



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: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Egress filtering – keeping the Internet safe from your systems

by Heather L. Flanagan
GSEC Practical Assignment Version 1.2c

Executive Summary

Egress filtering is the simple process of filtering of outbound traffic from your network. Why is this critical? Either through malicious intent or simple misconfiguration of a network, sites can flood the Internet with bogus packets. The massive DDOS attacks from February, 2000 were a perfect example of what happens when sites are hacked solely for their network connection to send bogus packets to other servers on the Internet. Common practice has sites filtering incoming traffic through their routers and firewalls, but often companies miss the simple procedure of egress filtering, leaving them open to being the source of Distributed Denial of Service (DDOS) attacks. DDOS attacks will not be controlled until more sites configure their networks for egress filtering.

This paper briefly discuss the benefits of egress filtering, gives examples for what common DDOS tools it can block, and directs the reader to sites with specific details on how to implement this filtering at your site.

Egress and Ingress filtering

The purpose behind egress filtering is to prevent any packets with invalid or incorrect addresses from leaving your site. These packets may be originating from a misconfigured router in your network or, more dangerously, from a compromised system hosting one of the many DDOS tools available.

Egress filtering generally occurs at the edge of a network, at the firewalls and border routers. At no time should your network send out any packets with addresses not legally assigned to you – to do so means either your firewall may be misconfigured to show the world your internal address space, or worse, that you are the home of one or more DDOS attack agents. There should be very little effect or loss of functionality to your network when implementing egress filtering – all legitimate traffic requires is your legal addresses, so blocking anything else will only break things that should not be sent in the first place! If your site has already been compromised, you may see your own connection to the Internet degrade as your firewalls and routers struggle to stop the traffic. Given the possibility of being brought into a lawsuit if your site is involved in a DDOS attack against another, finding and stopping bogus traffic from leaving your network at the cost of performance until you can clean up the DDOS agents is a small price to pay.

Similar to egress filtering is ingress filtering, described in detail in [RFC 2267](#). Ingress filtering is the filtering of “any IP packets with untrusted source addresses before they have a chance to enter and effect your system or network.” (See “5.2.3 Ingress and Egress Filtering”, Denial of Service Tools

Administration) This is best done at the ISP level where they can more cleanly handle the packets coming through their many networks. Unfortunately, for some of the larger ISP's like AT&T and sprintlink.net, they connect such a large number of networks that filtering for legal addresses is extremely difficult.

Ingress filtering has its limitations – for large ISP's, other companies with different addresses may be using their backbone. To prevent those addresses from going through the network would be its own form of denial of service attack. Keeping track of the many legitimate addresses that can go through a large ISP is next to impossible – it is better to have security as close to the source as possible, encouraging each site to perform their own egress filtering.

How egress filtering actually works

As discussed in SANS "Egress Filtering" document, the procedure behind this concept is quite simple. You want to make sure that your border routers and firewalls do not allow directed broadcast packets to be forwarded by default, and that only addresses assigned to you are allowed through. Directed broadcast packets are the result of ICMP packets being sent to a network's broadcast address. If there are 50 hosts on that network, and all respond to the broadcast, then suddenly there are 50 sets of replies being sent out in response to one set of ICMP packets. All common routers and firewalls are capable of this kind of filtering, and the "Egress Filtering" document describes the most popular routers and firewalls:

<http://www.incidents.org/protect/egress.php>

If you are not sure you have set up your egress filtering properly and would like another way to check, Mitre has a tool available called "Egressor" that checks your router configurations. See

<http://www.mitre.org/research/cyber/docs/TOOL.html>

In order to avoid the possible hit to your own network if you institute egress filtering, you should, if possible, make sure you are clean of any DDOS tools first. There are several free and commercial tools that will let you check your hosts for the most popular DDOS tools:

<http://www.nipic.gov/warnings/advisories/2000/00-055.htm> - According to the National Infrastructure Protection Center, the "find_ddos" tool will search for the following DDOS programs on your system: mstream, tfn2k client, tfn2k daemon, trinoo daemon, trinoo master, tfn daemon, tfn client, stacheldraht master, stacheldraht client, stacheldraht daemon and trn-rush client.

<http://theorygroup.com/Software/RID/> - "RID", another flexible DDOS searching tool, offered by the Theorygroup, is a C program that can be configured to search for many DDOS tools, including any remote software that elicits a

predefined response to a given set of packets such as Trinoo, TFN, and StachelDraht clients.

<http://www.iss.net/news/denialfaq.php#4.2> - RealSecure from ISS – detects floods on network & scans for presence of DOS tools

Details on DDOS attacks and prevention

There are several DDOS attack programs out there, and more evolving every day. Unfortunately, there are more DDOS attacks than there are good ways to deal with them! Some of the most popular DDOS programs are Trinoo, Tribal Flood Network (TFN) and TFN2K, and StachelDraht. Others that have evolved from that are Mstream, Shaft, Trinity and many others. The first well-known DDOS attack occurred in August 1999, featuring the use of Trinoo against the Internet Relay Chat (IRC) server at the University of Minnesota. This attack lasted for two days. DOS and DDOS attacks are quite common in the IRC world. If a user wants to take control of an IRC channel, they just have to work with a group of their friends to block the channel operator from the channel, and take over.

As mentioned earlier, in February 2000, DDOS attacks expanded from the IRC world into the prime Internet spotlight. Sites like CNN, Yahoo!, Amazon and others were made unavailable for two to four hours – for e-commerce sites, that means thousands if not millions of advertising and product dollars. These attacks used Trinoo and TFN. There was an immediate media outcry calling for more security on the systems, and complaining about having been robbed when they did not lock their doors. The DDOS tools that were involved in these attacks were just the beginning of a quick evolution of DDOS tools available to the hacker community.

Details for these attacks and more are available from several web sites, including:

<http://www.technotronic.com/>

Techtronic provides not only the tools to discover many of the DDOS programs on your system, it also provides a copy of the DDOS programs themselves so administrators can take them apart and examine them from the inside out.

http://www.sans.org/infosecFAQ/threats/understanding_ddos.htm

A research paper written by DeokJo Jeon, this goes over more information on the nature of DDOS attacks, and gives the locations of several more tools to combat them.

<http://www.w3.org/Security/Faq/wwwsf9.html>

There is a wealth of information in this document, which features specific questions and answers about all the common DDOS tools. This should be

reviewed by anyone involved in researching DDOS tools and ways to protect against them.

While varying in some important details, most of these programs can be stopped at the source using egress filtering. They all use IP spoofing, where a packet is artificially created to have an invalid (at least for that site) IP address. Spoofed packets are next to impossible to trace. Tracing requires a continuous flow of packets, and the cooperation of possibly several ISP's as they all try to trace the packets back one subnet at a time. Most tools only send packets out from their agents for a few minutes at a time.

An example of the use of spoofed packets is the basic DOS tool, Smurf. Most of the above mentioned tools can do a "Smurf" style attack, which is an ICMP attack relying on directed ICMP broadcasts. The originating server sends a packet with a spoofed source address, and the intermediate server – the destination of the ICMP, or ping, packet – tries to respond to the ping. Two systems are hit for the price of one packet – the intermediate server, and the owner of the spoofed IP address. If the attacker used the actual address of the machine originating the packet, s/he would quickly take themselves out of the attack as the originating system halted under the flow of return ICMP packets. More information specifically on Smurf is available at the WWW Security FAQ, question 97, "What is a "smurf attack" and how do I defend against it?"

While many DDOS attacks can use this ICMP directed broadcast, they may also use UDP packets (trino), SYN packets (TFN, TFN2K, StachelDraht) TCP ACK packets (mstream), or all of the above (TFN, TFN2k, StachelDraht, Shaft, Trinity). By preventing illegally and improperly addressed packets from leaving a network, a DDOS attack is halted before it can get started.

One DDOS attack tool can work its way around egress filtering – TFN2k. As described by Rick Farrow in his article, "Network Defense; DDOS is neither dead nor forgotten":

"TFN2k handlers can send a command to test source address spoofing by sending a test packet, with spoofed source address, to the handler. If the test packet isn't received, the agent is told to spoof only the lower eight bits of the network address, something which will not be blocked by egress filtering."

This does not mean egress filtering should be given up on! Egress filtering still narrows down the possibility of locating the source of the packets, and if the TFN2k agent is not configured properly, egress filtering will still catch the outgoing flood of packets.

Other tools to help control DDOS

Egress filtering is one giant step the Internet community can take to start controlling DDOS attacks. Other tools are available that can also help. Cisco offers a feature on their routers called "Committed Access Rate", or CAR, that limits the amount of bandwidth used by any particular packet type. This can

help deal with packet floods coming or going. Using the various scanning tools to check to see if your network is clean is another important step. And of course, keeping your system up to date with your vendors security patches can fix holes before they are exploited.

The use of true proxy servers is another way to limit DDOS attacks. One of the basic functions of a proxy server is to encapsulate all packets with the proxy server's external interface address. This does not help with any systems outside the proxy's control, such as external web or ftp servers. Those systems will still need egress filtering. Even if packets are created internally with spoofed IP addresses, the proxy server will change the address to the one legally assigned to that system.

For a more detailed discussion on securing your systems against DDOS attacks, there are several extremely useful web sites available:

http://www.sans.org/ddos_roadmap.htm

The SANS web site is a wealth of information on system and network security. This particular document was brought together by the Partnership for Infrastructure security, and details both the problems that bring about DDOS attacks and their solutions.

<http://www.sans.org/dosstep/index.htm>

Another SANS web site, which specifically discusses egress filtering and preventing broadcast amplification. A short document that gets straight to the point.

<http://www.cisco.com/warp/public/707/newsflash.html>

A very useful paper brought to us by Cisco, it is somewhat Cisco specific, but still provides several useful points, including logging information for law enforcement and what the characteristics of the common DDOS tools are from a network perspective.

Conclusion

Egress filtering can be thought of as the Internet's form of a Lojack™ car alarm – Lojack™ doesn't stop theft, but car's equipped with it can be tracked more easily by police, making it not nearly as worth stealing! Similarly, egress filtering, while it can be worked around by the attacker, makes the attack that much more difficult, and easy to trace back to the hosting systems. If enough networks are configured with egress filtering, we could have a bit more breathing room to figure out the next defense needed for the next evolution of attack.

References:

1. Ferguson, P. "RFC 2267 - Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing." January, 1998. URL: <http://www.landfield.com/rfcs/rfc2267.html> (27 April 2001)
2. Carter, Jeff. "Egress Filtering" Version 0.2. 29 February 2000. URL: <http://www.incidents.org/protect/egress.php> (27 April 2001)
3. "5.2.3 Ingress and Egress Filtering" Denial of Service Tools Administration. 5 March 2001. URL: http://www.tru64unix.compaq.com/faqs/publications/iass/OSIS_53/admin/DNSTLSXX.HTM (27 April 2001)
4. Livingston, Brian. "We can prevent those distributed denial of service attacks with 'egress filtering.'" 1 March 2000. URL: <http://www.cnn.com/2000/TECH/computing/03/01/prevent.ddos.idg/> (27 April 2001)
5. "Egressor: A tool for checking router configuration." 24 October 2000. URL: <http://www.mitre.org/research/cyber/docs/TOOL.html> (27 April 2001)
6. Farrow, Rick. "Network Defense; DDOS is neither dead nor forgotten." 7 January 2001. URL: <http://www.spirit.com/Network/net1200.txt> (27 April 2001)
7. Shankland, Stephen. "Expert Warns of Powerful New Hacker Tool." 1 May 2000. URL: http://www.infowar.com/hacker/00/hack_050100c_j.shtml (27 April 2001)
8. "Distributed Denial of Service; A white paper prepared by WatchGuard Technologies, Inc." February, 2000. URL: <http://www.watchguard.com/docs/ddos.pdf> (27 April 2001)
9. "Help Defeat Denial of Service Attacks: Step-by-Step" Revision 1.41. 23 March 2000. URL: <http://www.sans.org/dosstep/index.htm> (27 April 2001)
10. Jeon, DeokJo. "Understanding DDOS Attack, Tools, and Free anti-tools with Recommendation." 7 April 2001. URL: http://www.sans.org/infosecFAQ/threats/understanding_ddos.htm (27 April 2001)
11. Stein, Lincoln. "The World Wide Web Security FAQ." 24 March 2000. URL: <http://www.w3.org/Security/Faq/wwwsf9.html> (28 April 2001)
12. Lemos, Robert. "A year later, DDOS attacks still a major Web threat." 7 February 2001. URL: http://news.cnet.com/news/0-1003-201-4735597-0.html?tag=tp_pr (27 April 2001)
13. Radcliff, Deborah. "DDOS Attacks' Ultimate Lesson: Secure That Infrastructure." 21 March 2000. URL: http://www.info-sec.com/denial/00/denial_032100a_j.shtml (30 April 2001)
14. "Strategies to Protect Against Distributed Denial of Service (DDOS) Attacks." 17 February 2000. URL: <http://www.cisco.com/warp/public/707/newsflash.html> (30 April 2001)