

SIMULATING STRATEGIC FIREWALL PLACEMENT

James D. Box
boxjame@auburn.edu

Adam Hathcock
hathcah@auburn.edu
Department of Computer Science & Software Engineering
Auburn University
Auburn, AL 36849

Alan Hunt
huntala@auburn.edu

Maj. J. L. Humphries, Ph.D.
Jeffrey.Humphries@USAFA.af.mil
Department of Computer Science
U.S. Air Force Academy
Colorado Springs, CO 80840

J.A. Hamilton, Jr., Ph.D.
hamilton@eng.auburn.edu
Computer Science & Software Engineering
Auburn University
Auburn, AL 36849

Keywords: strategic firewall placement, network security, distributed denial of service

ABSTRACT

A distributed denial of service attacks is one of the most difficult security threats to defend against. Typically, a network administrator cannot stop a denial of service attack without contacting the Internet service provider. If an overwhelming amount of traffic is coming into the ISP and being dropped, then only the ISP can effectively filter out the attack traffic. [Chatam 2004] proposes a novel way to give more control over traffic filtering. In order to give the network more control of filtering, the firewall can be connected directly to the ISP via an Ethernet connection. The ISP can then send all traffic straight to the firewall to filter based on the network administrator's rules.

This paper describes a simulation of the proposed DDoS solution. We provide an analysis of the simulation results. We also propose some future research in this area.

INTRODUCTION

Over the past few years, organizations of all types have made doing business over the Internet as commonplace and indispensable as the telephone. Banks, academic institutions, and small businesses have become dependent on the Internet for even the most fundamental of daily functions. Therefore, the cost of a disruption in service and the subsequent recovery can be truly enormous. Malicious attempts by hackers to target and interrupt service are routine and are often attempted using a process called distributed denial of service.

Denial of service attacks present an extremely difficult research problem. Today's Internet technology provides

almost no traceable accountability for the actions of most Internet users. Any skilled hacker can gain control of a large number of computer systems, use them for malicious purposes, and still ensure that it will be virtually impossible for his identity to be discovered. Stopping the information used in a denial of service attack from being sent to a targeted system would not be feasible since it would require changing software in literally millions of independent computers. Therefore, current research focuses on modifying the technology that protects systems that could possibly be targeted. Instead of trying to control information sent to a system, most organizations are trying to control what kinds of information they choose to receive.

The strategic firewall placement method proposes a variation on the use of the Internet firewalls that most organizations have to protect themselves from unwanted Internet traffic. The firewall is usually placed on the communication line just before the point where it reaches the company's network. These firewalls prevent most unwanted information from actually reaching the company's network resources, but denial of service attacks can easily overcome this by taking advantage of the fact that the communication lines between a company and its Internet service provider has relatively low bandwidth. Most companies use Internet service providers that have a huge amount of bandwidth, but the service provider can only reserve a small amount of that bandwidth for each company it services because it must transmit information over such long distances. When a company's communication line is full, the router also serves to store information meant for that company until the line is able to process it. With this setup, an attacker can still overload the communication line between the service provider and the company, back up information in the router, and deny the company access to the Internet before his unwanted information reaches the company's firewall where it is filtered out.

This work was partially funded by the Academy Center for Information Security, at the United States Air Force Academy under research contract # F05611-03-D-0003.

The strategic firewall placement method calls for the company to place their firewall physically close to the service provider's router instead of placing it on the side of the line with the company's computer system. Since the distance between the router and firewall is so short, the service provider can easily produce a very high-speed connection line between these two devices. Then, the router can send much larger quantities of information directly to the firewall, where unwanted information can be immediately discarded instead of being allowed to back up in the router. Of course this assumes adequate processor speed/memory space on the devices. This way, the attacker's information is disposed of before it reaches the much longer, slower part of the communication line.

1. TRADITIONAL FIREWALL TOPOLOGY

A firewall is a perimeter defense component of a typical secure network topology. They are, in essence, a blockade between an internal network, that is assumed to be secure and trusted, and an external network that is not trusted. The most basic types of firewalls are packet filters. This type of firewall operates at the network layer to provide a routing device that includes access control for system addresses and communication sessions. Packet filter firewalls maintain speed and flexibility due to the fact that they are not usually involved in layers above the network layer (layer 3 in the OSI model). More complex firewalls such as application-proxy gateways deal with high-level access at the application layer.

Firewall technology provides a number of benefits to network security including scalability, portability, low cost, and the flexibility needed to focus on individual user needs [Smith et al 2003]. While traditional firewalls prove to be a useful addition to the security of networks, their ability to protect against certain types of threats is virtually non-existent.

2. DISTRIBUTED DENIAL OF SERVICE ATTACKS

Denial of service attacks consume the bandwidth available to the target network. This keeps legitimate traffic from getting through, effectively shutting down the network. A distributed denial service (DDoS) attack involves using multiple machines to blast traffic at the target. An attacker usually takes over other machines remotely and uses their resources to attack the target. The compromised machines are referred to as "zombies." Many tools such as Back Orifice and Sub7 are used to exploit other machines. This software is freely available on the Internet for anyone to download.

A firewall cannot defend against a DDoS if traffic is built up at the ISP and dropped. In this case, the victim cannot control the filtering of incoming traffic. For such an attack, the company must contact the ISP in order to filter out the attack traffic. This is not a very efficient approach to countering DDoS attacks. Since the company cannot control the bottleneck link, it cannot defend against a DDoS attack.

3. STRATEGIC FIREWALL PLACEMENT

In the strategic firewall placement method, the company's firewall, which would normally be located on the premises of the company, is moved to a location physically close to the Internet Service Provider's router. The firewall remains under the control of the company. Then the firewall is connected by an Ethernet cable to the ISP's router using an Ethernet link that is available on many routers. The connection between the ISP and the company will typically have much lower bandwidth due to the much greater distance. However, when the firewall is located on the ISP's premises, it is quite feasible to connect the firewall to the router via Ethernet.

The purpose of relocating the firewall is to allow the company to filter packets before they reach the bottleneck link. In a denial-of-service attack against a company with a typical firewall setup, the attack packets sent by an attacker have to go through the bottleneck link before the firewall can even attempt to filter them. The bottleneck line quickly becomes overwhelmed which will exhaust the buffer at the ISP which will begin dropping all incoming traffic. However, with strategic firewall placement, the filtering burden is shifted from the ISP to the company firewall connected via Ethernet. An attacker can still run a denial of service attack, but he would need a huge amount of bandwidth to be able to clog the Ethernet line. Therefore, this setup makes running an attack much harder.

Since company employees administer the firewall themselves, they are able to quickly update the firewall's security policies when necessary to thwart attacks. In a typical setup, if the company finds that an attack is coming from a specific IP address, they would have to call the ISP and ask them to block that IP to stop the attack. With the new setup, the company can simply change its own firewall settings and block the attack much faster. Of course, it is more common for DOS attacks to come from a variety of IP addresses. In this case, all traffic can be default denied except for mission-critical traffic from known addresses. Although this scheme is vulnerable to spoofing, this would allow for continuity of operations under many scenarios.




Figure 1. Simulation Topology

While moving the firewall is helpful, the company will also have to update its firewall configuration in a couple of different ways to have a secure defense against denial of service attacks. First, the firewall must keep a policy of denying packets by default. The firewall must keep track of requests from within the company to connect to the Internet. When a company employee wants to connect to a certain site, the company server will send out a packet requesting synchronization with that web site. When the site receives this packet, called a TCP SYN packet, the site responds with an acknowledgement packet called a SYN/ACK packet, which allows the connection to be completed. The firewall would only forward these acknowledgement packets if it finds that someone inside the company has specifically requested that connection.

This helps prevent a specific form of denial of service attacks called reflected denial of service. In a reflected denial of service attack, an attacker spoofs the victim's IP address and sends out requests for connections to many different sites. Each of these sites thinks that the victim company has requested the connection, so the sites then flood the company's connection with acknowledgement packets. However, if the firewall settings are changed in the suggested way to only allow SYN/ACK packets that are known to be legitimate, it would be much harder to flood the company's communication lines using this method.

The second change is how the firewall handles inbound connection requests. When the company receives a TCP SYN packet requesting a connection to their server, the firewall sends an acknowledgement packet back as if it had come from the server. This SYN/ACK packet is sent to the IP that the TCP SYN packet appears to have originated from, and the firewall will not allow a connection through unless it receives a response from the IP address that sent

the connection request. This alleviates unwanted traffic on the bottleneck link. Unless the attacker actually had control of the IP address that he was spoofing, he would be unable to respond to the firewall's SYN/ACK packet, and he would not be able to gain access to the T1 line. As a result, attackers would not be able to spoof multiple IP addresses at once and spam the company's communication lines with useless packets from all of them.

Managing these properties of the firewall would require a large amount of computational power, and the firewall by itself may not be able to handle the entire job. However, the functionality of the firewall could be distributed across multiple machines. The computers would do the filtering and send legitimate packets along to the firewall, and the firewall would be set to only allow packets through that came from those specific computers' MAC addresses. This setup would alleviate any problems with the firewall not being able to keep up with the traffic sent to it.

4. SIMULATION

We simulated the strategic firewall placement strategy with NS-2. Our is a simple model to evaluate the effects of firewall placement. Figure 1 shows the topology of the model. The bottom node on the far left represents legitimate traffic attempting to communicate with the target node. The top node on the far left represents an attacker sending bogus traffic in order to shut down the site by choking the bottleneck. Then we have the have the firewall, router, and the target site. The node representing the legitimate traffic is connected to the router with a 1.5 Mbps link. The links between the legitimate source and the router and the firewall and the router are variable. The link between the firewall and the target is 1.5 Mbps.




Figure 2 Network Animation of the Simulator

The attack node blasts UDP traffic to the target node. We use UDP traffic instead of TCP traffic to avoid flow control by the sender. We collected throughput on the choke point between the router and the firewall while varying the link capacities of the choke point and the data rate of the attack traffic.

Figure 2 is a NAM (Network Animator) screenshot of the simulation. The blue shapes are the offending packets and the red shapes are the legitimate traffic. The packet stack rising above the link between nodes 2 and 3 is the queue on that link. The packets then fall from the top of the queue in the animation to show that they are dropped.

		Attack Traffic			
		100 Mbps	50 Mbps	10 Mbps	1.5 Mbps
Bottleneck Link	100 Mbps	1.24 Mbps	1.24 Mbps	1.24 Mbps	1.24 Mbps
	50 Mbps	1.24 Mbps	1.24 Mbps	1.24 Mbps	1.24 Mbps
	10 Mbps	816 bps	32 Kbps	57 Kbps	1.23 Mbps
	1.5 Mbps	0 bps	0 bps	816 bps	6.5 Kbps

Figure 3 Simulation results with various link speeds.

Figure 3 shows the results of the simulations with various link data rates. The top row of the table shows the data rate that the attacker transmits to the target. The far left column is the link capacity of the bottleneck. There is no change in throughput when we reduce the bottleneck from 100 Mbps to 50 Mbps. However, when the bandwidth is reduced to 10 Mbps, we see a dramatic drop in throughput from 1.24 Mbps to 816 bps when the attack traffic is at 100 Mbps. Even when the attack traffic is reduced to 50 or even 10 Mbps, the throughput still decreases considerably as the bottleneck bandwidth decreases.

This shows how critical the link between the ISP router and the company firewall is. As the capacity of the bottleneck link decreases, the more traffic is dropped at the router. So in the case of low bandwidth between the ISP router and the firewall, only the ISP has control of filtering packets. By increasing the bandwidth and shifting

responsibility to the firewall, the company gains control of the bottleneck link.

5. CONCLUSION

The goal of this paper was to simulate this particular technique for mitigating distributed denial of service attacks in order to provide additional confirmation of its effectiveness. The results of our simulation provide further support for Chatam’s conclusion about strategic firewall placement. For the particular type of DDoS in which flooding is used, a firewall can only be of use if it acts as a gateway to the bottleneck link.

6. FUTURE WORK

Looking at this as a point solution for a large network will be the subject of future research. This kind of

relocation of firewalls (and possibly using distributed firewalls) in conjunction with the Secure Overlay Service architecture could provide a very robust defense against distributed denial of service attacks.

REFERENCES

Chatam, J. W. *Using Strategic Firewall Placement to Mitigate the Effects of Distributed Denial of Service Attacks*, Masters Thesis. Auburn University, 2003.

S. Gibson, "Distributed Reflection Denial of Service. Description and analysis of a potent, increasingly prevalent, and worrisome Internet attack," February 22, 2002, <http://grc.com/dos/drDOS.htm>

Smith, R.; Chen, Y; and Bhattacharya, S., "Cascade of Distributed and Cooperating Firewalls in a Secure Data Network," *IEEE Transactions on Knowledge and Data Engineering*, IEEE Educational Activities Department, vol 40, no 5, (September): pp 1307 – 1315, 2003.