



Security Matters

Andrew Cormack

Technical Guide

UKERNA Technical Guides

UKERNA Technical Guides are one of a series of user guides available to JANET customers. Technical Guides contain detailed technical information and are intended for technical support staff at JANET sites, network specialists or those with a particular interest in the specialist area.

If you have any queries or comments about the Guide or would like to obtain copies, please contact:

JANET Customer Service

UKERNA

Atlas Centre, Chilton, Didcot

Oxfordshire, OX11 0QS

Tel: 0870 850 2212

Fax: 0870 850 2213

E-mail: service@janet.ac.uk

Further details of the documents in this series are available at:

<http://www.ja.net/documents/>

Contents

Introduction	5
1 Security Issues Raised by Connection to JANET	7
1.1 Implications of Connecting Using IP Protocols	7
1.2 Implications of Connecting Workstations	8
1.3 Implications of Connecting to the Global Internet	8
2 Methods of Countering the Threats	11
3 Host Security	13
3.1 Controlling Access to the Host	13
3.1.1 IP Access	13
3.1.2 Registration in the DNS	13
3.2 Restriction on the Services Provided	14
3.2.1 Trusted Hosts	14
3.2.2 FTP – Public Access to Files	15
3.2.3 World Wide Web	15
3.2.4 Electronic Mail Services	16
3.2.5 Domain Name Service	16
3.2.6 Remote Login	17
3.3 Security Measures on all Systems	18
4 Network Security	19
4.1 Use of Filters Within Routers for Network Access Control	19
4.1.1 Benefits of Network Filtering	19
4.1.2 Limitations of Network Filtering	20
4.1.3 Recommendations	21
4.2 Designing Network Security Measures	21
4.2.1 The Importance of Policy	21
4.2.2 Partitioning the Network	22
4.2.3 Designing Access Controls	22
4.2.4 Implementing Access Controls	23
4.3 Router Security	24
5 Security Measures Within the JANET Backbone	27
5.1 Policy Filtering in the Backbone Routers	27
5.2 Operational Filtering in the Backbone Routers	27
5.3 Source Level Routing	27
6 Requirements of an Organisation Connecting to JANET	29
6.1 JISC Policies	29
6.2 Legal Requirements	29
6.3 Provision of Filtering Information	29

7 Related Information	31
7.1 The Threat and How to Protect Your Organisation	31
7.2 Host Security Measures	31
7.3 Network Security Measures	31
Glossary	33

Introduction

The security of computers, data and networks should be a matter of importance to everyone who uses them. Isolated individual computers are relatively secure provided their physical well-being is ensured; regular backups should be sufficient to ensure the integrity of the data they hold. However, once a computer is connected to a network, whether local or wide area, there are many other concerns that must be addressed. Connecting to a network provides many advantages – for example, it facilitates information sharing – but it can also expose the computers that connect to it to threats against their proper operation and the safety and privacy of the data they hold.

Any organisation connecting its computers to a network should take measures to protect them, and their users, against attack. The particular measures required will vary according to local conditions and which services the organisation wishes to access and offer across the network. On a local network such measures represent good practice; when connecting to a large network such as JANET they may be essential. JANET and the worldwide Internet contain a huge and diverse community of users. Both the opportunities and threats presented by this community need to be respected.

This guide discusses some of the more important issues that arise as soon as an organisation connects to JANET. The topic of security is complex and this Guide cannot be comprehensive; consequently it identifies the main threats and points to other published material that contains detailed information on implementing countermeasures. The security of individual hosts is discussed as well as the security measures that can be taken at the network level, using routers or firewalls. These techniques have complementary strengths and weaknesses and any effective security implementation is likely to involve a combination of different approaches. We hope that this guide will provide the information needed by service managers and administrators to improve the safety of their own section of the Internet.

1 Security Issues Raised by Connection to JANET

1.1 Implications of Connecting Using IP Protocols

JANET is an IP (Internet Protocol) network. This means that most communication across the network uses the Internet Protocol and the application protocols built on top of it. These protocols were designed for a network where every host was trustworthy, an assumption that is almost certainly no longer true on the global Internet. Some of the authentication methods used by hosts offering IP services may therefore be insufficient on their own for use on an untrusted network.

Not all IP services use the traditional authentication method of providing a login (or other) password. Indeed for services that are supplied by one computer to another, rather than by a computer to a human, such authentication is rather hard to imagine. Instead, many services have the concept of a 'trusted host', with authentication of the host based on either the supply of an address or (less commonly) a name. In the case of an address, this will be taken from the calling or 'source address' field of the incoming IP datagram and checked to see whether it is an address trusted by the receiving system. The address may also be checked against the NDS (Domain Name Service), to see whether the address has been registered with a DNS nameserver. If a host name is used to establish trust, the name will be supplied as a field in the application protocol, and will be checked against an internal table of trusted names.

Both these methods of authentication are vulnerable to careless configuration of the relevant tables and mechanisms on the recipient system (and especially to an intruder who manages to gain write access to the configuration files); and to the faking of the source address in the IP datagram and/or calling name in the application packet by an intruder wishing to impersonate a trusted system. The faking of addresses is relatively difficult when success requires that returned datagrams arrive at the impersonator's system rather than the impersonated system. The faking of names is relatively easy. However, there are a number of well-documented ways of faking the source address using the characteristics of the IP and/or TCP (Transmission Control Protocols); see 5.3 for a discussion of one of these. Faking of addresses, usually known as 'spoofing', is commonly used to obscure the origin of a denial of service attack. Some attacks also use a large number of packets from faked addresses to hide the one genuine packet which actually effects the security breach.

Implementations of the IP protocol suite normally include many services that rely on trusted hosts for some degree of authentication. Another technique used in UNIX® systems is the concept of a privileged TCP or UDP (User Datagram Protocol) port. The port is the addressing mechanism used in both TCP and UDP. It identifies the process on a system addressed by the IP address to which the TCP segment or UDP datagram should be delivered. Therefore the destination port usually identifies the service being called. The source port may also identify the calling process. Certain ports are deemed to be privileged, in that only a process having super-user privilege may use them. Some services rely upon this mechanism to determine whether the calling process is legitimate, usually in conjunction with one of the authentication methods described above. A request coming from a privileged port on a trusted host is assumed to have increased authority. As mentioned above, the reliance on host addresses is already suspect. Unfortunately the privileged port assumption is equally flawed. Even within a homogeneous UNIX® network it must be assumed that the users and owners of single-user workstations can gain super-user privilege (see section 1.2). Where other operating systems are present on the network these may not even have a concept of privileged ports, so any user can generate a request from such a port number.

A number of software systems that provide stronger authentication and authorisation techniques have been developed. However, the large-scale implementation of any of these

across a wide area network still causes problems. Furthermore, most use cryptographic techniques so may have restrictions on international use. It is therefore necessary to accept that systems connected to JANET will have less than perfect authentication mechanisms for some time.

1.2 Implications of Connecting Workstations

Most of the user computers connected to JANET are now workstations, with processing power far greater than the servers of only a few years ago. The distinction between a 'terminal' and a 'server' is now extremely unclear in both hardware and operating systems regardless of whether they are running Windows® or UNIX®. Most desktop workstations can and do offer services to their local area network, whether their users are aware of this or not. File sharing and web serving are the most common services offered, but are by no means the only ones. Once the local area network is connected to JANET, these mini-servers may pose a large security risk.

A fundamental weakness of most workstation operating systems is that they provide an insufficient number of levels of privilege. Often there are essentially only two: user privilege and administrator (super-user) privilege. This means that any server process must either run as a normal user process or as a super-user process. The former is likely to require relaxation of security to some degree to allow such an unprivileged process to operate. This in turn means that other user processes may have an unwarranted level of access to parts of the system. Conversely a server process running as super-user will often have more privilege than is required to perform its function. In theory this is not necessarily a problem, but in practice very many successful attacks have been based upon this weakness. It has proved to be very difficult to write server programs that are secure against attack by an intruder and, once breached, their super-user privilege means that they can be perverted to perform almost any action on behalf of the intruder.

The problem of super-user status is also exacerbated by the fact that each user/owner of a single-user workstation may potentially become a super-user. Much of the security philosophy of common network operating systems is still based on the concept that service requests from a remote system that purport to originate from a process having super-user privilege are to be taken at face value. Although the situation is improving slowly there is still a significant risk that the super-user of a remote system may not be a trustworthy colleague. By compromising one system, an attacker may gain the trust of many others.

A further different issue is the problem of configuration. Most vendors of operating systems supply them configured in such a way as to be almost completely 'open' on an IP network in terms of access to services and resources. The argument used to support this policy is that customers wish to have their systems delivered in this way, and can tighten up security to the level they require. The problem with this is that many user/owners of single-user workstations either do not appreciate that this is the case, or do not care. From the point of view of ensuring a reasonable level of security, the policy of supplying systems relatively 'closed' to the network, which have to be opened by the customer, would be far preferable. (However it would also probably involve the vendors in an increased level of after-sales support in helping their customers do this!)

1.3 Implications of Connecting to the Global Internet

An obvious consequence of connecting an organisation's LAN (Local Area Network) to JANET is that its IP networks then become part of the global Internet. This has implications in terms of the enormous number of people and systems 'out there' which will immediately have access to systems within the institution. It is estimated that there are millions of Internet users in the UK and hundreds of millions worldwide. Both these figures are increasing rapidly. Furthermore, the user population of the Internet is considerably different to that of JANET – whereas the latter is largely made up of professionals and students under the (presumed) control of their institutions, the Internet has a user base that is rapidly becoming indistinguishable from the general public in its variety.

The number and variety of systems and people with access to the Internet increases the likelihood of attack. Any large community will have malicious members. These may choose to attack any other member of the same community. Motivations vary, and may not involve dislike, or even knowledge, of the organisation under attack. Random attacks may be performed to gain access to a particular service without caring where it is located; to obtain control of systems for use in future attacks such as distributed denial of service attacks; or merely to score 'points' for the number of systems broken into. As part of the Internet, JANET may also be a victim of any attacks against that network as a whole.

A further implication of connection to the Internet arises from the fact that IP is not a 'point-to-point' protocol involving a relatively small number of links between switches and end-systems that can usually be trusted to be in responsible hands. Any owner of a PC, or a single-user UNIX® workstation, potentially has access to an IP implementation that may be subverted. The nature of Internet connectivity means that such a system can be directly connected to another system without the need for the network administrator's involvement. This has implications for the sort of attack outlined in section 1.1 above. Although this is also a risk from IP systems situated within an organisation, the potentially large, less controllable Internet community does increase the statistical chance of attack in this way.

2 Methods of Countering the Threats

There are two possible technical approaches by which any organisation can address security issues:

- secure configuration of host systems, and encouragement of responsible working practices by their users (host security)
- implementation of security measures at routers, firewalls and other network control devices (network security).

The latter is likely to involve at least the devices that connect the organisation's network, or networks, to JANET and devices within the local area network can also play a valuable role. Both of these must also be supported by non-technical measures, such as policies and user education.

It is likely that most organisations will need to use a combination of both approaches. The ideal would be to ensure that every host was secure in every aspect of its use. Unfortunately the cost of achieving this is extremely high, both in terms of effort required and inconvenience to users. However it should be borne in mind that the effective application of host security also brings benefits for the local network, as only host security can be effective against internal attackers and external attackers who have managed to compromise an internal system as a bridgehead. Network security can be used to partition the network, giving additional protection to systems that are at high risk either because of their management or the value of the activities in which they are involved. It is common practice to separate administrative functions, such as payroll and examinations, from the general network. The same approach can also be useful elsewhere, both within and at the boundaries of the organisation's network. In particular, using appropriate controls at the network level can limit the spread of future incidents as well as allowing effort to be concentrated on securing the small number of hosts which, to carry out their function, have to be most exposed to possible attack.

Section 3 deals with some issues of host security. Some practical advice is given, although this has been kept brief as there are now a reasonable number of books and other material available to assist with the implementation of security on individual hosts. Details of where to obtain such material, much of which is in the public domain, are given in section 7. It cannot be emphasised too strongly how important it is that an organisation connected, or planning to connect, to JANET ensures that it regularly audits its own security in the light of the practical recommendations made in these or other documents.

Section 4, Network Security deals with security implemented via the organisation's network control devices, particularly the router or firewall that provides the organisation's connection to JANET. It covers the advantages and disadvantages of security measures that use network devices, and some practical considerations that may arise.

Finally, it should be remembered that, although this guide deals only with the technical aspects of the provision of security, there are a number of other considerations, both administrative and legal, which need to be taken into account. Indeed, there is little point in devising technical solutions to security threats if there is no organisational computer security policy to support the measures to be implemented. This policy should set out clearly the responsibilities and authorities of all users and providers of the service to protect the security of the organisation: the interconnected nature of an IP network means that insecure behaviour by one user can cause significant and widespread damage to others. In formulating a security policy, items that need to be addressed include establishing acceptable behaviour rules for members of the organisation, enabling members of the organisation to use computers in a secure way and defining procedures to be followed when violations of security are detected. Advice on these issues is available from JISC and UCISA to assist sites in implementing their own local policies. The Internet Engineering Task Force has also published a Site Security Handbook (RFC 1244) which gives guidelines on the establishment of a site security policy in the context of the Internet (but does not give

technical information on the implementation of specific security measures). References to these and other sources of information about security policies are given in section 7.

3 Host Security

There are two things relevant to this method of security that need to be considered:

- what level of access to JANET and hence the Internet should be permitted from and to the host;
- what security measures need to be in place on the host as a consequence of this decision.

3.1 Controlling Access to the Host

3.1.1 IP Access

It is possible to isolate a host completely from JANET by preventing IP datagrams from the external network from reaching the host, and/or datagrams emanating from the host reaching JANET.

One method of doing this is to configure a router between the host and JANET to block datagrams to and from the host's IP address. This is dealt with in section 4.1. The other is to configure the host so that it does not have a route to the JANET router. In this case datagrams from JANET will reach the host but reply datagrams and datagrams originated by the host will not reach JANET, as the host will not be able to 'find' the JANET router. Note that if this technique is adopted then it is necessary to ensure that the system is not running a routed process as this may pick up routes advertised by other routers and hence subvert the restriction. Also, any technique which only blocks return datagrams will not be effective against attacks, such as denial of service, where the return datagrams are irrelevant to the success of the attack.

A typical situation where this type of restriction might be applicable is a Windows® or NFS (Network File System) fileserver that only serves clients on a single, local, network (for example, a single department). If there is no route to the server system from JANET, and vice versa, there can be no possibility of unauthorised mounts of the file systems it provides. The restriction could still be circumvented if any of the clients re-exports the file system. Of course, configuring the fileserver in this way would only be possible if it does not provide other services that require access to and from JANET.

As an extension to this technique, it is possible to prevent access to whole sub-networks in the organisation. This has to be accomplished through a router, rather than a host, and is also dealt with in section 4.1.

3.1.2 Registration in the DNS

Even if a host system has IP access, it is possible to choose not to register its name-to-address mapping in the DNS. This means that the host is able to make calls to JANET but will have difficulty in receiving them. Although this makes it more difficult for an intruder to find the address of the host, there are now tools that will scan the whole of a Class B or C network using numeric addresses alone. These will find all the hosts on a network in an hour or so.

There are also a number of disadvantages of not registering in the DNS. For statistics gathering and network management there is often the need to convert addresses to names. If no name is available then confusion can arise. More importantly, services such as web, FTP and e-mail often use the DNS as a means of checking the bona fides of an incoming call; calls from a host that is not registered may be rejected.

Apart from these disadvantages, this technique also really falls into the category of 'security by secrecy' and as such must ultimately fail: therefore it is not recommended.

3.2 Restriction on the Services Provided

Even if a host needs to have full access to the Internet, it is straightforward to reduce the risk of attack from the network. Most systems, when delivered, are configured to offer all the IP-based services of which they are capable. This is unlikely to be appropriate to the system's actual use, particularly if it is only required to perform a single function such as a desktop workstation or mailhub. In almost all cases some or all services can be turned off, resulting in both a reduction in the security risk and an improvement in performance. Services can only be removed if there is no requirement to run them at all. If a service is required for internal use but is not to be offered to JANET, then network control devices should be used instead to restrict the flow of traffic.

For example, a single-user workstation is unlikely to need to offer incoming *telnet* and FTP services, nor to be running a listener for printservicing or for SMTP (Simple Mail Transfer Protocol) mail. All of these can be disabled with immediate benefit. There are also a number of services that constitute a different type of security risk because, without demanding prior authentication or authorisation, they provide potentially useful information to an intruder. These include *finger* and *rusers*. The need to provide these services should be considered carefully.

When setting up a system, its purpose should be clearly defined and only those services that are necessary for the purpose should be run. All others should be turned off. For incoming services the means by which this is done will vary depending on the service being inhibited. Many UNIX® services run under the *inetd* daemon. These are turned off by editing the file */etc/inetd.conf*. On Windows® systems the Services control panel provides a similar function. Other services run as stand-alone processes which must be removed from the system startup. For example, if the system is not acting as an NFS server then the *nfsd* and *mountd* daemons should not be run.

There are a number of common services that most organisations will want to provide to the outside network. Each of these services may be targeted by attackers, so extra care should be taken with any host that provides them. These services should usually only be enabled on hosts intended for the purpose and should not normally be run on interactive workstations. The following sections consider the specific security implications of each of them. As well as disabling or restricting the services on individual hosts, most of these services can and should be controlled by filtering at the site router to protect other hosts which may offer them by accident.

3.2.1 Trusted Hosts

As pointed out in section 1.1, trusting another host solely on the basis of its claimed identity may be dangerous. Even if the legitimate users of the trusted system are benevolent, an attacker who compromises the security of a trusted system immediately gains access to all those that trust it. Trust therefore places the security of the system in the hands of others. Moreover, an attacker may not even need to compromise the trusted host. For some attacks it may be sufficient to impersonate it.

In practice it will always be necessary to trust other systems for some services – for example when sharing or exporting file systems – so a balance between usability and security needs to be found. New trust relationships should be approached with care, always configured to trust the minimum number of other hosts and with the least impact if the trust is abused. Many systems are delivered with an extremely open configuration, effectively trusting the whole of the Internet. These should, of course, be turned off until the actual requirement for trust has been determined. Some of the more notorious excesses are associated with the */etc/hosts.equiv* and *~/.rhosts* files, which are used to allow access without a password to UNIX® systems, and with file sharing on both UNIX® and Windows®.

File sharing is not just of concern for privacy of information. If it is possible to write to a shared file system then this may be used to alter the security settings on the host machine.

Whenever possible, file systems should therefore be shared for reading only, and to a limited set of hosts. These options are set in the `/etc/exports` file on UNIX® and in the ‘file sharing’ dialogues on Windows®. Neither file sharing system is suited for use over public WANs (Wide Area Networks), as their protocols are insufficiently robust. Other methods for exchanging files over WANs are described in section 3.2.2.

There are many other services that are commonly offered to some other networked hosts, but not to the whole Internet. In each case consideration should be given to limiting the hosts that have access. Although restrictions based on IP address cannot provide perfect security, they are still much better than no restrictions at all. Some applications can be configured to reject calls from untrusted addresses: others can be ‘wrapped’ by installing small address checking programs in front of them. Whenever using this type of restriction it is essential to ensure that addresses that do not match the trusted list are indeed denied.

Host-based access controls will often duplicate the network-based controls described in section 4.1 and this is no bad thing. Network controls may be harder to subvert but only host controls can protect against malicious local users.

3.2.2 FTP – Public Access to Files

There is often a requirement to provide public access to files, for example an archive of software or documents. Allowing remote readers to copy selected files across the network provides this function, but it is important to ensure that it does not also risk the security of the system. Probably the least attractive approach is to provide all possible users with an account and password. Such an account is likely to have too much access to both read and write files, and once its password is known to many users security is inevitably compromised.

Most implementations of Internet FTP (File Transfer Protocol) provide a restricted, ‘anonymous’ service to address these problems. A user who logs in using a conventional username (usually ‘anonymous’ or ‘ftp’) has access to a limited part of the system file space, often with only read access. No password is required but it is common to ask the user to enter an e-mail address for tracking purposes. (NB this address cannot be identified.) Such a system is reasonably secure provided it is configured properly. The manufacturer’s documentation should provide the necessary information.

It is also possible to allow anonymous users to write files to an FTP server, but this requires particular care. It should never be possible for the same file space to be both written and read by anonymous users. Any system that allows this is likely to be taken over as a distribution point for pirated software or other illegal material. Uploaded files should always be checked for content, viruses, etc. by a trusted user before being placed in a public readable area. Public writable directories should at least have disk quotas, and possibly be on separate disk partitions to reduce the impact on the rest of the system when a malicious or misguided user tries to upload many gigabytes of files.

3.2.3 World Wide Web

The World Wide Web can also be thought of as a file sharing system, so has many of the same issues as FTP described above. To prevent leakage of private information, the web server should only be allowed to serve files from a limited area of the system file space. There have been problems in the past with the controls provided by server programs, so operating system controls should be used in addition. One control that should always be used is to run the server as an unprivileged user. This reduces the immediate consequences of a vulnerability in the server or its associated programs.

As well as simply serving files to clients, a web server may also run programs in response to requests from remote users. It is vital to realise that these programs are run on the server host under the control of unknown, possibly hostile, users. They are therefore as exposed to attack as the operating system itself and should be implemented, tested and reviewed with

appropriate care. Unfortunately a high proportion of successful attacks against web servers use vulnerabilities in these scripts.

The original web protocol included the ability to publish files to a server as well as read them. However this has never been widely implemented. Furthermore a web server which is able to create, modify and delete files needs to have some increased privilege, so is at increased risk if a vulnerability is found. Such a server should be simplified as much as possible to reduce the opportunities for attack. In particular it should not have any additional programs or scripts installed. In the absence of a standard method for publishing pages onto web servers, vendors have adopted a number of public and proprietary methods. Some of these use their own web server scripts but these have the same problems of server privilege as described above. It is better to separate the functions of writing and reading web pages by using either a proprietary publishing server or a standard FTP server. Whichever server is chosen should be configured in as secure a fashion as possible. It should be given additional protection by network devices and all log files should be monitored for signs of misuse.

Web servers are high profile targets so especial care should be taken in securing the hosts they run on. Additional information on securing the server programs themselves is available from the suppliers of the software.

3.2.4 Electronic Mail Services

Electronic mail is one of the most widely used services on the Internet and beyond but there are far more hosts offering mail services than are required or desirable. A host only needs to run a mail service if it is used to store delivery mailboxes. Since most workstations only process mail under the direct control of a user – for example when moving messages from an inbox to local folders – they do not need to run mail server software. Unfortunately most UNIX® systems are delivered with the ‘sendmail’ server installed and running. All too often this will be an old version with known security or configuration problems. The first task in securing e-mail is therefore to disable all these unnecessary services and install extra protection at the network level for those which will undoubtedly spring up in future. For those hosts that do need to provide a mail service, there are less powerful alternatives to sendmail which may be sufficient for many situations while being easier, and therefore less error-prone, to set up.

A common problem in setting up mail servers, both legitimate and not, is uncontrolled relaying. Most servers are intended to transfer mail between an organisation and the Internet, so for any valid message either the sender or the recipient, or both, should be a local system within the organisation. However some servers are prepared to accept messages from external sources for external destinations. This behaviour is known as relaying. While there are good reasons for allowing some servers to relay – to provide backup for another organisation or to support mobile users for example – the behaviour needs to be controlled by careful configuration. Open relays, which allow any combination of origin and destination, are frequently abused by advertisers and others to distribute bulk e-mail. This will usually overload the server, affecting its ability to handle legitimate mail, and often leaves the organisation with a flood of complaints and error messages to deal with. Sites that are frequently abused as relays may be added to blacklists used by many network operators and ISPs (Internet Service Providers) to reject all e-mail and other traffic. Advice on preventing relaying is available from the MAPS (Mail Abuse Prevention System) web site, referenced in section 7.2.

3.2.5 Domain Name Service

The DNS is another service that is provided by far more hosts than necessary. Each Internet domain, such as camford.ac.uk, needs one primary DNS server and a small number of secondary servers. If separate sub-domains are used, these also need primary and secondary servers but each server host can handle a number of domains. Control of a DNS server gives the ability to alter the Internet name of any host in the domains it serves, so the servers must

be kept secure. For both security and performance it is recommended that DNS servers be dedicated machines performing only this task.

Requests to DNS servers can be of two types. The most common are simple queries, requesting the translation of a single Internet name or address. These requests may be made by any host on the Internet. The other type, known as zone transfers, request all the available information about a domain and are used by secondary servers to obtain the information they will publish. The list of hosts from which any server should expect a zone transfer request is therefore both short and known in advance. It is recommended that DNS servers are configured to log and reject zone transfer requests from other hosts. Zone transfers can also provide useful information to intruders about the identity, function and type of hosts on the network.

3.2.6 Remote Login

A group of services are used by users to access their files and other resources. Common examples are *telnet* and *rlogin*, used to obtain interactive terminal sessions, and POP and IMAP (Post Office Protocol and Internet Message Access Protocol), used to access mailboxes. These services usually require a username and password to be sent over the network to authenticate the user. On internal networks, and especially on JANET and other untrusted external networks, there is a risk that these usernames and passwords may be read off the network by others. Most local area networks use Ethernet, a broadcast medium where it is possible to 'listen in' to communications between other hosts. When packets are sent across the Internet, the routes they take cannot be predicted. Sensitive information such as passwords is then very likely to pass through systems that are not controlled by either the sending or receiving organisation. So called 'password sniffing' on local networks is one of the most common techniques used by intruders to gain access to more machines. Information theft from wide area networks requires more equipment, but the rewards are correspondingly greater. A thief who steals a password can impersonate its owner perfectly, whether he chooses to send malicious e-mail, read files, or create web pages. Once the identity has been stolen such activities are very hard to control since, as far as the computer is aware, they are being performed by a legitimate, authorised user.

Surprisingly, in view of these risks, many passwords are still transmitted across networks in clear text. Some LAN operating systems such as Novell® NetWare® and Windows NT® can use encryption to protect their passwords: however, these still have plain text options for backwards compatibility, which can be selected by careless configuration or malicious activity. If at all possible, such options should be disabled on servers. For logins across the WAN the situation is even worse with the majority of logins using plain text passwords. These will, of course, travel across parts of the LAN at each end.

A number of software systems can be used to encrypt passwords and other sensitive information flowing across the network. Most of these make no distinction between passwords and the rest of the traffic, simply providing an encrypted 'tunnel' down which information can be sent without risk of it being read in transit. Netscape's Secure Sockets Layer (SSL, also known as TLS) is often used by e-commerce web servers to protect credit card information sent across the network and is equally useful to protect passwords when logging in to web servers or gateways. SSL has the advantage of being available transparently on most popular web browsers. It is also available on servers though these require a little more setting up. *Telnet* and *rlogin* can be replaced by the encrypted SSH (Secure Shell). This can also be configured to provide an encrypted channel for POP, IMAP and most other well-behaved TCP protocols. Finally, many commercial companies offer software or hardware to set up secure VPNs (Virtual Private Networks) across a public wide area network. Anyone logging in across a WAN, or a LAN with a mixed user population, should consider using one of these options.

The X-Windows system is a particular case of remote login, since useful information can be obtained by monitoring the network traffic, but it also has additional security problems for the user machine. As with many other systems, X-Windows can be used in a secure fashion, but is often supplied with an insecure configuration by default. An X-Windows

session involves an application program (confusingly known as the 'client') and a display device, which is usually a workstation. It is possible for more than one application, running on different hosts, to use the same display device at the same time. If the device allows this, one application may read keystrokes, and write into windows which are intended for others. Applications to monitor the typing of usernames and passwords are available among the intruder community. To prevent this, X-Windows provides the '*xhost*' command to limit the hosts from which applications are allowed to connect. Workstations should be configured with the least permissive *xhost* options and users educated to permit new hosts only when required. Once an application has been started, it will not be affected by removing its host from the accepted list. If the host on which the application is run is a multi-user machine, restrictions based only on host will not prevent other malicious users from starting applications directed at the same display. This can be prevented by setting the system to exchange tokens between the display device and the application host to control which applications can be used. For the best security, the X-Windows traffic can be encrypted using SSH.

Theft is not the only way to obtain passwords. If passwords are not well used then simple guessing or social engineering (for example phoning a user, claiming to be a system administrator and asking for the password) can be very effective. The only solution to these problems is to educate the users in choosing good passwords and using them carefully. The scope for intrusion can also be reduced by limiting the number of services against which guessing attacks can be tried and ensuring that repeated login failures cause an account to be disabled, though some servers may not provide this function. There is no technical solution to users who leave their terminals logged in.

3.3 Security Measures on All Systems

Documents are now available for most operating systems suggesting good practice for running them in a networked environment. A selection of these is included in the references in section 7.2 of this guide. Mailing lists can also be useful sources of up-to-date information. It is particularly recommended that those responsible for systems subscribe to the alert lists maintained by their manufacturers. Some of the checklists for maintaining secure systems have been implemented as scripts or programs. These can be used to improve the default security level of a system – removing obviously insecure permissions and poor configurations, for example – but they do not remove the need to check and update systems as new problems are found. If a script of this kind is available for your system from a trusted source, then it may be worth using it as a starting point.

All operating systems software has bugs in it, and a few of those bugs will have implications for the security of the system. Exploiting bugs in network services is the most common form of attack by intruders. The easiest way to reduce the effect of these bugs is not to run unnecessary services – if a service is not running then its bugs do not matter. For those services and operating system components that must be used and exposed to attack, it is essential to ensure that their configurations are secure and that they are updated to deal with any new security problems that may be discovered. Manufacturers usually provide patches to correct bugs. An informed decision should be made as to whether to install each new patch. Patches will often restore configuration settings to their defaults so it is important to check after applying the patch that it has not reduced the security of the service. Every system visible from JANET should be actively maintained.

Finally, it should be remembered that even though an action is prohibited by the JANET AUP (Acceptable Use Policy), this does not guarantee that no-one will attempt to perform it! The JANET AUP is a document setting out policy, not a series of enforced controls on the JANET backbone (see section 5.1).

4 Network Security

4.1 Use of Filters Within Routers for Network Access Control

Although the main purpose of a router is to transfer IP datagrams from one network segment to another, most routers can also make decisions as to which datagrams 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 datagrams which will be transmitted, or a set which will be blocked. Rules may be based on a number of different aspects of the datagram, for example the source or destination host IP address (or network number), or a specific source or destination port, or flags within the UDP datagram or TCP segment which the IP datagram 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 hosts, since the listener for a particular service will usually be found at the same 'well-known' port number. For example, the *telnet* service is deemed by convention to operate on port 23 – therefore any *telnet* listener process expecting to accept an incoming call will do so via TCP port 23 and any *telnet* initiator will specify TCP port 23 when making its call. A router which discards IP datagrams with a TCP destination port of 23 will therefore block normal *telnet* calls 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 *telnet* 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.

These 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 hosts 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 systems 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 machines that do not need to be visible from the external network can be blocked at the router, immediately reducing the risk of external attack against those machines. The effort required to enforce good security on individual hosts can then be concentrated, at least in the first instance, on the small subset of hosts and services which are provided to the outside world.

Network filtering provides protection against external attack for host 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. As mentioned above, services such as filesharing may use only weak authentication methods so should probably not be made available to 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 new bug in an existing service to gain privileged access to the host system. A network filter can make it harder for attackers to perform these new 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 host. 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.

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, host 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 hosts 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 datagram 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 filtering can control which hosts 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 a firewall 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.

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 host 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 hosts within the institution and services provided outside the institution 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 machines. This requires planning in the initial allocation of addresses – machines with similar functions and security requirements should be grouped into address blocks – and can also penalise the users of properly configured, ‘safe’ hosts. So long as a sub-network contains some ‘unsafe’ hosts, the filtering rules will be constrained by the need to protect those machines and the others 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. Good host 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 host 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 hosts 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 hosts connected to the network. It is most effective in increasing overall security on diverse networks when used to protect and isolate hosts 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 hosts. Filtering which relies instead on blocking services and hosts which are known to be insecure is easy to subvert 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 host security. Hosts which are visible to the Internet will need to be secured individually in any case, and even internal hosts 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 host security can protect against attacks either from, or routed through, other hosts on the same network.
- Any security measures, whether implemented by host 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 hosts. 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 subvert the controls.

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 (but see for example the JTAP (JISC Technology Applications Programme) report referenced in section 7.3) as they will vary from one organisation to another, but there are some common themes which should apply to all security projects.

4.2.1 The Importance of Policy

JANET and the institutional networks connected to it have traditionally been very open, since this was thought to be necessary for academic work. This assumption has been questioned by experience. The openness of the networks has made them very easy targets for attackers who commonly use compromised academic hosts to attack other, often commercial, sites. To reduce the effort involved in tracking down such incidents, many institutions have decided that it is worth reducing the openness of their networks and, in particular, their connections to JANET and the Internet. These sites have found that, provided users and management are informed of the change and the need for it, this has not been detrimental to their normal work. Furthermore the number of security incidents has been greatly reduced. Even mis-users within the institution seem to be discouraged by the change in the security culture. It is, of course, necessary to respond promptly and fairly to those cases where an individual has genuine problems with the new arrangements: however it has usually been possible to find acceptable ways of meeting all legitimate requirements without reducing the overall security level.

With the expansion of JANET to include connections to non-academic organisations, and with the benefit of the experiences of those institutions that have already reduced their 'openness', it is clear that the adoption of a security policy is vital. This policy should be reasoned and must have the support of the management of the institution. Such a policy makes it clear to most users that the changes are for their benefit. There is little point in having the freedom to store or transmit any information on a network if the lack of security means that information may be unavailable or unreliable to anyone else. The policy also avoids the need to continually make exceptions for individuals, although there should always be a review process to allow considered changes when the policy is found to be inappropriate. Finally, a clear policy makes the implementation of security measures much simpler since it separates the policy and technical decisions. Implementation then becomes a case of translating the written policy into correct technical measures.

References to information about security policies designed for educational networks can be found in section 7.1.

4.2.2 Partitioning the Network

The traditional arrangement for a network security system is 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 across 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 privileges is clearly not a complete solution.

Most organisations have already begun to partition their internal networks by 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. 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. 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 more preferable approach, recommended in all the literature and by recent 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 host 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. A very good description of this process can be found in the JTAP report referenced in section 7.3 of this guide.

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 section 3.2, which deals with individual services.

Filters that are based on the Internet addresses of local hosts 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 host to avoid filtering is considerably harder and may well prevent traffic being routed to the host. Filters can be based on individual host addresses or on address ranges; the latter are usually expressed as sub-network numbers with the host portion being treated as 'don't care'.

The official Internet list of well-known port numbers is currently contained in RFC 1700 (Assigned Numbers). Most hosts will have a subset of these definitions in a local 'services' file, which may also include local or manufacturer extensions to the official list.

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.

A particular problem is caused by the RPC (Remote Procedure Call) mechanism, which is used by the NFS and other applications. These services use varying port numbers within the user range, and their location may change whenever the service is restarted. To use such a service the client first calls the portmapper service, which has the well-known port number 111. The portmapper will then return information, including the UDP or TCP port number,

about the requested RPC service. The client can then make a normal connection to the service. 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 is possible to control access to the portmapper service 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. Some routers and firewalls can control traffic based on the RPC service it is requesting, but in the absence of such a system there is no satisfactory way of controlling RPC services across a network. Any hosts that offer such services should therefore have especially good host security measures.

Many server programs can be configured to reject or accept connections based on the calling IP address. These facilities should be used as a supplement to network-based controls. It may also be possible to use a preliminary program to add this option to some services which do not provide it themselves (typically programs which are invoked once for each connection, for example using the UNIX® *inetd* daemon). Such preliminary programs are known as wrappers. Their sole function is to check the origin of each call, optionally log it and then, if the call is permitted, pass the connection on to the server program. The TCP wrappers program is now installed on many UNIX® systems. The program and a version of portmapper with similar functionality can also be downloaded from the Internet.

4.3 Router Security

Even if routers are being used 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.

Most network devices can now be managed remotely across the network. This is normally done using the *telnet* protocol but others may also be used. Whatever protocol is used, it is essential that the ability to login and hence configure the device is protected by at least a password. Some routers (e.g. Cisco®) 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.

If a router can be managed over a network connection, then it is possible for hostile attackers to try to access the management function, just as legitimate administrators would. Such attacks may come from outside or inside the organisation. Wherever possible, the router should be set to refuse connection requests which do not come from pre-configured IP addresses. These should normally be addresses within the organisation, since other filtering rules should already restrict the ability to forge them. If there is a requirement to manage network devices from outside the network then additional security measures such as encryption (described below) should be considered.

Many routers and other devices also allow management to be carried out using other protocols. Most of these allow aspects of the devices' configuration and logging to be read remotely. Some also allow management parameters to be set. The most common protocols used are SNMP (Simple Network Management Protocol) and the World Wide Web protocol, HTTP (HyperText Transfer Protocol).

SNMP has an authentication mechanism that uses 'community strings'. In effect these are 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. Anyone who can learn or guess a community string can gain access to those management functions on the device. Many network devices are delivered with default community strings. These should always be changed when a new device is installed and should be managed in the same way as any other privileged password. Each SNMP request

sent over the network includes the community string as authorisation. Although newer versions of the protocol make it possible to encrypt the community string, this is not yet widely supported so there is a risk that the community string will be intercepted. Unless encrypted community strings can be used, it is recommended that network devices be configured to ensure that SNMP can only be used to read information and not to update it.

Similar considerations apply to web-based interfaces to network devices. Like SNMP these transmit unencrypted passwords to authorise changes to the configuration of the device. Web browsers use SSL to provide encrypted communications, but unless this is also supported by the network device it cannot be used.

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. The systems used to manage the routers should therefore be designed to ensure the best possible security. The protocols through which a device can be managed should be known and controlled so that, if possible, management requests will only be accepted from fixed IP addresses. Protocols that will not be used should be disabled. Furthermore, if management commands are to be sent across untrusted networks (which may well include the organisation's own LAN) then any systems available should be used to prevent the communication being intercepted. Each of the common protocols has the possibility of encryption – SSH in place of *telnet*, SNMP version 3, SSL for web interfaces – and these should be chosen whenever possible.

5 Security Measures Within the JANET Backbone

The JANET network itself has a number of routers both within the network and at the connection points to other networks. There is therefore the possibility of implementing some network filtering within the backbone network itself. Such filtering would affect all sites connected to the network. This section discusses the security measures that have been considered for implementation at this level, and the decisions that have been reached.

5.1 Policy Filtering in the Backbone Routers

As noted in the previous section, any filtering decisions that were made by the backbone routers would inevitably apply to all sites connected to the network. Any decisions to filter would therefore need the agreement of the JANET community. It is unlikely that such agreement could be reached for any service, and the benefits would be extremely limited. In particular, filters applied at the backbone routers would not, in most cases, protect organisations from each other. Sites would therefore need to duplicate the filtering at their own routers to protect them from hostile users elsewhere on the JANET network.

The JANET AUP prohibits certain types of use of the network: however, these prohibitions are not of the kind that can be enforced by network filtering. In any case, the AUP has exceptions for activities performed in the cause of legitimate academic research so it is hard to make a definitive statement that any particular traffic flow is contrary to the AUP.

For these reasons, it has been decided that it is not appropriate to use filtering in the backbone routers as a measure to implement security policy for sites connected to the JANET network. Filters may be used for operational reasons, as described in the next section.

5.2 Operational Filtering in the Backbone Routers

Although filtering for policy reasons is not appropriate, there are good operational reasons to implement filters in the JANET backbone routers. The most obvious of these is to protect the routers themselves against attack. Filters may also be implemented if they are an effective way to resolve operational problems, for example if a Denial of Service attack is in progress, either directed against hosts connected to the JANET network, or using those hosts to attack a third party. These filters will usually be temporary and may be installed at the request of sites, the JANET CERT (Computer Emergency Response Team) or the JANET NOSC (Network Operations and Service Centre).

5.3 Source Level Routing

Source Level Routing (SLR) is an IP routing technique that allows the originator of an IP datagram to specify the precise route that delivery should take through the network (as opposed to letting the routers on the network determine the route based on their routing tables). It is implemented by placing in each IP datagram a list of addresses of routers to be traversed to get to the destination. SLR is intended for use in circumstances where there are exceptional routing arrangements that may not be known to intermediate routers.

At present there seems to be no legitimate need for SLR to be used across the JANET backbone. There is, however, a possibility that the technique might be used in an attack to assist in impersonating a trusted host. The JANET backbone routers have therefore been configured to reject any packets that attempt to use SLR. This decision will be reviewed in future if it is believed that SLR could be of significant operational use within JANET. Sites that are concerned about the possible hostile use of SLR should also configure their own routers to reject such packets, since a packet will only be recognised by the routers named in its delivery route. If an SLR packet passes through a JANET router on the way between two designated routers then it will not be recognised and will be discarded.

6 Requirements of an Organisation Connecting to JANET

6.1 JISC Policies

Any organisation connecting to JANET is required to abide by three policies set by JISC to support the role of the network in enabling education and research. These are the Acceptable Use Policy, the Connections Policy and the Security Policy. Each of them places responsibilities on the connected organisation (primary, sponsored or proxy) to ensure that its use of JANET does not harm the community with which the network is shared. In extreme cases, persistent failure to respect these policies can lead to a decision by JISC to suspend an organisation's connection to the network. Compliance with the policies may require some technical measures, but above all the policies require that each connected organisation has systems in place to encourage proper use of the network, to identify those responsible for any misuse, and to deal with misuse under appropriate disciplinary procedures.

6.2 Legal Requirements

There are also legal requirements on organisations' use of computers and networks, arising in particular out of the Data Protection Act 1998, the Human Rights Act 1998 and the Regulation of Investigatory Powers Act 2000. Other Acts also apply to individuals within organisations. Organisations that do not discourage harmful activity by their users, or do not address situations that may facilitate it, may be liable for civil damages for any harm caused to others.

6.3 Provision of Filtering Information

Sites are required to assist in the operation of the network and the investigation of any problems. In particular it is a requirement that all organisations provide, if requested, information to the JANET NOSC about any filtering that is used to block access from JANET to particular hosts, networks or services. This will help to prevent waste of network engineers' time in investigating apparent 'faults' which are in fact the result of filtering decisions.

7 Related Information

A great deal has been written about security on the Internet. This section lists only a few sources which have been found useful in designing and implementing security measures on hosts connected to JANET. Most of these references are to information freely available in electronic form. This has the advantage that such information can be kept up to date, but the disadvantage that it can prove less permanent than printed material. We hope that the sites referenced here will prove to be durable.

7.1 The Threat and How to Protect Your Organisation

A good background article on the risks of connecting computers to networks was written by staff of the CERT-CC® (CERT Co-ordination Center) for an encyclopedia of telecommunications. The article covers the reasons why computer insecurity is a problem, the types of threats and what can be done to protect against them. The article is available at:

http://www.cert.org/encyc_article/tocencyc.html

To preserve the JANET network for its intended use in supporting education and research, all sites connecting to the JANET network are required to abide by the JANET Acceptable Use Policy, the JANET Security Policy and the JANET Connections Policy. These policies are set JISC, which funds the network. The current versions of the policies are available from the JANET web site at:

http://www.ja.net/documents/policy_documents.html

Local security policies are essential to support any technical measures that may be put in place. Without a security policy, staff have little authority to implement technical measures or even to respond to security incidents. UCISA (the Universities and Colleges Information Systems Association) developed a set of model regulations to guide organisations in writing their own local rules for the use of computers and networks. These are available at:

<http://www.ucisa.ac.uk/resources/docs/library/modelreg.htm>

A paper on Developing an Information Security Policy, written for the education sector, is available from JISC at:

http://www.jisc.ac.uk/index.cfm?name=jcas_papers_security

There is also a British and International Standard, BS7799/ISO17799, for Information Security Management. A JISC pilot study in 1999 found that the Standard was a good reference when drawing up a policy for an organisation but that formal certification was unlikely to be worth the cost unless it was required by other partners. JISC and UCISA are doing further work to develop an Information Security Policy Starter Kit, based on the 2000 version of the Standard, which will be available from:

<http://www.ucisa.ac.uk/resources/>

A site security handbook containing recommendations for all sites connected to the Internet was published as RFC1244. This is available from:

<http://www.ietf.org/rfc/rfc1244.txt>

7.2 Host Security Measures

Information on securing host computers is available from most vendors of operating system software. For example Microsoft® and Sun®/Windo have a wide range of security guides and checklists for their products at:

<http://www.microsoft.com/technet/security/default.msp>

<http://www.sun.com/bigadmin/collections/security.html>

Information on securing e-mail servers is available from the MAPS Transport Security Initiative at:

http://www.mail-abuse.com/support/an_sec3rdparty.html

Detailed information on specific tools is available in the Security Improvement Modules produced by the CERT-CC® at Carnegie Mellon University. These include a number of public domain tools that can be used to check and monitor systems as well as the facilities built into the systems themselves. The modules cover Windows® and Solaris™, though the techniques used can often be applied to other flavours of UNIX®. These modules are available from:

<http://www.cert.org/security-improvement/>

CERT-CC® also publishes advisories, incident notes and Tech Tips on items of current interest on its web site at:

<http://www.cert.org/>

The TCP wrappers and many other security tools are available from:

<ftp://ftp.porcupine.org/pub/security/index.html>

Please ensure that you check the PGP signatures of any security package that you download from the Internet, since these are obvious targets for attackers to try to replace.

Information relating to host security is particularly subject to change as new vulnerabilities and solutions are discovered. Anyone who is responsible for the security of a computer should subscribe at least to the mailing lists run by CERT-CC® and most vendors to distribute new security alerts. A list of vendor security sites is maintained by JANET-CERT at:

<http://www.ja.net/CERT/JANET-CERT/prevention/patches.html>

7.3 Network Security Measures

Denial of service attacks commonly use packets with forged source addresses, not to try to gain trust, but simply to obscure the origin of the attack. Preventing forged packets, as described in section 4.1.1, is therefore a useful defence against your site being used in such an attack. These precautions are described by the SANS (Systems Administrator, Networking and Security) Institute in their Denial of Service roadmap at:

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

Filtering of individual services is described in a CERT-CC® Tech Tip at:

http://www.cert.org/tech_tips/packet_filtering.html

Information can also be found in the book ‘Building Internet Firewalls’ by Zwicky, Cooper & Chapman, (publisher O’Reilly UK, ISBN 1-56592-871-7). Specific instructions are also available from some router vendors, for example Cisco®’s pages at:

<http://www.cisco.com/warp/public/707/21.html>

CERT teams also provide regularly updated advice based on current activity by intruders.

The use of firewalls is described in the Chapman and Zwicky book, and also in ‘Firewalls and Internet Security 2nd Edition’ by Cheswick, Bellovin and Rubin (publisher Addison-Wesley, ISBN 0-201-63466-X). The first edition of this book is available online at:

<http://www.wilyhacker.com/1e/>

A report on the successful implementation of firewalls within a UK Higher Education institution was produced by Southampton University for the JISC Technology Applications Programme (JTAP project 631). This report, ‘Use of Firewalls in an Academic Environment’ covers all aspects of choosing and configuring a firewall, as well as how to make it acceptable to the institution, and is strongly recommended. It may be found at:

http://www.jisc.ac.uk/index.cfm?name=project_firewalls

Glossary

AUP	Acceptable Use Policy.
CERT-CC®	Computer Emergency Response Team Co-ordination Centre. US funded co-ordination centre for Incident Response Teams.
DNS	Domain Name Service. Name to address translation mechanism used in the IP environment.
FTP	File Transfer Protocol. The standard protocol used on the Internet to transfer files.
HTTP	HyperText Transfer Protocol.
IMAP	Internet Message Access Protocol. A protocol for retrieving e-mail messages.
IP	Internet Protocol. The network layer protocol for the Internet.
ISP	Internet Service Provider.
JANET®	The private computer network for the UK's education and research community.
JANET-CERT	JANET Computer Emergency Response Team. A task force established to co-ordinate counter measures against computer hacking.
JANET NOSC	JANET Network Operations and Service Centre.
JISC	Joint Information Systems Committee.
JISC-ASSIST	JISC Activities, Services and Special Initiatives Support Team.
LAN	Local Area Network.
MAPS	Mail Abuse Protection System – an initiative to isolate the originators of junk e-mail.
NFS	Network File System.
POP	Post Office Protocol.
RFC	Request for Comments. The technical documents that define the Internet.
RPC	Remote Procedure Call. A protocol which allows a program running on one host to cause code to be executed on another host.
SANS Institute	System Administrator, Networking and Security Institute.
SLR	Source Level Routing.
SMTP	Simple Mail Transfer Protocol.
SNMP	Simple Network Management Protocol.
SSL	Secure Sockets Layer, also known as TLS. A protocol developed by Netscape that encrypts data which is to be transmitted via the Internet.
SSH	Secure Shell. A program which provides strong authentication and secure communications over insecure channels.

TCP	Transmission Control Protocol.
UCISA	The Universities and Colleges Information Systems Association.
UDP	User Datagram Protocol.
VPN	Virtual Private Networks.
WAN	Wide Area Network. A network where services may be a long way apart, e.g. JANET.

Tell us what you think

Technical Guides are produced and published by UKERNA for use within the JANET Community. We welcome your comments on all aspects of this document and on any other UKERNA publication. Please direct feedback to JANET Customer Service, at the address below, or e-mail us directly at:

documentation@ukerna.ac.uk

UKERNA manages the networking programme on behalf of the higher and further education and research community in the United Kingdom. JANET, the United Kingdom's education and research network, is funded by the Joint Information Systems Committee (JISC).

For further information please contact:

JANET Customer Service
UKERNA
Atlas Centre, Chilton, Didcot
Oxfordshire, OX11 0QS

Tel: 0870 850 2212
Fax: 0870 850 2213
E-mail: service@janet.ac.uk

Copyright:

This document is copyright The JNT Association trading as UKERNA. Parts of it, as appropriate, may be freely copied and incorporated unaltered into another document unless produced for commercial gain, subject to the source being appropriately acknowledged and the copyright preserved. The reproduction of logos without permission is expressly forbidden. Permission should be sought from JANET Customer Service.

Trademarks:

JANET®, SuperJANET® and UKERNA® are registered trademarks of the Higher Education Funding Councils for England, Scotland and Wales. The JNT Association is the registered user of these trademarks.

Cisco® is a registered trademark of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

Microsoft® is a registered trademark of Microsoft Corporation in the United States and/or other countries.

Netscape® is a registered trademark of Netscape Communications Corporation.

Novell® and Netware® are registered trademarks of Novell, Inc. in the United States and other countries.

Sun® and Solaris™ are trademarks of Sun Microsystems.

Tandberg® is a registered trademark in the US and certain other countries. All other trademarks are property of their respective owners.

UNIX® is a registered trademark of the Open Group.

Windows® and Windows NT® are registered trademarks of the Microsoft Corporation in the United States.

Disclaimer:

The information contained herein is believed to be correct at the time of issue, but no liability can be accepted for any inaccuracies.

The reader is reminded that changes may have taken place since issue, particularly in rapidly changing areas such as internet addressing, and consequently URLs and e-mail addresses should be used with caution.

The JNT Association cannot accept any responsibility for any loss or damage resulting from the use of the material contained herein.

Availability:

Further copies of this document may be obtained from JANET Customer Service at the above address.

This document is also available electronically from: http://www.ja.net/documents/technical_guides.html



© The JNT Association 2004

JISC

