



SANS Institute

Information Security Reading Room

Egress Filtering FAQ

Chris Brenton

Copyright SANS Institute 2021. Author Retains Full Rights.

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

Egress Filtering FAQ

(Heavily borrowed from the SANS 502 Perimeter Security track)

Chris Brenton

Last revision: 2.0 April 19, 2006

Please direct all questions or suggestions to [cbrenton at chrisbrenton.org](mailto:cbrenton@chrisbrenton.org)

Preface

This paper will teach you some of the basics of egress filtering and firewall rules. If you are interested in learning more about this subject, we recommend taking the [SANS SEC502 Firewalls, Perimeter Protection and VPNs course](#), available both online and via live classroom training.

This FAQ covers the benefits of performing egress filtering on the end points of your perimeter. Egress filtering is not only beneficial to your own network security, but to the rest of the Internet as well. Deployment is usually relatively simple, provided you understand what traffic should be permitted to leave your network.

What is egress filtering?

Egress filtering is the control of traffic leaving your network. Far too often firewall administrators create a policy rule that essentially says “let my internal network transmit any and all traffic patterns out to the Internet”. Egress filtering limits this traffic flow to a reduced subset.

Why should I perform egress filtering?

Egress filtering prevents you from sending unwanted traffic out to the Internet. This could include leaking out private address space or stopping compromised systems attempting to communicate with remote hosts. Egress filtering can also help prevent information leaks due to misconfiguration, as well as some network mapping attempts. Finally, egress filtering can prevent internal systems from performing outbound IP spoofing attacks.

My hosts go through a NAT device. Should I perform egress filtering?

Absolutely! Network address translation devices can sometimes leak out the private address space located behind them. This can occur when the device experiences high utilization or can be due to an attack. If your private IP addresses leak, an attacker may be able to use that information to enumerate the layout of your internal network. Egress filtering can ensure that only packets stamped with your legal IP address range are permitted out to the Internet.

Which TCP/UDP ports should I consider blocking?

There are quite a few communication patterns that the average network would never want to let out to the Internet. Here are some examples:

MS RPC (TCP&UDP 135), NetBIOS/IP (TCP&UDP 137–139), SMB/IP (TCP/445)

When communicating with remote hosts, Windows systems love to fall back on sending queries via their default protocols. This can not only leak out information, but can easily be mistaken for malicious behavior by the target system. It is best to ensure these protocols remain within your

network.

Trivial File Transfer Protocol – TFTP (UDP/69)

When an attacker exploits a system, the first thing he does is go looking for some way to move his toolkit onto the system. TFTP is the tool of choice since it permits the attacker to transfer the file without any interactive prompting. Not only should you block outbound access to TFTP, but you should also alert on this traffic pattern since it is usually an indication that an internal system has already been compromised. As a bonus feature, blocking TFTP will prevent the transfer of the toolkit, thus making system recovery that much easier.

Syslog (UDP/514)

Syslog is used to transfer log information to a centralized server. Needless to say log files can contain critical information regarding our environment. Given the importance of this data, an egress filter insures that a mis-configured system never accidentally sends log entries out to the Internet.

Simple Network Management Protocol – SNMP (UDP 161–162)

SNMP is another protocol that can reveal critical information regarding your infrastructure. Again, best practice is to ensure that it never leaks out past the perimeter.

SMTP from all IP's but our mail server (TCP/25)

Many systems are compromised for the sole purpose of being turned into SPAM relays. Attackers make money by taking control of thousands of systems across the Internet and using them to transmit unsolicited e-mail. Having this e-mail originate from your network is a great way to end up on one or more black lists. By blocking outbound SMTP from all systems but your legitimate mail servers, you can help prevent this from occurring.

Internet Relay Chat – IRC (TCP 6660–6669)

IRC is a network of meeting areas where folks can communicate via text-based messaging. Unfortunately, it is not uncommon for compromised system to “call home” by reporting in to a specific IRC chat channel. This allows the attacker to keep track of the compromised systems as well as to send the bot commands without requiring a direct connection to the system. While IRC can run on any port, the most commonly used range is TCP/6660 – TCP/6669. This is another set of ports that you not only want to block, but you want to trigger an alert if it is detected since it could be an indication of a compromised system.

Should I be filtering outbound ICMP as well?

Along with certain TCP and UDP ports, there are specific ICMP type/codes you may wish to consider blocking as well:

ICMP Echo-Replies (type 0 code 0)

Echo-reply packets are returned by a system in response to receiving Echo-Request packets. This is usually an indication of someone running the Ping utility. Echo-Reply packets can also be used as a covert communication channel. For example, Loki uses Echo-Request/Echo-Reply packets in order to create a covert communication channel for Telnet like communications. Echo-Replies have also been

used to covertly communicate with compromised systems. If you have already implemented rules preventing inbound Echo-Requests, it's a good idea to block outbound Echo-Replies as well.

ICMP Host Unreachables (type 3 code 1)

Sometimes attackers find very creative ways to sneak in probe packets. For example if you have implemented a firewall that does not support full stateful inspection for ICMP Echo-Request/Echo-Reply packets, a savvy attacker can generate Echo-Reply packets in an attempt to map your network. While a host receiving an unsolicited Echo-Reply will not respond, if the host is off-line, the upstream routing device will return an ICMP host unreachable. Thus an attacker can ID which systems are on-line simply by checking which IP addresses do not cause a host unreachable to be generated. While we could try to block the inbound Echo-Replies, this could interrupt our legitimate use of the Ping utility. With this in mind, the easiest way to solve this problem is to simply block the host unreachables attempting to head out towards the Internet.

ICMP Time Exceeded in Transit (type 11 code 0)

Network mapping tools such as traceroute, tracert, Firewalk and tcptraceroute map all of the routers between a source and a target host by generating packets with an artificially low Time To Live (TTL) value within the IP header. This causes the routing devices along the path to return ICMP time exceeded in transit error messages. The more advanced tools like Firewalk and tcptraceroute target open ports (such as UDP/53 on a DNS server or TCP/80 on a Web server). This makes it impossible to block the inbound stimulus unless your firewall supports filtering packets based on the TTL value. Since most firewalls do not support this feature, we usually can't go after the stimulus. Our next best option is to filter out the returning reply packets. Filtering outbound time exceeded in transit errors will do just that.

Instead of blocking all these ports, why don't I just let out what I need and block everything else?

Certainly a restrictive policy is going to provide the highest level of security. The fewer ports you have open, the better. This policy, however, is administratively prohibitive in many environments. If you can go with a more restrictive policy, that should be your first choice. If not, this document identifies the minimum level of filtering you should consider performing.

Where should I implement egress filtering?

IP address filters should be implemented as close to the end of your perimeter as possible. Typically this will be the border router leading up to your ISP. The same is true for ICMP based filtering. The TCP and UDP ports can be filtered at the firewall or the router, although the router is preferred as this will provide a layer of protection for the router and firewall as well.

Can you provide some example rules?

Here is an example of the above filter recommendations applied to a Cisco router running IOS 11.3 or higher. It is assumed that the legal IP address space allocated to the network is 1.2.3.0/24 and the legitimate SMTP server is located at 1.2.3.4. The filter is applied outbound on the external interface of the router so that it not only controls the traffic leaving the network, but controls the traffic leaving the router as well. You could optionally choose to add the "log" parameter to any/all of these filter rules in order to record these events when they occur.

```
interface Serial 0
    ip access-group filter_outbound out

ip access-list extended filter_outbound
deny icmp any any time-exceeded unreachable echo-reply
deny tcp any any range 135 139
deny udp any any range 135 139
deny tcp any any 445
deny udp any any 69
deny udp any any 514
deny udp any any range 161 162
deny tcp any any range 6660 6669
permit tcp 1.2.3.4 any 25
deny tcp any any 25
permit ip 1.2.3.0 0.0.0.255 any
deny ip any any
```

Does egress filtering break any protocols?

The only applications that will break with egress filtering are the ones using the service ports specified above.

What about VPN connections?

A VPN tunnels traffic through some form of secure communication protocol such as IPSec, SSL, SSH or PPTP. If your VPN tunnel terminates inside of your egress filtering point, your egress filters will have no effect on traffic passing through the VPN tunnel. This means that you can continue to use RPC, SMB/IP, etc. through your VPN and your egress filters will have no effect on this traffic. Of course this also means your egress filters will provide no protection if you inadvertently leak information through your VPN.

Should every network perform egress filtering?

While egress filtering is considered an important best practice, clearly every network is somewhat unique to the business needs of the organization. The above filters are typically safe for an overwhelming majority of the networks communicating on the Internet. With this in mind, the implementation of egress filtering should be strongly considered.

This paper covers some of the basics of egress filtering and firewall rules. If you are interested in learning more about this subject, we recommend taking the [SANS SEC502 Firewalls, Perimeter Protection and VPNs course](#), available both online and via live classroom training.