5.2. Denial of Service

A DoS attack is an attempt by a hacker to flood a user's or an organization's system. As a

CEH, you need to be familiar with the types of DoS attacks and should understand how DoS and DDoS attacks work. You should also be familiar with robots (BOTs) and robot networks (BOTNETs), as well as smurf attacks and SYN flooding. Finally, as a CEH, you need to be familiar with various DoS and DDoS countermeasures.

There are two main categories of DoS attacks:

- Attacks sent by a single system to a single target (simple DoS)
- Attacks sent by many systems to a single target (distributed denial of service, or DDoS).

The goal of DoS isn't to gain unauthorized access to machines or data, but to prevent legitimate users of a service from using it. A DoS attack may do the following:

- ❖ Flood a network with traffic, thereby preventing legitimate network traffic.
- ❖ Disrupt connections between two machines, thereby preventing access to a service.
- ❖ Prevent a particular individual from accessing a service.
- ❖ Disrupt service to a specific system or person.

Different tools use different types of traffic to flood a victim, but the result is the same: a service on the system or the entire system is unavailable to a user because it's kept busy trying to respond to an exorbitant number of requests.

A DoS attack is usually an attack of last resort. It's considered an unsophisticated attack because it doesn't gain the hacker access to any information but rather annoys the target and interrupts their service. DoS attacks can be destructive and have a substantial impact when sent from multiple systems at the same time (DDoS attacks).

DDoS attacks can be perpetrated by BOTs and BOTNETs, which are compromised systems that an attacker uses to launch the attack against the end victim. The system or network that has been compromised is a secondary victim, whereas the DoS and DDoS attacks flood the primary victim or target.

How DDoS Attacks Work

DDoS is an advanced version of the DoS attack. Like DoS, DDoS tries to deny access to services running on a system by sending packets to the destination system in a way that the destination system can't handle. The key of a DDoS attack is that it relays attacks from many different hosts (which must first be compromised), rather than from a single host like DoS. DDoS is a large-scale, coordinated attack on a victim system.

The services under attack are those of the primary victim; the compromised systems used to launch the attack are secondary victims. These compromised systems, which send the DDoS to the primary victim, are sometimes called *zombies* or *BOTs*. They're usually compromised through another attack and then used to launch an attack on the primary victim at a certain time or under certain conditions. It can be difficult to track the source of the attacks because they originate from several IP addresses.

Normally, DDoS consists of three parts:

- Master/handler
- Slave/secondary victim/zombie/agent/BOT/BOTNET
- Victim/primary victim

The *master* is the attack launcher. A *slave* is a host that is compromised by and controlled by the master. The *victim* is the target system. The master directs the slaves to launch the attack on the victim system.

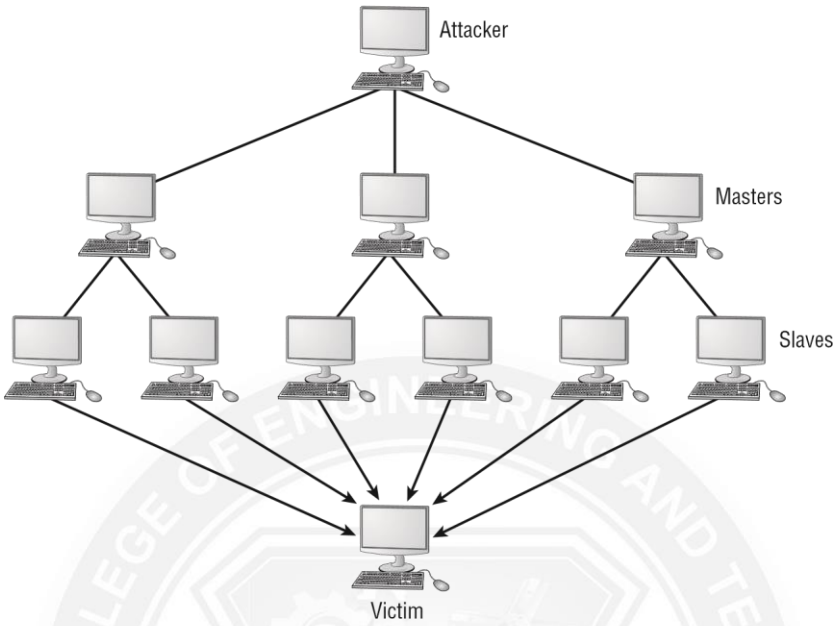

Fig: Master and Slaves in a DDoS Attack

DDoS is done in two phases. In the intrusion phase, the hacker compromises weak systems in different networks around the world and installs DDoS tools on those compromised slave systems. In the DDoS attack phase, the slave systems are triggered to cause them to attack the primary victim.

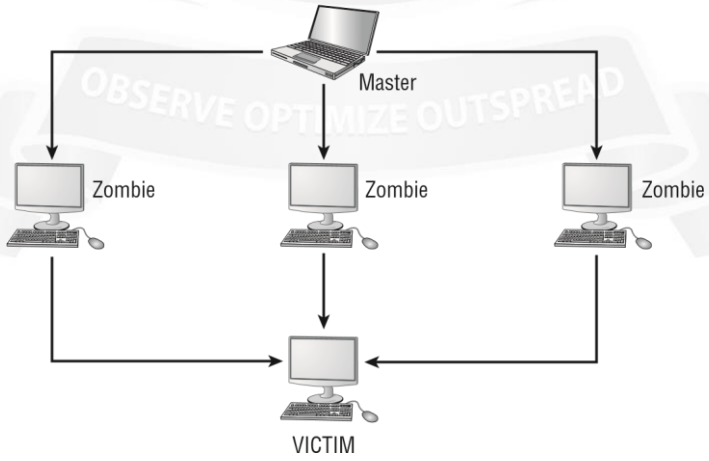

Fig: Bots or Zombie systems

How BOTs/BOTNETs Work

A BOT is short for *web robot* and is an automated software program that behaves intelligently. Spammers often use BOTs to automate the posting of spam messages on newsgroups or the sending of emails. BOTs can also be used as remote attack tools. Most often, BOTs are web software agents that interface with web pages. For example, web crawlers (spiders) are web robots that gather web page information.

The most dangerous BOTs are those that covertly install themselves on users' computers for malicious purposes.

Some BOTs communicate with other users of Internet-based services via instant messaging, Internet Relay Chat (IRC), or another web interface. These BOTs allow IRQ users to ask questions in plain English and then formulate a proper response. Such BOTs can often handle many tasks, including reporting weather; providing zip code information; listing sports scores; converting units of measure, such as currency; and so on.

A BOTNET is a group of BOT systems. BOTNETs serve various purposes, including DDoS attacks; creation or misuse of Simple Mail Transfer Protocol (SMTP) mail relays for spam; Internet marketing fraud; and the theft of application serial numbers, login IDs, and financial information such as credit card numbers. Generally a BOTNET refers to a group of compromised systems running a BOT for the purpose of launching a coordinated DDoS attack. See Figure 7.3.

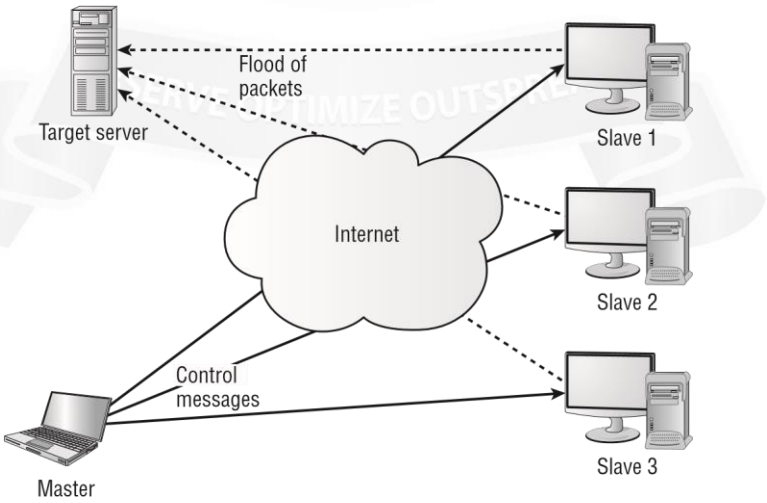

Fig: Anatomy of a Distributed DoS Attack

Smurf and SYN Flood Attacks

A *smurf* attack sends a large amount of ICMP Echo (ping) traffic to a broadcast IP address with the spoofed source address of a victim. Each secondary victim's host on that IP network replies to the ICMP Echo request with an Echo reply, multiplying the traffic by the number of hosts responding. On a multiaccess broadcast network, hundreds of machines might reply to each packet. This creates a magnified DoS attack of ping replies, flooding the primary victim. IRC servers are the primary victim of smurf attacks on the Internet.

A *SYN flood* attack sends TCP connection requests faster than a machine can process them. The attacker creates a random source address for each packet and sets the SYN flag to request a new connection to the server from the spoofed IP address. The victim responds to the spoofed IP address and then waits for the TCP confirmation that never arrives. Consequently, the victim's connection table fills up waiting for replies; after the table is full, all new connections are ignored. Legitimate users are ignored as well and can't access the server.

A SYN flood attack can be detected through the use of the netstat command. An example of the netstat output from a system under a SYN flood is shown in Figure 7.4.

Here are some of the methods used to prevent SYN flood attacks:

SYN Cookies SYN cookies ensure the server does not allocate system resources until a successful three-way handshake has been completed.

RST Cookies Essentially the server responds to the client SYN frame with an incorrect

SYN ACK. The client should then generate an RST packet telling the server that something is wrong. At this point, the server knows the client is valid and will now accept incoming connections from that client normally.

Micro Blocks Micro blocks prevent SYN floods by allocating only a small space in memory for the connection record. In some cases, this memory allocation is as small as 16 bytes.

Stack Tweaking This method involves changing the TCP/IP stack to prevent SYN floods. Techniques of stack tweaking include selectively dropping incoming connections or reducing the timeout when the stack will free up the memory allocated for a connection.

DoS/DDoS Countermeasures

There are several ways to detect, halt, or prevent DoS attacks. The following are common security features:

Network-Ingress Filtering All network access providers should implement networkingress filtering to stop any downstream networks from injecting packets with faked or spoofed addresses into the Internet. Although this doesn't stop an attack from occurring, it does make it much easier to track down the source of the attack and terminate the attack quickly. Most IDS, firewalls, and routers provide network-ingress filtering capabilities.

Rate-Limiting Network Traffic A number of routers on the market today have features that let you limit the amount of bandwidth some types of traffic can consume. This is sometimes referred to as *traffic shaping*.

Intrusion Detection Systems Use an intrusion detection system (IDS) to detect attackers who are communicating with slave, master, or agent machines. Doing so lets you know whether a machine in your network is being used to launch a known attack but probably won't detect new variations of these attacks or the tools that implement them. Most IDS vendors have signatures to detect Trinoo, TFN, or Stacheldraht network traffic.

Automated Network-Tracing Tools Tracing streams of packets with spoofed addresses through the network is a time-consuming task that requires the cooperation of all networks carrying the traffic and that must be completed while the attack is in progress.

Host-Auditing and Network-Auditing Tools File-scanning tools are available that attempt to detect the existence of known DDoS tool client and server binaries in a system. Network scanning tools attempt to detect the presence of DDoS agents running on hosts on your network.