



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials (Security 401)"
at <http://www.giac.org/registration/gsec>

Egress Filtering For a Better Internet

© SANS Institute 2005, Author retains full rights.

Jason Pierce
GSEC Practical Version 1.4c – Option 1
September 7, 2004

Abstract

During recent years, there has emerged a necessity for all internet users to try to stop inbound threats. Since most internet security is done from a defensive point of view, the question is left, "Can proactive internet security provide viable solutions to some of the most serious problems facing the internet today?"

The purpose of this paper is to recognize the most common problems affecting internet security and effectiveness, analyze the root cause of these problems and hopefully provide effective means of mitigation through egress filtering at the source of each problem.

Introduction

Type computer security in your favorite web search page and you will be inundated with information from firewalls and anti-virus to hardening your favorite OS. Each piece of information will provide more tips and procedures to insure your favorite OS or piece of software is as secure as possible before it is ever placed onto the internet. With all of this information so readily available, why do we continue to hear about the latest email worm spreading rapidly across the internet, another website being struck by a Distributed Denial of Service attack or the ever annoying floods of spam that everyone seems to have in their inbox.

To argue who or what is responsible for the ever increasing internet threat is beyond the scope of this document. However, what can be argued is that most of the security measures discussed, and practiced are defensive. For example, Microsoft constantly scrambles to release the latest OS patch in response to a new found security problem. How many internet service providers are taking action to stop the spread of the latest worm affecting that security problem? Should the internet community be responsible to limit the spread of worms, or spam? Is it even possible technologically or economically to expect large service providers to block the spread of these threats?

This paper will attempt to address some of the most pervasive threats that are posed today. The way these threats spread, what can be done by service providers through egress filtering to mitigate these threats, and whether these changes are possible economically and/or technologically?

Ingress/Egress Filtering – A Definition

Ingress- (noun) 1. entry into a place. ¹

Egress- (noun) 1. an exit from a place. ¹

What is egress filtering, and how does it compare to current security practices and technology?

All communications that take place on today's networks can be broken down into two basic categories. When discussing any given node on a network; traffic destined for that node is inbound or ingress, traffic leaving is outbound or egress.

Most security products today have the ability to monitor both ingress and egress traffic and apply rules based upon pre-determined criteria.

Distributed Denial of Service – (DDoS)

Overview

One common method for interrupting legitimate network traffic is called Denial of Service or DoS attacks. DoS attacks are not intended to gain unauthorized access to the system or systems being attacked.

To be stated another way, a service is any aspect of a computer system's functioning that provides benefits to a user. Any intervention that reduces or eliminates the availability of that service is called a Denial of Service, often abbreviated DoS².

There are three basic types of attack:

1. consumption of computational resources, such as bandwidth, disk space or CPU time³
2. disruption of configuration information, such as routing information³
3. disruption of physical network components³

It is also very common to launch DoS attacks from several different points on the internet. These attacks are called Distributed Denial of Service attacks or DDoS. DDoS attacks are usually carried out by a large number of slave hosts. The hosts are usually controlled by some form of virus that allows an attacker to direct attacks at unsuspecting services on a large scale in a coordinated fashion. With a properly launched DDoS attack some of the best designed and implemented web-sites have been denied service.

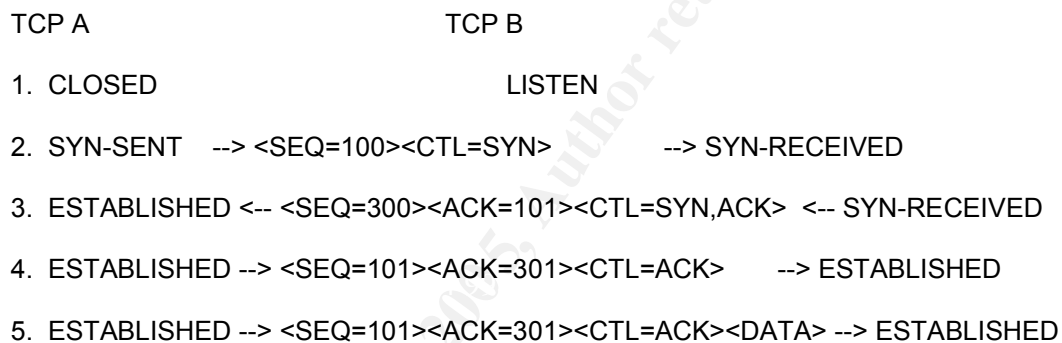
Analysis

There are several different types and methods for carrying out Denial of Service attacks. Some methods have been around for some time, others are still yet to be discovered.

Many Denial of Service attacks are successful because by design, they use legitimate traffic in large quantities to overwhelm the victim service. Because of this it becomes very hard, if not impossible in some cases to prevent or stop an

attack. For example, say website xyz.com has a server capable of accepting 1000 http requests per second. The question becomes what will happen to server xyz.com when a coordinated group of computers start sending the web server 10,000 or 100,000 requests per second. More than likely the result will be that some if not all legitimate traffic will be unable to contact the server, resulting in a Denial of Service.

To explain this further, take for instance the TCP SYN Flood DoS. This DoS is a result of how the TCP/IP protocol communicates and how the Three-way Handshake works. The three-way handshake is used to explain the method for session establishment in the TCP/IP protocol. For example let's say host A want to start a TCP session with host B. The three way handshake dictates that host A will send a SYN (synchronization) packet to host B, once received host B responds with a SYN/ACK (synchronize/acknowledge) packet. Host A then receives the SYN/ACK and responds to host B with an ACK packet. This in turn establishes the connection and data can be transferred between hosts.



Basic 3-Way Handshake for Connection Synchronization

Figure 1. ⁴

When a SYN packet arrives on a host it is stored in a queue, this queue is of a specified size and store. The TCP SYN Flood attacks this by continually sending SYN packets, but never acknowledging the SYN/ACK. As a result the queue fills up and depending on the operating system and patch level may start dropping legitimate SYN packets.

Filtering

If Denial of Service attacks are based upon legitimate traffic and legitimate services how can egress filtering help to solve the problem? The answer is that most DoS attacks are based upon spoofed or non-routable IP addresses. In other words when an attacker initiated a DoS whether from one host or distributed they forge the originating IP address. Sometimes just to another host, other times as an address reserved as non-routable by RFC 1597.

If filtering is configured at enough border routers or network aggregation points, that specify only packets with valid source address's for the originating network are permitted. It could be determined with much greater ease the source of the Denial of Service attack. Even though this may not be enough to prevent the attack from happening, if an attacker has to use only valid source addresses, it should give an administrator the information they need to block the offending network(s) and restore service to legitimate traffic.

It is strongly recommended that routers which connect enterprises to external networks are set up with appropriate packet and routing filters at both ends of the link in order to prevent packet and routing information leakage. An enterprise should also filter any private networks from inbound routing information in order to protect itself from ambiguous routing situations which can occur if routes to the private address space point outside the enterprise.⁵

Implementation

The steps required to implement filtering of spoofed source address's is fairly straight forward and can be accomplished with nearly any enterprise router of operating system.

Cisco Access List Example

```
access-list 100 deny ip 10.0.0.0 0.255.255.255 any log
access-list 100 deny ip 172.16.0.0 0.15.255.255 any log
access-list 100 deny ip 192.168.0.0 0.0.255.255 any log
access-list 100 deny ip any host 127.0.0.1 log
access-list 100 permit ip any [network IP address] [network mask] est
access-list 100 deny ip any any log
```

One of the most important things to note is that in order to be effective this type of filtering needs to be adopted on a large scale. As seen in RFC 1597 this is not a new idea, however with the rapid expansion of internet threats, administrators have the leverage required now to start pushing for the implementation of filtering at borders and aggregate network nodes.

Spam

Overview

What is spam? It is not hard to find news or information regarding spam, but the question remains for a lot of people what is spam, and why can't it be stopped.

Spam is nothing more than a slang term to explain unsolicited bulk or unsolicited commercial email (UBE or UCE). The recent influx of spam can be directly tied to economics. For direct marketers spam is a very inexpensive way to reach millions of potential customers. More often than not spammers use other persons servers, and bandwidth to send out massive amounts of email.

Spam is no longer just a nuisance it is currently costing businesses around the world billions of dollars, in wasted employee productivity. Figuring it takes 4.4 seconds on average to deal with a message, the messages add up to \$4 billion in lost productivity for U.S businesses each year⁶. A study from Symantec Corporation states that 65% of all email in July of 2004 was spam⁷.

Spammers can make a lucrative living even though only 50 in every million people respond to unsolicited commercial email⁸.

The question remains why can't spam be stopped? Countries all over the world have been drafting anti-spam legislation, software development companies are releasing products to filter spam, and almost everywhere you look someone is talking about the spam problem. To answer this question we first must analyze how spam works.

Analysis

Spam like any email relies on the SMTP protocol to travel between mail servers on the internet. SMTP the simple mail transfer protocol is a protocol that utilizes port 25 to send email from email clients to servers and then between servers.

When a user sends an email the message is sent to their mail providers email server to port 25. The mail server then determines where the message should go by comparing the email address usually to a DNS MX record. From the DNS record the mail server is able to determine the IP address of the destination mail server. It then forwards the message to that server again with destination port 25

Mail client → mail server → internet → destination mail server

Spammers are able to manipulate some of the original features in SMTP to send spam practically anonymously on the internet. When SMTP was originally developed it was built on the basis of a trusted network.

When a spammer sends a message they rarely will do so from a SMTP server owned by the spammer. Instead they use what is called open-relays to send messages. An open relay is a mail server configured to allow any host on the internet to send email through it. Open relays can be anything from an improperly configured enterprise mail server, to broadband end-users running poorly configured mail servers. Most of the time the owners of open mail relays are unaware there machines are being used to send spam.

Another feature with SMTP is the ability to forge the mail from address. Basically with SMTP the person sending the email can change the mail from and reply to address to any address they may choose.

These two aspects of SMTP make it possible for good spammers to hop throughout the internet spamming millions of addresses, possibly remaining untraceable the entire time.

Without a major overhaul to the SMTP protocol, which is unlikely to be adopted due to the vast number of machines that would have to be upgraded, these problems are likely to remain. However there is another approach that some major ISP's are starting to adopt which uses egress filtering to block SMTP to certain IP blocks.

Filtering

Comcast recently adopted egress filtering to block spammers from using port 25. However many people believe that is not enough. Major ISP's have the ability to block SMTP from all servers that are not mail servers operated by them. This approach would effectively block spammers from using huge portions of the internet.

One Comcast engineer estimated the daily email flow on the company's network at about 800 million messages, with only 100 million originating from its servers⁹.

It is my belief that internet service providers should take this approach to another level. Most legitimate email comes from mail servers that are registered with valid DNS names and MX records. If service providers were to take this one step further and only allow registered mail servers the ability to send outbound email, than spammers would be greatly limited in their ability to send mail anonymously.

Servers could be registered automatically by DNS records or another more stringent registration process could be adopted. This would not stop end-users from managing their own email server instead it would make them responsible by requiring them to register to send email on those servers.

Once only registered servers are allowed to send email, what is to stop spammers from registering there own mail server. First with a stringent registration process spammers will be more susceptible to litigation or legal problems. Second upstream providers could monitor each registered mail server for spamming practices. There is some argument of the ability for major internet service provider's ability to monitor all mail servers for acceptable use policies. This argument is valid but it is possible for providers to automatically monitor for spam. One example below is an actual email sent to an ISP from a major leased line provider.

Subject: xxxx – Abuse Documentation report – xxx.xxx.xxx.xxx

Message sent from x.x.x.x tripped an automated spam filter and was identified as an open relay. Since this is against our network AUP (acceptable use policy) this ip has been blocked until such time as are notified that the problem has been taken care of. Once this has been resolved you can contact us at xxx-xxx-xxxx to get the block removed

If all IP addresses are being automatically monitored now for spam activity then the question remains would it be more feasible for ISPs or upstream providers to only monitor groups of registered email servers for violations.

Worms

Overview

A worm is a program that makes copies of itself; for example, from one disk drive to another, or by copying itself using email or another transport mechanism¹⁰.

Worms have hit the mainstream in the last few years becoming one of the biggest problems facing the internet today. Some big name worms like Code Red, Nimda, Slammer, Blaster, and SoBig have caused billions of dollars in lost productivity and bandwidth by some accounts.

Worms can use many different methods to spread their payload. Mass-mailing worms like SoBig are very common as they do not rely on an insecure service to propagate. Other worms like Code Red and Slammer spread by taking advantage of security vulnerabilities in network services.

Mass Mailers --

Mass mailing worms make use of human nature as a means to propagate, in conjunction with an improperly configured email clients usually Microsoft Outlook, and Outlook Express. Worm writers use emails with catchy subjects or file names to entice users to execute the worm payload within an attachment.

These worms then use different means to create a list of users to email to, in turn spreading the worm to more unsuspecting victims.

SoBig.F was one such mass-mailing worm that spread across the internet recently. SoBig.F spread by finding email addresses in files on the victims' computer, and then used its own SMTP engine to propagate¹¹.

Filtering for mass-mailing worms can be directly tied to filtering for spam. Mass mailing worms rely on the fact that every computer they infect is able to send large amounts of email to different email servers on the internet. If the same filtering techniques are adopted that are mentioned to help deal with spam. These mass-mailing worms would be unable to propagate as they are designed today. It is inevitable that worm writers will always work to find ways around any filtering or security measure put in place, but every layer of defense that can be adopted will increase overall security.

Other Worms –

Worms are very dynamic and often make use of several different methods of propagation. Some worms that spread via email may also incorporate means to spread via P2P or say network shares. Because of these traits filtering for worm traffic is a very daunting task.

As worms become more destructive in nature, it is going to be even more important that service providers are able to stop or slow down the spread before massive amounts of damage can be done.

Some worms like Code Red used the HTTP service in Microsoft's IIS to spread. Since every person who uses the internet needs to have access to port 80 HTTP, filtering for that port is not a valid solution.

That doesn't mean that the worm is unstoppable without patches to the vulnerable system? What it does mean is each worm has to be analyzed by its payload and an appropriate filter built to stop its spread. Such a filter was created for Code Red shortly after it appeared in the wild and is displayed below.

Cisco Router config –

```
Class-map match-any http-hacks
Match protocol http url "*default.ida*"
Match protocol http url "*/mod_ssl:error:*"
Match protocol http url "*cmd.exe*"
Match protocol http url "*root.exe*"
Match protocol http url "*readme.eml*"
Match protocol http url "*nsiislog.dll*"
```

This example also blocks attempts by the Sadmin virus and the W32/Nimble worm. ¹²

Filtering on a per worm basis would prove to be very difficult for large organizations with current router OS technology. However this solution is viable for smaller organizations when applicable.

Sometimes though a worm replicates so fast it is impossible for human intervention to stop its spread as was the case with the Slammer worm. Slammer (sometimes called Sapphire) was the fastest computer worm in history, infecting 90 percent of vulnerable hosts within 10 minutes¹³.

Slammer spread through a vulnerability in Microsoft SQL. The worm infected at least 75,000 hosts, perhaps considerably more, and caused network outages and unforeseen consequences such as canceled airline flights, interference with elections, and ATM failures¹³.

Filtering for Slammer as a means of stopping the spread of the worm would have made no difference. For enterprises egress filtering of the MS SQL ports may be a viable option depending on the need of the service. For ISPs to filter MS SQL would be a little harder to do without causing some problems for customers.

Although worms provide a new set of challenges to internet security, there are steps that can be taken to limit the lasting effect of the worms on the internet.

Conclusion

Egress filtering may not be the panacea for all problems facing the internet today, however it can be proven to be a very valuable weapon in the fight to make the internet more secure.

It can also be said that it is the responsibility of everyone in the internet community to do their part to share the load of egress filtering. From the smallest ISP to the largest service provider, all can help eliminate some current threats as we know them today by adding a few simple steps to their network security procedures.

© SANS Institute 2005, Author retains full rights.

Reference:

1. Encarta Dictionary. <http://encarta.msn.com/>
2. Denial of Service Attack on the Rice. <http://www.safesite.com/WhitePapers/DoSanDDoS.asp>
3. Wikipedia – The Free Encyclopedia. http://en.wikipedia.org/wiki/Denial-of-service_attack/
4. Noureldien, Noureldien A., Osman, Izzeldin. 'A method for Defeating DoS/DDoS TCP SYN Flooding Attack The SYNDEF'. <http://www.first.org/events/progconf/2002/d4-05-noureldien-paper.pdf>
5. Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. 'RFC 1597 – Address Allocation for Private Internets'. <http://www.faqs.org/rfcs/rfc1597.html>
6. Associated Press. 'Study: Spam costs businesses \$13 billion'. <http://www.cnn.com/2003/TECH/biztech/01/03/spam.costs.ap/>
7. Symantec Corp. 'Spam Statistics'. <http://www.brightmail.com/spamstats.html>
8. The Register. 'The economics of spam'. http://www.theregister.co.uk/2003/11/18/the_economics/of/spam/
9. ZD Net. 'Comcast takes hard line against spam'. <http://zdnet.com.com/2100-1104-5230615.html>
10. Symantec Corp. http://securityresponse.symantec.com/avcenter/expanded_threats/
11. Symantec Corp., 'W32.Sobig.F@mm'. <http://securityresponse.symantec.com/avcenter/venc/data/w32.sobig.f@mm.html>
12. [WWW.JLSNET.CO.UK](http://www.jlsnet.co.uk). 'Cisco Configs – Block Code Red Attacks'. http://www.jlsnet.co.uk/index.php?tab=2&page=cc_codered
13. Moore, David, Paxson, Vern, Savage, Stefan, Shannon, Colleen, Staniford, Sturt, Weaver, Nicholas. 'Inside the Slammer Worm'. <http://www.computer.org/security/v1n4/j4wea.html>

© SANS Institute 2005, Author retains full rights.