



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.

1. Block "spoofed" addresses-- packets coming from outside your company sourced from internal addresses or private (RFC1918 and network 127) addresses. Also block source routed packets.

a. Block spoofed internal address:

Inbound traffic should never have an internal IP source. This would indicate that a system outside of the internal network had attempted to 'masquerade' as a trusted internal system in order to infiltrate the internal network to possibly do harm. The basis for this is the fact that many system-to-system interactions are authenticated merely by IP address. For example, when a user with an internal IP address attempts to access an internal server, the IP address is inspected, and when found to be from the internal network, accepted as a trusted host. Now, imagine that an outside source has spoofed this same internal IP address. The Internal server would again assume this was an internal trusted host and allow access to potentially sensitive and costly information.

b. Block private address schemes:

Again, this is an inbound rule which explicitly denies traffic coming from outside of the internal network that presents itself with a 'Private IP address' (i.e. 10.0.0.0 – 10.255.255.255; 172.16.0.0 – 172.31.255.255; 192.168.0.0 – 192.168.255.255) which are reserved for internal, non-internet use only. No traffic originating from outside your firewall should ever have these addresses assigned, as indicated in RFC 1918.

c. Block loopback address:

Any loopback address, (the 127.0.0.0 network), is reserved for host configuration testing. If a loopback address is used, it is solely for IP stack testing and will never leave the host. An IP address on the 127.0.0.0 network, would never be a valid IP address. In fact, every host has the same loopback address assigned programmatically by default as part of their operating system. For example, when anyone needs to verify that their system's IP stack is properly configured, they would ping their loopback address and the IP stack is traversed from the Application layer to the network layer and then back to the Application layer. At no point does this request reach or leave the physical NIC. There should never be inbound or outbound traffic on the network with a source or destination IP in this address range.

d. Block Source Routing:

Using IP loose source routing may allow an external system to re-route a malicious packet to a remote location by bypassing the normal access list. Source routing has to be disabled

Filters used:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Service/Protocol</u>	<u>Action</u>
1a	Inbound	Internal	Any	Any	Deny
1b	Inbound	Private	Any	Any	Deny
1c	Either	Loopback	Any	Any	Deny

1a – For inbound traffic, having an internal IP address, going to any destination on the network, for any service/protocol, the packet will be dropped.

1b - For inbound traffic, having a private IP range address, going to any destination on the network, for any service/protocol, the packet will be dropped.

1c - For either inbound or outbound traffic, having a loopback address, going to any destination on the network, for any service/protocol, the packet will be dropped.

To test this filter, attempt to connect with any of the aforementioned IP addresses from an external host through your firewall. Any attempts will result in the packets being dropped.

2. Login services-- telnet (23/tcp), SSH (22/tcp), FTP (21/tcp), NetBIOS (139/tcp), rlogin et al (512/tcp through 514/tcp)

Services such as telnet, SSH, FTP, Netbios and the Berkeley “r” commands all allow connections and interaction with remote hosts and servers. The fear of a system being compromised by any of these services is a valid one. Each has it’s own unique weakness. The issue for all of these services is, do we have trusted external sources with which we can allow any these services to be utilized.

a. Telnet allows a user to connect over TCP port 23 to any telnet server listening for a request. Telnet uses clear text as a default communications method. This means that even with authentication, the user exposes his system to potential hijacks. Anything traveling over a telnet session is vulnerable to anyone using a sniffer to monitor traffic. Once connected, the user can issue a number of commands to the server, which in turn are performed with no further authentication requests. For example, this is a commonly used ploy to hijack a mail server and spoof a trusted source within the network. However, by prohibiting Telnet over TCP port 23, be aware that you are also limiting your own ability to troubleshoot a system from a remote site (i.e. outside of the firewall).

b. SSH allows encrypted sessions to take place over telnet or through rep using TCP port 22. The advantage of this is the ability to still use remote management without

the fear of infiltration due to clear text communications. Unfortunately, if your system is compromised in a way that allows an intruder to gain root access to a machine, SSH can also be undermined. The issue here, is whether or not there is an overriding need exists to allow users to copy/transfer remote files, or to allow external users to gain access to internal resources. Without some outstanding need, this service should be disallowed.

c. FTP is a service that allows the transfer of data between systems over TCP port 21. Like telnet, FTP uses clear text communication. This includes authentication. That said, it would be possible for a hacker with a destructive intent to sniff the FTP traffic and gain access to potentially sensitive data with the authentication information they uncover. If we had a DMZ, FTP could be used to update web sites, however in our scenario, this is not stated nor assumed to be necessary.

d. Netbios is a Microsoft Windows service that uses runs over TCP port 139. It is used to establish sessions for file and print sharing and is enabled by default on Windows NT operating systems. A user indicates that they wish to 'share' a folder by setting this active in the object's properties. Once enabled, the default permission is to allow everyone access the share with full control of the object. Passwords and user permissions can be set, but the average user may not be savvy enough to implement this properly and merely accepts the default. Without these 'shares' being protected, any outside user can scan port 139 and find the open resources. Once found, they are fair game to any intruder who wishes to read, alter or delete any files or folders held within the share. By blocking port 139 both inbound and outbound, you can eliminate the this risk.

e. Berkeley "r" commands include rlogin, rsh, rdist, etc. These remote service commands can allow a user to accomplish remote command execution and terminal access to exchange or manipulate files on a local host from a remote location. These commands often use IP addresses for authentication so a spoofed address could technically become mistakenly authorized to access a host. Any host attempting to use these services, in either direction, should be denied. It is necessary then to prohibit the use of TCP over port 512, 513 and 514

FiltersUsed:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Service/Protocol</u>	<u>Action</u>
2a	Either	Any	Any	Telnet	Deny
2b	Either	Any	Any	FTP	Deny
2c	Either	Any	Any	SSH	Deny
2d	Either	Any	Any	NBT	Deny
2e	Either	Any	port 512-514	TCP	Deny

2a – For either inbound or outbound traffic, attempting to connect from any source, going to any destination on the network, for telnet services, the packet will be dropped.

2b - For either inbound or outbound traffic, attempting to connect from any source, going to any destination on the network, for FTP services, the packet will be dropped.

2c - For either inbound or outbound traffic, attempting to connect from any source, going to any destination on the network, for SSH services, the packet will be dropped.

2d - For either inbound or outbound traffic, attempting to connect from any source, going to any destination on the network, for Netbios services, the packet will be dropped

2e - For either inbound or outbound traffic, attempting to connect from any source, going to port 512 - 514 on the network, for TCP services, the packet will be dropped

To test this filter, attempt to connect through any of the aforementioned services from an external host through your firewall or from an internal host to the outside. Any attempts will result in the packets being dropped.

3. RPC and NFS-- Portmap/rpcbind (111/tcp and 111/udp), NFS (2049/tcp and 2049/udp), lockd (4045/tcp and 4045/udp)

a. RPC (remote procedure calls) is a client/server process that is built on top of TCP and UDP, port 111(for both). It is a technology that is used by many client/server applications that need to send small requests from a client for work to be done on a server and then results returned. Because there is a need to track each RPC request from a client, and to determine what server service to send the request to for processing, RPCs use a portmapper (Unix), or RPC locator service in NT. The portmapper, or locator service, tells the client where to send its request. Ports are dynamically assigned for each service upon startup, so this can be a tricky service to control or audit. Use of RPC inbound or outbound may represent an external source attempting to use an RPC server on an internal host to proxy an attack. Because of these issues, it is recommended to block all UDP, which 'dangerous' RPCs generally run over, however, with this approach, a second rule must be implemented above this one to allow UDP port 53 for DNS resolutions to be processed. TCP RPCs can also be prohibited inbound, and you may also want to monitor TCP port 111 if any traffic is allowed outbound.

b. NFS is a protocol that enables access to files on a remote system. NFS typically uses port 2049 over TCP and UDP to listen for client requests. Here, again, we have weak authentication on the remote file system. Client systems are allowed access to file resources based on IP address. Spoofed IP addresses could gain access to a seemingly secure system. Blocking all inbound and outbound NFS traffic is recommended.

c. *lockd* is a method used by NFS that prevents two users from making changes to the same file at the same time. If the server on which the lock exists is restarted, the clients with locks on files are required to resubmit them. During this time, files can be corrupted or data lost. A service, *statd* will attempt to save the information regarding *lockd* under its own root authority to avoid any data losses, without re-authenticating through the client who initiated the lock. The vulnerability here, is that *statd* can potentially allow a remote program to access file resources with authority of the root user. It is recommended that *lockd* running over TCP and UDP port 2049 be blocked inbound to avoid an outside source from hijacking a session.

Filters used:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Service/Protocol</u>	<u>Action</u>
3a	Outbound	Port >1023	Trusted DNS- Port 53	UDP	Allow
3a	Inbound	Trusted DNS- Port 53	Port >1023	UDP	Allow
3a	Either	Any	Any	UDP	Deny
3b	Either	Any	Any	NFS	Deny
3c	Inbound	Any	Any	lockd	Deny

3a – Because we will prohibit UDP to eliminate RPC risks, we must first allow DNS resolution from client to server and back. Thus UDP port 53 must be available to a client to send a request outbound, and UDP port 53 must be accepted as a source to send the response to the client initiated query. I am using Source and destination as Trusted DNS in reference to policy number 6. Lastly, we block either inbound or outbound traffic, attempting to connect from any source, going to any destination on the network, for UDP.

3b – For either inbound or outbound traffic, attempting to connect from any source, going to any destination on the network, for NFS services, the packet will be dropped.

3c - For inbound traffic, attempting to connect from any source, going to any destination on the network, for *lockd* requests, the packet will be dropped.

To test these rules, attempt a DNS query from any client to insure it is properly processed. Then attempt any RPC request, or NFS access from a client. Also, from outside of your firewall, attempt spoofed access to an NFS file utilizing *lockd* port 4045.

4. NetBIOS in Windows NT -- 135 (tcp and udp), 137 (udp), 138 (udp), 139 (tcp). Windows 2000 - earlier ports plus 445(tcp and udp)

As previously mentioned in regards to policy number 2, Netbios shares are a high exposure feature found in Windows products. It uses several ports that enable various services to communicate with one another over the network. Windows uses Netbios to establish 'shares' in order to allow users to access internal network resources. However, this is particularly vulnerable to attack because of the default nature of the service in Windows NT and the default share level, which is completely open to anyone (i.e. Everyone, Full Control). Other open features of Netbios involve broadcasting available services and shares in plain text within any subnet and also enumerating those shares and services to any requester utilizing the browser services on their network. Another exposure is the use of WINS, the Windows Internet Naming Service. It is similar to DNS as it is responsible for resolving NetBIOS computer names to IP addresses. However, this is a dynamic database of information that also enumerates available services and domain authenticators and other valuable information regarding the network resources. Having a WINS server compromised by an outside source could have devastating repercussions to the entire network.

- a. Port 135* running over TCP and UDP is used to administer the WINS service remotely. Access to a WINS server through this port allows a user to manage the WINS database potentially corrupting the database and causing severe problems with name resolution in a routed network. Essentially, an intruder could edit the WINS database and redirect NetBIOS traffic to its host IP address instead of the valid host's IP address and therefore hijacking potentially sensitive information from another system on the network.
- b. UDP Port 137* is used by client systems to query for name resolution through a WINS server. The potential exists for an outside source to gain information about hosts on a network through this process.
- c. UDP Port 138* is used for the NetBIOS datagram service. Datagrams are sent from one NetBIOS computer name to another and enables systems to transmit messages to a unique name or to a group name, or group systems running the same service. Again, information that may help a hacker infiltrate a network.
- d. TCP Port 139* is used by the NetBIOS to create a TCP session. This allows for communication between different NT services and administrative tools from a remote system. It essentially supports remote server administration in NT.
- e. Port 445* in Windows 2000 supports direct hosting, which means the systems can support communication without NetBIOS. However, we still do not want outside hosts to be able to attach to our internal servers.

It is a good idea to block the use of NetBIOS ports at the firewall. There is seldom a need to support this service between internal and external systems.

Filters used:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Service/Protocol</u>	<u>Action</u>
4a-d	Either	Any	Any	NBT	Deny
4e	Inbound	Any	Port 445	TCP	Deny
4e	Outbound	Port 445	Any	TCP	Deny
4e	Inbound	Any	Port 445	UDP	Deny
4e	Outbound	Port 445	Any	UDP	Deny

4a-d – For either inbound or outbound traffic, attempting to connect from any source, going to any destination on the network, for NetBIOS services, the packet will be dropped.

4e - For inbound traffic, attempting to connect from any source, going to port 445 on the network, for TCP or UDP, the packet will be dropped. For outbound traffic, attempting to connect from port 445, going to any host on the network, for TCP or UDP, the packet will be dropped

To test these rules, you can use a NetBIOS scanner and attempt to search for available NetBIOS shares. If the rule is processed properly, this scan should not find any available shares. Also, test for WINS resolutions by adding the internal WINS server to an external host and then attempting to query the internal system for information. This too should yield no results. You could also attempt to replicate with the WINS server and should have no success.

5. X Windows -- 6000/tcp through 6255/tcp

X-Windows is a GUI system used in UNIX. X Windows distributes the processing of applications by specifying a client/server relationship at the application level. The *what to do* part of the application is called an X client and is separated from the *how to do* part, the display, called the X server. X clients typically run on a remote machine that has excess computing power and displays on an X server. Any system listed in the trusted host list of an X-Windows server can connect to the client list. At that point an intruder can capture and dump client screen content, read the keyboard strokes as username and password are typed, or control applications on a client system.

A system can run multiple X-Window servers the first of which runs at TCP port 6000. Each additional client connection requires a new port on the order of $6000 + n$. So you will have some servers with multiple instances of X-Windows. To block access to these servers and prevent X-Windows exploits, it is best to deny access inbound to the server and from the server to the client.

Filters used:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Service/Protocol</u>	<u>Action</u>
5	Inbound	Any	Port 6000-6255	TCP	Deny
5	Outbound	Port 6000-6255	Any	TCP	Deny

5 – For inbound traffic, attempting to connect from any source, going to port 6000 – 6255, using TCP, the packet will be dropped. For outbound traffic, attempting to connect from port 6000 – 6255 to any source, using TCP, the packet will be dropped

To test this process, attempt to establish an x-windows session to a server outside of the firewall, and conversely, from outside the firewall to a client within the internal network. Also you could run a program like xwatchwin to attempt to view contents on an internal system.

6. Naming services-- DNS (53/udp) to all machines which are not DNS servers, DNS zone transfers (53/tcp) except from external secondaries, LDAP (389/tcp and 389/udp)

a. *DNS* communicates with a client to resolve a name query using UDP port 53. A client to DNS query uses a client source port above 1023 with a destination port of 53 for the DNS. The server to client response uses the reverse method, that is, port 53 returning the information to the client port indicated in the query request. If there is a server-to-server query or response, the source and destination ports are both 53 and use UDP as a transport. This is fairly harmless (never say never) and can typically be used without incident. The only issues here may arise when another DNS system attempts to ‘poison’ the DNS cache on a trusted server by offering invalid IP mappings when resolving client requests.

Another common type of DNS traffic is caused by a zone transfer. A zone transfer occurs when a primary DNS sends a copy of its database to a secondary server. The DNS zone transfer uses TCP with the primary DNS server listening on port 53 for the request from the secondary DNS server. Secondary servers are used to provide load balancing and redundancy for DNS in a network.. Because a zone transfer contains information about entire networks of systems, including identifiers of other DNS systems, mail servers, WINS servers, routers, etc., and their corresponding IP addresses, this information can be extremely valuable to any malicious user attempting to map a network prior to compromising key systems and servers. In fact, the DNS systems are typically the first and main system targeted by an attacker attempting to do harm or gain access to a target network..

It is always prudent to block TCP, port 53 both inbound and outbound to prevent hackers from gaining access to your internal system maps through an unauthorized zone transfer. We can limit access to the port only to known or “trusted” secondary DNS servers.

b. LDAP is a client/server protocol used to access a directory server. For example, it can be used to list public address books and public key certificates. An LDAP server can also maintain sensitive information, for example user authentication information, and could be a target for a hacker who wants to get user account names and passwords. Again, just as with telnet and many other services, the data is not encrypted and can be easily exposed with a sniffer on the network. LDAP servers use port 389 over TCP it is best to prohibit any inbound traffic destined for port 389 over TCP.

Filters used:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Service/Protocol</u>	<u>Action</u>
6a	Inbound	Port >1023	Port 53	UDP	Deny
6a	Outbound	Port 53	Port >1023	UDP	Deny
6a	Inbound	Trusted DNS	Port 53	TCP	Allow
6a	Outbound	Port 53	Trusted DNS	TCP	Allow
6b	Inbound	Any	Port 389	TCP	Deny

6a– For either inbound or outbound traffic, attempting to use DNS resolution outside of the network, access is denied. This assumes that all queries will be handled either by a trusted DNS server (as indicated in policy 3 rules) externally or by an internal host. Also, zone transfers are allowed through the firewall to trusted DNS.

6b - For inbound traffic, attempting to connect from any source, going to port 389 on the network, using TCP, the packet will be dropped.

To test, attempt to use Nslookup with an untrusted host. For LDAP, attempt to connect to an internal LDAP server from an external host.

7. Mail-- SMTP (25/tcp) to all machines, which are not external mail relays, POP (109/tcp and 110/tcp), IMAP (143/tcp)

a. SMTP handles e-mail transmissions between mail servers. Is a store and forward system that either delivers the mail to a local system or forwards it on to a system on another network. SMTP sends clear text traffic over the network with header information, destination and source email addresses, being easily accessible to anyone using a sniffer on the network. SMTP uses TCP port 25. Blocking unwanted SMTP access while still allowing inbound mail necessitates using an intermediary trusted bastion host, and from the bastion host SMTP can be directed to internal mail servers.

Conversely, outbound mail can be directed to the bastion host and then on to external systems.

b. The Post Office Protocol (POP) supports clients who connect to their mail server only long enough to download the mail that the server is holding for them. The mail is then deleted from the server. POP uses TCP port 110. Once again, we have a process that uses client authentication and also sends messages in clear text. A user name and password could easily be discovered using a sniffer. Users often attempt to make their life easier by using the same password throughout their network enterprise, thus if a password is discovered through this method there is a chance that the intruder will have a valid user account and password that can be used to access the entire network. Worse, if sensitive information is being discussed via email, the malicious user can access that information as well.

c. Internet Message Access Protocol (IMAP) also allows users to retrieve mail from a mail servers. IMAP allows a user to download messages or to leave them on the mail server. IMAP uses TCP port 143. Like POP, it too transmits authentication and email messages in clear text.

It is advisable to not allow users to transmit email over the Internet with either POP or IMAP. If it is necessary, make users encrypt sensitive mail messages prior to transmission and try to avoid letting users repeat passwords and user names that they use internally.

Filters used:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Service/Protocol</u>	<u>Action</u>
7a	Inbound	trusted host- >1023	mailserver- Port 25	SMTP	Allow
7a	Outbound	mailserver- Port 25	trusted host- >1023	SMTP	Allow
7a	Inbound	Any	Port 25	SMTP	Deny
7a	Outbound	Port 25	Any	SMTP	Deny
7b	Outbound	Any	Port 110	POP	Allow
7b	Inbound	Any	Port 110	POP	Deny
7c	Outbound	Any	Port 143	IMAP	Allow
7c	Inbound	Any	Port 143	IMAP	Deny

7a – For inbound and outbound traffic, allow SMTP traffic to be passed between a trusted bastion host, and the internal mail server. Deny any other attempts to connect inbound or outbound, using SMTP, port 25.

7b - For outbound traffic, from any source, attempting to connect to port 110, using POP is allowed. For inbound traffic, from any source, attempting to connect to port 110 using POP, packet will be dropped.

7c - For outbound traffic, from any source, attempting to connect to port 143, using IMAP is allowed. For inbound traffic, from any source, attempting to connect to port 143 using IMAP, packet will be dropped.

8. Web-- HTTP (80/tcp) and SSL (443/tcp) except to external Web servers, may also want to block common high-order HTTP port choices (8000/tcp, 8080/tcp, 8888/tcp, etc.)

The real risk with internal web servers, is that system administrators often feel safe to post sensitive data or materials on the server. The belief is that only trusted, internal sources will have access. However, we know this is a naïve stance, and not always the case.

HTTP, using default TCP port 80, can expose an internal web server to several malicious attacks. For example, a malicious client could upload content to the web server and execute it, and then use the web server to distribute the malicious code to all internal clients. To avoid these dangers, one can take some simple steps. For 99.9% of users, access to the server should be read only. This would prohibit a malicious user from uploading a damaging program explicitly to your web folders. Also, take advantage of the host file system security by controlling access to files with specific entries in access control lists. Try to limit the content placed on the web server to non-sensitive information. This is not a good place to store any information you can't afford to lose. Also, using TCP ports 8000, 8080, and 8888, as non-standard ports to provide specialized HTTP data needs to be addressed. Try to avoid using them or any other non-standard ports for clients to get HTTP content.

With the idea of protecting the internal Web server, make sure to block all inbound TCP traffic destined for port 80, or any other non-standard ports, i.e. 8000. It is a good idea to make sure your internal Web servers are also not responding to any external requests that may have unknowingly evaded your perimeter security by blocking any outgoing traffic originating from the web server using these ports.

Filters Used:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Protocol</u>	<u>Action</u>
8a	Inbound	Any	80, 8000, 8080, 8888	HTTP	Deny
8b	Outbound	80, 8000, Any 8080, 8888		HTTP	Deny

8a – For inbound traffic, attempting to connect from any source, going to port 80, 8000, 8080, 8888 using HTTP, the packet will be dropped.

8b - For outbound traffic, attempting to connect from port 80, 8000, 8080, 8888 to any source, using HTTP, the packet will be dropped.

To test, attempt to use any browser to connect with HTTP from an external host to an internal web server using port 80, 8000, 8080, or 8888.

9. "Small Services"-- ports below 20/tcp and 20/udp, time (37/tcp and 37/udp)

TCP and UDP services such as Echo, Discard, Chargen, and Daytime are rarely used and may be the target of a potential exploit by a creative hacker. It is generally recommended that these services be disabled and their ports blocked. The following are small services that should be blocked Echo, Discard, Chargen, Qotd, and time service.

Filters Used:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Protocol</u>	<u>Action</u>
9a	Inbound	Any	Port 7	TCP	Deny
9a	Inbound	Any	Port 7	UDP	Deny
9b	Inbound	Any	Port 9	TCP	Deny
9b	Inbound	Any	Port 9	UDP	Deny
9c	Inbound	Any	Port 19	TCP	Deny
9c	Inbound	Any	Port 19	UDP	Deny
9d	Inbound	Any	Port 17	TCP	Deny
9d	Inbound	Any	Port 17	UDP	Deny
9e	Inbound	Any	Port 37	TCP	Deny
9e	Inbound	Any	Port 37	UDP	Deny

9a. For any inbound traffic, from any source, using destination port 7 (Echo), over TCP or UDP, drop the packet.

9b. For any inbound traffic, from any source, using destination port 9 (Discard), over TCP or UDP, drop the packet.

9c. For any inbound traffic, from any source, using destination port 19 (Chargen), over TCP or UDP, drop the packet.

9d. For any inbound traffic, from any source, using destination port 17 (Qotd), over TCP or UDP, drop the packet.

9e. For any inbound traffic, from any source, using destination port 37 (Time), over TCP or UDP, drop the packet.

To test, attempt to connect to any of the above listed services from outside of the network. All attempts should be blocked.

10. Miscellaneous-- TFTP (69/udp), finger (79/tcp), NNTP (119/tcp), NTP (123/tcp), LPD (515/tcp), syslog (514/udp), SNMP (161/tcp and 161/udp, 162/tcp and 162/udp), BGP (179/tcp), SOCKS (1080/tcp)

a. Designed for diskless workstations, and router configuration, TFTP, using UDP port 69, provides no authentication to the service. If a file the client requests is available, the TFTP server provides a copy of the file, including routing table information. Unwanted access cannot be prevented so TFTP should be blocked.

b. The finger service is an information lookup service that provides information about users recently logged onto the system and the name of the system where they logged on. This can be useful for a hacker attempting to gather user names and system names in recon attempts. The finger service running at TCP port 79 should always be blocked.

c. NNTP, using TCP port 119, transfers Usenet information from one system to another across the Internet. As with SMTP, it is preferable to run an NNTP server outside of your firewall on a bastion host and not on an internal network. Therefore all external traffic bound for an internal NNTP server should be blocked. Additionally, to prevent internal users from downloading malicious content from NNTP servers from outside of your network, limit internal clients to connect to and get information only from trusted NNTP servers.

d. Network Time Protocol (NTP) enables users to set the clocks on their systems according to an authoritative time source such as an atomic clock. It operates using UDP port 123. The issue with NTP is that there it can be easily modified to give packets the wrong time settings because it has not authentication method. As a result it is recommended that you create an internal timeserver through which internal systems synchronize their times. All NTP traffic going to an external timeserver and the responses from the external timeserver should be blocked.

e. *LPD, Line Printer Daemon*, is the print server service using TCP port 515. Unless you are required to support printing for clients from across the Internet, this port should be blocked to prevent any creative exploitations.

f. *Syslog*, using UDP port 514, listens for log messages sent by other systems across the network. It is a method of centrally handling logs and usage information from multiple machines on the network. Syslog is a target of hackers who will flood port 515 with messages in an attempt to erase any evidence of their being on a particular site as the server runs out of disk space and ceases recording any new messages. It is recommended that syslog only be run internally and any external calls inbound to port 515 be blocked.

g. *Simple Network Management Protocol (SNMP)* is used to remotely manage systems. Using UDP ports 161 and 162. This service works by using an agent service on the client hosts to gather pertinent information about the host system configuration. The information is then sent to a management system where it can be collected and stored for lookup information. Trivial security is provided in authenticating managers in agents in the form of a common community name. Unfortunately, it is a common practice for administrators to leave the default community name of “public” as their community. This allows anyone using the same community to connect to a system posing as a manager and then to gather this sensitive system information. It is recommended that external connections using SNMP be blocked as well as responses from internal agents back to management systems. SNMP also allows management systems and agents to communicate in the form of traps. A trap is a form of a threshold or alert, that when exceeded causes the agent system to send information back to a management system. If subverted, an external management system could gather valuable information about systems running an SNMP agent. So it is recommended to block this type of activity from an external system management system. Normal SNMP management queries use port 161 with traps using port 162.

h. *BGP, Border Gateway Protocol* systems/routers, are located at the perimeter of an autonomous system (AS). Their function is to facilitate the routing of all packets having a destination network address outside of the AS. Most Autonomous Systems have only a single BGP router making a single point of failure and a likely target for malicious activity. As this is used for external interaction, it is not an option to prohibit traffic going in or out of these routers. BGP routers are vulnerable to TCP SYN flooding, RST attacks, DATA insertion attacks, and “Session Hijacking” attacks. Attempts have been made to secure BGP connections through encryption, but the algorithm has its flaws. No other action can be taken to secure this particular known vulnerability.

i. *SOCKS* acts as a proxy between a client on the internal network and a server across the Internet. It is designed to allow systems sitting behind a firewall to access the internet without exposing the system to attack. SOCKS handles the connection request, the Proxy circuit establishment, relaying the application data, and authentication. SOCKS uses TCP port 1080.

Filters Used:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Protocol</u>	<u>Action</u>
10a	Either	Any	Port 69	UDP	Deny
10b	Either	Any	Port 79	TCP	Deny
10c	Inbound	Any	Port 119	TCP	Deny
10c	Outbound	Any	Trusted NNTP- Port 119	TCP	Allow
10d	Outbound	Any	Port 123	UDP	Deny
10d	Inbound	Port 123	Any	UDP	Deny
10e	Outbound	Any	Port 515	TCP	Deny
10e	Inbound	Port 515	Any	TCP	Deny
10f	Inbound	Any	Port 514	UDP	Deny
10g	Either	Any	Any	SNMP	Deny

10a For any inbound or outbound traffic, from any source, to port 69 , using UDP, drop the packet.

10b. For any inbound or outbound traffic, from any source, to port 79, using UDP, drop the packet.

10c. For any inbound traffic, from any source, using destination port 119, using TCP, drop the packet. For any outbound traffic, from any source, using destination port 119 on a trusted NNTP server, over TCP, allow the packet.

10d. For any outbound traffic, from any source, using destination port 123, using UDP, drop the packet. For any inbound traffic, from port 123, to any destination, using UDP, drop the packet

10e. For any outbound traffic, from any source, using destination port 515, using TCP, drop the packet. For any inbound traffic, from port 515, to any destination, using TCP, drop the packet.

10f. For any inbound traffic, from any source, to port 514 , using UDP, drop the packet

10g. For any inbound or outbound traffic, from any source, to any destination , using SNMP, drop the packet

10h. No action was taken.

10i. No action was taken.

To test, attempt to connect to any of the above listed services from outside of the network. All attempts should be blocked.

11. ICMP-- block incoming echo request (ping and Windows traceroute), block outgoing echo replies, time exceeded, and unreachable messages

ICMP is considered one of the most vulnerable protocols that is utilized between IP host systems. It's purpose is to aid in troubleshooting and to deliver error messages in a IP environment. Unfortunately, because of the reporting and 'helpful' nature of this protocol, it is often a source of information and exploitation by malicious uses. Made famous for the Denial of Service attacks of late, such as the SMURF amplification attacks against Yahoo and others. This protocol should never be allowed outside of your trusted network.

ICMP Echo requests and replies, often referred to as the 'Ping' service for ICMP, determine whether a host is available on a network. The ping command is often used to troubleshoot when a connection cannot be made to another host on a network. By means of NetBIOS or name resolution of DNS or to test whether a default gateway is available and so on. Hackers soon realized that ping could also be used to map a network of hosts available. For example, by sending echo requests to a network segment of host ids, the initiator can determine which hosts are in use by the replies received through ICMP. If the host is available, an echo reply is sent, if not a 'host unreachable message may be sent. Ping can also be used to overflow the buffer on a victim host and render the system useless.

b. Traceroute is a UNIX command line utility used to display the routers a source host traversed to reach a destination host. While there is no harm in using this tool as an aid to determine where a problem router might exist when a host is unreachable, traceroute could also be used to determine internal system's router addresses by an outside source. By combining TTL values, with traceroute, one could determine exactly when and where the hop count is exceeded and use this information for mapping.

c. Time Exceeded is a message sent by ICMP whenever a packet's TTL reaches 0. This can be used to determine how many hops a packet takes, that is, how many routers it traverses, before it reaches a host. Again, information that can be used to map a network form the outside. Hackers will decrement the TTL until the Time Exceeded messages

begin to arrive. Then, they document this information as they collect a map of your internal network structure.

d. Host/Domain Unreachable is another ICMP message that can be used in the mapping process. This message will indicate whether a host is listening at certain ports or if it is available at all. Again, this information can be used to map resources in a vulnerable domain.

It is a good idea to not allow ICMP to travel into or outside of your network. For this reason, inbound and outbound ICMP should be blocked. Other values that may be set specifically to filter out echo requests and replies, Destination Unreachable messages, and Time exceeded messages, are to use content filters to determine the type of ICMP being used. For example: echo reply uses message type 0, echo requests use type 8; Destination Unreachable is message type 3 and Time exceeded is type 11. It is possible to only filter this particular type of traffic, but for simplification of the matter, and also because ICMP is an easy protocol to manipulate or hijack, I have disallowed all ICMP.

Filters Used:

<u>Rule</u>	<u>Direction</u>	<u>Source</u>	<u>Destination</u>	<u>Protocol</u>	<u>Action</u>
11a-d	Either	Any	Any	ICMP	Deny

11a-d – For either inbound or outbound traffic, attempting to connect from any source, going to any destination on the network, using ICMP, the packet will be dropped.

To test, attempt to ping, and use traceroute both from inside the firewall to an outside source, or from an external host to a host inside the firewall.

In conclusion, this rule base was developed without the benefit of a business needs analysis. Some rules would have to be modified, and perhaps changed to alerts instead of a simple logging scheme in light of possible business needs that require any of the aforementioned services. For instance, the knowledge of a DMZ existing would lean us towards split DNS and other management needs that may permit us to use SSH. It clearly shows the necessity of documenting a full security policy **before** implementing your perimeter defense.

Thank you for the opportunity to increase my knowledge of these common vulnerabilities. This has been a rich learning experience from which I have greatly benefited.