



Network Presence, LLCSM

*Innovative Security Solutions*SM

Understanding, Planning For, and Responding To Denial of Service Attacks

SANS Network Security Conference 2001

Robert Brown
Vice President
rjb@netpr.com

Barrett Lyon
Security Consultant
blyon@netpr.com

**5959 W Century Blvd, Ste 734
Los Angeles, CA 90045
Phone 310.216.6850
FAX 310.216.6905
www.netpr.com**

Table of Contents

ABSTRACT	1
INTRODUCTION	1
UNDERSTANDING DENIAL OF SERVICE ATTACKS.....	1
GOALS OF DOS ATTACKS.....	1
TYPES OF DENIAL OF SERVICE ATTACKS.....	2
<i>Single Source vs. Distributed Source Attack</i>	2
<i>Flood Based Attacks</i>	2
<i>Crash based Attacks</i>	4
VULNERABILITY MANAGEMENT	4
PLANNING FOR DENIAL OF SERVICE ATTACKS	5
CASE STUDY OF THESHELL.COM	5
MAPPING EXISTING TRAFFIC PATTERNS.....	5
ASSESSING RISK OF DOS BASED ON EXISTING TRAFFIC PATTERNS.....	5
CREATING ROUTER FILTERS	5
GETTING YOUR ISP INVOLVED.....	6
INTRUSION DETECTION AND ITS ROLE IN DOS ATTACKS.....	6
COMMERCIAL TOOLS	6
PACKET SNIFFERS.....	7
CREATING AN INCIDENT RESPONSE PLAN	7
RESPONDING TO DENIAL OF SERVICE ATTACKS	8
IDENTIFYING THE ATTACK.....	8
PSYCHOLOGY AND THE ATTACKER.....	8
TAKING ACTION TO STOP THE ATTACK.....	9
LEGAL ISSUES	9
CONCLUSIONS.....	10
APPENDIX A. LISTING OF THESHELL.COM OUTAGES	11
APPENDIX B. SAMPLE ISP-SIDE FILTER FOR A CISCO ROUTER	12
ISP INGRESS ACCESS LIST (ISP INGRESS, YOUR EGRESS).....	12
ISP EGRESS ACCESS LIST (ISP EGRESS, YOUR INGRESS).....	12
ISP EGRESS RATE LIMITING FILTERS	12
APPENDIX C. SAMPLE ISP CONTACT POLICY	14
REFERENCES	15

Abstract

Denial of Service attacks occur on a daily basis. Are you prepared? This paper provides a foundation for understanding and responding to Denial of Service (DoS) attacks with specific focus on Distributed (DDoS) attacks. We relate specific experience setting up and maintaining a network that continues to attract at least one denial of service attack per day. The presentation material is organized in a logical fashion and is designed to provide the reader with a step-by-step guide to creating and maintaining an incident response policy for Denial of Service attacks.

Introduction

Computers and networks need certain things to operate: network bandwidth, memory, disk space, CPU time, data structures, access to other computers and networks, and certain environmental resources such as power, cool air, or even water. Without these ingredients, the system will not perform as expected. A DoS attack is designed to impact one or more of the above requirements to cause the system or network to refuse to deliver the expected result.¹ It is estimated that there are at least 4,000 distinct DoS attacks per week on the global Internet – and the number is growing every day.² Organizations large and small are being attacked on a regular basis. Recent news coverage has focused on attacks against Microsoft, CERT, the FBI, E-Trade, and Yahoo, although there are many other attacks being carried out against lower-profile sites. Despite the threat of attack, very few organizations have an in-depth understanding of the threat and are adequately prepared to defend their network.

Our paper is organized to offer the system and network administrator a guide for understanding DoS attacks, creating a response plan, and taking swift action in the event of an actual attack against their network. We seek to educate the reader on understanding what an attack looks like, identifying an attack in progress, and responding to the attack to restore service to the network. The majority of the research for this paper comes from first-hand experience fighting DoS attacks at an Internet Service Provider (ISP), TheShell.com. We introduce this work as a case study in the “Planning For Denial of Service Attacks” section of the paper.

Understanding Denial of Service Attacks

Goals of DoS Attacks

One of the most important things to understand is what the attacker is attempting to do to your network. Are they trying to take a specific machine offline? How about a specific service? Is the goal to take down the entire network? Perhaps they desire to drive up the cost of your ISP bill by artificially increasing your traffic (many colocation companies charge by the amount of traffic used.) Other goals include using a DoS attack in conjunction with or as a decoy for another type of exploit. An excellent example is an attack against your domain name system (DNS) infrastructure. Here is a sample attack scenario:

Goal: To hijack all incoming mail to victimcompany.com.

Scenario: Determine what name servers host the DNS for victimcompany.com. Use the nslookup tool to gather the NS records for the domain. Use the dig command to request the version.bind string from each of the remote name servers.

¹ Denial, “Merriam-Webster’s Collegiate Dictionary.” 2001.

² Moore, D. and Voelker, G., “Inferring Internet Denial of Service Attacks.” 2001.

Our preliminary information has indicated that victimcompany.com hosts their own primary DNS and their ISP hosts their secondary DNS.

Attack: Their ISP is running an outdated version of the BIND software that is vulnerable to a remote exploit. We launch a denial of service attack against victimcompany.com's DNS server to take it (or the network link) offline thereby making the ISP's DNS server the only working server. We compromise that server, change the MX (mail exchanger) to point to our collection point, and wait.

Result: Emails that were destined for victimcompany.com are now redirected through a new site. While all of this is occurring, victimcompany.com's IT staff is grappling with the DoS attack and trying to restore service to their network.

Does this sound like an implausible scenario? It is not. This is just one example of how a DoS attack can be leveraged for additional gain. Understanding the motivation of the attacker is an important factor in grasping the overall concept of Denial of Service attacks.

Types of Denial of Service Attacks

The most common DoS attacks will target the computer's network bandwidth or connectivity. These attacks flood the network with such a high volume of traffic or number of connections that all available network resources are consumed and legitimate user requests cannot get through. Common types of DoS attacks are broken down into three categories: flood based attacks, crashed based attacks, and physical attacks. All of the attacks can either be from a single source or constitute a distributed attack.

Single Source vs. Distributed Source Attack

There are two main categories of attacks that will be launched against a network: single source attacks and distributed attacks.

Single source attacks are launched from a single IP address against a target network. The source is limited by the bounds of the computing power of that system and the network link it resides on. While these attacks are common, they are much less destructive than distributed attacks.

Distributed attacks (DDoS attacks) are launched from a large number of systems against a target network. The systems used in the attack are usually machines that have been remotely compromised by the attacker. After a compromise, the attacker will install software on the system to allow remote control and enable the use of the system (and its network link) to attack a target of his choosing. These attacks are very common, difficult to defend against, and are generally misunderstood. This paper focuses on responding to distributed attacks, however the concepts apply to single source attacks as well.

Flood Based Attacks

Flood-based attacks are usually distributed attacks executed against network connectivity. The goal is to prevent hosts or networks from communicating with legitimate client computers. The common types are SYN flooding, SYN-ACK flooding, ICMP flooding, and UDP flooding.

SYN Flooding

A widely used SYN flooding method gained popularity in late 1996. Cert released advisory CA-96.21.tcp_syn_flooding³. In this type of attack, the attacker begins the process of establishing a TCP

³ CA-96.21.tcp_syn_flooding - <http://www.cert.org/advisories/CA-1996-21.html>

connection to the victim machine, but only sends the first part of the TCP 3-way handshake, the SYN. The SYN packet is usually spoofed so the victim machine returns the SYN-ACK to the host it thinks is trying to initiate the TCP connection. In the meantime, the victim machine is waiting for the ACK packet to return and has reserved one of a limited number of data structures required to complete the impending connection. The result is that legitimate connections are denied while the victim machine is waiting to complete bogus “half-open” connections.

SYN-ACK Flooding

Much like the SYN flood, hosts can be used to further mask an attacker by sending specifically spoofed packets to relay points. By forging a packet with the victim’s IP address as the source of a SYN packet, a site can be used to relay SYN-ACK packets at the victim’s site. This hurts not only the victim but the relay site’s resources as well. We seldom see or hear about this attack being used in the wild.

ICMP Flooding

ICMP floods vary in many different ways. ICMP flooding can be implemented using the common command ‘ping’. On most systems –f (the ‘flood’ feature) can be used to flood a network with ICMP echo (type 8) and ICMP echo reply (type 4). This is the most basic form of the ICMP flood, but there are many variations of the attack.

The basic ICMP packets can have forged source addresses to reduce the chances of tracking the attacker back to his true host address. The packets can be randomly spoofed to create a huge level of complexity in tracking the attack to the true source. There are even attacks that randomize the ICMP type with the hopes of evading different types of Access Control Lists (ACLs).

ICMP flooding is also used in the most basic form of broadcast or DDoS attack as seen in the Smurf attack⁴: CERT describes the Smurf attack as follows:

On IP networks, a packet can be directed to an individual machine or broadcast to an entire network. When a packet is sent to an IP broadcast address from a machine on the local network, that packet is delivered to all machines on that network. When a packet is sent to that IP broadcast address from a machine outside of the local network, it is broadcast to all machines on the target network (as long as routers are configured to pass along that traffic).

In the “smurf” attack, attackers are using ICMP echo request packets directed to IP broadcast addresses from remote locations to generate denial-of-service attacks. There are three parties in these attacks: the attacker, the intermediary, and the victim (note that the intermediary can also be a victim).

ICMP flooding, both via standard echo/echo reply and via broadcast storms, is one of the most common ways of creating a DoS attack against a network.

UDP Flooding

The user datagram protocol (UDP) is a stateless protocol that does not require a handshake or connection to be established with a remote site. Common uses of the protocol are for audio/video traffic and other services that require high bandwidth and can tolerate lost packets in the stream.

⁴ CERT Advisory CA-1998-01: Smurf IP Denial-of-Service Attacks. Available: <http://www.cert.org/advisories/CA-1998-01.html>

A DoS attack that can be commonly found in the wild employs the use of UDP packets to flood a target network. This attack is a bit more difficult to detect and prevent since it can be using a random source or destination port. Perhaps the attackers feel UDP is a bit more inconspicuous – it is not quite as obvious because it will not show up in standard “netstat” commands or other network troubleshooting tools. Since the DNS service uses UDP, it is not possible to disable it completely on your network.

Crash based Attacks

Crash-based attacks are designed to exploit a programming flaw in a system, service, or protocol to create unexpected results. In practice, crash-based attacks have the advantage of not requiring the use of distributed attack clients however the resulting attack is usually limited to a single system or service. This can still be a powerful attack if the service is your core Internet router or production web server.

There are many different types of denial of service attacks that exploit software errors that will cause a machine to crash. The CERT advisory, CA-1997-28⁵ describes two common crash-based attacks:

Teardrop – Some implementations of the TCP/IP IP fragmentation re-assembly code do not properly handle overlapping IP fragments. Teardrop is a widely available attack tool that exploits this vulnerability.

Land – Some implementations of TCP/IP are vulnerable to packets that are crafted in a particular way (a SYN packet in which the source address and port are the same as the destination--i.e., spoofed). Land is a widely available attack tool that exploits this vulnerability.

Vulnerability Management

Managing your vulnerabilities is a key concept in preventing Denial of Service attacks. As stated above, many of the distributed attacks are carried out from clients that have been compromised via a known vulnerability. The ISC domain survey⁶ shows 109 million hosts advertised in DNS. With the large number of new vulnerabilities and the ever-growing Internet population, vulnerability management has become a key factor in preventing the spread of attack.

⁵ CA-1997-28: IP Denial of Service Attacks. Available: <http://www.cert.org/advisories/CA-1997-28.html>

⁶ The Internet Software Consortium Domain Survey is available at: <http://www.isc.org/ds/>

Planning for Denial of Service Attacks

The key to thwarting a DoS attack is simple: be prepared for it. If you are aware of the issues and have a plan in place ahead of time, the attack will be much easier to identify and respond to. In this section, we present the case study of TheShell.com and describe the methodology we developed and implemented in response to the high number of DoS attacks against TheShell's network.

Case Study of TheShell.com

TheShell.com has been providing Unix and Unix-like shell access to the Internet community for over five years. Most of their customers use IRC (Internet Relay Chat) from their shell accounts. As a result TheShell.com has become a prime target for Denial of Service attacks over the years. Because TheShell.Com aims to provide a quality user environment they have been forced to develop techniques and policies to prevent and reduce attacks both proactively and reactively. As of mid-2001, TheShell is experiencing at least one DoS attack per day and has logged 19 serious multiple-hour outages in the past 14 months. A listing of the outages is supplied in Appendix B.

Mapping Existing Traffic Patterns

The first step in preventing DoS attacks is to know your existing traffic patterns. If you don't know what normal traffic looks like, you will never be able to spot an anomaly. Network monitoring is a key aspect of DoS protection. Tools such as MRTG⁷ and NISCA⁸ provide an excellent way of visually displaying your network traffic statistics. Numerous commercial packages exist and can be used for this purpose, including enterprise management tools HP OpenView, IBM Tivoli, and CA Unicenter. Additionally, packet sniffers such as the aptly named "Sniffer" from Sniffer Technologies (a division of Network Associates, formerly Network General) are an excellent choice for temporary network monitoring and mapping.

When mapping your traffic patterns, you should focus on both ingress and egress traffic. Ingress is traffic destined for your network and egress is traffic originating from your network but destined for some point on the other side of your router, in this case the Internet. The focus of your mapping efforts should be to learn the normal peaks and valleys of your usage as well as the types of protocols and ports being used on your network. This is critical in configuring effective router access control lists.

Assessing Risk of DoS Based on Existing Traffic Patterns

The previous step discussed mapping your traffic patterns. Now that you know what your normal traffic looks like, the goal is to prune protocols and services that you do not use and limit the bandwidth on the ones that you do use.

The two most commonly used protocols in a DDoS attack are UDP and ICMP. If you can effectively limit their use within your network, the threat of DDoS attacks will be minimized. It is to your advantage to develop aggressive filters for these protocols. If your network functionality restricts your ability to completely filter them out, restrict them as much as possible.

Creating Router Filters

The key to creating effective router filters is to employ them on your ISP's side of the network. Filtering on your side is also a good idea but it will not help you during a massive DDoS attack. A sample Cisco router filter we designed for TheShell.com is presented in Appendix B.

⁷ Multi Router Traffic Grapher. Available: <http://people.ee.ethz.ch/~oetiker/webtools/mrtg/mrtg.html>

⁸ NISCA. Available: <http://nisca.sourceforge.net/>

The filters we designed limit the incoming amount of TCP SYN, UDP, and ICMP packets that can be sent to our network at any given time. Your network will never use 1 megabit of ICMP traffic, so limit it at the ISP.

In practice, your filters should be designed and implemented slowly. This will ensure that your existing traffic will not be disrupted while the new measures are put in place. Additionally, it is a good idea to test the filter on your side of the link prior to asking your ISP to implement it on their end. This will ensure it behaves as you expect it to.

Getting Your ISP Involved

The involvement of your ISP is key to protecting your network from DoS attacks. Since your Internet pipe is smaller than that of your ISP, you will most likely rely on the ISP to assist you in fighting the attack.

When soliciting a dedicated circuit from a network service provider, one must be aware of the service contract and service level agreement (SLA) that is involved. As usual, getting promises in writing is very important. When having a dedicated circuit installed, one should understand the role that the network provider will have during a denial of service attack. It is best to ask the question, “What would I do if I were to receive a huge attack that renders my services offline?” You need to know the policy (if any) that is in place for handling denial of service attacks. You also need to ensure that they will apply specific filters on their end of the network to provide your network with the needed proactive protection that is required. To facilitate this, you may have to have them add a special provision to your contract that specifically requires them to add filters to their routers. It is very important to negotiate this ahead of time – not during an attack. Use this as a bargaining tool when shopping for a new ISP.

Co-located servers are also a big issue. Many colocation companies bill their fees on an escalating scale based on the amount of bandwidth used. One of the standard billing practices is to calculate a monthly average based on the ninety-fifth percentile of usage. However, if an attack lasts longer than 24 hours, it would cause that percentile to rise. This would result in a larger bandwidth bill. Make absolutely certain (have it put in your contract) that you will not be billed for a DDoS attack. If your network is attacked, it would be prudent to contact an account manager to ensure that the attack is not billed.

If a dedicated server or co-located server is used, one should have the ability to identify the attack and to notify the provider’s data center. The traffic should be identified and logged in a trouble ticket.

Intrusion Detection and its role in DoS attacks

Network-based intrusion detection software (IDS) can be used to detect and respond to a DoS attack. The IDS software typically monitors the network for known patterns of attack and takes a predefined response if a pattern is matched. IDS can be an effective tool if deployed on your external network segment. It is not foolproof – the technology is known to have certain flaws and may not be able to detect attacks that it does not have a signature for.

Commercial Tools

There are numerous companies developing and deploying new technologies for detecting DDoS attacks. Some of the players in the market are Mazu Networks, Asta Networks, Arbor Networks, and Captus Networks. Additionally, companies such as Radware, Nortel, and Cisco have some measure of protection built-in to their high availability product lines. Many of the commercial firewall products have limited protection for SYN floods and other types of basic DoS attacks. None of the available tools will protect

you from a dedicated attack, however they may make detection and response easier. Each tool should be evaluated in the context of your network and with respect to the cost/benefit of the solutions.

Packet Sniffers

The next section of the paper, *Responding to DoS Attacks*, makes at least two critical assumptions: first that you have a packet sniffer available and second that you have the ability to plug it into your external switch. In terms of Cisco equipment, this means that you should have a SPAN or monitor port configured on your switch that will allow you to see all traffic to and from your ISP-connected router. This is vital in detecting an attack. You do not want to be trying to locate a sniffer, configure it, and then try to figure out how to plug it in during an attack. This simple step will save you a great deal of grief during a DoS attack.

In terms of what sniffer to use, we prefer the commercial Sniffer product from Sniffer Technologies. Open-source sniffers such as Sniffit and TCPdump can also be useful for the budget-minded IT shop. Cisco's IOS also has sniffer capabilities that are worth investigating.

Creating An Incident Response Plan

The incident response plan is perhaps the most important aspect of planning for an attack. The policy should contain a list of people in your organization and their role during an attack. This includes personal phone numbers, pager numbers, and any other desirable means of contact. The most important people on the list are your Senior Network Engineer and your Information Security Officer (ISO). The ISO will likely be directing the response effort while the engineer will be working to diagnose the attack and to create filters to stop it.

The response plan should also list the relevant ISP contact information: phone numbers for your ISP Network Operations Center (NOC) and IP team, circuit IDs, and cell phone, pager, and home phone numbers of your senior technicians. Do not wait until the last minute to get to know your ISP's senior engineers! Call them, introduce yourself, explain what you are trying to accomplish, and send them a case of beer and a pizza. They will likely give you their contact information. This is invaluable when trying to thwart an attack at 2 AM on a Saturday morning. A sample ISP contact form is included in Appendix C.

Responding to Denial of Service Attacks

Identifying The Attack

In the above section, we mapped our network traffic and deployed monitoring tools, created filtering rules at our ISP, documented an internal response plan, prepared a packet sniffer and monitoring port, and possibly deployed intrusion detection and other commercial tools. The next step is to take all of the above preparation and leverage it into an effective response to a DoS attack. To do that, you must first realize that an attack is taking place.

The first indication of attack will usually come from either a performance problem or, if you are proactively using enterprise management tools, your network monitoring infrastructure. If you suspect an attack, it is imperative to start making quick decisions and implementing your response plan. Contact the relevant people, get them in position to login to your network hardware, contact your ISP to open a trouble ticket, and start deciphering the attack. Your network engineer should log on to the router that connects you to your ISP. Dump the configuration and network statistics to see CPU usage, packet flow, network usage, and the count of packets that have hit your filters. The goal is to gain insight into the overall usage of the network. The attack may just jump out at you. If not, there are additional steps to take.

After reviewing your external router, connect your sniffer to the network to view the actual traffic that is coming from your router. Analyze the trends and attempt to pinpoint a pattern in the attack. Is the attack targeted at one specific host or service? What are the source addresses that are sending the data? The goal of this step is to effectively analyze the trend to first determine the intent of the attacker and then create a strategy to mitigate the attack while still allowing legitimate connections through. Note that not all attacks are constant; there is a trend that sends short bursts of data to a target network. This type of attack may be more difficult to detect and prevent than a constant flood.

Psychology and the Attacker

To stop an attack, you must understand the intent of the attacker. Put yourself in the position of the attacker and ask the question: What is the goal of this attack? The sniffer traces and other data should give you an understanding of what the attacker is trying to accomplish. The best way to respond is to give in to the attacker immediately. If they are attacking a web server but only a specific IP, remove the virtual IP address. If your link is already dead, there is no harm in taking down a specific machine or service. In many cases the attacker will feel that they have prevailed and will stop the attack. Sacrificing one machine or service to allow the rest of the network may be a tradeoff that you are willing to make.

Another psychological aspect of responding to DoS attacks plays into the stupidity of the attacker. While you are monitoring the network with your sniffer, pay special attention to ICMP echo/echo reply packets and UDP/ICMP packets that could be used for traceroute. In many cases the attacker will be monitoring your network to determine the effectiveness of his attack. Our experience shows that many of the attackers ping and traceroute from their actual home Internet connection. This could prove to be valuable if you wish to locate and prosecute the attacker at a later date. This is a much easier route to the attacker than attempting to trace the attack back via ISPs and compromised hosts.

Taking action to stop the attack

The next step in our response plan is to take our knowledge of the attack, create a customized filter, and “climb the ladder.” We define the ladder as follows:

- ISP’s Router
- Your Border Router
- Local Network Segment (switches/routers)
- Host IP Stack
- Port/Service

The goal is to push the attack up the ladder to your ISP’s side of the network link. Start by implementing rate limiting and source/destination deny filters on your router to limit the attack to that point. Once you have created a working filter, the attack will be limited to your border router. The individual hosts and services on your network should be available locally. Now, take your new filter, contact your ISP, and ask them to implement it on their router. You should make sure to test the filters on your side **prior** to asking your ISP to implement the changes on their router. This reduces the time the ISP engineer has to spend on the call and will make the process much easier for them to handle – especially if it is 2 AM. If the attack effectively stops, you have now pushed the problem up the ladder to your ISP. Now they can pick up the attack and try to trace it back and filter it out of their network.

Alternatively, if the attack targets one host, ask your ISP to add a “null route” for that address. This will prevent traffic for that IP from being routed to your network, effectively stopping the attack at the ISP. This is a much easier solution for your ISP and they generally prefer null routing over filtering due to the lower management overhead involved.

Legal issues

The attack has now been stopped – how do you find out who did it? It is nearly impossible to trace the source of a distributed attack. It involves quickly involving all of the ISPs that may have carried the DoS traffic to trace it to its source. At that point, you must contact the organization that was the source of the attack. In many cases, you will be telling them that one of their hosts was compromised. They may or may not help you depending on culture, language, time, and legal barriers. Even if you get this far, the compromised host must be analyzed to try to trace the attacker back to his source. It is a difficult problem and is more likely to waste your time than to lead you to the attacker.

In some cases, you may be able to go after the organizations whose systems attacked you. We are not aware of case law dealing with third-party compromise and attack using those hosts, however it could be argued that the organizations were negligent for not applying known security fixes to the attacking hosts. Liability laws differ from state to state, but in most cases proving negligence is difficult as it requires all of the following to be true: 1) The duty of care to the person being damaged, 2) The person or organization must be negligent (What would a reasonable person do under the circumstances?), 3) There must be damages, 4) The damages must be proximately caused (having an actual cause and being foreseeable). Prosecution is also difficult because information can travel across many jurisdictional boundaries.

An additional legal issue involves your liability to your customers. Will they hold you liable for downtime as a breach of service? If you are providing service level agreements or are in a business requiring timely dissemination of information (such as online stock trading), there may be some risk involved.

In all cases, we suggest contacting your legal department to discuss these issues during the implementation of your response plan. The most important aspect is discussing with your attorney how to reach a figure for damages caused by DoS attacks. This is very important since it is the first or second question the FBI will ask you after calling to report an attack.

Conclusions

This is a wide-ranging problem involving information security, network engineering, legal, and ethical problems. As of today, it is not a problem that can be solved by technology. Just like the rest of information security and business in general, it requires process and planning for effective management. You must understand the issues surrounding the problem, be prepared to deal with them, and be able to quickly respond in the event of an actual attack. With the concepts presented in this paper, we have been able to respond to the attacks against TheShell.com and keep the network operational. Our sincere hope is that your organization will see the value in creating a response plan and will implement some of the ideas presented in this paper. It will make for a safer and more reliable Internet for all of us.

Appendix A. Listing of TheShell.com Outages

(5-2000 to 2-2001)

05-19-2000 "Dropped offline"

05-31-2000 "DDoS directed at oxygen.theshell.com"

06-21-2000 "Lost routing and dropping packets"

06-23-2000 "Hard fail at 2:30 AM"

06-29-2000 "Packet loss caused by UDP flood"

07-06-2000 "Dead circuit"

07-10-2000 Unknown "DoS"

07-17-2000 "DoS attack to specific host"

07-21-2000 1:34 AM "Unknown loss of routing, looks like a DoS attack"

07-23-2000 9:39 PM "DoS"

07-24-2000 10:30 PM "circuit dead"

08-07-2000 10:02 PM "circuit offline, possible DoS"

08-20-2000 "DoS attack to specific host, ticket with above.net"

08-21-2000 "Qwest vs. Above.net peering problem"

09-19-2000 "DoS attack that lasted all night 12:01 AM to 10:00 AM"

10-02-2000 "DoS attack"

02-02-2001 "DoS attack that lasted for 2 days to specific host"

02-09-2001 "DoS attack that lasted for a day to specific host"

02-18-2001 "DoS attacks to router interfaces, so we put null routes in for the Ethernet and serial interfaces"

Appendix B. Sample ISP-side filter for a Cisco router

DISCLAIMER: The following access lists were designed for TheShell.com only. Blindly using these access lists on your network could cause problems by restricting traffic that you really do want to pass. This appendix should be looked at as an example, not something to take verbatim and drop into your router.

ISP Ingress Access List (isp ingress, your egress)

This filter is designed to prevent your ISP from accepting spoofed traffic from your site. All ISPs SHOULD be doing this although many do not. This access list works because all Cisco ACLs end in an implicit deny.

```
permit ip <your network> <your subnet> any
```

ISP Egress Access List (ISP egress, your ingress)

Note that the order of the rules is important. Everything that does not include layer 4 information is presented at the top of the list. The order then proceeds based on the highest expected match occurrence (TCP, UDP, ICMP). The TCP section is still ordered for expected highest match first for established packet flows. In this case, our network is 192.168.1.0/24. The first rule is for anti-spoofing.

```
deny ip 192.168.1.0 0.255.255.255 any
permit tcp any any established
permit tcp any any gt 1023
permit tcp any any eq www
permit tcp any any eq 443
permit tcp any any range ftp-data telnet
permit tcp any any eq ident

permit udp any any range talk 518
permit udp any any eq domain
permit icmp any any echo-reply
permit icmp any any unreachable
permit icmp any any echo
permit icmp any any time-exceeded
```

ISP Egress Rate Limiting Filters

Rate limiting filters are designed to allow only a set limit of a certain type of traffic. This is a very basic way to implement packet shaping however it is adequate for our purposes. The following rate limits match packets by access list. The first list is for SYN flood protection, the second is for UDP, and the third is for ICMP. Note that the SYN rule works because anything other than a SYN packet will be using a port greater than 1023 and it will have the established bit set.

Input Rules

```
matches: access-group 140
params: 600000 bps, 600000 limit, 600000 extended limits
```

matches: access-group 141
 params: 296000 bps, 300000 limit, 300000 extended limit
matches: access-group 142
 params: 120000 bps, 120000 limit, 120000 extended limit

Output Rules

matches: access-group 140
 params: 600000 bps, 600000 limit, 600000 extended limit
matches: access-group 141
 params: 296000 bps, 300000 limit, 300000 extended limit
matches: access-group 142
 params: 120000 bps, 120000 limit, 120000 extended limit

Access Lists

Extended IP access list 140
 permit tcp any any syn
Extended IP access list 141
 permit udp any any
Extended IP access list 142
 permit icmp any any

Appendix C. Sample ISP Contact Policy

TheShell.com

Qwest Communications

NOC : 1-800-222-3333 Press: 1,#,2,2

IP Team : 888-333-4444

Tony : 408-555-6677

Tony Cell : 703-455-6677

CORE : 98765432

ACCT : 44566789

Circuit : 1234567890

email : support@qwestip.net

: cmc1@qwest.com

References

1. (2001) “Merriam-Webster’s Collegiate Dictionary.” Available: <http://www.m-w.com>
2. Moore, David and Voelker, Geoffrey M. (2001) “Inferring Internet Denial-of-Service Activity.” Available: <http://www.caida.org/outreach/papers/backscatter>, July 2001.
3. CERT Advisory CA-96.21.tcp_syn_flooding (1996) Available: <http://www.cert.org/advisories/CA-1996-21.html>
4. CERT Advisory CA-1998-01 (1998) Smurf IP Denial-of-Service Attacks. Available: <http://www.cert.org/advisories/CA-1998-01.html>
5. CERT Advisory CA-1997-28: (1997) IP Denial of Service Attacks. Available: <http://www.cert.org/advisories/CA-1997-28.html>
6. The Internet Software Consortium Domain Survey. (2001) Available: <http://www.isc.org/ds>
7. Multi Router Traffic Grapher. (2001) Available: <http://people.ee.ethz.ch/~oetiker/webtools/mrtg/mrtg.html>
8. NISCA. (2001) Available: <http://nisca.sourceforge.net>