

Ahmed Koshok

CS756, Spring 2004

Project report: A survey of denial of service attacks and countermeasures

May 2, 2004

Abstract

As the proposal suggested, this project is practical survey of denial-of-service attacks and known countermeasures. The class presentation summarized the initial investigation's findings, critical for which were references 1-5. This report summarizes such findings along with continued research into attacks targeting the Internet's infrastructure along with recent DoS technology and exploitations. As to not stay only practical, there are selective analytical dives into certain topics. The main tone for the report is a practical one. We make some conclusions from patters we observe.

Introduction

Denial-of-service (DoS) attacks increasingly gained fame over the past few years. As the Internet becomes more ubiquitous, the threat of the denial-of-service attacks becomes more real and important for individuals, businesses, governmental organizations, and even countries. This is true for two main reasons. First, the Internet's resources are finite and its security model is interdependent.

It is important to note that the most common association made with DoS attacks is that of bandwidth attacks, where victims' system(s) are flooded with packets that exhaust finite resources such as CPU, network capacity or memory. Such attacks typically rely on known problems first exposed using logical DoS attacks. We provide a brief taxonomy of DoS attacks and practical defense pointers.

Further, we look at DoS trends and try to observe patterns in order to make some predictions on their evolution and perhaps even preemptive recommendations for defense.

Finally, we do some analysis on a very recent (late April, 2004) and potentially devastating TCP implementation flaw to the Internet's infrastructure and further analyze the prevalence of DoS attacks on the Internet.

DoS Taxonomy

By definition a Denial-of-service attack is any explicit attempt by an attacker to deem a resource or service unavailable for legitimate users. We can identify an attack as logical or exhaustive. A logical attack is one where the attacker uses precise methods to disrupt service. An example is the so-called 'Windows NT ping of death' attack, which exploits software flaws, knocking out entire systems as a result. An exhaustive attack on

the other hand relies on overwhelming target(s) for extended periods of time to make them unavailable, or perform poorly.

Interestingly, addressing flaws exposed to logical attacks does not eliminate them from becoming exposed to exhaustive attacks. The case in point is the TCP SYN attack, which remains a popular attack method despite fixes to the initial TCP implementations for the flaw permitting the attack. This is true understandably as there are no complete solutions to most discovered flaws. The same is true for UDP floods and ICMP echo floods. The latter is also referred to as amplification attack as the attacker relies on other hosts to continue the attack from him via broadcast protocols. An interesting variation of such attack is that of TCP SYN attack, where an amplifier replies multiple times to a SYN packet to the IP address of the a target [3], or worse broadcast the requests to multiple target, possibly including its own network as in the case of a Smurf attack.

A DoS attack can take on a physical nature. I.E. an attack may involve the physical destruction of equipment or property [2], such as a destruction of a routers or disruption of power.

Attacks can further be understood to use system break-ins to modify configurations, delete resources or take systems off-line. This characterization is not commonly associated with DoS attacks although.

Attacks can be classified as single source, or distributed. The latter case is also referred to as DDoS (Distributed Denial-of-Service) attacks. An interesting variation of which is the DRDoS (Distributed Reflection Denial-of-Service) attack.

An attack does not have to only be of one type. In fact, the most effective attacks diversify their methods as using a combination of logical, exhaustive and even physical methods. The attack with the richest arsenal of tools and tricks is the most effective.

DoS Tools and Trends

DoS attacks began with the attackers carefully choosing the machines from which they will conduct their attacks. The platform of choice was Unix and machines the attacks involved many long manual tasks and planning. DDoS attacks required even more work and coordination. Targets were also carefully analyzed for exploitations in order for the attack to be effective. [1]

The single most consistent trend is the shift from single source to multiple source, aka DDoS, attacks. DDoS attackers show a steadily increasing degree of sophistication in terms of deployment, propagation, immunity and degree of stealth of their attacks and the tools used. Figure 1 shows the phases of a DDoS attack. Their impact is also increasing to the point where the effects of an attack towards a certain site are felt in neighboring hosts. I.E. an attack's blast zone is getting larger. The growth in the amount of bandwidth attackers are able to employ in attacks is larger than the growth of available bandwidth [1]. In essence, attacks are becoming more effective. Figure 2 compares a DoS and a DDoS attack.

There is a shift toward using Windows as a platform of choice for the so-called "zombies", compromised machines used by attackers to flood targets. There are increasingly more unprotected Windows based machines with high-speed broadband connections. Windows provides a larger pool for attackers to exploit than the traditional Unix based systems. Further, newer versions of the Windows operating system are more 'attacker-friendly' in terms of the available APIs to manipulate network layer functionality [4].

Attackers increasingly automate the discovery of vulnerable machines from which attacks may be conducted. There is further automation in the exploitation of the vulnerabilities discovered and the deployment of the necessary agent logic to make a machine perform commands from the attacker, i.e. to become a Zombie. Exploiting flaws in a compromised machine's networks, or via social engineering, attackers can carry out propagation for DDoS attacks. It has been noticed that the success of social engineering has been effective to the point that it alone can cause significant degradation in network performance before even an attack starts. The case in point is the Love letter worm. In a sense this is 'problem' for attackers as the propagation phase has lost its stealthiness.

The final phase of a DDoS attack is the instruction of an agent to execute the attack. This area has also seen significant progress as attackers continue to use more stealth and resilient means. Initially, DDoS attacks used open sockets with known hosts and control channels. Later communication became encrypted and finally relied on IRC-

based networks. As such, discovering a node in a DDoS network does not yield discovery of other nodes.

Attacks show an ever-decreasing time to exploit. That is attackers are able to use a security hole faster than the vendor can adequately address it and make it available to its users. At the time of this writing the Sasser virus showed to be an anecdotal reminder.

A more disturbing trend is the use of legitimate traffic by attackers to simply overwhelm the targets. Previously attackers continuously spoofed the source IP of packets, and consistently varied the IP headers so as to confused routers from filtering packet streams. ISPs now employ filters on routers in order to protect their customers from acting as effective Zombies. Evidently this was an effective, yet an ephemeral solution.

Practical Defense Mechanisms & Countermeasures

For an individual, business owner or an organization, the first step is to acknowledge that the DoS attacks are a real threat. The second step is to realize that such attacks are not completely avoidable, and when they arrive they are not survivable as demonstrated in [3], [4] and [5].

A fascinating account of a determined administrator to defend his site is available in [3] and [4]. It is not reasonable to assume a similar kind of tenacity and knowledge for average administrators. However, we clearly see how the administrator understood the two steps above.

With such basic acknowledgment CERT [5] makes the following recommendations:

- Design for survivability. I.E. have a plan when your systems scum to a DoS attack.
- Take steps to ensure critical services continue in spite of attacks.
- Be a good netizen (net citizen)

The most important recommendation is the third one as if well implemented it could eliminate one of the fundamental reasons for DoS attacks, interdependent security.

The targeting of the Internet's infrastructure

Continuing the survey, I researched trends in DoS technology specifically targeting the Internet's infrastructures components. Reference [6] provides a good summary of attacks on routers and router specific protocols. Routers are good targets for attackers because:

- Compromises are much harder to find and fix
- Compromised devices can be used to further amplify attacks
- Routers tend to be softer targets, as they are not monitored as user machines.

Most of the security concerns shown are addressed by turning off protocols after they are no longer needed. A router can be made safer by being configured correctly and constantly updated. Quite a bit of the holes in routers need local network access, which stresses the importance of perimeter security. I.E. Even if a router is not publicly accessible, does not mean it will remain safe. As such, most logical threats on routers do not pose an immediate and critical concern.

However, there is continued research in malicious hackers circles in exploiting more advanced routing protocols. While there has been no high-profile successful attack based on such ideas, one can imagine the effect once such exploitations are applied to core Internet routers. Sadly, the presumption is that it is a question of "when" rather than "if". Attempts have been made to disable core Internet services before, namely the DNS system. And, [3] shows hackers awareness and brilliance in using the core Internet routers for a reflection amplification attacks, Figure 3.

A recent advisory (late April, 2004) in [7] was too interesting to ignore for this project. The threat in such advisories has been known for about 20 years but widely dismissed as being impractical to implement, which proves to be a false assertion [9].

It is a weakness in most TCP/IP implementations of ISN number generation, which could allow an attacker to terminate, hijack, or outright spoof a TCP connection by guessing a sequence within a window. For the attack to work, the attacker needs to know 2 ends of a TCP connection. Then, the attacker needs only to spoof TCP packets, acting as either end. Doing that requires gathering some information about the state of the connection and depending on the operating system, previous connections made at the

ends. While this attack can be applied to any TCP connection, it is of particular concern to ISPs.

BGP (Border Gateway Protocol) sessions, by which routers communicate, are long TCP connections with end points that are usually known. This makes them particularly vulnerable [8]. Luckily, the BGP protocol allows devices the use of an MD5 digital signature as an optional implementation feature of RFC2358. Without the use of such layered technology in TCP to insure the authenticity of the TCP packets, routers are particularly vulnerable to such attack. The ramifications of such attacks could be disastrous [7].

There are a few immediate conclusions here. First, we again learn there can be no 100% safe protocol. Second, we learn that the diligent use of multiple lines of defense, certainly optional ones, is critical to build secure systems. Such layered approach is also known as hardening. Unfortunately, the same cannot be said to all protocols other than BGP. So while we acknowledge that IPSec addresses much vulnerability in TCP headers, we understand that applying it is not practical from either an administrative or a performance perspective. Given the alternative of an unusable Internet, most will sacrifice performance for security vs. imminent threats.

RFC1948 provides guidelines on how to make ISN numbers harder/impractical to guess. However, some contend that it is still theoretically possible to guess an ISN even with the RCF. So while routers are protected for now [9] by using the MD5 extension in RFC2358, we look into this BGP vulnerability further [8].

We assume that an attacker wants to reset a BGP connection. He needs to know the source and destinations IPs. This is no problem. He needs to know the source and destination ports. The destination port is not a problem. And with a packet sniffer, nor is the source port. Next, the attacker needs to guess a sequence number within a window. As such, the number of needed guesses to get one acceptable packet is $2^{32}/\text{window size}$. One packet is enough to rest a connection, as each guess requires a packet. We can determine the amount of time needed to exhaust the search space based on the number of packets per sec an attack is cable of [8]. I.E.:

$$\text{Time to REST} = (2^{32}/\text{window size}) / \text{packets per second}$$

As presented in [8], this formula is not complete! There is a missing component. And that is the possibility that the window is no longer valid. We assume exponential distribution for the time at which the attack begins in a life of a window T . Then the time of opportunity for an attack = Window timeframe – T . After such time, the window is no longer a valid target. As such, the search space for attacks grows inversely proportional to window life span. The threat is valid as presented if the time to exhaust a search space is less than a window life span. Simple attacks from low bandwidth sources may not be as successful as claimed. Nor are more complex attacks. The author over estimates the threat.

Prevalence of DoS attacks on the Internet

The authors in [10] present a technique called backscatter analysis for detecting DoS attacks on the Internet. The idea is quite a simple one; DoS attacks spoof the source IP address of packet floods such that the receiver is not able to trace packets back to the attacker. Further, the receiver cannot filter incoming packets. Assuming random distribution of the source IP address, with enough servers listening in for TCP ack responses, one can determine if hosts are being attacked.

The authors assume one response for each packet sent with a spoofed IP address where realistically more than one response is generated as the victim machine TCP implementation thinks the source establishing the message failed to receive its response, so it tries up to 4 times in most implementations. This is a key factor to most of their work. If we apply it at the load of an attack to a victim, we could compute a load 4 times the observed load. That is, of course, assuming reliable delivery. Which in a case of a victim's network is a rarity. As such, there is a certain underestimation of the load on victims through their analysis. The authors acknowledge this to a degree.

Another assumption the authors make is that the spoofed IP addresses follow random distribution. Aside from the ingress filtering, the authors do not take into account how clever malicious hackers could black list known server from the spoofed IP address range in order to be stealth. Or, if hackers deliberately forward packets to the host monitoring DoS attacks if such a system becomes a reality.

A final observation is that the backscatter analysis will not be an accurate measure of the prevalence of DoS attacks on the Internet as attackers are predictably using legitimate traffic through Zombies. The backscatter method underestimates the volume and number of DoS attacks.

Summary and Conclusions

As a survey this project is; the goal is to better understand the topic rather than to propose solutions or theories. Despite that we can make a few conclusions.

First, more often than not, the threats of DoS attacks and exploitations are over estimated. It is very difficult to get information about traffic on the Internet, so there is a temptation to assume things are worse than they are. This is true partly because of the point-to-point nature of the Internet and, more importantly, because of the unwillingness of ISPs to share data.

Second, there is an increasing sophistication to the techniques of the attacks that is much faster than the increase of awareness of the average Internet user. With the growth of the Internet, this trend is expected to continue.

Some argue that the Internet works because no one is in charge. In fact, the Internet users are in charge. And without accountability for the Internet's users actions, vulnerabilities such as DoS attacks are inevitable. However, the establishment and enforcement of rules on the Internet questionable at best, at least from a legal perspective. Security, usability and performance remain competing requirements to any system. Despite its imperfections, the Internet reaches the right balance, or is at least close to it.

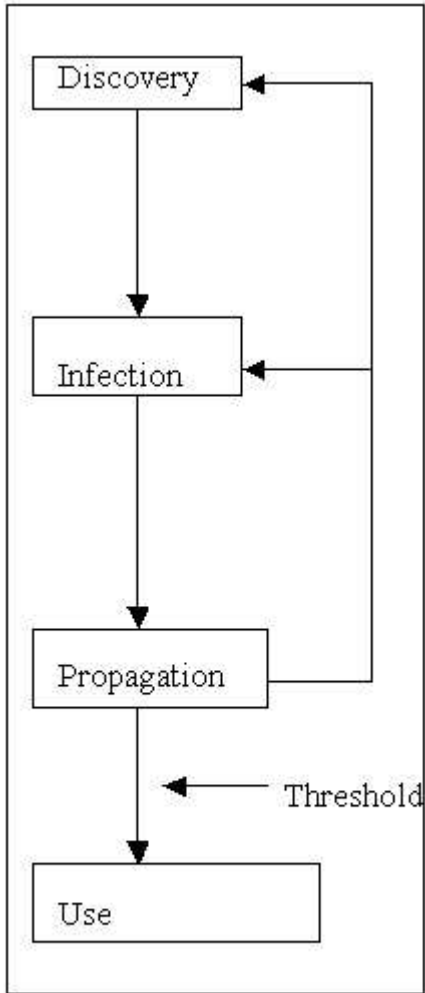


Figure 1

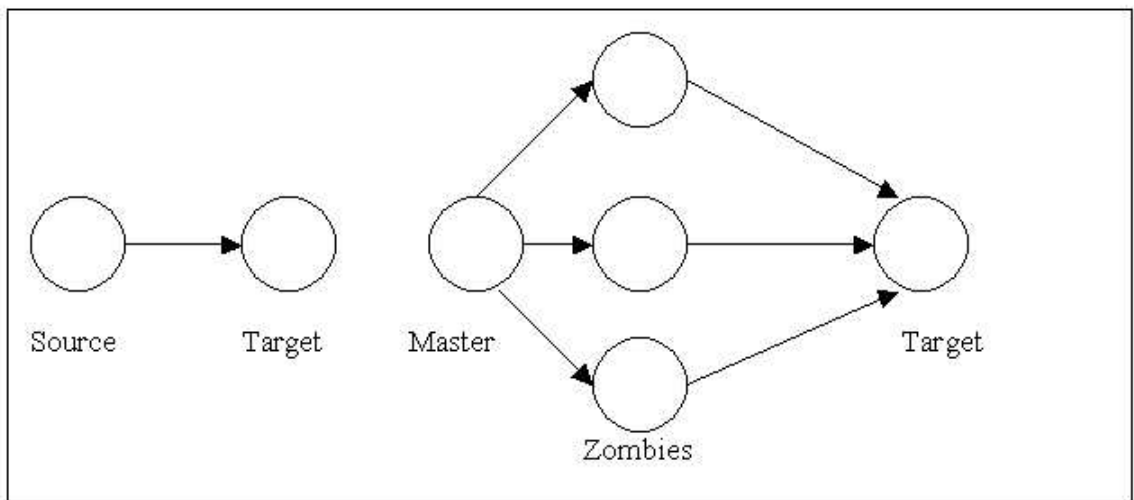


Figure 2 DoS vs DDoS

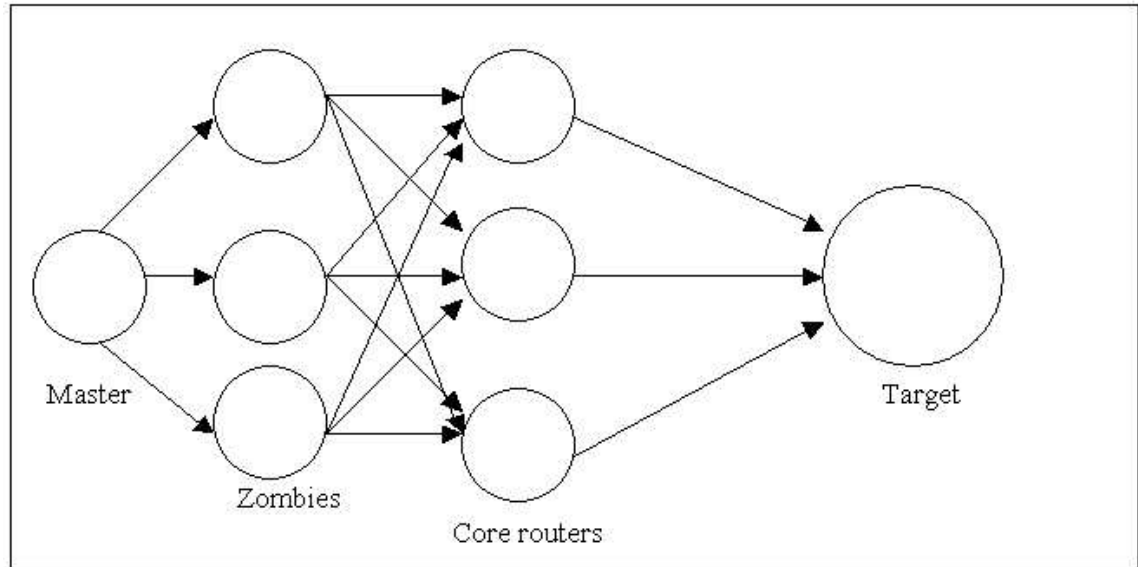


Figure 3. Distributed Reflection/Amplification DoS attack

A master uses zombies to send packets to reflection servers with an IP address of the target as the source. The reflection servers both amplify the attack and make it difficult to trace its source.

References

- 1) K. J. Houle, G. M. Weaver, N. Long, and R. Thomas, "Trends in denial of service attack technology." CERT Coordination Center, October 2001.
- 2) "Denial of Service Attacks." http://www.cert.org/tech_tips/denial_of_service.html
CERT Coordination Center
- 3) Steve Gibson, "Distributed Reflection Denial of Service."
<http://grc.com/dos/drDOS.htm> Gibson Research Corporation, Feb 2002.
- 4) Steve Gibson, "The Strange Tale of the Denial Of Service Attacks Against GRC.COM." <http://grc.com/dos/grcdos.htm> Gibson Research Corporation, Oct 2003.
- 5) Rob Thomas, Allen Householder, Art Manion, Linda Pesante, George M. Weaver, "Managing the Threat of Denial-of-Service Attacks." CERT Coordination Center, October 2001.
- 6) N. Fischback, S Lacoste-Seris, "Layer 2, routing protocols, router security & forensics." <http://www.defcon.org/html/links/defcon-media-archives.html> March 2003.
- 7) "NISCC Vulnerability Advisory 236929."
<http://www.uniras.gov.uk/vuls/2004/236929/index.htm> NISCC, April 2004
- 8) P. Watson "Slipping in the Window: TCP Reset Attacks"
<http://www.frame4.com/php/article2615.html> www.terrorist.net April 2004
- 9) S. Gallagher. "TCP Flaw No Cause for Alarm"
<http://www.eweek.com/article2/0,,1571169,00.asp> Eweek, April 2004.
- 10) D. Moore, G. Voelker, and S. Savage, "Inferring Internet denial-of-service activity," in USENIX Security Symposium, (Washington D.C.), August 2001.