Denial of Service Attacks, What They are and How to Combat Them

John P. Pironti, CISSP
Genuity, Inc.
Principal Enterprise Solutions Architect
Principal Security Consultant

Version 1.0 – November 12, 2001

Where do Denial-of-Service Attacks come from?

A Denial-of-Service (DOS) attack is not only one of the simplest attacks that an attacker can carry out, but also is one of the one of the most effective. The goal of a DOS attack is not to penetrate or modify a computer on a network, but to simply deny access to it from its intended users. DOS attacks have recently gained recognition due to high profile attacks on Internet web sites such as Yahoo and Amazon.com. In reality these kinds of attacks have existed since the inception of wide area networks such as the Internet.

Although DOS attacks may not compromise the data within a particular host, they still can be devastating to the organization that is owns the system. This is particularly evident when you are dealing with hosts with Internet connected hosts that are components of a web environment. If an on-line stock-trading web site becomes unavailable to its user community the financial impact to both the end user and the financial institution operating the site can be tremendous. Every minute that these web sites are unreachable could result in the potential loss of millions of dollars. DOS attacks also can create havoc for the public relations staff of these companies that now need to regain customer confidence in the service that is being offered.

How they are carried out

There are many ways for a malicious attacker to carry out a Denial-of-Service attack. The two most common ways that they are carried out are via the *network layer* or the *application layer*. There is also a distinction between single source and distributed Denial-of-Service attacks. Each of these attacks have only one purpose, to prevent a host on a network from being able to be reached by its intended user community.

3 types of network layer attacks

Network layer Denial-of-Service attacks come in many forms and flavors. The three most common attack types are single user, server, and bandwidth.

In a *single user attack* malformed TCP/IP packets (e.g. fragmented packets) are transmitted to a single host in the hopes that it will cause the receiving host to crash due to a failure in processing the packet correctly. Common tools such as “frag router”, can be downloaded from hacker web sites, and then be used to carry out this kind of attack. These tools do not require the attacker to be extremely knowledgeable about the attack method, and in some cases even have a graphical user interface.

A network level server Denial-of-Service attack occurs when an attacker transmits streams of packets to the host that do not contain meaningful data. The TCP/IP protocol stack on the host will become overwhelmed by the number of packets that it is processing, preventing legitimate user’s packets from being processed correctly by the host.

A common attack of this type is known as a TCP SYN flood. A TCP SYN flood occurs when an attacker transmits large quantities of TCP/IP packets that only contain the SYN command. The SYN TCP command is the first stage of a normal TCP/IP connection between hosts, and is used to establish communication between the initiating host and the receiving host. A host will have a certain number of continuous connections available before it can no longer accept any new connections. Once this limit is reached it will begin to reject any new sessions until it is able to terminate other sessions.

The TCP/IP stack on a given host will also have an established waiting period for connection completion prior to dropping the connection attempt, which will allow another user’s session to be initiated. The attacker will attempt to send enough TCP SYN packets to the host to force it to reach its limit of session initiations that it is capable of processing.

A bandwidth attack occurs when an attacker attempts to flood a site with an excessive amount of TCP/IP data packets with random data payloads. The goal of this kind of an attack is to use all of the available bandwidth at the site where the host exists. This will effectively stop any users within the site from being able to utilize their Internet connection and also prevent any outside users from being able to reach the site. A bandwidth attack can be created by a single host on the Internet, which can set off attacks known as "smurf" attacks that will amplify packets as they travel across multiple routers until they reach their final destination.

Understanding application layer DOS attacks

Application layer Denial-of-Service attacks occur when a host is flooded with large quantities of data. They differ from network layer attacks in that they do not necessarily attempt to disable the entire site, or attempt to utilize malformed or fragmented packets to reach their goal. Application layer attacks occur when large quantities of valid data connections are opened on a particular host or hosts, and are directed at a specific application such as a web server.

Application layer attacks are not necessarily malicious in nature. If a web environment is not properly designed to handle the peak user demand then a Denial-Of-Service attack will result during peak periods of usage. Examples of this can type of attack can be seen following events where web sites are well advertised, such as the Super Bowl.

In 1998 a well-known women's fashion company advertised an online fashion show to be held two days after the Super Bowl. This advertising resulted in an enormous amount of network traffic, well beyond the environments peak usage design limits, being directed at the environment at one time. The result was a limited number of users being able to access the show, and a large number of users being turned away with the infamous "http 404 error" or server not available.

Application layer attacks also have the ability to defeat network and host-based intrusion detection systems (IDS). They do this by sending large amounts of valid data streams which the host will potentially perceive as normal site activity. An IDS system can however help you determine when you are under attack from an application layer Denial-of-Service attack. You can tune both the network and host based elements of the IDS to trigger alarms when TCP connection levels reach specific threshold points as well. This will allow you to further investigate the situation to see if the traffic is valid, or if you are under a DOS attack.

Single Source Vs. Distributed Denial-of-Service attacks

A single source Denial-of-Service attack (figure 1.0) is an attack whose origin is from a single IP address on the Internet. These kinds of attacks can be the result of a simple tool being used to launch a SYN flood, or a packet fragmentation attack. Single source attacks are more annoying than effective because filtering easily defeats them.

Distributed Denial-of-Service (DDOS) (figure 2.0) attacks pose a new challenge to the individuals who are charged with putting countermeasures in place. Distributed Denial-of-Service attacks are attacks that originate from multiple hosts most likely on multiple networks. They can also be originated by groups of attackers who have formed alliances to attack specific environments. They can also originate from hosts which have been previously compromised, and have had attack tools loaded on them without their owner's knowledge or permission.

The now infamous attacks on Yahoo.com and Amazon.com were the result of DDOS attacks. Trojan programs were installed on numerous compromised hosts around the Internet. The attacks were initiated from an attacker on a remote host who directed all of the other compromised hosts and directed them to attack these web sites.

Trojan programs once were little more than simple SYN flood tools that were launched with a remote command. Recently there has been a surge of new smart trojans being created by attackers. These Trojans now include multiple forms of attack capabilities including fragmentation attacks, SYN floods, ping floods, and http floods. They also have intelligence included within them, which allows them to learn when they should rotate their attack technique due to the current one being counteracted.

Leverage your ISP

When you suspect an Internet based DOS attack is in progress it is important to immediately inform your Internet Service Provider (ISP). An ISP can then confirm that the attack is not a malfunction of a router upstream from your node, and that your environment is actually under a DOS or DDOS attack. An ISP can then put temporary filtering in place in their backbone routers to help alleviate the effects of the attack on your environment. They can also begin to trace the attack to its origin in order to take appropriate actions if you request they do so.

It is also important to understand the limitations that an ISP has when assisting you in an attack scenario. There is a common misunderstanding that an ISP will prosecute an individual if they are found to be attacking one of their customers. This is not true. An ISP will usually only assist law enforcement in their investigation if you as their customer decide to have a criminal investigation begun for a specific event. They cannot initiate a legal investigation on your behalf though. They also cannot release information about an attack without a legal subpoena, which requires them to do so.

It is also important to recognize that ISPs can only put filters and access control lists in place within their networks. They can request that other Internet service providers put filters in place in their respective networks, but cannot mandate it. Some ISPs are friendly to these requests while others will not take action unless it is affecting their own customers. This also means that your ISP can only positively trace attacks to the edge of its own networks. In cases where attacks are being carried out over multiple networks it is harder for your ISP to investigate the source address of the attacker.

Intrusion Detection solutions you can turn to

DOS and DDOS attacks can be detected and controlled by the use of network based and host based intrusion detection systems as well as virus detection software. Network-based intrusion detection solutions utilize network sniffers to monitor traffic across specific network segments. They can also be tuned to trigger alarms when they detect activity outside of the normal realm of use for a particular host or set of hosts.

Once the activity is detected the IDS can then generate alerts to warn system administrators of an event, or it can automatically put filtering within the network to help mitigate the effects of the attack. While automatic filtering is an attractive option for overworked network administrators, they are not always the best solution. Automatic filtering can result in denying access to the environment from valid users, which can have an even bigger impact on the environment than the DOS attack itself.

Host-based intrusion detection systems will examine packets once they have reached their intended host. While not as effective as networked based IDS solutions in detecting and preventing DOS and DDOS attacks, they have the benefit of not being defeated by encryption. A host based IDS solution can also detect zombie programs on a host, which act as attack agents in a DDOS attack. These agents will act as packet generators against victim hosts, and are usually controlled by a central program on a separate host.

A host-based IDS can also be tuned to detect large amounts of open TCP connections which are common in SYN flood attacks as well as a abnormal amount of ICMP connections which are common in ping flood attacks. Once the IDS has detected these events it can be programmed to take appropriate actions in the same way that a networked based IDS can. The preferred action is to alert an administrator to review the situation instead of taking any automatic action.

You can also rely on anti-virus software

Anti-Virus software can also be a very effective tool in preventing a host from becoming an attack agent in a DDOS attack. DDOS attackers will usually distribute their attack agents as trojan software to as many hosts as possible prior to launching their attack. Anti-Virus software can detect these Trojans once they have been identified by an Anti-Virus vendor's software and added to their virus signature database. It is important to update the anti-virus software regularly to insure that the databases are current and are detecting as many virus signatures and DDOS trojan programs as possible.

Final thoughts

Denial-of-Service attacks come in many sizes and types but serve only one purpose. They prevent intended users from being able to interact with a specific environment. These attacks are not technically challenging, nor are they considered elite by the underground. Unfortunately, they are some of the simplest and most effective attacks in an attacker's arsenal though. The best defense against them is early detection and proper event management procedures for response. These event management procedures should include filtering mechanisms, software patching, and contact methods for Internet Service Providers.

Your ISP can make the difference between an attack that is perceived as an annoyance versus and attack that renders your environment completely unavailable. They have a wide range of experience with DOS attacks, and have tools and techniques that will help to reduce the impact of these attacks on your environment. Although these attacks may never be eradicated, with proper precautions in place a system administrator can appropriately mitigate the risk Denial-of-Service and Distributed Denial-of-Service attacks pose to their environment.

Figure 1.0

Single Source Denial of Service

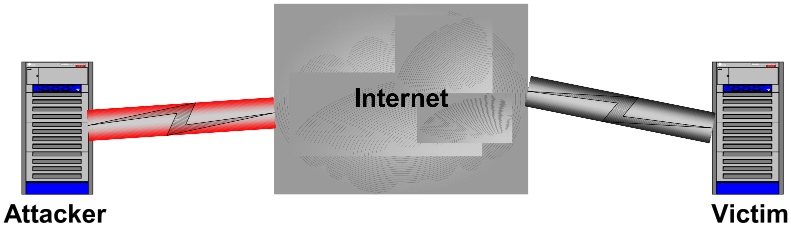

Figure 2.0

Distributed Denial of Service Attack

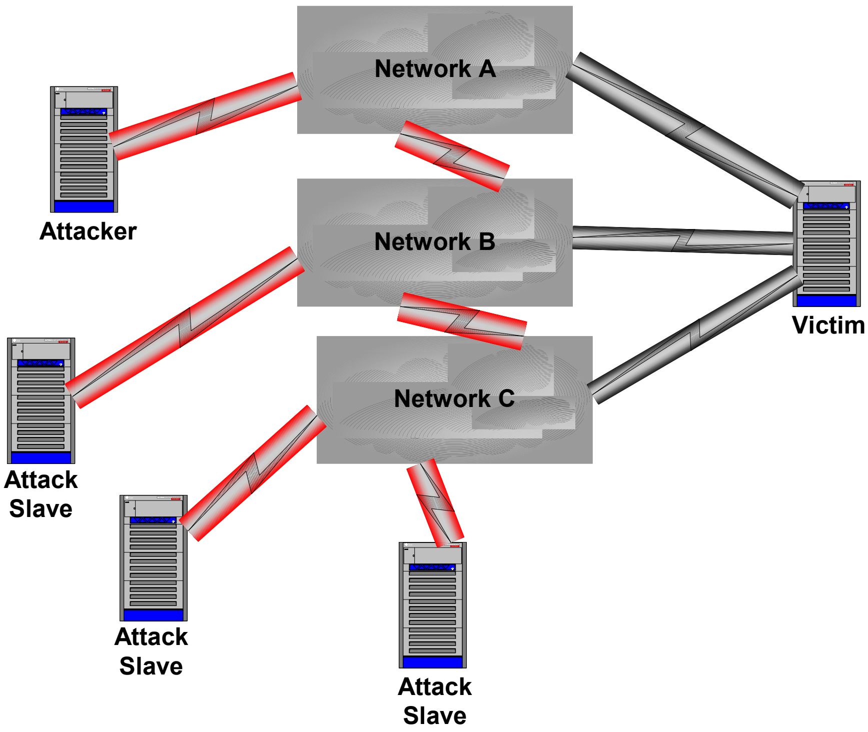
Network A
Attacker
Network B
Victim
Network C
Attack Slave
Attack Slave
Attack Slave