communications requests, such that it cannot respond to legitimate traffic, or responds so slowly as to be rendered effectively unavailable. Such attacks usually lead to a server overload. In general terms, DoS attacks are implemented by either forcing the targeted computer(s) to reset, or consuming its resources so that it can no longer provide its intended service or obstructing the communication media between the intended users and the victim so that they can no longer communicate adequately.

Denial-of-service attacks are considered violations of the IAB's Internet proper use policy, and also violate the acceptable use policies of virtually all Internet service providers. They also commonly constitute violations of the laws of individual nations. When the DoS Attacker sends many packets of information and requests to a single network adapter, each computer in the network would experience effects from the DoS attack.

Review of Literature

The United States Computer Emergency Readiness Team (US-CERT) defines symptoms of denial-of-service attacks to include:

1. Unusually slow network performance (opening files or accessing web sites).
2. Unavailability of a particular web site.
3. Inability to access any web site.
4. Dramatic increase in the number of spam emails received—(this type of DoS attack is considered an e-mail bomb)

Denial-of-service attacks can also lead to problems in the network 'branches' around the actual computer being attacked. For example, the bandwidth of a router between the Internet and a LAN may be consumed by an attack, compromising not only the intended computer, but also the entire network.

If the attack is conducted on a sufficiently large scale, entire geographical regions of Internet connectivity can be compromised without the attacker's knowledge or intent by incorrectly configured or flimsy network infrastructure equipment

Types of Attack

A "denial-of-service" attack is characterized by an explicit attempt by attackers to prevent legitimate users of a service from using that service. There are two general forms of DoS attacks: those that crash services and those that flood services.

A DoS attack can be perpetrated in a number of ways. The five basic types of attack are:

1. Consumption of computational resources, such as bandwidth, disk space, or process or time.
2. Disruption of configuration information, such as routing information.
3. Disruption of state information, such as unsolicited resetting of TCP sessions.
4. Disruption of physical network components.
5. Obstructing the communication media between the intended users and the victim so that they can no longer communicate adequately.

A DoS attack may include execution of malware intended to:

1. Max out the processor's usage, preventing any work from occurring.
2. Trigger errors in the microcode of the machine.
3. Trigger errors in the sequencing of instructions, so as to force the computer into an unstable state or lock-up.
4. Exploit errors in the operating system, causing resource starvation and/or thrashing, i.e. to use up all available facilities so no real work can be accomplished.
5. Crash the operating system itself.

ICMP flood

A smurf attack is one particular variant of a flooding DoS attack on the public Internet. It relies on misconfigured network devices that allow packets to be sent to all computer hosts on a particular network via the broadcast address of the network, rather than a specific machine. The network then serves as a smurf amplifier. In such an attack, the perpetrators will send large numbers of IP packets with the source address faked to appear to be the address of the victim. The network's bandwidth is quickly used up, preventing legitimate packets from getting through to their destination. To combat Denial of Service attacks on the Internet, services like the Smurf Amplifier Registry have given network service providers the ability to identify misconfigured networks and to take appropriate action such as filtering.

Ping flood is based on sending the victim an overwhelming number of ping packets, usually using the "ping" command from unix-like hosts (the -t flag on Windows systems has a far less malignant function). It is very simple to launch, the primary requirement being access to greater bandwidth than the victim.

Ping of death is based on sending the victim a malformed ping packet, which might lead to a system crash.

SYN Flood

A SYN flood occurs when a host sends a flood of TCP/SYN packets, often with a forged sender address. Each of these packets is handled like a connection request, causing the server to spawn a half-open connection, by sending back a TCP/SYN-ACK packet (Acknowledge), and waiting for a packet in response from the sender address (response to the ACK Packet). However, because the sender address is forged, the response never comes. These half-open connections saturate the number of available connections the server is able to make, keeping it from responding to legitimate requests until after the attack ends.

Teardrop Attacks

A Teardrop attack involves sending mangled IP fragments with overlapping, over-sized payloads to the target machine. This can crash various operating systems due to a bug in their TCP/IP fragmentation re-assembly code. Windows 3.1x, Windows 95 and Windows NT operating systems, as well as versions of Linux prior to versions 2.0.32 and 2.1.63 are vulnerable to this attack.

Around September 2009, a vulnerability in Windows Vista was referred to as a "teardrop attack", but the attack targeted SMB2 which is a higher layer than the TCP packets that teardrop used.

Low-rate Denial-of-Service attacks

The Low-rate DoS (LDoS) attack exploits TCP's slow-time-scale dynamics of retransmission time-out (RTO) mechanisms to reduce TCP throughput. Basically, an attacker can cause a TCP flow to repeatedly enter a RTO state by sending high-rate, but short-duration bursts, and repeating periodically at slower RTO time-scales. The TCP throughput at the attacked node will be significantly reduced while the attacker will have low average rate making it difficult to be detected.

Peer-to-Peer Attacks

Attackers have found a way to exploit a number of bugs in peer-to-peer servers to initiate DDoS attacks. The most aggressive of these peer-to-peer-DDoS attacks exploits DC++. Peer-to-peer attacks are different from regular botnet-based attacks. With peer-to-peer there is no botnet and the attacker does not have to communicate with the clients it subverts. Instead, the attacker acts as a "puppet master," instructing clients of large peer-to-peer file sharing hubs to disconnect from their peer-to-peer network and to connect to the victim's website instead. As a result, several thousand computers may aggressively try to connect to a target website. While a typical web server can handle a few hundred connections per second before performance begins to degrade, most web servers fail almost instantly under five or six thousand connections per second. With a moderately large peer-to-peer attack, a site could potentially be hit with up to 750,000 connections in short order. The targeted web server will be plugged up by the incoming connections.

While peer-to-peer attacks are easy to identify with signatures, the large number of IP addresses that need to be blocked (often over 250,000 during the course of a large-scale attack) means that this type of attack can overwhelm mitigation defenses. Even if a mitigation device can keep blocking IP addresses, there are other problems to consider. For instance, there is a brief moment where the connection is opened on the server side before the signature itself comes through. Only once the connection is opened to the server can the identifying signature be sent and detected, and the connection torn down. Even tearing down connections takes server resources and can harm the server.

This method of attack can be prevented by specifying in the peer-to-peer protocol which ports are allowed or not. If port 80 is not allowed, the possibilities for attack on websites can be very limited.

Asymmetry of resource utilization in starvation attacks

An attack which is successful in consuming resources on the victim computer must be either:

1. Carried out by an attacker with great resources, by either: controlling a computer with great computation power or, more commonly, large network bandwidth controlling a large number of computers and directing them to attack as a group. A DDOS attack is the primary example of this.
2. Taking advantage of a property of the operating system or applications on the victim system

which enables an attack consuming vastly more of the victim's resources than the attacker's (an asymmetric attack). Smurf attack, SYN flood, Sockstress and NAPTHA are all asymmetric attacks.

An attack may utilize a combination of these methods in order to magnify its power.

Permanent Denial-of-Service Attacks

A permanent denial-of-service (PDoS), also known loosely as phlashing, is an attack that damages a system so badly that it requires replacement or reinstallation of hardware. Unlike the distributed denial-of-service attack, a PDoS attack exploits security flaws which allow remote administration on the management interfaces of the victim's hardware, such as routers, printers, or other networking hardware. The attacker uses these vulnerabilities to replace a device's firmware with a modified, corrupt, or defective firmware image—a process which when done legitimately is known as flashing. This therefore "bricks" the device, rendering it unusable for its original purpose until it can be repaired or replaced.

The PDoS is a pure hardware targeted attack which can be much faster and requires fewer resources than using a botnet in a DDoS attack. Because of these features, and the potential and high probability of security exploits on Network Enabled Embedded Devices (NEEDs), this technique has come to the attention of numerous hacker communities. PhlashDance is a tool created by Rich Smith (an employee of Hewlett-Packard's Systems Security Lab) used to detect and demonstrate PDoS vulnerabilities at the 2008 EUsecWest Applied Security Conference in London.

Application-Level Floods

Various DoS-causing exploits such as buffer overflow can cause server-running software to get confused and fill the disk space or consume all available memory or CPU time. Other kinds of DoS rely primarily on brute force, flooding the target with an overwhelming flux of packets, oversaturating its connection bandwidth or depleting the target's system resources. Bandwidth-saturating floods rely on the attacker having higher bandwidth available than the victim; a common way of achieving this today is via Distributed Denial of Service, employing a botnet. Other floods may use specific packet types or connection requests to saturate finite resources by, for example, occupying the maximum number of open connections or filling the victim's disk space with logs.

A "banana attack" is another particular type of DoS. It involves redirecting outgoing messages from the client back onto the client, preventing outside access, as well as flooding the client with the sent packets.

An attacker with shell-level access to a victim's computer may slow it until it is unusable or crash it by using a fork bomb.

Nuke

A Nuke is an old denial-of-service attack against computer networks consisting of fragmented or otherwise invalid ICMP packets sent to the target, achieved by using a modified ping utility to repeatedly

send this corrupt data, thus slowing down the affected computer until it comes to a complete stop.

A specific example of a nuke attack that gained some prominence is the WinNuke, which exploited the vulnerability in the NetBIOS handler in Windows 95. A string of out-of-band data was sent to TCP port 139 of the victim's machine, causing it to lock up and display a Blue Screen of Death (BSOD).

Prevention and Response

Firewalls

Firewalls have simple rules such as to allow or deny protocols, ports or IP addresses. Some DoS attacks are too complex for today's firewalls, e.g. if there is an attack on port 80 (web service), firewalls cannot prevent that attack because they cannot distinguish good traffic from DoS attack traffic. Additionally, firewalls are too deep in the network hierarchy. Routers may be affected even before the firewall gets the traffic. Nonetheless, firewalls can effectively prevent users from launching simple flooding type attacks from machines behind the firewall.

Some stateful firewalls, like OpenBSD's pf(4) packet filter, can act as a proxy for connections: the handshake is validated (with the client) instead of simply forwarding the packet to the destination. It is available for other BSDs as well. In that context, it is called "synproxy".

Switches

Most switches have some rate-limiting and ACL capability. Some switches provide automatic and/or system-wide rate limiting, traffic shaping, delayed binding (TCP splicing), deep packet inspection and Bogon filtering (bogus IP filtering) to detect and remediate denial of service attacks through automatic rate filtering and WAN Link failover and balancing.

These schemes will work as long as the DoS attacks are something that can be prevented by using them. For example SYN flood can be prevented using delayed binding or TCP splicing. Similarly content based DoS can be prevented using deep packet inspection. Attacks originating from dark addresses or going to dark addresses can be prevented using Bogon filtering. Automatic rate filtering can work as long as you have set rate-thresholds correctly and granularly. Wan-link failover will work as long as both links have DoS/DDoS prevention mechanism.

Routers

Similar to switches, routers have some rate-limiting and ACL capability. They, too, are manually set. Most routers can be easily overwhelmed under DoS attack. If you add rules to take flow statistics out of the router during the DoS attacks, they further slow down and complicate the matter. Cisco IOS has features that prevent flooding, i.e. example settings.

Application Front End Hardware

Application front end hardware is intelligent hardware placed on the network before traffic reaches the servers. It can be used on networks in conjunction with routers and switches. Application front end hardware analyzes data packets as they enter the system, and then identifies them as priority, regular, or dangerous. There are more than 25 bandwidth

management vendors. Hardware acceleration is key to bandwidth management.

IPS Based Prevention

Intrusion-prevention systems (IPS) are effective if the attacks have signatures associated with them. However, the trend among the attacks is to have legitimate content but bad intent. Intrusion-prevention systems which work on content recognition cannot block behavior-based DoS attacks.

An ASIC based IPS can detect and block denial of service attacks because they have the processing power and the granularity to analyze the attacks and act like a circuit breaker in an automated way.

A rate-based IPS (RBIPS) must analyze traffic granularly and continuously monitor the traffic pattern and determine if there is traffic anomaly. It must let the legitimate traffic flow while blocking the DoS attack traffic.

DDS Based Defense

More focused on the problem than IPS, a DoS Defense System (DDS) is able to block connection-based DoS attacks and those with legitimate content but bad intent. A DDS can also address both protocol attacks (such as Teardrop and Ping of death) and rate-based attacks (such as ICMP floods and SYN floods).

Like IPS, a purpose-built system, such as the well-known Top Layer IPS products, can detect and block denial of service attacks at much nearer line speed than a software based system.

Blackholing and Sinkholing

With blackholing, all the traffic to the attacked DNS or IP address is sent to a "black hole" (null interface, non-existent server). To be more efficient and avoid affecting your network connectivity, it can be managed by the ISP.

Sinkholing routes to a valid IP address which analyzes traffic and rejects bad ones. Sinkholing is not efficient for most severe attacks.

All traffic is passed through a "cleaning center" via a proxy, which separates "bad" traffic (DDoS and also other common internet attacks) and only sends good traffic beyond to the server. The provider needs central connectivity to the Internet to manage this kind of service.

Side Effects of DoS Attacks

In computer network security, backscatter is a side-effect of a spoofed denial of service (DoS) attack. In this kind of attack, the attacker spoofs (or forges) the source address in IP packets sent to the victim. In general, the victim machine cannot distinguish between the spoofed packets and legitimate packets, so the victim responds to the spoofed packets as it normally would. These response packets are known as backscatter.

If the attacker is spoofing source addresses randomly, the backscatter response packets from the victim will be sent back to random destinations. This effect can be used by network telescopes as indirect evidence of such attacks. The term "backscatter analysis" refers to observing backscatter packets arriving at a statistically significant portion of the IP

address space to determine characteristics of DoS attacks and victims.

Research Methodology

It shows the simulation and results of the model which is presented.

Simulation

We use three parameters: the attack rate, the attack duration and the rule processing time. They showed that a larger matching probability (that is to say rules that are easier to match) means a reduced response time. Hence, they encouraged Cloud defenders to put those rules at the top of the rules list so as to increase users' satisfaction. They demonstrated, both analytically and experimentally, a direct correlation between the response time and the number of rules and attack rates. To estimate the cost of their system, they rented 20 VMs from Amazon EC2. In the end, running their clustered firewall turned out to cost 38 US/day and 266 US/week, while keeping in mind that long attacks are extremely rare, given that they are easily detected.

Conclusion

DDoS attacks are rising as a threat. Over the last few years, these attacks have grown in intensity and now have traffic volumes of up to 400 Gbps. These attacks are easy to carry out and do not require great knowledge or access to zero-day vulnerabilities. The duration of the attacks is often just a few hours or even minutes, but this can be enough to inflict a lot of damage at the target site. Currently, amplification or reflection attacks are the most popular attack. These attacks use DNS or NTP servers to amplify the attack traffic by a factor of 50-100 times. This allows small botnets to conduct huge volumetric attacks. Many initiatives can help to protect reflection servers, but there are still more than enough open amplifiers that can be misused. In 2014, we have noticed an increase in compromised Unix servers being used to launch attacks. They are of great interest to the

attacker, since they provide a large bandwidth. DDoS botnets can be rented as a service starting at \$5 for small attacks. Application-layer attacks, which target the Web application, are gaining in importance as well as they are difficult to mitigate. They will become even more important in the future as often, attackers adapt their methods during an attack in an attempt to bypass any short term defense mechanism. In the future, we might see more DDoS attacks coming from mobile devices or even the Internet of Things, but this is currently not happening on a large scale. The motivation of the attacker can vary widely, with hacktivism, profit, and disputes being the main reasons. Considering the ease of conducting large DDoS attacks, Symantec expects that the DDoS growth trend will continue in the future. The likelihood of being targeted by short but intensive DDoS attacks is rising. Some companies try to over-provision bandwidth resources to defend themselves against potential DDoS attacks. However, this arms race is very expensive to win. It is more important to be prepared for DDoS attacks and have an incident response plan ready. Talk to the upstream provider and ensure that they are aware of this threat and check what benefits the utilization of DDoS protection services can bring.

References

1. <https://www.sciencedirect.com/science/article/pii/S1877050915007541> by RV Deshmukh 2015
2. <https://www.cloudflare.com/learning/ddos/what-is-a-ddos-attack/>
3. <https://ieeexplore.ieee.org/document/7524500/>
4. <https://www.cisco.com/c/en/us/about/security-center/guide-ddos-defense.htm>
5. <https://businessinsights.bitdefender.com/amplified-ddos-attacks-are-here-to-stay-exper>
6. https://www.verisign.com/en_IN/security.../ddos...a-ddos-attack-work/index.xhtml