



Interested in learning
more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

A Summary of DoS/DDoS Prevention, Monitoring and Mitigation Techniques in a Service Provider Environment

The frequency and sophistication of Denial of Service (DoS) and Distributed Denial of Service attacks (DDoS) on the Internet are rapidly increasing. Service providers are under mounting pressure to prevent, monitor and mitigate DoS/DDoS attacks directed toward their customers and their infrastructure. The Internet is part of the critical national infrastructure but is unique in that it has no customary borders to safeguard it from attacks. Attacks that are seen everyday on the Internet include direct attacks, remote co...

Copyright SANS Institute
Author Retains Full Rights

AD

DEEPARMOR®

A Summary of DoS/DDoS Prevention, Monitoring and Mitigation Techniques in a Service Provider Environment

Michael Glenn

August 21, 2003

GSEC Practical Version 1.4b

Option 1

© SANS Institute 2003, Author retains full rights

Table of Contents

1	Abstract	1
2	Introduction	1
3	DoS/DDoS Attacks	2
3.1	Direct Flooding Attacks	2
3.2	Remote Controlled Network Attacks	3
3.3	Reflective Flooding Attacks	3
3.3.1	Smurf and Fraggle Attacks	4
3.3.2	ICMP	5
3.3.3	TCP SYN	5
3.3.4	UDP attacks	6
3.3.5	TTL Expiration	6
3.3.6	DRDoS	6
3.4	Worms	7
3.5	Viruses	8
3.6	Protocol Violation Attacks	8
3.7	Fragmentation Attacks	8
3.8	Network Infrastructure	9
3.8.1	Control Plane Attacks	9
3.8.2	Management Plane Attacks	9
3.8.3	Network Services Attacks	10
3.9	Other Attacks	10
4	Prevention	10
4.1	Policies and Procedures	11
4.2	New Product/Upgrade Design and Testing	11
4.3	Patch Management	12
4.4	Scanning/Auditing	12
4.5	uRPF	12
4.5.1	Cisco Implementation	13
4.5.2	Juniper Implementation	15
4.6	Management and Control Plane Protection	16
4.6.1	Router Access	16
4.6.2	Router Engine Protection	17
4.6.3	Prefix Filtering	18
4.7	FW/IDS/IPS	18
4.7.1	DNS Considerations	18
4.7.2	Other Services	19
5	Monitoring	19
5.1	Customer/Peer Notification	19
5.2	Sinkhole	19
5.3	Backscatter Techniques	20
5.4	Netflow Monitoring	21
5.4.1	Netflow DoS/DDoS Considerations	21
6	Mitigation	22

6.1	ACLs/Rate Limiting	22
6.1.1	Rate Limiting	23
6.1.2	Other	23
6.2	Destination based Black Hole Filtering	23
6.2.1	Customer Initiated Black Hole Filtering	24
6.2.2	Source based Black Hole Filtering	24
6.2.3	Black Hole Shunting	25
6.2.4	Advanced BGP Filtering	25
6.3	Attack Distribution and/or Isolation – Anycast	25
7	Conclusion	26
8	References	27

© SANS Institute 2003, Author retains full rights

Figures

Figure 1 - Direct Attack	3
Figure 2 - Example Reflective TCP SYN Flood Attack	4
Figure 3 - DRDoS Attack	7

© SANS Institute 2003, Author retains full rights

1 Abstract

The frequency and sophistication of Denial of Service (DoS) and Distributed Denial of Service attacks (DDoS) on the Internet are rapidly increasing. Service providers are under mounting pressure to prevent, monitor and mitigate DoS/DDoS attacks directed toward their customers and their infrastructure. The Internet is part of the critical national infrastructure but is unique in that it has no customary borders to safeguard it from attacks. Attacks that are seen everyday on the Internet include direct attacks, remote controlled attacks, reflective attacks, worms, and viruses. Specific attacks directed at a service provider's infrastructure can be very damaging and cause wide spread outages. This paper covers these attacks and discusses techniques to prevent attacks including good security policies, new/updated product security testing, patch management, spoofed packet dropping (uRPF) and firewall/IDS/IPS deployment in a service provider environment. Protection of the provider's infrastructure is another key aspect and is addressed in this paper.

Attack monitoring and mitigation is a crucial part of a service provider's operation. DoS/DDoS and DRDoS monitoring techniques are reviewed and practical mitigation techniques are discussed. Widespread deployment of remotely triggered black hole filtering is a quick and effective way of mitigating many of these attacks. New techniques that combine uRPF, rate limiting and granular filtering lists with black hole filtering are providing service providers with a new arsenal of tools to keep up with the ever escalating arms race on the Internet.

2 Introduction

The number Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks on the Internet has risen sharply in the last several years. Service providers are routinely expected to prevent, monitor and mitigate these types of attacks which occur daily on their networks. This paper discusses the most common types of DoS/DDoS attacks seen on the Internet and ways that service providers can prevent or mitigate damages from the attack threats.

DoS/DDoS attacks have become more sophisticated in the last several years as the level of attack automation has increased. Sample and fully functional attack software is readily available on the Internet. Precompiled and ready to use programs allow novice users to launch relatively large scale attacks with little knowledge of the underlying security exploits. The advent of remote controlled networks of computers used to launch attacks has changed the landscape and methods that a service provider must use. In the past year, Black Hats have taken theoretical optimizations in worm propagation and applied them to the fastest spreading worm today, Slammer. [STA01 6-11]

Slammer has changed the tools required for service providers to effectively deal with worm propagation. With 90% of the vulnerable hosts infected within the first 10 minutes

of release and an infection doubling time of 8.5 seconds [MOO01 1], service providers must have semi-automated techniques at their disposal to mitigate a large scale Internet event in a matter of minutes, instead of hours or days.

Expectation levels for service providers are also increasing as companies revenues are directly tied to having reliable connectivity to the Internet. The financial industry is especially susceptible to DoS/DDoS attacks as millions of consumers move to electronic bill payments, purchases and on-line banking. DoS/DDoS monitoring and black hole filtering are becoming entry level requirements for service providers to sell Internet services in the financial industry.

3 DoS/DDoS Attacks

DoS attacks can be classified as logic attacks and resource exhaustion flooding attacks. Logic attacks exploit security vulnerabilities to cause a server or service to crash or significantly reduce performance. Resource exhaustion flooding attacks cause the server's or network's resources to be consumed to the point where the service is no longer responding or the response is significantly reduced [MOO02 3]. Logic attacks will be evaluated based on their effect on the network infrastructure and critical network services (DNS, BGP, RADIUS, etc). A complete discussion of logic attacks is very broad and outside the scope of this paper.

Flooding attacks can be evaluated by their amplification factor. The amplification factor is the amount each source packet is multiplied by before reaching the victim. For example, in a direct flooding attack, for each source packet transmitted by the attacker, one packet is received at the victim's site. In a smurf reflective attack, each packet is reflected off a set of hosts that send multiple packets to the victim site. A smurf attack can achieve an amplification factor in the hundreds. In other words, for each source attack packet sent, hundreds of packets are received by the victim.

The duration between a publicly announced vulnerability to the time that an exploit is released in the wild is decreasing. Zero day exploits are exploits that are released very close to the same time that the vulnerability is announced. A recent example is the Cisco IOS vulnerability for 4 IP protocols [CIS04]. The exploit for this vulnerability was so trivial that exploits were available the day it was publicly announced. Steps taken at the request of Cisco by the tier 1 ISPs before the public announcement prevented the vulnerability from causing widespread disruption of Internet traffic.

3.1 Direct Flooding Attacks

The simplest case of a DoS attack is the direct flooding attack. In this case, the attacker sends packets directly from his computer(s) to the victim's site. In the attack, the source address of the packets may be forged. There are many tools available to allow this type of attack for a variety of protocols including ICMP, UDP and TCP. Some common tools include stream2, synhose, synk7, synsend, and hping2. This type of attack usually has

an amplification factor of 1 to 1. That is, for each packet sent by the attacker, one packet is received by the victim.




Figure 1 - Direct Attack

3.2 Remote Controlled Network Attacks

Remote controlled network attacks involve the attacker compromising a series of computers and placing an application or agent on the computers. The computer then listens for commands from a central control computer. The compromise of computers can either be done manually or automatically through a worm or virus. Typical control channels include IRC channels, direct port communication or even through ICMP ping packets [COL01 191]. Other versions can operate almost completely stealth. They can spoof the from and to addresses. The zombie listens passively (non promiscuous) for TCP SYN packets on different destination ports in a specific order. When the ports are matched, either from a specific IP address or any IP address, a user defined function is called. The attacker could use the packet header fields to determine what command to run and what IP address to attack. Cd00r.c is a working example of this [PHE01].

Attacks can be launched from the compromised computers either directly at a target or through a reflective media described below. Remote controlled attacks are very difficult to trace to the original control computer. A distributed reflective DoS attack is especially difficult to trace and is explained in detail later in this section.

3.3 Reflective Flooding Attacks

Reflective attacks forge the source address of the IP packets with the victim's IP address and send them to an intermediate host. When the intermediate host sends a reply, it is sent to the victim's destination address, flooding the victim. Depending on the type of protocol used and the application and configuration involved, amplification factors of 3 to several hundred are possible.

Reflective attacks can be difficult to trace to the original attacker because the flood packets are actually sent from intermediate servers. In many types of reflective attacks, the intermediate servers are usually well known, public servers such as www.amazon.com, www.cnn.com, etc. The victim's service provider cannot block access to these sites and many times end up blocking all the traffic to the victim's site to allow other network traffic to get through.

The speed of the reflective media (e.g. servers, routers, etc) is an important consideration for this type of attack. Paxson has an excellent paper discussing different types of reflective attacks, protocols involved, attack identification and effective defensive measures. He identifies three particular types of reflectors that make excellent reflective media: DNS servers, Gnutella servers and TCP servers (web servers for example) [PAX01]. ASICs based hardware can also provide an excellent reflective media. This would include misconfigured routers and SSL accelerators.




Figure 2 - Example Reflective TCP SYN Flood Attack

3.3.1 Smurf and Fraggle Attacks

One of the earliest, reflective attacks was the smurf attack. The smurf attack is performed by sending an ICMP echo request ping packet with the victim's address as the source address to a network's broadcast address. The fraggle attack is similar except it uses UDP packets. If the network router and network servers are configured to respond to a network address ping packet, all the servers on that subnet will respond to the forge source IP address, flooding the victim's site.

The attack is easily prevented by configuring the router to block broadcast packets that did not originate from that network. On a Cisco router, the command is "**no ip directed-broadcast.**" Beginning with Cisco IOS version 12.0, the default was to drop directed

broadcasts [CIS01]. The Juniper default is also to drop directed broadcast packets. Individual hosts can be protected by dropping ICMP broadcast packets; however, if the packets are blocked at the router, there is no need for individual host protection.

The amount of amplification that is achieved varies with the number of hosts located on the network. In general, it is difficult to find networks that will respond to ICMP or UDP broadcast packets. However, if a network is found, it usually provides a large amplification factor to the attacker.

3.3.2 ICMP

The reflective ICMP attack uses public sites that respond to ICMP echo request packets to flood the victim's site. Most well known public sites block ICMP to their networks as a result. However, routers respond very efficiently to ICMP and if not properly rate limited, can be an excellent reflective media. This attack by itself does not amplify the packets sent to the victim's site. If used in conjunction with a remote controlled network of computers, this attack can be very difficult to trace.

3.3.3 TCP SYN

The TCP SYN flood attack is a protocol violation attack that is used in several variations. In the simplest case, an attacker sends the first packet (with the SYN bit set) of the well known TCP three way handshake. The victim responds with the second packet back to the source address with SYN-ACK bit set. The attacker never responds to the reply packet, either on purpose or because the source address of the packet is forged. In the original attack, the victim's TCP receive queues would be filled up, denying new TCP connections.

Most modern UNIX and Windows implementations have fixed this issue by increasing the queue size and rate limiting the number of TCP SYN packets allowed. TCP SYN cookies are another way to mitigate this type of attack using cryptographic techniques to create the server's initial sequence number [BER01]. SYN cookie TCP stack implementations are available for many popular operating systems.

A variation to this attack uses public servers as a reflective media to flood the victim with TCP SYN ACK packets. In this case, the attacker spoofs the source address of the TCP SYN packet with the victim's address. The packet is sent to a public server that provides a public TCP service (such as HTTP). The server sends a TCP SYN ACK packet to the victim's host. The victim, having not sent the original packet either ignores the packets or sends a TCP RST packet. The technique can achieve 3 to 5 times amplification factors by retry packets sent from the reflection servers.

3.3.4 UDP attacks

The UDP protocol can be very efficient for DoS/DDoS attacks. UDP is a stateless protocol and does not have any acknowledgement mechanism by design. PROTOS, the SNMP test suite, and other SNMP tools have been used successfully to launch application level DoS attacks. The Slammer worm was extremely fast because it did not require a response from the compromised computer.

3.3.5 TTL Expiration

The TTL expiration attack relies on ICMP control messages to flood the victim. In this attack, the source address is forged to match the victim's address. The TTL for the packet is set to a low value that will expire in transit at a high speed router. When the TTL of the packet reaches zero, the router drops the packet and sends an ICMP TTL expired message to the source address, in this case the victim's site. Since TTL expiration is often done on the line card in ASIC, this can be an extremely fast reflective media.

The best defense for this type of attack is rate limiting ICMP to all routers in the service provider's network. Some network equipment vendors are now offering the ability to turn off TTL expiration processing, with the side effect of breaking traceroute.

3.3.6 DRDoS

A distributed reflective DoS (DRDoS) attack uses a remote controlled collection of computers to spray spoofed packets to a reflective media, typically servers or routers in one of the reflective attacks described above. Figure 3 shows a DRDoS UDP and TCP flood attack. As the diagram shows, this attack is especially difficult to trace because the controlling computer is 2 layers hidden from the packets received at the victim's site.




Figure 3 - DRDoS Attack

A DRDoS attack can use a variety of reflective attacks from the zombie computers through the reflective hosts to the victim. The traffic from the master computer to the zombie computers can be very low, even hidden with normal ICMP echo reply ping packets which are almost always allowed into a network [SANSTK01], allowing control of zombie computers behind firewalls. Programs available on the Internet include TFN2K, Trinoo, Stacheldraht, and Shaft [DIT01].

3.4 Worms

Worms are distinguished from viruses in the fact that a virus requires some form of human intervention to infect a computer where a worm does not. Worms have had the ability to significantly disrupt the normal operation of the Internet since the Morris worm in 1988. Since that time, significant worms that have caused significant network interruptions include Ramen(2000), LiOn(2000), Adore/redWorm(2001), Code Red I (2001), Code Red II (2001), Nimba (2001), Slapper (2002) and Slammer (2003). The initial DDoS against Microsoft by the MS Blaster (2003) worm was averted (at least at the time of this writing) by errors in the worm itself that allowed Microsoft to pull the DNS record (A record) of windowsupdate.com and prevent the attack against its site. The irony is that the worm succeeded in the ultimate DoS since the site is off the air with the name of windowsupdate.com. The vulnerable host scanning in both MSBlaster.A and MSBlaster.D affected enterprises, broadband and dial providers, and home users.

Worms can create Internet wide events based on scanning and infection traffic volumes (Code Red, Slammer), automated DDoS events (MS Blaster), or by creating zombie networks used to launch large scale DDoS attacks. Worm propagation technology has advanced significantly in the past several years [STA01, MOO01]. Slammer was able to infect 90% of the vulnerable hosts within the first 10 minutes of release and had an infection doubling time of 8.5 seconds [MOO01 1]. Slammer's spread was slowed only because of bandwidth limitations.

3.5 Viruses

Viruses have had a lesser but significant impact on network providers. They are often used today to build large zombie networks. Original research on viruses was occurring in 1983 and 1984 but only much later would they have a significant impact on Internet operations. Significant Internet viruses included Melissa (1999), Love Letter (2000), Nimda (2001 – a combination worm and virus), and SoBig (2003).

An interesting side effect is the worm and virus hoaxes. These are usually dire warnings that tell the person to notify all their friends about a fictitious worm, virus, or other situation. Although never a significant Internet problem, these have clogged enterprise e-mail systems and continue to circulate today. SPAM has far outstripped these hoaxes in causing problems on the Internet with e-mail.

3.6 Protocol Violation Attacks

All attacks could be considered protocol attacks in the sense that the attacker is sending packets in a manner not originally intended. Sometimes this is beneficial to the community as when Van Jacobson developed the traceroute program using ICMP return codes from the routers. In many situations, however, this is not the case.

Protocol violation attacks are generally referring to attacks that use IP protocols that are not valid or are reserved. Protocol 255 is reserved and protocols 135-254 are unassigned according to the Internet Assigned Numbers Authority (IANA) [IAN01]. Theoretically, traffic on those ports should never be seen on the Internet. Caution should be stated here, many protocols are in use by vendors using proprietary algorithms that have not been formally approved by the IANA. Blocking all non-approved protocols could cause outages for service provider customers.

3.7 Fragmentation Attacks

Packet fragmentation can be used in two distinct areas: evasion of IDS detection and as a DoS mechanism. As a DoS mechanism, fragmentation is used to exhaust a system's resources while trying to reassemble the packets. These types of attacks have occurred against CheckPoint firewalls, Cisco routers and Windows computers [SEC01].

3.8 Network Infrastructure

Attacks directed at network infrastructure can have a serious impact on the overall operation of the Internet. These attacks can create regional or global network outages or slowdowns. Recent attacks against the Internet's root nameservers caused enough concern for an FBI investigation into the attack. It sent a warning signal to the root nameservers operators to fortify the robustness of their infrastructure [VIX01 2]. Backbone services can cause significant network outages. This would include DNS and to a lesser extent RADIUS.

Traffic on network elements can be classified into the data plane, control plane and the management plane. The data plane comprises packets that are forwarded through the router to another destination. The control plane consists of the routing protocols that allow the network to properly function. The management plane addresses the tools and protocols used to manage the network elements [GRE03 7].

Control and management plane traffic are sent directly to the router's processing engine. Route engines are not designed to handle the large volumes of traffic that pass through the data plane. Without proper protection, these CPUs are easily overwhelmed.

3.8.1 Control Plane Attacks

Control plane attacks are attacks that are directed against the control plane of network elements, such as routers and switches. Attacks are usually directed at dynamic routing protocols such as BGP, OSPF, and EIGRP. ISIS is not as vulnerable to public attacks because it operates using the OSI protocol stack instead of the TCP/IP protocol stack and is an IGP routing protocol.

Direct DoS/DDoS attacks against the routing protocols can lead to regional outages. Another form of attack, malicious route injection, can lead to DoS attacks, traffic redirection, prefix hijacking, and AS hijacking. Prefix and AS hijacking are rare but becoming more common with hardcore SPAM operators. The control plane is often used in reflective attacks [GRE04].

3.8.2 Management Plane Attacks

The management plane allows network operators the ability to configure the network elements. This includes protocols such as telnet, SSH, SCP, HTTP, HTTPS, SNMP, NTP, FTP, and TFTP. If these protocols are vulnerable to attack, open to the public for access or used in an insecure manner, network elements are at risk of being broken into or DoS'd through resource exhaustion.

3.8.3 Network Services Attacks

DNS is a critical network service for operation of the Internet. Without DNS almost all end-user applications that run on the Internet stop functioning including HTTP traffic, e-mail, voice over IP (VoIP), electronic file transfers, and streaming media. The Internet DNS infrastructure is highly redundant and robust, operating as the largest distributed database in the world [FON01].

As a public service, DNS in a service provider's environment is subject to several types of attacks including direct DoS/DDoS attacks, being used as a reflective media in DRDoS attacks, through software and operating system vulnerabilities and through protocol vulnerabilities. BIND, the most widely used DNS software on the Internet has had numerous vulnerabilities discovered in its code. Cache poisoning is a well known protocol vulnerability that is mitigated by separating different DNS services. Secure DNS protocols have been proposed but they have not yet seen widespread adoption.

RADIUS, to a lesser extent, is another critical backbone service. Most service providers have extensive dialup and DSL networks that require RADIUS authentication for network access. DoS attacks against RADIUS servers could cause service outages for hundreds of thousands of consumers and small businesses.

Signaling System 7 (SS7) is becoming more common on the Internet, used in VoIP installations and ISP dial networks. SS7 is a set of protocols used for out of band signaling in the public switched telephone network (PSTN). Current SS7 to IP proprietary protocols are vulnerable to many types of attacks and firewall protection is required to protect most of them. IPSEC tunnels or out of band circuits should be considered to protect spoofing attacks between the SS7 gateway and the network access servers (NAS).

3.9 Other Attacks

Although not a direct attack, SPAM is quickly becoming a security issue on the Internet. As service providers aggressively take actions against SPAM operators, the SPAM operators are using more illegal or fringe activities to send SPAM including using worms and viruses to create mail proxy relays, prefix and AS route hijacking, asymmetrical traffic routing with spoofed IP addresses, using bogus accounts and stolen credit cards for dialup accounts and many other innovative techniques.

4 Prevention

Knowledge of your enemies' tactics and methods is a fundamental key in implementing methods to prevent attacks. No service provider will be able to prevent all attacks. The goal is to raise the bar for people to launch attacks.

4.1 Policies and Procedures

Security policies and procedures should be developed and in place to ensure that company and best practices are followed. Security policies are a very important part of a service provider's overall security architecture and are critical for stopping abusive users. A service provider's Acceptable Use Policy (AUP) is a key tool for removing abusive customers from their network. A more thorough discussion of security policies can be found at SANS [SANS01].

Service providers should also establish an Incident Response Team (IRT) that is responsible for responding to attacks. The IRT should develop procedures concerning:

- Who should be notified?
- What data needs to be collected (for possible law-enforcement action, later)?
- What responsive measures should be employed to protect the infrastructure or service?
- What is the escalation path for critical decisions?

The complete topic of incident response is outside the scope of the paper.

4.2 New Product/Upgrade Design and Testing

The first line of defense is security design and thorough testing of new or significantly upgraded products, services or platforms before a system is deployed in the production network. Security should be considered from the start of the system design. Things to consider include:

- Operating system lockdown and removal of any unnecessary processes, services and software. This should be done via scripts or by checklists preferably developed using industry best practices.
- Review of system protocols to ensure communication paths are properly authenticated and if necessary encrypted.
- Scanning of the systems to confirm and mitigate, if necessary, any security risks found.
- If software source code is available, security source code reviews should be performed to eliminate buffer overflows and other vulnerabilities.

Once a final architecture has been determined, a complete security review and penetration testing should be performed before the system is deployed.

4.3 Patch Management

Manual or automated procedures should be in place to address the ever increasing load of patch management on servers and network elements. Clear escalation procedures should be in place to address critical patch deployment or mitigation policies. Care needs to be taken as installation of patches can leave a system open to new or previously mitigated vulnerabilities when configuration files are replaced that were previously secured.

4.4 Scanning/Auditing

On-going scanning and auditing of servers and network elements is a critical part of network security management. Configuration management is a difficult task in a large network with hundreds of people making changes on different parts of the network. Today, many configuration changes are still done manually, allowing for human errors and inconsistent configurations.

Scanning should begin by focusing on the most critical network elements and servers from an outsider's vantage point, from outside any firewalls or router ACLs. This will allow operation personnel to address the most critical vulnerabilities first. The number of network elements included in the initial scans should be limited to a reasonable size to allow personnel to fix any issues before expanding the scanning. Scanning should occur at least once a quarter.

Typical scanning tools include nmap and nessus. Scanning should include both TCP and UDP ports; however, UDP scanning can take considerably longer, especially if the scanning is done through a firewall. All scanning techniques should be tested in a lab environment for the specific version of code before performing it on production equipment. Scanning has been known to break services and even stop servers and network elements from functioning properly.

4.5 uRPF

Unicast Reverse Path Forwarding (uRPF) is a technique developed to implement BCP 38/ RFC 2827 *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing* [GRE01 1]. URPF can drop IP packets with spoofed source addresses at the service provider customer's edge, depending on the method of deployment. URPF only works on the ingress interface.

URPF can be deployed in two basic modes, strict mode and loose mode. In strict mode, the router takes the source address and validates that it has a prefix that exists in

the forwarding information base (FIB) for the interface on the router that the packet came in on. It also validates that the route for that address is in the interface adjacency table. If a route to the source does not exist, the packet is dropped [GRE02 43-46], [JUN01 94-95]. Special considerations must be taken into account if the customer is multi-homed and traffic can have asymmetrical routing patterns.

Loose mode was designed for use at the ISP-ISP edge for peering applications where best path selection limits the applicability of strict mode. Loose mode only does the first check, ensuring that the source address of the packet has a route prefix in the FIB. This allows packets that have any valid route in the FIB to pass.

ACLs or filters can be configured if the uRPF check fails to allow packets to be counted, logged, or passed based on special criteria. If DHCP or BOOTP packets are forwarded on the interface with uRPF turned on, a special filter must be used to pass these packets for Juniper implementations [JUN01 97]. Cisco automatically passes these packets in IOS versions 12.0 and later [CIS02 SC-438]. Juniper has additional configuration options to allow only active paths or all feasible paths to be considered in the uRPF verification decision.

4.5.1 Cisco Implementation

The Cisco implementation of uRPF requires the use of Cisco Express Forwarding (CEF). Strict mode uRPF is available in IOS versions 11.1(17)CC and 12.0 and later. It is not available in IOS Release 11.2 or 11.3 [CIS02 SC-438]. Below is a sample configuration of a strict mode Cisco configuration.

```
ip cef
!
interface serial X/X/X
ip verify unicast reverse-path <acl>
```

The optional *acl* is available on the 10K platform beginning with 12.0.25S.

Input ACLs are first processed on an incoming packet. Only after the input ACLs does the router perform the uRPF check. Then a CEF lookup is done to determine the output interface for the packet. The output interface ACLs are then checked and if passed, then the packet is forwarded. If any check fails, the packet is discarded [CIS03 77].

Loose mode uRPF works by only verifying the FIB entry. This will prevent bogus addresses from passing the interface but other spoofed addresses will not be blocked.

The new IOS syntax supporting loose mode is:

```
ip verify unicast source reachable-via (rx|any) [allow-default] [allow-self-ping] [<list>]
```

The *rx* option checks that the packet has a route back on the interface it came in on (strict mode). *Any* is loose mode and does not perform the interface check. However, there must be a route to a real interface, thus routes to null0 will be discarded. The *allow-default* option allows the lookup to use the default route for verification. The *allow-self-ping* is used to allow the router to ping one of its own interfaces [GRE01 7]. This became available for the GSR platform in the 12.0.22S code release.

A sample configuration for loose mode on an interface would be:

```
ip verify unicast source reachable-via any
```

For multi-homed customers with asymmetrical routing issues, Cisco recommends either using loose mode uRPF or using BGP weights to create symmetrical routing paths.

Below is a copy of Cisco's recommended implementation methods for different ISP configurations [GRE01 16]:

Deployment Situation	Type of uRPF to use	Config Notes
Lease Line Customer	Strict Check	
Multi-homed Lease Line Customer (same ISP)	Strict Check or Loose Check	Remember to use BGP Weights on Strict Check
Multi-homed Lease Line Customer (different ISPs)	Strict Check or Loose Check	Remember to use BGP Weights on Strict Check
Dial-up Customers	Strict Check	
DSL Customers	Strict Check	
Cable Modem Customers	Strict Check	
IXP Connection – no private peering	Strict Check	
IXP Connection w/ private peering	Loose Check	
Private Peering – Dedicated Router	Strict Check	Symmetry should be expected between the routes advertised and source addresses sent by the peering ISP.
Private Peering /w several ISPs on the same router	Loose Check	
Co-Location Provider's Edge Routers	Loose Check	

4.5.2 Juniper Implementation

Juniper routers support strict and loose mode configuration per interface. In addition, Juniper supports the concept of active paths and feasible paths. Active or feasible paths option is set for the entire router in the routing options section and is not configurable per interface.

Four combinations of options are available on the Juniper platform:

Juniper Combination	Cisco Equivalent	Comments
Strict mode – Active paths	Strict mode	
Strict mode – Feasible paths	Roughly equivalent to Cisco's strict mode with BGP weights	Could be used with multi-homed customers without having to use BGP weights if customer announces all prefixes on all interfaces.
Loose mode – Active paths	Loose mode	
Loose mode – Feasible paths	None	Only applicable for routers configured with feasible path strict mode.

Juniper does not provide any implementation guidelines for using uRPF on their platforms. Thorough lab testing is required before deployment of uRPF on Juniper M and T series routers to fully understand the operational impacts on complicated customer configurations.

Active or feasible path options are set in the routing options section. This applies to all interfaces on the router and determines how the forwarding table is built. Below is a sample implementation [JUN02]:

```
[edit]
routing-options {
  forwarding-table {
    unicast-reverse-path (active-paths | feasible-paths);
  }
}
```

Use the command below to turn on uRPF on a specific interface:

```
[edit interfaces interface-name unit logical-unit-number family (inet |
inet6)]
  rpf-check <fail-filter filter-name> {
```

```
mode (loose|strict);  
}
```

Strict mode is the default if a mode is not specified.

The option not to have uRPF check the default route if one is set on the router is not available on Juniper routers. Thus, if a default route is set and uRPF is turned on in loose mode, all packets will be accepted. Juniper advises not to use loose mode with a default route. In strict mode, a Juniper will accept all packets if the default route is set on the interface using uRPF, contrary to their documentation [JUN01 96].

Juniper allows failed filters for special cases using uRPF. For example, if DHCP or BOOTP packets need to be forwarded through the interface using strict mode uRPF, a fail filter is required to allow the packets.

4.6 Management and Control Plane Protection

Protection of the management and control planes is critical for the successful operation of an ISP. It is easier to discuss both topics together because the router configuration to protect both is similar in many ways. Authenticated and encrypted protocols are preferred for router management. Protocols must be accepted only from trusted hosts. Steps to protect the control plane include: protection of the route engine using filters, authentication and integrity verification of routing protocol updates, rate limiting of diagnostic protocols [PEJ01 5-19] and filtering of routing prefix updates sent from customers and peers [GRE04].

4.6.1 Router Access

Management access to the router is one of the most important aspects of router security. Access can be done either through in-band communication paths, out-of-band paths, or a combination of both. In-band router management is very common but extra care must be given when using this communication path. Systems that are primarily managed in-band must have a backup, out-of-band (OOB) method to update the router configuration. Usually, this is a dial-up modem connection to the routers' console ports. OOB systems should also have back connection methods, such as dial-up modems.

For medium to large service providers, a good authentication, authorization and auditing (AAA) system is essential for network operations. The AAA system will provide central account management, different authorization levels and an audit trail for system logins. Two factor authentication using one-time passwords provides a very secure form of authentication and is required for secure operations of clear text management protocols (telnet for example). RADIUS and TACACS+ are good protocols used by AAA systems and are supported by most router and server vendors. Both support common 2 factor authentication methods.

If possible, secure shell and secure copy should be used for router access and configuration updates. If telnet is used, one time passwords must be used for authentication into the router. If TFTP is used to upload or download router configurations, very tight access lists must be in place to prevent unauthorized access.

SNMP write access should only be configured if SNMP version 3 is used with encryption. SNMP version 1 can be used for read only access if control lists are placed on the routers to limit the polling hosts. Weak or default passwords should not be used since the configurations of the routers usually contain weakly encrypted passwords.

NTP should be configured on all routers to correlate events and audit logs. NTP can be configured using MD5 hash authentication and integrity checking and should be used to accept updates only from trusted sources [PEJ01 20].

4.6.2 Router Engine Protection

Router engines have limited bandwidth and resources compared with the data plane that they control. The router engine should be protected from mistrusted sources to limit resource exhaustion attacks on the router itself and to limit reflective attacks from the router. Only required services and protocols should be turned on.

Dynamic routing protocols should be configured to use strong authentication and integrity mechanisms such as MD5 hash whenever possible. Several study groups are evaluating enhancements to better secure BGP but these have not seen any large deployments to date.

Protocols that should be filtered to the router engine include dynamic routing protocols (BGP, OSPF, RSVP, etc), management protocols and services (SSH, telnet, SNMP, NTP, DNS, TFTP, etc), and diagnostic protocols (ICMP, traceroute, etc). ICMP and TCP SYN packets should be rate limited where possible to protect against route engine flooding attacks. Other filters can be set to guard against fragmentation attacks, inappropriate IP options, and using the routers for reflective attacks [PEF01 11-16].

Juniper filters packets to the route engine by applying firewall filters to the loop back interface [PEJ01 11]. Trusted source lists should be created for the different protocols and services. Firewall filters and rate limiters can then be created based on those lists.

Cisco has implemented receive ACLs (rACL) to protect the routing engine. Currently this is only available on their GSR platform. RACLs perform the same basic function as firewall filters to the loop back address on Junipers but currently it cannot rate limit packets. Cisco intends to add rate limiting and extend support for rACLs to the 7500, 10K and other Cisco platforms [GRE03 126-134].

An innovative technique was proposed to limit the BGP packets accepted by a high valued TTL [GIL01]. Since most BGP updates come from adjacent routers, a TTL filter on the router would only accept packets that had a very limited TTL range (253-255 for example). Router vendor support is needed before this technique can be implemented in practice.

4.6.3 Prefix Filtering

Prefix filtering is a requirement for ISP operations both at the customer edge and the peering edge to protect the control plane. The best practice is both ingress and egress filtering. Prefix filtering prevents the ISP from accepting unexpected routes and redirecting traffic. Due to the nature of routing protocols, a longer prefix takes precedence over a shorter prefix. If a customer or peer is either accidentally or maliciously able to inject an unauthorized route, legitimate traffic can be disrupted. The lack of prefix filtering caused the 1997 outage on most of the Internet [FLI01].

For malicious intentions, traffic from a well known web site can be redirected to the attacker's location with the injection of a longer prefix route, accepted by the upstream provider. Spamming operators are using prefix hijacking to route unused IP blocks over their links for bulk unsolicited e-mail delivery.

Prefix filtering should include all IP addresses that should not be routed on the Internet. This includes private address space, test and reserved addresses and unallocated blocks, otherwise known as bogons. RFC 3330 documents the current special use IPv4 addresses [RFC01]. In addition, the good people at Cymru maintain a list of current bogons in a wide variety of formats [CYM01]. Example ingress prefix templates for both Cisco and Juniper routers are available on the Internet [GRE05] [GRE06].

4.7 FW/IDS/IPS

Firewalls, Intrusion Detection Systems (IDS), and Intrusion Prevention Systems (IPS) can be useful devices for protecting backbone services. All servers exposed on the Internet should have all non-essential services turned off and some type of host based firewall installed on the system. Separate network based firewalls can also be installed but the cost of the systems can outweigh the benefit. If firewalls are deployed, an IDS behind the firewall should be considered to monitor for unauthorized activity.

4.7.1 DNS Considerations

One use for IPS systems is to protect DNS servers as the rate of false positive reactions will be much lower than for multi-purpose systems. An IPS can also be used to deny legitimate but unwanted traffic to DNS servers to reduce the load on the servers. An example of this type of filtering would be to stop dynamic DNS update requests from Windows 2000 and XP computers.

Another useful, but currently unavailable, tool to help mitigate DoS/DDoS attacks would be a rate limiting function based on flow. Flow rates that exceed a specified value for a specified time could be rate limited lower until their flow dropped below a threshold. This way service could be limited, but not cut off completely, protecting other customers' use of the service.

4.7.2 Other Services

Both RADIUS and SS7 Gateway servers should be protected with firewalls and possibly IDS solutions. Packet and host based firewalls should be installed as a minimum to protect the servers. If performance will allow, network based firewalls can add another layer of protection.

Many SS7 gateways run insecure, proprietary protocols which are subject to many attacks listed above. VPNs should be considered to protect the traffic at these gateways for large dial networks.

5 Monitoring

The next step in DoS/DDoS protection is monitoring for attacks. It is difficult to mitigate an attack without good information about the characteristics of the attack. The mitigation techniques used will depend on the level of pain and inconvenience your customer is willing to put up with.

5.1 Customer/Peer Notification

The most basic monitoring is customer or peer notification. Many times, the notification from customers is very incomplete and inaccurate. The customer may only know their connectivity is having trouble and they want their ISP to fix the problem. Notification from peers usually has better attack information. Mailing lists have recently evolved to allow ISPs to coordinate their efforts and find operational contacts to mitigate attacks in progress. These lists include service providers and vendors.

5.2 Sinkhole

Sinkholes are used to direct and trap traffic in a service provider's network. A sinkhole can be used to monitor the scanning of bogon or "dark" IP space for worms, viruses or probing activity. Sinkholes can also be used to redirect an attack against a customer to the sinkhole for traffic analysis. Sinkhole techniques can be used in conjunction with black hole filtering to analyze spoofed DoS/DDoS attacks.

The most basic sinkhole consists of a router in the service provider's network that advertises out a /32 network block for the customer under attack. This advertisement routes all the traffic destined to that single IP address to the sinkhole. The traffic for the

sinkhole can be analyzed by a traffic sniffer to determine the type of attack. The router should be configured to prevent backscatter traffic from leaving the sinkhole. Backscatter is the traffic generated by a router or server in response to the incoming traffic. These responses can include ICMP echo replies, unreachables, etc.

Bogon address space can be advertised by the sinkhole to monitor for worms, viruses or network probing. There are some important considerations to remember advertising out bogon address to a sinkhole. First, it is important the bogon advertisement does not leave your network. Bogons should be advertised in the IGP, most commonly using iBGP. The BGP “no-export” community and egress prefix filtering policies will prevent bogon advertisements from reaching the Internet at large. Second, your sinkhole should be designed to handle a large in-flood of traffic [GRE07 57-78].

A service provider can also advertise the default route for the network to the sinkhole. When this is done, all the unwanted traffic without a specific route will head down the sinkhole including customer traffic when circuits flap, network scans to bogon address space, failed attacks, backscatter from spoofed attacks, misrouted traffic and more [GRE07 62]. Anycast, discussed later in this paper, can be used to distribute the load of a sinkhole to many sites [GRE07 62, 87-95].

5.3 Backscatter Techniques

Backscatter analysis can be used to monitor the level of spoofed DoS/DDoS activity on the Internet. A study performed by the CAIDA organization in 2001 recorded 12,805 attacks in one week of monitoring. The technique watches for return packets from the victim to the spoofed addresses. These include ICMP echo response, port unreachable, timestamp reply and TCP ACK and RST packets [MOO02 3, 6].

Backscattering trace back is a technique used to quickly determine the entry points of a DDoS into an ISP's network without the expense of special monitoring systems or netflow data collectors. Developed by UU Net, the technique uses a combination of sinkhole monitoring and black hole filtering [MOR01]. Black hole filtering is a technique to drop all network traffic at the ingress points into the service provider's network and is discussed below in the mitigation section.

The backscatter technique uses the sinkhole to watch the backscatter traffic generated from the routers where the spoofed attack traffic was black holed. The source address of the backscatter traffic is the ingress point. The technique requires some preparation in the setup of the black hole filtering and the sinkhole. Once the setup is complete, the origin of a spoofed attack can be traced back in minutes instead of hours. The technique is limited to spoofed attacks. If the attack is not spoofed, nothing will show up at the sinkhole. A DRDoS attack cannot be traced because the reflectors sending the attack are not spoofing their source address.

5.4 Netflow Monitoring

Netflow is a very useful tool in monitoring traffic patterns and DoS/DDoS attacks. Developed by Cisco in 1996, a flow is defined as having the following seven unique attributes [CIS05 5]:

- Source IP address
- Destination IP address
- Source port
- Destination port
- Layer 3 protocol type
- TOS byte (DSCP)
- Input logical interface (ifIndex)

Each unique flow is counted in the router. The flow data can be exported to a separate collection and correlation system. Netflow is unidirectional and is currently only available on the router ingress interface. To monitor traffic in both directions all router interfaces must be monitored, including uplinks to the core routers.

Netflow comes in two varieties, regular netflow and sampled netflow. Regular and sampled netflow are available on all Juniper M and T series routers, with a limitation of roughly 7,000 netflow packets per second if run in software. Regular netflow is available on all Cisco IOS platforms. Sampled netflow is not available on many platforms including the 7200 and 6500 but is available on the GSR.

There is concern about the statistical accuracy of sampled netflow versus the performance and bandwidth requirements. Internet Protocol Flow Information eXport (IPFIX) is an IETF Working Group focused on defining a standard netflow export format and with the statistical accuracy of sampled data. One goal of the group is to “consider the notion of IP flow information export based upon packet sampling.” [IPF01]

Netflow data can be exported in several formats. Cisco has defined formats for version 1, 5, 7, 8 and 9. Version 5 is the most common format used today. Version 5 is a fixed format and is not extensible. Version 9 was introduced to allow extensibility to define the characteristics of flows, the exported data format and allowing extensibility for MPLS, multicast and BGP next hop. Cisco has submitted the format to the IETF for consideration as a standard [CIS05 13-30].

5.4.1 Netflow DoS/DDoS Considerations

When used to monitor DoS/DDoS attacks, netflow is most useful around the perimeter of the service provider’s network. Netflow monitoring on peering and edge links allows the operator to determine the ingress point of spoofed DoS/DDoS attacks. Detailed flow

characteristics can be determined to write ACLs to block only the attack, with minimal impact on legitimate traffic. Flow data from multiple collection points can be correlated to identify the network ingress points for DDoS attacks and quickly determine the attack characteristics.

Several companies build DoS/DDoS correlation systems using netflow data including Arbor Networks and Mazu Networks. Arbor is currently the leader in passive, scalable DDoS/DoS collection, correlation and reporting. The Arbor system will correlate DoS/DDoS attacks across all the netflow collection points to a centralized reporting system. Traffic analysis can also be performed by combining BGP information with the netflow data.

The “show ip cache flow” command can be used on Cisco IOS routers to show the flows with the highest volumes when directly on a router. Cflowd is a free software program available to analyze netflow data.

6 Mitigation

When a customer or the network infrastructure is under attack, monitoring is important for quick identification of the attack characteristics and entry points but the next question that immediately follows is, “What are you going to do to stop it?” Good mitigation techniques are a required part of a service provider’s security architecture.

6.1 ACLs/Rate Limiting

Access control lists (ACL) or firewall filters are the first line of defense for a service provider. For a simple DoS attack directed at a single customer, deployment of an egress ACL on the customer’s edge router is an easy way to stop the attack. The problem with this technique is scaling both from a router performance perspective and as the number of attacks managed increases.

Operation personnel deploying the ACLs must know the performance limitation of the routers they are using. ASIC based ACLs will perform better than ACLs processed in software. Different ASICs can and do have different performance characteristics based on the packet size, interface speed and other features turned on in the router and interface cards. Juniper has less hardware performance differences across their platform than Cisco does today. However, as their product line matures and hardware ages, similar differences in performance are bound to occur.

The management of a large number of temporary ACLs that may have performance impacts on different router hardware and software is non-trivial and can be very labor intensive and error prone. Most service providers have home grown scripts for their router configuration and ACL management.

ACLs can and do play an important role in mitigating Internet scale attacks and vulnerabilities. Most major service providers deployed ACLs blocking UDP port 1434 on their core and peering routers to shutdown the spread of the Slammer worm. ACLs played a critical role in preventing exploitation of the Cisco IOS vulnerability for IP protocols 53, 55, 77 and 103, preventing significant Internet outages [CIS04].

From an operational perspective, network wide ACLs are easier to deploy (although it may not seem like it at the time) than to remove from the network. Once an ISP's customer base is protected by ACLs, it is prudent to provide customer notifications with patch information and lead times before removing the ACLs. Traffic loads must be monitored as the ACLs are removed to ensure that worm traffic from unpatched customers does not have a significant impact on other customers or the provider's backbone. Legitimate customer traffic may also be blocked by the ACLs and support organizations must be notified and prepared to answer customer's questions and complaints.

6.1.1 Rate Limiting

Rate limiting on the data plane is another important mitigation tool. Rate limiting can become very important when all traffic to a site cannot be blocked. The same issues discussed above also apply to rate limiting. Remote triggered rate limiting is another possibility and available on a limited number of Cisco platforms.

6.1.2 Other

Different queuing mechanisms are available to prioritize traffic. Care must be taken to ensure that these don't become avenues of DoS attacks with forged information.

6.2 Destination based Black Hole Filtering

Black hole filtering is an effective, quick and simple technique for dropping attack traffic destined toward a victim. Using iBGP as a trigger mechanism, black hole filtering can be remotely triggered across the entire perimeter of a provider's network. This technique is used when more harm is done by the attack filling up a customer's circuit than by the loss of an individual site. Many times, traffic can be redirected to a different IP address through DNS.

Preparation is the key to remote triggered black hole filtering and setup takes some effort. Each perimeter router that will participate needs a one time configuration setup to statically route a small bogon block of IP addresses to the null or drop interface on the router as a next hop. The route should be tagged to match a specific BGP community string. A trigger router is configured to inject static routes into iBGP. The router is configured with the BGP "no-export" and "IGP" parameters.

When an attack occurs, a static route is added to the trigger router to route the /32 IP address under attack to the bogon address block configured in the perimeter routers. The route is injected into iBGP and distributed to all routers in the network. The traffic for the attack is black holed at each ingress router to the network, effectively stopping the attack. This type of black hole filtering is only good to drop traffic based on the destination address [GRE 11-29].

Several variations of remotely triggered black hole filtering can be setup. By using different community strings, remote triggers can be setup for different types of routers such as edge and border. Community strings can be setup for different geographic regions or POPs in a provider's network. This flexibility allows the provider to identify the ingress points of the attack and only block traffic at those locations.

6.2.1 Customer Initiated Black Hole Filtering

Another variation of black hole filtering developed and in use at UU Net is customer initiated, remotely triggered, black hole filtering. In this variation, a BGP peering customer initiates the black hole by a BGP announcement from their AS. Ingress filters at the customer's edge router must be configured to let the bogon prefix announcement into the provider's network and allow /32 route announcements for that BGP community. The customer must be sophisticated enough to identify the attack and know the proper mitigation steps [MOR02].

This method allows customers to mitigate their own attacks at the service provider's ingress points without any direct contact with the service provider's network operation center (NOC).

6.2.2 Source based Black Hole Filtering

A promising technique developed by Cisco is source based, remotely triggered, black hole filtering. It combines uRPF with destination based black hole filtering. URPF must be deployed on all the routers needed to black hole the attack. This includes peering routers. The technique works by a modification to uRPF. When uRPF checks the source address to make sure a prefix exists in the router's FIB, uRPF checks that the route back is on a real interface. If the route is to the null0 interface, the packet is dropped.

The setup and trigger mechanism for source based, black hole filtering is identical to destination based black hole filtering with the addition that uRPF (either strict or loose mode) must be configured on the ingress interface of the perimeter router [GRE07 30-44].

6.2.3 Black Hole Shunting

Black hole shunting is another variation on the black hole filtering configuration. The difference is that instead of sending the traffic to the null0 or drop interface, the traffic is sent out a different physical interface. A data scrubber residing on the alternate data path can filter out the attack traffic from the good customer traffic and send the clean traffic to the customer. The Riverhead Guard is a data scrubber that uses this technique to mitigate DoS/DDoS attacks [RIV01].

6.2.4 Advanced BGP Filtering

A draft RFC, co-authored by Juniper, Cisco, Verio and Arbor Networks could provide much finer granularity in BGP propagated traffic filtering. The draft specifies detailed packet information that can be put into BGP to allow filtering of complex DDoS attacks [MAR01]. The detailed attack ACLs could be defined in conjunction with a traffic monitor, sinkhole or IDS that would be distributed to the ingress perimeter routers and the attack dropped at the provider's edge. Other legitimate traffic would pass without interruption. The ACLs would be centrally managed at the BGP injection router, similar to remote triggered, black hole filtering. Software and possibly hardware upgrades would be required to implement this feature from the router vendors.

6.3 Attack Distribution and/or Isolation – Anycast

IPv4 anycast implementations have been in use on the Internet for at least the past 10 years. Particularly suited for single response UDP queries, DNS anycast architectures are in use in most tier 1 Internet providers' backbones. Anycast implementations can be used for both DNS authoritative and recursive implementations. Several root name servers are implementing anycast architectures to mitigate DDoS attacks [ALB01]. Black hole filtering is a specialized form of anycast. Sinkholes can use anycast to distribute the load of an attack across many locations [GRE07 86-97].

Many DNS anycast implementations are done using eBGP announcements. Anycast networks can be contained in a single AS or span multiple AS's across the globe. Anycast provides two distinct advantages in regard to DoS/DDoS attacks. In a DoS attack, anycast localizes the effect of the attack. In a DDoS attack, the attack is distributed over a much larger number of servers, distributing the load of the attack and allowing the service to better withstand it.

The main disadvantage of an anycast implementation is brownout conditions. This is where the server is still functioning but running at full capacity. Some legitimate queries go unanswered due to resource exhaustion. This may be due to a DoS/DDoS attack or failure of a neighboring anycast server without adequate reserve capacity. If this resource is taken fully off-line and queries are redirected through anycast to the next server, a cascading effect can result taking down the entire service. To prevent this

from occurring, a true secondary anycast system is needed, separate from the primary anycast. This allows one area to failover to an independent anycast system.

Due diligence is needed when setting up and maintaining an eBGP anycast system. All BGP routing parameters are set the same for each anycast site. If a configuration error is made on a site to lower its routing preference relative to the others, it will act as a magnet for the traffic and the entire service can go down (as long as that route is advertised).

7 Conclusion

Attack techniques continue to advance and the number of software vulnerabilities continues to increase, without regard to the dot com bubble bursting. Internet worms that previously took days or weeks to spread now take minutes. Service providers and vendors are quickly adapting to the new landscape. Defense in depth must be practiced by service providers as zero day exploits are released.

Prevention is always the best measure. Wide scale deployment of uRPF across service providers' networks will prevent many attacks seen today from happening. Hardening of the service providers' infrastructure with full security testing is required for continued operation. Regular scanning and auditing will prevent configuration errors from exposing infrastructure to known attacks.

Automated DoS/DDoS monitoring and reporting will become the standard for service providers as reaction times have gone from days to minutes. Preparation is the key for service providers to mitigate attacks as they happen. The Internet is maturing as companies become more dependent on its use. Customers are beginning to expect the same reliability from the Internet as other critical infrastructures: PSTN, power and water. Vendors and service providers are meeting the challenge head on with a high level of cooperation and innovation.

8 References

[ABL01] Abley, Joe. "ISC Technical Note Series, Hierarchical Anycast for Global Service Distribution." 2003. Internet Software Consortium, 17 Aug. 2003. <<http://www.isc.org/tn/isc-tn-2003-1.html>>.

[BER01] Bernstein, Daniel. "SYN Cookies." 29 July 2003. <<http://cr.yp.to/syncookies.html>>.

[CIS01] "Cisco IOS IP Command Reference, Volume 1 of 4: Addressing and Services, Release 12.3." Cisco IOS Release 12.3 Configuration Guides and Command References. 28 July 2003. <http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipras_r/ip1_i1g.htm#1081245>.

[CIS02] "Configuring Unicast Reverse Path Forwarding." Cisco IOS Security Configuration Guide, Release 12.2. 17 Aug. 2003. <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_c/fot_hersf/scfrpf.pdf>.

[CIS03] "Cisco ISP Essentials, Essential IOS Features Every ISP Should Consider, Version 2.9." 17 Aug 2003. <<http://www.cisco.com/public/cons/isp/documents/IOSEssentialsPDF.zip>>

[CIS04] "Cisco Security Advisory: Cisco IOS Interface Blocked by IPv4 Packets, Document ID: 44020, Revision 1.13." 2 Aug. 2003. Cisco Product Security Advisories and Notices. 17 Aug. 2003. <<http://www.cisco.com/warp/public/707/cisco-sa-20030717-blocked.shtml>>.

[CIS05] "Netflow Overview." Feb. 2003. 17 Aug. 2003. <http://www.cisco.com/application/vnd.ms-powerpoint/en/us/guest/tech/tk362/c1482/ccmigration_09186a0080182b50.ppt>.

[COL01] Cole, Eric, Mathew Newfield, and John M. Millican. GSEC, Security Essentials Toolkit. Que Publishing, 2002.

[CYM01] "The Bogon Reference Page." 18 Aug. 2003. <<http://www.cymru.com/Bogons/>>.

[DIT01] Dittrich, Dave. "Distributed Denial of Service (DDoS) Attacks/tools." 12 Aug. 2003. 19 Aug. 2003. <<http://staff.washington.edu/dittrich/misc/ddos/>>.

[FLI01] "The Day the Internet Died - Courtesy of the Florida Internet Exchange." 18 Aug. 2003. <<http://flix.flirble.org/>>.

[FON01] "DNS Is Busting Out All Over." 13 Aug. 2003. TechWorldNews. 18 Aug. 2003. <<http://www.technewsworld.com/perl/story/31330.html>>.

[GIL01] Gill, Vijay, John Heasley, and David Meyer. "The BGP TTL Security Hack (BTSH)." May 2003. 18 Aug. 2003. <<http://www.ietf.org/internet-drafts/draft-gill-btsh-02.txt>>.

[GRE01] Greene, Barry Raveendran, Neil Jarvis. "Unicast Reverse Path Forwarding (uRPF) Enhancements for the ISP-ISP Edge, Version 1.5." 18 Feb. 2001. 31 July 2003 <<ftp://ftp-eng.cisco.com/cons/isp/security/URPF-ISP.pdf>>.

[GRE02] Greene, Barry Raveendran. "Phase 1 – Preparation for the Attack, Securing the Network and Data Plane." ISP Security Bootcamp Singapore 2003. 31 July 2003 <<ftp://ftp-eng.cisco.com/cons/isp/security/ISP-Security-Bootcamp-Singapore-2003/G-Preparation-DataPlane-BCP38-v3-0.pdf>>.

[GRE03] Greene, Barry Raveendran. "Phase 1 – Preparation for the Attack, Securing the Router and the Management Plane." ISP Security Bootcamp Singapore 2003. 31 July 2003 <<ftp://ftp-eng.cisco.com/cons/isp/security/ISP-Security-Bootcamp-Singapore-2003/E-Preparation-MgtPlane-v3-0.pdf>>.

[GRE04] Greene, Barry Raveendran. "Phase 1 – Preparation for the Attack, Securing the Router Protocol and the Control Plane." ISP Security Bootcamp Singapore 2003. 31 July 2003 <<ftp://ftp-eng.cisco.com/cons/isp/security/ISP-Security-Bootcamp-Singapore-2003/F-Preparation-CtlPlane-v3-0.pdf>>.

[GRE05] Greene, Barry. "Loose Mode Ingress Prefix Filter Template, Version 4." 8 Apr. 2003. 18 Aug. 2003. <<ftp://ftp-eng.cisco.com/cons/isp/security/Ingress-Prefix-Filter-Templates/T-ip-prefix-filter-ingress-loose-check-v4.txt>>.

[GRE06] Greene, Barry, and Stephen Gill. "JUNOS Loose ISP Prefix Filter Template, v. 1.5." 8 May 2003. 18 Aug 2003. <<http://www.qorbit.net/documents/junos-isp-prefix-filter-loose.pdf>>.

[GRE07] Greene, Barry Raveendran. "Phase 1 – Prepare the Tools and Techniques, Using IP Routing as a Security Tool." ISP Security Bootcamp Singapore 2003. 31 July 2003 <<ftp://ftp-eng.cisco.com/cons/isp/security/ISP-Security-Bootcamp-Singapore-2003/H-Preparation-Tools-v3-0.pdf>>.

[HAC01] Hackpalace Web Site. 19 Aug. 2003. <<http://www.hackpalace.com/usa/>>.

[IAN01] "Protocol Numbers." 13 Jan. 2003. Internet Assigned Numbers Authority. 17 Aug. 2003. <<http://www.iana.org/assignments/protocol-numbers>>.

[IPF01] "IP Flow Information Export (ipfix)" 08 July 2003. IPFIX Working Group Charter. 17 Aug. 2003. <<http://www.ietf.org/html.charters/ipfix-charter.html>>.

[JUN01] "JUNOS Internet Software Configuration Guide, Interfaces and Class of Service, Release 5.5." 17 Aug. 2003. <<http://www.juniper.net/techpubs/software/junos/junos55/swconfig55-interfaces/download/swconfig55-interfaces.pdf>>.

[JUN02] "Enable Unicast Reverse Path-Forwarding Check." JUNOS 5.5 Internet Software Configuration Guide: Routing and Routing Protocols. 17 Aug 2003. <<http://www.juniper.net/techpubs/software/junos/junos55/swconfig55-routing/html/routing-generic-config11.html>>.

[MAR01] Marques, Pedro, Nischal Sheth, Robert Raszuk, Jared Mauch, and Danny McPherson. "Dissemination of flow specification rules, draft-marques-idr-flow-spec-00.txt." June 2003. The Internet Engineering Task Force, Internet-Drafts. 14 Jul. 2003. <<http://www.ietf.org/internet-drafts/draft-marques-idr-flow-spec-00.txt>>.

[MOO01] Moore, David, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, and Nicholas Weaver. "The Spread of the Sapphire/Slammer Worm." 2003. Cooperative Association for Internet Data Analysis (CAIDA). 14 July 2003. <<http://www.caida.org/outreach/papers/2003/sapphire/sapphire.html>>.

[MOO02] Moore, David, Geoffrey M. Voelker, and Stefan Savage. "Inferring Internet Denial-of-Service Activity." Proceedings of the 10th Usenet Security Symposium, Washington, D.C. 13-17 Aug. 2001. USENIX Association. 14 July 2003. <<http://www.usenix.org/publications/library/proceedings/sec01/moore/moore.pdf>>.

[MOR01] Morrow, Chris, and Brian Gemberling. "BlackHole Route Server and Tracking Traffic on an IP Network." 19 Aug. 2003. <<http://www.secsup.org/Tracking/>>.

[MOR02] Morrow, Chris, and Brian Gemberling. "How to Allow Customers to BlackHole their own traffic." 19 Aug. 2003. <<http://www.secsup.org/CustomerBlackHole/>>.

[PAX01] Paxson, Vern. "An Analysis of Using Reflectors for Distributed Denial-of-Service Attacks." 19 Aug. 2003. <<http://www.icir.org/vern/papers/reflectors.CCR.01.pdf>>.

[PEJ01] Peymani, Pejhan, and Matt Kolon. "JUNOS Router Security, Best Common Practices for Hardening the Infrastructure." May 2003. 18 Aug. 2003. <http://www.juniper.net/solutions/literature/app_note/350013.pdf>.

[PHE01] FX of Pheneolit. "cdoor.c, packet coded backdoor, version 1.3." 13 June 2000. 19 Aug. 2003. <<http://www.phenoelit.de/stuff/cd00r.c>>.

[RFC01] "Request for Comments: 3330, Special-Use IPv4 Addresses." Sep. 2002. 17 Aug. 2003. <<http://www.ietf.org/rfc/rfc3330.txt>>.

[RIV01] "Riverhead Guard." 20 Aug. 2003. <<http://www.riverhead.com/pr/guard.html>>.

[SAN01] "The SANS Security Policy Project." 19 Aug. 2003. <<http://www.sans.org/resources/policies/>>.

[SEC01] "FW-1 IP Fragmentation vulnerability (remote DoS)." 6 June 2000. Beyond Security. 17 Aug. 2003. <<http://www.securiteam.com/securitynews/5NP010A1YI.html>>.

[STA01] Staniford, Stuart, Vern Paxson, and Nicholas Weaver. "How to Own the Internet in Your Spare Time." Proceedings of the 11th Usenet Security Symposium, San Francisco, CA. 5-9 Aug. 2002. USENIX Association. 14 July 2003. <<http://www.cs.berkeley.edu/~nweaver/cdc.web/cdc.pdf>>.

[VIX01] Vixie, Paul, Gerry Sneeringer, and Mark Schleifer. "Events of 21-Oct-2002." 18 Aug. 2003. <<http://f.root-servers.org/october21.txt>>.

© SANS Institute 2003, All rights reserved.



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS San Antonio 2017	San Antonio, TXUS	Aug 06, 2017 - Aug 11, 2017	Live Event
SANS Boston 2017	Boston, MAUS	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS Hyderabad 2017	Hyderabad, IN	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS Prague 2017	Prague, CZ	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS New York City 2017	New York City, NYUS	Aug 14, 2017 - Aug 19, 2017	Live Event
SANS Salt Lake City 2017	Salt Lake City, UTUS	Aug 14, 2017 - Aug 19, 2017	Live Event
SANS Chicago 2017	Chicago, ILUS	Aug 21, 2017 - Aug 26, 2017	Live Event
SANS Adelaide 2017	Adelaide, AU	Aug 21, 2017 - Aug 26, 2017	Live Event
SANS Virginia Beach 2017	Virginia Beach, VAUS	Aug 21, 2017 - Sep 01, 2017	Live Event
SANS San Francisco Fall 2017	San Francisco, CAUS	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS Tampa - Clearwater 2017	Clearwater, FLUS	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS Network Security 2017	Las Vegas, NVUS	Sep 10, 2017 - Sep 17, 2017	Live Event
SANS Dublin 2017	Dublin, IE	Sep 11, 2017 - Sep 16, 2017	Live Event
SANS Baltimore Fall 2017	Baltimore, MDUS	Sep 25, 2017 - Sep 30, 2017	Live Event
Data Breach Summit & Training	Chicago, ILUS	Sep 25, 2017 - Oct 02, 2017	Live Event
SANS London September 2017	London, GB	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Copenhagen 2017	Copenhagen, DK	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS SEC504 at Cyber Security Week 2017	The Hague, NL	Sep 25, 2017 - Sep 30, 2017	Live Event
Rocky Mountain Fall 2017	Denver, COUS	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Oslo Autumn 2017	Oslo, NO	Oct 02, 2017 - Oct 07, 2017	Live Event
SANS DFIR Prague 2017	Prague, CZ	Oct 02, 2017 - Oct 08, 2017	Live Event
SANS Phoenix-Mesa 2017	Mesa, AZUS	Oct 09, 2017 - Oct 14, 2017	Live Event
SANS October Singapore 2017	Singapore, SG	Oct 09, 2017 - Oct 28, 2017	Live Event
SANS AUD507 (GSNA) @ Canberra 2017	Canberra, AU	Oct 09, 2017 - Oct 14, 2017	Live Event
Secure DevOps Summit & Training	Denver, COUS	Oct 10, 2017 - Oct 17, 2017	Live Event
SANS Tysons Corner Fall 2017	McLean, VAUS	Oct 14, 2017 - Oct 21, 2017	Live Event
SANS Tokyo Autumn 2017	Tokyo, JP	Oct 16, 2017 - Oct 28, 2017	Live Event
SANS Brussels Autumn 2017	Brussels, BE	Oct 16, 2017 - Oct 21, 2017	Live Event
SANS Berlin 2017	Berlin, DE	Oct 23, 2017 - Oct 28, 2017	Live Event
Security Awareness Summit & Training 2017	OnlineTNUS	Jul 31, 2017 - Aug 09, 2017	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced