

Network security

4.1 Use of Filters Within Routers for Network Access Control

Although the main purpose of a router is to transfer IP packets from one network segment to another, most routers can also make decisions as to which packets will be allowed to pass. These decisions generally use some sort of rule or pattern. Routers may be configured with a list of rules to select either a set of packets which will be transmitted, or a set which will be blocked. Rules may be based on a number of different aspects of the packet, for example the source or destination host IP address (or network number), or a specific source or destination port, or flags within the UDP or TCP headers that the IP packet contains. Firewalls are advanced systems that perform the same routing and filtering tasks but with more sophisticated rules and more processing power.

These mechanisms can be used to block or permit access to specific services or specific computers, since the listener for a particular service will usually be found at the same 'well-known' port number. For example, the *http* service is deemed by convention to operate on port 80 - therefore any *http* server process expecting to accept an incoming request will do so via TCP port 80 and any *http* client will specify TCP port 80 when making its call. A router which discards IP datagrams with a TCP destination port of 80 will therefore block normal *http* requests between hosts on opposite sides of the router. It is important to note, however, that this block can easily be circumvented if the host and client agree that the *http* service will be provided on some other, non-standard, port. Also, a router block can only be effective if the traffic does in fact have to pass through the router.

Note also that since most computers now support two different versions of the Internet Protocol - IPv4 and IPv6 - any network controls need to be applied to both protocols otherwise security problems will merely switch protocol. IPv6 systems may even create their own tunnels, potentially bypassing firewalls and other network controls, if they are unable to communicate. It may well be better to provide some local, managed, support for IPv6 connectivity rather than have devices create their own, unknown, alternatives.

Packet filtering techniques can be used on any router or firewall connected to the institution's network to control the flow of traffic around that network. If all computers and users on the network have perfect security measures in place, as described in previous sections, then the only need for such filtering is to control the quantity of traffic. However, in the more common situation where the number of computers and users makes it difficult to guarantee that correct system management is in place everywhere, filtering the network can also be useful in restricting access to many common vulnerabilities. One possible place to apply such filters is on the router which connects the organisation's network to Janet, since this can restrict access from external hosts to internal and vice versa. The advantages and disadvantages of

this approach are described in the following sections.

4.1.1 Benefits of Network Filtering

The main benefit of filtering using network control devices, whether routers or firewalls, is that in a diverse network it can provide single, centrally-managed, points of control. Filters installed on an organisation's Janet router can control traffic between external hosts and all internal hosts without the need to visit each individual machine. Services and computers that do not need to be visible from the external network can be blocked at the router, immediately reducing the risk of direct external attack against those computers. The effort required to enforce good security on individual computers can then be concentrated, at least in the first instance, on the small subset of computers and services which are provided to the outside world.

Network filtering provides protection against external attack for computer systems that cannot be protected themselves. This may be because insufficient effort is available to configure and maintain them properly, or because they are required to provide - to a limited, internal audience - services that are known to have fundamental security problems. For example, some services such as directory access may use only weak authentication methods so should probably not be made available across a wide area network. However, within a trusted section of the network, such as a research group, such services can be used with reasonable safety provided access from the rest of the network is blocked by a network filter.

Even services that are thought to be secure can benefit from the protection offered by network filters. New bugs are discovered frequently, and some of these may have security implications. A common means of attack is to exploit a previously unknown bug in an existing service to gain privileged access to the computer system. A network filter can make it harder for attackers to perform these "zero-day" attacks.

Since many IP services use the IP address of the client machine as part, or all, of the authentication process, intruders may attempt to gain access by 'forging' IP datagrams to make it appear as if they are a trusted computer. Network filtering can, and should, be used to prevent this form of attack from succeeding. Most routers, whether within a site or connecting to Janet, should have a well-defined set of IP addresses on at least one side of them. If a packet from one of these addresses arrives from the 'wrong' side of the router then it should be treated as suspicious and discarded. Likewise, if a packet arrives from the well-defined side of the router with an address other than those expected, it should be discarded and investigated. There are also some IP addresses that should never be seen on the public Internet and should be dropped if they attempt to enter or leave an organisation: Team Cymru publish an up to date list of these '[bogons](#) ^[1]'.

4.1.2 Limitations of Network Filtering

Filtering using network control devices is not a complete solution and should never be regarded as a substitute for good, and improving, computer security. Relying on filtering alone places considerable demands on the configuration and operation of the filtering device. Any error here could leave the network with no protection at all. In the past, bugs have been discovered in filtering software that allowed attackers to avoid blocks. If filtering is not backed up by other forms of protection, then the discovery of such a bug would represent a very serious threat to all the vulnerable computers and users that were thought to be protected by

the filters.

Filtering rules can quickly become complex and hard to manage. Since rules may interact, it is often hard to determine the correctness of a given configuration and whether it is actually implementing the intended security measures. The application of filtering rules also imposes an extra processing load on the device which implements the rules since each packet must be examined to determine whether it should be transmitted or not. At high traffic levels and with complex sets of rules, the performance of the router may be affected. Firewalls are designed to handle more complex rule sets and may have additional tools to make configuration easier: however they too have limits on the complexity and traffic that they can handle.

Network filtering works at the IP and TCP/UDP levels of the traffic flow and therefore does not have sufficient information to make some types of filtering decisions. Electronic mail is a well-known example. Simple packet filtering can control which computers can send and receive mail but cannot make decisions based on the individual sender or recipient. World Wide Web traffic is similarly difficult to control, especially if web caches are used. For decisions requiring access to higher-level protocols, the most common solution is to run a program called a 'proxy' on either the firewall or a dedicated device alongside it and use router filters to force all traffic to pass through the proxy. Proxies are written to understand a particular high-level protocol and can make decisions based on all the information contained within the communication. However that specialisation means that it is usually necessary to run a different proxy for each service. In some cases recent firewalls are able to inspect higher level protocols and make filtering decisions based on their parameters, but this functionality is usually not as great as that from a dedicated proxy.

Any filters that rely on port numbers to identify particular services will only work if the convention on the use of 'well-known' ports is followed. It is possible for the administrator of a service to decide to run a service at a different port and advertise this port number either widely or just to the intended audience. A service that has been relocated in this way will not be identified correctly by filters based on port number. This may have undesirable effects, either permitting access to a service that is supposed to be blocked or preventing access to a legitimate service. Both situations may arise with respect to services provided by computers within the organisation and services provided outside the organisation which internal users wish to access. Where a block is circumvented by running a service on a non-standard port, it is very unlikely that this will be detected by the network managers unless there are other reasons for suspecting what is happening.

Efficient filters will usually be based on ranges of IP addresses, rather than addresses of individual computers. This requires planning in the initial allocation of addresses - computers with similar functions and security requirements should be grouped into address blocks - and also requires all computers in each address block to achieve the same level of security. So long as a sub-network contains some 'unsafe' hosts, the filtering rules will be constrained by the need to protect those computers and the other, well-configured, computers will be unable to use services that might, on an individual basis, be an acceptable risk. Pressure to relax such restrictions can lead in time to a reduction in the protection offered by filtering. Either the filter rules will become more complex due to a greater number of exceptions or else important controls will be reduced in effectiveness.

Network filtering cannot, of course, protect against attacks that do not need to pass through the filtering device, or from many attacks that exploit services that are allowed to pass through

the filter. Good computer security is the only effective protection against attacks from within the network, either by malicious local users or by outsiders who have succeeded in compromising a computer somewhere on the local area network. Poorly-configured modems, wireless networks and other systems which provide alternative routes into the LAN can also allow filtering rules to be bypassed. If computers are secured against these kinds of attack then they will automatically become secure against the same attacks from the external network. Network filtering then becomes a safety net to protect against mistakes in configuration and newly discovered vulnerabilities.

4.1.3 Recommendations

From the considerations above, the following recommendations can be made.

- Filtering in network control devices can be an effective step in reducing the number of successful attacks against computers connected to the network. It is most effective in increasing overall security on diverse networks when used to protect and isolate computers or sub-networks that are hard to protect by other means. The most effective protection will be achieved by filtering to permit only secured services and computers. Filtering which relies instead on blocking services and computers which are known to be insecure is easy to evade and very hard to sustain against newly-discovered threats.
- Filtering in routers or firewalls should never be relied upon for security, and should not be regarded as an alternative to adequate computer security. Computers which are visible to the Internet will need to be secured individually in any case, and even internal computers should be secured individually as far as possible. A security breach on an internal server immediately renders all local network filtering ineffective, since the point of control can be bypassed. Only adequate computer security can protect against attacks either from, or routed through, other hosts on the same network.
- Any security measures, whether implemented by computer security or by network filtering, must be supported by organisational policy and user understanding. Without these, there will be continuing pressure from users to enable access to an increasing number of services and computers. If these requests are accepted, the security measures will become increasingly complex and hard to maintain. If requests are refused without sufficient reason then users will inevitably find ways to evade the controls. Especially in student accommodation and classrooms, the filtering of traffic has been seen to lead to an arms race where users try to find the next means of bypassing filters before system managers discover it.

4.2 Designing Network Security Measures

Wherever possible, security measures should be designed, rather than implemented suddenly as a reaction to an event. A planned implementation is likely to be more effective in preventing attacks, since it should be technically more consistent and more acceptable to users. It should also avoid unnecessary disruption to legitimate use. The details of the process cannot be included here as they will vary from one organisation to another, but there are some common themes which should apply to all security projects. For more details see our guide to [firewall implementation](#) [2].

4.2.1 The Importance of Policy

The security measures adopted by each organisation need to strike a balance between two risks:

- the risk that Information and Communications Technology will not be sufficient to support the organisation's activities and
- the risk that Information and Communications Technology will be misused to harm those same activities.

Or, in more traditional security terminology, the organisation needs to find the right balance between confidentiality, integrity and availability of information and information systems. If security is too open then a major security incident could damage the organisation's research and education functions for days or weeks, and its reputation for months or years: if security is too closed then activities needed to carry out those same functions may be prevented.

The best way to find the right balance is through an organisational information security policy based on an assessment of these risks. A policy that balances risks and benefits can guide and justify decisions on what security controls are needed and the best tools (whether technical, physical, procedural or human) to use to implement them in the particular organisational context. Such a policy should identify those information services that the organisation needs and ensure they are provided in a secure way to the individuals who need them. Policies and services may well need to change over time as new technologies, new requirements, new risks and new expectations emerge, but making these changes within an overall information security process should ensure that the right balance of risks and opportunities is maintained.

Information security policies should be viewed by an organization as an integral part of their activities and not simply as rules handed down by the Information Technology function. It is important that the policy has the explicit support of senior management within the organization, so that its support of the organization's objectives is clear. If information security is seen purely as an IT function then resourcing and enforcement of policy can be an issue. Information security policies should be managed in the same way as other organization wide policies.

Clearly it is possible to implement technical security measures without an information security policy, but doing so increases the likelihood that an important security control will be missed, creating an unnecessary risk, or that security measures will interfere with a part of the organisation's function that their designers were not aware of. Both problems tend to be dealt with by ad hoc measures, rapidly producing a system that is both unmanageable and unusable; less efficient and less secure than it could be. Security measures that are not backed by policy may well end up more costly than those that are.

Developing an information security policy to inform and justify security measures should be a high priority.

4.2.2 Partitioning the Network

The traditional arrangement for a network security system was to have a single control device at the point of entry to the local network. This was designed for commercial organisations where there are clear differences in security between 'inside' and 'outside' the security

perimeter. A brief inspection of the physical security at most university or college campuses will show that the requirements of educational sites are usually more complex. There may be little or sometimes no security at the main entrance, but departments and work areas will often have their own internal security precautions. The same complexity is likely to apply to the organisation's network. A firewall at the perimeter will certainly help to deter attacks from the external network, but it leaves the internal network as a single security domain. Anything that treats students, staff and administration as having equal security needs and risks is clearly not a complete solution.

Most education organisations already partition their internal networks at least to the extent of separating the administrative functions from the main network. Between the two there may be a firewall, a filtering router, or no connection at all. Such internal partitions can normally use cheaper solutions than those at the organisation's connection to the public network, since the traffic flows will generally be smaller and the security policies simpler. Such devices should still be managed centrally. Other partitions within the organisation can be envisaged where different security controls might apply, for example main servers, staff offices, public terminal areas, student residences, research groups, etc. Separating wireless and other networks where personal computers can be connected may also be helpful as these computers cannot be assumed to have even the basic security measures discussed in chapter 3; their networks should be considered as unsafe as the public internet. Systems using the IEEE 802.1X protocol can be used to require users to authenticate before being able to send and receive packets or even to run health checks of computers with the possibility of placing them into a quarantine network if problems are detected.

Even if there is no immediate plan to enforce network controls between these areas, it may be a good idea to take the possibility of doing so into account when allocating IP addresses or installing network cabling.

4.2.3 Designing Access Controls

One common approach to designing access controls is to exclude services that are known to have problems. However, as noted above, such controls can easily be avoided by running services on non-standard ports and may well be powerless to prevent attacks on newly discovered vulnerabilities. Such an approach implicitly assumes that 'everything else' is safe and condemns all system administrators to an endless round of ensuring that this is in fact the case.

A much better approach, recommended in all the literature and by reports from within the educational community, is to design controls to permit only traffic that is reasonably safe. Safety may be assumed either because the service itself has no known security problems, or because the servers to which traffic is allowed have adequate and well-maintained computer security. New hosts and services are then protected by the access controls until an informed decision is made that they are inherently safe. This is much better than the alternative of assuming everything is safe until it is proved to be otherwise, usually by a destructive attack.

It is therefore recommended that any new access controls be designed with a view to 'what can be safely allowed', rather than 'what needs to be prevented'. This approach is usually known as 'default deny'. Where an existing system has a 'default permit' policy, transition can be achieved by first observing all the existing traffic flows and setting up 'allow' rules for those which are legitimate; once all legitimate traffic is covered by these rules, the default can be

changed from 'allow' to 'deny'.

4.2.4 Implementing Access Controls

This section contains some general observations on the implementation of access controls. These should be read in conjunction with Computer security [3], which deals with individual services.

Filters that are based on the Internet addresses of local computers are in general more reliable than those which depend on port numbers. As noted above, a service can easily be provided on a port number other than its standard one and this will change the effect of a port number filter. Changing the Internet address of a computer to avoid filtering is considerably harder and may well prevent traffic being routed to the computer. Filters can be based on individual IP addresses or on address ranges; the latter are usually expressed as sub-network numbers with the host portion being treated as 'don't care'.

When the Internet Protocols were designed, a distinction was made between port numbers less than 1024, known as service ports, and those greater than 1024, known as user ports. All well-known services were intended to use service ports and on many operating systems only processes running with enhanced privilege were allowed to offer services in this port range. All other processes, for example clients making connections to services, were intended to use the remaining port numbers from 1024 to 65535. Both of these intentions have now broken down. Servers commonly have well-known port numbers above 1024 (for example web caches can be found at 3128 or 8080) and some operating systems will allow client programs to use ports in the service range. This makes designing filtering rules more difficult, since it is no longer possible to distinguish between clients and servers, privileged and unprivileged, based on the port number. Any attempt to block a service in the user range may occasionally prevent connections by clients, while clients running in the service range may be able to gain unintended access through filtering routers. The official Internet list of well-known port numbers is maintained by IANA [4]. Most computers will have a subset of these definitions in a local 'services' file, which may also include local or manufacturer extensions to the official list.

Another problem with filtering using port numbers is that an increasing number of services never run at fixed port numbers, but instead choose their service port number(s) dynamically and advertise them in a local or remote directory service. The chosen port numbers may change during operation or when the service or system is rebooted. To contact the service, the client first obtains the current port number from the directory and then makes the appropriate TCP or UDP connection request. Since the port at which the service will appear is not known in advance, it cannot be either blocked or permitted using a filtering rule based on port number. It may be possible to control access to the directory service (if it is running locally) using its well-known port number, but this control can be avoided by an attacker who simply tries each port in the user range until he finds the service. In practice, services tend to start within a small range of ports so this search can be much less time-consuming than might be expected. Any computers that offer such services should therefore have especially good computer security measures.

Many server programs can be configured to reject or accept connections based on the calling IP address - for example to reject connections that come from outside the local network. These facilities should be used as a supplement to network-based controls. For servers that do not provide this function it may be possible to use a preliminary wrapper program or a host-

based firewall to check the origin of each request to the service, decide whether or not to establish the connection to the real server program and, optionally, log a record of this decision. As discussed in the [Computer Security](#) [3] section, host-based firewalls are particularly useful on clients and workstations where services are not normally provided but may appear as part of other software.

4.3 Router Security

Routers and firewalls are a critical part of an organisation's network infrastructure and are therefore an obvious target for attackers wishing to cause disruption. Whether routers and firewalls are used to implement filtering or just to transfer IP traffic, it is imperative to ensure that their security is not compromised. If an intruder manages to obtain control of a router or firewall then he or she will be able to remove (temporarily or permanently) any rules used to protect the network, to read or re-direct any traffic passing through the device or simply to create havoc by breaking the organisation's local and wide area connections. Even being able to read the device configuration or logs could provide valuable information for attacks on the router or other devices on the network.

Most network devices can now be managed remotely across the network, using protocols such as SNMP (Simple Network Management Protocol), HTTP (HyperText Transfer Protocol) and SSH (Secure Shell). These interfaces may allow the configuration to be both read and altered, as well as providing access to logs. Hostile attackers may try to access the management function, just as legitimate administrators would. Remote management functions must therefore be protected, ideally using network controls, encryption and authentication. Clearly, if remote management through a particular protocol is not required, the service providing it should be disabled and the associated network port(s) blocked.

If a router or firewall can be managed over a network connection then, whenever possible, it should be set only to accept connections to the management function when they come from the known IP addresses of network management workstations. These should normally be addresses within the organisation, so rules on the perimeter router/firewall should prevent them being forged from external sources. If off-site access is required then technologies such as Virtual Private Networks (VPNs) should be considered to ensure that only legitimate administrators can access it. It may be simplest to place all network management stations and device management interfaces in a single (virtual) LAN segment with very limited access to it from the rest of the internal and external networks.

The traffic between the management workstations and the router/firewall is likely to include sensitive information such as configuration details, commands, logs and passwords. It is important to protect this traffic against both reading and replaying by unauthorised people who may have access to the network. Each of the protocols commonly used for management interfaces offers the possibility for encryption - SSH in place of *telnet*, SNMP version 3, SSL for web interfaces - and these should be used whenever possible. If an encrypted version is not supported it may be safer to disable access using the unencrypted protocol or at least limit it to read-only management functions.

It is essential that the ability to login and hence configure the device is protected by at least a password. Some routers have two levels of privilege, each protected by a separate password: the lower level giving access to 'read only' functions, the higher to 'read/write' functions. Both

levels must be adequately protected via passwords and these passwords should be chosen and managed with at least as much care as any other passwords to privileged accounts on other IT facilities. Where an organization has more than a small number of network devices it makes administrative sense to control access to network devices through technologies such as TACAS+ or RADIUS where authentication can be centrally controlled. Such central authentication technologies must, of course, be managed securely themselves. It is also best practise to consider logging all administrative authentications on network devices.

SNMP's authentication mechanism uses 'community strings', but these function as passwords, so should only be known to the controlled device and those who are authorised to control it. Different community strings may be used for different groups of management functions. Finally, many network devices are delivered with default passwords or community strings. These should always be changed before a new device is installed and should be managed in the same way as any other privileged password.

Remember also that the workstations used by network managers are a critical point in the organisation's defences. If an intruder can compromise one of these systems or an account on it then they are likely to be able to gain access to organisation's network infrastructure. They, too, need to be configured and used in accordance with the best available security practice.

4.4 Protecting Bootstrap Protocols

Finally, there are a number of key network protocols that can cause security problems if they are accidentally or deliberately misused. Many of these are used by workstations or mobile devices when they join a network, to find out what address(es) they should use, how packets should be routed and where domain names can be resolved. If computers receive incorrect information at this stage then they will not be able to use the network as intended; they may also disrupt other traffic if, for example, two devices try to use the same IP address. If computers receive malicious information then this may make it easy for intruders to read their traffic or direct them to other malicious services.

DHCP is the protocol most commonly used to provide this network information; on IPv6 networks router advertisements are used first to discover the network connectivity. In each case there will be authoritative servers for these protocols and, where possible, network equipment should be configured to only allow these protocols to flow to and from those authoritative servers. Where this is not possible, network monitoring can be used to detect rogue servers offering network configuration information and the unexpected traffic flows that may result.

Source URL: <https://community-stg.jisc.ac.uk/library/janet-services-documentation/network-security>

Links

[1] <http://www.team-cymru.org/Services/Bogons/>

[2] <https://community.ja.net/library/advisory-services/firewall-implementation-janet-connected-organisations>

[3] <http://community.ja.net/library/janet-services-documentation/host-security>

[4] <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml>