

Denial of Service – Attacks from the Internet (1)

Attacks on company networks endanger business processes, and the number of "denial of service" (DoS) attacks is growing. The first part of this article sketches the goals and effects of DoS attacks and covers some types of attacks. The second half will appear in the next issue. In that article, the author treats additional types of attacks and what should be done to counter these attacks from the Internet.

Attacks on company networks enable unauthorized access to network resources with greater or lesser levels of protection. For example, to gain credit card information, hackers use the Internet to break into a company's customer database. In another example, users within a company manipulate internal applications to gain unauthorized access to protected data.

But this type of attack is not the only danger to company networks. Denial of service (DoS) attacks increasingly pose a danger. Such attacks aim at compromising network services, making them unavailable to other users. This goal involves all the business processes handled by computers, along with all the systems and resources needed for the processes.

Unavailable Systems

During a DoS attack, systems are no longer available for their intended functions. For example, a successful attack on a company's central network computer means that no one in the company can access shared data, such as the customer database or inventory levels. In short, a large portion of internal processing is interrupted, orders cannot be processed, and employees sit at their desks – unable to work.

The danger of becoming the victim of a DoS attack over the Internet is not a theoretical problem, as some recent, familiar attacks made clear. For example, a DoS attack made it impossible to consult the Web site of Consors, a German stock broker, for several hours. Hackers had attacked the Internet server at Consors with counterfeit requests for so long that it no longer had any resources available. Cloud Nine, an English Internet service provider, fell victim to an attack with serious consequences. The attack on its most important systems lasted several months and ultimately led to the collapse of the company. The list of victims reads like a Who's Who of Internet companies. In addition to familiar firms, such as Amazon, Yahoo, eBay, and Buy.com, non-commercial facilities were also attacked. For example during a local election in Hessen (Germany) in 2001, an attack began on the voting computer. Although the attack did not lead to major damage, it shows that attacks can also arise from political motivation.

DoS Attacks and their Results

Of course, many companies can still process work manually by taking inventory directly in the warehouse rather than performing a query on an online system. But companies that do business on the Internet come to a complete standstill. During an attack, customers can no longer connect to a company's systems. The result? No incoming orders, angry customers, and an article in the "breaking news" section of a wire service.

DoS attacks are seldom identified as such within a company network. Few system administrators connect the need to restart a server several times a day with an incoming attack. When a company does discover such attacks, it usually tries not to admit it publicly because it fears damage to its public image.

To implement effective, defensive measures, it's important to recognize various options for detecting DoS techniques.

Overloading Resources

When companies implement network-based applications, such as Internet portals, online warehouses, or company databases, they first perform sizing. Sizing defines the capacity of individual components, such as working memory, processors, disk space, and the performance of network connections. The capacity of the components determines how well the system can perform any tasks assigned to it. No matter how generously a system has been configured, it always has an upper limit to its load. Once the limit is reached, the system cannot handle any more queries.

The expected number of system queries for a given unit of time is an important metric for sizing. For example, if a system is planned and installed to process up to ten thousand queries simultaneously, there are no problems when the system faces an average of only six hundred parallel queries. An attacker, however, can manipulate queries to reach the maximum load and thus block access by other users. For example, an attacker might regularly create such a high number of connections to an online warehouse system that the connections consume all of the network bandwidth available for the online system. Were such an attack to be conducted over a long period, no resources would be available to the actual customers of the system.

Attacks based upon the excessive use of resources usually mean that the attacker has access to a high-performance system. An attacker who connects to the Internet over an analog modem will find it difficult to bring a powerful system (one connected to the Internet with a dedicated, high-speed connection) to its knees by consuming the entire bandwidth available to the system. Nonetheless, resourceful minds have developed techniques that enable attackers with limited resources to initiate successful DoS attacks. This approach employs a third system to attack the target. Such attacks have become familiar as distributed denial of service (DDoS) attacks. Here, a hacker places a software module, called an agent, in several, poorly secured systems on the Internet. The hacker uses control software to instruct the agents to attack the target simultaneously with manipulated data packets. The agents overload the target system and its services fail. A DDoS attack shows the importance of security for publicly accessible systems. Even when these systems do not contain any sensitive, confidential data, they can be misused to attack other systems.

Attacks on Transport Protocols

The TCP/IP family of protocols consists of various protocols that can be combined with each other and has become the standard to handle communications within company networks and over the Internet. The individual protocols include TCP, IP, UDP, and ICMP. In networks, IP creates packets based upon sender and recipient addresses; it functions much like a post office does in delivering mail. To fulfill its tasks, each protocol consists of usage data, such as the contents of a Web site, and a header that contains data specific to the protocol, such as the IP address of the parties attempting to communicate with each other.

DoS attacks do not aim at discovering information or receiving authorization; they manipulate the availability of online resources. The motivation for such attacks ranges from spontaneous actions by frustrated hackers, to intentional damage to a business, and ultimately to electronic warfare and terrorism. Network-based attacks have two advantages not available in the real world: the victims can be reached from anywhere, and there are several ways for the attacker to remain anonymous. In many online attacks, the threshold for engaging in criminal acts seems lower, because the attacker does not need to be physically present at the scene of the crime.

The motivation of DoS-Attacks

Several types of DoS attacks manipulate the contents of the header to trigger a failure of the module that evaluates network data when it attempts to interpret the contents. Since the module is usually implemented deep in the operating system of the server, the failure usually leads to a collapse of the entire system.

ICMP recognizes errors in networks. A "ping" is one element of this protocol. A ping determines if a given system can be reached over the network. The "ping of death" is one of the best-known attacks. For a long time, it allowed an attacker to send manipulated ping packets that brought the target system to a standstill.

TCP is based upon the packet-creation function of IP, and helps network applications transfer data between client and server. An older, but effective type of attack, the "SYN flood" attack, manipulates protocol headers. This attack takes advantage of a weakness in the design of TCP to render a system incapable of communication. Services that communicate over TCP begin with a three-way handshake. The handshake ensures that the sender and recipient actually want to communicate with each other and that the communication is synchronized. Simply put, the client sends a packet with a request for a connection; the server confirms that it wants to participate in the communication. The client then sends a packet to the server and documents the confirmation. Client and server begin the actual communication only after the entire handshake has concluded successfully.

However, data packets traveling along the Internet might experience a delay. Once a server receives a communications request confirmation, it waits a while for the receipt of the third packet. If the third packet does not arrive in a timely manner, the server severs communications, because it assumes that it's dealing with an erroneous transmission.

Author: Oliver Karow - Oliver@greyhat.de

A server can, however, accept several requests for communications simultaneously, so it places the semi-open requests in a queue. The queue is limited to a specific number of requests. An attacker who repeatedly sends communications requests to a server, without ever ending a previously initiated three-way handshake, fills the queue in the target system, which means that it cannot accept any more requests.