

Security

Isolated individual computers are relatively secure as long as their physical well-being is ensured and regular backups are carried out to protect the integrity of the data held. However, once computers are connected to a LAN or WAN, they become exposed to threats which may jeopardize their proper operation and the safety and privacy of the data held.

An organisation connecting its computers to a network should therefore take measures to protect its equipment, and 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 LAN such measures represent good practice; when connecting to a WAN such as Janet, they are essential. Janet and the Internet contain a huge and diverse community of users. Both the opportunities and threats presented by this community need to be considered.

Each organisation connecting to Janet should consider how to address issues such as

- developing local security and acceptable use policies
- improving users' awareness of their responsibility and ability to protect themselves and others
- configuring networks and systems to improve security
- use of technical measures such as firewalls, filters and intrusion detection systems
- developing processes, tools and skills to detect, investigate and recover from security incidents.

Security Policy

Organisations connected to Janet are required to comply with the Janet Security Policy, to protect the security of Janet and of their own internal networks. Further details are available in Section 9.

JISC requires organisations with a Primary Connection to take responsibility for both the security of their own connection to Janet and the security of any Sponsored or Proxy Connections they provide. The Security Policy can be obtained from the Janet Service Desk or from <https://community.ja.net/library/janet-policies/security-policy> ^[1].

Appendix 9 on Packet and Content Filtering provides information on some forms of control of Internet usage and blocking the ingress of undesirable material into a network.

It is essential that organisations with a Primary Connection have at least one nominated security contact: see Section 9.

Developing an Information Security Policy

UCISA has published an Information Security Toolkit, and a Janet Training course is available on using it.

Raising Awareness

People are the most important component in any security system. Uninformed or unthinking users or administrators can make decisions for their own convenience that nullify all technical security measures. Aware and observant users and administrators can, however, reduce the likelihood and impact of security incidents even where only a few technical measures have been taken. Ideally organisations should aim to combine secure technology and security-conscious people into a truly robust system.

Janet Training holds courses and events centred on Security Awareness.

Janet has also produced a range of publications on security issues.

Firewalls

A firewall is a system that implements and enforces an access control (or security) policy between two networks, for example between an internal private network and an external public network. Essentially, a firewall connects two or more networks but only allows specified forms of traffic to flow between them.

Firewalls are used to restrict traffic between different parts of the network, thereby providing protection against some types of attack. This kind of system is often used at the entrance point of an organisation's network, where the LAN joins the WAN, to protect an organisation from hazards on the Internet. However, controls within an organisation can also be useful, for example to separate different departments, working areas or networks.

Sites that do not have a firewall may be subject to attack from hackers. If the hackers gain control of the network then the organisation may suffer:

- financial loss, such as loss of income or possibly fines/compensation imposed by a court
- loss of reputation, if embarrassing material is revealed or forged e-mails are sent
- denial of access to resources, if a key piece of network or server equipment has been rendered unserviceable.

Before implementing a firewall, an organisation must have a defined security policy. The firewall may then be used to enforce some aspects of that policy and vulnerable assets can be protected against attack from outside. A default deny firewall can also protect against unexpected forms of attack, since only predefined traffic is accepted. Without a policy, a firewall is unlikely to be effective since there is no defined basis for making decisions about which traffic should be permitted and which denied.

It is not possible to keep all of an organisation's networks entirely inside a firewall. Public servers, such as e-mail and web, must be exposed to the outside world in order to perform their functions. There is always a risk in running such services, but with careful configuration and maintenance, as well as a suitable firewall, that risk can be minimised.

Firewalls can be bought 'off the shelf' as dedicated devices or can be constructed from individual components by those with the necessary skills. A choice between these options should be based on convenience, flexibility and cost. A dedicated package is likely to be easier to configure and support but may prove inflexible and expensive, while custom-built firewalls require high levels of technical expertise but are infinitely flexible. Many routers also provide basic, but useful firewall facilities.

Further details about the process of choosing and implementing a firewall can be found in the Janet Technical Guide Firewall Implementation at Janet-connected Organisations and the report *The Use of Firewalls in an Academic Environment*.

System Configuration and Maintenance

General-purpose computer systems as supplied are not designed to be connected to hostile networks. The Internet outside the organisation should certainly be regarded as hostile and for some purposes parts of the internal organisation should also be viewed in this light. This means that many of the computers in the organisation need additional configuration and maintenance to reduce the likelihood of them falling victim to an attack across the network.

Most computers are configured to provide far more services than are required to perform their intended function. Workstations do not need to run web, database or nameserver programs; servers do not need to run web browsers or word processors. The first step in securing any computer is therefore to remove, or at least disable, any software or options that are not needed for its intended purpose. The software that is needed should then itself be configured to remove any unnecessary options that present an unacceptable risk. New threats are being discovered continually, so systems must also be maintained to take account of them. This

should be a continuous process throughout the lifetime of the computer and may include installing new software versions or patches, enabling security options and disabling insecure functions. The Janet Factsheet Securing Networked Computers provides more details.

Additional programs and services can be installed to detect and prevent attacks. One of the most effective of these is anti-virus software, which can be installed centrally, on end-user systems, or both. See the Janet Factsheet Computer Viruses - Don't Click Here.

Secure configuration and maintenance must be included in the organisation's security policy and procedures. Without guidance from these documents, this vital work will not happen. The first priority should be those computers that are intended to provide a service to an external network, for example web, mail and nameservers and proxies. These machines will appear in public directories so are likely to be the most obvious targets for attack. They will often also be less protected by firewalls than purely internal systems. Other computers that may connect directly to the untrusted Internet, for example laptops that connect from time to time to other networks, and computers that participate in Grids or other peer-to-peer systems, should also be a high priority since these cannot always be protected by an organisation's gateway firewall. It may be appropriate to run host-based software firewalls on these machines. Internal systems that are particularly important, for example file servers or administrative machines, should also be secured as a high priority, and appropriate protective software installed on other computers that handle messages from the outside world. Nor should the threat from within the organisation be forgotten: students and staff are not always well-intentioned towards the organisation. In time the aim should be to have all computers on the network significantly better protected than when they came out of their packing cases.

Monitoring and Response

Janet recommends that organisations connected to Janet carry out their own internal monitoring of their network connection. On a simple level the Janet Netsight system can highlight abnormal traffic levels on a site's access link that may be a result of illegal activity. The Janet Factsheet Unusual Traffic gives examples of how Netsight can be used to detect these kinds of problems.

Janet also recommends that organisations record sufficient information about the use of their networks and maintain tools which enable them to investigate and deal with problems. Note that any logging or monitoring that results in data about individual users will be subject to the Data Protection Act 1998, while any actions that reveal the content of communications are subject to the Regulation of Investigatory Powers Act 2000. Such actions must be properly authorised, controlled and notified to users, or else they may be criminal offences.

The Data Protection Act 1998 can be found at
<http://www.hms0.gov.uk/acts/acts1998/19980029.htm> [2]

The Regulation of Investigatory Powers Act 2000 can be found at
<http://www.hms0.gov.uk/acts/acts2000/20000023.htm> [3]

The Employment Practices Data Protection Code Part 3: Monitoring at Work issued by the Information Commissioner deals with both logging and interception. It is available from
http://www.ico.gov.uk/Home/what_we_cover/data_protection/guidance/codes_... [4]

The Janet Technical Guide Logfiles discusses the legal and technical issues involved in the recording of usage information.

There is more information in Section 9 on Janet Policies and Legal Requirements.

Active monitoring of networks can be used to obtain more detailed information. This can range from scanning networks to identify the machines present and the services they run, to full penetration testing using all the tools of intruders (technical and social) to assess the preparedness of a network and its defences. Such activities should be carefully planned with clear objectives, or a great deal of time, effort and money can be wasted. In addition to the laws mentioned above, active monitoring is also likely to be subject to other legislation including the Computer Misuse Act 1990. Monitoring must only be done with the appropriate authority on networks and systems that the organisation controls.

Janet CSIRT is able to provide advice on monitoring tools for all sites. FE and specialist colleges may also be able to receive assistance from their JISC RSC.

Detecting security problems will have little effect unless the organisation also has processes, tools and skills available to investigate and remedy the problem. The Janet Guidance Note on Effective Incident Response contains practical ideas and case studies on how this can be done at Janet sites.

Wireless Security

While wired networks tend to rely, at least in part, on physical restrictions on connection to the network to protect the privacy of communications and the accountability of messages sent by them, wireless traffic must be assumed to be 'public' since radio signals leak beyond the physical bounds of buildings and effective remote eavesdropping equipment is readily available. As a result, wireless LANs typically implement a higher standard of security (including data encryption and audit trail) than wired infrastructure (for more information see the Wireless Security Factsheet). Furthermore, as a network service aimed particularly at mobile devices and users who may have made use of other networking contexts with less protections in place, WLANs must also be resilient against inimical software on the client nodes themselves (e.g. Trojans and viruses). For this reason, it is rarely appropriate to connect wireless and wired networks together without some form of filtering and access control at the junction between the two network technologies (see the Connecting Wired and Wireless Networks Factsheet). Even the most secure WLANs rely to a degree upon responsible user behaviour to retain their integrity, both when accessing the local infrastructure and when using potentially less secure networks elsewhere (see the Factsheets Safe Use of Web Redirect Wireless Networks and Safe Use of 802.1x Wireless Networks).

Network Authentication Methods

Given the potential hostility of the wireless environment, a robust audit trail is essential to fulfil both network management and some legal obligations. The first step in this trail is to identify the connecting user reliably. Currently, there are two main authentication methods for accessing wireless networks:

- 802.1x-based
- Web-based redirect

802.1x is a port based IEEE OSI layer 2 authentication method between a mobile node and

an access control device, either a switch on a wired network or an access point in a wireless context. Robust encryption and mutual authentication make 802.1x the current leading security option for controlling network access at the edge. By providing a framework able to support a number of authentication technologies, 802.1x can accept various proofs of identity, be they token-based, conventional username and password, or certificates. 802.1x also allows the access control device to grant different types of network access depending on who the authenticated user is, or the patch level and anti-virus status of their device (e.g. unpatched systems could be assigned to a quarantine VLAN until the condition is remediated).

When users attach to a network that uses web redirection authentication, they get a docking IP address (and associated local network configuration) via DHCP, but are initially unable to receive and send traffic outside a restricted domain, typically gaining access only to web pages about the organisation or service and to a web-based SSL-encrypted login interface. To gain access beyond this, users must launch a web browser which will be redirected automatically to the authentication web page. Once username and passwords have been entered and authentication is successful, users are then granted external access in accordance with the organisation's policy (e.g. by client-specific dynamic access control on an authentication appliance or by VLAN reassignment).

Janet Roaming

Both 802.1x and web-based redirect typically rely on existing separate authentication servers, so local users can authenticate using their normal login credentials. However, Janet Roaming can also be used, if organisations wish, to allow guests from other participating Janet-connected organisations to authenticate and gain access to Janet using their home login credentials. Janet Roaming provides the means to tunnel an access authentication request securely from the visited organisation's network access server to the guest's home organisation for evaluation, and to return a response. By handing off the authentication in this way, the visited organisation is spared the administrative burden of identifying the user and managing temporary accounts, and receives a guarantee from the home organisation that the visitor is a current member in good standing by virtue of any 'access-accept' response returned.

The Policy of Janet Roaming requires that guests must respect the policy of the local site they are visiting as well as abiding by the Janet Policies and those of their home organisation.

Janet may only be used to provide network access for guests who are visiting the organisation for educational or research purposes. Organisations that wish to provide network access to members of the public, for example delegates at commercial conferences or users of other facilities of the organisation, must not use Janet for this. Other options are described in the Factsheet on Guest and Public Access.

Janet CSIRT (Computer Security Incident Response Team)

Janet provides Janet customers with help and advice on computer security and incident handling. Janet CSIRT exists to warn organisations of potential threats to computer security, to suggest how to protect against these threats and, in the last resort, to advise on rebuilding

a compromised system. The team has observed that most compromises could have been prevented if sufficient care had been taken to protect the computer systems affected.

The team also provides training and maintains a comprehensive database of relevant literature. The Janet CSIRT web site is a national source of security advice, tools and documents. As well as local information, the site has pointers to other security sites around the world.

Obtaining Advice

Janet sites with a Primary Connection should contact Janet CSIRT if they would like advice on setting up a suitable security system for their organisation or have any problems or queries.

FE and specialist colleges may also contact their JISC RSC for advice on appropriate security measures.

Security Mailing Lists

Janet CSIRT provides a mailing list service for organisation security contacts and issues early warnings of new risks and threats. When an organisation first connects to Janet, the details of the nominated security contact are forwarded to Janet CSIRT by the Janet Service Desk.

There are two mailing lists, one for announcements, the other a discussion list. Organisation security contacts will be added automatically to the announcements list but can choose whether or not to subscribe to the discussion list. They may also nominate other individuals to be added to the mailing lists.

Reference Material

Janet has published a number of Technical Guides, Guidance Notes and Factsheets on security issues such as PGP, digital signatures, firewalls, backups and viruses.

Source URL: <https://community.jisc.ac.uk/library/janet-services-documentation/security-1>

Links

[1] <https://community.ja.net/library/janet-policies/security-policy>

[2] <http://www.hmso.gov.uk/acts/acts1998/19980029.htm>

[3] <http://www.hmso.gov.uk/acts/acts2000/20000023.htm>

[4] http://www.ico.gov.uk/Home/what_we_cover/data_protection/guidance/codes_of_practice.aspx