Global Information Assurance Certification Paper

Distributed Denial of Service Attacks: Threats, Motivations & Management

GSEC Practical Assignment version 1.4b - Option 1
Tom Simcock
November 5, 2002

Abstract

Distributed Denial of Service (DDoS) attacks are a threat to anyone who relies on the Internet for computer-based business or services. These attacks can deny users access to services by causing Internet based systems to become unavailable and have the potential to cause great economic and other injury. DDoS attacks have evolved from single system attacks, to attacks using distributed networks of machines resulting in a potent attack mechanism. These attacks exploit vulnerabilities in computer systems, protocol behaviour and software logic. DDoS attacks can be launched by anyone due to the advent of easy to use DDoS attack tools, which use automated processes and employ a variety of attack methods. Anyone with the motivation to launch a DDoS attack can download freely available attack tools and launch an attack. DDoS attacks are difficult to stop but certain actions such as system hardening and risk management planning can be used to mitigate their effects. The future holds promise with advances in technology and improved software development methodologies.

Introduction

Distributed Denial of Service attacks have become an increasingly important threat to today's online world. Some businesses are now totally Internet based and any disruption to their online presence can mean a loss in potential revenue and in some cases going out of business (Warner, 2002). According to the Yankee Group, the DDoS attacks of February 2000 caused an estimated cumulative revenue loss of US$1.2 billion to e-commerce organisations including Amazon, Yahoo and eBay (Murphy, 2000). The latest major DDoS attack (October, 2002) was aimed at the thirteen root DNS servers of the Internet in an attempt to cause major disruptions (McGuire & Krebs, 2002). DDoS attacks have the potential to disrupt any computer related service connected to the Internet and therefore should not be ignored.

A denial of service (DoS) is a situation where a computer or network is prevented from providing one or more services that would normally be available to authorised users (TechTarget, 2001). These services could comprise service that an authorised user would have access to, from serving up web pages, making online transactions or using email, through to the availability of life critical systems. A DoS attack is an attempt to deny services to authorised users by creating a situation where computer systems become unavailable for use. If a system is unavailable then authorised users do not have access to the services they may require. A 'distributed DoS' attack

(DDoS) is a DoS attack *en masse* using a large network of computer systems to create a denial of service situation. The denial of service situation occurs because the network or systems are either overwhelmed, or can no longer provide services effectively (NIPC Watch, 2001).

History

Classic DoS attacks involved an attacker using a single system to attack a target as opposed to a distributed multi-system approach used in DDoS attacks. An example of a classic attack is the 'Ping of Death' attack (Figure 1) where an attacker crafts an illegal ICMP (Internet Control Message Protocol) ping packet that exceeds the legal size of 65,535 bytes. The oversized ping packet is then sent to a target system, which would not have a way of reassembling it. As a result, the target system would freeze up or need to be rebooted causing a denial of service situation. These sorts of single system attacks can be thwarted through firewall and router configuration. Once the attackers IP address has been established, rules can be added to drop any packets originating from the attackers IP address (Sans Notes, 2001). In contrast, blocking a single IP address will not stop a DDoS attack, as the attacks may originate from thousands of machines with unique IPs making them significantly harder to defend against.

Figure 1. Ping of Death Attack
(A) Attacker sends illegally sized ICMP ping packet to a target machine.

DDoS attacks deny service by either exhausting bandwidth or network resources, by flooding the target system with massive amounts of data. DDoS attacks can also be used to attack vulnerabilities in software logic as in the 'Ping of Death' example. If the DDoS attack is successful the target system becomes overwhelmed and is effectively taken offline. In addition, some attackers may 'spoof' or forge the source IP address of the computer from which the attack is originates to make it appear to have originated from somewhere else.

What Allows DDoS Attacks to Transpire

Weaknesses in the Internet Protocols

DDoS attacks are possible due to the lack of security associated with the TCP/IP protocols and the large number of easily compromised systems

connected to the Internet. The TCP/IP protocols function well in allowing connectivity between computers, but they were not designed with security in mind. A technical weakness in the TCP/IP protocols is that IP address fields are not validated at any stage, allowing an attacker to forge source IP addresses. A large number of data packets can be directed towards a server and will be accepted as legitimate traffic. A concerted attack will eventually overwhelm a target.

Vulnerabilities in Systems

A number of machines is required to launch a DDoS attack, and there are possibly millions of vulnerable unsecured machines connected to the Internet, making the massing of a distributed army relatively easy (SANS Institute, 1999-2000). These systems can be compromised due to the ever present vulnerabilities found in software and operating systems. Many attackers gain access to these systems using well known exploits that users or systems administrators have left unsecured, where fixes are often available. This reflects the naivety of many users in relation to Internet security and either untrained or negligent systems administrators.

In software development today the demand for features and speed are often a higher priority than security. To make things worse, programmers with little to no training in secure software writing are being employed to code increasingly sophisticated programs (SANS Institute, 1999-2000). Furthermore, studies have shown that programmers introduce about one error for every ten lines of code written (Tepper, 2002). Its no wonder there are new vulnerabilities being found everyday. Current software development procedures need to be enhanced in order to achieve improved security for computer systems.

A logical solution is to secure every computer connected to the Internet. If all systems on the Internet were 100% secure, systems could not be compromised and DDoS attacks would be more difficult to launch. While this would go a long way to solving the problem, it is highly impractical as vulnerabilities are always likely to exist in addition to the many people who don't secure their systems. Many users with poorly secured systems lack the understanding of the potential dangers associated with connecting to the Internet. Some users may desire security but have no idea how to implement it. This lack of understanding leads to systems being compromised and used for malicious purposes.

The Cost of Security

Security is not a top priority for some users including businesses and others cannot afford it. As Allan Paller stated "only the richest can defend themselves against this type of attack, and most of them cannot withstand a concerted attack" (McGuire & Krebs, 2002). According to Paul Vixie, chairman of the Internet Software Consortium, ISPs have not made DDoS security a priority. ISPs have had the ability to filter packets with forged IP addresses through router configuration for a number of years, but have been reluctant to do so because of the additional work and lack of financial benefit (Fisher, 2002).

The number one priority for most businesses is to make a profit. Software companies often sacrifice speed and functionality over security and employ programmers to write secure software even though they have had little experience in secure code writing (Cert, 2001). Software is released to the market quickly, to generate cash flow, without being adequately tested.

Anonymity

The ability to remain anonymous allows an attacker to launch DDoS attacks and often escape detection. If an attacker can remain anonymous then breaking the law will not be a deterrent. The use of spoofed IP addresses can make it very difficult and time consuming to locate points of attack, often allowing enough time for DDoS attacks to succeed. Attackers often launch attacks from stolen Internet accounts increasing their chances of remaining undetected (CERT, 2001).

The Anatomy of a DDoS Attack

A DDoS attack is comprised of four parties:

- Attacker - The attacker and his or her local machine from which they orchestrate the DDoS attack.

- Masters (aka Handlers) - A small number of machines that the attacker controls and directly communicates with.

- Zombies (aka Slaves, Agents or Bots) - Often numbering in the thousands, these machines take orders from the attacker via the Master machines, and actually carry out the attack. Daemon processes are installed on to these machines, which run silently listening for commands from the Masters.

- Target / Victim - The unlucky recipient who must try and prevent their system from being taken offline.

The *attacker* must first mass an army of masters and zombies to carry out the attack (Figure 2). This is accomplished by compromising poorly secured systems and installing daemon or master programs. The *attacker* initiates a scan looking for systems with remotely exploitable vulnerabilities that can be used to gain access into those systems. This process is automated, allowing attackers to compromise and install the DDoS attack programs in less than five seconds. Reiteration of this process using multiple threads means a large number of systems can be compromised quickly, e.g. several thousand / hour (Cisco, 2000). Now that the attacker has assembled the DDoS army, he/she is ready to launch the attack.

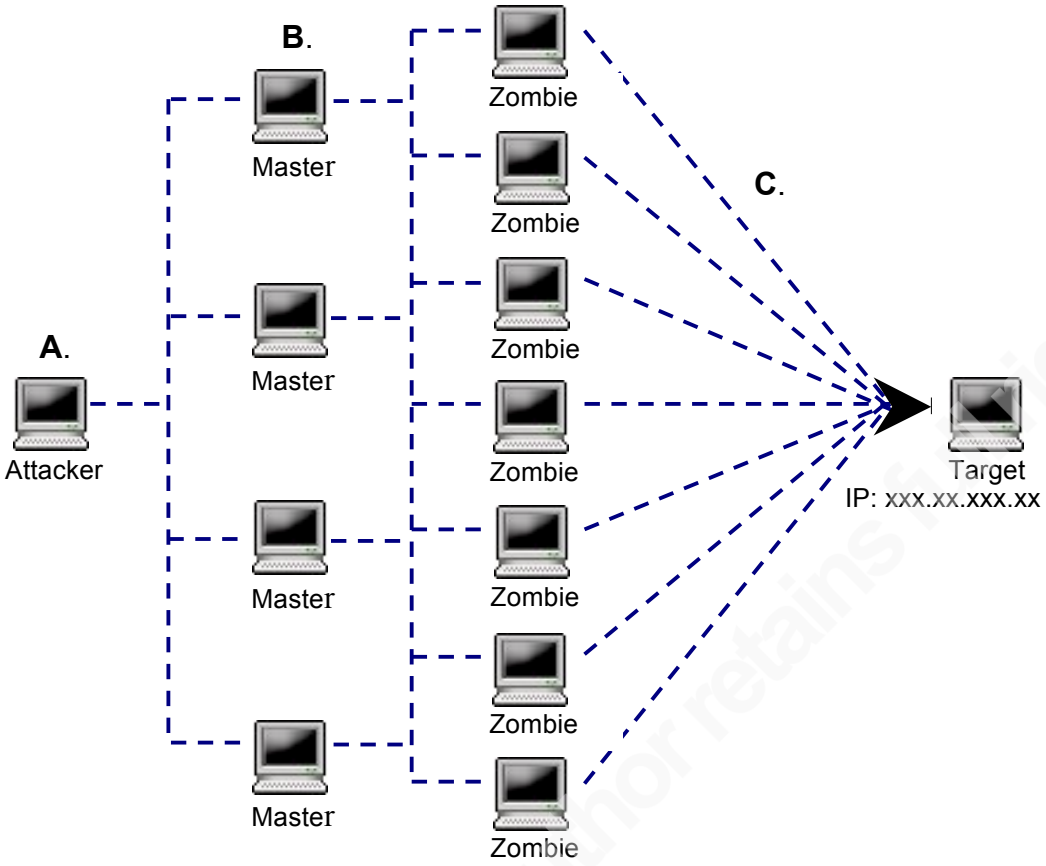

Figure 2. DDoS Attack
(A) Attacker instructs masters to launch an attack on IP xxx.xx.xxx.xx
(B) Master machines relay attackers instructions to the zombie machines.
(C) Zombie machines flood target with data packets.

The *attacker* does not communicate directly with the target. Instead the *attacker* sends out the *target's* IP address to the *master* machines, which instruct the *zombie* machines to initiate the attack. The *zombies* then flood the victim's network/system with massive amounts of unwanted traffic in an attempt to overwhelm the *target* and if successful, the system/network will be taken offline.

DDoS Attack Methods

There are many ways to create a denial of service situation. Following are some of the common methods used in DDoS attacks.

Bandwidth Attacks

Bandwidth attacks use UDP, ICMP or TCP to bombard servers or network equipment (routers, firewalls, etc.) with large numbers of packets at high packet rates (measured in packets per second). The resources of the server or network equipment become overwhelmed creating a situation where services can no longer be provided (CERT, 2001).

Looping UDP (User Datagram Protocol) Ports Attack

This attack causes a denial of service by consuming CPU resources using two targets. A spoofed UDP packet is sent to the echo service port of target 1 (Figure 3) appearing to be from the chargen service (character generator) port from target 2. A loop develops because both services respond to the others data. The chargen service from target 2 keeps replying with streams of data to the echo service of target 1, whose echo service keeps responding with an echo reply. The UDP traffic keeps bouncing back and forth consuming system resources, preventing services from being provided.

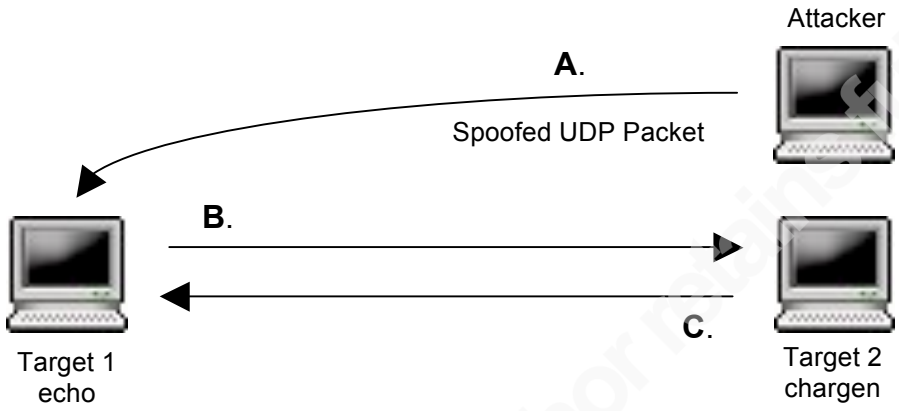

Figure 3. Looping UDP Ports Attack.
(A) Attacker sends forged UDP packet to echo service of target 1 appearing to come from target 2.
(B) Echo service responds with an echo reply addressed to target 2.
(C) Chargen service of target 2 receives echo reply and sends back its own data packet.
A loop develops between target 1 and target 2 consuming system resources and preventing either system from doing anything useful.

ICMP Attack (or Smurf Attack)

This attack floods a target by sending a number of ICMP echo requests (ping packets) to an IP broadcast address and forges the source address to create the illusion that they are originating from the intended victims machine. If enough servers listening to the IP broadcast respond with an echo reply, a large amount of traffic is sent to the target machine and can cause a denial of service. A similar attack (known as a Fraggle attack) uses the same attack method but uses UDP packets instead of ICMP.

TCP-SYN Attacks (Half-Open Attack)

TCP-SYN attacks are targeted at systems providing services like web, mail and file servers. TCP *(Transmission Control Protocol)* is a connection-orientated protocol, which is used to establish connections between two computers (Figure 4). In order for users to use services via TCP, connections must be established between the users machine and the server providing the service. This connection is established by what is known as the 'three way handshake'.

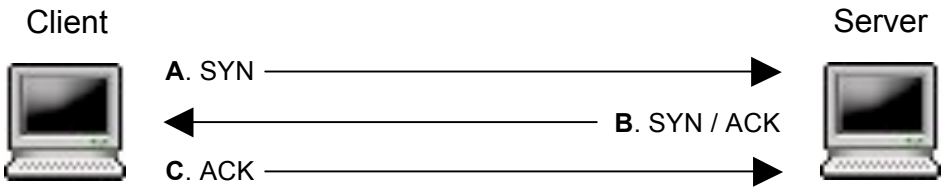

Figure 4. The TCP Three Way Handshake
(A) A client machine sends a SYN (connection request) packet to the server.
(B) The server responds by sending back a SYN/ACK (connection request / acknowledgment) packet.
(C) The users machine sends back a SYN packet to complete the connection.
(SANS Notes, IP Concepts 2, 2002).

In a DDoS TCP-SYN attack, compromised machines are directed to flood the target with TCP SYN packets. The SYN packets have spoofed IP source addresses appearing to come from an unreachable host. The victims machine reserves system resources for each connection request but the 'three way handshake' never completes because the unreachable host never responds with the ACK packet. Once the server's resources fill up with incomplete connections, the server can no longer serve new requests, creating a denial of service.

Tools of Mass Destruction

There are a number of publicly available DDoS attack tools (Table 1) that are capable of the sorts of 'attack methods' mentioned above. No doubt, there are also a number of privately held attack tools, kept hidden to conceal their source code, methods, and attack patterns. DDoS capable tools (created for non-malicious reasons) have been commercially available for some time for the purposes of testing a networks bandwidth capacity and the reliability of services under heavy traffic (Mixter, 2002). DDoS attack tools started appearing publicly ca.1997 (Navratilova, 2000). These tools automated the process of launching a DDoS attack and made it possible for anyone to use. With a DDoS attack tool, "nothing more than the whim of a 13-year old hacker is required to knock any user, site, or server right off the Internet," stated Steve Gibson (2002), reflecting on his website being taken offline by an angry thirteen year old.

Before DDoS attack tools were available, attackers would have to telnet (or remotely login) into each machine that was going to be used as a part of the DDoS attack and manually launch a flooding tool against their target. This manual process could take some time, so to speed up this process attackers wrote automated DDoS attack tools. "DDoS tools just make an old concept easier" (Mixter, 2002, author of two DDoS attack tools) by automating the process of compromising and controlling systems and being able to launch an attack with a single command.

Current DDoS tools are becoming more sophisticated than ever, integrating automated processes like information gathering, remote control and automatic

update functions (Hyunwoo, 2001). Some DDoS attacks tools are using encryption for the communications between the DDoS controlled systems and the attacker in an attempt to remain discrete. Other technologies like viruses, Internet worms and Trojan horses are also being used as mediums for propagating DDoS control programs onto vulnerable machines.

DDoS Tool	Attack Methods	Employs IP Spoofing	Uses Encryption
Trinoo	UDP	no	no
TFN	UDP, ICMP Echo, TCP SYN, Smurf	yes	no
Stacheldracht & variants	UDP, ICMP, TCP SYN, Smurf	yes	varies
TFN 2K	UDP, ICMP Echo, TCP SYN, Smurf	yes	yes
FAPI	UDP, ICMP, TCP SYN, TCP ACK	yes	no
Shaft	UDP, ICMP, TCP SYN	optional	no
Mstream	TCP ACK	yes	no
Trinity	UDP, TCP(Fragment, SYN, RST, RandomFlag, ACK), Establish, Null	-	no

Table 1. Various known DDoS tools and their attack methods (Adapted from Riverhead, 2002).

Motivations Behind DDoS Attacks

What motivates a person to launch such attacks?

There could be any number of reasons why someone would want to launch a DDoS attack against another party or organisation. Some general motivations of attackers can be categorized using the following categories:

Script Kiddies, Packet Monkeys and Ankle Biters:

These derogatory terms are some of the names used to describe those with little to no technical prowess, often associated with young teenagers, these attackers like to try out easy to use hacker programs and scripts without really knowing what they are doing. Their motivation may be to impress their peers or just muck around and often lack an appreciation for the damage they cause (TechTarget, 2001).

Information Warfare:

In a desire to gain a strategic advantage over military or business adversaries one may be motivated to indulge in a DDoS attack as a means of information warfare. Information Warfare refers to the offensive and defensive use of information and computer based systems to exploit and deny an adversary's computer based services and systems (Adapted from Goldberg, 2002).

In a DDoS attack information is being used to the attacker's advantage. Hackers create DDoS tools from their knowledge of the weaknesses in the TCP/IP protocols and known exploits in poorly secured systems. They can also incorporate IP source spoofing as a means of defence. The DDoS attack uses computer-based systems (masters and zombies) as the means to carry out the attack in an attempt to deny the adversary of their computer / information systems and services. The users of the DDoS attack tools only need to know where to get the tools and how to run them.

In order to gain an edge over a competitor a rival company could launch a DDoS attack against its competitor. The motivation could be to damage the image of the opposing company or damage its business in an attempt to secure more of a market share. If a company cannot provide services to its customers, then those customers may take their business elsewhere. Other companies whose machines had been used as participants in the DDoS attack may also suffer. Some people will feel reluctant to entrust a company whose machines have been hacked into, out of fear that their investments and personal data could be compromised. In relation to national security, attacks on critical communications and computer dependant critical infrastructure are a major concern and must be protected against.

Disgruntled Employee:

The motivation of a disgruntled employee is usually revenge. The employee or ex-employee (as is often the case) bears a grudge against their workplace due to having lost their job or some other perceived injustice and feels that they would gain some degree of satisfaction by punishing their company.

Hacker/Political:

These individuals or groups are out to make a point. Those who strongly disagree with certain comments that have been published on a website or what a company stands for may attempt to launch a DDoS attack in order to take the website offline.

Unintentional:

It may be that a huge number of people are visiting a website, perhaps for a world sports event or calamity and too many requests are made of the web server(s). The webserver becomes overwhelmed by the large number of requests and can no longer service new clients

Defending Against Distributed Denial of Service Attacks

What Users Can Do:

1) Harden System

Keeping systems secure requires a continued focus. Users can and must harden their systems so they cannot be compromised and used in a DDoS attack.

Recommendations:

- Scan systems for the existence of known zombies and remove them. There are many tools available to perform these tests.
- Keep systems updated with the latest patches and anti-virus signatures.
- Create hardening scripts.
- Run a vulnerability scanner such as Nessus to help identify areas that need to be tightened up.
- Turn off all services that are not needed.
- Personal firewalls like Zone Alarm (by Zone Labs) give users control over what programs have access to the Internet (egress filtering) at the host level. This could potentially stop the system from being used in a DDoS attack.

2) Plan for Disaster – Risk Management

How long can a person or a business afford to have their online capabilities go offline? Hardening systems may prevent them from being used as masters or zombies, but it will not mitigate the chances of being targeted by an attacker. Therefore a risk management strategy needs to be devised before being attacked or the chances of dealing with the attack will be limited. The CERT® Coordination Center provided some useful ideas that can be helpful in designing a strategy. These are can be broken into three areas: protecting, detecting and reacting.

Protecting

- Identify mission critical assets that need to be protected.
- Identify the interdependencies of different services; a critical asset may rely on another service for its proper functioning, making the dependant asset just as critical.
- Invest in additional system capacity, which could be used to help absorb an attack (e.g. additional bandwidth and servers). Using additional webservers that all contain the same content coupled with a load-balancer, overwhelming a single server becomes much more difficult. When one webserver nears overload and load balancer simply passes requests to the next webserver in the line (Savetz, 2002).

- Design network / systems to be able to degrade or disable non-critical services during an attack, to try and keep critical services functioning during the attack.
- In some situations the best option may be to shut down all services for the duration of the attack.
- Filter unwanted traffic by using ingress and egress filtering.

Detecting

- Establish a baseline for what is considered normal traffic. This baseline can be helpful in detecting abnormal traffic patterns and alerting the user that he/she may be under attack.
- Intrusion detection systems can be useful in alerting the user if there is unusual traffic. A downside is that during a DDoS attack such systems can cause bottlenecks by producing a large number of logs and consuming system resources.

Reacting

- Develop a relationship with the ISP. Organise someone to contact at the ISP during an attack and decide the specific steps that the ISP will take to help minimize the attack.
- Find out what measures the ISP has taken to help protect, detect and react to DDoS attacks, including measures taken to disable IP spoofing attacks and denying traffic from broadcast / multicast addresses.
- Create a checklist or set of actions to be followed in the event of being attacked.

(CERT, 2001)

The degree to which measures can be implemented will vary depending on the available budget for security.

Future

According to security technologist Bruce Schneier, the solution to DDoS attacks may rest in the redesigning of the Internet, which would be a major undertaking but a possibility for the future (Sage, 2002). For now though upcoming technologies may help in the fight against DDoS attacks. The next version of the Internet Protocol IPv6 offers a number of useful features that may mitigate some of the threat. These features include ECN (congestion control), IPSEC (authentication, integrity and confidentiality of IP packets) and the development of trace-back technology to help identify an attackers origin. These technologies will reduce traffic congestion, prevent IP address spoofing and increase the general level of IP packet security (CERT, 2001).

If secure software development procedures are employed in future software releases, and general security awareness can be increased in the Internet community, then the number of vulnerable computer systems connected to the Internet may be reduced. However, new and improved attack tools will

continue to be developed by hackers and new threats will have to be addressed.

Conclusion

There are a myriad of sophisticated DDoS attacks tools and methods to wreak havoc on the Internet. These tools have become available for anyone to download and use. If a system is connected to the Internet, there is nothing to stop it from being targeted by an attacker. However, through risk management and tightening organisational systems, the effects of such attacks can be lessened. A company's legal liability may also be lessened, by ensuring that its systems are not used as zombies in a DDoS attack. The future holds hope in technological advances but as long as the motivation exists, the discovery of vulnerabilities and ways to exploit them will always remain. Therefore, the adoption of a security-conscious attitude is imperative in order to stay on top of the evolving threats of attackers.

References

CERT® Coordination Center. “Managing the Threat of Denial-of-Service Attacks.” October 2001.
URL: http://www.cert.org/archive/pdf/Managing_DoS.pdf

Cisco. “Strategies to Protect Against Distributed Denial of Service (DDoS) Attacks.” 2000.
URL: http://www.cisco.com/warp/public/707/newsflash.html

Fisher, Dennis. “ICANN Targets DDoS Attacks.” October 28, 2002.
URL: http://www.eweek.com/article2/0,3959,651686,00.asp

Gibson, Steve. Gibson Research Corporation. 2002.
URL: http://grc.com/dos/grcdos.htm

Goldberg, Dr. Ivan, Institute for the Advanced Study of Information Warfare, 5/1/2002.) URL: http://www.psycom.net/iwar.1.html

Lee, Hyunwoo. “Korea Opportunities: Analysis on recent Internet worm & malicious agent.” 2001.
URL: http://www.securitymap.net/sdm.docs.virus/Virus-Bulletine.txt

McGuire, D & Krebs, Brian. “Attack On Internet Called Largest Ever.” 2002.
URL: http://www.washingtonpost.com/wp-dyn/articles/A828-2002Oct22.html

Mixter. “Tribe Flood Network 3000.” 2002.
URL: http://mixter.void.ru/tfn3k.txt

Murphy, Dave. “DDOS Outages Cost $1.2 Billion.” February 14, 2000.
URL: http://itrain.org/itinfo/2000/it000214.html

Navratilova, Viki. “A Brief History of Distributed Denial of Service Attacks.” Blue Meteor. August 22, 2000.
URL: http://www.uniforum.chi.il.us/slides/ddos/tsld011.htm

NIPC Watch. “Overview of Scans and DDoS Attacks Executive Summary.” 2001. URL: www.nipc.gov/ddos.pdf

Riverhead Networks. “Known DDoS Tools.” 2002.
URL: http://www.riverhead.com/library/tools.html

Sage, John. “Distributed Denial of Service Attacks.” 2002.
http://www.finchhaven.com/pages/computers/ddos.html

SANS. “Help Defeat Denial of Service Attacks: Step-by-Step.” 23/3/2002.
URL: http://www.sans.org/dosstep/index.htm

SANS. “IP Concepts 2.” SANS Course Notes. 2002.

SANS. “Information Warfare.” Sans Notes pg 5-18, 2000.

SANS Institute. “Consensus Roadmap for Defeating Distributed Denial of Service Attacks.” 1999-2000. URL: http://www.sans.org/ddos_roadmap.htm

Savetz, Kevin. “Managing Traffic Spikes.” New Architect. 2002. URL: http://www.newarchitectmag.com/documents/s=7652/na1102b/index.html

TechTarget. “distributed denial-of-service attack.” Jun 03, 2001. URL: http://searchsecurity.techtarget.com/sDefinition/0,,sid14_gci557336,00.html

TechTarget. “script kiddy” Jul 31, 2001. URL: http://searchsecurity.techtarget.com/sDefinition/0,,sid14_gci550928,00.html

Tepper, Michelle. “Why Software Still Stinks.” Networker, Volume 6.3, September 2002.

Warner, Bernhard. “Hacker attack shuts down British ISP CloudNine.” 2/1/2002. URL: http://news.cnet.com/investor/news/newsitem/0-9900-1028-8672877-0.html