



**NISCC Technical Note 01/2006
Issued 20 April 2006**

Egress and Ingress Filtering

This NISCC Technical Note supplements NISCC Technical Note 10/04, "Understanding Firewalls". Its purpose is to assist those thinking about implementing or changing packet filtering devices at their network boundaries. The advice may be applied directly by those organisations managing their own network infrastructures or can be used in discussion with outsourced managed service providers as necessary. The advice is aimed at technical managers and security officers and those persons responsible for managing outsourced security or network management contracts. Architectural decisions aside, the advice is vendor independent.

Reference to any specific commercial product, process or service by trade name, trademark manufacturer, or otherwise, does not constitute or imply its endorsement, recommendation, or favouring by NISCC. The views and opinions of authors expressed within this document shall not be used for advertising or product endorsement purposes.

NISCC shall also accept no responsibility for any errors or omissions contained within this document. In particular, NISCC shall not be liable for any loss or damage whatsoever, arising from the usage of information contained in this document.

**National Infrastructure
Security Co-ordination Centre**
PO Box 832
London, SW1P 1BG

Tel: 0870 487 0748
Fax: 0870 487 0749
Email: enquiries@nisc.gov.uk
Web: www.nisc.gov.uk

Key Points

- This note is an introduction to packet filtering.
- Packet filtering remains a useful component of a *defence in depth* strategy.
- Egress filtering is just as important as ingress filtering.
- Businesses should consider carefully both the type and the destination of traffic they permit outbound.
- Businesses are strongly recommended to consider blocking outbound web connectivity to geographic areas with which they have no legitimate need to communicate.
- Packet filters can be employed to enforce business management decisions.

Introduction

1. This NISCC Technical Note supplements NISCC Technical Note 10/04, "Understanding Firewalls". Its purpose is to assist those thinking about implementing or changing packet filtering devices at their network boundaries. The advice may be applied directly by those organisations managing their own network infrastructures or can be used in discussion with outsourced managed service providers as necessary. The advice is aimed at technical managers and security officers and those persons responsible for managing outsourced security or network management contracts. Architectural decisions aside, the advice is vendor independent.

2. This note is not intended to be exhaustive in its discussion of network filtering products or applications, nor does it give advice on host based security. Instead, it concentrates on the capabilities of network boundary devices to mediate traffic based on the contents of TCP/IP packet headers. Typically such devices will be screening routers at network boundaries rather than application level firewalls (though such firewalls will usually also include packet filtering capability). Packet filtering should be applied as part of a strategy of *defence in depth* in conjunction with appropriate host based security mechanisms as well as application level or stateful filtering firewalls. Boundary packet filtering can relieve inbound application layer firewalls of much unnecessary and unwanted traffic.

3. This note should be read in conjunction with earlier NISCC technical notes and advice and guidance. In particular:

- Understanding Firewalls – NISCC Technical Note – December 2004 (revised February 2005).
- Spam Mitigation Techniques – NISCC Technical Note – February 2004
- Protecting your computer network – Guidance on securing LANs, WANs and Internetworks – NISCC Technical Note – March 2002
- Mitigating the risk of Malicious Software – NISCC guidance note – October 2004

Background and definitions

4. Packet filtering limits the flow of information across a network boundary based on attributes of the network packets. In the case of IPv4 based networks, these attributes are usually confined to source and destination IP address and destination (or more rarely) source port numbers. Packet filters may also inspect the network layer IP headers to determine the transport protocol encapsulated or any IP options set (such as source routing). Packet filters let an administrator

specify rules to limit protocol specific traffic to one network segment or domain. Packet filters do not inspect application level data encapsulated in the transport layer, such inspection is performed by application level firewalls and proxy relays.

5. RFC 3704¹ notes that there are at least five ways one can implement packet filtering as described in RFC 2827² of which ingress filters are but one. RFC3704 describes an ingress filter as one “that checks the source address of every message received on a network interface against a list of acceptable prefixes, dropping any packet that does not match the filter.” Whilst RFC2827 covers only source address spoofing, the principles it extols are more widely applicable. This NISCC note concentrates on the business benefit which may be derived from packet filtering generally. In particular, it recommends consideration be given to filtering whole classes of networks, or indeed netblocks³, where there is no business need for communication.

Direction of traffic flow (ingress v. egress)

6. We consider ingress filtering to apply to all traffic inbound to the organisation from outside the corporate boundaries. Egress filtering applies to all traffic outbound from the organisation. Note that this definition of ingress and egress is the opposite of that used in RFC3013⁴ describing measures to be applied by ISPs, where ingress is taken to mean *from* the edge site (corporate customer) *to* the Internet.

7. Whilst typically the filters discussed here will be applied at the boundary between an organisation and a wider public network (e.g. the Internet), they may also be applied at intra-organisational boundaries where little or no trust can be assumed. Filtering can also apply at greater levels of granularity within an organisation depending on the policy stance taken. For example, a research department might be treated by the wider organisation as a potentially hostile environment because it uses an unmediated connection to the Internet.

Defence in depth

8. Packet filtering has performance advantages over application level inspection mechanisms but it is necessarily limited in the intelligence or granularity of decision making it can apply. More granular decisions may be applied by boundary firewalls. Packet filtering should be seen as part of a cohesive and comprehensive security policy.

¹ RFC 3704 - Ingress Filtering for Multihomed Networks – March 2004

² RFC 2827 - Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing – May 2000

³ A range of consecutive IP addresses, usually belonging to the same organisation.

⁴ RFC3013 - Recommended Internet Service Provider Security Services and Procedures – November 2000

9. Whilst RFC3013 is aimed at ISPs, it recognises that filtering at the interface between the customer and the Internet may not always be possible and recommends that the customer should be encouraged to implement his or her own filtering within their networks. The RFC recommends that, in general, filtering should be done as close to the actual hosts as possible. This reinforces the implied need for defence in depth. Ideally all Internet facing hosts within an organisation will be located on a separate network usually known as the *Demilitarized Zone* (DMZ). Access Control Lists (ACLs) applying to that network can then reinforce the security policy at the closest point possible.

10. Defence in depth does not mean simply applying the same kind of filtering at different points in the network. No organisation should rely for its network protection on packet filtering alone. Any filters applied at boundary routers will inevitably need to be reinforced by policy applied at other network choke points. Typically such choke points will be firewalls or other relays or proxies which apply policies based on checks of content in layers of the protocol hierarchy above the network or transport layer. For further discussion of such choke points, see NISCC Technical Note 10/04 – Understanding Firewalls.

Applying packet filters – what to filter and where

11. The organisational boundary router is best placed to enforce the widest possible applicable filters – i.e. which **networks** to allow traffic to or from. RFC 1918⁵ reserves blocks of IP addresses for private networks. None of these addresses should be routed across the Internet. Most organisations now allocate internally IP addresses drawn from RFC1918 when building their networks. Network address translation (NAT) is then applied at the boundary in order to map the internal unroutable network addresses to those which are publicly addressable. But RFC1918 addresses are not the only ones which are reserved. RFC3330⁶ widens the list to include any special use IPv4 address.

Ingress filters

12. Because RFC3330 network addresses cannot legitimately be routed, they are commonly known as bogons (from “bogus networks”). ACLs applied at the organisational boundary should deny inbound any traffic which purports to come from a bogon address. Such addresses are often used in denial of service attacks.

⁵ RFC 1918 - Address Allocation for Private Internets – February 1996

⁶ RFC3330 – Special-Use IPv4 Addresses – September 2002

13. Bogon filters are listed by a number of organisations and are published on the web⁷. The network security researchers known as Team Cymru⁸ have published the bogon lists in a variety of formats which may be used to build ACLs.

14. Where an organisation does not use RFC1918 addresses internally, or uses some publicly routable addresses for certain networks (as is typically the case for networks hosting DNS, mail, web or other public servers on a DMZ) then the ingress filter should also deny all traffic which purports to have such an internal source address. Such traffic cannot be legitimate if it appears inbound at the external router interface.

15. Ingress filters can also be used to reinforce *default deny* firewall policies by dropping all inbound traffic except that specifically required by the business. Thus for example, if the business offers public web services on a service network (or DMZ), then traffic destined for port 80 on the web server may be allowed, but all other traffic with a destination port 80 should be dropped as potentially hostile. Similarly, if an organisation hosts its own SMTP server, then traffic to port 25 on that server, and only that server, may be allowed. Ingress filters should be carefully designed to support the business security policy. The intention must be to limit inbound traffic to the minimum necessary for proper business functionality.

Egress filters

16. Whilst almost all Internet connected organisations now apply fairly comprehensive ingress filters, many fail to apply egress filters at all, or apply only limited filtering. This is unfortunate because filtering outbound traffic can benefit both the organisation itself and the wider network community. For example, the January 2003 Slammer worm attack would have been much less damaging if appropriate egress filters had been in place on all networks hosting MS SQL servers.

17. Egress filtering is aimed at blocking a local host's outbound connectivity, or limiting that connectivity to activity which is essential to the business. Hosts or clients on the organisational network which do not need access to the Internet should be denied all connectivity. Egress filtering is designed to reinforce the security policy by assuming the worst – i.e. that some part of the network has been compromised by a failure elsewhere and the damage caused must be limited. For example, there is no need for any desktop client machine to make outbound connections direct to port 25 to any destination other than the local mail relay. Spam email is often generated by compromised client systems. A simple *default deny* egress filter which only permitted the mail relay to connect

⁷ See for example <http://www.completewhois.com/bogons/index.htm> and <http://www.cymru.com/Documents/bogon-list.html>

⁸ <http://www.cymru.com/>

out would prevent a compromised corporate system from generating successful spam. Similarly, no internal system other than the corporate proxy server has any need to connect out to a web server. The same *default deny* stance would prevent desktop clients making direct outbound connections.

18. To complement the ingress filters which deny inbound traffic purporting to have internal or RFC3330 source addresses, the egress filters applied at the organisational boundary should drop any packet having a source address which is not a legitimate internal routable address. A NAT firewall should be inbound of the filtering router and will handle translation to/from reserved addresses. If all organisations denied outbound non-legitimate traffic, it would greatly ease the load on the network posed by traffic with spoofed source addresses. Furthermore, denying outbound traffic except from legitimate internal addresses will hinder malicious software which may have bypassed other filters (such as spam or virus/Trojan filters) and infected the network.

19. Malicious software will often attempt to open communication channels to external servers. Such activity is commonly associated with botnet⁹ software which 'calls home' to a command and control centre, often over channels associated with IRC. Here again, a properly enforced default deny egress filter would prevent such connectivity.

20. In most organisations, outbound traffic will normally only be necessary to a very limited range of services and from a very limited range of hosts. Typically this will comprise:

- from the corporate DNS server to external DNS servers on UDP and TCP port 53¹⁰;
- from the corporate web proxy server to external web servers on TCP ports 80 and 443 (and possibly others such as 8000, or 8080 depending on local policy);
- from the corporate mail server to external mail servers (or to an upstream relay hosted at the ISP) on TCP port 25.

21. All other outbound connection attempts should be dropped. Note that there is normally no legitimate reason for a corporate desktop client to make any direct external connection.

Broader business decisions

22. In addition to the default filtering policy outlined above, all businesses should consider applying additional policy blocks on connectivity to networks or

⁹ See NISCC Briefing 11a/05 – Botnets, the threat to the Critical National Infrastructure

¹⁰ See <http://homepages.tesco.net/~J.deBoynePollard/FGA/dns-shaped-firewall-holes.html> for some advice on DNS filtering.

domains with which they have no need of communication. For example, compromised home PCs, usually on high bandwidth 'always on' DSL or cable networks are the source of much spam email. Security company Sophos claims that 60% of all spam emanates from compromised 'zombie' computers on DSL connections. Furthermore, some areas of the world are more likely than others to host such compromised machines. Indeed, Sophos'¹¹ survey of its spam traps for April to September 2005 indicated that the United States alone was responsible for over 26% of all worldwide spam. The top three countries (US, South Korea and China) between them accounted for around 62% of all spam. Whilst this does not mean that those countries are the ultimate source of the email, it does nevertheless mean that compromised machines in those countries are potentially hostile and could usefully be denied inbound access to corporate mail servers.

23. Given that the vast majority of malicious email emanates not from legitimate mail servers, but rather from compromised PCs on domestic DSL or cable networks, ingress filters could be employed to block inbound mail access from netblocks known to host such services. Subscription to real-time blocking services (DNSRBL) such as Spamhaus¹² can assist in this process, but such filtering takes place at the application level (SMTP). Packet filtering of whole netblocks known to contain only domestic or dynamic DSL networks can be equally effective, but must be implemented with care. Sites such as SORBS¹³ and Epaxsys¹⁴ can be useful in determining such dynamic netblocks, whilst Spamstopshere¹⁵ also lists "country of origin" filters to allow administrators to block whole geographic domains. Where there is any doubt about the legitimacy of a particular network address or addresses, they can be checked against lists of known malicious sites through aggregate spam checking sites such as RBLs¹⁶ before they are included in filters.

24. NISCC Briefing Note 08/05¹⁷ on Trojan activity noted that the Trojans often communicated back to the attackers using standard application ports such as TCP port 80. Since the majority of internet connected organisations will permit outbound web connectivity, this makes blocking the attacks problematic. However, as we also noted in the same Briefing Note, IP addresses associated with the attacks are often associated with the Far East.

25. We recognise that there may be a considerable management overhead in maintaining boundary ACLs where the network addresses to be blocked are non-contiguous, or change rapidly over time. Nevertheless, businesses are strongly recommended to consider the feasibility of applying egress filters which will block

¹¹ See SOPHOS website at

http://www.sophos.com/pressoffice/news/articles/2005/10/sa_dirtydozooct05.html

¹² <http://www.spamhaus.org/>

¹³ <http://www.nl.sorbs.net/>

¹⁴ <http://www.epaxsys.net/dnsbl/>

¹⁵ http://www.spamstopshere.com/antispam_howitworks.aspx

¹⁶ <http://rbls.org/>

¹⁷ NISCC Briefing 08/05 – Targeted Trojan Email Attacks

outbound web connectivity to geographic areas with which they have no legitimate need to communicate. So for example, if your organisation has no business requirement to connect to web servers in the Far East, an outbound policy block on the relevant IP address ranges should be applied. **Annex A** attached lists some resources which may be useful in mapping Geographic locations to IP addresses. Note that this information changes over time so businesses are encouraged to subscribe to one or more Geographic IP database provider in order to ensure that they have the most up to date information.

Annex A – Geolocation Reference Sites

Free IP Geolocation Sites

- <http://www.32tech.com/ip-to-country/>
- <http://ip.ludost.net>
- <http://ip-to-country.webhosting.info>
- <http://www.maxmind.com/app/country>
- <ftp://ftp.apnic.net/pub/stats/apnic/apnic-latest>
<ftp://ftp.apnic.net/pub/stats/ripe-ncc/ripenncc.latest>
<ftp://ftp.apnic.net/pub/stats/arin/arin.<date>>
<ftp://ftp.apnic.net/pub/stats/lacnic/lacnic.<date>>

Commercial IP Geolocation services

- <http://www.activetarget.com>
- <http://www.ip2location.com>
- <http://www.ip2country.net>
- <http://www.quova.com>